



Essays on Information and Debt

Citation

Hebert, Benjamin. 2015. Essays on Information and Debt. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:17467323>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Essays on Information and Debt

A dissertation presented by

Benjamin Michael Hebert

to

The Committee for the PhD in Business Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Business Economics

Harvard University

Cambridge, Massachusetts

April 2015

© 2015 - Benjamin Hebert

All rights reserved.

Dissertation Advisors:

Professor Emmanuel Farhi

Professor David Laibson

Author:

Benjamin Michael Hebert

Essays on Information and Debt

Abstract

These essays attempt to explain why debt contracts are so common, and to explore the consequences resulting from the use of debt contracts. In the first essay, “Moral Hazard and the Optimality of Debt,” I use tools from information theory to study a novel form of moral hazard, and show that debt contracts are the optimal security design in this setting. In the second essay, “Generalized Rational Inattention,” written with Michael Woodford, we develop a generalized version of rational inattention, based on an axiomatic characterization, using the same theorems employed in the first essay. In the third essay, “The Costs of Sovereign Default: Evidence from Argentina,” written with Jesse Schreger, we estimate the costs that Argentina’s 2014 sovereign default imposed on Argentine firms.

Contents

1	Moral Hazard and the Optimality of Debt	1
1.1	Introduction	1
1.2	Model Framework	8
1.3	The Benchmark Model	14
1.4	The Non-Parametric Model with Invariant Divergences	21
1.5	Parametric Models with Invariant Divergences	30
1.6	Dynamic Moral Hazard	35
1.7	A Mean-Variance Approximation for Continuous Time Models	43
1.8	Conclusion	45
2	Generalized Rational Inattention	46
2.1	Introduction	46
2.2	Information Costs	49
2.3	Static Problems of Rational Inattention	57
2.4	Rational Inattention and Dynamic Asset Allocation	61
2.4.1	Fixed Mean-Variance	67
2.4.2	The Tradeoff Between Mean-Variance Optimality and Learning	68
2.5	Conclusion	70
3	The Costs of Sovereign Default:	
	Evidence from Argentina	71
3.1	Introduction	71
3.2	Argentina’s Sovereign Debt Saga	74
3.2.1	The Argentine Default of 2001 and the Restructurings of 2005 and 2010	74
3.2.2	Argentina vs. the “Vultures”	75
3.2.3	A Simple Interpretation	77
3.3	Data and Summary Statistics	78

3.3.1	Stock Market and CDS Data	78
3.3.2	Definition of Events and Non-Events	85
3.3.3	Summary of Events and Non-Events	87
3.3.4	Exchange Rates	89
3.4	Framework	90
3.4.1	OLS Estimates	94
3.4.2	Case Study: Announcement	96
3.4.3	Event Study	98
3.4.4	Standard Errors and Confidence Intervals	103
3.4.5	Variance-based Analysis	105
3.5	Cross-Sectional Evidence	108
3.6	Interpretation	118
3.6.1	The Options Available to Argentina	120
3.6.2	How Much Would Argentina Have to Repay?	122
3.6.3	Are the Legal Rulings Exogenous?	125
3.6.4	Interpreting Stock Returns	126
3.6.5	Was Argentina Already in Default?	127
3.7	Conclusion	128
4	References	129
5	Appendix	141
5.1	Appendix: Moral Hazard and the Optimality of Debt	141
5.1.1	Additional Figures	141
5.1.2	Timing Conventions	143
5.1.3	Calibration	144
5.1.4	Contracting with a Rationally Inattentive Seller	146
5.1.5	Proof of proposition 1.1	151

5.1.6	Proof of lemma 5.1	152
5.1.7	Proof of proposition 1.2	153
5.1.8	Proof of claim 5.1	154
5.1.9	Proof of proposition 1.3 and proposition 1.4	157
5.1.10	Proof of corollary 1.1	162
5.1.11	Proof of lemma 1.1	164
5.1.12	Proof of corollary 1.2	165
5.1.13	Proof of proposition 1.5	171
5.1.14	Proof of proposition 5.1	175
5.1.15	Proof of proposition 1.6	181
5.1.16	Proof of lemma 1.2	184
5.1.17	Proof of proposition 1.7	185
5.1.18	Proof of lemma 5.2	195
5.2	Appendix: Generalized Rational Inattention	200
5.2.1	Proof of theorem 2.1	200
5.2.2	Proof of lemma 2.1	202
5.2.3	Proof of theorem 2.2	203
5.2.4	Proof of lemma 5.3	209
5.2.5	Proof of lemma 2.2	211
5.2.6	Proof of lemma 2.3	212
5.2.7	Proof of lemma 2.4	213
5.2.8	Proof of lemma 5.4	214
5.2.9	Proof of theorem 2.3	216
5.3	Appendix: The Costs of Sovereign Default	221
5.3.1	GMM and Returns-IV	221
5.3.2	Mexico and Brazil	224
5.3.3	Risk-Neutral Default Probabilities	224

5.3.4	Econometric Model	226
5.3.5	Figures and Tables	229

List of Tables

3.1	Firms Included in Analysis	80
3.2	OLS Results	95
3.3	Summary Statistics	98
3.4	Standard Event Study: Index	100
3.5	Heterogeneous-Window Event Study: Index	102
3.6	IV-Style Event Study	104
3.7	CDS-IV	107
3.8	Cross-Section: Industry Returns	114
3.9	Cross-Section: Long-Short Portfolios	117
5.1	Timing Conventions During Period One	143
5.2	Returns-IV	222
5.3	GMM	223
5.4	Regressions for Brazil and Mexico	224
5.5	Events and Non-Events	230
5.5	Events and Non-Events (Continued)	232
5.5	Events and Non-Events (Continued)	233
5.5	Events and Non-Events (Continued)	234
5.5	Events and Non-Events (Continued)	235
5.5	Events and Non-Events (Continued)	236

List of Figures

1.1	The Utility of Various Security Designs	27
-----	---	----

3.1	Default Probability Change and Equity Returns during Events and Non-Events . . .	88
3.2	Official, Onshore Blue, and ADR Exchange Rates (ARS/USD)	90
3.3	Event Data from June 16, 2014, 9:30-11:30am EST	97
3.4	Estimated Response to Default Shocks: Industries	113
3.5	Estimated Response to Default Shocks: Long-Short	116
5.1	Possible Security Designs	141
5.2	Second-Order Optimal Security Designs	142
5.3	Change in Default Probability and other Financial Variables on Event and Non- Event Days	231

Acknowledgments

I am greatly indebted to my advisors, without whom this dissertation would not have been possible. I would like to thank Emmanuel Farhi for many helpful conversations during my time in graduate school, and in particular for asking some of the initial questions that led me to write “Moral Hazard and the Optimality of Debt.” I am grateful to Philippe Aghion for introducing me to contract theory, and to both Philippe and Oliver Hart for the advice and encouragement they have offered me over the years. I would like to thank John Campbell for his thoughtful feedback on virtually all of my projects, and David Scharfstein for his insistence that my theorizing remain connected with reality.

I am particularly grateful to David Laibson, who has been a wonderful advisor, mentor, and collaborator. My PhD would not have been possible without David, who, along with James Choi, Brigitte Madrian, and John Beshears, gave me an opportunity in economics. David also introduced me to Mike Woodford, which led indirectly to the second essay in this dissertation.

I would also like to thank Lars Hansen for inviting me to visit to the Becker-Friedman Institute, and for his advice during my time at the BFI. I am also grateful to all of the feedback and comments on my papers, from Alp Simsek, Alex Edmans, Luis Viceira, Jeremy Stein, Yao Zeng, Ming Yang, Sam Hanson, Adi Sunderam, Guillaume Pouliot, Yuliy Sannikov, Zhiguo He, Roger Myerson, Gabriel Carroll, Drew Fudenberg, Scott Kominers, Eric Maskin, Mikkel Plagborg-Møller, Ed Glaeser, Paul Goldsmith-Pinkham, Alex Roth, Mary Bryce Millet Steinberg, Mira Frick, Divya Kirti, Ran Shorrer, and many seminar participants.

I would like to thank my co-authors, Jesse Schreger and Mike Woodford. They are a pleasure to work with, and I am grateful for their stimulating conversations and tremendous patience.

Finally, I would like to thank my family, and especially my wife, Jean Lu. Without her support, encouragement, feedback, editing, and patience, I would never have made it through.

Preface

In Chapter 1, I show that debt securities minimize the welfare losses from the moral hazards of excessive risk-taking and lax effort. For any security design, the variance of the security payoff is a statistic that summarizes these welfare losses. Debt securities have the least variance, among all limited liability securities with the same expected value. The optimality of debt is exact in my benchmark model, and holds approximately in a wide range of models. I study both static and dynamic security design problems, and show that these two types of problems are equivalent. The models I develop are motivated by moral hazard in mortgage lending, where securitization may have induced lax screening of potential borrowers and lending to excessively risky borrowers. My results also apply to corporate finance and other principal-agent problems.

In Chapter 2, Michael Woodford and I introduce a generalized version of the rational inattention framework, based on axioms that are intuitive for many applications in economics. These axioms include standard assumptions about continuity, differentiability, and convexity, and an axiom regarding the Blackwell ordering. We show that any rational inattention cost function satisfying these axioms, not just mutual information, exhibits an invariance property. This invariance property allow us to approximate any rational inattention cost function with a Taylor expansion, which is accurate whenever the signals acquired by an agent are not very informative. The approximate cost functions we study resemble mutual information, the standard rational inattention cost function, except that it can be relatively easier or harder for the agent to learn about particular states, and there can be complementarities or substitutabilities with respect to learning about different states. These relative costs, complementarities, and substitutabilities are characterized by a matrix, defined over pairs of states. This matrix can be thought of as a primitive of the problem. We then introduce a dynamic model, in which an agent repeatedly solves a rational inattention problem. We study the continuous time limit of this dynamic model. This limit ensures that our approximation to the cost function is accurate, and allows us to make sharp predictions about the agent's behavior, regardless of the information gathering technology employed by the agent.

In Chapter 3, Jesse Schreger and I estimate the causal effect of sovereign default on the equity returns of Argentine firms. We identify this effect by exploiting changes in the probability of Argentine sovereign default induced by legal rulings in the case of *Republic of Argentina v. NML Capital*. Because the legal rulings affected the probability of Argentina defaulting on its debt, independent of underlying economic conditions, these rulings allow us to study the effect of default on firm performance. Using both standard event study methods and a heteroskedasticity-based identification strategy, we find that an increase in the probability of sovereign default causes a significant decline in the Argentine equity market. A 1% increase in the risk-neutral probability of default causes a 0.55 percent fall in the US dollar value of index of Argentine American Depositary Receipts (ADRs). Extrapolating from these estimates, we conclude that the recent Argentine sovereign default episode caused a cumulative 33 percent drop in the ADR index from 2011 to 2014. We find suggestive evidence that banks, exporters, and foreign-owned firms are particularly affected.

1 Moral Hazard and the Optimality of Debt

1.1 Introduction

Debt contracts are widespread, even though debt encourages excessive risk taking. A leading example of this phenomenon occurs in residential mortgage securitization. Prior to the recent financial crisis, mortgage lenders sold debt securities, backed by mortgage loans, to outside investors. The issuance of these securities may have weakened the incentives of mortgage lenders to lend prudently. Even before the financial crisis, the incentive problems associated with debt were well-known. Why did mortgage lenders sell debt securities? Would some other security design, such as equity, have created better incentives?

In this chapter, I show that debt is the optimal security design in a model in which both reduced effort and excessive risk-taking are possible. In the model, the seller of the security (e.g. the mortgage lender) can alter the probability distribution of outcomes in arbitrary ways. I show that, to minimize the welfare losses arising from this moral hazard, the security's payout must be designed to minimize variance. Debt securities are optimal because, among all limited-liability securities with the same expected value, they have the least variance. This theory offers an explanation for why mortgage lenders found it optimal to sell debt securities, even though debt encourages excessive risk taking.

In my model, the seller can choose any probability distribution for the value of the assets (e.g. the mortgage loans) backing the security, which allows for both reduced effort and risk-shifting. Effort refers to actions that change the mean asset value, and risk-shifting refers to actions that change other moments of the distribution. The security is the portion of the asset value received by the outside investors. If the seller retains a levered equity claim, she¹ has sold a debt security.

There are gains from trade, meaning that the outside investors value the security more than the seller does, holding the distribution of outcomes fixed. Both the outside investors and the seller are risk-neutral. There is a “zero-cost” distribution, which the seller will choose if she has no stake

¹Throughout this chapter, I will use she/her to refer to the seller and he/his to the buyer of the security. No association of the agents to particular genders is intended.

in the outcome. If the seller chooses any other probability distribution, she incurs a cost. In my benchmark model, the cost to the seller of choosing a probability distribution p is proportional to the Kullback-Leibler divergence² of p from the zero-cost distribution. Under these assumptions, the optimal limited liability security design is a debt contract.

In this model, the combined effects of reduced effort and risk-shifting can be summarized by one statistic, the variance of the security payoff. The gains from trade are proportional to the mean security payoff. I show that debt securities maximize a mean-variance tradeoff, which means that debt optimally balances the problems of reduced effort and risk-shifting against the gains from trade. This illustrates a key distinction between my model and the existing security design literature: I argue that debt is optimal precisely because both effort and risk-shifting are possible.

The classic paper of Jensen and Meckling (1976) argues that debt securities are good at providing incentives for effort, but create incentives for risk-shifting, while equity securities avoid risk-shifting problems, but provide weak incentives for effort. A natural conjecture, based on these intuitions, is that when both risk-shifting problems and effort incentives are important, the optimal security will be “in between” debt and equity. This chapter shows, contrary to this intuition, that a debt security is optimal.

The argument of Jensen and Meckling (1976) that debt is best for inducing effort relies on a restriction to monotone security designs. The “live-or-die” result of Innes (1990) (see the appendix, section §5.1.1, Figure 5.1) shows that when the seller can supply effort to improve the mean of the distribution, it is efficient to give the seller all of the asset value when the asset value is high, and nothing otherwise.³ I show that when both risk-shifting and effort are important, debt is optimal. Debt is “in between” the live-or-die security, which induces optimal effort, and equity, which avoids risk-shifting. I formalize this idea in section §1.4.

I also show that, with cost functions other than the KL divergence, the variance of the security design approximates the costs of reduced effort and risk-shifting. In these models, debt securities are approximately optimal. I then demonstrate that this mean-variance tradeoff applies to models

²The KL divergence is also known as relative entropy. It is defined in section §1.2.

³This result assumes a monotone likelihood ratio property in effort.

in which the seller can choose only a few parameters, rather than the entire probability distribution (“parametric” models). Finally, I show that this mean-variance tradeoff applies to continuous-time models of moral hazard, in which the seller can dynamically choose to exert effort. Taken together, these results show that debt is optimal or approximately optimal in a wide range of models.

With more general cost functions, debt securities are approximately optimal. The approximation I use applies when the moral hazard and gains from trade are small relative to scale of the assets. Debt is first-order optimal, meaning that when this approximation is accurate, debt securities are a detail-free way to achieve nearly the same utility as the optimal security design. A mixture of debt and equity is second-order optimal. This can be interpreted as a “pecking order,” in which the security design grows more complex as the size of both the moral hazard problem and gains from trade grow, relative to the scale of the assets.

The literature on security design with moral hazard typically studies parametric models of moral hazard. One example of a parametric model is when the seller controls the mean and variance of a log-normal distribution. In these models of moral hazard, the design of the optimal contract depends on the set of controls available to the seller. I show that regardless of the controls available to the seller, the mean-variance tradeoff provides an approximate lower bound on the security designer’s utility, and debt maximizes this lower bound.⁴ This lower bound applies to problems in which the standard “first-order approach” fails,⁵ and to most single-parameter moral hazard problems. Because this lower bound does not depend on the actions available to the agent, it is similar in spirit to the robust contracting models of Carroll (2013) and Chassang (2013).

I also provide a micro-foundation for the security design problem with the KL divergence cost function. I show that a continuous-time moral hazard problem, similar to Holmström and Milgrom (1987), is equivalent to the static moral hazard problem. The equivalence of the static and dynamic problems provides an intuitive explanation for how the seller can create any probability distribution of outcomes. The key distinction between the dynamic models I discuss and the principal-agent models found in Holmström and Milgrom (1987) is limited liability. In Holmström and Milgrom

⁴The result that debt is approximately optimal in parametric models is a generalization of Bose et al. (2011).

⁵For a discussion of the “first-order” approach to principal-agent problems, see Jewitt (1988); Mirrlees (1999).

(1987), linear contracts for the seller (agent) are optimal, because they induce the seller to take the same action each period. In my model, because of limited liability, the only way to always implement the efficient action is to offer the seller a very large share of the asset value. However, offering the seller a large share of the asset value limits the gains from trade. It is preferable to pay the seller nothing in the worst states of the world, and then at some point offer a linear payoff. Even though this design does not induce the seller to take the efficient action in every state, it achieves a much larger level of gains from trade. This design for this retained tranche, levered equity, corresponds to selling a debt security.

The models I present are motivated by the securitization of mortgage loans, and other asset-backed securities.⁶ In this setting, the seller sells a security, backed by those loans, to outside investors. I refer these investors as a single agent, the “buyer.” The buyer bears some or all of the risk associated with these loans.⁷ If the buyer does not observe the quality of the loans, there is a moral hazard problem, because the seller might not carefully screen borrowers when making loans. To mitigate this problem, the seller can retain exposure to the performance of the loans. One question that naturally arises is what the shape of this retained exposure should be. This is a security design problem, called “tranche retention” or “risk retention.” Tranche retention is the subject of new regulations after the recent financial crisis, in both the U.S. and E.U. (Fender and Mitchell (2009), Geithner (2011)). The issue of tranche retention by the seller is distinct from the question of how the buyer’s claims might subsequently be divided into tranches. This chapter is silent about the latter issue.

In my model, the seller acts to maximize the value of the portion of the asset that she does not sell to the buyer. The most obvious interpretation of this assumption is that the seller retains the tranche not sold to the buyer, as is common practice in asset securitizations (Begley and Purnanandam (2013); Gorton (2008); Gorton and Metrick (2012)). Alternatively, imagine that there are two

⁶In this context, the asset in question (the mortgage loan) is also a debt, between the homeowner and the lender. The issues of effort and risk-shifting emphasized in this chapter may also be relevant to the relationship between the homeowner and mortgage lender.

⁷This is a very stylized account of securitization. In particular, I model the originator of the mortgage loans and the sponsor of the securitization (usually an investment bank) as a single entity. See Gorton and Metrick (2012) for an in-depth description of the process.

kinds of buyers, informed and uninformed, and that informed buyers require higher returns on their investments. The seller sells the security to the uninformed buyer, who cannot observe the seller's actions, and the remaining portion of the cash flows to the informed buyer, who can observe the seller's actions. This market structure will also cause the seller take actions that maximize the value of the "retained" tranche, even though she is selling that tranche to a third party. Two recent papers provide evidence that the size of the first loss tranche in residential and commercial mortgage-backed securities deals affected the subsequent default rates on those deals, consistent with the assumption that the design of the retained tranche altered the incentives of the originator (Begley and Purnanandam (2013) and Ashcraft et al. (2014)).

In my benchmark model, debt is optimal regardless of whether the moral hazard occurs before or after the security is sold. With mortgage origination, it is natural to assume that the mortgage loans are made before the security is sold, and that this is where the bulk of the moral hazard exists. An example of the opposite timing is the so-called "Bowie bonds," in which the musician David Bowie sold debt securities backed by song royalties to the insurance company Prudential Financial Inc. (Sylva (1999)). One possible concern in such a deal is that the actions taken by Bowie to promote himself and his music might be altered by the transaction, because of his reduced incentive to maximize the value of the song royalties. In that deal, David Bowie retained the rights to the song royalties, as long as the interest on the bonds was paid, mitigating this moral hazard problem. My models apply to this version of the tranche retention problem as well.

My models also could be applied to many other settings. One possible application in finance is the design of collateralized debt obligations, in which the security designer could choose a wide range of assets to put into the CDO. The models can be applied to corporate finance, when considering the capital structure of firms, and to principal-agent problems more generally.

I focus on one particular aspect of these problems: the flexible nature of the seller's moral hazard problem. When a mortgage originator makes loans, she can screen borrowers along many different, unobservable dimensions. The benchmark model in this chapter takes this idea to one extreme, allowing the seller to create any probability distribution of outcomes, subject to a cost.

This approach to moral hazard problems was introduced by Holmström and Milgrom (1987). It is conceptually similar to the notion of flexible information acquisition, emphasized in Yang (2012).⁸ In contrast, much of literature on security design with moral hazard allows the seller to control only one or two parameters of the probability distribution. These papers do not find that debt is optimal.⁹

Several of these models are examples of the parametric models discussed previously. Closest to this chapter is Hellwig (2009), who has a two-parameter model with continuous choices for risk-shifting and effort, and finds that a mix of debt and equity are optimal. In his model, risk-shifting is costless for the agent. Fender and Mitchell (2009) have a model of screening and tranche retention, which is a single-parameter model. The model of Gorton and Pennacchi (1995) is also a single-parameter model that fits into my framework.

Innes (1990) advocates a moral-hazard theory of debt, but debt is optimal only when the seller controls a single parameter, and the security is constrained to be monotone. If the security does not need to be monotone, or if the seller controls both the mean and variance of a log-normal distribution, the optimal contract is not debt.¹⁰ In the corporate finance setting, one argument for monotonicity is that a manager can borrow from a third party, claim higher profits, and then repay the borrowed money from the extra contract payments. In addition to the accounting and legal barriers to this kind of “secret borrowing,” the third party might find it difficult to force repayment.¹¹ Moreover, in the context of asset-backed securities, where cash flows are more easily verified, secret borrowing is even less plausible.¹²

⁸One key distinction between the approach of this chapter and Holmström and Milgrom (1987), and the approach in Yang (2012), concerns the interpretation of the costs. In this chapter, the cost of choosing a probability distribution should be interpreted as a cost associated with the actions required to cause that distribution to occur (underwriting or not underwriting mortgage loans, for example). In the rational inattention literature, which Yang (2012) builds on, gathering or processing information (as opposed to taking actions) is costly. This distinction is blurred in the rational inattention micro-foundation in the appendix, section §5.1.4. There is also a methodological relationship between the two approaches (see the appendix, sections 5.1.5 and 5.1.15).

⁹In Acharya et al. (2012), bank managers can both shift risk and pursue private benefits, but do this by choosing amongst three possible investments. In Biais and Casamatta (1999), there are three possible states and two levels of effort and risk-shifting. Edmans and Liu (2011), who argue that is efficient for the agent (not the principal) to hold debt claims, also have a binary project choice.

¹⁰For brevity, I have omitted this result from this chapter. It is available upon request.

¹¹This argument against secret borrowing was suggested to me by Oliver Hart.

¹²Another argument in favor of monotonicity from corporate finance concerns the possibility of the buyer (principal, outside shareholders) sabotaging the project. In the context of securitization, the buyer exerts minimal control over the securitization trust and sabotage is not a significant concern.

Two strands of the existing security design literature have focused on adverse selection as a justification for debt. In these models, either the seller¹³ or buyer¹⁴ is endowed with, or can acquire, information about the asset's value before trading the security. In the context of mortgage origination, there is empirical evidence for lax screening by originators who intended to securitize their mortgage loans, which suggests that moral hazard is a relevant issue.¹⁵ There are also mechanisms to mitigate adverse selection by the seller, such as the inability to retain loans and random selection of loans into securitization (Keys et al. (2010)). For these reasons, I focus on moral hazard as an explanation for debt, while noting that moral hazard and buyer-side adverse selection are mutually compatible explanations for debt.¹⁶ Several other recent papers also model moral hazard and securitization, but focus on the intertemporal aspect of security design.¹⁷

In corporate finance, there are many theories to explain the prevalence of debt. I argue that most of these theories are not applicable to the tranche retention problem in securitization, which motivates my search for an alternative explanation. For example, there are usually no tax benefits to retaining one type of tranche over another. The cash flows underlying the security are readily verifiable, so arguments based on costly state verification (Townsend (1979); Gale and Hellwig (1985)) are not applicable. The servicer of the asset-backed security, who maintains the assets post-contracting, has limited ability to pursue non-pecuniary benefits or expand the scale of the assets. Explanations based on control or limiting investment (Aghion and Bolton (1992), Jensen (1986), Hart and Moore (1994)) are therefore not relevant for this context. Other explanations

¹³See Nachman and Noe (1994), DeMarzo and Duffie (1999), and DeMarzo (2005).

¹⁴See Gorton and Pennacchi (1990), Dang et al. (2011), and Yang (2012). The buyer might also anticipate adverse selection by a third-party in the future.

¹⁵See Demiroglu and James (2012); Elul (2011); Jiang et al. (2013); Keys et al. (2010); Krainer and Laderman (2014); Mian and Sufi (2009); Nadauld and Sherlund (2013); Purnanandam (2011); Rajan et al. (2010). Some of this evidence is disputed (see Bubb and Kaufman (2014)). Some of this evidence is consistent with information asymmetries, but cannot distinguish between moral hazard and adverse selection.

¹⁶In earlier versions of this chapter, I showed that the moral hazard model I develop can be combined with the model of Yang (2012), and the the optimal security design is a debt. In contrast, combining a parametric model of moral hazard with the model of Yang (2012) would not generally result in debt as the optimal security design. I view buyer-side adverse selection theories of debt as being focused on why the senior (AAA) securities were designed as debt, whereas this chapter focuses on the shape of the tranche retained by the seller. Following a different approach, Vanasco (2013) combines a seller-side adverse selection model with moral hazard.

¹⁷See Hartman-Glaser et al. (2012), and Malamud et al. (2013).

include standardization or regulatory treatment, but these explanations do not account for the wide prevalence of debt.

I begin in section §1.2 by explaining the benchmark security design problem, whose structure is used throughout the chapter. I then show in section §1.3 that for a particular cost function, debt is optimal, and explain how this relates to a mean-variance tradeoff. Next, I show in section §1.4 that this mean-variance tradeoff holds approximately for a wider class of cost functions. I will then discuss higher-order approximations, and in section §1.5 discuss parametric models. Finally, in section §1.6 and section §1.7, I provide micro-foundations for the non-parametric models, from a continuous time model.

1.2 Model Framework

In this section, I introduce the security design framework that I will discuss throughout the chapter. The problem is close to Innes (1990) and other papers in the security design literature. There is a risk-neutral agent, called the “seller,” who owns an asset in the first period. In the second period, one of $N + 1$ possible states, indexed by $i \in \Omega = \{0, 1, \dots, N\}$, occurs.¹⁸ In each of these states, the seller’s asset has an undiscounted value of v_i . I assume that $v_0 = 0$, v_i is non-decreasing in i , and that $v_N > v_0$.

The seller discounts second period payoffs to the first period with a discount factor β_s . There is a second risk-neutral agent, the “buyer,” who discounts second period payoffs to the first period with a larger discount factor, $\beta_b > \beta_s$. Because the buyer values second period cash flows more than the seller, there are “gains from trade” if the seller gives the buyer a second period claim in exchange for a first period payment. I will refer to the parameter $\kappa = \frac{\beta_b - \beta_s}{\beta_s}$ as the gains from trade.

I assume there is limited liability, so that in each state the seller can credibly promise to pay at most the value of the asset. I also assume that the seller must offer the buyer a security, meaning that the second period payment to the buyer must be weakly positive. In this sense, the seller must offer the buyer an “asset-backed security.” When the asset takes on value v_i in the second period,

¹⁸Using a discrete outcome space simplifies the exposition, but is not necessary for the main results.

the security pays $s_i \in [0, v_i]$ to the buyer. Following the conventions of the literature, I will say that the security is a debt security if $s_i = \min(v_i, \bar{v})$ for some $\bar{v} \in (0, v_N)$.

I also assume that the seller designs the security and makes a “take-it-or-leave-it” offer to the buyer at price K . If the buyer rejects the offer, the seller retains the entire asset. In this case, it is as if the seller had offered the “nothing” security at a zero price. I have given all the bargaining power to the seller, which simplifies the exposition and does not alter the main results.

I will briefly discuss several possible timing conventions for the sequence of decisions by the seller during the first period. In that period, the seller designs the security, sells it to the buyer (assuming the buyer accepts), and takes actions that will create or modify the assets backing the security. The timing convention refers to the order in which these three steps occur. In the first timing convention, the “shelf registration” convention (using the terminology of DeMarzo and Duffie (1999)), the security is designed before the assets are created, but sold afterward. In the second timing convention, the “origination” convention, the security is designed and sold after the assets are created. In the third timing convention, the “principal-agent” convention, the security is designed and sold before the seller takes her actions. In this last convention, it is natural to assume that the asset exists before the security is designed, but its payoffs are modified by the seller’s actions after the security is traded.

There are asset securitization examples for each of these timing conventions. For some asset classes, such as first-lien mortgages, the security design is standardized, and the “shelf registration” timing convention is appropriate. For more unusual assets, the security design varies deal-by-deal, and the “origination” timing convention is appropriate. In some cases, such as the Bowie bonds mentioned previously, maintaining incentives post-securitization is important, and the principal-agent timing convention applies.

Most of the results in this chapter hold regardless of the timing convention. Intuitively, because the model is a “pure” moral hazard model, there is no information to signal. This contrasts with models based on adverse selection by the seller, such as DeMarzo and Duffie (1999), in which

the timing of events is crucial. For a more detailed description of the timing conventions and equilibrium concept, see section §5.1.2 in the appendix.

The moral hazard problem occurs when the seller creates or modifies the asset. During this process, the seller will take a variety of actions, and these actions will alter the probability distribution of second period asset values over the sample space. Following Holmström and Milgrom (1987), I model the seller as directly choosing a probability distribution, p , over the sample space Ω , subject to a cost $\psi(p)$. I will consider two versions of the model: one in which any probability distribution p can be chosen—which I will call “non-parametric”—and one in which p must belong to a parametric family of distributions.¹⁹

I will make several assumptions about the cost function $\psi(p)$. First, I assume that there is a unique probability distribution, q , with full support over Ω , that minimizes the cost. Second, because I will not consider participation constraints for the seller, I assume without loss of generality that $\psi(q) = 0$. I also assume that $\psi(p)$ is strictly convex and smooth. Below, I will impose additional assumptions on the cost function, but first will describe the moral hazard and security design problems.

The moral hazard occurs because the seller cares only about maximizing the discounted value of her retained tranche. When the value of the asset is v_i , the discounted value of the retained tranche is

$$\eta_i = \beta_s(v_i - s_i).$$

Because of that assumption that $v_0 = 0$, and limited liability, it is always the case that $\eta_0 = s_0 = 0$. Denote the probability that state $i \in \Omega$ occurs as p^i , under probability distribution p . The moral hazard sub-problem of the seller can be written as

¹⁹Because the sample space Ω is a finite set of outcomes, even in the “non-parametric” case, the choice of p can be expressed as a choice over a finite number of parameters. I am using the terms non-parametric and parametric to denote whether the set M of feasible probability distributions is the entire simplex, or a restricted set.

$$\phi(\eta) = \sup_{p \in M} \left\{ \sum_{i>0} \eta_i p^i - \psi(p) \right\}, \quad (1.1)$$

where M is the set of feasible probability distributions and $\phi(\eta)$ is the indirect utility function. In the non-parametric case, when M is the set of all probability distributions on the sample space, the moral hazard problem has a unique optimal p for each η . In the parametric case, there may be multiple $p \in M$ that achieve the same optimal utility for the seller. In that case, I assume that the seller chooses from this set of optimal p to maximize the buyer's utility.²⁰

The buyer cannot observe p directly, but can infer the seller's choice of p from the design of the seller's retained tranche η . At the security design stage, the buyer's valuation of a security is determined by both the structure of the security s and the buyer's inference about which probability distribution the seller will choose, $p(\eta)$. Without loss of generality, I will define the units of the seller's and buyer's payoffs so that $\beta_s \sum_i v_i p^i(\beta_s v_i) = 1$. If the seller retains the entire asset, and takes actions in the moral hazard problem accordingly, the discounted asset value is one.²¹

Let $s_i(\eta)$ be the security design corresponding to retained tranche η . The security design problem can be written as

$$U(\eta^*) = \max_{\eta} \left\{ \beta_b \sum_{i>0} p^i(\eta) s_i(\eta) + \phi(\eta) \right\} \quad (1.2)$$

subject to the limited liability constraint that $\eta_i \in [0, \beta_s v_i]$. From the seller's perspective, when she is designing the security, she internalizes the effect that her subsequent choice of p will have on the buyer's valuation, because that valuation determines the price at which she can sell the security. The security serves as a commitment device for the seller, providing an incentive for her to choose a favorable p . This commitment is costly, however, because allocating more of the available asset value to the retained tranche necessarily reduces the payout of the security, reducing the gains from trade.

²⁰This assumption is made for expository purposes and does not affect any of the chapter's results.

²¹I use this convention to ensure that the units correspond to an empirically observable quantity—the value of the assets, if those assets are retained by the seller. This convention is useful when calibrating the model (see the appendix, section §5.1.3).

The structure of this model generalizes Innes (1990). In that paper, the model is parametric, with a single parameter, “effort.” The convex cost of effort can be rewritten to depend on p , fitting into the framework described above. The probability distribution that would result from zero effort corresponds to q . More complex parametric models (such as the ones in Fender and Mitchell (2009) and Hellwig (2009)) can also fit into this framework.

From the earlier assumptions about the cost function, it follows that the cost function can be written as a divergence²² between p and the zero-cost distribution, q , defined for all $p, q \in M$:

$$\psi(p) = D(p||q).$$

In section §1.3, I begin the chapter by focusing on a particular divergence, the Kullback-Leibler divergence. The KL divergence, also called relative entropy, is defined as

$$D_{KL}(p||q) = \sum_{i \in \Omega} p^i \ln\left(\frac{p^i}{q^i}\right).$$

The KL divergence has the assumed convexity and differentiability properties, and also guarantees that the p chosen by the seller will be mutually absolutely continuous with respect to q . The KL divergence has been used in a variety of economic models, notably Hansen and Sargent (2008), who use it to describe the set of models a robust decision maker considers. It also has many applications in econometrics, statistics, and information theory, and the connection between the security design problem and these topics will be discussed later in the chapter. I will show that when the cost function is proportional to the KL divergence, debt is the optimal security design.

I will also discuss more general class of divergences, known as the “ f -divergences” (Ali and Silvey (1966), Csiszár (1967)). This class of divergences, which include the KL divergence, can be written as

²²A “divergence” is similar to a distance, except that there is no requirement that it be symmetric between p and q , or that it satisfy the triangle inequality.

$$D_f(p||q) = \sum_{i \in \Omega} q^i f\left(\frac{p^i}{q^i}\right), \quad (1.3)$$

where $f(u)$ is a convex function on \mathbb{R}^+ , with $f(1) = 0$ and $f(u) \geq 0$. I will limit my discussion to smooth f -functions, for mathematical convenience, and use the normalization that $f''(1) = 1$. Commonly used f -divergences include the Hellinger distance $f(u) = 4(\frac{1}{2} - \sqrt{u} + \frac{1}{2}u)$, the χ^2 -divergence $f(u) = \frac{1}{2}(u^2 + 1 - 2u)$, and the KL divergence $f(u) = u \ln u - u + 1$.

The f -divergences are themselves part of a larger class of divergences, the “invariant divergences.” In section §1.4 and section §1.5, I analyze the case of an arbitrary invariant divergence cost function, for non-parametric and parametric models, respectively. Invariant divergences are defined by their invariance with respect to sufficient statistics (Chentsov (1982), Amari and Nagaoka (2007)).²³ The invariant divergences include all of the f -divergences, and also the Chernoff and Bhattacharyya distances (which are not f -divergences), among others. I will show that with an invariant divergence cost function, debt is approximately optimal.

The KL divergence, and the broader class of invariant divergences, are interesting because they are closely related to ideas from information theory. In the appendix, section §5.1.4, I illustrate this in a model based on rational inattention (Sims (2003)), in which the cost function is related to the KL divergence. The KL divergence cost function can also be micro-founded from a dynamic moral hazard problem. In section §1.6, I show that a continuous time problem is equivalent to the static moral hazard problem with the KL divergence cost function. In section §1.7, I extend this analysis to a more general class of continuous time problems and show that they are related, in a certain sense, to static moral hazard problems with invariant divergence cost functions.

In this section, I have introduced the framework that I will use throughout the chapter. In the next section, I analyze the benchmark model, in which the cost function is the KL divergence.

²³The exact definition of an invariant divergence is rather technical, and not required to comprehend the results in this chapter: one can interpret the results as applying to all f -divergences, albeit with some loss of generality. I therefore refer interested readers to the aforementioned sources for the definition of an invariant divergence.

1.3 The Benchmark Model

In this section, I discuss the non-parametric version of the model, in which the set M of feasible probability distributions is the set of all probability distributions on Ω . I assume that the cost function is proportional to the KL divergence between p and q ,

$$\psi(p) = \theta D_{KL}(p||q).$$

where θ is a positive constant that controls how costly it is for the seller to have p deviate from q . Larger values of θ make deviations more costly, and therefore reduce the moral hazard. The KL divergence cost function guarantees an interior solution to equation (1.1), for all η . It follows that the indirect utility function $\phi(\eta)$ is the convex conjugate (Fenchel-Legendre transform) of $\psi(p)$.

This observation suggests that the retained tranche η can also be thought of as a set of dual coordinates for the probability distribution p , in the sense of Amari and Nagaoka (2007). As shown by Holmström and Milgrom (1987), for each η there is a unique p that the seller will choose, and for each p there is a unique η that will cause the seller to choose p . The notion of η as a set of probability coordinates will be emphasized later in the chapter. In this section, several results related to this duality will allow me to characterize the optimal security design.

Because each η results in a unique p , and the function $p(\eta)$ is differentiable and interior, the “first-order approach” for the security design problem is valid. The Lagrangian for the security design problem (taking into account the limited liability constraints) is

$$\mathcal{L}(\eta, \lambda, \omega) = \beta_b \sum_{i>0} p^i(\eta) s_i(\eta) + \phi(\eta) + \kappa \sum_{i>0} \lambda^i \eta_i + \kappa \sum_{i>0} \omega^i (\beta_s v_i - \eta_i), \quad (1.4)$$

where I have scaled the multipliers λ and ω to simplify later equations. To derive the optimal security design, I will use several properties related to this duality. Using the envelope theorem, it follows from the moral hazard problem that

$$\frac{\partial \phi(\eta)}{\partial \eta_i} \Big|_{\eta=\eta^*} = p^i(\eta^*).$$

Intuitively, because the seller is maximizing her welfare over p , small changes in the security design affect her utility only through their direct impact on the cash flows received, and the indirect effects of the changes in $p(\eta)$ can be ignored. With this observation, we can write the first-order condition (FOC) for the Lagrangian:

$$\kappa(p^i(\eta^*) - \lambda^i + \omega^i) - \beta_b \sum_{j>0} \frac{\partial p^j(\eta)}{\partial \eta_i} \Big|_{\eta=\eta^*} s_j(\eta^*) = 0. \quad (1.5)$$

This equation states that there are two welfare-relevant components of any infinitesimal security design change. First, there is a direct effect on the gains from trade, proportional to the constant κ . Second, there is a change in the buyer's valuation of the security, due to his anticipation of the seller's changing incentives. To analyze this equation further, I will use several additional properties related to the duality of η and p .

The first-order condition of the moral hazard problem is, for $i > 0$,

$$\eta_i = \frac{\partial \psi(p)}{\partial p^i} \Big|_{p=p^*(\eta)},$$

where I use the notation that $p \in \mathbb{R}_N^+$ (not \mathbb{R}_{N+1}^+) and $p^0 = 1 - \sum_{i>0} p^i$. We can differentiate both sides by η_j and see that

$$\delta_i^j = \sum_{k>0} \frac{\partial \psi(p)}{\partial p^i \partial p^k} \Big|_{p=p(\eta^*)} \frac{\partial p^k(\eta)}{\partial \eta_j} \Big|_{\eta=\eta^*},$$

where δ_i^j is the Kronecker delta. In matrix terms, the inverse of the Hessian of the cost function is the matrix $[\frac{\partial p^k(\eta)}{\partial \eta_j} \Big|_{\eta=\eta^*}]_{kj}$.²⁴ Using this fact, we can rewrite our first-order condition as

$$\beta_b s_j(\eta^*) = \kappa \sum_{i>0} [p^i(\eta^*) - \lambda^i + \omega^i] \frac{\partial^2 \psi(p)}{\partial p^i \partial p^j} \Big|_{p=p(\eta^*)}. \quad (1.6)$$

This expression for the optimal security holds for any convex cost function that guarantees an interior solution to the moral hazard problem, not just the KL divergence. I will return to it when

²⁴By assumption, the Hessian of the cost function is positive-definite, and therefore invertible.

I consider more general cost functions. The equation demonstrates the close connection between the shape of the optimal security and the Hessian of the cost function, recognizing that $p(\eta^*)$ is itself endogenous and determined in part by the cost function.

When the cost function is proportional to the KL divergence, its Hessian is proportional to the Fisher information matrix:

$$\frac{\partial^2 \psi(p)}{\partial p^i \partial p^j} \Big|_{p=p(\eta^*)} = \theta g_{ij}(p(\eta^*)),$$

where $g_{ij}(p)$ is the (i, j) element of the Fisher information matrix, evaluated at p . To understand the security design first-order condition, consider the $N = 2$ case (three second-period states, including the $v_0 = 0$ state), for which the Fisher information matrix is

$$[g_{ij}(p)] = \begin{bmatrix} \frac{1}{p^1} & 0 \\ 0 & \frac{1}{p^2} \end{bmatrix} + \begin{bmatrix} \frac{1}{p^0} & \frac{1}{p^0} \\ \frac{1}{p^0} & \frac{1}{p^0} \end{bmatrix}.$$

The inverse Fisher information matrix, which I will denote $[g^{ij}(p)]$, can be written in the $N = 2$ case as

$$[g_{ij}(p)]^{-1} = [g^{ij}(p)] = \begin{bmatrix} p^1 & 0 \\ 0 & p^2 \end{bmatrix} - \begin{bmatrix} p^1 \\ p^2 \end{bmatrix} \begin{bmatrix} p^1 & p^2 \end{bmatrix}.$$

The inverse Fisher information matrix can be used to compute the variance. For any security design s_i , the variance of s_i under probability distribution p is

$$V^P[s_i] = \sum_i \sum_j s_i s_j g^{ij}(p).$$

In words, in the non-parametric model, the Cramér-Rao bound is an equality. This observation can be used to connect the first order condition of the security design problem with moral hazard (equation (1.5) above) to a different security design problem. I will first discuss the connection between these two security design problems, and then demonstrate that debt is optimal.

Suppose that the asset value has a fixed probability distribution, \bar{p} . Consider the set of limited liability securities, and ask the following question: of all securities with the same expected value under \bar{p} , which security has the least variance in its payout? Formally, the problem is

$$\min_{s_i} V^{\bar{p}}[s_i],$$

subject to the limited liability constraints and that $E^{\bar{p}}[s_i] = C$, for some constant C . Here, $V^{\bar{p}}[\cdot]$ and $E^{\bar{p}}[\cdot]$ denote the variance and mean, respectively, under the probability distribution \bar{p} . The first-order condition of this problem is

$$\hat{\kappa}(\bar{p}^i - \lambda^i + \omega^i) - \beta_b \sum_{j>0} g^{ij}(\bar{p})s_j = 0,$$

where $\hat{\kappa}$ is the multiplier on the constraint that $E^{\bar{p}}[s_i] = C$. This equation is exactly the same as the first-order condition in the moral hazard problem, if $\bar{p} = p(\eta^*)$ and $\hat{\kappa} = \kappa$.

The optimal contract in the security design problem with moral hazard problem is the one that minimizes the variance of its payout, among all limited liability securities with the same expected value. Intuitively, because debt securities are “flat” wherever possible, they minimize variance²⁵, and are therefore optimal. Examining the equity, live-or-die, and debt securities shown in the appendix, Figure 5.1, it is clear why the debt security minimizes the variance of the payout, among all limited-liability securities. The proof of proposition 1.1 below shows both that the variance-minimizing security is a debt contract, and that this is optimal in the security design problem with moral hazard.

Proposition 1.1. *In the non-parametric model, with the cost function proportional to the Kullback-Leibler divergence, the optimal security design is a debt contract,*

$$s_j(\eta^*) = \min(v_j, \bar{v}),$$

²⁵It seems likely that the variance-minimizing property of debt has been pointed out by previous authors, but I am not aware of any papers mentioning this fact. Shavell (1979) emphasizes that flat contracts minimize variance, in a context without limited liability.

for some $\bar{v} > 0$. If the highest possible asset value is sufficiently large ($v_N > \sum_i q^i v_i + \frac{\kappa}{\beta_b} \theta$), then $\bar{v} < v_N$.

Proof. See appendix, section 5.1.5. □

The result in proposition 1.1 shows that debt is optimal, for any full-support zero-cost distribution q .²⁶ The condition that v_N be “high enough” is weak. If it was not satisfied for some sample space Ω and zero-cost distribution q , one could include a new highest value v_{N+1} in Ω , occurring with vanishingly small probability under q , such that the condition was satisfied. Intuitively, the sample space must contain high enough values to observe the “flat” part of the debt security.

The optimality of debt runs contrary to the themes of the existing literature, which emphasize that risk-shifting undoes the optimality of debt (Acharya et al. (2012); Biais and Casamatta (1999); Fender and Mitchell (2009); Hellwig (2009); Jensen and Meckling (1976)). The intuition for the optimality of debt comes from its mean-variance optimality. Higher mean values of the security (under a fixed distribution) are desirable because they drive the gains from trade. Higher variance in the security payout is problematic for two reasons. First, variance in the security payout reduces the seller’s incentive to increase the mean value of the asset, relative to the zero-variance (“sell nothing”) security. Second, variance in the security payout gives the seller an incentive to inefficiently shift cash flows towards her retained tranche, holding the mean asset value constant. Every security is subject to these two effects, although which of these effects is most prevalent depends on the security design. For equity securities, there is no risk-shifting, and the variance of the security design summarizes reduced effort. For the live-or-die security design, effort is high, possibly excessive, and there is also large amounts of risk-shifting. For debt securities, both reduced effort and risk-shifting occur. Under the KL divergence cost function, at the margin, for any security design, these two effects are exactly summarized by one statistic: the variance of the security payout. Debt securities are optimal because they are mean-variance optimal, minimizing the combined losses due to reduced effort and risk-shifting.

²⁶The level of the debt, \bar{v} , depends on q indirectly, through the endogenous probability distribution $p(\eta^*)$.

A different way to view this intuition is to consider small perturbations to the retained tranche, $\eta_i = \eta_i^* + \tau \varepsilon_i$, where for small enough τ , the retained tranche still satisfies limited liability. Changing the retained tranche will cause the seller to change her behavior in the moral hazard problem. If the perturbation ε provides incentives that go “in the same direction” as the existing incentives η , the seller will move the probability distribution p further away from q , and increase $D_{KL}(p||q)$. For small perturbations, the amount by which $D_{KL}(p||q)$ will change is proportional to the covariance of η and ε under the endogenous probability distribution $p(\eta^*)$. The perturbation ε might also have a benefit, in that it causes the seller to increase the expected value of the asset under the endogenous probability distribution $p(\eta)$. The extent to which $E^{p(\eta)}[\beta_s v_i]$ changes, for small perturbations, is $Cov^{p(\eta^*)}[\beta_s v_i, \varepsilon_i]$. Intuitively, if the seller’s incentives become more aligned with the value of the asset, the seller will act to increase the asset value. The combined effects of the change in moral hazard costs and the change in asset value is $Cov^{p(\eta^*)}[\beta_s s_i, \varepsilon_i]$, or the extent the perturbation is aligned with the security design. This is also the opposite of the change in the variance of the security design, which explains why the variance of the security design summarizes the moral hazard problems.

Several of the assumptions in the benchmark model can be relaxed without altering the debt security result of proposition 1.1.²⁷ The lowest possible value, v_0 , can be greater than zero. The seller and the buyer can Nash bargain over the price of the security, and the result will hold as long as the seller has some bargaining power. The buyer can be risk-averse, with any increasing, differentiable utility function. Finally, although the model is set up as a security design problem, similar results could be obtained in principal-agent and investor-entrepreneur contexts.

One additional aspect of these results worth emphasizing is the issue of redundant states. In the setup of the model, there can be two states, i and j , such that $v_i = v_j$. It is a result of proposition 1.1 that the security payoff will be identical in these states. One way of interpreting this result relates to the practice of “pooling” multiple assets together into a single securitization structure. We can think of the redundant states v_i and v_j as two states in which the combined value of all the assets

²⁷These assumptions are useful later in the chapter, for both parametric models and non-parametric models without the KL divergence cost function.

in the securitization pool is the same, but the composition of the asset value is different. It is a result of the model that the optimal security designs pays equally in those two states. Intuitively, while the security design could reward the seller for increasing the value of some assets in the pool relative to other assets, it is not efficient to do so in the presence of this form of moral hazard.

The optimal security described in proposition 1.1 has an interesting comparative static. Define the “put option value” of a debt contract as the discounted difference between its maximum payoff \bar{v} and its expected value. For the optimal security design²⁸,

$$P.O.V. = \beta_b \bar{v}(\eta^*) - \beta_b \sum_{i>0} p^i(\eta^*) s_i(\eta^*) = \kappa \theta. \quad (1.7)$$

When the constant θ is large, meaning that it is costly for the seller to change the distribution, the put option will have a high value. Similarly, when the gains from trade, κ , are high, the put option will have a high value. For all distributions q , a higher put option value translates into a higher “strike” of the option, \bar{v} , although the exact mapping depends on the distribution q and the sample space Ω . Restated, when the agents know that the moral hazard is small, or that the gains from trade are large, they will use a large amount of debt, resulting in a riskier debt security.²⁹

Before generalizing this theory, I will briefly discuss the ways in which the model could be applied to data. Suppose that we had a reasonable estimate of θ . To compute the riskiness of the optimal debt security, as measured by the “put option value,” the only additional information required is the gains from trade κ . There are several empirical papers that attempt to estimate this parameter directly (see the appendix, section §5.1.3). Armed with estimates of θ and κ , one could compare actual security designs against the model’s predictions.

This approach would require an estimate of θ . I will discuss two distinct approaches, and describe them in more detail in the appendix, section §5.1.3. The first approach could be thought of as an “experimental” approach. If an experimenter randomly assigned different sellers one of two retained tranches, η and $\hat{\eta}$, and observed the resulting endogenous probability distributions

²⁸See lemma 5.1 in the appendix.

²⁹The model also has comparative statics for the zero-effort distribution q , which are related to the “informativeness principle.” I have omitted them for the sake of brevity.

$p(\eta)$ and $p(\hat{\eta})$, the experimenter could estimate the parameter θ (assuming that the cost function was the KL divergence).³⁰ Unfortunately, because empirical work necessarily observes ex-post outcomes, rather than ex-ante probabilities, and the outcomes of different pools of loans are not independent, this approach is difficult to implement credibly. The second approach is to infer θ from the design of securitizations (although one cannot then test whether these designs were optimal). This approach is somewhat more successful, insofar as it can be carried out using estimates from the empirical literature, but far from conclusive. For details, see the appendix, section §5.1.3.

In the next section, I consider alternative cost functions, and show that the mean-variance intuition for the optimality of debt is a general phenomena.

1.4 The Non-Parametric Model with Invariant Divergences

In this section, I analyze the general case of invariant divergence cost functions. First, I will show that among a special subset of invariant divergences known as f -divergences, the Kullback-Leibler divergence is the only divergence that always results in debt as the optimal security design. Second, I will return to the general case of invariant divergence cost functions, and show that debt is approximately optimal for all invariant divergence cost functions. Third, I will discuss higher-order approximations, under which the optimal security design is a mixture of debt and equity.

I begin by assuming that the cost function is proportional to a smooth f -divergence (defined earlier, see equation (1.3)):

$$\psi(p) = \theta D_f(p||q).$$

These divergences are analytically tractable because they are decomposable, meaning that the cost of choosing some p^i is not affected by value of p^j , $j \neq i$, except through the adding-up constraint on probability distributions. In some cases, such as the Hellinger distance or KL divergence, the seller's choice of p is guaranteed to be interior, but this is not true for all f -divergences.

³⁰ In fact, the experimenter could even test the assumption of the KL divergence cost function.

Whenever the solution is interior, the optimal security design is characterized by equation (1.6). The Hessians of the f -divergences have a simple structure. Using this structure, I find the following result:

Proposition 1.2. *In the non-parametric model, with an f -divergence cost function, if the optimal security design is debt for all sample spaces Ω and zero-cost probability distributions q , then that f -divergence is the Kullback-Leibler divergence.*

Proof. See appendix, section 5.1.7. □

The statement of proposition 1.2 shows that the KL divergence is special, in the sense that it is the only smooth f -divergence that always results in debt as the optimal security design. The structure of the proof uses the “flat” part of the debt security to derive this result. I first show that to have a flat optimal contract, when the limited liability constraints don’t bind, the $f(u)$ function must be equal to $u \ln u - u + 1$ for the values of u that (endogenously) correspond to the flat part of the contract. I then show that for debt to be optimal for all sample spaces and zero-cost distributions, $f(u) = u \ln u - u + 1$ for its entire domain.

What does this negative result imply about the mean-variance intuition discussed previously? I interpret proposition 1.2 as showing that the marginal losses due to reduced effort and risk-shifting are not perfectly summarized by the variance of the security payout, unless the cost function is the KL divergence. However, the intuition that security payout variance is costly, because it reduces effort and increases risk-shifting, still holds. This leads me to investigate whether the utility in the security design problem can be approximated as a mean-variance tradeoff, and therefore whether debt is “approximately optimal.”

The approximation I consider is a first-order expansion of the utility function in the security design problem. I approximate the utility of using an arbitrary security design s , relative to selling nothing, to first order in θ^{-1} and κ . When θ^{-1} is small, and therefore θ is large, it is difficult for the seller to change p . When κ is small, the gains from trade are low. One possible justification for this approximation is that the time period in question is short. Another possible justification

is that legal recourse or other forms of monitoring deter the seller from taking extreme actions, explaining why θ is high. I take this approximation around the limit point $\theta^{-1} = \kappa = 0$. This approximation applies when θ^{-1} and κ are small but positive. The limit point itself is degenerate; because there is no moral hazard and no gains from trade, the security design does not matter. Using this approximation, I will show that debt securities achieve approximately the same utility as the optimal security design, for any invariant divergence cost function. Moreover, only debt contracts have this property, and it arises through the same mean-variance intuition discussed in the previous section.

The relevance of the approximation will depend on whether θ^{-1} and κ are small enough, relative to the higher order terms of the utility function, for those terms to be negligible. This is a question that can only be answered in the context of a particular application. In the appendix, section §5.1.3, I discuss a calibration of the model relevant to mortgage origination, for which the approximation is accurate. It is important to distinguish between the relevance of the approximation and the economic importance of the problem. In the calibration for mortgage origination, debt achieves almost the same utility as the optimal security design, while the utility difference between the best debt contract and selling nothing is about 0.73% of the total asset value. A single mortgage securitization could involve billions of dollars in loans, and a single investment bank could sponsor many such deals each year. In this calibration, even though the moral hazard and gains from trade are small enough that debt is nearly optimal, the private gains for those involved in securitization are very large.³¹

With the KL divergence cost function, debt is the optimal security design because the KL divergence's Hessian, with respect to p , is the Fisher information matrix. There is a broader connection between all invariant divergences with continuous second derivatives (which include all smooth f -divergences) and the Fisher information matrix. At the point q , the Hessian for all such divergences is proportional to the Fisher information matrix (Chentsov (1982)):

³¹I emphasize that the private gains are large, because I do not model foreclosure or other externalities that may be relevant for social welfare.

$$\frac{\partial^2 D(p||q)}{\partial p^i \partial p^j} \Big|_{p=q} = c \cdot g_{ij}(q),$$

where c is a positive constant. I will assume the cost function is proportional to an invariant divergence, with the previously mentioned convexity and differentiability properties,

$$\psi(p) = \theta c^{-1} D(p||q),$$

ensuring that the Hessian of the cost function is the positive constant θ times the Fisher matrix. When θ is large, the seller will endogenously choose a p that is close to q . In the neighborhood of q , the cost function “looks like” the KL divergence, in the sense of having the Fisher matrix as its Hessian. Putting these two ideas together, it is not surprising that the optimal security design resembles a debt contract.

I consider a first-order asymptotic expansion of the security design problem utility, $U(s; \theta^{-1}, \kappa)$, around the point $\theta^{-1} = \kappa = 0$. This approximation holds β_s fixed as it changes κ , implicitly changing β_b . I consider the utility of an arbitrary security s , relative to the “sell-nothing” security³², and derive the following result:

Proposition 1.3. *In the non-parametric model, with a smooth, convex, invariant divergence cost function, the difference in utilities achieved by an arbitrary security s and the sell-nothing security is*

$$U(s; \theta^{-1}, \kappa) - U(0; \theta^{-1}, \kappa) = \kappa E^q[\beta_s s_i] - \frac{1}{2} \theta^{-1} V^q[\beta_s s_i] + O(\theta^{-2} + \theta^{-1} \kappa). \quad (1.8)$$

Proof. See appendix, section 5.1.9. □

The intuition of the security design problem as a mean-variance optimization problem holds (approximately) for all invariant divergences, not just the KL divergence. The proof argues that

³²There are two natural benchmark security designs, selling nothing and selling everything. The choice of benchmark does not affect the optimal security design. I use selling nothing as the benchmark, because it clarifies why debt is approximately optimal (the mean-variance tradeoff).

when θ is large, p will be close to q , and then uses the aforementioned fact that the Hessian of the invariant divergence, at point q , is the Fisher information matrix at point q .

Two related results show that as both the moral hazard and gains from trade become small, the optimal contract converges to debt, and a debt contract is first-order optimal.

Corollary 1.1. *Define s_{debt} as the unique solution (which is a debt security) to the mean-variance security design problem suggested by equation (1.8), ignoring the $O(\theta^{-2} + \theta^{-1}\kappa)$ terms:*

$$s_{debt}(\theta^{-1}, \kappa) = \arg \max_s \kappa E^q[\beta_s s_i] - \frac{1}{2} \theta^{-1} V^q[\beta_s s_i],$$

subject to the limited liability constraints. Under the conditions of proposition 1.3, defining $s^(\theta^{-1}, \kappa)$ as the optimal contract given θ^{-1} and κ ,*

1. $\lim_{\theta^{-1} \rightarrow 0^+} s^*(\theta^{-1}, \bar{\kappa} \theta^{-1}) = s_{debt}(1, \bar{\kappa})$ and
2. $U(s^*(\theta^{-1}, \kappa); \theta^{-1}) - U(s_{debt}(\theta^{-1}, \kappa); \theta^{-1}, \kappa) = O(\theta^{-2} + \theta^{-1} \kappa)$.

Proof. See appendix, section 5.1.10. □

The accuracy of the approximation that both the moral hazard and gains from trade are small will vary by application. The generality of proposition 1.3, which holds for all sample spaces, zero-cost distributions, and invariant divergences, suggests that as long as the moral hazard is not too large, the agents can neglect the details of the cost function. In this case, the mean-variance intuition for security design holds, and a debt security is approximately optimal.

The result of corollary 1.1 shows that when the gains from trade and moral hazard are small, but not zero, debt is approximately optimal in a way that other security designs are not. In Figure 1.1, I illustrate this idea. I plot the utility of the optimal debt contract, optimal equity contract, and the optimal contract, relative to selling everything, for different values of θ , with $\kappa = \bar{\kappa} \theta^{-1}$.³³ As θ becomes large, all security designs converge to the same utility. For intermediate values of θ , the optimal debt contract achieves nearly the same utility as the optimal contract, which is what the

³³The results of proposition 1.3 reference the utility of a security relative to selling nothing, instead of selling everything. In Figure 1.1, I plot utilities relative to selling everything for visual clarity.

approximation results show. For low values of θ , the high-stakes case, the gap between the optimal debt contract and optimal contract grows.

These results hold for any divergence whose Hessian, at the point $p = q$, is proportional to the Fisher information matrix. The earlier discussion of security design with KL divergence cost functions emphasized the connection between debt securities and the Fisher information matrix. In light of that discussion, the mean-variance results presented here are not surprising, conditional on the assumption that the Hessian of the cost function is the Fisher information matrix. The less intuitive aspect of the result is the connection between that assumption and the invariant divergences, which are defined by certain desirable information-theoretic properties (Chentsov's theorem).

I next turn to the question of higher order approximations. The mean-variance characterization of the utility function, up to order θ^{-1} , naturally raises the question of what the $O(\theta^{-2} + \kappa\theta^{-1})$ terms look like. Fortunately, the theorem of Chentsov (1982) also characterizes the third derivative of all invariant divergences. Using this theorem, I derive the following lemma:

Lemma 1.1. *For every invariant divergence with continuous third derivatives, there exists a real number α such that*

$$\frac{\partial^3 D_f(p||q)}{\partial p^i \partial p^j \partial p^k} \Big|_{p=q} = c \left(\frac{3 + \alpha}{2} \right) \partial_i g_{jk}(p) \Big|_{p=q},$$

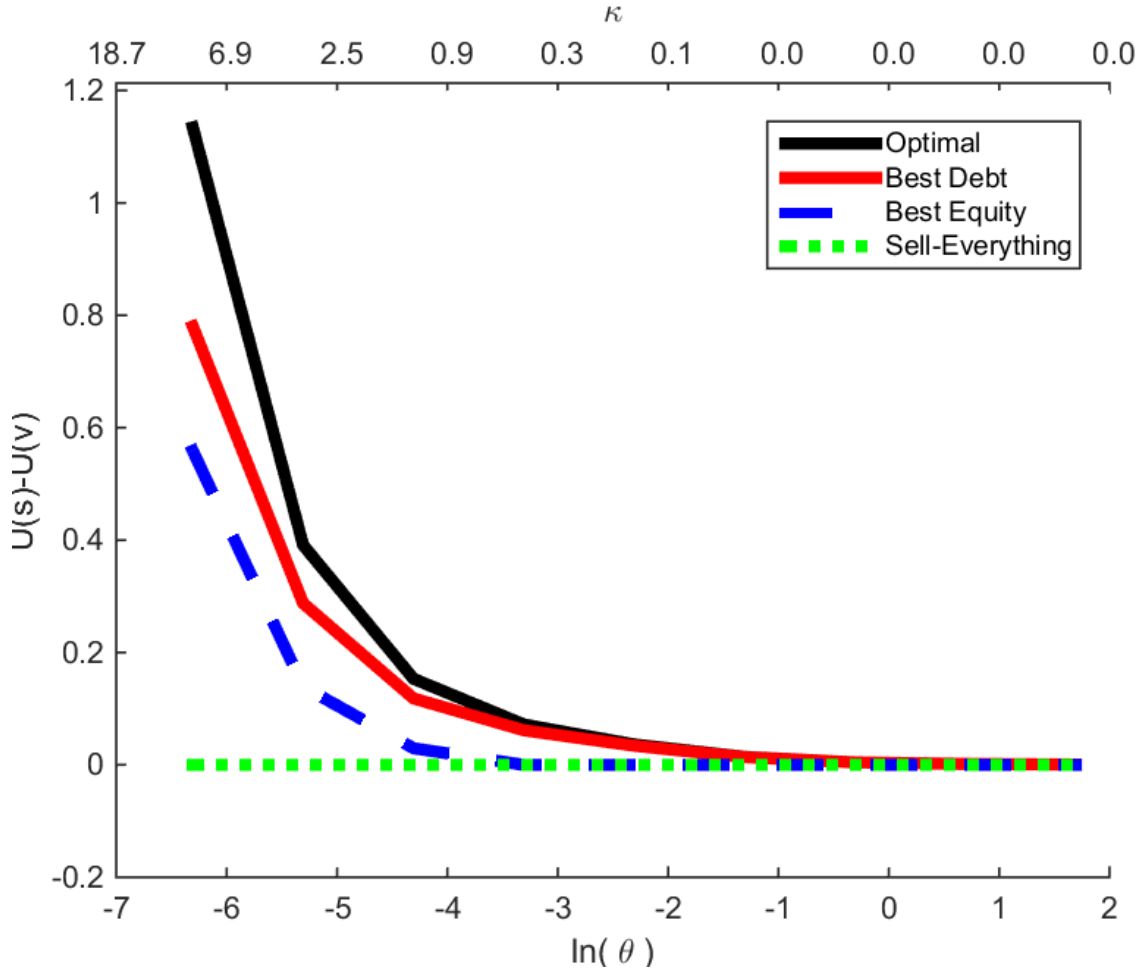
where c is a positive constant such that $c \cdot g_{ij}(q) = \frac{\partial^2 D_f(p||q)}{\partial p^i \partial p^j} \Big|_{p=q}$.

Proof. See appendix, section 5.1.11. □

The parameter α controls how rapidly the curvature of the Hessian of the cost function changes, as p moves away from q . For the KL divergence, $\alpha = -1$, while the Hellinger distance and χ^2 -divergence have α values of zero and negative three, respectively. Unsurprisingly, the parameter α influences the design of the second-order optimal security.

Proposition 1.4. *Under the assumptions of proposition 1.5, the difference in utilities between an arbitrary security s and selling nothing is*

Figure 1.1: The Utility of Various Security Designs



Notes: This figure compares the utility of several security designs (debt, equity, and the optimal security design) relative to the utility of selling everything, for different values of θ . The bottom x-axis is the value of $\ln(\theta)$, the top x-axis is the value of κ , and the y-axis is the difference in security design utility between the security (debt, equity, etc.) and selling everything. For each θ and corresponding κ , the optimal debt security, equity security, and the optimal security are determined. Then, the utility of using each of the four securities designs, given θ and κ , is computed. The cost function is a α -divergence, with $\alpha = -7$, which was chosen because the optimal security design is easy to characterize (it is a mix of debt and equity) and it is sufficiently different from the KL divergence to ensure that debt is not always optimal. The gains from trade, κ , vary as θ changes, with $\kappa = \bar{\kappa}\theta^{-1}$, $\bar{\kappa} = 0.0171$. This parameter was chosen to be consistent with the calibration in the appendix, section §5.1.3. The discounting parameter for the seller is $\beta_s = 0.5$. The zero-cost distribution q is a discretized, truncated gamma distribution with mean 2, 0.3 standard-deviation, and an upper bound of 8. The outcome space v is a set of 401 evenly-spaced values ranging from zero to 8. The utilities are plotted for nine different values of θ , ranging from $2\exp(-7)$ to $2\exp(1)$, and linearly interpolated between those values.

$$\begin{aligned}
U(s; \theta^{-1}, \kappa) - U(0; \theta^{-1}, \kappa) &= \kappa E^q[\beta_s s_i] - \frac{1}{2} \theta^{-1} V^q[\beta_s s_i] + \\
&\quad \kappa \theta^{-1} \text{Cov}^q[\beta_s s_i, \eta_i(s)] - \frac{3 + \alpha}{12} \theta^{-2} K_3^q[\eta_i(s), \beta_s s_i, \beta_s v_i] - \\
&\quad \frac{3 + \alpha}{6} \theta^{-2} K_3^q[\beta_s s_i, \beta_s s_i, \beta_s s_i] + O(\theta^{-3} + \kappa \theta^{-2}),
\end{aligned}$$

where $\text{Cov}^q[X, Y]$ is the covariance under probability distribution q and $K_3^q[X, Y, Z]$ is the third cross-cumulant under q .

Proof. See appendix, section 5.1.9. □

The covariance term in proposition 1.4 represents the interaction between the gains from trade and the moral hazard, while the third cross-cumulant terms can be thought of as accounting for the difference between $p(\eta)$, the endogenous probability distribution, and q . The economic intuition behind this result is harder to grasp than the mean-variance intuition, but fortunately the optimal security design helps clarify matters. Taking the first-order condition of this expression, I derive the following corollary:

Corollary 1.2. *Define $s_{debt-eq}$ as the optimal security design for the problem suggested by proposition 1.4, ignoring the $O(\theta^{-3} + \kappa \theta^{-2})$ terms. Assuming that $\alpha < (1 + \frac{2}{\kappa})$, the second-order optimal security design, for some constant $\bar{v} > 0$, is*

$$s_{debt-eq,i}(\theta^{-1}, \kappa) = \begin{cases} v_i & \text{if } v_i < \bar{v} \\ \max[-\frac{\kappa(1+\alpha)}{2+\kappa(1-\alpha)}(v_i - \bar{v}) + \bar{v}, 0] & \text{if } v_i \geq \bar{v}. \end{cases}$$

Under the conditions of proposition 1.4, defining $s^*(\theta^{-1}, \kappa)$ as the optimal contract given θ^{-1} and κ ,

$$U(s^*(\theta^{-1}, \kappa); \theta^{-1}, \kappa) - U(s_{debt-eq}(\theta^{-1}, \kappa); \theta^{-1}, \kappa) = O(\theta^{-3} + \kappa \theta^{-2}).$$

Proof. See appendix, section 5.1.12. □

The higher-order optimal security design can be thought of as a mixture of debt and equity (at least when $\alpha \leq -1$), whose slope is determined by the gains from trade parameter κ and the change of curvature parameter α . For the KL divergence, $\alpha = -1$, the contract is always a debt contract, regardless of κ . As the moral hazard and gains from trade become small, in the limit considered earlier, the optimal contract converges to a debt contract, regardless of the parameter α . The restriction that $\alpha < (1 + \frac{2}{\kappa})$ is without loss of generality in the limit as $\kappa \rightarrow 0$. It also worth noting that for any $\alpha > -1$, the optimal contract is non-monotonic. For any invariant divergence cost function with such an α , imposing the restriction that the security design be monotonic results in a debt contract as the (second-order) optimal contract. Viewed in this light, it is not surprising that many authors in the security design literature find that debt is optimal when assuming monotonicity.

One way to interpret this result is to think of α as controlling the balance between effort and risk shifting. When α is large and negative, the second-order optimal security design resembles an equity contract, because concerns about risk-shifting dominate concerns about effort. When α is large and positive, the second-order optimal security design approaches a “live-or-die” contract, suggesting that concerns about effort dominate concerns about risk-shifting. The live-or-die, debt, and equity security designs exist on a continuum, indexed by α , that formalizes what it means for debt to be “in between” the live-or-die and equity contracts. In Figure 5.2, in the appendix, I illustrate the different second-order optimal security designs associated with varying values of α .

The results for first-order and second-order optimal security designs can be summarized as a type of “pecking order” theory. When the moral hazard and gains from trade are small, the agents can use debt contracts. As the stakes grow larger, so that both the moral hazard and gains from trade are bigger concerns, the agents can use a mix of debt and equity. For very large stakes, the security design will depend on the precise nature of the moral hazard problem.

Having analyzed non-parametric models, I next turn to parametric models. I will show that the mean-variance intuition and the approximate optimality of debt also applies for these models.

1.5 Parametric Models with Invariant Divergences

The moral hazard models of Innes (1990), Hellwig (2009), and Fender and Mitchell (2009) can all be thought of as restriction the seller's choice of probability distributions to a parametric model. In this section, I discuss a parametric moral hazard problem with an invariant divergence cost function. This discussion will also allow me to consider “almost non-parametric” models, in which there is a single dimension of aggregate risk that is not controlled by the seller.

I assume that the set of feasible probability distributions, M_ξ , is a curved exponential family, smoothly embedded into the set of all distributions on Ω .³⁴ The parameters, ξ , are coordinates for the curved exponential family. I will rewrite the parametric moral hazard sub-problem as

$$\phi(\eta; M_\xi, \theta^{-1}) = \max_{p \in M_\xi} \left\{ \sum_{i>0} p^i \eta_i - \psi(p; \theta^{-1}) \right\}. \quad (1.9)$$

I assume the seller controls at least one parameter that alters the mean asset value, so the problem is not trivial. As noted earlier, in the parametric model there is no guarantee that there is a unique p which maximizes equation (1.9). I assume for all retained tranches η , for each optimal p , the parameters ξ corresponding to that p are interior. As in the previous section, I will assume that

$$\psi(p) = \theta c^{-1} D(p||q)$$

where $D(p||q)$ is a invariant divergence and c is defined as in Chentsov's theorem, and continue to assume that D is smooth and convex in p . I also assume that $q \in M_\xi$, meaning that the “zero-cost” distribution is feasible. Let $U(s; M_\xi, \theta^{-1})$ be the utility in the security design problem,

$$U(s; M_\xi, \theta^{-1}, \kappa) = \beta_s (1 + \kappa) \sum_{i>0} p^i(\eta(s); M_\xi, \theta^{-1}) s_i + \phi(\eta(s); M_\xi, \theta^{-1}),$$

where $p(\eta; M_\xi, \theta^{-1})$ is the endogenous probability the seller would choose given η , M_ξ , and θ^{-1} . I have written $p(\eta; M_\xi, \theta^{-1})$ as a function, assuming an arbitrary rule for choosing between

³⁴Because the set of all probability distributions on a discrete outcome space Ω is an exponential family, this is almost without loss of generality. The key restriction is that the sub-manifold M_ξ is smoothly embedded.

different optimal p , in the event that there are multiple maximizers of equation (1.9). Unlike the non-parametric models considered previously, the mapping between securities and probability distributions is many-to-one. Moreover, even under the approximation considered in the previous section, the optimal contract does not approach debt. For example, in the Innes (1990) problem, the optimal contract in this limit is still a “live-or-die” contract. Nevertheless, I use the approximation to derive an analog to proposition 1.3 for parametric models:

Proposition 1.5. *In the parametric model, with a smooth, convex, invariant divergence cost function, the difference in utilities achieved by an arbitrary security s and the sell-nothing security, for sufficiently small θ^{-1} and κ , is bounded below:*

$$U(s; M_\xi, \theta^{-1}, \kappa) - U(0; M_\xi, \theta^{-1}, \kappa) \geq \kappa E^q[\beta_s s_i] - \frac{1}{2} \theta^{-1} V^q[\beta_s s_i] + O(\theta^{-2} + \kappa \theta^{-1}; M_\xi).$$

The lower bound, to first order, does not depend on M_ξ , only on q . There exist parameters ξ and corresponding set of feasible probability distributions M_ξ such that the lower bound is tight. The notation $O(\theta^{-2} + \kappa \theta^{-1}; M_\xi)$ indicates terms of order θ^{-2} , $\kappa \theta^{-1}$, or lower that may depend on M_ξ .

Proof. See appendix, section 5.1.13. □

In the parametric model, there is still a mean-variance tradeoff between the gains from trade and the losses due to both reduced effort and risk-shifting. However, because the model is parametric, there are ways that the security can vary that do not alter the incentives of the seller, because there is no action she can take to respond to this variation. The seller’s limited ability to respond to changed incentives explains why the result in proposition 1.5 is a lower bound.³⁵

This lower bound applies to a large class of problems in the existing security design literature. For problems with a single choice variable, such as Innes (1990), the restriction to invariant diver-

³⁵The proof uses the monotonicity property of the Fisher information metric, and is closely related to the Cramér-Rao bound.

gence cost functions is almost without loss of generality, because all f -divergences are invariant, and the associated f -function can be any smooth, convex function. The lower bound even applies to problems where the standard “first-order-approach” is not valid (Jewitt (1988); Mirrlees (1999)). Although in general parametric moral hazard sub-problems do not have unique or everywhere-differentiable solutions, in the neighborhood of $\theta^{-1} \rightarrow 0^+$ the policy function $p(\eta; M_\xi, \theta^{-1})$ becomes unique and differentiable for all smooth manifolds M_ξ . The lower bound can therefore be thought of as a tractable approach to approximating problems that are otherwise difficult to analyze.

The lower bound is tight, in the sense that there exists an M_ξ such that the utility difference is equal to the mean-variance objective, up to order θ^{-1} . The example that illustrates this provides an interesting economic intuition. Take the security s and corresponding retained tranche η as given. Suppose that M_ξ is an exponential family,

$$p^i(\xi) = q^i \exp(\xi_1 \eta_i + \xi_2 s_i - A(\xi)),$$

where η and s are the sufficient statistics of the distribution³⁶, and $A(\xi)$ is the log-partition function that ensures $p^i(\xi)$ is a probability distribution for all ξ . The zero-cost distribution, q^i , corresponds to $\xi_1 = \xi_2 = 0$. This example is a “worst-case scenario” for the agents, over the set of possible actions, holding the security design fixed. It is the worst case because the exponential family can also be expressed as a function of its dual coordinates, τ , with $\tau_1(\xi) = \sum_i p^i(\xi) \eta_i$ and $\tau_2(\xi) = \sum_i p^i(\xi) s_i$ (Amari and Nagaoka (2007)). In effect, the seller separately controls the value of her retained tranche and the security. The seller will use only the action that increases the value of her retained tranche (τ_1) and will not use the action that increases the value of the security (τ_2). Relative to the case where the seller retained the entire asset, this situation exhibits both reduced effort and risk-shifting.

³⁶Informally, the sufficient statistics capture all of the information in a data sample. In this case, if an observer saw the value of both the retained tranche and security, there would be no additional information that would be useful when trying to infer the seller’s actions.

The lower bound in this example is tight, and there is no “costless variation,” because the restriction of the seller’s actions to this particular parametric family did not alter the probability distribution $p(\eta)$ that she would choose, relative to the non-parametric model. The proof of proposition 1.5 relies on the fact that the security s is a sufficient statistic of the distribution, and that sufficient statistics (as estimators) attain the Cramér-Rao bound.

It is interesting to note the contrast between the worst-case scenarios for debt securities and equity securities. For debt securities, the worst-case scenario is that the seller controls an option-like payoff, which increases both the mean and variance of the asset value. For equity securities, the worst-case scenario is that the seller controls the mean outcome.

The security design utility in the parametric problem, under the approximation, exhibits a “monotonicity” with respect to the set of feasible distributions M_ξ .³⁷ If $M_\xi \subseteq \hat{M}_\xi$, then for all s , and sufficiently small θ^{-1} and κ ,

$$U(s; M_\xi, \theta^{-1}, \kappa) \geq U(s; \hat{M}_\xi, \theta^{-1}, \kappa).$$

Earlier, I assumed that the seller could influence the mean of the probability distribution. It follows that giving the seller additional actions can only, under the first-order approximation, expand the scope for risk-shifting, and therefore reduce the utility that can be achieved in the security design problem. Eventually, as the set of feasible probability distributions approaches the entire probability simplex, the problem converges to the non-parametric case, and the lower bound is tight for all securities. In this case, as in the previous section, debt is optimal.

The lower bound in proposition 1.3 (the mean and variance terms) does not depend on the parameters ξ , which suggests an interpretation of debt as a robust security design. The robustness of debt is complementary to the result of Carroll (2013). The key difference between this model and Carroll (2013) is the cost of each potential probability distribution, p . In Carroll (2013), each probability p could have different, arbitrary costs. When there is no structure on the cost of potential actions, risk-shifting concerns dominate concerns about reduced effort, and it is crucial that

³⁷See appendix section 5.1.13, which proves this.

the buyer and seller's payoffs be exactly aligned. In this case, the optimal security is equity. By contrast, the invariant divergence cost assumption in this model imposes a structure on the cost that any probability distribution p would have, if that p were feasible. With the cost structure I have imposed, both reduced effort and risk-shifting are potentially important, the mean-variance intuition applies, and debt is approximately optimal. Together, these theories can explain the prevalence of both debt and equity securities. Debt securities are used when concerns about reduced effort and risk-shifting are both relevant. If concerns about risk-shifting dominate, equity securities are used.

In this discussion, I have emphasized the application of the parametric framework to models in which the seller controls a small number of parameters. However, we can also use the result to consider an "almost non-parametric" model. Assume that each state $i \in \Omega$ contains information about the idiosyncratic outcome v_i and an aggregate state. In the almost non-parametric model, the seller controls the conditional probability distribution of the asset value for each aggregate state, but does not control the probability distribution of the aggregate states.

If the cost function is the KL divergence, the optimal security design will be an aggregate-state contingent debt security (effectively, there is a separate security design problem for each aggregate state). Nevertheless, the lower bound results of proposition 1.5 apply, and a non-contingent debt security maximizes this lower bound. These results suggest that it is possible to bound the utility losses of using a non-contingent debt security instead of an aggregate-state contingent debt security.

One application of this idea relates to currency choice.³⁸ If we assume that the cost function is a pecuniary cost, it follows that the non-contingent debt security is a debt in the currency of the cost function. For example, if a mortgage originator creates loans in the United States, and sells the loans to a euro zone bank, the security design maximizing proposition 1.5 is a dollar-denominated debt, not a euro-denominated (or yen-denominated, ...) debt. The optimal security design likely depends on aggregate outcomes, including exchange rates, but a non-contingent, dollar-denominated debt may achieve nearly the same utility.

³⁸This interpretation was suggested to me by Roger Myerson.

Finally, in the appendix, section §5.1.4, I provide a micro-foundation for a version of the parametric model, using a model of rational inattention. The rational inattention model motivates my study of parametric models with the KL divergence cost function. The exact solution for the model cannot be characterized analytically, whereas the approximately optimal debt contract is simple to describe.

In the next two sections of the chapter, I will discuss continuous time models of effort. I show that the optimality of debt, and the intuitions about mean-variance tradeoffs, apply in continuous time as well. These sections provide a micro-foundation for the non-parametric static models introduced earlier.

1.6 Dynamic Moral Hazard

In this section, I will analyze a continuous time effort problem. This problem is closely connected to the static models discussed previously. The role of this section is to explain how an agent could “choose a distribution,” and show that the mean-variance intuition and optimality of debt discussed previously apply in dynamic models.

I will study models in which the seller controls the drift of a Brownian motion. The contracting models I discuss are similar to those found in Holmström and Milgrom (1987), Schaettler and Sung (1993), and DeMarzo and Sannikov (2006), among others. The models can be thought of as the continuous time limit of repeated effort models³⁹, in which the seller has an opportunity each period to improve the value of the asset. Two recent papers are particularly relevant. The models I discuss are a special case of Cvitanić et al. (2009). The contribution of this chapter is to note that, in the special case I discuss, debt securities are optimal, and that the reason debt securities are optimal is that the mean-variance intuitions discussed earlier still hold. To demonstrate this, I use the results of Bierkens and Kappen (2012), who study a single-agent control problem (effectively,

³⁹See Biais et al. (2007); Hellwig and Schmidt (2002); Sadzik and Stachetti (2013) for analysis of the relationship between discrete and continuous time models.

the seller's moral hazard problem) and show that it is equivalent to a relative entropy minimization problem.

The timing in these models follows the principal-agent convention discussed earlier. At time zero, the seller and buyer trade a security. Between times zero and one, the seller will apply effort (or not) to change the value of the asset. At time one, the asset value is determined and the security payoffs occur.

Between times zero and one, the seller controls the drift of a Brownian motion. Define W as a Brownian motion on the canonical probability space, $(\Omega, \mathcal{F}, \tilde{P})$, and let \mathcal{F}_t^W be the standard augmented filtration generated by W . Denote the asset value at time t as V_t , and let \mathcal{F}_t^V be the filtration generated by V . The seller observes the history of both W_t and V_t at each time, whereas the buyer observes (or can contract on) only the history of V_t . This information asymmetry is what creates the moral hazard problem.

The initial value, $V_0 > 0$, is known to both the buyer and the seller. The asset value evolves as

$$dV_t = b(V_t, t)dt + u_t \sigma(V_t, t)dt + \sigma(V_t, t)dW_t,$$

where $b(V_t, t)$ and $\sigma(V_t, t) > 0$ satisfy standard conditions to ensure that, conditional on $u_t = 0$ for all t , there is a unique, everywhere-positive solution to this SDE.⁴⁰ The seller's control, u_t , should be thought of as effort. Let \mathcal{U} be the set of admissible effort strategies, which I will define shortly. Effort is costly. There is a flow cost of effort, a general form of which is $g(t, V_t, u_t)$. The function $g(\cdot)$ is weakly positive, twice-differentiable, and strictly convex in effort. For all t and V_t , $g(t, V_t, 0) = 0$. Effort always improves the expected value of the asset. For all t and V_t , the expectation of V_1 is increasing in u_s , for all $s > t$.⁴¹

⁴⁰The following conditions are sufficient. For all $V \in \mathbb{R}^+$ and $t \in [0, 1]$, $\sigma(V, t) > 0$ and $|b(V, t)| + |\sigma(V, t)| \leq C(1 + |V|)$ for some positive constant C . For all $t \in [0, 1]$, $V, V' \in \mathbb{R}^+$, $|b(V, t) - b(V', t)| + |\sigma(V, t) - \sigma(V', t)| \leq D|V - V'|$, for some positive constant D . For all $t \in [0, 1]$, $\lim_{v \rightarrow 0^+} \sigma(v, t) = 0$, and $\lim_{v \rightarrow 0^+} b(t, v) \geq 0$.

⁴¹This assumption places restrictions on the functions $b(V_t, t)$ and $\sigma(V_t, t)$.

The retained tranche, $\eta(V)$, is an \mathcal{F}_1^V -measurable random variable that can depend on the entire path of the asset value. I continue to assume limited liability, meaning that $\eta(V) \in [0, \beta_s V_1]$ for all paths V . The seller's indirect utility function can be written as

$$\phi_{CT}(\eta) = \sup_{\{u_t\} \in \mathcal{U}} \phi_{CT}(\eta; \{u_t\}) = \sup_{\{u_t\} \in \mathcal{U}} \{E^{\tilde{P}}[\eta(V)] - E^{\tilde{P}}[\int_0^1 g(t, V_t, u_t) dt]\},$$

where $E^{\tilde{P}}$ denotes the expectation at time zero under the physical probability measure. Alternatively, if the buyer and seller were risk-averse, but shared a common risk-neutral measure \tilde{P} , the problem would be identical. The key assumption in that case would be that the problem is small, in the sense that the outcome of this particular asset and security does not alter the common risk-neutral measure. To guarantee that utility is finite, I will make some additional assumptions, which I will explain shortly.

The security design problem is almost identical to the security design problem in the static models. The seller internalizes the effect that the security design will have on the price that the buyer is willing to pay, and solves

$$U_{CT}(s^*) = \sup_{s \in \mathcal{S}} \{\beta_b E^{\tilde{P}}[s(V)] + \phi_{CT}(\eta(s))\},$$

where \mathcal{S} is the set of \mathcal{F}_1^V -measurable limited liability security designs. Here, as in the static models, $\eta(V) = \beta_s(V_1 - s(V))$.

This model, which is a version of the one discussed by Schaettler and Sung (1993) and Cvitanić et al. (2009), uses the “strong” formulation of the moral hazard problem. There is a second approach to the moral hazard problem, known as the “weak” formulation, discussed by those authors. The weak formulation is closely related to the static problems discussed earlier, and is equivalent to the strong formulation for the purpose of determining the optimal security design.

In the weak formulation, let X_t be a stochastic process that evolves as

$$dX_t = b(X_t, t)dt + \sigma(X_t, t)dB_t,$$

where B is a Brownian motion on the probability space (Ω, \mathcal{F}, Q) , with standard augmented filtration \mathcal{F}_t^B , and the $b(\cdot)$ and $\sigma(\cdot)$ are identical to the functions discussed above. Consider an \mathcal{F}_t^B -adapted control strategy u_t such that the stochastic exponential, $Z_t = \exp(\int_0^t u_s dB_s - \frac{1}{2} \int_0^t u_s^2 ds)$, is a martingale. For such a strategy, Girsanov's theorem holds, and we can define a measure, P , that is absolutely continuous with respect to Q , such that $\frac{dP}{dQ} = Z_1$. Under the measure P , $B_t^P = B_t - \int_0^t u_s ds$ is a Brownian motion, and the process X_t evolves as

$$dX_t = b(X_t, t)dt + u_t \sigma(X_t, t)dt + \sigma(X_t, t)dB_t^P.$$

Under this measure P , the stochastic process X has the same law as the asset price V does under measure \tilde{P} in the strong formulation of the problem. Because the assumed effort strategies are \mathcal{F}_t^B -adapted, they can be written as functionals $u_t = u(X, t)$. For any such control strategy, it follows that

$$\phi_{CT}(\eta; \{u_t\}) = E^P[\eta(X)] - E^P\left[\int_0^1 g(t, X_t, u(X, t))dt\right],$$

meaning that the indirect utility in the strong formulation is the same as the indirect utility in the weak formulation. Here, I interpret X as the asset price, and require that η be $\mathcal{F}_1^X = \mathcal{F}_1^B$ -measurable. By the same logic, $E^{\tilde{P}}[s(V)] = E^P[s(X)]$ under these effort strategies.

I define the set of admissible strategies \mathcal{U} as the set of \mathcal{F}_t^B -adapted, square-integrable controls such that $E[Z_1^4] < \infty$ (similar to the assumptions of Cvitanić et al. (2009)). Defining the set \mathcal{U} in this way ensures that Girsanov's theorem can be applied. It follows that

$$\phi_{CT}(\eta) = \sup_{\{u_t\} \in \mathcal{U}} \{E^P[\eta(X)] - E^P\left[\int_0^1 g(t, X_t, u_t)dt\right]\},$$

meaning that the optimal strategies in the weak and strong formulations, assuming they exist and are unique, are identical.

The restriction that u_t be \mathcal{F}_t^B -adapted has economic meaning. The filtrations generated by B_t and X_t are identical, and X_t is interpreted as the asset price. This restriction states that the seller's

control strategy does not depend on the seller's private information (B^P) but only the seller and buyer's common knowledge, X (and t). Under the form of the flow cost function I have assumed, the optimal strategy will satisfy this restriction, even if it is not imposed (Cvitanic et al. (2009)).

The weak formulation can be rewritten to emphasize the idea of “choosing a distribution.” The expected values of the security and the retained tranche depend only on the Radon-Nikodym derivative $\frac{dP}{dQ}$, and not directly on the control strategy.

Define the set M as the set of measures P that are absolutely continuous with respect to Q , and for whom $E^Q[(\frac{dP}{dQ})^4] < \infty$. Bierkens and Kappen (2012) (and the sources cited therein) show that any $P \in M$ can be created through some control strategy $\{u_t\}$. I combine results found in Cvitanic et al. (2009) and Bierkens and Kappen (2012) into the following lemma:

Lemma 1.2. *For any effort strategy $u(X, t) \in \mathcal{U}$, the stochastic exponential $Z_t = \exp(\int_0^t u(X, s)dB_s - \frac{1}{2} \int_0^t u(X, s)^2 ds)$ is an everywhere-positive martingale, and the measure defined by $\frac{dP}{dQ} = Z_1$ is a measure in M . Conversely, for any measure $P \in M$, there exists an effort strategy $u(X, t) \in \mathcal{U}$ such that $\frac{dP}{dQ} = \exp(\int_0^1 u(X, s)dB_s - \frac{1}{2} \int_0^1 u(X, s)^2 ds)$. The effort strategy $u(X, t)$ is unique up to an evanescence.*

Proof. See appendix, section 5.1.16. □

We can define a divergence,

$$D_g(P||Q) = \inf_{\{u_t\} \in \mathcal{U}} E^P[\int_0^1 g(t, X_t, u_t) dt],$$

subject to the constraint that $\frac{dP}{dQ} = \exp(\int_0^1 u_s dB_s - \frac{1}{2} \int_0^1 u_s^2 ds)$. By the uniqueness result in lemma 1.2, all strategies $\{u_t\} \in \mathcal{U}$ that satisfy this constraint are identical for our purposes. Note that, because $g(t, X_t, u_t) = 0$ if and only if $u_t = 0$, and is otherwise positive, $D_g(P||Q)$ satisfies the definition of a divergence.

The moral hazard can be written as

$$\phi_{CT}(\eta) = \sup_{P \in M} \{E^Q[\frac{dP}{dQ} \eta(X)] - D_g(P||Q)\}.$$

I have now rewritten the continuous time moral hazard problem as a static problem, in which the seller chooses a probability measure subject to a cost that is described by a divergence. In light of the results for static models, two questions immediately arise. First, is there a $g(\cdot)$ function such that $D_g(P||Q)$ is the Kullback-Leibler divergence? Second, are there $g(\cdot)$ functions such that $D_g(P||Q)$ is an invariant divergence?

The answer to the first question comes from the work of Bierkens and Kappen (2012) and the sources cited therein, who show that quadratic costs functions, $g(t, X_t, u_t) = \frac{\theta}{2}u_t^2$, lead to the KL divergence.⁴² In the remainder of this section, I will analyze quadratic cost models, and show that debt is the optimal security design. In the next section, I will explore the question of whether there are $g(\cdot)$ functions that lead to invariant divergences.

As mentioned earlier, some additional assumptions are necessary to ensure that utility in the moral hazard problem is finite. With quadratic costs, it is sufficient to assume that $E^Q[\exp(4\theta^{-1}X_1)] < \infty$ (see Cvitanic et al. (2009)). This should be thought of as a restriction on the functions $b(\cdot)$ and $\sigma(\cdot)$ discussed previously. For a given $P \in M$, the KL divergence can be defined as

$$D_{KL}(P||Q) = E^P[\ln(\frac{dP}{dQ})].$$

For any $u(X, t) \in \mathcal{U}$ such that $\frac{dP}{dQ} = \exp(\int_0^1 u_s dB_s - \frac{1}{2} \int_0^1 u_s^2 ds)$, it follows that

$$\begin{aligned} D_{KL}(P||Q) &= E^P[\int_0^1 u(X, s) dB_s - \frac{1}{2} \int_0^1 u(X, s)^2 ds] \\ &= E^P[\int_0^1 u(X, s) dB_s^P + \frac{1}{2} \int_0^1 u(X, s)^2 ds] \\ &= E^P[\frac{1}{2} \int_0^1 u(X, s)^2 ds]. \end{aligned}$$

For quadratic cost functions $g(t, X_t, u_t) = \frac{\theta}{2}u_t^2$, we have $D_g(P||Q) = \theta D_{KL}(P||Q)$. The moral hazard problem can be rewritten as

⁴²Note that this formulation rules out time discounting of the effort costs. One possible justification for this assumption is that the time period in question is short.

$$\phi_{CT}(\eta) = \sup_{P \in \mathcal{M}} \{E^P[\eta(X)] - \theta D_{KL}(P||Q)\},$$

and the security design problem as

$$U_{CT}(s^*) = \sup_{s \in \mathcal{S}} \{\beta_b E^{P^*(\eta(s))}[s(X)] + \phi_{CT}(\eta(s))\},$$

where $P^*(\eta)$ is the measure chosen in the moral hazard problem. The model is a continuous sample space version of the non-parametric model discussed earlier. It is not surprising that debt is the optimal security design.

Proposition 1.6. *In the continuous time model, with the quadratic cost function, the optimal security design is a debt contract,*

$$s_j(\eta^*) = \min(v_j, \bar{v}),$$

for some $\bar{v} > 0$.

Proof. See appendix, section 5.1.15. The result is a special case of Cvitanić et al. (2009). □

Debt is the optimal security design in the continuous time model for same reasons it is optimal in the non-parametric model: it minimizes the variance of the security payout, holding fixed the mean value. The mean-variance intuition arises from concerns about reduced effort and risk-shifting, balanced against the gains from trade. Risk-shifting is possible in this model because of the seller can selectively alter her effort level, based on the path that the asset value has taken up to the current time. The surprising result of proposition 1.6 is that this sort of risk-shifting is equivalent to the risk-shifting in the static problem discussed earlier.

The intersection of these results with Holmström and Milgrom (1987) is intuitive. In the principal-agent framework, when the asset value Ito process is an arithmetic Brownian motion and the flow cost function is quadratic, without limited liability, a constant security for the prin-

cial is optimal. With limited liability, in the security design framework, optimal security simply reduces the constant payoff where necessary, and debt is optimal.

The debt security design may or may not be renegotiation-proof. Suppose that at some point, say time 0.5, the seller can offer the buyer a restructured security. Assume that at this time, there are no gains from trade.⁴³ If the current asset value is low enough, the debt security provides little incentive for the seller to continue putting in effort in the future. In this state, the buyer might agree to “write down” the debt security, even though he cannot receive any additional payments from the seller, because the buyer’s gains from increased effort by the seller could more than offset the loss of potential cash flows. In this model, write-downs can be Pareto-efficient if the time-zero expected value of the debt, $E^P[s(X)]$, is greater than θ .⁴⁴ Write-downs will never be Pareto-efficient when κ and θ^{-1} are both small, but could occur if both the gains from trade at time zero and the moral hazard were large.

One possible extension of this model concerns the timing of the seller’s actions and payment. In the context of mortgage origination, it might be reasonable to assume that the mortgage originator (seller) controls the drift of V_t from time zero to time one, when she is making the mortgage loans. However, the mortgage security might be constrained to pay both the originator and security buyers at time $T > 1$, and the drift of V_t between times one and T is not controlled by the originator. It turns out that this version of the model is related to the parametric problems discussed earlier. We can still think of the seller as controlling the time T Radon-Nikodym derivative, $\frac{dP}{dQ}$, but subject to the constraint that the derivative be \mathcal{F}_1^B -measurable. One consequence of the invariance property of the KL divergence is that, if $\frac{dP}{dQ}$ is \mathcal{F}_1^B -measurable, then the KL divergence between P and Q , projected onto \mathcal{F}_1^B , is the same as the KL divergence evaluated at time T . Similar to the parametric models discussed earlier, the seller cannot choose any distribution over paths from time zero to time T , but instead only a subset of all such distributions. Using arguments from that section, one can show that, under the approximation used earlier, a mean-variance tradeoff provides a lower bound on the utility any security design can achieve, and that debt maximizes this lower bound.

⁴³Otherwise, if the asset value has increased, the seller will “lever up” and sell more debt to the buyer.

⁴⁴Proof available upon request.

In the final section of the chapter, I will discuss the second question mentioned earlier. Are there cost functions $g(\cdot)$ for which $D_g(P||Q)$ is an invariant divergence?

1.7 A Mean-Variance Approximation for Continuous Time Models

For the static models discussed earlier, invariant divergence cost functions lead to models in which debt was approximately optimal. I will not directly answer the question of whether there functions $g(\cdot)$ such that $D_g(P||Q)$ is invariant. Instead, I will show that for all $g(t, X_t, u_t) = \theta \psi(u_t)$, where $\psi(u_t)$ is a convex function, debt is approximately optimal. The approximations used in this section are identical to the ones discussed previously, in section §1.4. I consider problems in which both the moral hazard and gains from trade are small, relative to the scale of the assets. I show that the utility of arbitrary security designs can be characterized, to first-order, by a mean-variance tradeoff.

In this section, I will continue to analyze the models introduced in the previous section, with several small modifications. I will assume that the control is bounded, $|u_t| \leq \bar{u}$. This automatically ensures that all \mathcal{F}_t^B -adapted processes u_t are in \mathcal{U} , and simplifies the discussion of integrability conditions. I assume ψ satisfies the conditions required for g in the previous section, and in addition that for all $|u| \leq \bar{u}$, $\psi''(u) \in [K_1, K_2]$ for some positive constants $0 < K_1 < 1 < K_2$. I also normalize $\psi''(0) = 1$. I assume that, for bounded control strategies $|u_t| \leq \bar{u}$, the asset value has a finite fourth moment. That is, $E^P[X_1^4] < \infty$ for all P for created by a bounded control strategy.

There is a sense in which any convex cost function $\psi(u_t)$ resembles the quadratic cost function, as u_t becomes close to zero, because their second derivatives are the same. Similarly, in static models, all invariant divergences resemble the KL divergence, because their Hessian is the Fisher information matrix. I apply this idea to the divergences D_ψ induced by the convex cost functions ψ .

First, I define the Fisher information. Consider a model with a finite number of parameters, τ , and associated probability distribution $p(\omega; \tau)$, over sample space Ω . The Fisher information matrix defined as

$$I_{ij} = E^{p(\omega; \tau)} \left[\left(\frac{\partial \ln p(\omega; \tau)}{\partial \tau_i} \right) \left(\frac{\partial \ln p(\omega; \tau)}{\partial \tau_j} \right) \right].$$

I will define the Fisher information for the continuous time model, which has an infinite number of parameters, in an analogous way. Let $\frac{dP}{dQ}(\gamma, \tau) = \exp(\int_0^1 (\gamma u_s + \tau v_s) dB_s - \frac{1}{2} \int_0^1 (\gamma u_s + \tau v_s)^2 ds)$, where γ and τ parametrize a perturbation in the direction defined by the square-integrable, predictable processes u_s and v_s .⁴⁵ The Fisher information, in the directions defined by u_s and v_s , is defined as

$$I(u, v) = E^{P(\gamma, \tau)} \left[\left(\frac{\partial}{\partial \gamma} \ln \left(\frac{dP}{dQ}(\gamma, \tau) \right) \right) \left(\frac{\partial}{\partial \tau} \ln \left(\frac{dP}{dQ}(\gamma, \tau) \right) \right) \right]_{\gamma=\tau=0}.$$

Plugging in the definition of $\frac{dP}{dQ}(\gamma, \tau)$, and using the Ito isometry, $I(u, v) = E^Q[\int_0^1 u_s v_s ds]$. Next, consider the second variation of $D_\psi(P(\gamma, \tau)||Q)$,

$$\begin{aligned} \frac{\partial}{\partial \tau} \frac{\partial}{\partial \gamma} D_\psi(P(\gamma, \tau)||Q)|_{\gamma=\tau=0} &= \theta \frac{\partial}{\partial \tau} \frac{\partial}{\partial \gamma} E^{P(\gamma, \tau)} \left[\int_0^1 \psi(\gamma u_s + \tau v_s) ds \right]_{\gamma=\tau=0} \\ &= \theta E^Q \left[\int_0^1 \psi''(0) u_s v_s ds \right]. \end{aligned}$$

It follows that $\frac{\partial}{\partial \tau} \frac{\partial}{\partial \gamma} D_\psi(P(\gamma, \tau)||Q)|_{\gamma=\tau=0} = \theta I(u, v)$.⁴⁶ For any cost function ψ , the second variation in the directions u_s and v_s is equal to the Fisher information in those directions. The divergences D_ψ resemble, locally, the KL divergence, in exactly the same way that all invariant divergences resemble the KL divergence.

I consider the same approximation discussed earlier, in which both θ^{-1} and κ are small.⁴⁷ As the cost of effort rises, the seller will choose to respond less and less to the incentives provided by the retained tranche. Regardless of the cost function ψ , the divergence $D_\psi(P||Q)$ will approach $D_{KL}(P||Q)$, and debt will be approximately optimal. To make this argument rigorous, I use Malliavin calculus in a manner similar to Monoyios (2013) to prove the following theorem:

⁴⁵These perturbations are the Cameron-Martin directions used in Malliavin calculus.

⁴⁶I am omitting several technicalities from this discussion. These technicalities are addressed in the proof of proposition 1.7.

⁴⁷The limit I consider is related to the ‘‘large firm limit’’ discussed by Sannikov (2013).

Proposition 1.7. *For any limited liability security design s , the difference in utilities achieved by an arbitrary security s and the sell-nothing security is*

$$U(s; \theta^{-1}, \kappa) - U(0; \theta^{-1}, \kappa) = \kappa E^Q[\beta_s s] - \theta^{-1} \frac{1}{2} V^Q[\beta_s s] + O(\theta^{-2} + \theta^{-1} \kappa).$$

Proof. See appendix, section 5.1.17. □

In the continuous time effort problem with an arbitrary convex cost function, debt securities are first-order optimal. The same mean-variance intuition that I discussed in static models applies to continuous time models. The variance of the security payoff is again a summary statistic for the problems of reduced effort and risk shifting associated with the moral hazard problem.

1.8 Conclusion

In this chapter, I offer a new explanation for the prevalence of debt contracts. Debt arises as the solution to a moral hazard problem with both effort and risk-shifting. Debt is exactly optimal in the non-parametric model, when the cost function is the KL divergence, for which I provide a continuous-time micro-foundation. Debt is approximately optimal for all invariant divergence cost functions, again because it is the mean-variance optimal contract. For parametric models, such as the rational inattention model, debt maximizes an approximate lower bound on the level of utility. In all of these results, debt is desirable because it minimizes the variance of the security payout, balancing the need to provide incentives for effort, minimize risk-shifting, and maximize trade. Taken together, these results offer a new explanation for why debt is observed in a wide range of settings.

2 Generalized Rational Inattention

2.1 Introduction

The rational inattention theory, surveyed by Sims (2010), has found applications in a wide range of economic problems. There are two core elements of the theory: first, that an agent faced with a decision making problem can choose an arbitrary signal structure to gather information prior to making a decision, and second, that the cost of each possible signal structure is described by the mutual information cost function. The mutual information cost function plays a central role in information theory (Cover and Thomas (2012)). However, it is not clear that the axioms used to derive mutual information, in the context of information theory, are appropriate for many economic applications. In particular, the mutual information cost function imposes a type of symmetry across different states of nature, so that it is equally easy or difficult for an agent to learn about two equally probable states.

In this chapter, we study alternative rational inattention cost functions, other than mutual information, that do not share this property. The cost functions we study capture the idea that it may be particularly easy or difficult to learn about certain states, and that there might be complementarities or substitutabilities in learning about different states of the world. We show that the mutual information cost function has a special structure, which is not generic, and that the rational inattention model can be solved with many other cost structures. We then apply our methods to a delegated portfolio management problem. This problem illustrates that, despite the generality of our theoretical framework, it is possible to derive concrete predictions about an agent's behavior.

We derive the properties of the cost functions we study from several axioms. These axioms are, for the most part, standard, and based on existing work in the literature (De Oliveira et al. (2013), Caplin and Dean (2014)). The critical axiom for our purposes is an axiom related to Blackwell's ordering (Blackwell (1953)). The axiom states that if one signal structure Blackwell-dominates another signal structure, it must be weakly more costly. Blackwell's ordering is incomplete; most

signal structures do not Blackwell-dominate each other. Nevertheless, the axiom enables us to characterize the local properties of the cost functions we study.

The key observation is that any cost function respecting Blackwell's ordering exhibits certain type of invariance. If one signal structure is equivalent to another, in the sense that the two signal structures convey the same information, expressed in different alphabets, then the two signal structures must be equally costly. This sort of invariance may sound trivial, but in fact it places strong restrictions on the local properties that any rational inattention cost function must have. We are able to characterize, up to a second order approximation, the entire class of cost functions satisfying our axioms.

For any particular state, the cost of gathering a small amount of information is approximately proportional to the Fisher information of the signal in that state. The constants of proportionality associated with this cost function, for each state, form the diagonal of a matrix, defined over pairs of states, that is central to our theory. We refer to this matrix as the "information cost matrix." The off-diagonal elements of the information cost matrix characterize the complementarities or substitutabilities between gathering information in different states of nature. For the mutual information cost function, this matrix is also the Fisher information matrix. This illustrates the essential role that the Fisher information plays in our analysis. Our axioms agree with the mutual information cost function, in the sense that, holding the state fixed, the second derivative of the cost function is proportional to the Fisher information matrix. However, the information cost matrix that characterizes relative costs across states does not need to be the Fisher matrix; there are many other matrices that would satisfy our axioms.

Armed with this result, we solve a static rational inattention problem. The behavior of the agent in the static model depends on the details of the information cost matrix, and in general may not resemble the behavior of an agent with the mutual information cost function. The properties of rationally inattentive behavior with mutual information (the invariant likelihood ratio property and the locally invariant posteriors property described in Caplin and Dean (2013)) do not apply in the general case. Our solution uses an approximation, which is that the agent's ability to gather infor-

mation is limited, or equivalently that the stakes for the agent are small relative to their information processing capabilities. This approximation ensures that our characterization of the second-order properties of the information cost function is sufficient to determine the agent's behavior.

The dependence of the agent's behavior on the information cost matrix opens the door to a wide variety of predicted behaviors. The matrix could be identified in experimental settings (as in Caplin and Dean (2014)), but the external validity of such experiments is an open question. The information cost matrix might differ across agents, economic environments, and time. To circumvent these concerns, we develop a dynamic model, which illustrates that it is possible to derive sharp predictions using our general framework.

We develop a model of delegated portfolio management, in which the agent (a portfolio manager, trader, or the like) repeatedly both gathers information and selects an asset allocation. We use the framework developed in the first two sections of the chapter to model the information acquisition costs of the agent. We show that this leads to a stochastic process for the assets under management that is simple to characterize, and (crucially) is not sensitive to the details of the information cost matrix. The stochastic process resembles the one described in Woodford (2014), and the results in Cvitanić et al. (2009) and Hébert (2014) allow us solve for the agent's value function and optimal behavior.

This chapter builds on the rational inattention literature, surveyed in Sims (2010). In its use of axioms to characterize general rational inattention cost functions, it is particularly close to Caplin and Dean (2013), Caplin and Dean (2014), and De Oliveira et al. (2013). The Chentsov (1982) theorems used to characterize the properties of general rational inattention cost functions were also used by Hébert (2014), in a different context. We also use techniques developed by Kamenica and Gentzkow (2011) and Matejka and McKay (2011). We develop a repeated version of the rational inattention model, and consider its continuous time limit, as in Woodford (2014). Our application to dynamic asset allocation and delegated portfolio management is related to the literature on that subject, surveyed by Stracca (2006). The application is similar in spirit to the model of Dybvig

et al. (2010). Our application is closely related to Van Nieuwerburgh and Veldkamp (2010), who also consider general learning technologies in the context of asset allocation.

Section 2.2 discusses our axioms and main theorem, characterizing the local structure of rational inattention cost functions. Section 2.3 solves a static model of rational inattention. Section 2.4 discusses the dynamic asset allocation application, which is a repeated version of the static model. Section 2.5 concludes.

2.2 Information Costs

The rational inattention model of Sims (2010) studies decision making by an agent who faces constraints on the quantity of information she can process. To describe the quantity of information contained in a particular signal, Sims uses the mutual information between that signal and the underlying state. Mutual information is widely used in information theory, to describe noisy communication channels (Cover and Thomas (2012)). In the context of the noisy channel coding theorems, it is the only “correct” measure of the quantity of information transmitted (Csiszár (1972)). Sims developed the rational inattention theory to model an agent with cognitive constraints on information processing, and made reference to these coding theorems. However, it is not clear that mutual information is the only reasonable way to model cognitive constraints. Moreover, many economists have subsequently used mutual information to model the “physical” cost of acquiring information, and the relevance of the coding theorems is not clear in this context.

Despite these caveats, the rational inattention theory and mutual information have found many applications in economics. The mutual information cost function has several properties that are desirable in these applications. Mutual information is also well-defined, in the sense that an economist can take the framework and apply to it new problems without substantial modification. In this section, we discuss a large class of information costs, which include mutual information. These information costs are characterized by a set of conditions that are relevant or useful in economic applications. We show that all of the information costs in this class are approximately the same, in a sense, for signals that convey very little information. In the subsequent sections, we will

analyze dynamic problems, in which very little information is gathered each period. We will rely on these approximations to generate predictions, regardless of which information cost from this class is relevant.

We begin by describing the rational inattention framework of Sims (2010). Let $x \in X$ be the underlying state of the nature, and $a \in A$ be the action taken by the agent. The agent's utility from taking action a in state x is $u(x, a)$. However, the agent does not perfectly observe the state x . Instead, the agent will receive a signal, $s \in S$, which can convey information about the state. For simplicity, we assume that X , A , and S are finite sets. We define $\mathcal{P}(\Omega)$ as the probability simplex over finite set Ω .

Let $q(x)$ denote the agent's prior belief (before receiving a signal) about the probability of state x . Define $p(s|x)$ as the probability of receiving signal s in state x . The set of conditional probability distributions $\{p(s|x)\}$ define a "signal structure." After receiving signal s , the agent will take an action, consistent with Bayes' rule,

$$a^*(s; \{p(s|x)\}) \in \arg \max_a \sum_x \frac{q(x)p(s|x)}{\sum_{x' \in X} q(x')p(s|x')} u(x, a). \quad (2.1)$$

The rational inattention problem maximizes the expected utility of the agent over possible signal structures, taking into account the cost of each signal structure. For now, take the alphabet of signals, S , as given.⁴⁸ When the cost of a particular signal structure is proportional to mutual information, with constant of proportionality $\theta > 0$, the agent solves

$$\max_{\{p(s|x) \in \mathcal{P}(S)\}} \sum_{x \in X} q(x) \sum_{s \in S} p(s|x) u(x, a^*(s; \{p(s|x)\})) - \theta I(q(x), \{p(s|x)\}). \quad (2.2)$$

Here, $I(\cdot)$ denotes the mutual information between the signal and the state.

There are several equivalent definitions of mutual information. For our purposes, the most convenient formulation uses the Kullback-Leibler divergence, also known as relative entropy. The KL divergence between a probability distribution $p(s)$ and another probability distribution $r(s)$ is

⁴⁸A standard result in this literature is that it is without loss of generality to define S as equal to the set of possible actions, A .

defined as

$$D_{KL}(p||r) = \sum_{s \in S} p(s) \ln\left(\frac{p(s)}{r(s)}\right).$$

The mutual information between the signal s and the state x can be defined as

$$I(q(x), \{p(\cdot|x)\}) = \sum_{x \in X} q(x) D_{KL}(p(\cdot|x) || \sum_{x' \in X} q(x') p(\cdot|x')).$$

Mutual information has several properties worth noting. First, the cost of a particular signal structure $\{p(s|x)\}$ depends on the prior. Two rationally inattentive agents with different priors about the state would generically face different costs when observing the same signal structure. Second, the cost does not depend on the economic meaning of the states X , only the cardinality of the set. If two states x and x' are equally likely under the prior, swapping the conditional distribution of signals $p(s|x)$ and $p(s|x')$ does not change the mutual information.

We consider generalized information cost functions which do not necessarily share these properties. Define the cost of a signal structure as $C(\{p(s|x)\})$, which may depend on the definition of the states, X , the signal alphabet S , and the agent's prior, $q(x)$. We assume four conditions that characterize the family of information cost functions we consider. All of these conditions are satisfied by mutual information, but also by many other cost functions. The first three conditions characterize “canonical” rational inattention cost functions, in the terminology of De Oliveira et al. (2013).

Condition 1. Signal structures that convey no information ($p(s|x) = p(s|x')$ for all $x, x' \in X$) have zero cost. All other signal structures have a cost greater than zero.

This condition ensures that the least costly strategy for the agent is to acquire no information, and make her decision based on the prior. The requirement that gathering no information has zero utility cost is a normalization.

Condition 2. The cost of information is convex in $p(s|x)$ for each $x \in X$. That is, for all $\lambda \in (0, 1)$,

$$C(\{\lambda p_1(s|x) + (1 - \lambda)p_2(s|x)\}) \leq \lambda C(\{p_1(s|x)\}) \\ + (1 - \lambda)C(\{p_2(s|x)\}).$$

This condition is useful as an economic assumption because it encourages the agent to minimize the number of distinct signals employed. It also helps ensure that there is a unique signal structure that solves the rational inattention problem.

The next condition uses Blackwell's ordering. Consider two signal structures, $\{p_1(s|x)\}$, with signal alphabet S , and $\{p_2(s'|x)\}$, with alphabet S' . The first signal structure Blackwell dominates the second signal structure if, for all utility functions $u(a, x)$ and all priors $q(x)$,

$$\sup_{a(s)} \sum_{x \in X} \sum_{s \in S} q(x) p_1(s|x) u(a(s), x) \geq \sup_{a(s')} \sum_{x \in X} \sum_{s' \in S'} q(x) p_2(s'|x) u(a(s'), x).$$

If signal structure $\{p_1(s|x)\}$ Blackwell dominates $\{p_2(s'|x)\}$, it is weakly more useful for every decision maker, regardless of that decision maker's utility function and prior. In this sense, it conveys weakly more information. This ordering is incomplete; most signal structures do not dominate each other in this sense. However, when a signal structure does Blackwell dominate another signal structure, we assume the dominant signal structure is more costly.

Condition 3. If the signal structure $\{p_1(s|x)\}$ for $s \in S$ is more informative, in the Blackwell sense, than $\{p_2(s'|x)\}$, for $s' \in S'$, then

$$C(\{p_1(s|x)\}) \geq C(\{p_2(s'|x)\}).$$

The fourth condition we assume, which is not imposed by De Oliveira et al. (2013), is a differentiability condition that will allow us to characterize the local properties of our cost functions.

Condition 4. The information cost function is continuously twice-differentiable in $p(s|x)$, for each $x \in X$.

From these four conditions, we derive a result about the second-order properties of the cost function. The first three conditions are innocuous, in the sense that, for any stochastic choice data, there is a cost function satisfying those properties that is consistent with conditions 1-3 (theorem 2 of Caplin and Dean (2014)). Condition 4 is not completely general; for example, it rules out the case in which the agent is constrained to use only signals in a parametric family of probability distributions, and the cost of other signal distributions is infinite⁴⁹. Mutual information, as mentioned above, satisfies each of these four conditions. However, it is not the only cost function to do so. To introduce these other cost functions, it is first useful to recall Blackwell’s theorem.

Theorem 2.1. (*Blackwell (1953)*) *If, and only if, the signal structure $\{p_1(s|x)\}$ for $s \in S$ is more informative, in the Blackwell sense, than $\{p_2(s'|x)\}$, for $s \in S'$, then there exists a Markov transition matrix $\Pi(s',s)$ such that*

$$p(s'|x) = \sum_{s \in S} \Pi(s',s)p(s|x), \tag{2.3}$$

for all $x \in X$.

This Markov transition matrix is known as the “garbling” matrix. Another way of interpreting condition 3 is that garbled signals are (weakly) less costly than the original signal.

There are certain kinds of garbling matrices that don’t really garble the signals. These garbling matrices have left inverses that are also Markov transition matrices. If we define a signal structure $p(s|x)$, and another signal structure $p(s'|x)$ using one of these matrices, via 2.3, then $p(s|x)$ is more informative than $p(s'|x)$, but $p(s'|x)$ is also more informative than $p(s|x)$. It follows that the cost of these two signal structures must be equal, by condition 3.

These matrices are called Markov congruent embeddings by Chentsov (1982). Chentsov (1982) studied tensors and divergences that are invariant to Markov congruent embeddings (we will say “invariant” for brevity). The KL divergence, used earlier to define mutual information, is invariant.

⁴⁹However, this situation can still be analyzed in our framework. Details available upon request.

Let Π be a Markov congruent embedding from S to S' . For any probability distributions p and r over $\mathcal{P}(S)$,

$$D_{KL}(p||r) = D_{KL}(\Pi p||\Pi r).$$

There are many other invariant divergences. Let $D(\cdot)$ be an arbitrary invariant divergence that is convex in its arguments. We can define an alternative version of mutual information,

$$I_D(q(x), \{p(s|x)\}) = \sum_{x \in X} q(x) D(p(s|x) || \sum_{x' \in X} q(x') p(s|x')).$$

This alternative version satisfies conditions 1-3 above, and is therefore also a “canonical” rational inattention cost function. More generally, there are cost functions that satisfy the conditions above, but cannot be written as the the expectation of a divergence. Examples of such costs functions have been proposed, for the purpose of studying rational inattention problems, by Caplin and Dean (2013). However, our axioms rule out other proposed alternatives, such as the channel capacity constraint⁵⁰ suggested by Woodford (2012).

We proceed by considering properties that are common to all of these information costs. By condition 3, all canonical information costs functions are invariant to Markov congruent embeddings. Let Π be a Markov congruent embedding. It necessarily follows that

$$C(\{p(s|x)\}) = C(\{\Pi p(s|x)\}).$$

The argument is that the signal structure $\{p(s|x)\}$ is more informative than $\{\Pi p(s|x)\}$, but the reverse is also true, due to the existence of a Markov left inverse for Π .

All convex, invariant divergences with continuous Hessian matrices have the property that, at the point $p = r$, their Hessian matrix is proportional to the Fisher information matrix, with some constant of proportionality $c > 0$ (Chentsov (1982)). The KL divergence has this property, and it

⁵⁰The channel capacity constraint does not satisfy the differentiability axiom we impose; however, it may be possible to apply our methods to generalized versions of the channel capacity.

can be interpreted as a statement about the costs of a small amount of information. In states x and x' , if the conditional distribution of signals is close to the unconditional distribution of signals, very little information has been gained about the relative likelihood of x vs. x' .

This property of invariant divergences is a corollary of Chentsov's more general results. Chentsov proved the following results⁵¹:

1. Any continuous function that is invariant over the probability simplex is equal to a constant.
2. Any continuous, invariant 1-form tensor field over the probability simplex is equal to zero.
3. Any continuous, invariant quadratic form tensor field over the probability simplex is proportional to the Fisher information matrix.

These results will allow us to characterize the local properties of rational inattention cost functions, via a Taylor expansion. Consider a signal structure $p(s|x; \varepsilon, \nu) = r(s) + \varepsilon \tau(s|x) + \nu \omega(s|x)$. Here, $\tau(s|x)$ and $\omega(s|x)$ represent informative signal structures, and ε and ν are perturbations in those directions, away from an uninformative signal structure. By condition 1, $C(\{p(s|x; 0, 0)\}) = 0$. The first order term is

$$\frac{\partial}{\partial \varepsilon} C(\{p(s|x; \varepsilon, \nu)\})|_{\varepsilon=\nu=0} = \sum_{x \in X} C_x(\{p(s|x; 0, 0)\}) \cdot \tau(s|x),$$

where C_x denotes the derivative with respect to $p(s|x)$. $C_x(\{p(s|x; 0, 0)\})$, for $r \in \mathcal{P}(S)$, forms a continuous 1-form tensor field over the probability simplex $\mathcal{P}(S)$. By the invariance of $C(\cdot)$, it also follows that C_x is invariant, meaning that

$$C_x(\{p(s|x; 0, 0)\}) \cdot \tau(s|x) = C_x(\{\Pi p(s|x; 0, 0)\}) \cdot \Pi \tau(s|x).$$

Therefore, by Chentsov's theorem, it is equal to zero.

We repeat the argument for the second derivative terms. Those terms can be written as

⁵¹Lemma 11.1, Lemma 11.2, and Theorem 11.1 in Chentsov (1982). See also proposition 3.19 of Ay et al. (2014), who demonstrate how to extend the Chentsov results to infinite sets X and S .

$$\frac{\partial}{\partial \mathbf{v}} \frac{\partial}{\partial \boldsymbol{\varepsilon}} C(\{p(s|x; \boldsymbol{\varepsilon}, \mathbf{v})\})|_{\boldsymbol{\varepsilon}=\mathbf{v}=0} = \sum_{x' \in X} \sum_{x \in X} \langle \boldsymbol{\omega}(s|x') | C_{xx'}(\{p(s|x; 0, 0)\}) | \boldsymbol{\tau}(s|x) \rangle,$$

where the notation $\langle V_1 | h | V_2 \rangle$ denotes the inner product of the tangent vectors V_1 and V_2 with respect to the quadratic form tensor h :

$$\langle \boldsymbol{\omega}(s|x') | C_{xx'}(\{p(s|x; 0, 0)\}) | \boldsymbol{\tau}(s|x) \rangle = \sum_{s \in S}$$

By the invariance of $C(\cdot)$, the quadratic form $C_{xx'}(\cdot)$ is invariant, and therefore is proportional to the Fisher information matrix.

We can define the matrix $k(x, x')$ as the constants of proportionality associated with various terms of the derivative. That is,

$$\frac{\partial}{\partial \mathbf{v}} \frac{\partial}{\partial \boldsymbol{\varepsilon}} C(\{p(s|x; \boldsymbol{\varepsilon}, \mathbf{v})\})|_{\boldsymbol{\varepsilon}=\mathbf{v}=0} = \sum_{x' \in X} \sum_{x \in X} k(x, x') \langle \boldsymbol{\omega}(s|x') | g(r(s)) | \boldsymbol{\tau}(s|x) \rangle.$$

In the case of mutual information, this matrix is the inverse Fisher information matrix,

$$k(x, x') = \delta(x, x')q(x) - q(x)q(x'),$$

where $\delta(x, x')$ is the Kronecker delta.

We are now in position to discuss our approximation of the information cost function. We use Taylor's theorem to approximate the cost function and its gradient up to order Δ (we use Δ because in future sections, we will be looking at small time intervals).

Proposition 2.1. *Suppose that a signal structure $\{p(s|x)\}$ is described by the equation*

$$p(s|x) = r(s) + \Delta^{\frac{1}{2}} \boldsymbol{\tau}(s|x) + o(\Delta^{\frac{1}{2}}).$$

Let $C(\{p(s|x)\})$ be a rational inattention cost function that satisfies conditions 1-4. Then

1. The cost of this signal structure is, for some matrix $k(x, x')$,

$$C(\{p(s|x)\}) = \frac{1}{2}\Delta \sum_{x' \in X} \sum_{x \in X} k(x, x') < \tau(s|x')|g(r(s))|\tau(s|x) > + o(\Delta).$$

2. The gradient of this cost function with respect to $\tau(s|x)$, for a given $x \in X$,

$$\nabla_x C(\cdot) = \Delta^{0.5} \sum_{x' \in X} k(x, x') < \tau(s|x')|g(r(s))| + o(\Delta^{\frac{1}{2}}).$$

3. The matrix $k(x, x')$ is positive semi-definite and symmetric, with a single eigenvector in its null space, and satisfies $\sum_{x'} k(x, x') = 0$.

Proof. See appendix, section 5.2.1. □

The results of theorem 2.1 characterize the cost of a small amount of information, for any rational inattention cost function satisfying our conditions. The theorem substantially restricts the local structure of the cost function, relative to the most general possible alternatives (which would not satisfy our conditions). Potential information structures $\{p(s|x)\}$ can be represented as vectors of dimension $D = (|S| - 1) \times |X|$. Under the assumptions of conditions 1, 2, and 4 (but not the Blackwell's ordering condition, condition 3), the cost function would locally resemble an inner product with respect to a positive semi-definite, $D \times D$ matrix. By imposing condition 3, the results of theorem 2.1 show that we can restrict this matrix to the $k(x, x')$ matrix, an $|X| \times |X|$ matrix. If the agent were only allowed binary signals ($|S| = 2$), this restriction would be trivial. When the agent is allowed to contemplate more general signal structures, the restriction is non-trivial.

In the next section, we analyze a standard rational inattention problem, using this result.

2.3 Static Problems of Rational Inattention

We begin with the standard rational inattention problem, described in equation (2.2). Our plan is to use the approximation we have derived in the previous section to consider a problem in which the utility consequences of the decision are small (or, equivalently, the cost of acquiring information is

large). We parametrize the scale of the utility function by the parameter $\Delta^{\frac{1}{2}}$, and consider the limit as $\Delta \rightarrow 0^+$. Our choice of the notation Δ , and the scaling of the utility by $\Delta^{\frac{1}{2}}$, are motivated by the problem in the next section, in which Δ will have be the length of time period⁵².

The agent solves

$$\max_{\{p(s|x) \in \mathcal{P}(S)\}} \Delta^{\frac{1}{2}} \sum_{x \in X} q(x) \sum_{s \in S} p(s|x) u(x, a^*(s; \{p(s|x)\})) - C(\{p(s|x)\}), \quad (2.4)$$

with the optimal action $a^*(\cdot)$ defined as in equation (2.1). We will use the standard simplification⁵³ that it is without loss of generality to consider a signal alphabet S that is the one-to-one with the set of possible actions A . That is, instead of optimizing over conditional probability distributions $\{p(s|x)\}$, we will optimize over the conditional probabilities of actions, $\{p(a|x)\}$. We assume that the utility function $u(x, a)$, over states $x \in X$ and actions $a \in A$, has full row rank, meaning that no action's payoffs can be perfectly replicated by a linear combination of the other actions⁵⁴.

We consider an approximation as Δ becomes small. In the lemma below, we show that the optimal policy converges to an uninformative signal structure.

Lemma 2.1. *Let $A_+ \subseteq A$ be the set of actions that are ex-ante optimal. The optimal signal structure in the static rational inattention problem converges to an uninformative signal that always selects an ex-ante optimal action. That is, under an arbitrary norm $\|\cdot\|$,*

$$\lim_{\Delta \rightarrow 0^+} \min_{r \in \mathcal{P}(A_+)} \|p_{\Delta}^*(a|x) - r(a)\| = 0.$$

Proof. See appendix, section 5.2.2. □

Lemma 2.1 introduces the set of ex-ante optimal actions, A_+ , and argues that any optimal policy becomes arbitrary close to an uninformative signal as $\Delta \rightarrow 0^+$. The lemma is written to avoid the

⁵²For the static problem, scaling the utility by $\Delta^{\frac{1}{2}}$, as opposed to Δ or some other alternative, is arbitrary and without loss of generality.

⁵³Sims (2010); De Oliveira et al. (2013)

⁵⁴This assumption is made for mathematical convenience. We intend to analyze the general case, without this restriction, in future work.

assumption, at this point, that there is a unique probability distribution over the ex-ante optimal actions to which the optimal signal structure must converge. Intuitively, it seems possible that, if the agent can gather almost no information, he will randomize over the set of ex-ante optimal actions, and that there are many equally optimal ways of randomizing. We show in the next lemma that this intuition is essentially correct.

Proposition 2.2. *The set of optimal policies in the rational inattention problem, as $\Delta \rightarrow 0$, satisfy*

$$p_{\Delta}^*(a|x) = r^*(a) + \phi_{\Delta}(a) + \Delta^{\frac{1}{2}} r^*(a) \sum_{x' \in X_0} q(x') k^+(x, x') [u(x', a) - \sum_{a' \in A} r(a') u(x', a')] + o(\Delta^{\frac{1}{2}}),$$

where $r^*(a) \in \mathcal{P}(A_+)$ and $k^+(x, x')$ is the pseudo-inverse of $k(x, x')$. The function $\phi_{\Delta}(a)$ satisfies $\sum_{a \in A} \phi_{\Delta}(a) = 0$, $\phi_{\Delta}(a) \neq 0 \rightarrow a \in A_+$, and $\lim_{\Delta \rightarrow 0^+} \phi_{\Delta}(a) = 0$.

The function $r^*(a)$, which is the limit of every sequence of optimal policies, is the unique solution to the problem

$$\begin{aligned} & \max_{r(a) \in \mathcal{P}(A_+)} \sum_{a \in A_+} r(a) m(a, a) \\ & - \sum_{a \in A_+} \sum_{a' \in A_+} r(a') m(a, a') r(a), \end{aligned}$$

where the matrix $m(a, a')$ is defined, for $a, a' \in A_+$, as

$$m(a, a') = \sum_{x \in X_0} \sum_{x' \in X_0} q(x) k^+(x', x) q(x') u(x', a) u(x, a').$$

Proof. See appendix, section 5.2.3. □

This theorem generalizes common results from a static rational inattention problem with the mutual information cost function. Under the mutual information cost function, the matrix $q(x)q(x')k^+(x, x')$ is the inverse Fisher information matrix, evaluated at the prior $q(x)$, and the matrix $m(a, a')$ is the covariance matrix of the ex-ante utility of the various ex-ante optimal actions.

As $\Delta \rightarrow 0^+$, the quantity of information acquired by the agent converges to zero, and she can randomize, in essentially arbitrary ways, over the set of ex-ante optimal actions. In theorem 2.2, the function $\phi_\Delta(a)$ embodies this result. The uniqueness of the limit $r^*(a)$ is vacuous— the convergence of $\phi_\Delta(a)$ to zero can occur arbitrarily slowly. However, $r^*(a)$, but not $\phi_\Delta(a)$, also plays a role in the $\Delta^{\frac{1}{2}}$ -order optimal signal structure. The uniqueness of $r^*(a)$ arises from the need to maximize the agent's ability to learn.

Up to order $\Delta^{\frac{1}{2}}$, the agent always chooses an action that is ex-ante optimal. Among this set of actions, the agent prefers actions whose payoffs differ from the payoffs of the other ex-ante optimal actions. In the case of the mutual information cost function, it is desirable to sometimes choose actions whose utility varies the most across states (because one can choose the action more frequently when the utility is high than when it is low). In the general case, a similar intuition holds, but the variance is computed under a distorted probability measure (equivalently, something other than the variance is the relevant statistic).

The matrix $k(x, x')$, which determines the relative difficulty of learning about each state of the world, is the new ingredient in these formulas, relative to the standard rational inattention problem. In the standard rational inattention problem, it depends on the prior. For the purpose of modeling cognitive constraints, this dependence may be reasonable; after all, depending on one's prior, one may have an easier or harder time processing a particular signal. However, for the purposes of modeling the costs of a non-cognitive information acquisition technology, it is appealing to use a prior-independent cost. Two agents, faced with a similar set of options regarding gathering information, should face the same cost for acquiring the same signal, even if they have different prior beliefs about the environment.

Because the matrix $q(x)q(x')k^+(x, x')$ is not the inverse Fisher information matrix, the general cost functions we derive differ in their predictions from the mutual information cost function in several respects. Neither the invariant likelihood ratio property or the locally invariant posteriors property described in Caplin and Dean (2013) hold, consistent with those authors' rejection of the invariant likelihood ratio property in a laboratory setting.

In the next section, we explore a repeated version of this problem. We will view the matrix $k(x, x')$ as both fixed over time and exogenous. However, it may be fruitful to model investment in information acquisition technologies as a change in $k(x, x')$, or to view it as dependent on some time-varying aspect of the problem.

2.4 Rational Inattention and Dynamic Asset Allocation

In this section, we apply the methods developed in the previous section to a model of asset allocation. An agent (a trader, portfolio manager, or the like) manages money, perhaps on behalf of a principal, from time $t = 0$ to $t = 1$. After every period of length Δ , the agent can choose a new asset allocation, and then receive the (stochastic) returns from that asset allocation. In the framework of the previous section, a is an asset allocation, chosen from the feasible set A , and $x \in X$ is the stochastic state that determines the return of each possible asset allocation. This problem is a repeated version of the one studied by Van Nieuwerburgh and Veldkamp (2010). It is more general, in the sense of using arbitrary learning technologies and avoiding assumptions about normality, but also less general, in the sense that it requires decisions be made frequently (Δ becomes small).

The agent's log wealth under management, w_t , evolves as

$$w_{t+\Delta} = w_t + \mu(a)\Delta + \Delta^{\frac{1}{2}}f(a, x),$$

where $f(a, x)$ is the surprise log return of asset allocation a in state x , and $\mu(a)$ mean log-return.

The agent's prior about the likelihood of each state occurring in each period is $q(x)$ (the states are IID over time). Under this prior, the expected surprise log return is zero for each possible asset allocation: $\sum_{x \in X} q(x)f(a, x) = 0$ for all $a \in A$. The agent's ability to gather information allows her to achieve positive excess returns, in expectation, by determining which state is most likely before to making her asset allocation decision. The agent's value function at time t , with log wealth under management w_t and period length Δ , is denoted $V_\Delta(t, w_t)$. At time $t = 1$, the agent's

payoff (in utility terms) is $V_\Delta(1, w_1) = \bar{V}(w_1)$, where $\bar{V}(w_1)$ is an increasing, twice-differentiable, Lipschitz-continuous function⁵⁵.

Before time $t = 1$, the agent's Bellman equation is

$$V_\Delta(t, w_t) = \max_{\{p_{t,\Delta}(a|x) \in \mathcal{P}(A)\}} -C(\{p_{t,\Delta}(a|x)\}) + \sum_{x \in X} q(x) \sum_{a \in A} p_{t,\Delta}(a|x) V_\Delta(t + \Delta, w_t + \mu(a)\Delta + \Delta^{0.5} f(a, x)).$$

The agent chooses the conditional probability of each potential asset allocation a , given state x .

We analyze this equation as the time period Δ grows short.

We must impose an additional condition on the cost function, to ensure that the agent's utility remains finite. We assume a new condition, "strong convexity for informative signals," that is a strengthening of conditions 1 and 2.

Condition 5. The cost function is strongly convex for informative signals. Let $\{p(s|x)\}$ be a signal structure, define $r(s) = \sum_{x \in X} q(x)p(s|x)$, and let $m > 0$ be a positive constant. For all $\lambda \in (0, 1)$,

$$\begin{aligned} C(\{(1 - \lambda)r(s) + \lambda p(s|x)\}) &\leq \lambda C(\{p(s|x)\}) \\ &\quad - \frac{1}{2} \lambda (1 - \lambda) m \|\{p(s|x) - r(s)\}\|^2, \end{aligned}$$

where $\|\cdot\|$ is an arbitrary norm.

Condition 5 strengthens conditions 1 and 2 by assuming that strong convexity holds globally (for all $p(s|x)$), instead of holding only in the neighborhood of uninformative signals. The mutual information cost function satisfies this condition.

Next, we below characterize the optimal choice of $p_{t,\Delta}(a|x)$.

⁵⁵The assumptions that the function is increasing and twice-differentiable are made for mathematical convenience. The assumption of Lipschitz continuity ensures that utility remains finite, in the limit as $\Delta \rightarrow 0^+$. An alternative method to ensure that the agent's utility remained finite would be to include some type of decreasing returns to scale, in which case the restriction of Lipschitz continuity could be relaxed.

Lemma 2.2. *In the asset allocation problem, the optimal policy $p_{t,\Delta}^*(a|x)$ is characterized, for some $r_t^*(a) \in \mathcal{P}(A)$, as*

$$p_{t,\Delta}^*(a|x) = r_t^*(a) + \phi_\Delta(a) + \Delta^{\frac{1}{2}} V_{\Delta,w}(t, w_t) r_t^*(a) \sum_{x' \in X} q(x') k^+(x, x') [f(a, x') - \sum_{a' \in A} r_t(a') f(a', x')] + o(\Delta^{\frac{1}{2}}),$$

where $V_{\Delta,w}(t, w_t) = \frac{\partial V_\Delta(t, w_t)}{\partial w_t}$ and the function $\phi_\Delta(a)$ satisfies the properties described in theorem 2.2.

Proof. See appendix, section 5.2.5. □

The agent's optimal policy is identical to the one from the static problem of rational inattention, except that the utility has been replaced by the log excess return, scaled by the marginal utility of wealth under management. When the agent has a higher marginal utility of wealth under management, he will put more effort into information gathering. We emphasize that an agent's marginal utility of wealth under management could be a function of her preferences, formal incentives, market incentives, and other forces, all of which are embodied in our model by the terminal payoff function $\bar{V}(w_1)$.

We will next show that this extra effort results in positive expected excess returns. First, we will define three useful statistics, based on the unconditional probability of each action chosen by the agent (r_t). The first is the variance of the log excess return under the prior, denoted as $\sigma^2(r_t)$:

$$\sigma^2(r_t) = \sum_{x \in X} \sum_{a \in A} r_t(a) q(x) f(a, x)^2.$$

The second is related to a conditional variance, and to the information cost matrix. We define the quantity $\sigma_c^2(r_t)$ as

$$\sigma_c^2(r_t) = \sum_{a \in A} \sum_{x \in X} \sum_{x' \in X} r_t(a) q(x) q(x') k^+(x, x') f(x, a) [f(x', a) - \sum_{a' \in A} r_t(a') f(x', a')].$$

In the case of the mutual information cost function, this can be interpreted as the expectation (over the set of states) of the conditional variance of the log excess return, conditioned on that state and

choosing an asset allocation with probability $r_t(a)$. For our purposes, the important features of this equation are that it is a function of r_t , and it is weakly positive. Finally, define

$$\mu(r_t) = \sum_{a \in A} r_t(a) \mu(a)$$

as the mean log return, given asset allocation $r_t(a)$.

Let $\sigma_{c,t}^2 = \sigma_c^2(r_t^*)$, $\sigma_t^2 = \sigma^2(r_t^*)$, and $\mu_t = \mu(r_t^*)$ denote these statistics under the optimal policy.

In the next lemma, we use these statistics to describe the evolution of log wealth.

Lemma 2.3. *Under the optimal policy, and resulting statistics σ_t^2 and $\sigma_{c,t}^2$, the expected change in log wealth is*

$$E_t[w_{t+\Delta} - w_t] = \Delta\mu_t + \Delta V_{\Delta,w}(t, w_t) \sigma_{c,t}^2 + o(\Delta)$$

and the variance of the change in log wealth is

$$\text{Var}_t[w_{t+\Delta} - w_t] = \Delta\sigma_t^2 + o(\Delta).$$

Proof. See appendix, section 5.2.6. □

The upward drift in wealth is generated by the agent's ability to gather information before making her asset allocation decision. The higher her marginal utility of wealth under management, the more information she will gather, and the stronger this upward drift will be. The conditional variance (in the case of mutual information) controls the ability of the agent to use the information he can gather in a productive way. If there are both good and bad asset allocations in each state, information is very useful. Conversely, if all of the feasible asset allocations are highly correlated, there is a little the agent can do, because she does not control the distribution of states. The matrix $k^+(x, x')$ determines exactly which statistic (the conditional variance or something else) represents this effect.

The next step is to Taylor-expand the Bellman equation up to order Δ . The approximate Bellman equation will resemble a continuous time HJB equation, and will allow us to determine the optimal choice of r_t .

Lemma 2.4. *The Bellman equation can be approximated as*

$$0 = \max_{e_t, r_t} -\frac{1}{2}\Delta\sigma_c^2(r_t)e_t^2 + \Delta V_{\Delta, w}(t, w_t)e_t\sigma_c^2(r_t) + \Delta V_{\Delta, w}(t, w_t)\mu(r_t) \\ + \frac{1}{2}\Delta V_{\Delta, ww}(t, w_t)\sigma^2(r_t) + \Delta V_{\Delta, t}(t, w_t) + o(\Delta).$$

The first-order condition for the choice variable e_t is $e_t^ = V_{\Delta, w}(t, w_t)$. The wealth process evolves as*

$$w_{t+\Delta} - w_t = \Delta\mu_t + \Delta e_t^* \sigma_{c,t}^2 + \Delta^{\frac{1}{2}}\sigma_t \varepsilon_{t+\Delta} + o(\Delta),$$

where $\varepsilon_{t+\Delta}$ is a mean zero, IID, unit variance shock.

Proof. See appendix, section 5.2.7. □

The choice variable e_t can be interpreted as effort in information gathering. If $e_t^* = 0$ (which would only occur if the agent had no incentives), the agent pays no information costs, but generates no expected excess returns. If e_t^* is positive, the agent will engage in privately costly information gathering activities, but generate an upward drift in log wealth.

The choice of r_t is more subtle. In the context of the model, one can think of the probability vector $r_t \in \mathcal{P}(A)$ as determining the set of asset allocations the agent ‘‘considers.’’ If an asset allocation a is not considered ($r_t^*(a) = 0$), it will never be chosen. When determining which assets to consider, the agent faces a tradeoff between controlling the variance, σ_t^2 , maximizing his ability to learn, $\sigma_{c,t}^2$, and maximizing his mean return, μ_t .

Suppose that the agent’s value function is concave⁵⁶, so that the agent would like to decrease the variance, σ_t^2 . If there is a unique asset allocation in the set A that is mean-variance optimal,

⁵⁶This is by no means guaranteed– it will depend on the agent’s terminal utility function. If the agent has a convex value function, she will trade off maximizing variance against her ability to learn and maximizing returns.

the agent could choose that asset allocation one hundred percent of the time. However, this would leave the agent with no ability to learn about the relative returns of different assets; if the agent is committed to choosing one and only one asset allocation, there is no point in acquiring any information. This manifests itself, in the equations above, as $\sigma_{c,t}^2 = 0$.

Conversely, there could be asset allocations that are not desirable from a mean-variance perspective, but are negatively correlated with the other asset allocations the agent typically employs (“hedge asset allocations”). These hedges are very appealing from the agent’s perspective, because considering them ($r_t(a_{hedge}) > 0$ for some hedge asset allocation a_{hedge}) increases the ability of the agent to employ her information gathering abilities. In states where most of the asset allocations an agent normally chooses have poor performance, there is a high marginal return to considering the hedge asset allocation.

To characterize the behavior of the agent, we consider the continuous time limit of the system described in lemma 2.4.

Proposition 2.3. *In the limit as $\Delta \rightarrow 0^+$, the value function $V(t, w_t) = \lim_{\Delta \rightarrow 0^+} V_\Delta(t, w_t)$ satisfies the Hamilton-Jacobi-Bellman equation*

$$0 = \max_{e_t, r_t \in \mathcal{P}(A)} -\frac{1}{2} \sigma_c^2(r_t) e_t^2 dt + V_w(t, w_t) (e_t \sigma_c^2(r_t) + \mu(r_t)) dt + \frac{1}{2} V_{ww}(t, w_t) \sigma^2(r_t) dt + V_t(t, w_t) dt.$$

The wealth process evolves as

$$dw_t = (e_t^* \sigma_{c,t}^2 + \mu_t) dt + \sigma_t dB_t,$$

where B_t is a Brownian motion on the canonical space $(\Omega, \mathcal{F}, \tilde{P})$. The boundary condition is $V_\Delta(1, w_1) = \bar{V}(w_1)$.

Proof. See appendix, section 5.2.9. The proof is not yet complete. □

We divide our analysis of this problem into two parts. First, we consider the case where the mean, $\mu(r_t)$, and variance, $\sigma^2(r_t)$, are fixed. Second, we consider the general case, in which there is a tradeoff between μ_t , σ_t^2 , and $\sigma_{c,t}^2$.

2.4.1 Fixed Mean-Variance

The simplest version of our problem occurs when $\mu(r_t)$ and $\sigma^2(r_t)$ are not under control of the agent. This would occur when the set A contains only asset allocations with a particular mean and variance. A principal might choose to restrict the set A in this way, to encourage the agent to employ her information-gathering abilities⁵⁷.

In any case, assume that $\mu(r_t) = \bar{\mu}$ and $\sigma^2(r_t) = \bar{\sigma}^2$. Subject to this constraint, the agent will choose r_t to maximize $\sigma_{c,t}^2(r_t)$. Because the problem is identical at each time t and wealth w_t , this will result in a unique value⁵⁸ for $\sigma_{c,t}^2$, $\sigma_{c,t}^2 = \bar{\sigma}_c^2$, for all t and w_t . It is useful to define the parameter $\theta = \frac{\bar{\sigma}^2}{\bar{\sigma}_c^2}$, and redefine the control variable $u_t = \bar{\sigma}_c^2 e_t \bar{\sigma}^{-1}$. The HJB equation can be rewritten as

$$0 = \max_{e_t} -\frac{1}{2}\theta u_t^2 dt + V_w(t, w_t)(\bar{\sigma} u_t + \bar{\mu})dt + \frac{1}{2}V_{ww}(t, w_t)\bar{\sigma}^2 dt + V_t(t, w_t)dt,$$

subject to the stochastic process $dw_t = \mu dt + u_t \bar{\sigma} dt + \bar{\sigma} dB_t$. This is exactly the quadratic effort cost moral hazard problem analyzed in Cvitanic et al. (2009) and Hébert (2014), and we can use their results to quickly characterize the agent's behavior. The value function is $V(t, w_t) = \theta \ln(E_t[\exp(\theta^{-1}\bar{V}(w_1))|u_s = 0])$, where $E_t[\exp(\theta^{-1}\bar{V}(w_1))|u_s = 0]$ is the expectation of that exponential under the assumption of no further effort (that $u_s = 0$ for all $s > t$). The optimal information-gathering effort, e_t , is the “delta,” in the option pricing sense of the term, of the terminal value function $\bar{V}(x_1)$. Because of the tractable nature of the agent's behavior in this model, it is also

⁵⁷The construction of the set A is itself a crucial part of the delegated portfolio management contract (as in Dybvig et al. (2010)), and some sort of restrictions are necessary to prevent the agent from undoing the effects of whatever incentive scheme is implemented (Admati and Pfleiderer (1997)).

⁵⁸We assume this value is strictly positive.

possible to characterize the optimal contract between this agent and a principal (see Cvitanić et al. (2009) and Hébert (2014) for details).

2.4.2 The Tradeoff Between Mean-Variance Optimality and Learning

In this section, we return to the more general model, in which it is possible for the agent to trade off $\mu(r_t)$, $\sigma_c^2(r_t)$ and $\sigma^2(r_t)$. Formally, the agent solves

$$\max_{r_t \in \mathcal{P}(A)} V_w(t, w_t) \mu(r_t) + \frac{1}{2} V_w(t, w_t)^2 \sigma_c^2(r_t) + \frac{1}{2} V_{ww}(t, w_t) \sigma^2(r_t)$$

at each time t . Returning to the definitions of $\sigma_c^2(r_t)$ and $\sigma^2(r_t)$, this maximization can be written as

$$\begin{aligned} \max_{r_t} V_w(t, w_t)^2 \sum_{a \in A} \sum_{x \in X_0} \sum_{x' \in X_0} r_t(a) q(x) q(x') k_0^{-1}(x, x') f_0(x, a) [f_0(x', a) - \sum_{a' \in A} r_t(a') f_0(x', a')] + \\ V_{ww}(t, w_t) \sum_{x \in X} \sum_{a \in A} r_t(a) q(x) f(a, x)^2 + 2V_w(t, w_t) \sum_{a \in A} r_t(a) \mu(a) + \\ + 2V_w(t, w_t)^2 \lambda_t (1 - \sum_{a \in A} r_t(a)) + 2V_w(t, w_t)^2 \sum_{a \in A} v_t(a) r_t(a), \end{aligned}$$

where λ_t and $v_t(a)$ are the multipliers on the constraints that $r_t(a)$ sum to one and be weakly positive, respectively, scaled by $V_w(t, w_t)^2$ for convenience. Define the matrix (see theorem 2.2)

$$m(a, a') = \sum_{x \in X_0} \sum_{x' \in X_0} q(x) q(x') k_0^{-1}(x, x') f_0(x, a) f_0(x', a').$$

In the case of the mutual information cost function, this is simply the covariance matrix of the log excess returns of the various feasible asset allocations. The first-order condition for the maximization problem can be written as

$$\sum_{a' \in A} m(a, a') r_t(a') = \frac{1}{2} m(a, a) + \frac{\mu(a)}{V_w(t, w_t)} + \frac{V_{ww}(t, w_t)}{2V_w(t, w_t)^2} \sum_{x \in X} q(x) f(a, x)^2 - \lambda_t + v_t(a).$$

Assuming that there are no redundant asset allocations in the set A , the matrix $m(a, a')$ will be positive definite, and there will be a unique optimal $r_t(a)$.

In general, it is difficult to determine which asset allocations will be chosen with positive probability. If the agent was not able to learn, she would choose the mean-variance optimal portfolio only. However, with the ability to learn, she also considers asset allocations that might not be mean-variance optimal, and in general will not always choose the same asset allocation.

In the case of mutual information, when $m(a, a)$ is the covariance matrix of asset returns, the choice of r_t that maximizes $\sigma_c^2(r_t)$ has an interesting interpretation. The r_t that maximizes $\sigma_c^2(r_t)$ is the “portfolio” of asset allocations that maximizes the variance reduction of diversification. That is, it is the “portfolio” for which the difference between the weighted average variance of its holdings and the variance of the portfolio return is largest.

The simplest version of this model is a two-asset model, in which one of the assets is the risk-free asset. In this case, the first-order condition for $r_t(a_{risky})$ (assuming it is interior) is

$$r_t(a_{risky}) = \frac{1}{2} + \frac{\frac{1}{V_w(t, w_t)} \mu(a_{risky}) + \frac{V_{ww}(t, w_t)}{2V_w(t, w_t)^2} \sum_{x \in X} q(x) f(a_{risky}, x)^2}{\sum_{x \in X_0} \sum_{x' \in X_0} q(x) q(x') k_0^{-1}(x, x') f_0(x, a_{risky}) f_0(x', a_{risky})}. \quad (2.5)$$

When the agent is indifferent between the risky and risk-free asset allocations (the numerator in equation (2.5) is zero), he will choose the risky asset allocation half of the time, and the risk-free asset allocation half of the time. The intuition is that, by choosing each alternative half of the time, the agent maximizes her ability to learn. This is not a “mixed strategy,” in which any randomization between the two portfolios is equally desirable; the agent optimally chooses a 50/50 probability, not any other probability. This strategy is also not the same as always choosing a portfolio that is composed of half the risky asset and half the risk-free asset; the agent diversifies over time, but always holds entirely one asset or the other.

As the risky asset becomes more or less attractive, in a mean-variance sense, she chooses the risky or safe asset more frequently. The numerator in equation (2.5) reflects the mean-variance desirability of the risky asset, relative to the risk-free asset. The denominator captures the inverse of the agent’s ability to learn about the returns of the risk free asset. When the denominator is

high (learning is easy), the agent is reluctant to deviate from the 50/50 strategy, even if either the risk-free or risky asset is much better in a mean-variance sense. When the denominator is low, the agent will quickly cease to randomize, if one asset allocation comes to dominate the other in a mean-variance sense.

Another interesting case is when the mean returns are all zero ($\mu(a) = 0$) and, for each asset allocation, there is a “short” asset allocation with the opposite returns. This case might be used to describe a derivatives trader working at a bank, for example. In this case, the agent will exhibit “specialization,” in the sense described by Van Nieuwerburgh and Veldkamp (2010). The agent will choose only the largest variance asset, and be “long” that asset half of the time, and “short” the other half of the time. If there are multiple assets with maximal variance, the agent will choose only one (“specialization”). This case would reduce to the fixed volatility case described above.

2.5 Conclusion

In this chapter, we have introduced a new generalization of rational inattention models. This generalization is built from axioms that are appropriate for many economic problems. Despite the generality of our framework, we derive concrete predictions about behavior. Our approach employs an approximation, that the amount of information acquired by the agent is small, that may or may not be applicable in various settings. In a dynamic model of asset allocation, we illustrate how to use our approach to characterize behavior, despite the generality of our framework. Our solution to the dynamic problem does not require specific knowledge about the agent’s information acquisition technology.

3 The Costs of Sovereign Default: Evidence from Argentina

3.1 Introduction

A fundamental question in international macroeconomics is why governments repay their debt to foreign creditors, given the limited recourse available to those creditors. The seminal paper of Eaton and Gersovitz (1981) argues that reputational concerns alone are sufficient to ensure that sovereigns repay their debt. Because a default leads to a loss of international reputation, defaulting countries are excluded from sovereign bond markets and can no longer share risk. Eaton and Gersovitz (1981) argue that countries repay their debt to maintain their international reputation and access to credit markets. In a famous critique, Bulow and Rogoff (1989b) demonstrate that reputational contracts alone cannot be sustained in equilibrium without some other type of default cost or punishment. Following this critique, hundreds of papers have been written trying to identify these costs of default. The fundamental identification challenge is that governments usually default in response to deteriorating economic conditions, which makes it hard to determine if the default itself caused further harm to the economy.

The case of *Republic of Argentina v. NML Capital* provides a natural experiment to disentangle the causal effect of sovereign default. Following Argentina's sovereign default in 2001, NML Capital, a subsidiary of Elliott Management Corporation, purchased defaulted bonds and refused to join other creditors in restructurings of the debt during 2005 and 2010. Instead, because the debt was issued under New York law, NML sued the Argentine state in US courts to receive full payment. To compel the Argentine government to repay the defaulted debt in full, the US courts blocked Argentina's ability to pay its restructured creditors until NML and the other holdout creditors were paid in full. The Argentine government resisted paying the holdouts in full, even though the required payments would not be particularly large relative to the Argentine economy. As a re-

sult, rulings in favor of NML raised the probability that Argentina would default on its restructured bonds, while rulings in favor of Argentina lowered this probability.

Because the court rulings were not responding to private information about the underlying economic circumstances in Argentina, we can use them to examine the effect of changing default probabilities on Argentine firms. We use credit default swaps (CDS) to measure the change in the risk neutral probability of default. Compiling rulings from the United States District Court for the Southern District of New York, the Second Court of Appeals, and United States Supreme Court, we identify sixteen rulings that potentially changed the probability of default. We find that, for every 1% increase in the 5-year cumulative default probability, the US dollar value of an index of Argentine American Depository Receipts (ADRs) falls 0.55%.⁵⁹ Between January 3, 2011, when our data starts, and July 30, 2014, when Argentina defaulted, the risk-neutral 5-year default probability increased from roughly 40% to 100%. Our estimates imply that this episode reduced the value of the Argentine firms in our index by 33%.

We begin our analysis by studying these legal rulings in an event study framework. We find economically significant negative returns for the ADRs of Argentine firms in response to legal rulings in favor of NML, and positive returns in response to rulings in favor of Argentina. We find these effects when using two-day event windows, and when using narrower windows that vary in size depending on the announcement time of the rulings. We also find that a measure of the “blue rate,” the unofficial exchange rate between Argentine pesos and US dollars, depreciates in response to rulings in favor of NML and appreciates in response to rulings in favor of Argentina.

The event study approach is subject to the concern that other factors may have changed during the relevant event windows. To alleviate this concern, following Rigobon (2003) and Rigobon and Sack (2004), we identify the effect of changes in the default probability on equity returns through heteroskedasticity. We assume that on days in which US courts rule on *Republic of Argentina v. NML Capital* and related court cases, the variance of shocks to the probability of default is higher than on other days. Using this identification strategy, we find results consistent with our

⁵⁹American Depository Receipts are shares in foreign firms that trade on US stock exchanges in US dollars.

event study approach. We interpret our results as providing evidence that sovereign default causes economically significant harm to corporations from the defaulting country.

To better understand how a sovereign default affects the economy, we examine which types of firms are harmed more or less by an increase in the probability of default. We sort firms along the dimensions suggested by the theoretical sovereign debt literature, as well as on some additional firm characteristics. We find suggestive evidence that banks, exporters, and foreign-owned firms are hurt more by increases in the probability of sovereign default than would be expected, given their “beta” to the Argentine market. Our results do not necessarily imply that these types of firms are hurt more in absolute terms by an increase in the probability of default than other firms, although this is indeed the case for banks. Instead, our results show that these firms are hurt more by an increase in the risk of a sovereign default than they would be by a “typical” shock that had the same impact on a broad index of Argentine stocks.⁶⁰

This chapter contributes to a large literature examining the costs of sovereign default. The question of the cost of sovereign default is surveyed in Borensztein and Panizza (2008). Using quarterly data, Yeyati and Panizza (2011) find that output generally falls in anticipation of a sovereign default and the default itself tends to mark the beginning of the recovery. Bulow and Rogoff (1989a) argue that default is costly because foreign lenders can disrupt trade, a channel for which Rose (2006), Borensztein and Panizza (2010), and Zymek (2012) find empirical support. Gennaioli et al. (2014), Acharya et al. (2014), Bocola (2013) and Perez (2014) present models of the disruptive effect of default on the financial system and the consequent disruption of macroeconomic activity. Mendoza and Yue (2012) present a general equilibrium strategic default model, building on the framework of Aguiar and Gopinath (2006) and Arellano (2008), where default is costly because it reduces the ability of domestic firms to import intermediate goods, reducing their productivity. Cole and Kehoe (1998) argue that a sovereign default causes the government to lose its reputation not just in regards to the repayment debt, but also more generally. Arteta and Hale (2008) observe that during

⁶⁰We calculate the effect of sovereign default conditional on the index because we do not want to imply that all high-beta firms are hurt more by sovereign default risk than other firms, simply because their value falls more than the broad index in response to an increase in sovereign risk. We find that exporters are more adversely affected by sovereign default than implied by their market beta, but they are not actually hurt more in absolute terms than non-exporters.

a sovereign default, external credit to the private sector is reduced. Schumacher et al. (2014) study sovereign debt litigation across a range of countries over the past 40 years. They find that creditor litigation is associated with a decline in international trade, sovereign exclusion from financial markets, and a longer time before the default is resolved. The Argentine case studied here differs from most of the cases studied in Schumacher et al. (2014) as this litigation also changed the probability of a new default, in addition to affecting the government's ability to resolve an ongoing default.

This chapter is structured as follows: section 3.2 discusses the case of *Republic of Argentina v. NML Capital*. Section 3.3 describes the data and presents summary statistics for the behavior of CDS and equity returns on event and non-event days. Section 3.4 presents our estimation framework, the identifying assumptions, and our results. Section 3.5 discusses industries and firm characteristics that are associated with larger responses to changes in the probability of sovereign default. Section 3.6 presents the interpretation of the results. Section 3.7 concludes.

3.2 Argentina's Sovereign Debt Saga

3.2.1 The Argentine Default of 2001 and the Restructurings of 2005 and 2010

Following decades of rampant inflation, in 1991 the Argentine government adopted the “convertibility plan,” introducing a currency board in an attempt to irrevocably fix the peso-dollar exchange rate at one-to-one. This meant that the government legally committed itself not to print any currency that was not backed one-to-one by a US dollar in reserves. While inflation fell following the convertibility plan, the government continued to run a deficit, largely financed through external dollar borrowing. In 2001, Argentina entered a deep recession, with unemployment reaching 14.7% in the fourth quarter.⁶¹ In December 2001, after borrowing heavily from the IMF, Argentina defaulted on over \$100 billion in external sovereign debt and devalued the exchange rate by 75%.⁶²

The Argentine government then spent three years in failed negotiations with the IMF, the Paris Club, and its private creditors. In January 2005, Argentina presented a unilateral offer to its private

⁶¹Data from Global Financial Data.

⁶²Daseking et al. (2005).

creditors, which was accepted by the holders of \$62.3 billion of the defaulted debt.⁶³ To strengthen its bargaining position, the Argentine legislature passed the “Lock Law,” prohibiting the government from reopening the debt exchange or making any future offers on better terms.⁶⁴ After the first round of restructuring, holdout creditors were still owed \$18.6 billion of principal, the Paris Club of creditors was owed \$6.3 billion, and the IMF was owed \$9.5 billion.⁶⁵ Despite the existence of the holdout creditors, S&P declared the end of the Argentine default in June 2005 and upgraded Argentina’s long-term sovereign foreign currency credit rating to B-. In 2006, Argentina fully repaid the IMF, and Argentina reached an agreement with the Paris Club creditors in May 2014.⁶⁶

In December 2010, Argentina offered another bond exchange to the holdout private creditors. Holdout private creditors who were owed \$12.4 billion of principal agreed to the exchange. Following the exchange, on December 31, 2010, the remaining holdout creditors were owed an estimated \$11.2 billion, split between \$6.8 billion in principal and \$4.4 billion in accumulated interest.⁶⁷ At this point, Argentina had restructured over 90% of its original debt.

3.2.2 Argentina vs. the “Vultures”

Following the 2010 debt exchange, the remaining holdout creditors, termed “vultures” by the Argentine government, continued their legal battle. One line of attack was on the Argentine government’s reserve assets, with the creditors arguing the country’s reserves, held at the Federal Reserve Bank of New York, should be subject to attachment. While a district court initially agreed with the creditors, in 2011 the appellate court overturned the ruling.⁶⁸ The second line of attack, focused on the *pari passu* clause, was the one that eventually culminated in Argentina’s recent default. The

⁶³Hornbeck (2013).

⁶⁴This Lock Law would feature prominently in Judge Griesa’s interpretation of the *pari passu* clause, presenting evidence that holdouts were not on the same footing as the holdout creditors.

⁶⁵Hornbeck (2013).

⁶⁶<http://www.reuters.com/article/2014/05/29/us-argentina-debt-parisclub-idUSKBN0E90JI20140529>

⁶⁷Hornbeck (2013).

⁶⁸Hornbeck (2013)

pari passu clause requires equal treatment of all bondholders. The creditors, led by NML Capital,⁶⁹ argued that the Argentine government breached this clause by paying the exchange bondholders and refusing to honor the claims of the holdouts. In addition, the holdouts asserted that the “Lock Law,” by making explicit the government’s policy of pledging not to re-open negotiations or pay any money, effectively subordinated them to the restructured bondholders.

The case took several years to work its way through the US courts, going from the United States District Court for the Southern District of New York (“Southern District”), to the United States Court of Appeals for the Second Circuit (“Second Circuit”), all the way to the United States Supreme Court. The numerous rulings that these three courts issued between December 2011, when Judge Thomas P. Griesa of the Southern District first ruled in favor of the holdouts on the *pari passu* issue,⁷⁰ until July 2014 when Argentina defaulted. For the purposes of this study, we view the various rulings as events that made it more or less likely that Argentina would be unable to pay the restructured bondholders, if it did not also repay the holdouts. Because of the Argentine government’s unwillingness to pay the holdouts in full, rulings in favor of NML increased the probability of a default on the restructured bonds, while rulings in favor of Argentina reduced the probability of default.⁷¹

Following Griesa’s initial ruling in December 2011, a year of legal wrangling ensued over what this ruling actually meant and how it would be enforced. Griesa clarified that Argentina was required to repay the holdouts as long as it was continuing to pay the exchange bondholders (using a “ratable” payment formula). Argentina was not willing to comply with this ruling, and continued to pay the exchange bondholders without paying the holdouts. Griesa then ordered the financial intermediaries facilitating Argentina’s payments to stop forwarding payments to the restructured bondholders until Argentina paid the holdouts. As a result, even if Argentina wanted to

⁶⁹Elliott Management Corporation, the parent company of NML, has a long history in litigating against defaulting countries. See Gulati and Klee (2001) for a discussion of Elliot’s litigation against Peru and Panizza et al. (2009) for an excellent literature review on the law and economics of sovereign default.

⁷⁰See the decision at <http://blogs.reuters.com/felix-salmon/files/2012/04/2011-12-07-Equal-Treatment-Liability-Order.pdf>

⁷¹We use the term default to refer to a “credit event,” as defined in the credit default swap contracts we study. Defaults come in many varieties, from a temporary cessation in payments to complete repudiation.

pay the restructured creditors, it could not do so without repaying the holdouts, as its trustee would not be allowed to disburse the funds delivered for the coupon payment. In late 2012, Griesa ordered Argentina to negotiate with the holdouts, but the holdouts and the courts rejected Argentina's offer of a deal comparable to the 2005 and 2010 bond exchanges. Argentina then twice appealed to the Supreme Court, with the Supreme Court declining to hear either appeal. Following the decline of the second appeal on June 16, 2014, Griesa's orders were implemented, and Argentina had only two weeks before a coupon to the restructured creditors was due. Against the court orders, Argentina actually sent this coupon payment to the bond trustee, Bank of New York Mellon (BNYM), but due the court order, BNYM did not forward to the payment to the restructured bond holders. Argentina legally missed the coupon payment on June 30, which began a 30-day grace period. After negotiations failed, Argentina entered default on July 30, 2014.

3.2.3 A Simple Interpretation

In the simplest interpretation of the unfolding court events, Argentina was forced to default by the US court system. This was the interpretation offered by a number of commentators in the financial press.⁷² Under this interpretation, Argentina could not pay its debts because the US courts forbade financial intermediaries from facilitating the coupon payment. As a result, the court rulings did nothing but change the probability of a default.

We also argue that these legal rulings do not reveal information about the underlying state of the economy (or other unobserved fundamentals), except insofar as they change the probability of default. The key assumption is that Judge Griesa (and the second circuit and Supreme Court) have no information advantage over the market with respect to the state of the Argentine economy.⁷³

⁷²For instance, Matt O'Brien of the Washington Post wrote "Argentina was forced to default now, because it wouldn't pay the bonds it had defaulted on in 2001" (<http://www.washingtonpost.com/blogs/wonkblog/wp/2014/08/03/everything-you-need-to-know-about-argentinas-weird-default/>).

⁷³In the event study literature that focuses on Federal Reserve monetary policy announcements, there is some concern that the Federal Reserve has more information than market participants about the state of the economy. These sorts of concerns are unlikely to apply in this chapter.

Under this interpretation, we can use credit default swaps to measure the market-implied changes in the probability of default following the court rulings. Any effect these rulings have on other variables, such as equity returns, is caused by the change in the probability of default. By comparing these two quantities (the change in default probability and the stock return), we can estimate the effect of default on the value of the firm.

This interpretation motivates our empirical strategy. We look at the stock returns in windows around each of these events, and estimate how changes in the risk of sovereign default are related to changes in the valuation of Argentine firms. We employ several different empirical methods, which are described in detail in the next two sections. After presenting our empirical results, we will discuss alternative interpretations of the legal rulings and our results. We will also discuss several important details about the Argentine debt situation that are relevant for the interpretation of our results.

3.3 Data and Summary Statistics

3.3.1 Stock Market and CDS Data

Our dataset consists of daily observations of financial variables from January 3, 2011 to July 29, 2014 (the day before Argentina most recently defaulted). We study the returns of US dollar-denominated ADRs issued by Argentine firms, which are traded in the United States, as well the Argentine peso-denominated equities traded in Argentina. The ADRs trade on the NYSE and NASDAQ, are relatively liquid, and can be traded by a wide range of market participants⁷⁴. However, using only the ADRs limits the number of firms that can be included in our analysis. To study the cross-sectional patterns of Argentine firms, we also examine the returns of firms traded only in Argentina. In order to ensure sufficient data quality, we limit our study of local Argentine equities to firms with a 2011 market capitalization at least 200 million pesos, and for which the

⁷⁴Several market participants have told us that capital controls and related barriers are significant impediments to their participation in local Argentine equity markets.

equity price changes on at least half of all trading days in our sample. The full list firms included in our analysis, along with select firm characteristics, can be seen in table 3.1.

Table 3.1: Firms Included in Analysis

Company	Ticker	Industry	Exports	Imports	Market Cap	ADR	Foreign
Aluar	ALU	Manufacturing	19.4	13.1	9443.0		
Banco Santander Rio	BIO	Banks	1.5	0.4	12786.1	Y	Y
Boldt	BOL	Telecoms			1537.5		
Banco Patagonia	BPT	Banks	1.5	0.4	3488.4	Y	Y
Banco Macro Bansud	BSU	Banks	1.5	0.4	9379.2	Y	
Carlos Casado	CAO	Real Estate	0.0	0.0	378.1	Y	
Endesa Costanera	CEC	Utilities	0.1	3.2	702.6		Y
Celulosa	CEL	Chemicals	11.2	19.2	760.3		
Central Puerto Rights	CEP	Utilities	0.1	3.2	1814.4		
Grupo Clarin	CLR	Telecoms	0.4	1.3	5633.4	Y	
Sociedad Comercial Del Plata	COM	Diverse	1.5	0.4	212.3		
Capex	CPX	Utilities	0.1	3.2	1087.8		
Cresud	CRE	Non-Durables	14.5	2.1	3495.9	Y	
Edenor	EDC	Utilities	0.1	3.2	1894.5	Y	
BBVA Banco Frances	FRA	Banks	1.5	0.4	7723.6	Y	Y
Pampa Energia	FRG	Utilities	0.1	3.2	3417.2	Y	
Gp Finance Galicia	GGA	Banks	1.5	0.4	7125.7	Y	
Hipotecario Naci	HPD	Banks	1.5	0.4	3540.0	Y	
Solvay Indupa	IND	Chemicals	11.2	19.2	1218.0		Y
IRSA	IRS	Real Estate	0.0	0.0	3350.5	Y	
Juan Minetti	JMI	Manufacturing	3.6	5.3	1633.5		Y
Ledesma	LED	Non-Durables	14.5	2.1	4004.0		
Metrogas	MET	Utilities	0.1	3.2	677.3	Y	
Mirgor	MIR	Manufacturing	10.1	11.2	512.0	Y	
Molinos Rio De La Plata	MOL	Non-Durables	19.5	2.8	8014.4		
Quickfood	PAY	Non-Durables	19.5	2.8	641.9		Y
Petrobras Energia	PER	Energy	25.5	2.1	8228.4	Y	Y
IRSA Propiedades Comerciales	SAM	Real Estate	0.0	0.0	2960.1	Y	
Moli Juan Semino	SEI	Non-Durables	19.5	2.8	325.5		
Siderar	SID	Manufacturing	19.4	13.1	10893.1		Y
Sa San Miguel	SMG	Non-Durables	19.5	2.8	491.1		
Telecom Argentina	TEC	Telecoms	2.7	8.1	21754.8	Y	Y
Transportadora De Gas Del Sur	TGS	Energy	25.5	2.1	2558.3	Y	
Transportadores De Gas Del Norte	TN4	Utilities	0.1	3.2	540.4		
Transener	TRA	Utilities	0.1	3.2	640.3		
YPF	YPF	Energy	14.2	8.4	74532.8	Y	

Notes: This table lists the 35 firms used in the analysis of local equities, and one firm (ticker SAM) whose ADR is included in our ADR sample, but whose local stock returns do not pass our data quality requirement. Ticker indicates the company's ticker in Datastream. Exports (Imports) denotes the percentage of exports (imports) of total output for the firm's primary industry. Both Exports and Imports are calculated by classifying the firm into one of the 37 industries in the OECD STAN Input-Output Table according the SIC code of the firm's primary industry. Market Cap is the firm's average end-of-quarter market capitalization in 2011 from Bloomberg, measured in Argentine pesos. ADR Sample indicates whether the ADR is included in our sample of ADRs. Foreign is an indicator for whether the firm is owned by a non-Argentine parent company.

Our primary measure of the performance of the Argentine equity market comes from the MSCI Argentina index, an index of six Argentine ADRs. As of November 28, 2014, the companies included in the index (with their index weights in parentheses), are: YPF (30.56%); Telecom Argentina (22.11%); Banco Macro (18.64%); Grupo Financiero Galicia (16.11%); BBVA Banco Frances (7.50%); and Petrobras Argentina (5.09%). In addition, we construct our own indices of ADRs covering different sectors of the Argentine economy. We classify Argentine firms by whether they are a bank, a non-financial firm, or a real estate holding company . The industry classifications are based on the Fama-French 12 industry classification and are described in detail at the end of this section. We give equal weighting to each ADR included in our three indices. The financial index is composed of the ADRs of BBVA Banco Frances (ADR ticker BFR); Banco Macro (BMA); and Grupo Financiero Galicia (GGAL). The industrial index is composed of Cresud (CRESY); Empresa Distribuidora y Comercializadora Norte (EDN); Pampa Energia (PAM); Petrobras Argentina (PZE); Telecom Argentina (TEO); Transportador Gas Sur (TGS); and YPF (YPF). The real estate index is composed of Alto Palermo (APSA, local ticker SAM) and IRSA (IRS). Alto Palermo is the only firm included in the ADR-based analysis but excluded from the local equity analysis. It is excluded because the local price data does not meet the data quality requirements. The classification of the 35 firms included in the analysis of local equities can be found in table 3.1. Because the MSCI Argentina Index is heavily tilted toward financial and energy firms, for our analysis of local returns, we construct an index equally weighting all of the returns of each of the 35 firms.

We use credit default swap (CDS) spreads to measure the market-implied risk-neutral probability of default. A CDS is a financial derivative where the seller of the swap agrees to insure the buyer against the possibility that the issuer defaults. Once a third party, generally the International Swaps and Derivatives Association (ISDA), declares a credit event, an auction occurs to determine the price of the defaulted debt. The CDS seller then pays the buyer the difference between the face and auction value of the debt. In appendix section §5.3.3, we provide details on how we impute

risk-neutral default probabilities from the term structure of CDS spreads using the ISDA Standard Model. We focus on the 5-year cumulative default probability, the risk-neutral probability that Argentina defaults within 5 years of the CDS contract initiation.

Our CDS data is from Markit, a commercial data provider. We use a “sameday” CDS spread as of 9:30 am EST, which we refer to as the “open,” and a composite end-of-day spread, which we refer to as the “close.”⁷⁵ The composite end-of-day spread is gathered over a period of several hours from various market makers, and is the spread used by those market makers to value their own trading books. The composite end-of-day spread uses updated expectations about the recovery rate⁷⁶, whereas the sameday spread is built under the assumption that the expected recovery rate has not changed from the previous day’s close. Markit uses a data cleaning process to ensure that both the sameday and composite end-of-day quotes are reasonable approximations of market prices.

Because we want to capture the abnormal variation in Argentine CDS and equity returns caused by changes in the probability of default, we need to account for other global factors that may affect both measures. To proxy for global risk aversion, we use the VIX index, the 30-day implied volatility on the S&P 500.⁷⁷ We use the S&P 500 to measure global equity returns and we use the MSCI Emerging Markets Asia ETF to proxy for factors affecting emerging markets generally. We use the Asian index to ensure that movements in the index are not directly caused by fluctuations in Argentine markets. To control for aggregate credit market conditions, we use the Markit CDX High Yield and Investment Grade CDS indices.⁷⁸ These controls are included in all specifications reported in this chapter, although our results are qualitatively similar when using a subset of these factors, or no controls at all.

In order to examine the channels through which a sovereign default can affect domestic firms, in section 3.5 we will sort firms along a number of dimensions. We begin by classifying firms

⁷⁵We have also run our results with a 3:30pm “sameday” quote, instead of the composite end-of-day. Our point estimates are similar, but the standard errors are larger.

⁷⁶Markit surveys CDS dealers at the end of each day to gather expected recovery rates.

⁷⁷See Longstaff et al. (2011b) for discussion of VIX and variation in sovereign CDS spreads.

⁷⁸We use the continuous on the run series from Thomson Reuters Datastream. More information on these indices can be found at <https://www.markit.com/news/Credit%20Indices%20Primer.pdf>.

according to their Fama-French industry classifications available on Kenneth French's website.⁷⁹ We sort firms into their corresponding Fama-French industries according the SIC code of their primary industry, available from Datastream. After this initial sort, we only have one firm, Boldt, classified as Business Equipment, and so we combine it with the telecommunications firms. The "Finance" Fama-French 12 industry classification is also too broad for our purposes, as it combines banks, holding companies, and real estate firms. We therefore split the nine firms initially classified as "Finance" according to their Fama-French 49 industry classification. This gives us six banks, two real estate firms, and one "Trading" firm, Sociedad Comercial del Plata. Because Sociedad Comercial del Plata is a diversified holding company, and is the only company in the Fama-French 49 industry classification of "Trading," we rename its industry "Diversified", and do not merge it with any other industry classification. After these modifications, we end up with six banks, two chemical firms, one diversified firm, three energy firms, four manufacturing firms, six non-durables firms, two real estate firms, three telecoms and eight utilities.

The next dimension along which we will sort firms is their exporter status. Bulow and Rogoff (1989a) posited that default was costly because foreign creditors have the ability to interfere with a country's trade. To test this channel, we examine if exporting firms are particularly hurt by an increase in the probability of default. Unfortunately, this task is complicated by the fact that publicly available data sources do not comprehensively report firm-level exports. We instead rely on industry-level measures. We use the OECD STAN Input-Output Tables for Argentina to calculate what share of each industry group's output is exported. The Input-Output Table covers 37 industries, each of which covers at least one two-digit ISIC industry, and some of which, such as "Agriculture, hunting, forestry and fishing", cover up to five two-digit ISICs. After we calculate the share of exports for each of these 37 industries, we classify our 35 firms into one of these industries according to the SIC code of its primary output. Unfortunately, the most recent Input-Output Table for Argentina uses data from 1995, so our export analysis assumes that the relative tradabil-

⁷⁹Classifications available [here](#). We use the versions formatted by Dexin Zhou.

ity of different products has not changed too much over the past 20 years.⁸⁰ When we construct a zero-cost long-short portfolio, going long exporters and short non-exporters, we will classify firms as exporters if exports accounted for at least 10% of their primary industry's revenues in our Input-Output table, and non-exporters otherwise.

In order to examine the channel proposed by Mendoza and Yue (2012) that a sovereign default is costly because it reduces the ability of firms to import intermediate goods for production, we calculate the share of intermediate inputs imported for each industry. We again use the OECD STAN Input-Output Tables to calculate the reliance on imported intermediate goods for 37 industries, and then match each of our firms to these industries using their primary SIC code. As with exports, we rely on the 1995 Input-Output Table. For portfolio construction, we classify firms as non-importers if imported intermediates are less than 3% of total sales in their primary industry, and as importers otherwise.

The next cut of the data divides firms among those that are subsidiaries of foreign corporations and those that are not. We classify firms as foreign-owned if the headquarters of their ultimate parent is any country other than Argentina in Bloomberg (Field `ULT_PARENT_CNTRY_DOMICILE`). There are a number of reasons the effect of sovereign default might be different for foreign-owned and domestic firms. For instance, if a sovereign default has a large effect on the domestic banking system, perhaps foreign affiliates might have a relatively easier time accessing finance than domestically owned firms. A similar effect is documented for multinational and local firms following an exchange rate depreciation in Desai et al. (2008), with multinational firms cutting investment less than domestic firms, presumably because external financing helps multinational mitigate the balance sheet effect. On the other hand, if defaulting costs Argentina its "general reputation," as in Cole and Kehoe (1998), it may be more inclined to seize the assets of foreign firms. In this case, we would expect foreign-owned firms to underperform. We use the most recent ownership of this

⁸⁰For those firms that report data on revenue from exports, there is a strong correlation between reported exports as a share of sales and the imputed share of exports from the 1995 input-output table. These results are available upon request.

variable and cannot account for the possibility that an Argentine firm was only recently purchased by a foreign parent.

The final variable we use to classify our local equities is an indicator for whether or not the firms have an ADR. The reason for this is that ADRs are a potentially important way for residents to evade the government's capital controls. We might expect firms with ADRs to outperform firms without ADRs because the former is a valuable vehicle for acquiring foreign currency offshore. This feature of ADRs is addressed in detail by Auguste et al. (2006)

3.3.2 Definition of Events and Non-Events

We build a list of legal rulings issued by Judge Griesa, the Second Circuit, and the Supreme Court. We have created this list using articles in media (the Wall Street Journal, Bloomberg News, and the Financial Times), LexisNexis searches, and publicly available information from the website of a law firm (Shearman) that practices sovereign debt law.

In appendix table 5.3.5, we list all of these events and links to the relevant source material. Unfortunately, for many of the events, we are unable to determine precisely when the ruling was issued. We employ several methods to determine the timing of rulings. First, we examine news coverage of the rulings, using Bloomberg News, the Financial Times, and LexisNexis searches. Sometimes, contemporaneous news coverage specifically mentions when the ruling was released. Second, we use the date listed in the ruling (usually next to the judge's signature). Third, many of rulings are released in the PDF electronic format, and have a "creation time" and/or "modification time" listed in the meta-information of the PDF file. In appendix table 5.3.5, we list the information used to determine the approximate time of each ruling.

For most of our analysis, we use two-day event windows. Consider the Supreme Court ruling on Monday, June 16th, 2014. The two-day event window, applied to this event, would use the CDS spread change from the close on Friday, June 13th to the close on Tuesday, June 17th. It would use

stock returns (for both ADRs and local stocks) from 4pm EDT on Friday, June 13th to 4pm EDT on Tuesday, June 17th⁸¹.

For one section of our analysis, we use narrower window sizes, when possible. We classify events into several types based on when they occurred. We classify events as close-to-close, open-to-open, close-to-open, and open-to-close. For the Supreme Court ruling on June 16th, 2014, the event occurred in the morning of the 16th, after the stock market opened. In the appendix, we classify this ruling as “open-to-close” meaning that we will use the CDS spread change from 9:30am EDT on Monday the 16th to roughly 4pm EST on Monday the 16th, and the ADR returns from 9:30am EDT on Monday the 16th to 4pm EDT on Monday the 16th. If we had instead classified the event as “close-to-close,” we would compare the 4pm EDT close on Friday the 13th to the 4pm EDT close on Monday the 16th. The “close-to-open” and “open-to-open” windows are defined in a similar way.

We choose our sample of non-events to be a set of two-day default probability changes and stock returns (based on closes), non-overlapping, at least two days away from any event, and at least two days away from any of the “excluded events.” “Excluded events” are legal rulings that we do not use, but also exclude from our sample of “non-events.” For three of the legal rulings, we could not find any contemporaneous media coverage, and are therefore unable to determine when the event was known to market participants. For one legal ruling, we could not find the ruling itself, only references to it in media coverage. One of the legal rulings was issued on the Friday in October 2012 shortly before “Superstorm Sandy” hit New York, and another the night before Thanksgiving.⁸² Finally, one of the legal rulings was issued at the beginning of an oral argument, in which Argentina’s lawyers may have revealed information about Argentina’s intentions. We exclude this day because it violates our identification assumptions. For the heteroskedasticity-based identification strategy we employ, removing these legal rulings increases the validity of our identifying assumption that the variance of shocks induced by legal rulings is higher on event

⁸¹For events occurring outside of daylight savings time in the eastern time zone, the local stocks close at 5pm ART (3pm EST), while the ADRs use 4pm EST. We make no attempt to correct for this.


⁸²The ruling issued the night before Thanksgiving is problematic in several ways (see the appendix for details).

days than non-event days. However, our results are robust to including these days in the set of non-events.

3.3.3 Summary of Events and Non-Events

In Figure 3.1, we plot the two-day change in the 5-year default probability and the two-day return of MSCI Argentina index over our sample. Small data points in gray/light are non-events and the maroon/dark dots cover event windows in which a US court made a legal ruling regarding Argentina's debt. In most of our analysis, and in this plot, we use two-day return windows. As a result, there is some risk that other shocks occurred during the event window. In Figure 3.1, the event labeled "1" is affected by this issue. In our analysis that uses small window sizes, this event is no longer an outlier. The details on each event can be found in Appendix A. In appendix 5.3.5, we present a similar figure for the different sectors of the Argentine economy, the exchange rate, and Mexican and Brazilian CDS changes and equity returns.

Figure 3.1: Default Probability Change and Equity Returns during Events and Non-Events



Event Number	Two-Day Window End Date	ΔD (%)	Equity Return (%)
1	November 27, 2012	4.47	1.49
2	November 29, 2012	-10.78	8.94
3	December 05, 2012	-6.44	1.45
4	December 07, 2012	-0.58	2.13
5	January 11, 2013	3.61	-0.78
6	March 04, 2013	-5.43	10.24
7	March 27, 2013	2.68	-2.32
8	August 26, 2013	2.39	-3.16
9	October 04, 2013	0.06	0.23
10	October 08, 2013	-1.55	0.58
11	November 19, 2013	0.01	-4.29
12	January 13, 2014	2.48	-0.39
13	June 17, 2014	12.70	-7.57
14	June 24, 2014	-5.75	2.24
15	June 27, 2014	6.10	-3.70
16	July 29, 2014	10.11	-0.91

Notes: This figure plots the change in change in the risk-neutral probability of default and returns on the MSCI Argentina Index on event and non-event two-day windows. Each event and non-event day is a two-day event or non-event as described in the text. The numbers next to each maroon/dark dot references each event-day in the table below the figure. The procedure for classifying events and non-events is described in the text.

3.3.4 Exchange Rates

Argentina has capital controls, and its official exchange rate has diverged from the “blue market” exchange rate. Argentina also imposes deposit requirements on foreigners who own local securities. One consequence of these capital controls is that is very unprofitable for foreigners to purchase local Argentine stocks. Instead, foreigners who wish to invest in Argentine companies purchase ADRs. Argentine citizens can also use ADRs, as a means of circumventing capital controls. By purchasing local shares, converting them to ADRs, and then selling them in dollars in the U.S., Argentine citizens can gain access to US dollar currency without government approval (Auguste et al. (2006)). The convertibility of ADRs effectively establishes a shadow exchange rate. We find (in unreported results) that the implied exchange rate computed using ADR and local stock market prices does not vary significantly across firms. In our results, we report an “ADR Blue Rate,” which computes the implied exchange rate for each of the six firms in the MSCI Argentina index, and weighs them using the weights of that index, described previously.⁸³

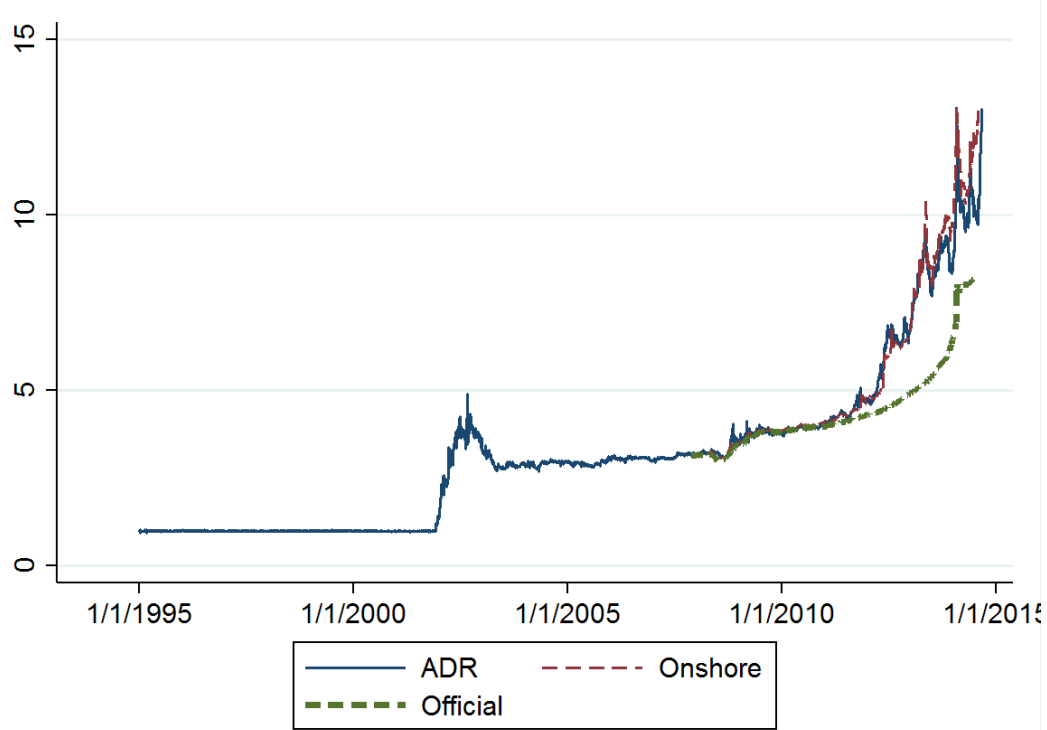
There is a second way to measure the “blue market” exchange rate, which is to poll currency dealers in Argentina. For Argentine households and firms who cannot purchase dollars from the government at the official rate, and cannot execute the ADR-based currency conversion, these dealers are one way to secure dollars. Datastream, a data provider, polls these dealers and computes a “blue rate” based on their responses.

In Figure 3.2, we show the time series of the official exchange rate, the ADR-based blue rate, and the “Onshore” blue rate computed by Datastream.

The recent divergence between the ADR-based blue rate and the onshore blue rate coincides with the rise in the default probabilities experienced by Argentina. In our empirical results, we

⁸³To compute the ADR blue rate, we need prices for both the ADRs and the corresponding locally traded Argentine stocks. As a result, the ADR blue rate is available only on days when both markets are open. This results in smaller sample in our regressions.

Figure 3.2: Official, Onshore Blue, and ADR Exchange Rates (ARS/USD)



Notes: This figure plots the three versions of the ARS/USD exchange rate. ADR refers to the version of the “blue market” rate calculated by comparing the ADR share price in dollars with the underlying local stock price in pesos, Onshore is the exchange rate available through Argentine FX bureaus, and Official is the government’s official exchange rate.

attempt to estimate whether increases in the default probability caused the blue rate to diverge from the official rate, and whether increases in the default probability caused the ADR blue rate to diverge from the onshore blue rate. We find statistically significant evidence that increases in the default probability cause the blue rate to diverge from the official rate, immediately after a legal ruling. We do not find statistically significant evidence that increases in default probability cause the ADR blue rate to diverge from the onshore blue rate.

3.4 Framework

Our goal is to estimate the causal effect of sovereign default on equity returns. The key identification concerns are that stock returns might have an effect on default probabilities, and that

unobserved common shocks might affect both the probability of default and stock returns. In our context, one example of the former issue is that poor earnings by large Argentine firms might harm the fiscal position of the Argentine government, and therefore alter the probability of default. An example of the latter issue is a shock to the market price of risk, which could cause both CDS spreads and stock returns to change.

We consider these issues through the lens of a simultaneous equation model (following Rigobon and Sack (2004)). While our actual implementation uses multiple assets and controls for various market factors, for exposition we discuss only a single asset, r_t , and the change in the risk-neutral probability of default, ΔD_t , and ignore constants.⁸⁴ For exposition, we will refer to this asset, r_t , as the equity market. The model we consider is

$$\Delta D_t = \gamma r_t + \kappa_D F_t + \varepsilon_t \quad (3.1)$$

$$r_t = \alpha \Delta D_t + \kappa F_t + \eta_t \quad (3.2)$$

where F_t is an unobserved factor that moves both the probability of default and equity returns, ε_t is a shock to the default probability, and η_t is a shock to the equity market return.⁸⁵ The goal is to estimate the parameter α , the impact of a change in the probability of default on equity market returns. If one were to simply run the regression in equation 3.2 using OLS, the coefficient estimate would be

$$\begin{aligned} \hat{\alpha} &= \frac{\text{cov}(\Delta D_t, r_t)}{\text{var}(\Delta D_t)} \\ &= \alpha + (1 - \alpha\gamma) \frac{\kappa(\kappa_D + \gamma\kappa)\sigma_F^2 + \gamma\sigma_\eta^2}{(\kappa_D + \gamma\kappa)^2\sigma_F^2 + \gamma^2\sigma_\eta^2 + \sigma_\varepsilon^2} \end{aligned} \quad (3.3)$$

⁸⁴This is equivalent to treating abnormal returns and abnormal default probability changes as observed. Abnormal returns are the excess returns after projecting the return on to observable factors, and abnormal default probability changes are defined similarly. In our econometrics, we account for the estimator error associated with this projection when computing standard errors.

⁸⁵We assume these shocks and unobserved factors are independent.

where σ_ε^2 is the variance of the default probability shock, σ_η^2 is the variance of equity return shock, and σ_F^2 is the variance of the common shock.⁸⁶ There are two sources of bias: simultaneity bias and omitted variable bias. The simultaneity bias exists if $\gamma \neq 0$ and $\sigma_\eta > 0$, and omitted variable bias exists if $\kappa \neq 0$, $\kappa_D \neq 0$, and $\sigma_F > 0$. In order for the OLS regression to be unbiased, equity market returns must have no effect on default probabilities and there must be no omitted common shocks. These assumptions are implausible in our context, but we present this OLS regression in section 3.4.1 for comparison purposes.

We can rely on more plausible assumptions by adopting an event study framework (see, for instance, Kuttner (2001) or Bernanke and Kuttner (2005)). We can make the identifying assumption that changes to Argentina's probability of default on during the event windows (time periods in which a US court makes a ruling in the case of the *Republic of Argentina v. NML Capital*) are driven exclusively by those legal rulings, or other idiosyncratic default probability shocks (ε_t).⁸⁷ Under this assumption, we can directly estimate equation 3.2 using OLS on these ruling days. We will pursue this strategy in section 3.4.3.

Finally, we will consider a heteroskedasticity-based identification strategy, following Rigobon (2003) and Rigobon and Sack (2004). This does not require the complete absence of common and idiosyncratic shocks during event windows. This strategy instead relies on the weaker identifying assumption that the variances of the common shocks F_t and equity return shocks η_t are the same on non-event days and event days, whereas the variance of the shock to the probability of default ε_t is higher on event days than non-event days. The variance of ε_t is assumed to be higher because of the impact of the legal rulings, which are modeled as ε_t shocks under the exclusion restriction. Under this assumption, we can identify the parameter α by comparing the covariance matrices of abnormal returns and default probability changes on event days and non-event days.

⁸⁶This expression is the one presented in Rigobon and Sack (2004).

⁸⁷Rigobon and Sack (2004) demonstrate that the event study makes the identification assumption that on event days, the ratio of the default shock variance σ_ε to both the equity return shock σ_η and the common shock σ_F is infinite. If this assumption holds, we can see from equation 3.3 that $\hat{\alpha}$ is an unbiased estimator of α .

In order to see how we can use this strategy to identify our key parameter of interest, we first solve for the reduced form of equations 3.1 and 3.2:

$$\begin{aligned} r_t &= \frac{1}{1-\alpha\gamma} ((\alpha\kappa_D + \kappa) F_t + \eta_t + \alpha\varepsilon_t) \\ \Delta D_t &= \frac{1}{1-\alpha\gamma} ((\kappa_D + \gamma\kappa) F_t + \gamma\eta_t + \varepsilon_t) \end{aligned}$$

We can then divide all days in our sample into two types of days, event (E) and non-event (N) days. For each of the two types of days $j \in \{E, N\}$, we can estimate the covariance matrix of $[r_t, \Delta D_t]$, denoted Ω_j :

$$\Omega_j = \begin{bmatrix} \text{var}_j(r_t) & \text{cov}_j(r_t, \Delta D_t) \\ \text{cov}_j(r_t, \Delta D_t) & \text{var}_j(\Delta D_t) \end{bmatrix}$$

Calculating these moments using the reduced form equations, we can then write the covariance matrix on day type j as

$$\Omega_j = \left(\frac{1}{1-\alpha\gamma} \right)^2 \begin{bmatrix} \alpha^2 \sigma_{\varepsilon,j}^2 + \sigma_{\eta}^2 + (\alpha\kappa_D + \kappa)^2 \sigma_F^2 & \alpha \sigma_{\varepsilon,j}^2 + \gamma \sigma_{\eta}^2 + ((\alpha\kappa_D + \kappa)(\gamma\kappa + \kappa_D)) \sigma_F^2 \\ \alpha \sigma_{\varepsilon,j}^2 + \gamma \sigma_{\eta}^2 + ((\alpha\kappa_D + \kappa)(\gamma\kappa + \kappa_D)) \sigma_F^2 & \sigma_{\varepsilon,j}^2 + \gamma^2 \sigma_{\eta}^2 + (\kappa_D + \gamma\kappa)^2 \sigma_F^2 \end{bmatrix}$$

We can then define the difference in the covariance matrices on event and non-event days as $\Delta\Omega = \Omega_E - \Omega_N$, which simplifies to

$$\Delta\Omega = \lambda \begin{bmatrix} \alpha^2 & \alpha \\ \alpha & 1 \end{bmatrix} \quad (3.4)$$

where $\lambda = \left(\frac{\sigma_{\varepsilon,E}^2 - \sigma_{\varepsilon,N}^2}{(1-\alpha\gamma)^2} \right)$. This provides us with a number of ways to estimate the coefficient of interest α that we will examine in section 3.4.5. Although we have described our framework where the only asset is the market, in appendix section 5.3.4 we demonstrate how an equivalent system can be derived in a multi-asset framework.

The heteroskedasticity-based approach is our preferred estimation procedure. If the identification assumptions required for the OLS or event study hold, the heteroskedasticity-based strategy will also be valid, but the converse is not true. However, the event study approach does have one advantage over the heteroskedasticity approach (as we have implemented it). For the heteroskedasticity approach, we use two-day event days, because those are the smallest uniformly-sized windows that all of our events can fit into. However, as discussed earlier, all of our events can in fact fit into smaller windows (open-close, open-open, close-open, or close-close), but those windows are not the same size for each event. Using the event study approach, we present results defined using these narrower windows. If the identification assumptions required for this event study hold, this approach may have more power than the heteroskedasticity-based approach.

We begin by presenting the OLS estimates, as point for comparison with our subsequent results.

3.4.1 OLS Estimates

In this section, we assume the OLS identifying assumption: $F_t = 0$ and $\gamma = 0$ in equations 3.1 and 3.2 above. The model can be written as

$$r_t = \alpha \Delta D_t + \eta_t$$

where α is the coefficient of interest, and $Cov(\Delta D_t, \eta_t) = 0$. We can estimate this equation with OLS.

In our actual implementation, we include a constant and the vector of controls X_t discussed in section 3.3.1. We estimate the OLS model for the returns of the MSCI Argentina Index, our three ADR industry groups, and our three measures of the exchange rate.

The results in table 3.2 imply that a 1% increase in the probability of default is associated with a 0.46% fall in the MSCI Argentina Index. In appendix table 5.4, we see increases in the probability of an Argentine default are associated with increases in Brazilian and Mexican CDS

Table 3.2: OLS Results

	(1)	(2)	(3)	(4)
	Index	Banks	Non-Financial	Real estate
ΔD	-46.41***	-49.89***	-35.74***	-11.53
Robust SE	(6.318)	(8.111)	(6.077)	(8.023)
95% CI	[-61.8,-28.7]	[-66.0,-34.0]	[-51.7,-17.4]	[-27.9,6.1]
Observations	413	413	413	413
R-squared	0.389	0.345	0.305	0.071

	(5)	(6)	(7)
	FX (ADR)	FX (On.)	FX (Official)
ΔD	30.26***	11.55***	3.863
Robust SE	(4.573)	(4.019)	(1.740)
95% CI	[18.1,43.3]	[3.4,19.4]	[-13.1,8.6]
Observations	368	413	413
R-squared	0.182	0.029	0.072

Notes: This table reports the results for the OLS regression of equity returns and foreign exchange (FX) rate on changes in the risk-neutral default probability (ΔD) and the covariates discussed in the text. The column headings denote the outcome variable. Index is the MSCI Argentina Index, Banks is our equally weighted index of Argentine bank ADRs, Non-Financial is our equally weighted index of Argentine non-financial ADRs, and Real Estate is our equally weighted index of Argentine real estate holding companies. FX (ADR) is the ARS/USD exchange rate derived from the ratio of ADR prices (in USD) to the price of the underlying equity (in ARS). FX (On.) is the ARS/USD exchange rate offered by onshore currency dealers. FX (Official) is the exchange rate set by the Argentine government. The coefficient on ΔD is the effect on the percentage returns of an increase in the 5-year risk-neutral default probability from 0% to 100%, implied by the Argentine CDS curve. Standard errors and confidence intervals are computed using the stratified bootstrap procedure described in the text. The underlying data is based on the two-day event windows and non-events described in the text. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

spreads and declines in the Brazilian and Mexican equity markets. This correlation points to the importance of omitted common factors. In our heteroskedasticity-based estimates presented below and in the appendix, we show that the legal rulings have no measurable impact on Brazilian and Mexican CDS or equity markets. The method we use to construct standard errors and confidence intervals is discussed below in section 3.4.4. For the OLS estimates, it is essentially equivalent to heteroskedasticity-robust standard errors and confidence intervals based on first-order approximations.

3.4.2 Case Study: Announcement

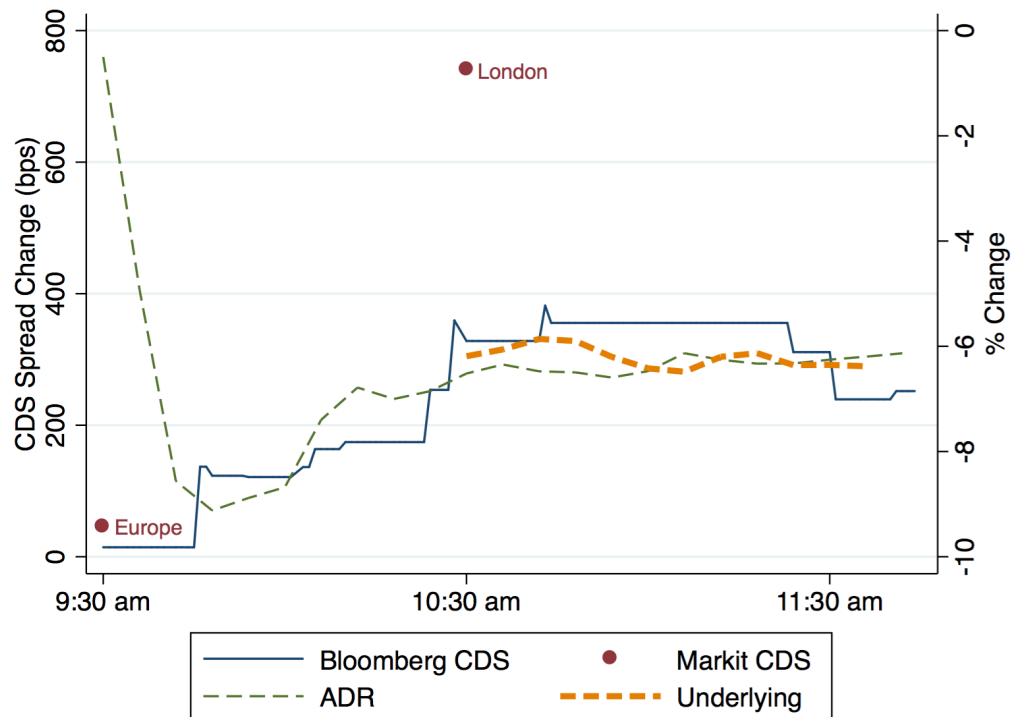
We begin our discussion of the event study approach with a single event. On June 16, 2014, the U.S. Supreme Court denied two appeals and a petition from the Republic of Argentina. This denial had several effects. First, it allowed the holdouts to pursue discovery against all of Argentina's foreign assets, not just those in the United States. Second, the court declined to review Judge Griesa's interpretation of the *pari passu* clause and his orders demanding equal treatment. The denial of Argentina's petition meant that Judge Griesa could prevent the Bank of New York, the payment agent on Argentina's restructured bonds, from paying the coupons on those bonds, unless Argentina also paid the holdouts. Because Argentina had previously expressed its unwillingness to pay the holdouts, this news meant that Argentina was more likely to default.

This event is ideal for our purposes because we are able to precisely determine the time the news was released. The Supreme Court announces multiple orders in a single public session, and simultaneously provides copies of those orders to the press. Prior to releasing the official opinion, the court announces the order. SCOTUSBlog, a well-known legal website that provides news coverage and analysis of the Supreme Court, had a "live blog" of the announcements on that day. At 9:33am EST, SCOTUSBlog reported that "Both of the Argentine bond cases have been denied. Sotomayor took no part."⁸⁸ At 10:09am, the live blog stated that Argentina's petition had been denied. At 10:11am, the live blog provided a link to the ruling.

In figure 3.3, we plot the returns of the Argentine ADRs, underlying equities and the percentage change in the sovereign CDS spread. The ADRs begin trading in New York at 9:30am but trading of the underlying local stocks does not begin in Argentina until 10:30am EST. To compare the returns on the underlying local stocks with the ADRs, we weight the return of the underlying stocks according to their weight in the MSCI Argentina Index of ADRs. Finally, we include 1-

⁸⁸See http://live.scotusblog.com/Event/Live_blog_of_orders_and_opinions__June_16_2014?Page=0

Figure 3.3: Event Data from June 16, 2014, 9:30-11:30am EST



Notes: This figure plots the change in the CDS spread over the previous night's close ("CDS", left axis, basis points), the change in the price of the MSCI Argentina Index against the previous night's close ("ADR", right axis, %), and the change in the price of an index of Argentine stocks with the same weightings as in the ADR based MSCI Index over the previous night's close ("Underlying", right axis, %). The Markit CDS points are labeled with the name of the reporting market, with European markets reporting at 9:30am EST and London Markets reporting at 10:30am EST.

minute interval data on the mid-price (halfway between the bid and ask) on 5-year Republic of Argentina Senior Credit Default Swaps, from Bloomberg.⁸⁹

From 9:30am to 10:30am, the MSCI index of ADRs fell 6% and the sameday 5-year CDS spread (measured by Markit) increased by 693 basis points (bps), implying a 9.8% increase in the risk-neutral probability of default over the next 5 years. When the Argentine stock market opened,

⁸⁹We believe that the CDS data ultimately comes from the "screen" of an inter-dealer broker. It is not clear that these rates represent the actual market in the CDS. We use the Bloomberg data only for this figure, and rely on Markit data for our regressions. During the one-hour interval from 9:30am to 10:30am, the Markit sameday CDS spread increased by substantially more than the CDS spread reported by Bloomberg, although both changes are large relative to typical hourly movements.

Table 3.3: Summary Statistics

Day Type	Mean ΔD (%)	SD ΔD (%)	Mean Equity (%)	Equity SD (%)
Event	0.88	6.13	0.26	4.50
Non-Event	-0.01	1.79	0.01	3.17

Day Type	Cov($\Delta D, Equity$)	# of Days
Event	-21.26	16
Non-Event	-2.37	397

Notes: This table reports the mean default probability change, the standard deviation of default probability changes, the mean MSCI Argentina Index return, the standard deviation of that return, and the covariance of default probability changes and that return during events and non-events. The underlying data is based on the two-day event windows and non-events described in the text.

the index of equities opened 6.2% lower than it closed the previous night. Under the standard event study assumptions, this implies that a 1% increase in the probability of default causes a 0.63% fall in ADR prices, and virtually no change in the ADR-based blue rate.

3.4.3 Event Study

Following the discussion in section 3.4, we present the results of three event studies. Each event study uses the same identification assumptions, outlined above. The first event study uses two-day windows around events. We begin by presenting summary statistics for the returns of the MSCI Argentina Index and the changes in 5-year risk-neutral default probabilities, during event windows and non-event windows.

Our event study methodology follows the one described in Campbell et al. (1997). Let N denote the set of non-event days, and let $L1 = |N|$. We first estimate the factor model on the non-event days,

$$r_{i,t} = \mu_i + \omega_i^T X_t + v_{i,t},$$

and generate a time series of abnormal returns, $\hat{r}_{i,t} = r_{i,t} - \hat{\mu}_i - \hat{\omega}_i^T X_t$, where X_t is the vector of controls discussed in section 3.3.1. We also estimate the variance of the abnormal returns associ-

ated with the factor model (assuming homoskedastic errors), $\hat{\sigma}_i^2 = \frac{1}{L1} \sum_{t \in N} \hat{v}_{i,t}^2$. We next estimate a factor model for the change in the probability of default, ΔD_t , and create a time series of abnormal default probability changes, \hat{d}_t . We then classify our event days into three categories, based on the abnormal default probability change during the event window. Let σ_d denote the standard deviation of the abnormal default probability changes. If the probability increases by at least σ_d , we label that day as an “higher default” event. If the probability decreases by at least σ_d , we label that event as a “lower default” event. If the default probability change is less, in absolute value, than σ_d , we label that as a “no news” event.

For each type of event, we report the cumulative abnormal return and cumulative abnormal default probability change over all events of that type (higher default, lower default, no news). We also report two statistics that are described in Campbell et al. (1997). In this event study (but not the next one we discuss), which does not aggregate returns across different ADRs, the two statistics are identical, up to a small sample size correction. Define $E_{\{h,l,n\}}$ as the set of event days of each type. The first statistic, $J1$, is computed, for event type j and ADR i , as

$$J1_{ij} = \frac{\sum_{t \in E_j} \hat{r}_{i,t}}{\sqrt{|E_j| \hat{\sigma}_i^2}}.$$

Under the null hypothesis that the events have no effect on the stock returns, $J1_{ij}$ is asymptotically distributed as a standard normal. However, because we have so few events in each category, asymptotic normality will be a poor approximation, if the abnormal returns are themselves far from normal. This is one reason we prefer the variance-based estimators discussed in the next section.

The second statistic, $J2$, is nearly identical to $J1$ for this event study (they will be different in the next event study we describe). For each event, we can define a standardized cumulative abnormal return,

$$z_{i,t} = \sqrt{\frac{|E_j| - 4}{|E_j| - 2}} \frac{\hat{r}_{i,t}}{\sqrt{\hat{\sigma}_i^2}},$$

where the first term represents a small-sample correction. The statistic $J2$ is defined as

Table 3.4: Standard Event Study: Index

Shock Type	# Events	CAR (%)	ΔD (%)	J_1	J_2
Higher Default	8	-18.12	44.21	-2.46**	-2.45**
No News	3	-2.56	-0.03	-0.57	-0.56
Lower Default	5	23.59	-29.48	4.05***	4.04***

Notes: CAR indicates cumulative abnormal return over the event windows, ΔD is the change in the risk-neutral probability of default, and the test statistics J_1 and J_2 are described in the text and in Campbell et al. (1997), pp. 162. A shock type of higher default indicates that this event raised the default probability by more than one two-day standard deviation, a shock type of lower default indicates that this event lowered the default probability by more than one two-day standard deviation, and a shock type of no news indicates a day with a legal ruling in which the default probability did not move at least one two-day standard deviation in either direction. The underlying data is based on the two-day event windows and non-events described in the text. The p-values are the p-values for a two-sided hypothesis test assuming normality. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

$$J2_{ij} = \frac{\sum_{t \in E_j} z_{i,t}}{\sqrt{|E_j|}}.$$

This statistic is also asymptotically standard normal under the null hypothesis, subject to the same caveat about return normality. In table 3.4, we present these two statistics for the MSCI Argentina Index.

The results of this event study are broadly similar to the OLS estimates. In the 8 event days where the default probability significantly increased, the cumulative increase in the default probability was 44.21% and the stock market experienced a cumulative abnormal return of -18.12%. Assuming a linear relationship between default probabilities and equity returns, this implies that a 1% increase in the probability of default causes a 0.41% fall in the stock market. During the 5 days where the default probability significantly declined, the cumulative fall in the default probability was 29.48% with a cumulative abnormal return of 23.6%. This implies a 1% fall in the probability of default causes an 0.80% rise in the stock market. Treating the movements symmetrically and adding together the absolute value of the change in default probability and cumulative abnormal returns, we find that a 1% increase in the probability of default causes a 0.57% fall in the equity market. While the large window sizes used in this study raise concerns about the validity of the identification assumptions, we will see that this estimate is very close to the results we find from our heteroskedasticity-based estimates.

The next event study we present uses four different window sizes, discussed earlier. Our data set includes one additional event (17 instead of 16), because one of the two-day windows in fact contained two separate legal rulings on consecutive days. Conceptually, the event study is almost identical, except that we must study each type of event (higher default, lower default, no news) for each window size. That is, we separately estimate abnormal returns and abnormal default probability changes for each window size $s \in S$, the set of window sizes. We classify events based on the standard deviation of abnormal default probability changes for the associated window size. Let E_{js} denote an event of type j (higher default, lower default, no news) with window size s (close-to-close, open-to-open, close-to-open, and open-to-close). The abnormal return $\hat{r}_{i,t,s}$ is the abnormal return for ADR i at time t with window size s , and $\hat{\sigma}_{is}^2$ is the variance of the abnormal returns for that window size. The $J1$ statistic is computed as

$$J1_{ij} = \frac{\sum_{s \in S} \sum_{t \in E_{js}} \hat{r}_{i,t,s}}{\sqrt{\sum_{s \in S} |E_{js}| \hat{\sigma}_{is}^2}}.$$

Asymptotically, subject to the same caveats mentioned previously, this statistic is distributed as a standard normal. The second statistic, $J2$, is constructed in a similar fashion. However, the standardized cumulative abnormal returns are now defined with respect to the event window size,

$$z_{i,t,s} = \sqrt{\frac{|E_{js}| - 4}{|E_{js}| - 2}} \frac{\hat{r}_{i,t,s}}{\sqrt{\hat{\sigma}_{is}^2}},$$

and the $J2$ statistic is

$$J2_{ij} = \frac{\sum_{s \in S} \sum_{t \in E_{js}} z_{i,t,s}}{\sqrt{\sum_{s \in S} |E_{js}|}}.$$

This statistic is also, subject to the same caveats, asymptotically standard normal. It is not the same as the $J1$ statistic, because of the heterogeneity in window size. If the cumulative abnormal returns occur mostly in shorter windows (which have smaller standard deviations), the $J2$ statistic will be larger in absolute value than the $J1$ statistic. If the reverse is true, the $J1$ statistic will be larger. The size of the window may depend in part on the court releasing the opinion, the urgency

Table 3.5: Heterogeneous-Window Event Study: Index

Shock Type	# Events	CAR (%)	ΔD (%)	J_1	J_2
Higher Default	5	-13.70	14.61	-3.71***	-3.34***
No News	7	-5.26	4.91	-1.06	-1.03
Lower Default	5	20.14	-32.58	5.31***	5.26***

Notes: CAR indicates cumulative abnormal return over the event window, ΔD is the change in the risk-neutral probability of default, and the test statistics J_1 and J_2 are described in the text and in Campbell et al. (1997), pp. 162. This study pools events across different window sizes (open-open, open-close, close-open, close-close). A shock type of higher default indicates that this event raised the default probability by more than one standard deviation, where the standard deviation is defined for non-events with the same window size. A shock type of lower default indicates that this event lowered the default probability by more than one standard deviation, and a shock type of no news indicates a day with a legal ruling in which the default probability did not move at least one standard deviation in either direction. The underlying data is based on the event windows and non-events described in the text, and uses the narrowest windows possible with our data and uncertainty about event times. The p-values are the p-values for a two-sided hypothesis test assuming normality. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

with which the opinion was required, and other endogenous factors. It is not obvious whether the J_1 or J_2 statistic should be preferred. Fortunately, the results presented in table 3.5 using the two statistics are similar.

In the 5 event days where the default probability significantly increased, the cumulative probability of default rose 14.61% and the stock market had a cumulative abnormal return of -13.7%. This estimate implies that a 1% increase in the probability of default causes a 0.94% fall in equity returns. During the 5 days where the default probability significantly declined, the cumulative fall in the default probability was 32.58% with a cumulative abnormal equity return of 20.14%. This implies a 1% fall in the probability of default causes an 0.62% rise in the stock market. When we again treat up and down movements symmetrically, we find that a 1% increase in the probability of default causes a 0.72% fall in the equity market.

Finally, we present “IV-style” event study. This study uses the two-day events and non-events described previously. The second stage equation we wish to estimate is equation 3.2, discussed above. The instrument we use is $1(t \in E)\Delta D_t$ (and $1(t \in E)$), where E is the set of event days and $\mathbf{1}(\cdot)$ is the indicator function. The first-stage regression is

$$\Delta D_t = \chi 1(t \in E)\Delta D_t + \rho 1(t \in E) + \mu_D + \omega_D^T X_t + \tau_t,$$

where τ_t is a composite of the three unobserved shocks (ε_t, F_t, v_t) on the non-event days. Under the event study assumptions, the unobserved shocks ε_t and F_t (in the second stage) are not correlated with the change in the default probability on event days.

The IV-style event study has the advantage that of offering an interpretable coefficient, $\hat{\alpha}$, that estimates the change in stock prices given a change in the default probability. It also takes into account the magnitude of the default probability changes on each event day, whereas the event studies discussed earlier treat each event in a category equally. However, it is not *a priori* clear that the impact of the default probability on stock returns should be linear, and therefore not obvious that this approach is superior to the two-day event study. The similarity of the two results suggests linearity is not a bad assumption. Because the IV-style event study uses two-day event windows, it requires stronger identification assumptions than the heterogeneous-window event study. The standard errors and confidence intervals for this approach are described in section 3.4.4, below.

Using this method, we find that a 1% increase in the probability of default causes a 0.55% fall in the MSCI Argentina Index, a 0.59% fall in financial stocks, a 0.33% fall in industrial stocks, and only a 3% fall in REIT-eligible stocks. While the coefficient differences are suggestive, we will defer a discussion of whether they are significantly different from one another until section 3.4.5. We also find that a 1% increase in the probability of default causes a 0.35% depreciation of the ADR blue rate, a 0.16% depreciation of the onshore blue rate, and has no effect on the official exchange rate.

3.4.4 Standard Errors and Confidence Intervals

To construct confidence intervals for our coefficient estimates, we employ the bootstrap procedure advocated by Horowitz (2001). The advantage of this procedure is that it offers “asymptotic refinements” for the coverage probabilities of tests, meaning that it is more likely to achieve the desired rejection probability under the null hypothesis. Our estimators (except for the OLS) are effectively based on a small number of the data points (the events), and therefore these refinements may provide significant improvements over first-order asymptotics. As a practical matter, our confidence

Table 3.6: IV-Style Event Study

	(1)	(2)	(3)	(4)
	Index	Banks	Non-Financial	Real Estate
ΔD	-55.18***	-59.14***	-31.61***	-3.124
Robust SE	(10.84)	(13.92)	(10.41)	(13.76)
95% CI	413	413	413	413
Observations	[-78.6,-28.4]	[-82.5,-32.3]	[-56.6,-7.2]	[-27.0,23.3]
1st Stage F-Stat	211.2	211.2	211.2	211.2

	(5)	(6)	(7)
	FX (ADR)	FX (On.)	FX (Official)
ΔD	35.07***	15.60***	-0.188
Robust SE	(7.577)	(6.896)	(3.000)
95% CI	368	413	413
Observations	[12.0,59.5]	[4.7,27.0]	[-4.1,4.1]
1st Stage F-Stat	212.2	211.2	211.2

Notes: This table reports the results for the variance-based estimator estimated as the ratio of $\lambda\alpha$ to λ . This estimator is called the “CDS-IV” estimator because it depends on the excess variance of the CDS spread on event days. The column headings denote the outcome variable. Index is the MSCI Argentina Index, Banks is our equally weighted index of Argentine bank ADRs, Non-Financial is our equally weighted index of Argentine non-financial ADRs, and Real Estate is our equally weighted index of Argentine real estate holding companies. FX (ADR) is the ARS/USD exchange rate derived from the ratio of ADR prices (in USD) to the price of the underlying equity (in ARS). FX (On.) is the ARS/USD exchange rate offered by onshore currency dealers. FX (Official) is the exchange rate set by the Argentine government. The coefficient on ΔD is the effect on the percentage returns of an increase in the 5-year risk-neutral default probability from 0% to 100%, implied by the Argentine CDS curve. Standard errors and confidence intervals are computed using the stratified bootstrap procedure described in the text. The underlying data is based on the two-day event windows and non-events described in the text. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

intervals are in almost all cases substantially wider than those based on first-order asymptotics. Nevertheless, these “asymptotic refinements” are still based on asymptotic arguments, and there is no guarantee that they are accurate for our data. We also find (in unreported results) that our confidence intervals for our coefficient of interest, α , are similar to confidence intervals constructed under normal approximations, using a bootstrapped standard error.

We use 1000 repetitions of a stratified bootstrap, re-sampling with replacement from our set of events and non-events, separately, so that each bootstrap replication contains 16 events and 397 non-events.⁹⁰ In each bootstrap replication, we compute the (asymptotically pivotal) t-statistic

⁹⁰The number of events and non-events listed apply to the ADRs. The exchange rates have a slightly different number of events and non-events, due to holidays, missing data, and related issues.

$t_k = \frac{\hat{\alpha}_k - \hat{\alpha}}{\hat{\sigma}_k}$, where $\hat{\alpha}$ is the point estimate in our actual data sample, $\hat{\alpha}_k$ is the point estimate in bootstrap replication k , and $\hat{\sigma}_k$ is the heteroskedasticity-robust standard deviation estimate of $\hat{\alpha} - \alpha$ from bootstrap sample k . We then determine the 2.5th percentile and 97.5th percentile of t_k in the bootstrap replications, denoted $\hat{t}_{2.5}$ and $\hat{t}_{97.5}$, respectively. The reported 95% confidence interval for $\hat{\alpha}$ is $[\hat{t}_{2.5}\hat{\sigma} + \hat{\alpha}, \hat{t}_{97.5}\hat{\sigma} + \hat{\alpha}]$, where $\hat{\sigma}$ is the heteroskedasticity-robust standard deviation estimate of $\hat{\alpha} - \alpha$ from our original data sample. We construct 90% and 99% confidence intervals in a similar fashion, and use them to assign asterisks in our tables.⁹¹ In the tables, we report the 95% confidence interval and the heteroskedasticity-robust standard error from our dataset ($\hat{\sigma}$).

3.4.5 Variance-based Analysis

Our final set of analysis is based on the difference between the covariance matrices in equation 3.4. There are several potential ways to estimate α based on $\Delta\Omega$. Two such estimators, which we call the CDS-IV and Returns-IV estimators, respectively, are defined as

$$\hat{\alpha}_{CIV} = \frac{\Delta\Omega_{1,2}}{\Delta\Omega_{2,2}} = \frac{\text{cov}_E(\Delta D_t, r_t) - \text{cov}_N(\Delta D_t, r_t)}{\text{var}_E(\Delta D_t) - \text{var}_N(\Delta D_t)}$$

$$\hat{\alpha}_{RIV} = \frac{\Delta\Omega_{1,1}}{\Delta\Omega_{1,2}} = \frac{\text{var}_E(r_t) - \text{var}_N(r_t)}{\text{cov}_E(\Delta D_t, r_t) - \text{cov}_N(\Delta D_t, r_t)}$$

As shown in Rigobon and Sack (2004), these estimators can be implemented in an instrumental variables framework. More generally, equation 3.4 provides us with three moment conditions.

$$\Delta\Omega_{1,1} - \lambda\alpha^2 = 0, \tag{3.5}$$

$$\Delta\Omega_{1,2} - \lambda\alpha = 0, \tag{3.6}$$

$$\Delta\Omega_{2,2} - \lambda = 0. \tag{3.7}$$

The GMM estimator uses all three moment conditions.

⁹¹These asterisks represent an “equal-tailed” test that $\alpha \neq 0$.

The Returns-IV estimator uses an “irrelevant instrument” under the null hypothesis that $\alpha = 0$. The estimator $\hat{\alpha}_{RIV}$ is the ratio of the sample estimates of $\Delta\Omega_{1,1}$ and $\Delta\Omega_{1,2}$, both of which are zero in expectation under the null hypothesis. The denominator, $\Delta\Omega_{1,2}$, is the covariance between the default probability, which is the variable being instrumented for, and the instrument. Under the null hypothesis, this covariance is zero, meaning that the instrument is irrelevant. As a result, the behavior of the $\hat{\alpha}_{RIV}$ estimator under the null hypothesis is not characterized by the standard IV asymptotics, and our confidence intervals will not have the correct coverage probabilities.⁹² The CDS-IV estimator does not suffer from this issue. The estimator $\hat{\alpha}_{CIV}$ is based on the ratio of the sample estimates of $\Delta\Omega_{1,2}$ and $\Delta\Omega_{2,2}$. Under the null hypothesis that $\alpha = 0$ and $\lambda > 0$, the CDS-IV instrument is still relevant, and the standard asymptotics for $\hat{\alpha}_{CIV}$ apply. The GMM estimator, $\hat{\alpha}_{GMM}$, which uses all three moments, can be thought of as a geometric average of the CDS-IV and Returns-IV estimators. When $\alpha \neq 0$, using all three moments is advantageous because it takes advantage of all available information and makes over-identifying tests possible. However, under the null hypothesis that $\alpha = 0$, using the Returns-IV estimator in any way is problematic. The two-step GMM procedure, implemented using standard asymptotics to estimate the optimal weighting matrix, would generally not correctly estimate the variances, because of the irrelevant instrument. As a result, the weight matrix might effectively place excessive weight on the Returns-IV estimator, relative to the CDS-IV estimator, and end up providing problematic results. For these reasons, we use the CDS-IV estimator as our preferred estimation procedure. We report the results for the other two methods in the appendix.

The CDS-IV instrument is relevant under the assumption that $\lambda > 0$. We formally test this assumption using a one-sided F-test of the ratio of $(\Omega_E)_{22}$ to $(\Omega_N)_{22}$, which is the ratio of the variance of changes in the default probability on event days and non-event days. We test the alternate hypothesis that this ratio is greater than 1 (implying $\lambda > 0$) against the null hypothesis that it is equal to one. In our sample, this F-statistic is 11.78, well above the 99th percentile,

⁹²When α is near, but not equal, to zero, weak identification asymptotics may be a better characterization of the sample distribution of $\hat{\alpha}_{RIV}$.

Table 3.7: CDS-IV

	(1)	(2)	(3)	(4)
	Index	Banks	Non-Financial	Real Estate
ΔD	-54.52***	-59.21***	-29.59**	-0.601
Robust SE	(11.60)	(14.88)	(11.15)	(14.73)
95% CI	[-81.1,-25.3]	[-84.3,-31.0]	[-55.9,-2.9]	[-28.3,31.4]
Observations	413	413	413	413
1st Stage F-Stat	171.9	171.9	171.9	171.9

	(5)	(6)	(7)
	FX (ADR)	FX (On.)	FX (Official)
ΔD	35.20**	16.17***	-0.579
Robust SE	(7.951)	(7.375)	(3.212)
95% CI	[7.5,62.1]	[4.3,28.7]	[-2.6,4.9]
Observations	368	413	413
1st Stage F-Stat	178.7	171.9	171.9

Notes: This table reports the results for the variance-based estimator estimated as the ratio of $\lambda\alpha$ to λ . This estimator is called the “CDS-IV” estimator because it depends on the excess variance of the CDS spread on event days. The column headings denote the outcome variable. Index is the MSCI Argentina Index, Banks is our equally weighted index of Argentine bank ADRs, Non-Financial is our equally weighted index of Argentine non-financial ADRs, and Real Estate is our equally weighted index of Argentine real estate holding companies. FX (ADR) is the ARS/USD exchange rate derived from the ratio of ADR prices (in USD) to the price of the underlying equity (in ARS). FX (On.) is the ARS/USD exchange rate offered by onshore currency dealers. FX (Official) is the exchange rate set by the Argentine government. The coefficient on ΔD is the effect on the percentage returns of an increase in the 5-year risk-neutral default probability from 0% to 100%, implied by the Argentine CDS curve. Standard errors and confidence intervals are computed using the stratified bootstrap procedure described in the text. The underlying data is based on the two-day event windows and non-events described in the text. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

one-sided, bootstrapped critical value of 1.98.⁹³ The relevance of the CDS-IV instrument is also suggested by the weak-identification F-test of Stock and Yogo (2005) (not to be confused with the F-test for $\lambda > 0$) shown in table 3.7. In table 3.7, we present the results of our CDS-IV estimation. The standard errors and confidence intervals use the bootstrap procedure described previously.

We find that increases in the 5-year risk-neutral default probability cause statistically and economically significant declines in the MSCI Argentina Index, bank ADRs, and non-financial ADRs. In contrast, we do not find a statistically significant effect on the ADRs of Argentine real-estate holding companies, although we cannot rule out economically significant effects. The point esti-

⁹³We use the bootstrapping procedure for pivotal statistics described by Horowitz (2001), and in our section on standard errors and confidence intervals.

mates in table 3.7 are very close to those reported in table 3.6, with a 1% increase in the probability of default causing a 0.55% fall in the broad index, a 0.59% fall in bank stocks, a 0.30% fall in non-financial stocks, and 0.006% fall in real estate stocks. Increases in the probability of default also cause significant depreciation of the peso blue rate, measured with ADRs or by polling onshore currency dealers. However, there was no corresponding same-day change in the official exchange rate.⁹⁴ The increase in the risk-neutral default probability from 40% to 100%, which is roughly what Argentina experienced, would cause more than a 30% fall in the ADR index, by our estimates.

We formally test whether financial ADRs fall more than industrial ADRs, whether bank ADRs fall more than real estate ADRs, and whether industrial ADRs fall more than real estate ADRs. We construct these one-sided t-tests using the same bootstrap procedure for pivotal statistics discussed earlier. We find that both the bank ADRs and industrial ADRs fall more than the real estate ADRs (at the 95% confidence level), but cannot reject the hypothesis that bank and industrial ADRs respond equally to changes in the default probability.

We also test whether the blue rates depreciate relative to the official rate, in response to increases in the default probability. We find that the difference between the onshore blue rate and the official exchange rate is significant at the 95% level, while the difference between the ADR-based blue rate and the official rate is significant at the 90% level. We cannot rule out the hypothesis that the two blue rates respond equally to increases in the default probability.

Our results are consistent with the hypothesis that Argentina's default would cause significant harms to Argentina's economy. In the next section, we examine which sectors of the Argentine economy are more adversely affected.

3.5 Cross-Sectional Evidence

In this section, we examine which firm characteristics are associated with larger or smaller responses to the default shocks. The cross-sectional pattern of responses across firms can help shed

⁹⁴We cannot rule out the possibility that the official exchange responds subsequently. This non-result is consistent with our identifying assumption that actions by Argentina's government, unrelated to the legal rulings, are not more likely on event days than other days.

light on the mechanism by which sovereign default affects the economy. First, we examine how different industries respond to default shocks. Second, we examine the heterogeneous firm responses to an increase in the probability of default, through the lens of different theories on the channel by which sovereign default affects the broader economy.

In their seminal contribution, Eaton and Gersovitz (1981) argue that the reason governments repay their debt is to maintain their reputation and ensure continued access to international bond markets. Because this access allows governments to smooth income fluctuations, it is valuable and is generally sufficient to guarantee repayment.⁹⁵ Because of the threat of attachment from outstanding creditors, Argentina had not issued a new international bond in thirteen years and was unlikely to do so soon. This suggests that the effect of default that we measure is different than the reputational mechanism posited in Eaton and Gersovitz (1981).⁹⁶ Instead, this points to the importance of alternative theories of sovereign default costs, examined in the literature following Bulow and Rogoff (1989b). We will attempt to examine the empirical relevance these hypothesized costs of sovereign default by examining whether four groups of firms are particularly affected by default: exporters, importers, banks, and foreign-owned companies.

First, motivated by Bulow and Rogoff (1989a), we will examine whether or not firms that are reliant on exports are particularly hurt. Bulow and Rogoff (1989a) argue that in the event of a sovereign default, foreign creditors can interfere with a country's exports. We would therefore expect exporters to underperform in response to increases in the probability of default. Using aggregate data, Rose (2006) and others have found support for this channel. Second, motivated by Mendoza and Yue (2012), we will examine whether or not firms that are reliant on imported intermediate goods are particularly hurt by default. Mendoza and Yue (2012) argue that a sovereign

⁹⁵Tomz (2007) provides a historical account to argue in favor of the reputational model of sovereign debt. English (1996) argues that the experience of US states in the 1840s provides evidence in favor of the reputational model of sovereign default by arguing that no direct sanctions were available to creditors. The Eleventh Amendment prevents foreign creditors from suing US states to receive payments on defaulted debt, constitutionally guaranteed interstate free trade prevents foreign creditors from locking defaulters out of trade markets, and the US federal government prevents foreign creditors from using force to collect on the debt. English demonstrates that defaulting states are unable to borrow again for a number of years, concluding that the concern for maintaining a reputation for repayment is therefore the only explanation for continued repayment.

⁹⁶We are not providing evidence against the importance of the the type of reputational concerns in Eaton and Gersovitz (1981), but rather arguing that this particular default is not likely to be affected by such concerns.

default reduces aggregate output because firms cannot secure financing to import goods needed for production, and so are forced to use domestic intermediate goods, which are imperfect substitutes. This would lead us to expect firms that are relatively more reliant on imported intermediate goods would underperform in response to a default shock. Third, motivated by Gennaioli et al. (2014), Acharya et al. (2014), Bolton and Jeanne (2011), Bocola (2013) and Perez (2014), we will examine whether financial firms are more adversely affected. While these papers are not explicitly about whether banks are hurt more than other firms, they posit that the aggregate decline in output following a sovereign default occurs because of the default's effect on bank balance sheets. This leads to a reduction in financial intermediation and a reduction in aggregate production. If this argument or something like it were correct, we would expect banks to be hurt disproportionately by an increase in the probability of default. Finally, motivated by Cole and Kehoe (1998), we examine whether foreign-owned firms underperform following an increase in the probability sovereign default. Cole and Kehoe (1998) argue that even if the loss of a reputation for repayment alone is not sufficient to motivate countries to repay their debt, if their "general reputation" is lost by defaulting on sovereign debt, foreigners would become less willing to trust the defaulting government. This theory would lead us to expect increases in the risk of sovereign default to cause foreign-owned firms to underperform, as foreigners perceive a higher risk that Argentina will act disreputably in other arenas, such as investment protection. Our empirical approach is similar to several papers in the literature studying the cross-section of firms' responses to identified monetary policy shocks, using an event study for identification. Bernanke and Kuttner (2005) study U.S. stock market data and find that the response of various industry portfolios to a monetary policy shock is proportional to that industry's CAPM beta. Put differently, the ensemble of shocks that generate returns outside of the event windows have a similar cross-sectional pattern of returns to the monetary policy shock. Gorodnichenko and Weber (2013) find that a measure of the stickiness of firms' prices is correlated with the squared magnitude of firms' response to squared monetary policy shocks. We apply similar strategies to our context. First, we explore the abnormal returns for various industries in response to a default probability shock, controlling for the abnormal return of the Argentine mar-

ket. Second, we form portfolios based on firm characteristics suggested by theory and then study the abnormal returns of those portfolios, again controlling for the abnormal return of the Argentine market.

Our procedures are motivated by a modified version of the model in equations 3.2 and 3.1. We derive both models from a single underlying system of equations, presented in the appendix, section 5.3.4. The modified version of the those equations has the return of the Argentine market index, $r_{m,t}$, on the right-hand side, in addition to the observable factors X_t and unobservable factors F_t . We denote the return of a particular stock or portfolio as $r_{i,t}$:

$$\Delta D_t = \mu_D + \omega_D X_t + \gamma r_{i,t} + \gamma_m r_{m,t} + \kappa_D F_t + \varepsilon_t \quad (3.8)$$

$$r_{i,t} = \mu_i + \omega_i X_t + (\alpha_i - \beta_i \alpha_m) \Delta D_t + \beta_i r_{m,t} + \kappa_i F_t + \eta_{i,t}. \quad (3.9)$$

The parameter α_m is the response of the Argentine market index, $r_{m,t}$, to the default shock. For the purposes of our study, this two equation system has exactly the same form as the system described in section 3.4. The Argentine market return, $r_{m,t}$, is an observable common factor, no different from the S&P 500 or other observable factors in X_t . The Rigobon (2003) procedure, applied to this system, identifies the coefficient $(\alpha_i - \beta_i \alpha_m)$, which can be interpreted as the excess sensitivity of the portfolio to the default shock, above and beyond what would be expected from the Argentine market's exposure to the default shock, and the sensitivity of the portfolio to the Argentine market. In this sense, our approach generalizes the CAPM-inspired analysis of Bernanke and Kuttner (2005) to a model with multiple exogenous shocks.

We begin by studying the response of industry portfolios to default shocks, controlling for the response of the Argentine market. To increase our sample size of firms, we use local Argentine stock returns, rather than ADRs. We convert the local stock returns, denominated in pesos, into dollars using the ADR-based blue rate described previously. For stocks with ADRs, the converted return will be nearly identical to the ADR return.⁹⁷ The use of the ADR-based exchange rate

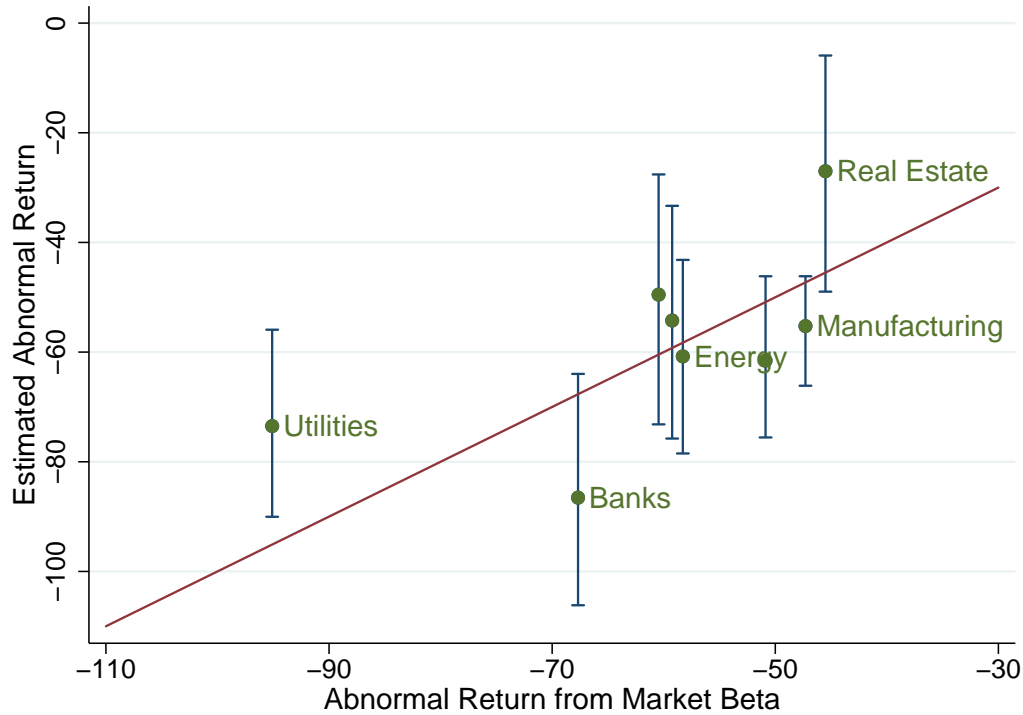
⁹⁷As mentioned previously, the implied exchange rate between various stock-ADR pairs does not vary substantially across firms.

requires that both the New York and Buenos Aires stock markets be open, which reduces the size of our sample. However, the events in our sample remain the same, with one exception.⁹⁸

We group these firms into equal-weighted industry portfolios, using the industry definitions described in section 3.3.1. We also construct an equal-weighted index of all of the firms in our sample, which is restricted to firms passing a data quality test also described in section 3.3.1. We use this equal-weighted index as our measure of the Argentine market return. In figure 3.4 and table 3.8 below, we display estimates of the excess sensitivity of the industry portfolios to the default shock, using the CDS-IV estimator and the bootstrapped confidence intervals described in the previous sections.

⁹⁸The treatment of the events around Monday, June 23, 2014 is different when using ADR and local stock data, as the result of an Argentine holiday on June 20th. The ADR data uses a two-day window from the close of June 20 to the close of June 24, whereas the local stock data uses the close of Jun 19 to June 23.

Figure 3.4: Estimated Response to Default Shocks: Industries



Notes: Industry classifications are based on the Fama-French 12 industry categories with the modifications described in section 3.3.1. The labels for chemical firms, diverse firms, non-durables producers, and telecoms are suppressed. On the the x-axis, we plot the expected abnormal return for each portfolio, calculated as the beta of each long-short portfolio on the index times α_M , the effect of an increase in the probability of default in the abnormal return of the index. On the y-axis, we plot the sum of the expected abnormal return and $(\alpha_i - \beta_i \alpha_M)$, the additional sensitivity of each portfolio to an increase in the probability of default. Values above the line indicates that the portfolio over-performed following increases in the probability of default, relative to what would be implied by the portfolio's market beta. Values below the line indicate under-performance. The ranges indicate bootstrapped 90% confidence intervals.

Three industries (banks, real estate, and utilities) stand out as over- or under-sensitive to default shocks. However, care must be taken when interpreting the results. First, the confidence intervals around these estimates are very wide. Our point estimates suggest that a 10% increase in the probability of default would cause bank stocks to fall by roughly 2% more than would be expected, given their beta to the Argentine index, and would cause real estate stocks and utilities to fall by 2% less than would be expected. However, the standard deviation of these estimates is almost 1%, and only the utilities' out-performance is significant at the 95% confidence interval. The

Table 3.8: Cross-Section: Industry Returns

	(1)	(2)	(3)	(4)	(5)	(6)
	Index	Banks	Chemicals	Diverse	Energy	Manufact.
ΔD	-63.91*** (8.173)	-18.86 (8.277)	5.006 (13.95)	10.94 (19.06)	-2.495 (9.932)	-7.965 (7.299)
Index Beta		1.059	.927	.946	.912	.74
95% CI	[-81.3,-44.3]	[-42.0,6.4]	[-21.0,30.3]	[-19.1,36.2]	[-23.7,18.4]	[-21.4,3.4]
Observations	356	358	358	358	358	358
F_Stat	259.5	173.5	173.5	173.5	173.5	173.5

	(7)	(8)	(9)	(10)	(11)
	Non-Durables	Non-Financial	Real Estate	Telecoms	Utilities
ΔD	-7.726 (6.884)	2.82 (1.852)	18.5 (9.692)	-10.67 (9.91)	21.61** (8.763)
Index Beta	.739	1.008	.712	.796	1.488
95% CI	[-25.2,6.8]	[-3.7,9.0]	[-6.7,43.0]	[-26.9,7.6]	[1.6,41.1]
Observations	358	358	358	358	358
F-Stat	173.5	173.5	173.5	173.5	173.5

Notes: This table reports the results for the “CDS-IV” estimator. The column headings denote the outcome variable. Index is the equal-weighted index of local equities in table 3.1. The industry classifications are based on Fama-French with modifications described in section 3.3.1. The coefficient on ΔD is the effect on the percentage returns of an increase in the 5-year risk-neutral default probability from 0% to 100%, implied by the Argentine CDS curve. Index beta is the coefficient on the equal-weighted index of Argentine local equities, as described in section 3.5. Standard errors and confidence intervals are computed using the stratified bootstrap procedure described in the text. The underlying data is based on the two-day event windows and non-events described in the text. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

uncertainty around our point estimates is driven by the small number of events we study, and the idiosyncratic variation in stocks’ response to the different legal announcements. Second, our confidence intervals have not been adjusted for multiple testing; the fact that one industry has significant over- or under-performance at the 95% confidence level is not surprising, given the number of tests being performed.

That said, our point estimates are economically large. Taken at face value, our results suggest that as Argentina went from a 40% to 100% probability of defaulting, its banks’ value fell by 11% more (in dollar terms) than would have been expected, given a 38% fall in the dollar value of the equal-weighted index. The excessive sensitivity of bank stocks to default risk is consistent with

the theories of Gennaioli et al. (2013, 2014), Bocola (2013), and Bolton and Jeanne (2011). We interpret our data as providing suggestive evidence for these theories⁹⁹.

We next consider which characteristics of non-financial firms are associated with over- or under-performance in response to default shocks. As discussed in section 3.3.1, we form zero-cost, long-short portfolios of non-financial firms based on the export intensity of their primary industry, the import intensity of their primary industry, whether they are a listed subsidiary of a foreign firm, and whether they have an associated ADR. An import-intensive industry is not the opposite of an export-intensive one; some industries are classified as neither import nor export intensive, whereas others are both import and export intensive.¹⁰⁰ Finally, we compare firms with and without ADRs.

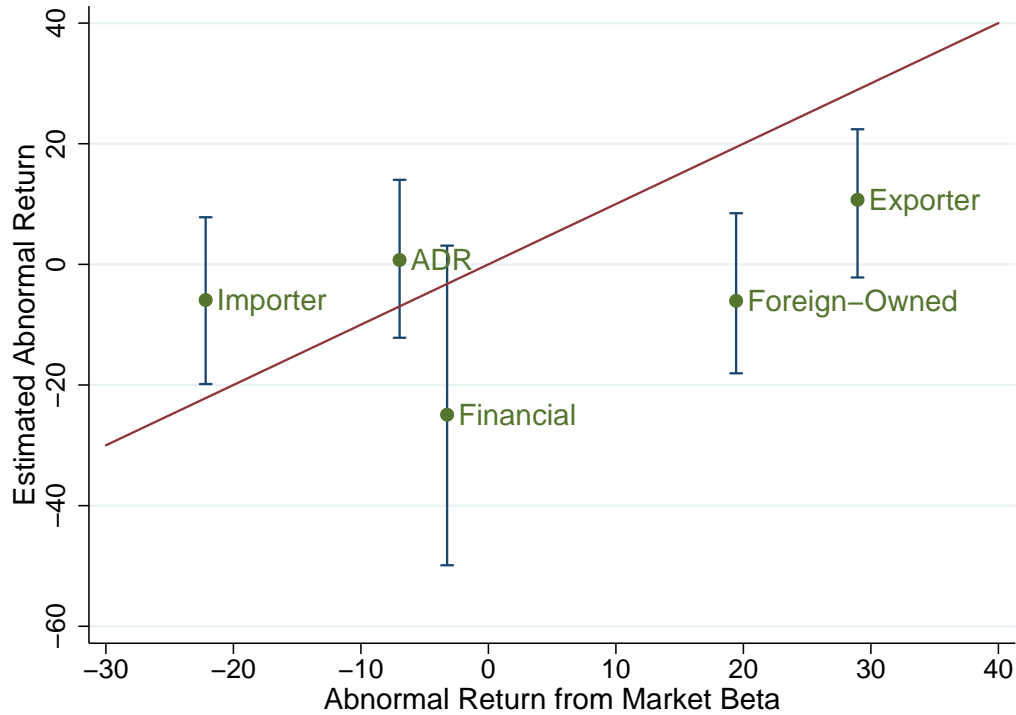
In these portfolios, we equally weight firms within the “long” and “short” groups. For example, we classify 12 of our 26 non-financial firms¹⁰¹ as high export intensity, and 14 of 26 as low export intensity. We equally weight these firms, so that the “long” portfolio has a 1/12 weight on each high export intensity firm, and the short portfolio has a 1/14 weight on each low export intensity firm. We then form the long-short portfolio, and determine whether the portfolio over- or under-performs after a default shock, using the CDS-IV estimator and bootstrapped confidence intervals discussed previously.

⁹⁹Regarding the outperformance of utilities, one market participant suggested to us that pressure on the Argentine government’s foreign reserves, exacerbated by the default, might lead them to liberalize utility prices. In the status quo, underpriced electricity (for example) leads to over-consumption, which results in excessive importation of utilities’ inputs. Excessive imports reduce the Argentine government’s foreign reserves position, and their inability to borrow makes it difficult to replenish these reserves. This story is one possible explanation for why utility companies could indirectly benefit from a sovereign default, relative to other companies.

¹⁰⁰The correlation in our sample of non-financial firms is 0.16.

¹⁰¹We in fact have 27 non-financial firms, but one is a technology firm. The technology firm’s industry classification did not exist when the input/output table we use to construct the data was generated.

Figure 3.5: Estimated Response to Default Shocks: Long-Short



Notes: Each label denotes a zero-cost long short portfolio. Exporter is a portfolio going long export-intensive non-financial firms and short non-export-intensive non-financial firms. Importer is defined equivalently for importers. Financial goes long banks and short non-financial firms. Foreign-owned firms goes long non-financial firms with a foreign parent and short domestically-owned non-financial firms. ADR goes long non-financial firms with an American Depository Receipt and short non-financial firms without one. The data sources are described in section 3.3.1. On the the x-axis, we plot the expected abnormal return for each portfolio, calculated as the beta of each long-short portfolio on the index times α_M , the effect of an increase in the probability of default in the index. On the y-axis, we plot the sum of the expected abnormal return and $(\alpha_i - \beta_i \alpha_M)$, the additional sensitivity of each portfolio to an increase in the probability of default. Values above the line indicates that the portfolio over-performed following increases in the probability of default, relative to the abnormal return implied by the portfolio's market beta. Values below the line indicate underperformance. The ranges indicate bootstrapped 90% confidence intervals.

In figure 3.5 and table 3.9, we find that firms whose primary industry is export-intensive underperform, while firms whose primary industry is import intensive over-perform expectations, given their exposure to the equal-weighted index and the index's response to the default probability shock. We find that our long-short exporter portfolio underperforms 0.18% more for each 1% increase in the risk-neutral probability of default than would be expected given the portfolio's

Table 3.9: Cross-Section: Long-Short Portfolios

	(1)	(2)	(3)	(4)	5)
	Exporter	Importer	Financial	Foreign-Owned	ADR
ΔD	-18.24**	16.27*	-21.68	-25.47***	7.698
	(9.401)	(7.912)	(10.01)	(7.974)	(7.527)
Index Beta	-.453	.347	.0507	-.304	.109
95% CI	[-32.8,-3.4]	[-.1,32.4]	[-50.8,10.3]	[-40.2,-8.0]	[-7.9,24.1]
Observations	358	358	358	358	358
F-Stat	173.5	173.5	173.5	173.5	173.5

Notes: This table reports the results for the “CDS-IV” estimator. The column headings denote the outcome variable. Each column is a zero-cost long short portfolio. Exporter is a portfolio going long export-intensive non-financial firms and short non-export-intensive non-financial firms. Importer is defined equivalently for importers. Financial goes long banks and short non-financial firms. Foreign-owned firms goes long non-financial firms with a foreign parent and short domestically-owned non-financial firms. ADR goes long non-financial firms with an American Depository Receipt and short non-financial firms without one. The data sources are described in section 3.3.1. The coefficient on ΔD is the effect on the percentage returns of an increase in the 5-year risk-neutral default probability from 0% to 100%, implied by the Argentine CDS curve. Index beta is the coefficient on the equal-weighted index of Argentine local equities, as described in section 3.5. Standard errors and confidence intervals are computed using the stratified bootstrap procedure described in the text. The underlying data is based on the two-day event windows and non-events described in the text. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

loading on the market index. However, our results about import intensive firms are not robust to changes in the portfolio formation threshold. In unreported results, we find that using a 4% threshold for import intensity, instead of 3%, results in all of the utilities being reclassified from high import intensity to low import intensity. Because the utilities responded far less to default shocks than their beta would predict, their reclassification is sufficient to change the sign of the results. In contrast, we find that the results for exporters are qualitatively robust to variations in the threshold.

The over- or under-performance of the export and import portfolios is not an ideal test of the theories. In the context of the Bulow and Rogoff (1989a) theory, if we do not observe that exporting firms under-perform, it may be because the firms we observe are not the ones whose exports would be seized, or because our export-intensive and non-export-intensive firms also differ on some other characteristic that predicts over- or under-performance¹⁰². The reverse is also true; a significant

¹⁰²Essentially, an omitted variables problem

result does not necessarily validate the theory, but might instead be found because of a correlation across firms between exporting and some other firm characteristic.

We also find that non-financial foreign subsidiaries, of which there are seven, substantially underperform relative to non-financial firms that are not foreign subsidiaries. The long-short portfolio falls 0.25% more in response to a 1% increase in the risk-neutral probability of default than would be expected given the portfolio's loading on the index. This result is consistent with the general reputation theory of Cole and Kehoe (1998), implying that foreign investors become more reluctant to invest, although there are many other possible interpretations. We do not find that non-financial firms with an ADR substantially under- or out-perform non-financial firms without ADRs.

We interpret this cross-sectional analysis as lending modest support to several of the theories in the existing literature that try to understand the costs of sovereign default. The theories are not exclusive; sovereign default may harm the financial system, impede trade, and weaken a country's reputation in many areas. Our estimates are insufficiently precise to reject any of these theories, or speak to their quantitative importance. Nevertheless, our approach does have the advantage over the existing literature that we can pinpoint the direction of causality, from sovereign default to performance, in a way that would be very difficult using aggregate or annual data.

3.6 Interpretation

We begin by describing an imaginary “ideal experiment” to identify the causal effect of default on economic activity. We will then discuss the ways in which our research design does and does not approach this ideal. We will offer alternative interpretations of the effect of the legal rulings, and discuss their implications for the interpretation of our results. We also discuss several additional aspects of Argentina's situation that are relevant.¹⁰³

The ideal experiment would randomly induce one of two otherwise-identical groups of countries to default on their debt. These groups of countries would have characteristics similar to those of typical sovereign borrowers. The treatment (default) would be randomly assigned, so that it

¹⁰³Alfaro (2014) examines the implications of the legal rulings on future sovereign debt restructurings.

would be uncorrelated with the underlying state of the countries' economies. The treatment would induce a country to default, but would otherwise neither encourage nor impair other actions by that country's government, firms, or households. The null hypothesis in this experiment is that default does not affect economic activity. The alternative hypothesis is that default impairs economic activity, through some unspecified channel.

We emphasize the idea of "inducing" a country to default because we view default as a choice of the government. For the purposes of understanding why sovereign borrowers repay their debts, we would like to understand the consequences of them choosing not to repay. These consequences include the effects of whatever mitigating actions a country might take after having decided to default. These consequences also include the effects of firms, households, and other agents changing their behavior as a result of the default. The government's actions could include renegotiating with creditors, finding other means to borrow, balancing budgets via taxes or reduced spending, taking actions that affect the convertibility of the currency, among other actions. When we refer to the causal effects of sovereign default, we include the anticipated effects of whatever policies the government is expected to employ as a result of having defaulted.

Our research design differs from this ideal experiment in a variety of ways. First, we study Argentina, a country whose experience with sovereign debt is very different from most other countries. Argentina is in some sense in default for the entirety of our sample, depending on the definition of "default." It has an unusual currency regime. Argentina defaulted for convoluted legal reasons. Additionally, the way in which Argentina acts to mitigate the consequences of its default might be different from the way other countries would respond in similar circumstances. Second, there is the issue of whether the default is exogenous to Argentina's economic circumstances. Third, our outcome variables are not perfect measures of economic activity. Fourth, these legal rulings might have effects on firms' stock prices, through channels other than changes in the likelihood of default (the exclusion restriction may not hold). If the legal rulings compelled Argentina to repay a large amount of money, relative to its economy or foreign reserves, then firms' stock prices

might fall due to the burdens of debt repayment and associated reduction in economic activity, rather than through any default-related effects.

In the remainder of this section, we will discuss each of these issues in more detail.

3.6.1 The Options Available to Argentina

It is not clear that Argentina was forced to default. Prior to these legal rulings, Argentina had several feasible courses of action with respect to its restructured debt and the holdouts. It could maintain the status quo, in which it was subject to attachment orders and other actions by the holdouts, while it continued to pay its restructured creditors. It could attempt to negotiate with the holdouts, and completely resolve its default. Finally, it could choose to default on its restructured creditors.

The cumulative effect of these legal rulings changed the menu of options available to Argentina. The status quo option, in which Argentina continued to pay its restructured bondholders without paying the holdouts, became infeasible.¹⁰⁴ Instead, Argentina could make payments on its debt, which would be divided between the restructured bondholders and the holdouts according to the “ratable payment” formula devised by Judge Griesa. Alternatively, it could attempt to negotiate with the holdouts, to avoid defaulting on its restructured bondholders. Finally, it could default on the restructured bondholders.

Argentina effectively chose the third option (default). It made a payment to the Bank of New York Mellon (BNYM), the trustee for its restructured bonds, that was sufficient to pay the restructured bondholders, without paying anything to the holdouts. Judge Griesa’s order prohibited BNYM from forwarding this payment to the restructured bondholders, and Argentina missed a coupon payment. After the 30-day grace period, Argentina was declared in default by the rating agencies.

As of this writing, how the situation will be resolved is unclear. One recent proposal involves replacing BNYM with another, non-U.S. trustee, who would not be subject to the U.S. courts’

¹⁰⁴This infeasibility might be temporary or permanent— it is not clear as of this writing.

orders, and could continue to pay the restructured bondholders. Another complication concerns the treatment of euro-denominated bondholders, whose coupon payments are included in the amount held by BNYM. These bondholders have argued that BNYM acted contrary to Belgian and U.K. law, and that they should continue to be paid.

The cumulative effect of the legal rulings raised the probability of default on the restructured bonds and/or payment of the holdouts, relative to the probability that the status quo would continue. If Argentine firms would be affected by payment of the holdouts, holding default or no default fixed, then the exclusion restriction of our experiment would not hold.

One possibility is that the legal rulings might change the probability or size of a settlement with the holdouts, and this could affect the firms. Under the null hypothesis, if the government somehow repaid the holdouts without fiscal consequences (say, using a gift from abroad), there would be no effect on firms. In reality, because the government would need to raise taxes, cut spending, or borrow to repay the holdouts, an increase in the probability or size of a settlement with the holdouts could harm firms.

To get a sense of whether this is reasonable, we consider the extent to which the bonds owned by the holdouts appreciated, on our event days. Based on preliminary findings, we believe that the increase in the expected value of the holdout bonds is dwarfed by the cumulative losses of the Argentine firms.¹⁰⁵ This suggests that if, in expectation, the entirety of the burden of repayment fell on these firms, that would only explain a small part of the stock market declines. A very large “multiplier” for the loss of equity value associated with the debt burden would be required for this argument to apply.

More generally, the legal rulings could have had other effects. However, Argentine corporations are legally independent from the Argentine government, and their assets cannot be attached by the holdouts.¹⁰⁶ The ruling affects them only to the extent that it changes the behavior of the Argentine government or other actors. This still leaves open several possible effects. The legal rulings could

¹⁰⁵These calculations are available upon request.

¹⁰⁶There was litigation regarding whether the Argentine central bank qualified as independent from a legal perspective, but no such litigation for any of the companies listed in the stock index.

have provoked the government of Argentina into a sequence of actions unrelated to sovereign default. They could have influenced the probability that the current government of Argentina stays in power in the next election. The legal rulings could have changed the law regarding sovereign debt generally.

We can muster evidence against this last effect. In the appendix, section 5.3.2, we show that the stock markets and sovereign CDS spreads of Brazil and Mexico did not respond to these legal rulings (our estimates are close to zero, and relatively precise). This is in contrast to the OLS estimates, which show that those financial variables are correlated with the Argentine risk-neutral probability of default, presumably due to common shocks affecting Latin America or emerging markets more generally. This evidence suggests that, whatever changes to sovereign debt law occurred as the result of these rulings, they did not materially impact other Latin American countries that issue debt in New York.

However, we cannot rule out every possible channel through which these rulings might have affected firms, other than via sovereign default. Ex-post, it appears that the primary response of the Argentine government to these rulings was default. We are unaware of any direct consequences for Argentine firms. Consistent with this interpretation, S&P did not downgrade any Argentine firms immediately upon the sovereign's default (Standard and Poor's 2014a). However, it subsequently downgraded a variety of firms, arguing that deteriorating economic conditions reduced those firms' credit quality (Standard and Poor's 2014b).

3.6.2 How Much Would Argentina Have to Repay?

To meet the precise demands of the courts, Argentina needed to pay its litigating creditors only \$1.5 billion. However, the \$1.5 billion owed to the litigating creditors was only around 10% of the estimated \$15 billion holdout debt outstanding.¹⁰⁷ Presumably, if Argentina paid NML and its co-litigants in full, the other holdout creditors would demand repayment on similar terms. Even

¹⁰⁷See Gelpern (2014a).

if we assume that Argentina would need to pay the full \$15 billion, that represented only 3% of GDP,¹⁰⁸ and 45% of foreign currency reserves.¹⁰⁹

However, it is possible that if Argentina did indeed pay the holdouts in full, it would then owe the restructured creditors a large payment as well. During its 2004-2005 debt restructuring, Argentina sought to convince its creditors that the unilateral offer it made was the best offer they would ever receive. Argentina included a “Rights Upon Future Offers” (RUFO) clause in the bond prospectus of the restructured debt.¹¹⁰ The RUFO clause entitled restructuring creditors to terms at least as good as anything holdouts would receive in the future:

Under the terms of the Pars, Discounts and Quasi-pars, if following the expiration of the Offer until December 31, 2014, Argentina voluntarily makes an offer to purchase or exchange or solicits consents to amend any Eligible Securities not tendered or accepted pursuant to the Offer, Argentina has agreed that it will take all steps necessary so that each holder of Pars, Discounts or Quasi-pars will have the right, for a period of at least 30 calendar days following the announcement of such offer, to exchange any of such holder’s Pars, Discounts or Quasi-pars for the consideration in cash or in kind received in connection with such purchase or exchange offer or securities having terms substantially the same as those resulting from such amendment process, in each case in accordance with the terms and conditions of such purchases, exchange offer or amendment process.¹¹¹

In other words, if Argentina made an offer to the holdouts that was better than what the restructured creditors received, the restructured creditors would be entitled to the better deal, provided the

¹⁰⁸The CIA World Factbook reports Argentina’s 2013 GDP as \$484.6 billion. However, this calculation uses the official exchange rate, which may overstate the size of Argentina’s economy.

¹⁰⁹The CIA World Factbook reports Argentina’s foreign exchange and gold reserves at \$33.65 billion as of December 31, 2013.

¹¹⁰Olivares-Caminal (2013) refers to this as the “most favored creditor clause.”

¹¹¹Full bond prospectus available at <http://www.sec.gov/Archives/edgar/data/914021/000095012305000302/y04567e424b5.htm>

offer occurred before December 31, 2014. Argentina claimed that this RUFO clause meant that it could not pay NML the \$1.5 billion owed without incurring hundreds of billions in additional liabilities. There is one crucial word in the RUFO that makes the whole matter more complicated: *voluntarily*. If Argentina offered the holdouts a better deal because US courts would otherwise have blocked its payments to the restructured bondholders, would that be voluntary or involuntary? Indeed, some observers noted that Argentina’s counsel told the Second Circuit Court of Appeals that Argentina “would not voluntarily obey” court rulings to pay the holdouts in full.¹¹² In addition, other commentators noted that the RUFO appeared to have some loopholes, allowing Argentina to potentially settle with the holdouts without triggering the clause.¹¹³ Finally, exchange bondholders could waive their right to exercise the RUFO, and because it takes 25% of exchange bondholders to trigger the clause, the whole issue could have been rendered moot if the exchange bondholders could be persuaded that this was preferable to having their coupon payments blocked.¹¹⁴ Of course, this possibility assumes Argentina would have paid any amount to the holdouts, a questionable proposition given the domestic politics surrounding the holdouts.¹¹⁵

For the purposes of interpreting our results, the RUFO clause complicates matters in two ways. First, if the RUFO clause was binding and could not have been easily circumvented, negotiation with the holdouts may not have been feasible. In this case, it would be correct to say that the legal rulings forced Argentina to default (the simple interpretation offered above). Second, if the RUFO clause was binding, but the legal rulings compelled Argentina to involuntarily pay the holdouts (and therefore circumvented the RUFO clause), they might have made renegotiation feasible when it was not previously feasible. Finally, if the RUFO clause was not binding, it does not alter the interpretation of the rulings discussed above.

The RUFO clause expired on December 31, 2014 but, as of the time of this draft, a settlement has not yet been reached.

¹¹²<http://ftalphaville.ft.com/2013/03/06/1411442/raising-the-rufo-in-argentine-bonds/>

¹¹³See the comment’s from Barclay’s reported in FT Alphaville: <http://ftalphaville.ft.com/2013/03/06/1411442/raising-the-rufo-in-argentine-bonds/>

¹¹⁴See Gelper (2014b).

¹¹⁵See Gelper (2014b).

3.6.3 Are the Legal Rulings Exogenous?

We argue that the rulings of the courts are not influenced by news about the Argentine economy. Formally, the interpretation of the laws in question does not depend on the state of the Argentine economy. Substantively, because the amount required to repay the holdouts in full was small relative to the Argentine economy, news about the Argentine economy's prospects would not materially change their ability to pay. Moreover, even if the judges were responding to economic fundamentals, under the null hypothesis that default does not affect fundamentals, the judges would have no information advantage over market participants. It follows that the effects of economic news on the judges' rulings would be anticipated by the market prior to those rulings, and any response by the market to the judges' rulings would not reflect news about fundamentals.

More subtle interactions between the state of the Argentine economy and the legal rulings might complicate the interpretation of our analysis. For example, if bad news about the Argentine economy causes the market response to the legal rulings to be larger than it otherwise would have been, our estimates will reflect some sort of average effect, where the averaging occurs over states of the economy. Relatedly, the underlying state of the economy might influence the Argentine government's decision about whether to negotiate with the holdouts or default, and therefore interact with the legal rulings to determine the extent to which the default probability and stock prices change. These issues emphasize that our estimates should be considered average treatment effects, relevant to Argentina.

It is important that our event study avoid announcements by the Argentine government, because such announcements might be responding to news about fundamentals, or affect corporations in ways other than through default. In the case of the Supreme Court decision discussed earlier, the Argentine government did not respond immediately to the ruling.¹¹⁶ More generally, we include as events only orders by a judge or judges. We exclude orders that were issued during oral argu-

¹¹⁶See the following Bloomberg story: [Link](#).

ments, because those events also include opportunities for lawyers representing Argentina to reveal information.

Our identifying assumption is that the variance of “legal shocks” is higher on days when a US court rules on the dispute between NML and Argentina while the variance of all other shocks remain the same. However, if in addition to shocks to economic fundamentals, and legal shocks, we imagine that there are “political shocks” which move the probability of Argentina defaulting on its debt, then it could be that the variance of these shocks are higher on event days because the government is more likely to make a pronouncement revealing how likely it is to default following a ruling by Judge Griesa.¹¹⁷

3.6.4 Interpreting Stock Returns

We view the response of stock prices to default shocks through the lens of the Campbell-Shiller decomposition. One reason ADR or local equity prices might decline is that default reduces the expected growth rate of corporate dividends, by harming the Argentine economy. Another reason that prices might decline is that higher default probabilities cause an increase in the required returns of Argentine firms. Because Argentina is small, relative to the world economy, and the ADRs are traded by investors in the U.S., there is no reason to expect that the legal rulings we identify alter the stochastic discount factor of the marginal investor. We also use controls for various risk factors known to predict financial market returns, to isolate the abnormal returns on Argentine ADRs that cannot be attributed to changes in the stochastic discount factor. Based on these arguments, we believe that the negative returns on Argentine ADRs associated with increases in default risk reflect reductions in expected dividend growth.¹¹⁸

¹¹⁷Based on news stories, we believe that such “political shocks” are no more likely on event days than non-event days. In the future, we hope to use news stories to determine a set of dates that correspond to political shocks, and test whether such events are more likely on event days than non-event days. Even if such political events are more likely one event days, our research design is valid. In this case, we would be identifying the causal effect of the rulings on default, inclusive of both the ruling’s direct effect and the Argentine government’s endogenous response. Alternatively, if the political events were more likely on event days but unrelated to the issue of default, or affected firms through some mechanism other than default, our identification would fail. However, there is no apparent reason for political events to be more likely on event days, unless they are related primarily to the legal rulings.

¹¹⁸Theoretically, the returns could also be caused by increases in the exposure of the ADRs to priced risk factors (an increase in “beta”).

Assume that the above arguments are correct, and adverse legal rulings cause reductions in expected ADR dividend growth by increasing the likelihood of default. The most straightforward interpretation of this result is that default will harm the Argentine economy, and this is what reduces expectations of dividend growth. However, there are other possible interpretations that we cannot rule out. Default may lead to changes in the corporate share of income in Argentina, without harming the overall economy. This story, and others like it, cannot be ruled out because our outcome variable is the price of a financial asset that may not be representative of the Argentine economy.

In order to examine the different responses of firms to increases in the probability of sovereign default, we also study the response of local equities. While we convert their prices to dollars using the ADR-implied blue rate, we cannot rule out that the returns may be affected by the risk of capital controls, particularly if only domestic residents own the securities. This would also affect the ADR-implied blue rate. Nevertheless, we do not find a significant difference in the response between how the local equity price of firms with and without ADRs response to increases in the probability of default. This makes us optimistic that the segmentation effects are relatively minor.

3.6.5 Was Argentina Already in Default?

Although the debt exchanges of 2005 and 2010 eventually achieved a participation rate of 91.3%, above the level generally needed by a sovereign to resolve a default and reenter capital markets,¹¹⁹ Argentina remained unable to borrow internationally. This is because the ongoing creditor litigation had resulted in an attachment order, which would allow the holdouts to confiscate the proceeds from a new bond issuance.¹²⁰ However, ratings agencies took a different view, and on June 1, 2005, S&P declared the end of the Argentine default and gave them a sovereign foreign currency credit rating of B-.

There are several complications arising from Argentina's ambiguous international standing. If the costs of default for Argentina were lower than that of a typical sovereign debtor, because Argentina was already unable to borrow in international markets, then our estimates understate

¹¹⁹Hornbeck (2013).

¹²⁰Hornbeck (2013).

the costs for the typical sovereign. On the other hand, because Argentina chose to default despite an ability to pay, the costs might be higher than is typical. Complicating the story further is that the Argentine government was still able to borrow in local markets, and via inter-country loans. In the aggregate, the country of Argentina was able to run a current account deficit, because its households, firms, and even local governments were able to borrow internationally, despite the inability of its federal government to do so. Therefore, even if the federal government of Argentina was in default for our entire sample, it is not clear that (as a country) it was locked out of international markets, before or after the latest default. These complications emphasize the uniqueness of Argentina's circumstances.

3.7 Conclusion

For several decades, one of the most important questions in international macroeconomics has been “why do governments repay their debts?” Using an identification strategy that exploits the timing of legal rulings in the case of *Republic of Argentina v. NML Capital*, we present evidence that a sovereign default significantly reduces the value of domestic firms. We provide suggestive evidence that exporters, banks, and foreign-owned firms are particularly hurt by sovereign default.

4 References

- ACHARYA, V., I. DRECHSLER, AND P. SCHNABL (2014): “A pyrrhic victory? Bank bailouts and sovereign credit risk,” *The Journal of Finance*, 69, 2689–2739.
- ACHARYA, V., H. MEHRAN, AND A. THAKOR (2012): “Caught between Scylla and Charybdis? Regulating bank leverage when there is rent-seeking and risk-shifting,” *CEPR Discussion Paper No. DP8822*.
- ADMATI, A. R. AND P. PFLEIDERER (1997): “Does it all add up? Benchmarks and the compensation of active portfolio managers*,” *The Journal of Business*, 70, 323–350.
- AGHION, P. AND P. BOLTON (1992): “An incomplete contracts approach to financial contracting,” *The Review of Economic Studies*, 59, 473–494.
- AGUIAR, M. AND G. GOPINATH (2006): “Defaultable debt, Interest Rates and the Current Account,” *Journal of International Economics*, 69, 64–83.
- ALFARO, L. (2014): “Sovereign Debt Restructuring: Evaluating the Impact of the Argentina Ruling,” *Harvard Business Law Review (Forthcoming)*.
- ALI, S. AND S. SILVEY (1966): “A general class of coefficients of divergence of one distribution from another,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 131–142.
- AMARI, S. AND H. NAGAOKA (2007): *Methods of information geometry*, vol. 191, American Mathematical Society.
- ARELLANO, C. (2008): “Default Risk and Income Fluctuations in Emerging Economies,” *American Economic Review*, 98, 690–712.
- ARTETA, C. AND G. HALE (2008): “Sovereign debt crises and credit to the private sector,” *Journal of International Economics*, 74, 53–69.

- ASHCRAFT, A. B., K. GOORIAH, AND A. KERMANI (2014): “Does Skin-in-the-Game Affect Security Performance? Evidence from the Conduit CMBS Market,” *Evidence from the Conduit CMBS Market (May 15, 2014)*.
- AUGUSTE, S., K. M. DOMINGUEZ, H. KAMIL, AND L. L. TESAR (2006): “Cross-border trading as a mechanism for implicit capital flight: ADRs and the Argentine crisis,” *Journal of Monetary Economics*, 53, 1259–1295.
- AY, N., J. JOST, H. VÂN LÊ, AND L. SCHWACHHÖFER (2014): “Information geometry and sufficient statistics,” *Probability Theory and Related Fields*, 1–38.
- BEGLEY, T. AND A. PURNANANDAM (2013): “Design of financial securities: Empirical evidence from private-label RMBS deals,” *Available at SSRN 2022380*.
- BERNANKE, B. S. AND K. N. KUTTNER (2005): “What explains the stock market’s reaction to Federal Reserve policy?” *The Journal of Finance*, 60, 1221–1257.
- BIAIS, B. AND C. CASAMATTA (1999): “Optimal leverage and aggregate investment,” *The Journal of Finance*, 54, 1291–1323.
- BIAIS, B., T. MARIOTTI, G. PLANTIN, AND J. ROCHET (2007): “Dynamic security design: Convergence to continuous time and asset pricing implications,” *The Review of Economic Studies*, 74, 345–390.
- BIERKENS, J. AND H. J. KAPPEN (2012): “Explicit solution of relative entropy weighted control,” *arXiv preprint arXiv:1205.6946*.
- BLACKWELL, D. (1953): “Equivalent comparisons of experiments,” *The annals of mathematical statistics*, 24, 265–272.
- BOCOLA, L. (2013): “The Pass-Through of Sovereign Risk,” *Manuscript, University of Pennsylvania*.

- BOLTON, P. AND O. JEANNE (2011): “Sovereign default risk and bank fragility in financially integrated economies,” *IMF Economic Review*, 59, 162–194.
- BORENSZTEIN, E. AND U. PANIZZA (2008): *The costs of sovereign default*, 8-238, International Monetary Fund.
- (2010): “Do sovereign defaults hurt exporters?” *Open Economies Review*, 21, 393–412.
- BOSE, A., D. PAL, AND D. SAPPINGTON (2011): “On the performance of linear contracts,” *Journal of Economics & Management Strategy*, 20, 159–193.
- BUBB, R. AND A. KAUFMAN (2014): “Securitization and moral hazard: Evidence from credit score cutoff rules,” *Journal of Monetary Economics*, 63, 1–18.
- BULOW, J. AND K. ROGOFF (1989a): “A Constant Recontracting Model of Sovereign Debt,” *The Journal of Political Economy*, 97.
- (1989b): “Sovereign Debt: Is to Forgive to Forget?” *The American Economic Review*, 43–50.
- CAMPBELL, J. Y., A. W. LO, AND C. MACKINLAY (1997): *The econometrics of financial markets*, Princeton University press.
- CAPLIN, A. AND M. DEAN (2014): “Revealed preference, rational inattention, and costly information acquisition,” Tech. rep., National Bureau of Economic Research.
- CAPLIN, M. AND A. DEAN (2013): “The Behavioral Implications of Rational Inattention with Shannon Entropy,” .
- CARROLL, G. (2013): “Robustness and linear contracts,” Tech. rep.
- CHASSANG, S. (2013): “Calibrated incentive contracts,” *Econometrica*, 81, 1935–1971.
- CHENTSOV, N. N. (1982): *Statistical decision rules and optimal inference*, 53, American Mathematical Soc.

- COLE, H. L. AND P. J. KEHOE (1998): “Models of sovereign debt: Partial versus general reputations,” *International Economic Review*, 55–70.
- COVER, T. M. AND J. A. THOMAS (2012): *Elements of information theory*, John Wiley & Sons.
- CSISZÁR, I. (1967): “Information-type measures of difference of probability distributions and indirect observations,” *Studia Sci. Math. Hungar.*, 2, 299–318.
- (1972): “A class of measures of informativity of observation channels,” *Periodica Mathematica Hungarica*, 2, 191–213.
- CVITANIĆ, J., X. WAN, AND J. ZHANG (2009): “Optimal compensation with hidden action and lump-sum payment in a continuous-time model,” *Applied Mathematics and Optimization*, 59, 99–146.
- DANG, T., G. GORTON, AND B. HOLMSTRÖM (2011): “Ignorance and the Optimality of Debt for Liquidity Provision,” Tech. rep., Working Paper, Yale University.
- DASEKING, C., A. R. GHOSH, T. D. LANE, AND A. H. THOMAS (2005): *Lessons from the Crisis in Argentina*, International monetary fund.
- DE OLIVEIRA, H., T. DENTI, M. MIHM, AND M. K. OZBEK (2013): “Rationally inattentive preferences,” *Available at SSRN 2274286*.
- DEMARZO, P. AND D. DUFFIE (1999): “A liquidity-based model of security design,” *Econometrica*, 67, 65–99.
- DEMARZO, P. M. (2005): “The pooling and tranching of securities: A model of informed intermediation,” *Review of Financial Studies*, 18, 1–35.
- DEMARZO, P. M. AND Y. SANNIKOV (2006): “Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model,” *The Journal of Finance*, 61, 2681–2724.

- DEMIROGLU, C. AND C. JAMES (2012): “How important is having skin in the game? originator-sponsor affiliation and losses on mortgage-backed securities,” *Review of Financial Studies*, 25, 3217–3258.
- DESAI, M. A., C. F. FOLEY, AND K. J. FORBES (2008): “Financial constraints and growth: Multinational and local firm responses to currency depreciations,” *Review of Financial Studies*, 21, 2857–2888.
- DI NUNNO, G., B. ØKSENDAL, AND F. PROSKE (2008): *Malliavin calculus for Lévy processes with applications to finance*, Springer.
- DYBVIG, P. H., H. K. FARNSWORTH, AND J. N. CARPENTER (2010): “Portfolio performance and agency,” *Review of Financial Studies*, 23, 1–23.
- EATON, J. AND M. GERSOVITZ (1981): “Debt with potential repudiation: Theoretical and empirical analysis,” *The Review of Economic Studies*, 289–309.
- EDMANS, A. AND Q. LIU (2011): “Inside debt,” *Review of Finance*, 15, 75–102.
- ELUL, R. (2011): “Securitization and mortgage default,” *Available at SSRN 1786317*.
- ENGLISH, W. B. (1996): “Understanding the costs of sovereign default: American state debts in the 1840’s,” *The American Economic Review*, 259–275.
- FENDER, I. AND J. MITCHELL (2009): “Incentives and tranche retention in securitisation: a screening model,” *CEPR Discussion Paper No. DP7483*.
- GALE, D. AND M. HELLWIG (1985): “Incentive-compatible debt contracts: The one-period problem,” *The Review of Economic Studies*, 52, 647–663.
- GEITHNER, T. (2011): *Macroeconomic Effects of Risk Retention Requirements*, DIANE Publishing.
- GELPERN, A. (2014a): “After Argentina,” Credit Slips Blog.

- (2014b): “Argentina: The RUFO Crazy,” Credit Slips Blog.
- GENNAIOLI, N., A. MARTIN, AND S. ROSSI (2013): “Banks, government bonds, and default: what do the data say?” .
- (2014): “Sovereign default, domestic banks, and financial institutions,” *The Journal of Finance*, 69, 819–866.
- GORODNICHENKO, Y. AND M. WEBER (2013): “Are sticky prices costly? Evidence from the stock market,” .
- GORTON, G. (2008): “The panic of 2007,” Tech. rep., National Bureau of Economic Research.
- GORTON, G. AND A. METRICK (2012): “Securitization,” *NBER Working Paper No. 18611*.
- GORTON, G. AND G. PENNACCHI (1990): “Financial intermediaries and liquidity creation,” *Journal of Finance*, 49–71.
- (1995): “Banks and loan sales marketing nonmarketable assets,” *Journal of Monetary Economics*, 35, 389–411.
- GULATI, G. M. AND K. N. KLEE (2001): “Sovereign piracy,” *The Business Lawyer*, 635–651.
- HANSEN, L. AND T. SARGENT (2008): *Robustness*, Princeton Univ Pr.
- HART, O. AND J. MOORE (1994): “A Theory of Debt Based on the Inalienability of Human Capital,” *The Quarterly Journal of Economics*, 109, 841–879.
- HARTMAN-GLASER, B., T. PISKORSKI, AND A. TCHISTYI (2012): “Optimal securitization with moral hazard,” *Journal of Financial Economics*, 104, 186–202.
- HE, H. (1990): “Convergence from discrete-to continuous-time contingent claims prices,” *Review of Financial Studies*, 3, 523–546.

- HEARON, J. Z. (1967): "Construction of epr generalized inverses by inversion of nonsingular matrices," *J. Res. Nat. Bur. Standards Sect. B*, 71, 57–60.
- HÉBERT, B. (2014): "Moral Hazard and the Optimality of Debt," *Available at SSRN 2185610*.
- HELLWIG, M. (2009): "A reconsideration of the Jensen-Meckling model of outside finance," *Journal of Financial Intermediation*, 18, 495–525.
- HELLWIG, M. F. AND K. M. SCHMIDT (2002): "Discrete-Time Approximations of the Holmström–Milgrom Brownian–Motion Model of Intertemporal Incentive Provision," *Econometrica*, 70, 2225–2264.
- HOLMSTRÖM, B. AND P. MILGROM (1987): "Aggregation and linearity in the provision of intertemporal incentives," *Econometrica*, 303–328.
- HORNBECK, J. (2013): "Argentina's Defaulted Sovereign Debt: Dealing with the 'Holdouts'," *Congressional Research Service*.
- HOROWITZ, J. L. (2001): "The bootstrap," *Handbook of econometrics*, 5, 3159–3228.
- INNES, R. (1990): "Limited liability and incentive contracting with ex-ante action choices," *Journal of Economic Theory*, 52, 45–67.
- JACOD, J. AND A. N. SHIRYAEV (2003): "Limit theorems for stochastic processes," .
- JENSEN, M. (1986): "Agency costs of free cash flow, corporate finance, and takeovers," *The American Economic Review*, 76, 323–329.
- JENSEN, M. AND W. MECKLING (1976): "Theory of the firm: Managerial behavior, agency costs and ownership structure," *Journal of Financial Economics*, 3, 305–360.
- JEWITT, I. (1988): "Justifying the first-order approach to principal-agent problems," *Econometrica: Journal of the Econometric Society*, 1177–1190.

- JIANG, W., A. A. NELSON, AND E. VYTLACIL (2013): “Securitization and Loan Performance: Ex Ante and Ex Post Relations in the Mortgage Market,” *Review of Financial Studies*, hht073.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KARATZAS, I. AND S. E. SHREVE (1991): “Brownian motion and stochastic calculus, vol. 113 of,” *Graduate texts in Mathematics*.
- KEYS, B., T. MUKHERJEE, A. SERU, AND V. VIG (2010): “Did securitization lead to lax screening? Evidence from subprime loans,” *The Quarterly Journal of Economics*, 125, 307–362.
- KRAINER, J. AND E. LADERMAN (2014): “Mortgage loan securitization and relative loan performance,” *Journal of Financial Services Research*, 45, 39–66.
- KUTTNER, K. N. (2001): “Monetary policy surprises and interest rates: Evidence from the Fed funds futures market,” *Journal of monetary economics*, 47, 523–544.
- LONGSTAFF, F. A., J. PAN, L. H. PEDERSEN, AND K. J. SINGLETON (2011a): “How Sovereign Is Sovereign Credit Risk?” *American Economic Journal: Macroeconomics*, 3, 75–103.
- (2011b): “How Sovereign is Sovieregn Credit Risk?” *American Economic Journal: Macroeconomics*, 3, 75–103.
- MALAMUD, S., H. RUI, AND A. WHINSTON (2013): “Optimal incentives and securitization of defaultable assets,” *Journal of Financial Economics*, 107, 111–135.
- MATEJKA, F. AND A. MCKAY (2011): “Rational inattention to discrete choices: A new foundation for the multinomial logit model,” *CERGE-EI Working Papers wp442, The Center for Economic Research and Graduate Education-Economic Institute, Prague*.
- MATĚJKA, F. AND A. MCKAY (2011): “Rational inattention to discrete choices: a new foundation for the multinomial logit model,” *CERGE-EI Working Paper Series*.

- MENDOZA, E. G. AND V. Z. YUE (2012): “A general equilibrium model of sovereign default and business cycles,” *Quarterly Journal of Economics*, 127.
- MIAN, A. AND A. SUFI (2009): “The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis,” *The Quarterly Journal of Economics*, 124, 1449–1496.
- MIRRLIIS, J. A. (1999): “The theory of moral hazard and unobservable behaviour: Part I,” *The Review of Economic Studies*, 66, 3–21.
- MONOYIOS, M. (2013): “Malliavin calculus method for asymptotic expansion of dual control problems,” *SIAM Journal on Financial Mathematics*, 4, 884–915.
- NACHMAN, D. C. AND T. H. NOE (1994): “Optimal design of securities under asymmetric information,” *Review of Financial Studies*, 7, 1–44.
- NADAULD, T. D. AND S. M. SHERLUND (2013): “The impact of securitization on the expansion of subprime credit,” *Journal of Financial Economics*, 107, 454–476.
- NADAULD, T. D. AND M. S. WEISBACH (2012): “Did securitization affect the cost of corporate debt?” *Journal of financial economics*, 105, 332–352.
- OLIVARES-CAMINAL, R. (2013): “The pari passu clause in sovereign debt instruments: developments in recent litigation,” *Panel 1–Policy panel*, 121.
- PAN, J. AND K. J. SINGLETON (2008): “Default and recovery implicit in the term structure of sovereign CDS spreads,” *The Journal of Finance*, 63, 2345–2384.
- PANIZZA, U., F. STURZENEGGER, AND J. ZETTELMAYER (2009): “The economics and law of sovereign debt and default,” *Journal of Economic Literature*, 651–698.
- PEREZ, D. J. (2014): “Sovereign Debt, Domestic Banks and the Provision Of Public Liquidity,” *Manuscript, Stanford University*.

- PRIVAULT, N. (2013): “Cumulant operators for Lie–Wiener–Itô–Poisson stochastic integrals,” *Journal of Theoretical Probability*, 1–30.
- PURNANANDAM, A. (2011): “Originate-to-distribute model and the subprime mortgage crisis,” *Review of Financial Studies*, 24, 1881–1915.
- RAJAN, U., A. SERU, AND V. VIG (2010): “The failure of models that predict failure: distance, incentives and defaults,” *Chicago GSB Research Paper*.
- RIGOBON, R. (2003): “Identification through heteroskedasticity,” *Review of Economics and Statistics*, 85, 777–792.
- RIGOBON, R. AND B. SACK (2004): “The impact of monetary policy on asset prices,” *Journal of Monetary Economics*, 51, 1553–1575.
- ROSE, A. K. (2006): “Size really doesn’t matter: In search of a national scale effect,” *Journal of the Japanese and International Economies*, 20, 482–507.
- SADZIK, T. AND E. STACHETTI (2013): “Agency Models with Frequent Actions,” Tech. rep., mimeo, NYU and UCLA.
- (2013): “Moral Hazard and Long-Run Incentives,” .
- SCHAETTLER, H. AND J. SUNG (1993): “The first-order approach to the continuous-time principal–agent problem with exponential utility,” *Journal of Economic Theory*, 61, 331–371.
- SCHUMACHER, J., C. TREBESCH, AND H. ENDERLEIN (2014): “Sovereign defaults in court,” *Available at SSRN 2189997*.
- SHAVELL, S. (1979): “Risk sharing and incentives in the principal and agent relationship,” *The Bell Journal of Economics*, 55–73.
- SHIRYAEV, A. N. (1996): “Probability, volume 95 of Graduate texts in mathematics,” .

- SIMS, C. (2003): “Implications of rational inattention,” *Journal of Monetary Economics*, 50, 665–690.
- SIMS, C. A. (2010): “Rational inattention and monetary economics,” *Handbook of Monetary Economics*, 3, 155–181.
- STANDARD AND POOR’S (2014a): “Argentine Nonsovereign Entities Ratings Unaffected Following The Sovereign’s Downgrade,” Tech. rep., Ratings Direct.
- (2014b): “Several Rating Actions Taken On 13 Argentine Companies; Off Watch Negative; Outlook Negative,” Tech. rep., Ratings Direct.
- STOCK, J. H. AND M. YOGO (2005): “Testing for weak instruments in linear IV regression,” *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 1.
- STRACCA, L. (2006): “Delegated portfolio management: A survey of the theoretical literature,” *Journal of Economic Surveys*, 20, 823–848.
- SYLVA, J. B. (1999): “Bowie Bonds Sold for Far More than a Song: The Securitization of Intellectual Property as a Super-Charged Vehicle for High Technology Financing,” *Santa Clara Computer & High Tech. LJ*, 15, 195.
- TOMZ, M. (2007): *Reputation and international cooperation: sovereign debt across three centuries*, Princeton University Press.
- TOWNSEND, R. (1979): “Optimal contracts and competitive markets with costly state verification,” *Journal of Economic Theory*, 21, 265–93.
- VAN NIEUWERBURGH, S. AND L. VELDKAMP (2010): “Information acquisition and underdiversification,” *The Review of Economic Studies*, 77, 779–805.
- VANASCO, V. M. (2013): “Incentives to Acquire Information vs. Liquidity in Markets for Asset-Backed Securities,” .


- WHITE, R. (2013): “The Pricing and Risk Management of Credit Default Swaps, with a Focus on the ISDA Model,” *OpenGamma Quantitative Research*.
- WOODFORD, M. (2012): “Inattentive valuation and reference-dependent choice,” *Unpublished Manuscript, Columbia University*.
- (2014): “An Optimizing Neuroeconomic Model of Discrete Choice,” Tech. rep., National Bureau of Economic Research.
- YANG, M. (2012): “Optimality of securitized debt with endogenous and flexible information acquisition,” *Available at SSRN 2103971*.
- YEYATI, E. L. AND U. PANIZZA (2011): “The elusive costs of sovereign defaults,” *Journal of Development Economics*, 94, 95–105.
- ZYMEK, R. (2012): “Sovereign default, international lending, and trade,” *IMF Economic Review*, 60, 365–394.

5 Appendix

5.1 Appendix: Moral Hazard and the Optimality of Debt


5.1.1 Additional Figures

Figure 5.1: Possible Security Designs



Notes: This figure illustrates several possible security designs: a debt security, an equity security, and the “live-or-die” security of Innes (1990). The x-axis, labeled $\beta_s v_i$, is the discounted value of the asset, and the y-axis, labeled $\beta_s s_i$, is the discounted value of the security. The level of debt, the cutoff point for the live-or-die, and the fraction of equity are chosen for illustrative purposes. The discount factor for the seller is $\beta_s = 0.5$. The outcome space v_i is a set of 401 evenly-spaced values ranging from zero to 8. The x-axis is truncated to make the chart clearer.

Figure 5.2: Second-Order Optimal Security Designs



Notes: This figure shows the second-order optimal security designs, for various values of the curvature parameter α . The x-axis, labeled $\beta_s v_i$, is the discounted value of the asset, and the y-axis, labeled $\beta_s s_i$, is the discounted value of the security. These securities are plotted with the same \bar{v} for each α (not an optimal \bar{v}). The value of κ used to generate this figure is roughly 0.17, which was chosen to ensure that the slopes of the contracts would be visually distinct (and not because it is economically reasonable). The outcome space v is a set of 401 evenly-spaced values ranging from zero to 8.

5.1.2 Timing Conventions

Table 5.1: Timing Conventions During Period One

<i>Principal-Agent Timing</i>	<i>Origination Timing</i>	<i>Shelf Registration Timing</i>
1. Security Designed	1. Actions Taken	1. Security Designed
2. Security Traded	2. Security Designed	2. Actions Taken
3. Actions Taken	3. Security Traded	3. Security Traded

The principal-agent timing convention is the simplest convention to analyze. In any sub-game perfect equilibrium, the seller takes actions that maximize the value of her retained tranche, because the price that she receives for the security has already been set. The buyer anticipates this, forming beliefs about the distribution of outcomes based on the design of the security. The buyer's beliefs affect the price that he is willing to pay for the security, and the seller internalizes this when designing the security. Multiple equilibria are possible if the seller's optimal actions for a particular retained tranche are not unique, or if there are multiple security designs that maximize the seller's utility. The moral hazard, in this timing convention, can occur either because the buyer is unaware of the seller's actions, or because he can observe those actions but is powerless to enforce any consequences based on them.

In the other two timing conventions, I use sequential equilibrium as the equilibrium concept and assume that the buyer does not observe the seller's actions. The buyer, because he is unaware of the actions the seller has taken, can accept or reject the security based only on its design and the price. In any equilibrium, the seller will always take actions that maximize the expected value of her retained tranche, where the expectations are taken with respect to the buyer's strategy and the state of nature. There is an equilibrium in which the buyer forms off-path beliefs by assuming that the seller acted in a manner consistent with the sub-game perfect equilibrium of the principal-agent timing, and therefore accepts or rejects the security design and price on that basis. In this

equilibrium, the seller will choose the same security design and price that would have been chosen in the principal-agent timing.

There are other sequential equilibria, which involve the same optimal security design, but different prices. These equilibria are all pareto-efficient, and equivalent to solving the principal-agent timing with Nash bargaining over prices, and different degrees of bargaining power for the seller and buyer. There are also sequential equilibria with other security designs, which are not pareto-efficient. I speculate that these equilibria are not proper equilibria.

5.1.3 Calibration

I will discuss the experimental approach first. This approach is consistent in spirit with the empirical literature on moral hazard in mortgage lending (Keys et al. (2010); Purnanandam (2011), others). In that literature, the quasi-experiment compares no securitization ($\eta_i = \beta_s v_i$) with securitization. If we assume securitization uses the optimal security design η^* , then θ can be approximated as

$$\theta^{-1} \approx E^{P(\beta_s, v)}[v_i] \cdot \frac{E^{P(\beta_s, v)}[v_i] - E^{P^*}[v_i]}{Cov^{P^*}(v_i, s_i^*)}.$$

This formula illustrates the difficulties of calibrating the model using the empirical work on moral hazard in mortgage lending. For the purposes of the model, what matters is the loss in expected value due to securitization, relative to the risk taken on by the buyers, ex-ante. The empirical literature estimates ex-post differences, and the magnitude of these differences varies substantially, depending on whether the data sample is from before or during the crash in home prices. Converting this an ex-ante difference would require assigning beliefs to the buyer and seller about the likelihood of a crash. Estimating the ex-ante covariance, which can be understood as a measure of the quantity of “skin in the game,” is even more fraught. For these reasons, I have not pursued this calibration strategy further.

The second calibration strategy, which is somewhat more promising, is to use the design of mortgage securities to infer θ . Essentially, by (crudely) estimating the other terms in the “put option value” equation (equation (1.7)), and assuming the model is correct, we can infer what the

security designers thought the moral hazard was. Rearranging that equation,

$$\underbrace{\frac{\beta_b \bar{v} - \beta_b E^{p^*}[s_i]}{\beta_b E^{p^*}[s_i]}}_{\text{Spread}} \underbrace{\frac{E^{p^*}[s_i]}{E^{p^*}[v_i]}}_{\text{Share}} \left(1 - \underbrace{\frac{E^{p(\beta_s v)}[v_i] - E^{p^*}[v_i]}{E^{p(\beta_s(v))}[v_i]}}_{\text{Moral Hazard}}\right) \kappa^{-1} = \theta.$$

The spread term should be thought of as reflecting the initial spread between the assets purchased by the buyer and the discount rate, under the assumption that the bonds will not default. Using a 90/10 weighting on the initial AAA and BBB 06-2 ABX coupons reported in Gorton (2008), I estimate this as 34 basis points per year. In a different setting (CLOs), the work of Nadauld and Weisbach (2012) estimates the cost of capital advantage due to securitization at 17 basis points per year. The “share” term is the ratio of the initial market value of the security to the initial market value of the assets. Begley and Purnanandam (2013) document that the value of the non-equity tranches was roughly 99% of the principal value in their sample of residential mortgage securitizations. Similarly, the moral hazard term is likely to be small. The estimates of Keys et al. (2010), whose interpretation is disputed by Bubb and Kaufman (2014), imply that pre-crisis, securitized mortgage loans defaulted at a 3% higher rate¹²¹ than loans held in portfolio. Assuming a 50% recovery rate, and using this as an estimate of the ex-ante expected difference in asset value, this suggests that the moral hazard term is roughly 1.5%, and therefore negligible in this calibration. Combining all of these estimates, I find θ of 2 is consistent with the empirical literature on securitization. This calibration assumed that the security design problem with the KL divergence was being solved. However, this formula also holds (approximately) under invariant divergences, conditional on the assumption that θ^{-1} and κ are small enough.

The value of $\theta = 2$ can be compared with the results of Figure 1.1. Under the assumptions used to generate that figure, which are described in its caption, I find that with $\theta = 2$ and $\kappa = 0.85$ (17 basis points per year times 5 years), debt would be achieve 99.96% of gains achieved by the optimal contract, relative to selling everything (and an even larger fraction of the gains relative to selling nothing). Under these parameters, the utility difference between the best debt security and

¹²¹After about one year, ~11% of securitized loans were in default, compared to ~8% of loans held in portfolio.

selling nothing would be roughly 0.73% of the total asset value. While that might seem like an economically small gain, for a single deal described in Gorton (2008), SAIL 2005-6, the private gains of securitization would be roughly \$16.4mm. In contrast, the utility difference between the best equity security and selling nothing is about 0.56% of the total asset value. The private cost of using the optimal equity contract, instead of the optimal debt contract, would be roughly \$4mm for this particular securitization deal.

The numbers discussed in this calculation depend on the assumptions used in Figure 1.1, some of which are ad hoc. Nevertheless, they illustrate the general point that it is simultaneously possible for debt to be approximately optimal, and for the private gains of securitization to be large.

5.1.4 Contracting with a Rationally Inattentive Seller

In this section, I consider the problem of contracting with a rationally inattentive seller. I will show that this model is very close to the parametric security design model discussed in section §1.5, and that the lower bound result in proposition 1.3 applies.

One possible economic motivation for this type of problem comes from mortgage origination. Suppose that there is a set of A alternative groups of borrowers that the seller (a mortgage originator) could lend to. These groups could be borrowers from different zip codes, with different credit scores, etc., and a single borrower might belong to multiple groups. The seller can choose only one of these groups to lend to, and create the asset by making loans to this group. I assume that the seller can freely observe to which groups a potential borrower belongs, but is uncertain about which group she should lend to. I also assume that, under the seller's and buyer's common prior, lending to each alternative group will result in an asset with the same expected value and the same variance of asset value.

The buyer observes which group the seller chose to lend to. The buyer can also observe (in the second period) the ex-post returns of each group, not just the group the seller picked. The set of possible states of the world is X , where each state $x \in X$ corresponds to a set of asset values for each of the alternative groups, if that group had been chosen by the seller. There are $N = |A| \cdot |X|$

possible combinations of states and choices by the seller. Let i be an index of these state-choice pairs, $i \in \{0, \dots, N-1\}$, and let v_i denote the value of the asset the seller actually created in state-choice i . I will assume that $v_0 = 0$.

The value of the security that the seller offers to the buyer can depend on i , the state-choice pair. In this sense, the security can be “benchmarked,” and the payoff can depend not only on the ex-post value of the asset the seller created, but also on the ex-post value of the other alternatives she might have chosen. However, limited liability still applies, based on the value of the asset the seller actually created. I will denote the value of the security in state-choice i as s_i , and require $s_i \in [0, v_i]$.

The moral hazard in this problem comes from the seller’s ability to acquire information about which states $x \in X$ are mostly likely, before choosing amongst the A alternatives. The quantity and nature of the information the seller acquires is not contractible, and this is what creates the moral hazard problem. The seller could choose an information structure to maximize the expected value of the assets she creates, but will instead choose an information structure to maximize the value of the retained tranche η . I will model the choice and cost of different information structures using the rational inattention framework of Sims (2003) and the results about discrete choice and rational inattention found in Matějka and McKay (2011).

As in previous sections, I will assume that the seller has a lower discount factor than the buyer, creating gains from trade. For expository purposes, I will use the “principal-agent” timing convention, but the optimal security design does not depend on the timing.¹²² The timing of the model is as follows. During the first period:

1. The seller designs the security s_i , subject to limited liability.
2. The seller makes a take-it-or-leave-it offer to the buyer, at price K .
3. The buyer accepts or rejects the offer.

¹²²In the “origination” and “shelf registration” timing conventions, one can show that equilibria with pooling across signal realizations are ruled out in the optimal signal structure. In effect, endogenous information acquisition leads to pure moral hazard problems, in situations where exogenous signal structures would lead to adverse selection.

4. The seller chooses a signal structure, unobservable to the buyer.
5. The seller receives a signal about the likelihood of each state in X .
6. The seller chooses an alternative $a \in A$.

The seller's choice of information structure and action is similar to a version of the parametric models discussed earlier. Let $g(x)$ denote the common prior over the possible states $x \in X$. Following the standard simplification result used in rational inattention problems, I will think of the seller as choosing the conditional probability distribution of alternatives $a \in A$, for each state x . This model is parametric, because the seller cannot create any probability distribution over the N state-choice combinations. Instead, she is constrained to create only those probability distributions over the N state-choices that have the marginal distribution over states $x \in X$ equal to $g(x)$. Economically, the mortgage originator does not control the amount that each group of mortgage borrowers will repay (the state). The best the mortgage originator can do is learn about which states are likely, and choose which group of borrowers to lend to accordingly.

The cost of the signal structure is described by mutual information, which can be rewritten as a KL divergence. The seller's moral hazard problem is

$$\phi_{RI}(\eta; \theta^{-1}) = \max_{p \in M_{RI}} \left\{ \sum_i p^i \eta_i - \theta D_{KL}(p || q(p)) \right\},$$

where M_{RI} denotes the set of probability distributions over state-choices with the marginal distribution over states equal to $g(x)$, and $q(p)$ is the joint distribution between states and choices, when choices are independent of states and the marginal distribution of choices in q is the same as the marginal distribution in p . The constant of proportionality θ is now interpreted as the cost per unit of information that the seller acquires. I will continue to assume that the solution is interior and unique. This assumption simplifies the proof and discussion, but can be relaxed (Matějka and McKay (2011) provide more primitive conditions under which the assumption would hold).

The key difference between this moral hazard problem and the parametric problem discussed previously is that the q distribution is endogenous, not fixed. This endogeneity is a property of the

mutual information cost function. If the seller were to choose p so that her action were independent of the state, there would be no information cost, regardless of the marginal distribution of actions. Economically, if the mortgage originator decided to randomly choose a group of borrowers to give loans to, independent of economic fundamentals, this would require no costly information gathering, regardless of how the mortgage originator decided to randomize.

The security design problem is almost identical to the parametric moral hazard model. For an arbitrary security s ,

$$U_{RI}(s; \theta^{-1}, \kappa) = \beta_s(1 + \kappa) \sum_{i>0} p^i(\eta(s); \theta^{-1}) s_i + \phi_{RI}(\eta(s); \theta^{-1}),$$

where $p(\eta; \theta^{-1})$ is the probability distribution over state-choices that the seller will choose, given retained tranche η and cost of information θ . Despite the aforementioned complication that q is endogenous, the lower bound result of proposition 1.5 applies.

Proposition 5.1. *In the rational inattention model, the difference in utilities achieved by an arbitrary security s and the sell-nothing security, for sufficiently small θ^{-1} and κ , is bounded below:*

$$U_{RI}(s; \theta^{-1}, \kappa) - U_{RI}(0; \theta^{-1}, \kappa) \geq \kappa E^{\bar{p}}[\beta_s s_i] - \frac{1}{2} \theta^{-1} V^{\bar{p}}[\beta_s s_i] + O(\theta^{-2} + \kappa \theta^{-1}),$$

where \bar{p} is the endogenous probability distribution that the seller would choose, if they could not gather information and the security was s .

Proof. See appendix, section 5.1.14. □

The statement of proposition 5.1 discusses the endogenous probability distribution \bar{p} , which is the probability distribution the seller would choose if they could only choose the unconditional distribution of actions, and the actions were independent of the state. The key intuition is that different security designs can change the seller's incentives to pick one alternative over another, even if the seller could not gather any information at all. This could be true because the security treats different alternatives differently, or because the security's valuation depends on third or

higher moments of the value distribution (only the mean and variance of each alternative under the prior are assumed to be identical). Debt maximizes the lower bound in proposition 5.1 only when small, symmetric perturbations to the debt security do not alter the unconditional distribution of actions.¹²³ This would hold true if the prior distribution of values were identical for each alternative.

In general, rational inattention problems will not result in debt as the optimal contract, because the seller does not control the state $x \in X$. If there are some states that offer higher payoffs than others, regardless of the alternative chosen, it is not efficient for the seller to receive higher cash flows in those “good” states. Additionally, the security design influences the unconditional probability that a particular alternative will be chosen. A debt contract, whose payoff does not vary by alternative chosen, only by outcome, may result in a sub-optimal choice of unconditional probabilities for different actions.

The lower bound is tight for symmetric contracts under two conditions.¹²⁴ The first is the exchangeable prior discussed by Matějka and McKay (2011). When the prior is exchangeable, there is no reason for the security payoff to vary by alternative chosen, because each alternative is ex-ante equivalent. The second condition is that the states are symmetric, and there is no such thing as a “good state” or “bad state,” only good and bad alternatives conditional on the state. One justification for this second condition is complete markets (see the discrete-outcome, complete asset market described in He (1990)). Under these two conditions, the lower bound is tight, and a debt security is approximately optimal.¹²⁵

Contracting problems with rationally inattentive agents are challenging to solve exactly, both analytically and computationally. The lower bound result of proposition 5.1 shows that, when the seller’s ability to gather information is weak, debt contracts are a tractable, detail-free way of guaranteeing a minimum utility level.

¹²³By symmetric perturbations, I mean those perturbations that, like the debt contract, depend only on the value of the outcome, and not the alternative chosen.

¹²⁴These conditions are sufficient, not necessary.

¹²⁵In fact, under these two conditions, a debt security is optimal, not just approximately optimal.

5.1.5 Proof of proposition 1.1

To prove that debt securities are optimal, I first re-characterize the security design first-order condition.

Lemma 5.1. *In the non-parametric model, with the cost function proportional to the Kullback-Leibler divergence, the optimal security design satisfies the following condition for all values of j :*

$$\beta_b s_j(\eta^*) = \beta_b \sum_{i>0} p^i(\eta^*) s_i(\eta^*) + \kappa \theta \left(1 - \frac{\lambda^j}{p^j(\eta^*)} + \frac{\omega^j}{p^j(\eta^*)}\right).$$

Proof. See appendix, section 5.1.6. □

To characterize the optimal security, I analyze the equation of lemma 5.1 and the complementary slackness conditions for the multipliers, as in Yang (2012). The complementary slackness conditions on the multipliers require that if $\lambda^j > 0$, $s_j = v_j$, and that if $\omega^j > 0$, $s_j = 0$. Define the endogenous positive constant

$$\bar{v} = \sum_{i>0} p^i(\eta^*) s_i(\eta^*) + \frac{\kappa}{\beta_b} \theta.$$

Consider the case where $\lambda^j > 0$. In this case, we must have $\omega^j = 0$ and $s_j = v_j$. Using lemma 5.1, in this case $v_j = \bar{v} - \frac{\kappa}{\beta_b} \theta \frac{\lambda^j}{p^j} < \bar{v}$. Next, consider the case where $\omega^j > 0$. In this case, we would have $\lambda^j = 0$, $s_j = 0$, and therefore $0 = \bar{v} + \frac{\kappa}{\beta_b} \theta \frac{\omega^j}{p^j}$, a contradiction. Finally, consider the case where $\lambda^j = \omega^j = 0$. In this case, $s_j = \bar{v}$. Combining these three cases, we see that if $v_j < \bar{v}$, we must have $\lambda^j > 0$ and $s_j = v_j$, and if $v_j \geq \bar{v}$, we must have $s_j = \bar{v}$. Therefore, s_j is a debt contract.

Assume that $v_N > \sum_i q^i v_i + \frac{\kappa}{\beta_b} \theta$, and suppose that $\bar{v} \geq v_N$. The contract would be a sell-everything contract, and the seller would choose the minimum cost distribution, $p^i(\eta^*) = q^i$. Therefore, $\bar{v} = \sum_i q^i v_i + \frac{\beta_b - \beta_s}{\beta_b \beta_s} \theta$, a contradiction.

The argument above shows that any security satisfying the first-order condition of the Lagrangian is a debt contract. The existence of an optimal contract is guaranteed by the existence of

a maximum of $U(\eta)$ on the closed and bounded set of \mathbb{R}^N corresponding to limited liability securities, due to the continuity of $U(\eta)$. Because a maximum exists and $U(\eta)$ is differentiable, the first-order condition for the Lagrangian is necessary for optimality. Therefore the optimal security must be a debt security.

5.1.6 Proof of lemma 5.1

The security design equation, specialized to the Fisher metric Hessian, is

$$\beta_b s_j(\eta^*) = \sum_{i>0} \theta \kappa [p^i(\eta^*) - \lambda^i + \omega^i] g_{ij}(p(\eta^*)),$$

where

$$g_{ij}(p(\eta^*)) = \left(\frac{\delta_{ij}}{p^i(\eta^*)} + \frac{1}{p^0(\eta^*)} \right).$$

Taking expected values,

$$\begin{aligned} \beta_b \sum_{j>0} p^j(\eta^*) s_j(\eta^*) &= \theta \kappa \sum_i [p^i(\eta^*) - \lambda^i + \omega^i] + \\ &\quad \theta \kappa \sum_i [p^i(\eta^*) - \lambda^i + \omega^i] \frac{1 - p^0(\eta^*)}{p^0(\eta^*)}. \end{aligned}$$

Simplifying,

$$\beta_b \sum_{j>0} p^j(\eta^*) s_j(\eta^*) = \frac{\theta \kappa}{p^0(\eta^*)} \sum_i [p^i(\eta^*) - \lambda^i + \omega^i].$$

Plugging this into the security design equation,

$$\beta_b s_j(\eta^*) = \beta_b \sum_{i>0} p^i(\eta^*) s_i(\eta^*) + \theta \kappa \sum_{i>0} [p^i(\eta^*) - \lambda^i + \omega^i] \frac{\delta_{ij}}{p^i(\eta^*)}.$$

which is equivalent to the desired result.

5.1.7 Proof of proposition 1.2

For any f -divergence, the Hessian is

$$\frac{\partial^2 \psi(p)}{\partial p^i \partial p^j} \Big|_{p=p(\eta^*)} = \theta \left[\frac{\delta_{ij}}{q^i} f''\left(\frac{p^i(\eta^*)}{q^i}\right) + \frac{1}{q^0} f''\left(\frac{p^0(\eta^*)}{q^0}\right) \right].$$

To prove proposition 1.2, I suppose that there is some f -divergence such that, for every sample space Ω and zero-cost distribution q , the optimal security design is debt. I will show that this f -divergence must be the KL divergence.

Take some Ω and q such that the optimal security design is a debt, with at least two distinct indices j and j' such that $s_j = s_{j'}$ and $0 < s_j < \min(v_j, v_{j'})$, and such that the solution to the moral hazard problem under the optimal contract is interior. The indices j and j' correspond to the “flat” part of the a debt contract. I construct an Ω and q with this property below, showing that such examples exist. Because the solution to the moral hazard problem is interior, from the security design equation,

$$s_j(\eta^*) = \sum_{i>0} \theta(\beta_s^{-1} - \beta_b^{-1}) [p^i(\eta^*) - \lambda^i + \omega^i] \left[\frac{\delta_{ij}}{q^i} f''\left(\frac{p^i(\eta^*)}{q^i}\right) + \frac{1}{q^0} f''\left(\frac{p^0(\eta^*)}{q^0}\right) \right].$$

For the indices j and j' associated with the flat part of the debt contract, neither multiplier binds, and therefore

$$\frac{p^j(\eta^*)}{q^j} f''\left(\frac{p^j(\eta^*)}{q^j}\right) = \frac{p^{j'}(\eta^*)}{q^{j'}} f''\left(\frac{p^{j'}(\eta^*)}{q^{j'}}\right).$$

If the optimal security is a debt for some Ω and q , this property must hold for all j and j' associated with the flat part of that debt contract. Below, I state a claim. I will first prove the rest of the proposition, assuming this claim, and then prove the claim.

Claim 5.1. For all pairs $u_1 \in (0, 1]$ and $u_2 \in [1, \infty)$, there exist Ω and q such that, under the optimal debt contract, there are indices j and j' , associated with the flat part of the debt contract, such that

$$\frac{p^j(\eta^*)}{q^j} = u_1$$

and

$$\frac{p^{j'}(\eta^*)}{q^{j'}} = u_2.$$

Assume the above claim is true. If the optimal security is a debt, then the optimal debt contract (the best contract in the class of debt contracts) is the optimal security (in the space of all securities). Because debt is optimal for all Ω and q , including the ones constructed in the above claim, it follows that for any u_1 and u_2 , we must have

$$u_1 f''(u_1) = u_2 f''(u_2).$$

When $u_1 = 1$, $u f''(u) = 1$ by the normalization assumption, and therefore for all $u \in (0, \infty)$,

$$u f''(u) = 1.$$

Solving the differential equation, using the normalization assumptions that $f'(1) = 0$ and $f(1) = 0$,

$$f(u) = u \ln u - u + 1,$$

proving the proposition. To complete the proof, I prove the above claim.

5.1.8 Proof of claim 5.1

First, note the first-order condition for the moral hazard problem, assuming an interior solution, requires that

$$f'\left(\frac{p^i(\eta^*)}{q^i}\right) - f'\left(\frac{p^0(\eta^*)}{q^0}\right) = \theta^{-1} \eta_i = \theta^{-1} \beta_s (v_i - s_i).$$

For the flat region for some j ,

$$f'\left(\frac{p^j(\eta^*)}{q^j}\right) - f'\left(\frac{p^0(\eta^*)}{q^0}\right) = \theta^{-1}\beta_s(v_j - \bar{v}).$$

Let $\Omega = \{0, v_1, v_2\}$. I will construct, under the optimal debt security,

$$\frac{p^1(\eta^*)}{q^1} = u_1,$$

$$\frac{p^2(\eta^*)}{q^2} = u_2,$$

for given $u_1 \in (0, 1]$ and $u_2 \in [1, \infty)$. Assume for now that the optimal debt contract has $\bar{v} < v_1$ as its maximum value. By the first-order conditions, we have (using the normalization condition)

$$f'(u_1) - f'(u_0) = \theta^{-1}\beta_s(v_1 - \bar{v}),$$

where $u_0 = \frac{p^0(\eta^*)}{q^0}$. It follows that

$$f'(u_2) - f'(u_1) = \theta^{-1}\beta_s(v_2 - v_1).$$

Because $f(u)$ is convex, this equation can be solved to pin down v_2 given u_2 and u_1 , and v_1 , noting that $v_2 \geq v_1$ by the convexity and normalization of f , and $u_2 \geq u_1$. Next, I will use the adding up constraints to choose a q . We must have

$$q^0 + q^1 + q^2 = 1$$

and

$$q^0 u_0 + q^1 u_1 + q^2 u_2 = 1.$$

Putting these together,

$$u_0(q^1, q^2) = \frac{1 - q^1 u_1 - q^2 u_2}{1 - q^1 - q^2}.$$

The solution will be interior if $u_0(q^1, q^2) > 0$. Finally, I define the optimal debt security. Denote the class of retained tranches associated with debt securities as $\eta(\bar{v})$. The security design FOC with respect to \bar{v} is

$$\left(1 - \frac{\beta_b}{\beta_s}\right) \sum_{i>0} p^i(\eta^*) \frac{\partial \eta_i(\bar{v})}{\partial \bar{v}} + \beta_b \sum_{i,j>0} \frac{\partial^2 \phi(\eta)}{\partial \eta_i \partial \eta_j} \Big|_{\eta=\eta^*} s_j(\eta^*) \frac{\partial \eta_i(\bar{v})}{\partial \bar{v}} = 0.$$

Define $\bar{v}_{full}(q^1, q^2)$ as the solution to this equation, given q^1 and q^2 . The Hessian of ϕ is the inverse of the Hessian of ψ . Using the Sherman-Morrison formula,

$$\frac{\partial^2 \phi(\eta)}{\partial \eta_i \partial \eta_j} \Big|_{\eta=\eta^*} = \theta^{-1} \frac{\delta_{ij} q^i}{f''(u_i)} - \theta^{-1} \frac{\frac{q^i q^j}{f''(u_i) f''(u_j)}}{\sum_i \frac{q^i}{f''(u_i)}}.$$

Assuming that $v_1 > \bar{v}$, this simplifies to

$$(\beta_b^{-1} - \beta_s^{-1})(1 - u_0 q^0) + \beta_b \bar{v} \sum_{i,j>0} \frac{\partial^2 \phi(\eta)}{\partial \eta_i \partial \eta_j} \Big|_{\eta=\eta^*} = 0.$$

Define $\bar{v}(q^1, q^2)$ as the solution to this equation, which differs from $\bar{v}_{full}(q^1, q^2)$ due to the assumption that $\bar{v} < v_1$ and associated simplifications. Summing,

$$\sum_{i,j>0} \frac{\partial^2 \phi(\eta)}{\partial \eta_i \partial \eta_j} \Big|_{\eta=\eta^*} = \theta^{-1} \sum_{i>0} \frac{q^i}{f''(u_i)} - \theta^{-1} \frac{\sum_{i,j>0} \frac{q^i q^j}{f''(u_i) f''(u_j)}}{\sum_i \frac{q^i}{f''(u_i)}}.$$

Because ϕ is convex, we can think of this as

$$\sum_{i,j>0} \frac{\partial^2 \phi(\eta)}{\partial \eta_i \partial \eta_j} \Big|_{\eta=\eta^*} = h(q^1, q^2),$$

where $h(q^1, q^2)$ is a positive-valued function. Therefore,

$$\bar{v}(q^1, q^2) = (\beta_s^{-1} - \beta_b^{-1})\theta \frac{q^1 u_1 + q^2 u_2}{h(q^1, q^2)}.$$

We can define v_1 from

$$v_1 = \bar{v}(q^1, q^2) + \theta \beta_s^{-1} (f'(u_1) - f'(u_0(q^1, q^2))),$$

which is greater than \bar{v} by the convexity and normalization of f , as long as $u_0 < u_1$, which can be maintained by choice of q_1 and q_2 . It remains to be shown that $\bar{v}(q^1, q^2)$ is the globally optimal contract, considering also contracts with $\bar{v} > v_1$. Note that in the limit as $q^1 \rightarrow 0$,

$$\lim_{q^1 \rightarrow 0^+} \bar{v}_{full}(q^1, q^2) = \lim_{q^1 \rightarrow 0^+} \bar{v}(q^1, q^2),$$

because whether $v_1 > \bar{v}$ or not does not alter the utility from the security design, since q^1 is small. Moreover, in this limit, $u_0(q^1, q^2)$ can be made arbitrarily small through a choice of q^2 , preserving the requirement that $u_0(q^1, q^2) \in (0, u_1)$. It follows that there exist q^1 and q^2 such that $u_0(q^1, q^2) < u_1$, $0 < \bar{v}(q^1, q^2) < v_1 < v_2$, and the desired result

$$\frac{p^1(\eta^*)}{q^1} = u_1,$$

$$\frac{p^2(\eta^*)}{q^2} = u_2,$$

holds.

5.1.9 Proof of proposition 1.3 and proposition 1.4

This section proves both proposition 1.4 and proposition 1.3, which is a special case. In this proof, I am using the summation convention. For example, $p^i s_i$ is a summation, $\sum_{i>0} p^i s_i$. We define the utility generated by a particular contract as

$$U(s; \theta^{-1}) = \beta_b p^i(\eta(s)) s_i + \phi(\eta(s)).$$

The first-order condition for the moral hazard problem is

$$\eta_i - \partial_i \psi(p(\eta)) = 0.$$

In the neighborhood of $\theta \rightarrow \infty$,

$$\lim_{\theta \rightarrow \infty} p(\eta) = q,$$

because otherwise the seller would receive unbounded negative utility. The solution to the moral hazard problem is guaranteed to be interior in this neighborhood, due to the assumption that q has full support. Expanding the function $\partial_i \psi(p(\eta))$ around $p = q$,

$$\eta_i - \partial^i \psi(q) - (p^j - q^j) \partial_i \partial_j \psi(q) - \frac{1}{2} (p^j - q^j) (p^k - q^k) \partial_i \partial_j \partial_k \psi(p^*) = 0,$$

where $p^* = q + a^*(p - q)$ for some $a^* \in (0, 1)$. Because $\psi(p)$ is invariant, we can simplify this to

$$p^j - q^j = \theta^{-1} g^{jj}(q) \eta_i - \frac{1}{2} g^{jj}(q) (p^l - q^l) (p^k - q^k) h_{ikl}(p^*),$$

where $h_{ikl}(p^*) = \theta^{-1} c \partial_i \partial_j \partial_k \psi(p^*)$. It follows that

$$p^j - q^j = O(\theta^{-1}).$$

Returning to the first-order condition, and expanding it up to order θ^{-2} ,

$$\theta^{-1} \eta_i g^{il}(q) - \frac{1}{2} (p^j - q^j) (p^k - q^k) h_{ijk}(q) g^{il}(q) + O(\theta^{-3}) = (p^l - q^l).$$

Plugging this equation into itself,

$$\theta^{-1}\eta_i g^{il}(q) - \frac{1}{2}\theta^{-2}\eta_m \eta_n g^{jm}(q)g^{kn}(q)h_{ijk}(q)g^{il}(q) + O(\theta^{-3}) = (p^l - q^l).$$

For ease of notation, define

$$h^{lmn}(q) = g^{jm}(q)g^{kn}(q)h_{ijk}(q)g^{il}(q).$$

It also follows that to second order, the cost function can be approximated as

$$\psi(p) = \frac{\theta}{2}(p^i - q^i)(p^j - q^j)g_{ij}(q) + \frac{\theta}{6}(p^i - q^i)(p^j - q^j)(p^k - q^k)h_{ijk}(q) + O(\theta^{-3}),$$

Using the first-order condition,

$$\begin{aligned} \psi(p) &= \frac{1}{2}(p^i - q^i)\eta_i - \\ &\quad \frac{1}{4}\theta^{-1}(p^i - q^i)\eta_m \eta_n g^{jm}(q)g^{kn}(q)h_{ijk}(q) + \\ &\quad \frac{\theta}{6}(p^i - q^i)(p^j - q^j)(p^k - q^k)h_{ijk}(q) + O(\theta^{-3}), \end{aligned}$$

and

$$\begin{aligned} \psi(p) &= \frac{1}{2}\theta^{-1}\eta_j \eta_i g^{ij}(q) - \\ &\quad \frac{1}{4}\theta^{-2}\eta_m \eta_n \eta_i h^{ilm}(q) - \\ &\quad \frac{1}{4}\theta^{-2}\eta_m \eta_n \eta_i h^{ilm}(q) + \\ &\quad \frac{\theta^{-2}}{6}\eta_m \eta_n \eta_i h^{ilm}(q) + O(\theta^{-3}). \end{aligned}$$

Simplifying,

$$\psi(p) = \frac{1}{2}\theta^{-1}\eta_j\eta_i g^{ij}(q) - \frac{1}{3}\theta^{-2}\eta_m\eta_n\eta_i h^{ilm}(q) + O(\theta^{-3}).$$

The utility given by an arbitrary security is

$$U(s; \theta^{-1}, \kappa) = \beta_s(1 + \kappa)p^i(\eta(s))v_i - \kappa p^i(\eta(s))\eta_i(s) - \psi(p(\eta(s))),$$

which under these approximations is

$$\begin{aligned} U(s; \theta^{-1}, \kappa) &= \beta_s(1 + \kappa)(p^i - q^i)v_i + \beta_s(1 + \kappa)q^i v_i - \\ &\quad \kappa(p^i - q^i)\eta_i - \kappa q^i \eta_i - \\ &\quad \frac{\theta}{2}(p^i - q^i)(p^j - q^j)g_{ij}(q) - \frac{\theta}{6}(p^i - q^i)(p^j - q^j)(p^k - q^k)h_{ijk}(q) + O(\theta^{-3} + \kappa\theta^{-2}). \end{aligned}$$

and the expression can be rewritten as

$$\begin{aligned} U(s; \theta^{-1}, \kappa) &= \beta_s(1 + \kappa)\theta^{-1}\eta_i g^{il}(q)v_l - \frac{\theta^{-2}}{2}\beta_s\eta_j\eta_k v_i h^{ijk}(q) + \\ &\quad + \beta_s(1 + \kappa)q^i v_i - \\ &\quad \kappa\theta^{-1}\eta_i g^{il}(q)\eta_l - \kappa q^i \eta_i - \\ &\quad \frac{\theta^{-1}}{2}\eta_i\eta_j g^{ij}(q) + \frac{\theta^{-2}}{3}\eta_i\eta_j\eta_k h^{ijk}(q) + O(\theta^{-3} + \kappa\theta^{-2}). \end{aligned} \quad (5.1)$$

For the zero security,

$$U(0; \theta^{-1}, \kappa) = \frac{1}{2}\beta_s^2\theta^{-1}v_i g^{il}(q)v_l - \frac{1}{6}\theta^{-2}\beta_s^3 v_j v_k v_i h^{ijk}(q) + \beta_s q^i v_i + O(\theta^{-3} + \kappa\theta^{-2}).$$

Taking the difference,

$$\begin{aligned}
U(s; \theta^{-1}, \kappa) - U(0; \theta^{-1}, \kappa) &= \beta_s(1 + \kappa)\theta^{-1}\eta_i g^{il}(q)v_l - \frac{\theta^{-2}}{2}\beta_s\eta_j\eta_k v_i h^{ijk}(q) + \\
&+ \beta_s \kappa q^i s_i - \frac{1}{2}\beta_s^2 \theta^{-1} v_i g^{il}(q)v_l + \frac{1}{6}\theta^{-2}\beta_s^3 v_j v_k v_i h^{ijk}(q) - \\
&\kappa\theta^{-1}\eta_i g^{il}(q)\eta_l - \\
&\frac{\theta^{-1}}{2}\eta_i \eta_j g^{ij}(q) + \frac{\theta^{-2}}{3}\eta_i \eta_j \eta_k h^{ijk}(q) + O(\theta^{-3} + \kappa\theta^{-2}).
\end{aligned}$$

Substituting out (most) of the η terms for s and v terms,

$$\begin{aligned}
U(s; \theta^{-1}, \kappa) - U(0; \theta^{-1}, \kappa) &= -\frac{1}{2}\beta_s^2 \theta^{-1} s_i g^{il}(q)s_l + \beta_s \kappa \theta^{-1} n_i g^{il}(q)s_l + \kappa \beta_s q^i s_i \\
&+ \frac{1}{6}\theta^{-2}\beta_s^2 \eta_j s_k v_i h^{ijk}(q) + \frac{1}{3}\theta^{-2}\beta_s^3 s_j s_k s_i h^{ijk}(q) + O(\theta^{-3}).
\end{aligned}$$

This expression is equivalent to the statement of the theorem, except that it remains to show that

$$s_i g^{il}(q)s_l = \text{Var}^q[s_i],$$

$$v_i g^{il}(q)s_l = \text{Cov}^q[v_i, s_i],$$

and

$$x_i y_j z_k h^{ijk}(q) = \left(\frac{3 + \alpha}{2}\right) \mathbf{K}_3^q(x_j, y_k, z_i).$$

The first two follow from the definition of the non-parametric Fisher information metric (see also theorem 2.7 in Amari and Nagaoka (2007)). Applying lemma 1.1,

$$x_i y_j z_k h^{ijk}(q) = \left(\frac{3 + \alpha}{2}\right) x_i g^{il}(q) y_j g^{jm}(q) z_k g^{kn}(q) \partial_l g_{mn}(p) \Big|_{p=q}.$$

This can also be written in terms of the derivative of the inverse Fisher metric,

$$x_i y_j z_k h^{ijk}(q) = -\left(\frac{3+\alpha}{2}\right) x_i g^{il}(q) \partial_l (y_j z_k g^{jk}(p))|_{p=q}.$$

Because of the non-parametric nature of the problem,

$$y_j z_k g^{jk}(p) = E^p[y \cdot z] - E^p[y] E^p[z].$$

Differentiating,

$$\partial_l (y_j z_k g^{jk}(p))|_{p=q} = y_l \cdot z_l - y_l E^q[z] - z_l E^q[y].$$

Therefore,

$$\begin{aligned} -\frac{2}{3+\alpha} x_i y_j z_k h^{ijk}(q) &= x_i g^{il}(q) \partial_l (y_j z_k g^{jk}(p))|_{p=q} \\ &= E^q[x \cdot y \cdot z] - E^q[x \cdot y] E^q[z] - E^q[x \cdot z] E^q[y] - \\ &\quad E^q[y \cdot z] E^q[z] + 2E^q[x] E^q[y] E^q[z], \end{aligned}$$

which is the definition of the third cross-cumulant.

5.1.10 Proof of corollary 1.1

As argued in the proof of proposition 1.3, in the limit as $\theta \rightarrow \infty$, the solution to the moral hazard problem is interior. It follows that the security design equation equation (1.6), for an invariant divergence, can be written (again, using the summation notation) as

$$\beta_s s_j^*(\theta^{-1}, \bar{\kappa} \theta^{-1}) = \frac{\bar{\kappa}}{1 + \bar{\kappa} \theta^{-1}} [p^i(\theta^{-1}) - \lambda^i(\theta^{-1}) + \omega^i(\theta^{-1})] \partial_i \partial_j D_f(p(\theta) || q).$$

Here, I am writing the endogenous probability distributions and multipliers as functions of θ^{-1} , to simplify by discussion of limits. Defining $\lim_{\theta^{-1} \rightarrow 0^+} \lambda^i(\theta^{-1}) = \lambda_{opt}^i$ and $\lim_{\theta^{-1} \rightarrow 0^+} \omega^i(\theta^{-1}) = \omega_{opt}^i$,

$$\lim_{\theta^{-1} \rightarrow 0^+} s_j^*(\theta^{-1}, \bar{\kappa}\theta^{-1}) = \frac{\bar{\kappa}}{\beta_s} [q^i - \lambda_{opt}^i + \omega_{opt}^i] g_{ij}(q).$$

It follows that $\lim_{\theta^{-1} \rightarrow 0^+} s_j^*(\theta^{-1}, \bar{\kappa}\theta^{-1})$ is a debt security, by the proof of proposition 1.1. Denote s_j^* as this limit. The definition of $s_{debt}(1, \bar{\kappa})$ is that it is the maximizer of

$$\bar{\kappa}E^q[\beta_s s_i] - \frac{1}{2}V^q[\beta_s s_i],$$

over the set of limited liability securities, which is unique by the positive-definiteness of the Fisher information metric. If $s_j^* \neq s_{debt}(1, \bar{\kappa})$, then

$$\bar{\kappa}E^q[\beta_s s_i^*] - \frac{1}{2}V^q[\beta_s s_i^*] < \bar{\kappa}E^q[\beta_s s_{debt,i}(1, \bar{\kappa})] - \frac{1}{2}V^q[\beta_s s_{debt,i}(1, \bar{\kappa})].$$

However, because $s^*(\theta^{-1}, \bar{\kappa}\theta^{-1})$ is the optimal security, we must have

$$U(s^*(\theta^{-1}, \bar{\kappa}\theta^{-1}); \theta^{-1}) - U(0; \theta^{-1}, \bar{\kappa}\theta^{-1}) \geq U(s_{debt}(1, \bar{\kappa}); \theta^{-1}, \bar{\kappa}\theta^{-1}) - U(0; \theta^{-1}, \bar{\kappa}\theta^{-1})$$

for all θ . It follows that

$$\begin{aligned} \lim_{\theta^{-1} \rightarrow 0^+} \theta \{U(s^*(\theta^{-1}, \bar{\kappa}\theta^{-1}); \theta^{-1}, \bar{\kappa}\theta^{-1}) - U(0; \theta^{-1}, \bar{\kappa}\theta^{-1})\} \\ \geq \theta \{U(s_{debt}(1, \bar{\kappa}); \theta^{-1}, \bar{\kappa}\theta^{-1}) - U(0; \theta^{-1}, \bar{\kappa}\theta^{-1})\} \end{aligned}$$

which would require, by proposition 1.3, that

$$\bar{\kappa}E^q[\beta_s s_i^*] - \frac{1}{2}V^q[\beta_s s_i^*] \geq \bar{\kappa}E^q[\beta_s s_{debt,i}(1, \bar{\kappa})] - \frac{1}{2}V^q[\beta_s s_{debt,i}(1, \bar{\kappa})],$$

contradicting the assertion that $s_j^* \neq s_{debt}(1, \bar{\kappa})$. Therefore, $s_j^* = s_{debt}(1, \bar{\kappa})$. Defining S as the set of limited liability securities, the second part of the corollary follows from the fact that

$$U(s^*(\theta^{-1}, \kappa); \theta^{-1}) - U(0; \theta^{-1}, \kappa) \leq \sup_{s \in S} \{ \kappa E^q[\beta_s s_i^*] - \frac{1}{2} \theta^{-1} V^q[\beta_s s_i^*] \} + \sup_{s \in S} \{ O(\theta^{-2} + \kappa \theta^{-1}) \},$$

and therefore

$$U(s^*(\theta^{-1}, \kappa); \theta^{-1}) - U(s_{debt}(\theta^{-1}, \kappa); \theta^{-1}, \kappa) \leq \sup_{s \in S} \{ O(\theta^{-2} + \kappa \theta^{-1}) \}.$$

Because the space of limited liability securities is compact and the utility function is continuous, it follows that

$$U(s^*(\theta^{-1}, \kappa); \theta^{-1}) - U(s_{debt}(\theta^{-1}, \kappa); \theta^{-1}, \kappa) = O(\theta^{-2} + \kappa \theta^{-1}).$$

5.1.11 Proof of lemma 1.1

The proof of this lemma uses Chentsov's theorem and several results from Amari and Nagaoka (2007). We have, for any monotone divergence,

$$\frac{\partial^3 D(p||q)}{\partial p^i \partial p^j \partial p^k} \Big|_{p=q} = ch_{ijk}(q) = c \partial_i g_{jk}(p) \Big|_{p=q} + \Gamma_{jk,i}^{(\alpha)},$$

where $\Gamma_{jk,i}^{(\alpha)}$ are the connection coefficients of the α -connection in the m-flat coordinate system.

Using results in Amari and Nagaoka (2007), p. 33 and 36,

$$\Gamma_{jk,i}^{(\alpha)} = \Gamma_{jk,i}^{(-1)} - \frac{1+\alpha}{2} T_{ijk} = -\frac{1+\alpha}{2} T_{ijk}$$

where T_{ijk} is a covariant symmetric tensor of degree three. Repeating the argument for the Riemannian connection,

$$\Gamma_{ij,k}^{(0)} = -\frac{1}{2}T_{ijk} = \Gamma_{ik,j}^{(0)}$$

and

$$c\partial_i g_{jk} = \Gamma_{ij,k}^{(0)} + \Gamma_{ik,j}^{(0)} = T_{ijk}.$$

It follows that

$$h_{ijk}(q) = \left(\frac{3+\alpha}{2}\right)\partial_i g_{jk}(p)|_{p=q}.$$

5.1.12 Proof of corollary 1.2

From section 5.1.9, the utility for a particular security design can be written as

$$\begin{aligned} U(s; \theta^{-1}, \kappa) &= \beta_s(1+\kappa)\theta^{-1}\eta_i g^{il}(q)v_l - \frac{\theta^{-2}}{2}\beta_s\eta_j\eta_k v_i h^{ijk}(q) + \\ &+ \beta_s(1+\kappa)q^i v_i - \\ &\kappa\theta^{-1}\eta_i g^{il}(q)\eta_l - \kappa q^i \eta_i - \\ &\frac{\theta^{-1}}{2}\eta_i\eta_j g^{ij}(q) + \frac{\theta^{-2}}{3}\eta_i\eta_j\eta_k h^{ijk}(q) + O(\theta^{-3} + \kappa\theta^{-2}). \end{aligned}$$

Taking the FOC with respect to η ,

$$\begin{aligned} \kappa(\lambda^i - \omega^i) &= \beta_s(1+\kappa)\theta^{-1}g^{il}(q)v_l - \theta^{-2}\beta_s\eta_k v_j h^{ijk}(q) - \\ &2\kappa\theta^{-1}g^{il}(q)\eta_l - \kappa q^i - \\ &\theta^{-1}\eta_j g^{ij}(q) + \theta^{-2}\eta_j\eta_k h^{ijk}(q) + O(\theta^{-3} + \kappa\theta^{-2}). \end{aligned}$$

where λ^i and ω^i are the scaled limited liability multipliers. Recall the lemma,

$$h_{ijk}(q) = \left(\frac{3+\alpha}{2}\right) \partial_i g_{jk}(p)|_{p=q}.$$

Define a new probability distribution,

$$\hat{p}(\eta) = q + \left(\frac{3+\alpha}{2}\right)(p(\eta) - q).$$

I will use \hat{p} and p instead of $\hat{p}(\eta)$ and $p(\eta)$ to keep the notation compact. Taylor expanding,

$$\begin{aligned} g^{ij}(\hat{p}) &= g^{ij}(q) - g^{ik}(q)(\partial_l g_{km}(p)|_{p=q})g^{mj}(q)(\hat{p}^l - q^l) + O(\theta^{-2}) \\ &= g^{ij}(q) - g^{ik}(q)h_{lkm}(q)g^{mh}(q)(p^l - q^l) + O(\theta^{-2}). \end{aligned}$$

The approximate first-order condition in the moral hazard problem was

$$\theta^{-1} \eta_i g^{il}(q) + O(\theta^{-2}) = (p^l(\eta) - q^l).$$

Putting these two together,

$$g^{ij}(\hat{p}) = g^{ij}(q) - \theta^{-1} h^{ijk}(q) \eta_k + O(\theta^{-2}).$$

Using these two results,

$$\begin{aligned} \kappa(\lambda^i - \omega^i) &= \beta_s \theta^{-1} g^{il}(\hat{p}) v_l + \\ &\quad \kappa \theta^{-1} g^{il}(q) (\beta_s s_l - \eta_l) - \kappa q^i - \\ &\quad \theta^{-1} \eta_j g^{ij}(\hat{p}) + O(\theta^{-3} + \kappa \theta^{-2}). \end{aligned}$$

$$\begin{aligned}\kappa(\lambda^i - \omega^i) &= \beta_s \theta^{-1} g^{il}(\hat{p}) s_l + \\ &\quad \kappa \theta^{-1} g^{il}(q) (\beta_s s_l - \eta_l) - \kappa q^i + O(\theta^{-3} + \kappa \theta^{-2}).\end{aligned}$$

$$\kappa(\hat{p}^l - q^l) = \left(\frac{3+\alpha}{2}\right) \kappa \theta^{-1} \eta_l g^{il}(q) + O(\kappa \theta^{-2})$$

$$\begin{aligned}\kappa(\lambda^i - \omega^i) &= \beta_s \theta^{-1} g^{il}(\hat{p}) s_l + \kappa \theta^{-1} g^{il}(q) \beta_s s_l \\ &\quad - \kappa \hat{p}^i + \left(\frac{\alpha+1}{2}\right) \kappa \theta^{-1} \eta_l g^{il}(q) + O(\theta^{-3} + \kappa \theta^{-2}).\end{aligned}$$

Note that

$$\kappa \theta^{-1} g^{ij}(\hat{p}) = \kappa \theta^{-1} g^{ij}(q) + O(\kappa \theta^{-2}).$$

Therefore,

$$\begin{aligned}\kappa(\lambda^i - \omega^i) &= \beta_s \theta^{-1} \left(1 + \kappa \frac{1-\alpha}{2}\right) g^{il}(\hat{p}) s_l \\ &\quad - \kappa \hat{p}^i + \beta_s \left(\frac{1+\alpha}{2}\right) \kappa \theta^{-1} v_l g^{il}(q) + O(\theta^{-3} + \kappa \theta^{-2}).\end{aligned}$$

Solving for the security design,

$$s_l \left[\beta_s + \frac{1-\alpha}{2} \beta_s \kappa\right] \theta^{-1} = \kappa(\hat{p}^i - \lambda^i + \omega^i) g_{il}(\hat{p}) - \frac{1+\alpha}{2} \beta_s \kappa \theta^{-1} v_l + O(\theta^{-3} + \kappa \theta^{-2}).$$

This expression can be rewritten (ignoring higher order terms) as

$$\beta_s s_l (1 + \kappa) + \frac{1 + \alpha}{2} \kappa \eta_l(s) = \kappa \theta (\hat{p}^i - \lambda^i + \omega^i) g_{il}(\hat{p}).$$

Taking expected values under the \hat{p} distribution,

$$\beta_s \hat{p}^l s_l (1 + \kappa) + \frac{1 + \alpha}{2} \kappa \hat{p}^l \eta_l(s) = \kappa \theta \hat{p}^l (\hat{p}^i - \lambda^i + \omega^i) g_{il}(\hat{p}).$$

Recalling that

$$g_{il}(\hat{p}) = \left(\frac{\delta_{il}}{\hat{p}^i} + \frac{1}{\hat{p}^0} \right),$$

$$\begin{aligned} \beta_s \hat{p}^l s_l (1 + \kappa) + \frac{1 + \alpha}{2} \kappa \hat{p}^l \eta_l(s) = \\ \kappa \theta \sum_i [\hat{p}^i - \lambda^i + \omega^i] + \\ \kappa \theta \sum_i [\hat{p}^i - \lambda^i + \omega^i] \frac{1 - \hat{p}^0}{\hat{p}^0}. \end{aligned}$$

Therefore,

$$\beta_s \hat{p}^l s_l (1 + \kappa) + \frac{1 + \alpha}{2} \kappa \hat{p}^l \eta_l(s) = \frac{\kappa \theta}{\hat{p}^0} \sum_i [\hat{p}^i - \lambda^i + \omega^i].$$

Plugging this back in,

$$\begin{aligned} \beta_s s_l (1 + \kappa) + \frac{1 + \alpha}{2} \kappa \eta_l(s) = \beta_s \hat{p}^l s_l (1 + \kappa) + \frac{1 + \alpha}{2} \kappa \hat{p}^l \eta_l(s) + \\ \kappa \theta \sum_{i>0} [\hat{p}^i - \lambda^i + \omega^i] \frac{\delta_{ij}}{\hat{p}^i}. \end{aligned}$$

Define

$$\bar{v} = \hat{p}^l s_l + \frac{\frac{1+\alpha}{2} \kappa \hat{p}^l \eta_l(s) + \kappa \theta}{\beta_s(1+\kappa)}.$$

Suppose that $\omega^i > 0$, and therefore $s_l = 0$, $\eta_l = \beta_s v_l$, $\lambda_l = 0$. Then

$$\frac{1+\alpha}{2} \beta_s \kappa v_l = \beta_s(1+\kappa) \bar{v} + \kappa \theta \frac{\omega^i}{\hat{p}^i}.$$

This can never occur if $\alpha \leq -1$ and $\bar{v} \geq 0$, by the requirement that $v_l > 0$. If $\alpha > -1$, we know that $\bar{v} > 0$, and the condition occurs only if

$$\frac{1+\alpha}{2} \beta_s \kappa v_l - \beta_s(1+\kappa) \bar{v} > 0,$$

which requires large values of v_l . Next, consider the $\lambda^i > 0$ case, when $s_l = v_l$ and $\eta_l = 0$ (and $\omega_l = 0$). In this case,

$$\beta_s v_l(1+\kappa) = \beta_s(1+\kappa) \bar{v} - \kappa \theta \frac{\lambda^i}{\hat{p}^i}.$$

For any $v_l < \bar{v}$, this condition can hold. Finally, consider the case when $\lambda^i = \omega^i = 0$. In this case,

$$\beta_s s_l(1+\kappa) + \frac{1+\alpha}{2} \kappa \eta_l(s) = \beta_s(1+\kappa) \bar{v}.$$

First, assume that $\alpha \leq -1$. In this case, it follows that if $v_l < \bar{v}$, this condition cannot hold, and otherwise it can. Therefore, there are two regions for $\alpha \leq -1$: a region where $s_l = v_l$, for $v_l < \bar{v}$, and a region where

$$s_l = \frac{(1+\kappa)}{(1+\kappa \frac{1-\alpha}{2})} \bar{v} - \frac{\frac{1+\alpha}{2} \kappa}{(1+\kappa \frac{1-\alpha}{2})} v_l$$

for $v_l > \bar{v}$. This simplifies to

$$s_l = \bar{v} - \frac{\kappa(1+\alpha)}{(2+\kappa(1-\alpha))} (v_l - \bar{v}).$$

It follows that

$$\eta_l = -\frac{\beta_s(1+\kappa)}{(1+\kappa\frac{1-\alpha}{2})}\bar{v} + \frac{\beta_s(1+\kappa)}{(1+\kappa\frac{1-\alpha}{2})}v_l$$

in this region. The equation defining \bar{v} becomes

$$\bar{v} \sum_{i:v_l < \bar{v}} \hat{p}^i = \sum_{i:v_l < \bar{v}} \hat{p}^i v_i + \frac{\kappa\theta}{\beta_s(1+\kappa)}$$

which is guaranteed to be greater than zero. Note that the “option” value of the debt is the same as in the KL divergence case, but under the \hat{p} probability distribution. Next, consider $\alpha \in (-1, 1 + \frac{2}{\kappa})$.

It follows that

$$\frac{1+\alpha}{2}\kappa \in (0, 1+\kappa)$$

and therefore $\bar{v} > 0$. The condition that

$$\frac{1+\alpha}{2}\beta_s\kappa v_l - \beta_s(1+\kappa)\bar{v} > 0$$

is satisfied only for some $v_l > v_{max} > \bar{v}$. Suppose that $v_l > \bar{v}$. The left hand side of the zero-multiplier condition below,

$$\beta_s s_l(1+\kappa) + \frac{1+\alpha}{2}\kappa\eta_l(s) - \beta_s(1+\kappa)\bar{v} = 0,$$

achieves its minimum when $\eta_l(s) = \beta_s v_l$, and therefore this cannot hold if $v_l > v_{max}$. By a similar argument, it achieves its maximum when $s_l = v_l$, and therefore cannot hold if $v_l < \bar{v}$. It follows that there are three regions: a low v_l region, in which $v_l < \bar{v}$, an intermediate $v_l \in (\bar{v}, v_{max})$ region in which

$$s_l = \bar{v} - \frac{\kappa(1+\alpha)}{(2+\kappa(1-\alpha))}(v_l - \bar{v}),$$

and a high v_l region in which $s_l = 0$. Therefore, $s_{debt-eq}$ is the second-order optimal security design.

The statement that

$$U(s^*(\theta^{-1}, \kappa); \theta^{-1}, \kappa) - U(s_{debt-eq}(\theta^{-1}, \kappa); \theta^{-1}, \kappa) = O(\theta^{-3} + \kappa\theta^{-2})$$

follows, from the fact that $s_{debt-eq}(\theta^{-1}, \kappa)$ maximizes the non- $O(\theta^{-3} + \kappa\theta^{-2})$ terms of $U(s; \theta^{-1})$ (see the proof of corollary 1.1).

5.1.13 Proof of proposition 1.5

This proof is similar in structure to proposition 1.3. I will continue to use the summation notation. We define the utility generated by a particular contract as

$$U(s; M_\xi, \theta^{-1}, \kappa) = \beta_s(1 + \kappa)p^i(\eta(s); M_\xi, \theta^{-1})s_i + \phi(\eta(s); M_\xi, \theta^{-1}).$$

It will also be useful to define $p^i(\xi)$, which is the probability distribution induced by the parameters ξ . Define

$$B_a^i(\bar{p}) = \frac{\partial p^i(\xi)}{\partial \xi^a} \Big|_{p=\bar{p}},$$

$$B_{ab}^i(\bar{p}) = \frac{\partial p^i(\xi)}{\partial \xi^a \partial \xi^b} \Big|_{p=\bar{p}},$$

The first-order condition for the moral hazard problem is

$$B_a^i(p(\xi))(\eta_i - \partial_i \psi(p(\xi))) = 0.$$

As mentioned in the text, there may be multiple ξ for which this FOC is satisfied. By assumption, however, the FOC holds at all ξ that maximize the seller's problem. In the neighborhood of $\theta \rightarrow \infty$,

$$\lim_{\theta \rightarrow \infty} p(\eta; M_\xi, \theta^{-1}) = q,$$

because otherwise the seller would receive unbounded negative utility. By assumption, $q \in M_\xi$, and therefore there are coordinates ξ_q such that $p(\xi_q) = q$. Expanding the functions $B_a^i(p(\xi))$ and $B_a^i(p(\xi))\partial_i\psi(p(\xi))$ around $\xi = \xi_q$, using the fact that $\partial_i\psi(q) = 0$,

$$B_a^i(p(\xi)) = B_a^i(q) + B_{ab}^i(q)(\xi^b - \xi_q^b) + B_{abc}^i(\hat{p})(\xi^b - \xi_q^b)(\xi^c - \xi_q^c),$$

$$\begin{aligned} B_a^i(p(\xi))\partial_i\psi(p(\xi)) &= B_a^i(q)B_b^j(q)\partial_i\partial_j\psi(q)(\xi^b - \xi_q^b) + \\ &B_a^i(p^*)B_{bc}^j(p^*)\partial_i\partial_j\psi(p^*)(\xi^b - \xi_q^b)(\xi^c - \xi_q^c) + \\ &\frac{1}{2}B_a^i(p^*)B_b^j(p^*)B_c^k(p^*)\partial_i\partial_j\partial_k\psi(p^*)(\xi^b - \xi_q^b)(\xi^c - \xi_q^c) + \\ &\frac{1}{2}B_{ab}^i(p^*)B_c^j(p^*)\partial_i\partial_j\psi(p^*)(\xi^b - \xi_q^b)(\xi^c - \xi_q^c), \end{aligned}$$

where $p^* = q + c^*(p - q)$ for some $c^* \in (0, 1)$ and \hat{p} similarly defined. Define $g_{ab}(q) = B_a^i(q)B_b^j(q)g_{ij}(q)$ and $g^{ab}(q)$ as its inverse. Define

$$m_{ab}(q, \eta) = \theta g_{ab}(q) - B_{ab}^i(q)\eta_i$$

and

$$\begin{aligned} m_{abc}(\hat{p}, p^*, \eta_i) &= B_a^i(p^*)B_{bc}^j(p^*)\partial_i\partial_j\psi(p^*) + \frac{1}{2}B_a^i(p^*)B_b^j(p^*)B_c^k(p^*)\partial_i\partial_j\partial_k\psi(p^*) + \\ &\frac{1}{2}B_{ab}^i(p^*)B_c^j(p^*)\partial_i\partial_j\psi(p^*) - B_{abc}^i(\hat{p})\eta_i. \end{aligned}$$

The FOC can be written, using the invariance property of f -divergences, as

$$B_a^i(q)\eta_i = m_{ab}(q, \eta)(\xi^b - \xi_q^b) + m_{abc}(\hat{p}, p^*, \eta)(\xi^b - \xi_q^b)(\xi^c - \xi_q^c).$$

Multiplying by $g^{ab}(q)\theta^{-1}$,

$$\begin{aligned} \xi^b - \xi_q^b &= \theta^{-1}g^{ab}(q)B_a^i(q)\eta_i + \theta^{-1}(\xi^c - \xi_q^c)g^{ab}(q)B_{ac}^i(q)\eta_i + \\ &\quad \theta^{-1}g^{ab}(q)m_{acd}(\hat{p}, p^*, \eta)(\xi^d - \xi_q^d)(\xi^c - \xi_q^c). \end{aligned}$$

Note that $\theta^{-1}m_{acd}(\hat{p}, p^*, \eta) = O(1)$ in θ . By the argument that $\lim_{\theta \rightarrow \infty} p(\eta; M_\xi, \theta^{-1}) = q$, we know that $\xi^b - \xi_q^b = o(1)$. It follows that

$$\xi^b - \xi_q^b = \theta^{-1}g^{ab}(q)B_a^i(q)\eta_i + O(\theta^{-2}).$$

This demonstrates that $\xi^b(\eta)$, and therefore $p(\eta; M_\xi, \theta^{-1})$, is locally unique when θ is large. It follows that $p(\eta; M_\xi, \theta^{-1})$ is differentiable in the neighborhood of the expansion, regardless of the arbitrary rule used to “break ties” when multiple p maximize the moral hazard sub-problem.

We now proceed as before, substituting in for the definition of β_b and $\phi(\eta)$,

$$U(s; M_\xi, \theta^{-1}, \kappa) = \beta_s(1 + \kappa)p^i(\eta(s); M_\xi, \theta^{-1})v_i - \kappa p^i(\eta(s); M_\xi, \theta^{-1})\eta_i(s) - \psi(p(\eta(s); M_\xi, \theta^{-1})).$$

Expanding the function $p(\xi)$ around $\xi = \xi_q$, and using the result above,

$$p^i(\eta; M_\xi, \theta^{-1}) - q^i = \theta^{-1}B_a^i(q)g^{ab}(q)B_b^j(q)\eta_j + O(\theta^{-2})$$

Expanding $\psi(p(\xi))$, again using the result from above,

$$\psi(p) = \frac{\theta^{-1}}{2}\eta_i\eta_j B_a^i(q)B_b^j(q)g_{ij}(q) + O(\theta^{-2}).$$

Plugging these two into the utility function,

$$U(s; M_\xi, \theta^{-1}, \kappa) = \beta_s(1 + \kappa)q^i v_i + \beta_s \theta^{-1} g^{ab}(q) B_a^i(q) B_b^j(q) \eta_i(s) v_j - \kappa q^i \eta_i(s) - \frac{1}{2} \theta^{-1} B_a^i(q) B_b^j(q) g^{ab}(q) \eta_i(s) \eta_j(s) + O(\theta^{-2} + \theta^{-1} \kappa).$$

The no-trade utility is

$$U(0; M_\xi, \theta^{-1}, \kappa) = \beta_s q^i v_i + \frac{1}{2} \beta_s^2 \theta^{-1} B_a^i(q) B_b^j(q) g^{ab}(q) v_i v_j + O(\theta^{-2} + \theta^{-1} \kappa).$$

Taking the difference,

$$\begin{aligned} U(s; M_\xi, \theta^{-1}, \kappa) - U(0; M_\xi, \theta^{-1}, \kappa) &= \beta_s \kappa q^i s_i + \beta_s \theta^{-1} B_a^i(q) B_b^j(q) g^{ab}(q) \eta_i(s) v_j - \\ &\quad - \frac{1}{2} \theta^{-1} B_a^i(q) B_b^j(q) g^{ab}(q) \eta_i(s) \eta_j(s) - \\ &\quad \frac{1}{2} \beta_s^2 \theta^{-1} B_a^i(q) B_b^j(q) g^{ab}(q) v_i v_j + O(\theta^{-2} + \theta^{-1} \kappa) \end{aligned}$$

which simplifies to

$$U(s; M_\xi, \theta^{-1}, \kappa) - U(0; M_\xi, \theta^{-1}, \kappa) = \beta_s \kappa q^i s_i - \frac{1}{2} \beta_s^2 \theta^{-1} B_a^i(q) B_b^j(q) g^{ab}(q) s_i s_j + O(\theta^{-2} + \theta^{-1} \kappa).$$

The lower bound,

$$U(s; M_\xi, \theta^{-1}, \kappa) - U(0; M_\xi, \theta^{-1}, \kappa) \geq \kappa E^q[\beta_s s_i] - \frac{\theta^{-1}}{2} \text{Var}^q[\beta_s s_i] + O(\theta^{-2} + \theta^{-1} \kappa),$$

follows from theorems 2.7 and 2.8 of Amari and Nagaoka (2007). Note that

$$\frac{\partial}{\partial \xi^a} E^p[s_i] \Big|_{p=q} = s_i B_a^i(q).$$

From theorems 2.7 and 2.8, we have

$$\text{Var}^q[s_i] = s_i s_j g_{ij}(q) \geq s_i s_j B_a^i(q) B_b^j(q) g^{ab}(q).$$

The tightness of the lower bound applies when M_ξ is an exponential family, and s_i is a linear combination of its sufficient statistics, by the definition of a sufficient statistic.

5.1.14 Proof of proposition 5.1

This proof follows the proof of proposition 1.5. I will use the summation convention. We can write the utility from a particular security as

$$U_{RI}(s; \theta^{-1}, \kappa) = \beta_s(1 + \kappa) p^i(\eta(s); \theta^{-1}) s_i + \phi_{RI}(\eta(s); \theta^{-1}),$$

$$\phi_{RI}(\eta; \theta^{-1}) = \max_{p \in M_{RI}} \{p^i \eta_i - \theta D_{KL}(p||q(p))\}.$$

Define a (for now) arbitrary distribution \bar{p} , with the property that actions are independent of states, and the marginal distribution of states is $g(x)$. We can write

$$\psi(p) = \theta D_{KL}(p||q(p)) = \theta D_{KL}(p||\bar{p}) - \theta D_{KL}(q(p)||\bar{p}),$$

where $q_a(p)$ refers to the marginal distribution over actions of q , and \bar{p}_a is the marginal distribution of actions over \bar{p} . The key to this idea is that $q(p)$ has actions independent of states, and $q_a(p) = p_a$, and therefore this equation holds for any full support \bar{p} .

As noted earlier, the set of feasible probability distributions M_{RI} is an $|X| \cdot (|A| - 1)$ dimensional space, not an N dimensional space. It is an exponential family embedded in the space of all probability distributions. Denote a flat coordinate system ξ , such that

$$\frac{\partial p^i(\xi)}{\partial \xi^b} = B_b^i.$$

Note that unlike the general parametric model, B_b^i is a constant matrix. Similarly, define

$$\frac{\partial q^i(p)}{\partial p^j} = C_j^i.$$

The first-order condition for the moral hazard problem is

$$B_b^i(\eta_i - \partial_i \psi(p(\xi))) = 0.$$

As in the parametric problem, there may be multiple ξ for which this FOC is satisfied. By assumption, however, the FOC holds at all ξ that maximize the seller's problem. In the neighborhood of $\theta \rightarrow \infty$,

$$\lim_{\theta \rightarrow \infty} p(\eta; \theta^{-1}) = \lim_{\theta \rightarrow \infty} q(p(\eta; \theta^{-1})),$$

because otherwise the seller would receive unbounded negative utility. This is different from the parametric problem, because even in the limit as θ becomes large, the unconditional distribution of actions might depend on the retained tranche. The only requirement in this limit is that actions are independent of states. However, by the assumption of uniqueness, for each η there exists a unique $\bar{p}(\eta)$ such that

$$\lim_{\theta \rightarrow \infty} p(\eta; \theta^{-1}) = \bar{p}(\eta).$$

By assumption, $\bar{p} \in M_\xi$, and therefore there are coordinates $\xi_{\bar{p}}$ such that $p(\xi_{\bar{p}}) = \bar{p}$. Expanding the function $\partial_i \psi(p(\xi))$ around $\xi = \xi_{\bar{p}}$, using the fact that $\partial_i \psi(\bar{p}) = 0$,

$$B_b^i \partial_i \psi(p(\xi)) = B_b^i B_c^j \partial_i \partial_j \psi(\bar{p}(\eta)) (\xi^c - \xi_{\bar{p}}^c) + \frac{1}{2} B_d^i B_b^j B_c^k \partial_i \partial_j \partial_k \psi(p^*) (\xi^d - \xi_{\bar{p}}^d) (\xi^c - \xi_{\bar{p}}^c),$$

where $p^* = q + c^*(p - q)$ for some $c^* \in (0, 1)$. Using the properties of the KL divergence,

$$B_b^i B_c^j \partial_i \partial_j \psi(\bar{p}(\eta)) = \theta B_b^i B_c^j [g_{ij}(\bar{p}(\eta)) - C_i^k C_j^l g_{kl}(\bar{p}(\eta))] = \theta m_{bc}(\bar{p}(\eta)),$$

$$B_b^i B_c^j B_d^k \partial_i \partial_j \partial_k \psi(p) = \theta B_b^i B_c^j B_d^k [h_{ijk}(p) - C_i^l C_j^m C_k^n g_{lmn}(p)] = \theta m_{bcd}(p).$$

The FOC can be rearranged to

$$m_{bc}(\bar{p}(\eta)) (\xi^c - \xi_{\bar{p}}^c) = \theta^{-1} B_b^i \eta_i - \frac{1}{2} (\xi^d - \xi_{\bar{p}}^d) (\xi^c - \xi_{\bar{p}}^c) m_{bcd}(p^*).$$

Note that unlike the previous proof, $m_{bc}(\bar{p})$ is positive-semidefinite but singular, and therefore not invertible. Nevertheless, this is sufficient to show that

$$m_{bc}(\bar{p}(\eta)) (\xi^c - \xi_{\bar{p}}^c) = \theta^{-1} B_b^i \eta_i + O(\theta^{-2}).$$

We now proceed as before, substituting in for the definition of β_b and $\phi(\eta)$,

$$U(s; \theta^{-1}, \kappa) = \beta_s (1 + \kappa) p^i(\eta(s); \theta^{-1}) v_i - \kappa p^i(\eta(s); \theta^{-1}) \eta_i(s) - \psi(p(\eta(s); \theta^{-1})).$$

Expanding the function $p(\xi)$ around $\xi = \xi_{\bar{p}}$, and using the result above,

$$p^i(\eta; \theta^{-1}) - \bar{p}^i(\eta) = \theta^{-1} B_b^i (\xi^b - \xi_{\bar{p}}^b) + O(\theta^{-2})$$

Expanding $\psi(p(\xi))$, again using the result from above,

$$\psi(p) = \frac{\theta}{2} m_{bc}(\bar{p})(\xi^b - \xi_{\bar{p}}^b)(\xi^c - \xi_{\bar{p}}^c) + O(\theta^{-2}).$$

Therefore,

$$\begin{aligned} U(s; \theta^{-1}, \kappa) &= \beta_s(1 + \kappa)\bar{p}^i(\eta(s))v_i + \beta_s B_b^i(\xi^b - \xi_{\bar{p}}^b)v_i - \\ &\quad \kappa\bar{p}^i(\eta(s))\eta_i(s) - \\ &\quad \frac{\theta}{2} m_{bc}(\bar{p}(\eta(s)))(\xi^b - \xi_{\bar{p}}^b)(\xi^c - \xi_{\bar{p}}^c) + O(\theta^{-2} + \theta^{-1}\kappa). \end{aligned}$$

When the security is trading everything,

$$U(v; \theta^{-1}, \kappa) = \beta_s(1 + \kappa)\bar{p}^i(0)v_i,$$

where $\bar{p}^i(0)$ is the arbitrary decision the seller will make about unconditional actions, when they have no incentive whatsoever. When the security is trading nothing,

$$\begin{aligned} U(0; \theta^{-1}, \kappa) &= \beta_s\bar{p}^i(\beta_s v)v_i + \beta_s B_b^i(\xi_0^b - \xi_{\bar{p}(0)}^b)v_i - \\ &\quad - \frac{\theta}{2} m_{bc}(\bar{p}(\beta_s v))(\xi_0^b - \xi_{\bar{p}(0)}^b)(\xi_0^c - \xi_{\bar{p}(0)}^c) + O(\theta^{-2} + \theta^{-1}\kappa), \end{aligned}$$

where ξ_0^b is the endogenous distribution for the zero security, and $\xi_{\bar{p}(0)}^b$ is the associated limiting distribution. Because the expected payoff is the same for all actions, when those actions are independent of state,

$$\bar{p}^i(\beta_s v)v_i = \bar{p}^i(\eta)v_i = \bar{p}^i(0)v_i,$$

and

$$C_j^i v_i = 0$$

for all j . We know that for any ξ_{ind}^b that generates independence between actions and states,

$$(\xi_{ind}^b - \xi_{\bar{p}}^b) m_{ab}(\bar{p}) (\xi_{ind}^c - \xi_{\bar{p}}^c) = 0,$$

and therefore that $\xi_{ind}^b - \xi_{\bar{p}}^b$ is in the null space of $m_{bc}(\bar{p})$. Therefore, the nullity of $m_{bc}(\bar{p})$ is at least $|A| - 1$. The rank of $B_b^i B_c^j g_{ij}(\bar{p})$ is $(|A| - 1) \cdot |X|$ (full rank), and the rank of $B_b^i B_c^j C_i^k C_j^l g_{kl}(\bar{p})$ is $|A| - 1$, so the rank of $m_{ab}(\bar{p})$ satisfies

$$(|A| - 1) \cdot |X| \leq \text{rank}(m_{bc}(\bar{p})) + |A| - 1.$$

It follows that the nullity of $m_{bc}(\bar{p})$ is exactly equal to $|A| - 1$, the dimension of marginal distributions of actions. This fact is useful later in the proof. Moreover, because $B_b^j C_j^i v_i = 0$, $B_b^j v_j$ lies entirely in the column-space of $m_{ab}(\bar{p})$. Therefore, there exists a vector $v^c(\bar{p})$ such that

$$B_b^j v_j = m_{bc}(\bar{p}) v^c(\bar{p}).$$

We can therefore rewrite

$$\begin{aligned} U(s; \theta^{-1}, \kappa) - U(v; \theta^{-1}, \kappa) &= \beta_s (\xi^b - \xi_{\bar{p}}^b) m_{bc}(\bar{p}(\eta(s))) v^c(\bar{p}(\eta(s))) - \kappa \bar{p}^i(\eta(s)) \eta_i(s) \\ &\quad - \frac{\theta}{2} m_{bc}(\bar{p}(\eta(s))) (\xi^b - \xi_{\bar{p}}^b) (\xi^c - \xi_{\bar{p}}^c) + O(\theta^{-2} + \theta^{-1} \kappa). \end{aligned}$$

Consider a generalized inverse of $m_{bc}(\bar{p})$, $m_+^{bc}(\bar{p})$, which has the property that

$$m_{bc}(\bar{p}) m_+^{cd}(\bar{p}) m_{de}(\bar{p}) = m_{be}(\bar{p}).$$

Rewriting the utility condition,

$$\begin{aligned}
U(s; \theta^{-1}, \kappa) - U(v; \theta^{-1}, \kappa) &= \beta_s (\xi^b - \xi_{\bar{p}}^b) m_{bc}(\bar{p}(\eta(s))) v^c (\bar{p}(\eta(s)) - \kappa \bar{p}^i(\eta(s))) \eta_i(s) \\
&\quad - \frac{\theta}{2} m_{bc}(\bar{p}(\eta(s))) m_+^{cd}(\bar{p}(\eta(s))) m_{be}(\bar{p}(\eta(s))) (\xi^b - \xi_{\bar{p}}^b) (\xi^e - \xi_{\bar{p}}^e) \\
&\quad + O(\theta^{-2} + \theta^{-1} \kappa).
\end{aligned}$$

Applying the first-order condition,

$$\begin{aligned}
U(s; \theta^{-1}, \kappa) - U(v; \theta^{-1}, \kappa) &= \beta_s \theta^{-1} B_b^j B_c^i m_+^{bc}(\bar{p}(\eta(s))) \eta_i v_j - \kappa \bar{p}^i(\eta(s)) \eta_i(s) \\
&\quad - \frac{\theta^{-1}}{2} m_+^{bc}(\bar{p}(\eta(s))) B_b^i B_c^j \eta_i \eta_j + O(\theta^{-2} + \theta^{-1} \kappa).
\end{aligned}$$

For the sell-nothing security,

$$U(0; \theta^{-1}) - U(v; \theta^{-1}) = \frac{1}{2} \beta_s^2 \theta^{-1} B_b^j B_c^i m_+^{bc}(\bar{p}(\beta_s v)) v_i v_j - \kappa \theta^{-1} \bar{p}^i(\beta_s v) v_i + O(\theta^{-2}).$$

Because

$$g_{ab}(\bar{p}) = m_{ab}(\bar{p}) + B_b^i B_c^j C_i^k C_j^l g_{kl}(\bar{p})$$

is non-singular, it follows from Hearon (1967) that a generalized inverse can be constructed as

$$m_+^{bc}(\bar{p}) = g^{bd}(\bar{p}) m_{de}(\bar{p}) g^{ce}(\bar{p}) = g^{bc}(\bar{p}) - g^{bd}(\bar{p}) B_d^i B_e^j C_i^k C_j^l g_{kl}(\bar{p}) g^{ce}(\bar{p}).$$

Because $B_b^i v_i$ is entirely in the column space of $m_{bc}(\bar{p})$, it is entirely in the row-space of $m_+^{bc}(\bar{p})$, and therefore

$$m_+^{bc}(\bar{p}) B_b^i B_c^j v_i v_j = g^{bc}(\bar{p}) B_b^i B_c^j v_i v_j.$$

Because the unconditional variance of asset values is identical across assets,

$$g^{bc}(\bar{p}(\beta_s v)) B_b^i B_c^j v_i v_j = g^{bc}(\bar{p}(\eta)) B_b^i B_c^j v_i v_j$$

for all η . It therefore follows that

$$\begin{aligned} U(s; \theta^{-1}, \kappa) - U(0; \theta^{-1}, \kappa) &= \beta_s \kappa \bar{p}^i(\eta(s)) s_i \\ &\quad - \frac{\theta^{-1}}{2} \beta_s^2 m_+^{bc}(\bar{p}(\eta(s))) B_b^i B_c^j s_i s_j + O(\theta^{-2} + \theta^{-1} \kappa). \end{aligned}$$

By the positive-definiteness of $g^{bd}(\bar{p}) B_d^i B_e^j C_i^k C_j^l g_{kl}(\bar{p}) g^{ce}(\bar{p})$, and the monotonicity of the Fisher metric,

$$U(s; \theta^{-1}, \kappa) - U(0; \theta^{-1}, \kappa) \geq \beta_s \kappa \bar{p}^i(\eta(s)) s_i - \frac{\theta^{-1}}{2} \beta_s^2 g^{ij}(\bar{p}^i(\eta(s))) s_i s_j + O(\theta^{-2} + \theta^{-1} \kappa),$$

which is the desired result.

5.1.15 Proof of proposition 1.6

The relaxed moral hazard problem can be written as

$$\phi_{CT}(\eta) = \sup_{\frac{dP}{dQ}} \{E^P[\eta(X)] - \theta D_{KL}(P||Q)\}.$$

A suitably modified version of the proof of proposition 1.1 could be applied to this problem, but the extension of the some of the shortcuts used in that proof is not straightforward. Instead, I will use a calculus-of-variations approach, similar to the earlier drafts of this chapter. The result is a specialized version of Cvitanić et al. (2009), and I will rely on that paper for the proof of the existence of an optimal security design. I will also assume that the space of allowed security designs is restricted to depend on the asset value at an arbitrary set of times, which includes the

final value. This assumption is necessary to use a theorem from the calculus of variations, but is not required for the proof of Cvitanić et al. (2009).

Let \hat{P} be an alternative measure on Ω that is absolutely continuous with respect to Q . Define

$$\frac{dP(\eta, \alpha)}{dQ} = (1 - \alpha) \frac{dP^*(\eta)}{dQ} + \alpha \frac{d\hat{P}}{dQ},$$

where $P^*(\eta)$ is the measure the maximizes the moral hazard problem. The retained tranche $\eta : \Omega \rightarrow \mathbb{R}$ is a \mathcal{F}_1^B -measurable function on the sample space Ω . To keep notation compact, I use η instead of $\eta(X)$ and $\frac{dP}{dQ}$ instead of $\frac{dP}{dQ}(B)$ when the meaning is not ambiguous.

$$\phi_{CT}(\eta, \alpha) = \int_{\Omega} \frac{dP(\eta, \alpha)}{dQ} dQ - \theta \int_{\Omega} \frac{dP(\eta, \alpha)}{dQ} \ln\left(\frac{dP(\eta, \alpha)}{dQ}\right) dQ + \theta \int_{\Omega} \left[\frac{dP(\eta, \alpha)}{dQ} - 1\right] dQ,$$

we must have

$$\frac{\partial \phi_{CT}(\eta, \alpha)}{\partial \alpha} \Big|_{\alpha=0} \leq 0$$

for all \hat{P} . It follows that, for all \hat{P} ,

$$\int_{\Omega} \left[\frac{d\hat{P}}{dQ} - \frac{dP^*(\eta)}{dQ}\right] [\eta - \theta \ln\left(\frac{dP^*(\eta)}{dQ}\right)] dQ \leq 0.$$

This will be satisfied only for

$$\frac{dP^*(\eta)}{dQ} = \exp(\theta^{-1} \eta - \lambda),$$

for some constant λ . Of course, $\frac{dP^*(\eta)}{dQ}$ must be a valid Radon-Nikodym derivative, and therefore

$$\lambda = \ln(E^Q[\exp(\theta^{-1} \eta)]) > 0.$$

The integrability assumption is $E^Q[\exp(4\theta^{-1}X_1)] < \infty$, and therefore $E^Q[\exp(\theta^{-1}\eta)]$ is finite for all limited liability η .

The security design problem can be written as

$$\begin{aligned} U_{CT}(\eta) &= \beta_b E^P[X_1 - \beta_s^{-1}\eta] + \phi_{CT}(\eta) \\ &= \beta_b E^P[X_1 - \beta_s^{-1}\eta] + \theta \ln(E^Q[\exp(\theta^{-1}\eta)]). \end{aligned}$$

This equation is identical to the one in Yang (2012) (proposition 3), and the proof (from this point) is essentially the same. I define

$$\eta(X, \varepsilon) = \eta^*(X) + \varepsilon \tau(X),$$

where $\tau : \Omega \rightarrow \mathbb{R}$ is another measurable function, restricted to the same set of payoff-relevant times.

Using the calculus-of-variations approach again,

$$\frac{\partial U_{CT}(\eta(X, \varepsilon))}{\partial \varepsilon} \Big|_{\varepsilon=0} \leq 0$$

for all τ for which there exists some $\varepsilon > 0$ such that $\eta(X, \varepsilon)$ is a limited liability security. We have

$$\frac{\partial}{\partial \varepsilon} \frac{dP^*(\eta)}{dQ}(B) \Big|_{\varepsilon=0} = \theta^{-1} \frac{dP^*(\eta^*)}{dQ}(B) (\tau(X(B)) - E^{P^*(\eta^*)}[\tau]),$$

and

$$\frac{\partial}{\partial \varepsilon} \phi(\eta) \Big|_{\varepsilon=0} = E^{P^*(\eta^*)}[\tau].$$

It follows that, for all τ , if an η^* that maximizes U_{CT} exists, then

$$\left(1 - \frac{\beta_b}{\beta_s}\right) \int_{\Omega} \frac{dP^*(\eta^*)}{dQ} \tau dQ + \theta^{-1} \beta_b \int_{\Omega} \frac{dP^*(\eta^*)}{dQ} (\tau - E^{P^*(\eta^*)}[\tau]) sdQ \leq 0.$$

This can be rearranged to

$$(1 - \frac{\beta_b}{\beta_s}) \int_{\Omega} \frac{dP^*(\eta^*)}{dQ} \tau dQ + \theta^{-1} \beta_b \int_{\Omega} \frac{dP^*(\eta^*)}{dQ} \tau (s - E^{P^*(\eta^*)}[s]) dQ \leq 0.$$

Because the set of payoff-relevant times is fixed and finite, these integrals can be rewritten as integrals over \mathbb{R}^M , where M is the number of payoff-relevant times. By the du Bois-Reymond lemma, for any interior $s(X)$,

$$\theta \frac{\beta_b - \beta_s}{\beta_b \beta_s} = s(X) - E^{P^*(\eta^*)}[s(X)].$$

Note that if $s(X)$ is zero, decreasing $\eta(X)$ (increasing $s(X)$) will increase utility, and therefore $s(X) > 0$ for all paths with $X_1 > 0$. By an argument similar to proposition 1.1, the optimal contract is a debt security. The security depends only on the time-one value of the asset.

5.1.16 Proof of lemma 1.2

The first part of the lemma is a restatement of lemma 7.1 in Cvitanić et al. (2009). The second part essentially restates a result found Bierkens and Kappen (2012), proposition 3.2. The uniqueness, up to an evanescence, of $u(X, t)$ is shown below.

For all $P \in M$,

$$E^Q[(\frac{dP}{dQ})^2] = E^P[\frac{dP}{dQ}] < \infty.$$

By the inequality that $\ln x < x - 1$ for all $x > 0$, and the absolute continuity of $\frac{dP}{dQ}$,

$$E^P[\ln(\frac{dP}{dQ})] + 1 < E^P[\frac{dP}{dQ}] < \infty.$$

Therefore, $D_{KL}(P||Q)$ is finite. It follows that the process $u(X, t)$ given by proposition 3.2 of Bierkens and Kappen (2012) is square integrable, and therefore in the set \mathcal{U} .

The semimartingale M_t that solves

$$Z_t = E\left[\frac{dP}{dQ} \middle| \mathcal{F}_t^B\right] = \exp(M_t - [M, M]_t)$$

is unique, up to an evanescence (see theorem 8.3 in Jacod and Shiryaev (2003)). Therefore, it has a version such that

$$M_t = \int_0^t u(X, s) dB_s,$$

and is a square integrable martingale. By the Ito representation theorem, $u(X, t)$ is the only square-integrable process for which this equation is satisfied.

5.1.17 Proof of proposition 1.7

I will proceed in four steps. In the first two steps, I will establish convergence results. In the third and fourth step, I will Taylor-expand the indirect utility function and buyer's security valuation (the two components of the security design utility function). The extra steps in this proof, relative to the static models, arise because of the need to ensure integrability, and to prove that certain limits and integrals can be interchanged.

Define the retained tranche as a function of the Q -Brownian motion, $\hat{\eta}(B) = \eta(X(B))$. By limited liability, $\hat{\eta}(B) \in [0, \beta_s^{-1}X(B)]$, and therefore $E^P[\hat{\eta}(B)^2] < \infty$. It follows that $\hat{\eta}$ is Hida-Malliavin differentiable (Di Nunno et al. (2008)). Define $h_t = \int_0^t u_s ds$. Following Monoyios (2013), the first-order condition for u^* to be optimal (assuming the bounds do not bind) is

$$\psi'(u_t^*) = \theta^{-1} E^Q[D_t \hat{\eta}(B + h^*) | \mathcal{F}_t^B],$$

where $h^* = \int_0^1 u_s^* ds$. If the bounds do bind, so that $|u^*| = \bar{u}$ at some time and state, then

$$|\psi'(u_t^*)| \leq |\theta^{-1} E^Q[D_t \hat{\eta}(B + h^*) | \mathcal{F}_t^B]|$$

By the mean value theorem,

$$\psi''(\hat{u}_t)u_t^* = \theta^{-1}E^Q[D_t\hat{\eta}(B+h^*)|\mathcal{F}_t^B],$$

for some $|\hat{u}_t| \leq |u_t^*|$, if the bounds do not bind, and

$$\psi''(\hat{u}_t)|u_t^*| \leq |\theta^{-1}E^Q[D_t\hat{\eta}(B+h^*)|\mathcal{F}_t^B]|,$$

if they do. Define e_t as

$$e_t = E^Q[D_t\hat{\eta}(B)|\mathcal{F}_t^B].$$

Additionally, define

$$f_t = E^Q[D_t\hat{\eta}(B+h^*)|\mathcal{F}_t^B].$$

Note that e_t does not depend on θ , whereas f_t depends on θ through its dependence on h^* .

The proof proceeds in several steps. First, I will show that u_t^* converges to zero, in the $L^2(Q \times [0, 1])$ sense, meaning that

$$\lim_{\theta^{-1} \rightarrow 0^+} E^Q\left[\int_0^1 (u_t^*)^2 dt\right] = 0.$$

For all $|u| \leq \bar{u}$, $\psi''(u) \geq K_1 > 0$, and therefore

$$(u_t^*)^2 \leq \theta^{-2}K_1^{-2}f_t^2,$$

regardless of whether the bounds on u_t^* bind. By the assumption that $\hat{\eta}(B+h^*)$ is in $L^2(Q)$, for all h , and the Clark-Ocone theorem for $L^2(Q)$ (theorem 6.35 in Di Nunno et al. (2008)), $f_t \in L^2(Q \times [0, 1])$.

By the Ito isometry,

$$E^Q[\int_0^1 f_t^2 dt] = E^Q[(\int_0^1 f_t dB_t)^2].$$

By the Clark-Ocone theorem,

$$\int_0^1 f_t dB_t = \eta(B + h^*) - E^Q[\eta(B + h^*)].$$

Putting these two together,

$$E^Q[\int_0^1 f_t^2 dt] = E^Q[(\eta(B + h^*) - E^Q[\eta(B + h^*)])^2]. \quad (5.2)$$

Because $E^Q[(\eta(B + h))^2] < \infty$ for all feasible u , and the set of feasible u is compact, it follows that

$$\lim_{\theta^{-1} \rightarrow 0^+} E^Q[\int_0^1 f_t^2 dt] < \infty.$$

Using this result, $\lim_{\theta^{-1} \rightarrow 0^+} \theta^{-2} K_1^{-2} E^Q[\int_0^1 f_t^2 dt] = 0$. By the squeeze theorem,

$$\lim_{\theta^{-1} \rightarrow 0^+} E^Q[\int_0^1 (u_t^*)^2 dt] = 0.$$

A similar application of the Ito isometry and Clark-Ocone theorem shows that

$$E^Q[\int_0^1 e_t^2 dt] = E^Q[(\eta(B) - E^Q[\eta(B)])^2], \quad (5.3)$$

which is useful later in the proof.

Next, I will show that $\lim_{\theta^{-1} \rightarrow 0^+} \theta u_t^* = \lim_{\theta^{-1} \rightarrow 0^+} f_t = e_t$, with convergence in $L^2(Q \times [0, 1])$.

First, note that

$$\lim_{\theta^{-1} \rightarrow 0^+} \theta u_t^* = \lim_{\theta^{-1} \rightarrow 0^+} \frac{1}{\psi''(\hat{u}_t)} f_t.$$

Because $|\hat{u}_t| \leq |u_t^*|$, it converges to zero in measure (L^2 convergence implies convergence in measure). Therefore, $\lim_{\theta^{-1} \rightarrow 0^+} \frac{1}{\psi''(\hat{u}_t)} = 1$, with convergence in measure. Because the measure $Q \times \mu([0, 1])$ is finite ($\mu([0, 1])$ is the Lebesgue measure), the product of two sequences that

converge in measure converges in measure to the product of the limits. Therefore, to show that $\lim_{\theta^{-1} \rightarrow 0^+} \theta u_t^* = e_t$, with convergence in measure, it is sufficient to show that $\lim_{\theta^{-1} \rightarrow 0^+} f_t = e_t$, with convergence in measure.

Therefore, if

$$\lim_{\theta^{-1} \rightarrow 0^+} E^Q[\eta(B+h^*)^2] = E^Q[\eta(B)^2],$$

it would follow that f_t converges to e_t in the $L^2(Q \times [0, 1])$ sense.

Define

$$Z_1 = \exp\left(\int_0^1 u_t^* dB_t - \frac{1}{2} \int_0^1 (u_t^*)^2 dt\right).$$

By Girsanov's theorem,

$$\begin{aligned} E^Q[\eta(B+h^*)^2] &= E^Q[Z_1 \eta(B)^2] \\ &= E^Q[\eta(B)^2] + E^Q[(Z_1 - 1)\eta(B)^2]. \end{aligned}$$

By the Cauchy-Schwarz inequality,

$$E^Q[\eta(B+h^*)^2] = E^Q[\eta(B)^2] + E^Q[(Z_1 - 1)^2]^{0.5} E^Q[\eta(B)^4]^{0.5}.$$

By limited liability, $E^Q[\eta(B)^4] \leq E^Q[X_1(B)^4]$, and by assumption is finite. By construction, $E^Q[Z_1] = 1$.

$$E^Q[Z_1^2] = E^Q\left[\exp\left(2 \int_0^1 u_t^* dB_t - \int_0^1 (u_t^*)^2 dt\right)\right].$$

Therefore,

$$\begin{aligned}
E^{\mathcal{Q}}[Z_1^2] &\leq E^{\mathcal{Q}}[\exp(2 \int_0^1 u_t^* dB_t)] \\
&\leq E^{\mathcal{Q}}[\exp(2 \int_0^1 (u_t^*)^2 dt)].
\end{aligned}$$

Using the inequality that $(u_t^*)^2 \leq K_1^2 \theta^{-2} f_t^2$,

$$\begin{aligned}
E^{\mathcal{Q}}[(Z_1 - 1)^2] &\leq E^{\mathcal{Q}}[\exp(2 \int_0^1 K_1^{-2} \theta^{-2} f_t^2 dt)] - 1 \\
&\leq E^{\mathcal{Q}}[\exp(2 \int_0^1 K_1^{-1} \theta^{-1} f_t dB_t)] - 1 \\
&\leq E^{\mathcal{Q}}[\exp(2K_1^{-1} \theta^{-1} (\eta(B + h^*) - E^{\mathcal{Q}}[\eta(B + h^*)]))] - 1 \\
&\leq E^{\mathcal{Q}}[\exp(2K_1^{-1} \theta^{-1} X_1(B + h^*))] - 1 \\
&\leq E^{\mathcal{Q}}[\exp(2K_1^{-1} \theta^{-1} X_1(B + \bar{u}))] - 1.
\end{aligned}$$

The first inequality follows from the inequality that $(u_t^*)^2 \leq K_1^2 \theta^{-2} f_t^2$. The second follows from the expectation of the stochastic exponential. The third applies the Clark-Ocone theorem. The fourth follows from limited liability. The fifth follows from the monotonicity of the asset value in effort.

By the assumption that, for all effort strategies, X_1 is square-integrable, it follows that the moment-generating function exists in some neighborhood around zero. Therefore,

$$\lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}}[\exp(\sqrt{2} K_1^{-1} \theta^{-1} X_1(B + \bar{u}))] = 1.$$

It follows that

$$\lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}}[(Z_1 - 1)^2] = 0. \tag{5.4}$$

I have shown that f_t converges to e_t , in the sense of $L^2(Q \times [0, 1])$ convergence. It follows that f_t and θu_t^* converge in measure to e_t . Moreover, $\theta^2(u_t^*)^2$ and f_t^2 converge to e_t^2 in measure, which will be useful below. Additionally, Z_1 converges in measure to 1.

The third step is to Taylor-expand the indirect utility function from the moral hazard problem, in terms of θ^{-1} . To start, consider the first-order term:

$$\frac{\partial}{\partial \theta^{-1}} \phi_{CT}(\eta; \theta) = \theta^2 E^P \left[\int_0^1 \psi(u_t^*) dt \right].$$

By assumption, $\forall |u| < \bar{u}$, $\psi''(u) \in [K_1, K_2]$ for some positive constants K_1 and K_2 . Therefore, for some $|\tilde{u}_t| < |u_t^*|$,

$$\begin{aligned} \theta^2 \psi(u_t^*) &= \frac{1}{2} \psi''(\tilde{u}_t) \theta^2 (u_t^*)^2 \\ &= \frac{1}{2} \frac{\psi''(\tilde{u}_t)}{(\psi''(\hat{u}_t))^2} f_t^2 \\ &\leq K f_t^2, \end{aligned}$$

for some finite positive constant $K = \frac{1}{2} \frac{K_2}{(K_1)^2}$. Therefore,

$$E^P \left[\int_0^1 \theta^2 \psi(u_t^*) dt \right] \leq E^Q \left[\int_0^1 K f_t^2 dt \right] + E^Q \left[(Z_1 - 1) \int_0^1 K f_t^2 dt \right].$$

To demonstrate that the second term on the right-hand side of this equation converges to zero, I prove the following lemma. The purpose of this lemma is only to establish that the second term converges to zero.

Lemma 5.2. *With $Z_1 - 1$ and f_t defined as above,*

$$E^Q \left[(Z_1 - 1) \int_0^1 f_t^2 dt \right] = \frac{1}{3} E^Q \left[(Z_1 - 1) (\eta(B + h^*) - E^Q[\eta(B + h^*)])^2 \right].$$

Proof. See appendix, section 5.1.18. □

Using this lemma and the Cauchy-Schwarz inequality,

$$E^Q[(Z_1 - 1)(\eta(B+h^*) - E^Q[\eta(B+h^*)])^2] \leq E^Q[(Z_1 - 1)^2]^{0.5} E^Q[(\eta(B+h^*) - E^Q[\eta(B+h^*)])^4]^{0.5}.$$

Therefore, by the assumption that $E^Q[\eta(B+h^*)^4] < \infty$, and equation (5.4),

$$\lim_{\theta^{-1} \rightarrow 0^+} E^Q[(Z_1 - 1) \int_0^1 f_t^2 dt] = 0.$$

Using equation (5.2) and Girsanov's theorem,

$$\begin{aligned} E^Q\left[\int_0^1 f_t^2 dt\right] &= E^Q[(\hat{\eta}(B+h^*) - E^Q[\hat{\eta}(B+h^*)])^2]. \\ &= E^Q[Z_1(\hat{\eta}(B) - E^Q[\hat{\eta}(B)])^2] \\ &= E^Q[(\hat{\eta}(B) - E^Q[\hat{\eta}(B)])^2] + \\ &\quad E^Q[(Z_1 - 1)(\hat{\eta}(B) - E^Q[\hat{\eta}(B)])^2]. \end{aligned}$$

By the Cauchy-Schwarz inequality and $\lim_{\theta^{-1} \rightarrow 0^+} E^Q[(Z_1 - 1)^2] = 0$,

$$\begin{aligned} \lim_{\theta^{-1} \rightarrow 0^+} E^Q[Z_1 \int_0^1 f_t^2 dt] &= \lim_{\theta^{-1} \rightarrow 0^+} E^Q\left[\int_0^1 f_t^2 dt\right] \\ &= E^Q[(\hat{\eta}(B) - E^Q[\hat{\eta}(B)])^2] \\ &= E^Q\left[\int_0^1 e_t^2 dt\right]. \end{aligned}$$

The second step follows from equation (5.3).

Consider the sample space $\Omega \times [0, 1]$, with the standard tensor-product sigma algebra and product measure $Q \times \mu([0, 1])$. The above result establishes convergence in L^1 of $Z_1 f_t^2$ to e_t^2 , and there-

fore convergence in measure. By theorem 5 in chapter 11.6 of Shiryaev (1996), $Z_1 f_t^2$ is uniformly integrable. I will next argue that $Z_1 \theta^2 \psi(u_t^*)$ is uniformly integrable.

By lemma 2 in chapter 11.6 of Shiryaev (1996), a necessary and sufficient condition for uniform integrability (which some authors use as the definition of uniform integrability) is that, for some random variable x_t that depends on θ , and for all θ ,

$$E^Q[\int_0^1 |x_t| dt] < \infty,$$

and, for all $\varepsilon > 0$, there exists a $\delta > 0$ such that, for any set $A \subseteq \Omega \times [0, 1]$ with $E^Q[\int_0^1 \mathbf{1}(A) dt] \leq \delta$, and all θ ,

$$E^Q[\int_0^1 |x_t| \mathbf{1}(A) dt] \leq \varepsilon.$$

By the inequality that $\psi(u) \in [0, K f_t^2]$, it follows that

$$\frac{1}{K} E^Q[\int_0^1 |Z_1 \theta^2 \psi(u_t^*)| dt] \leq E^Q[\int_0^1 |Z_1 f_t^2| dt] < \infty,$$

and

$$\frac{1}{K} E^Q[\int_0^1 |Z_1 \theta^2 \psi(u_t^*)| \mathbf{1}(A) dt] \leq E^Q[\int_0^1 |Z_1 f_t^2| \mathbf{1}(A) dt] \leq \varepsilon.$$

Therefore $\frac{1}{K} Z_1 \theta^2 \psi(u_t^*)$ is also uniformly integrable. It follows that $Z_1 \theta^2 \psi(u_t^*)$ is uniformly integrable. Taylor-expanding around $u_t^* = 0$, to second order,

$$Z_1 \theta^2 \psi(u_t^*) = \frac{1}{2} Z_1 \psi''(\tilde{u}_t) (\theta u_t^*)^2,$$

for some $|\tilde{u}_t| \leq u_t^*$. By the convergence in measure of Z_1 to 1, u_t^* (and therefore \tilde{u}_t) to zero, and θu_t^* to e_t , $Z_1 \theta^2 \psi(u_t^*)$ converges in measure to $\frac{1}{2} e_t^2$. Because $Z_1 \theta^2 \psi(u_t^*)$ converges in measure and

is uniformly integrable,

$$\lim_{\theta^{-1} \rightarrow 0^+} E^Q[Z_1 \int_0^1 \theta^2 \psi(u_t^*) dt] = E^Q[\int_0^1 e_t^2 dt] = \frac{1}{2} V^Q[\hat{\eta}(B)].$$

I have shown that the first-order term in the Taylor expansion converges to the variance. The zero-order term converges to the expected value of the retained tranche. Therefore, the Taylor expansion around the limit $\theta^{-1} \rightarrow 0^+$ is

$$\phi_{CT}(\eta; \theta) = E^Q[\hat{\eta}] + \theta^{-1} \frac{1}{2} V^Q[\hat{\eta}(B)] + O(\theta^{-1}).$$

The fourth step is to consider a Taylor expansion of the buyer's security valuation, again around $\theta^{-1} \rightarrow 0^+$. Define $\hat{s}(B) = s(X(B))$. The buyer's expected value is

$$E^Q[Z_1 \hat{s}(B)] = E^Q[\hat{s}(B)] + E^Q[(Z_1 - 1)\hat{s}(B)].$$

I will show that

$$\lim_{\theta^{-1} \rightarrow 0^+} E^Q[\theta(Z_1 - 1)\hat{s}(B)] = E^Q[(\hat{\eta}(B) - E^Q[\hat{\eta}(B)])\hat{s}(B)].$$

Using the stochastic logarithm,

$$Z_1 - 1 = \int_0^1 Z_t u_t^* dB_t,$$

where $Z_t u_t^*$ is in $L^2(\Omega \times [0, 1])$. By definition, $Z_t = E^Q[Z_1 | \mathcal{F}_t]$. By equation (5.4), it follows that Z_t converges in measure to 1, and therefore that $\theta Z_t u_t^*$ converges in measure to e_t .

Because $\hat{s}(B) \leq X_1(B)$ is in $L^2(\Omega)$, we can define

$$r_t = E^Q[D_t \hat{s}(B) | \mathcal{F}_t^B]$$

and see that it is in $L^2(\Omega \times [0, 1])$. By the Clark-Ocone theorem,

$$E^{\mathcal{Q}}[\theta(Z_1 - 1)(\hat{s}(B) - E^{\mathcal{Q}}[\hat{s}(B)])] = E^{\mathcal{Q}}[(\int_0^1 \theta Z_t u_t^* dB_t)(\int_0^1 r_t dB_t)].$$

Using the Ito isometry,

$$E^{\mathcal{Q}}[(\int_0^1 \theta Z_t u_t^* dB_t)(\int_0^1 r_t dB_t)] = E^{\mathcal{Q}}[\theta \int_0^1 Z_t u_t^* r_t dt].$$

By construction, $\theta Z_t u_t^*$ is square-integrable for all θ . Because $\theta Z_t u_t^*$ converges in measure, it converges weakly on the space $L^2(\Omega \times [0, 1])$, which is a Hilbert space, and therefore

$$\lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}}[\theta \int_0^1 Z_t u_t^* r_t dt] = E^{\mathcal{Q}}[\int_0^1 e_t r_t dt].$$

Reversing the use of the Ito isometry and Clark-Ocone theorem,

$$\lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}}[\theta(Z_1 - 1)(\hat{s}(B) - E^{\mathcal{Q}}[\hat{s}(B)])] = E^{\mathcal{Q}}[(\hat{\eta}(B) - E^{\mathcal{Q}}[\hat{\eta}(B)])(\hat{s}(B) - E^{\mathcal{Q}}[\hat{s}(B)])].$$

The term

$$\begin{aligned} \lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}}[\theta(Z_1 - 1)]E^{\mathcal{Q}}[\hat{s}(B)] &= \lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}}[\int \theta Z_t u_t^* dB_t]E^{\mathcal{Q}}[\hat{s}(B)] \\ &= 0. \end{aligned}$$

Therefore,

$$\lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}}[\theta(Z_1 - 1)\hat{s}(B)] = E^{\mathcal{Q}}[(\hat{\eta}(B) - E^{\mathcal{Q}}[\hat{\eta}(B)])\hat{s}(B)].$$

The zero-order term in the buyer's security valuation converges to

$$\lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}}[Z_1 \hat{s}(B)] = E^{\mathcal{Q}}[\hat{s}(B)].$$

The first-order term converges to

$$\lim_{\theta^{-1} \rightarrow 0^+} \frac{\partial}{\partial(\theta^{-1})} E^{\mathcal{Q}} \left[\frac{\partial Z_1}{\partial \theta} \hat{s}(B) \right] = \lim_{\theta^{-1} \rightarrow 0^+} E^{\mathcal{Q}} [\theta(Z_1 - 1) \hat{s}(B)].$$

Combining the results in the third and fourth steps, and also computing the first-order expansion of the buyer's security valuation in terms of κ , it follows that

$$\begin{aligned} U(s; \theta^{-1}, \kappa) &= (\beta_s + \kappa \beta_s) E^{\mathcal{Q}}[\hat{s}(B)] + \theta^{-1} \beta_s E^{\mathcal{Q}}[\hat{s}(B) \hat{\eta}(B)] - \\ &\quad \theta^{-1} \beta_s E^{\mathcal{Q}}[\hat{s}(B)] E^{\mathcal{Q}}[\hat{\eta}(B)] + E^{\mathcal{Q}}[\hat{\eta}(B)] + \\ &\quad \frac{1}{2} \theta^{-1} V^{\mathcal{Q}}[\hat{\eta}(B)] + O(\theta^{-2} + \theta^{-1} \kappa). \end{aligned}$$

This equation is equivalent to the first-order terms of equation (5.1) in section 5.1.9. The remainder of the proof is identical to the algebra in that section.

5.1.18 Proof of lemma 5.2

The purpose of this lemma is to derive an alternate form of a particular expression, which is shown to go to zero in the above proof. To accomplish this, I use a result regarding the relationship between Malliavin derivatives, Skorohod integrals, and cumulants, from Privault (2013). I specialize that paper's results to a simple case, involving adapted, square integrable processes that are themselves Malliavin derivatives. The proof relies heavily on the theorems of Malliavin calculus described in Di Nunno et al. (2008).

Using the martingale representation theorem, we can define

$$Z_1 - 1 = \int_0^1 z_t dB_t,$$

where z_t is a square-integrable, \mathcal{F}_t^B -adapted process. Note that f_t also has these properties. It follows that both z_t and f_t are themselves Hida-Malliavin differentiable.

Using lemma 4.2 of Privault (2013),

$$\begin{aligned}
E^{\mathcal{Q}}[(Z_1 - 1)(\int_0^1 f_t dB_t)^2] &= E^{\mathcal{Q}}[(Z_1 - 1) \int_0^1 f_t^2 dt] + E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t(Z_1 - 1))(D_s f_t) f_s ds dt] + \\
&E^{\mathcal{Q}}[(\int_0^1 f_t dB_t) \int_0^1 f_t (D_t(Z_1 - 1)) dt].
\end{aligned}$$

Using theorem 3.18 of Di Nunno et al. (2008), the fundamental theorem of Malliavin calculus,

$$\begin{aligned}
D_t(Z_1 - 1) &= D_t(\int_0^1 z_r dB_r) \\
&= \int_0^1 (D_t z_r) dB_r + z_t.
\end{aligned}$$

Therefore,

$$\begin{aligned}
E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t(Z_1 - 1))(D_s f_t) f_s ds dt] &= E^{\mathcal{Q}}[\int_0^1 \int_0^1 (\int_0^1 (D_t z_r) dB_r)(D_s f_t) f_s ds dt] + \\
&E^{\mathcal{Q}}[\int_0^1 \int_0^1 z_t (D_s f_t) f_s ds dt].
\end{aligned}$$

Analyzing the first term of this expression, using the integration by parts formula (theorem 6.15 of Di Nunno et al. (2008)), along with the Clark-Ocone theorem,

$$\begin{aligned}
E^{\mathcal{Q}}[\int_0^1 \int_0^1 (\int_0^1 (D_t z_r) dB_r)(D_s f_t) f_s ds dt] &= E^{\mathcal{Q}}[\int_0^1 (\int_0^1 (D_t z_r) dB_r) f_t (\hat{\eta}(B + h^*) - \bar{\eta}) dt] - \\
&E^{\mathcal{Q}}[\int_0^1 (\int_0^1 (D_t z_r) dB_r) \int_0^1 f_t f_s dB_s dt].
\end{aligned}$$

To minimize notation, I have used $\bar{\eta} = E^{\mathcal{Q}}[\hat{\eta}(B + h^*)]$ in the previous expression. Analyzing the first term of this expression,

$$\begin{aligned}
E^{\mathcal{Q}}\left[\int_0^1 \left(\int_0^1 (D_t z_r) dB_r\right) f_t(\hat{\eta}(B+h^*) - \bar{\eta}) dt\right] &= E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 (D_t z_r) D_r(f_t(\hat{\eta}(B+h^*) - \bar{\eta})) dr dt\right] \\
&= E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 (D_t z_r) f_t D_r(\hat{\eta}(B+h^*) - \bar{\eta}) dr dt\right] \\
&= E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 (D_t z_r) f_t f_r dt dr\right] - \\
&\quad E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 (D_t z_r) f_t \left(\int D_r f_s dB_s\right) dr dt\right] \\
&= E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 (D_t z_r) f_t f_r dt dr\right] - \\
&\quad E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 \int_0^1 (D_r f_s)(D_t z_r)(D_s f_t) ds dr dt\right] \\
&= E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 (D_t z_r) f_t f_r dt dr\right].
\end{aligned}$$

The first step follows from integration by parts, the second by the fact (for adapted processes z_r and f_t) that $(D_t z_r)(D_r f_t)$ integrates to zero (see remark 6.18 of Di Nunno et al. (2008)). The third step follows from the fundamental theorem, and the fourth from integration by parts. The fifth step applies the same step about adapted processes.

By the Ito isometry,

$$E^{\mathcal{Q}}\left[\int_0^1 \left(\int_0^1 (D_t z_r) dB_r\right) \int_0^1 f_t f_s dB_s dt\right] = E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 (D_t z_r) f_t f_r dt dr\right].$$

Putting these last two results together,

$$E^{\mathcal{Q}}\left[\int_0^1 \int_0^1 \left(\int_0^1 (D_t z_r) dB_r\right) (D_s f_t) f_s ds dt\right] = 0.$$

Next, consider the term

$$\begin{aligned}
E^{\mathcal{Q}}[(\int_0^1 f_t dB_t) \int_0^1 f_t (D_t(Z_1 - 1)) dt] &= E^{\mathcal{Q}}[(\int_0^1 f_t dB_t) \int_0^1 f_t (\int_0^1 (D_t z_s) dB_s) dt] + \\
&E^{\mathcal{Q}}[(\int_0^1 f_t dB_t) \int_0^1 f_t z_t dt],
\end{aligned}$$

where I have used integration by parts. Considering the first term in this expression,

$$\begin{aligned}
E^{\mathcal{Q}}[(\int_0^1 f_t dB_t) \int_0^1 f_t (\int_0^1 (D_t z_s) dB_s) dt] &= \int_0^1 E^{\mathcal{Q}}[(\int_0^1 (D_t z_s) dB_s) (f_t (\hat{\eta}(B + h^*) - \bar{\eta}))] dt \\
&= E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t z_r) D_r (f_t (\hat{\eta}(B + h^*) - \bar{\eta})) dr dt] \\
&= E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t z_r) f_t f_r dt].
\end{aligned}$$

The first step is Fubini's theorem and the Clark-Ocone theorem, the second is integration by parts, and the third step reuses the algebra from above.

Putting it together, and using the Clark-Ocone theorem on the last term,

$$\begin{aligned}
E^{\mathcal{Q}}[(Z_1 - 1) (\int_0^1 f_t dB_t)^2] &= E^{\mathcal{Q}}[(Z_1 - 1) \int_0^1 f_t^2 dt] + E^{\mathcal{Q}}[\int_0^1 \int_0^1 z_t (D_r f_t) f_r dr dt] + \\
&E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_r z_t) f_t f_r dr dt] + E^{\mathcal{Q}}[(\hat{\eta}(B + h^*) - \bar{\eta}) \int_0^1 f_t z_t dt].
\end{aligned}$$

Using the chain rule of Malliavin calculus,

$$\begin{aligned}
E^{\mathcal{Q}}[(Z_1 - 1) (\int_0^1 f_t dB_t)^2] &= E^{\mathcal{Q}}[(Z_1 - 1) \int_0^1 f_t^2 dt] + E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_r(z_t f_t)) f_r dr dt] + \\
&E^{\mathcal{Q}}[(\hat{\eta}(B + h^*) - \bar{\eta}) \int_0^1 f_t z_t dt].
\end{aligned}$$

Integrating by parts,

$$E^{\mathcal{Q}}[(Z_1 - 1)(\int_0^1 f_t dB_t)^2] = E[(Z_1 - 1) \int_0^1 f_t^2 dt] + 2E^{\mathcal{Q}}[(\hat{\eta}(B + h^*) - \bar{\eta}) \int_0^1 f_t z_t dt].$$

Finally, I figure out what the second term in this expression is.

$$\begin{aligned} \frac{1}{2}E^{\mathcal{Q}}[(\hat{\eta}(B + h^*) - \bar{\eta})^2(Z_1 - 1)] &= E^{\mathcal{Q}}[\int_0^1 \frac{1}{2}z_t D_t[(\int_0^1 f_s dB_s)^2]] \\ &= E^{\mathcal{Q}}[(\hat{\eta}(B + h^*) - \bar{\eta}) \int_0^1 f_t z_t dt] + \\ &E^{\mathcal{Q}}[(\eta(B + h^*) - \bar{\eta}) \int_0^1 z_t (\int_0^1 (D_t f_s) dB_s) dt]. \end{aligned}$$

The first step uses integration by parts and the Clark-Ocone theorem. The second uses the chain rule, Clark-Ocone theorem, and then fundamental theorem. Considering the second term,

$$\begin{aligned} E^{\mathcal{Q}}[(\hat{\eta}(B + h^*) - \bar{\eta}) \int_0^1 z_t (\int_0^1 (D_t f_s) dB_s) dt] &= E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t f_s) D_s(z_t (\hat{\eta}(B + h^*) - \bar{\eta})) ds dt] \\ &= E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t f_s) z_t D_s((\hat{\eta}(B + h^*) - \bar{\eta})) ds dt] \\ &= E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t f_s) z_t f_s dt ds] - \\ &E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t f_s) z_t (\int_0^1 D_s f_r dB_r) ds dt] \\ &= \frac{1}{2}E^{\mathcal{Q}}[\int_0^1 \int_0^1 (D_t f_s^2) z_t dt ds] - \\ &E^{\mathcal{Q}}[\int_0^1 \int_0^1 \int_0^1 (D_t f_s)(D_r z_t)(D_s f_r) ds dr dt] \\ &= \frac{1}{2}E^{\mathcal{Q}}[(Z_1 - 1) \int_0^1 f_s^2 ds]. \end{aligned}$$

The first step is integration by parts, and the second applies the previously mentioned fact about adapted processes. The third uses the fundamental theorem, and the fourth uses both the chain rule (in the first term) and integration by parts (in the second term). The last step uses the fact about adapted processes and integration by parts, along with the Clark-Ocone theorem. It follows that

$$E^Q[(\hat{\eta}(B+h^*) - \bar{\eta})^2(Z_1 - 1)] = 3E^Q[(\hat{\eta}(B) - \bar{\eta}) \int_0^1 f_t z_t dt].$$

Plugging in this into the earlier equation,

$$E^Q[(Z_1 - 1) \left(\int_0^1 f_t dB_t \right)^2] = E[(Z_1 - 1) \int_0^1 f_t^2 dt] + \frac{2}{3} E^Q[(\hat{\eta}(B+h^*) - \bar{\eta})^2(Z_1 - 1)].$$

Using the Clark-Ocone theorem,

$$E^Q[(\hat{\eta}(B+h^*) - \bar{\eta})^2(Z_1 - 1)] = 3E^Q[(Z_1 - 1) \int_0^1 f_t^2 dt],$$

which proves the lemma.

5.2 Appendix: Generalized Rational Inattention

5.2.1 Proof of theorem 2.1

Parts 1 and 2 of the theorem follow from a Taylor expansion of the cost function. Using the lemmas and theorem of Chentsov (1982), cited in the text, we know that for any invariant cost function with continuous second derivatives,

$$C(\{p(s|x)\}) = \frac{1}{2} \Delta \sum_{x' \in X} \sum_{x \in X} k(x, x') < \tau(s|x') |g(r(s))| \tau(s|x) > + O(\Delta^{1.5}).$$

The second claim follows by a similar argument.

We next demonstrate the claimed properties of $k(x, x')$. First, $k(x, x')$ is symmetric, by the symmetry of partial derivatives and the assumption of continuous second derivatives (condition 4). Recall the assumption that

$$p(s|x) = r(s) + \Delta^{0.5} \tau(s|x) + O(\Delta),$$

which implies that $\sum_{s \in S} r(s) = 1$ and $\sum_{s \in S} \tau(s|x) = 0$ for all $x \in X$. Consider a signal structure, for which $\tau(s|x) = \phi(s)v(x)$, with $\sum_{s \in S} \phi(s) = 0$. For this signal structure,

$$C(\{p(s|x)\}) = \frac{1}{2} \Delta \bar{g} \sum_{x' \in X} \sum_{x \in X} v(x)v(x')k(x, x') + O(\Delta^{1.5}),$$

where $\langle \phi(s)|g(r(s))|\phi(s) \rangle = \bar{g} > 0$ is the Fisher information of $\phi(s)$. The posteriors associated with this signal structure are

$$p(x|s) = \frac{p(s|x)q(x)}{\sum_{x' \in X} q(x')p(s|x')} = q(x) \frac{r(s) + \Delta^{0.5} \phi(s)v(x)}{r(s) + \Delta^{0.5} \phi(s) \sum_{x' \in X} q(x')v(x')}.$$

If $v(x)$ is constant, then $p(x|s) = q(x)$ for all $s \in S$, the signal structure is uninformative, and $C(\{p(s|x)\}) = 0$ by 1. In this case,

$$\sum_{x' \in X} \sum_{x \in X} k(x, x') = 0.$$

If $v(x)$ is not constant, then $p(x|s)$ is not equal to $q(x)$ for all s . In this case, the signal structure is informative, and $C(\{p(s|x)\}) > 0$. Therefore,

$$\sum_{x' \in X} \sum_{x \in X} v(x)v(x')k(x, x') > 0.$$

It follows that $k(x, x)$ is positive semi-definite. It has a single eigenvector (a vector of constants) in its null space, and all of its other eigenvectors are associated with positive eigenvalues.

5.2.2 Proof of lemma 2.1

We begin by considering the optimal policy $p_{\Delta}^*(a|x)$. The problem can be written as

$$\begin{aligned} \max_{\{p_{\Delta}(a|x) \in \mathbb{R}_{|A|}^+, \kappa(x), v(a|x)\}} \Delta^{0.5} \sum_{x \in X} q(x) \sum_{a \in A} p(a|x) u(x, a) - C(\{p(a|x)\}) \\ + \sum_{x \in X} q(x) \kappa(x) (1 - \sum_{a \in A} p(a|x)) \\ + \sum_{x \in X} q(x) \sum_{a \in A} p(a|x) v(a|x) \end{aligned}$$

where $\kappa(x)$ is the multiplier on the constraint that $\sum_{a \in A} p_{\Delta}^*(a|x) = 1$ and $v(a|x)$ is the multiplier on the constraint that $p_{\Delta}^*(a|x) \geq 0$.

The optimal policy must achieve a weakly higher utility than any policy involving an uninformative signal structure. Define a probability distribution

$$\phi^*(a) = \arg \max_{\phi(a) \in \mathcal{P}(A)} \sum_{x \in X} q(x) \sum_{a \in A} \phi(a) u(x, a)$$

that is an optimal policy, conditional on gathering no information. It must be the case that

$$\Delta^{0.5} \sum_{x \in X} q(x) \sum_{a \in A} p_{\Delta}^*(a|x) u(x, a) - C(\{p_{\Delta}^*(a|x)\}) \geq \Delta^{0.5} \sum_{x \in X} q(x) \sum_{a \in A} \phi^*(a) u(x, a).$$

Rearranging this,

$$\Delta^{0.5} \sum_{x \in X} q(x) \sum_{a \in A} [p_{\Delta}^*(a|x) - \phi^*(a)] u(x, a) \geq C(\{p_{\Delta}^*(a|x)\}).$$

Let $\bar{u}(x) = \max_{a \in A} u(x, a)$ and let $\underline{u}(x) = \min_{a \in A} u(x, a)$. It follows that

$$\Delta^{0.5} \sum_{x \in X} q(x) (\bar{u}(x) - \underline{u}(x)) \geq C(\{p_{\Delta}^*(a|x)\}).$$

From this, we see that $\lim_{\Delta \rightarrow 0^+} C(\{p_{\Delta}^*(a|x)\}) = 0$. By condition 1, it must be the case that

$$\lim_{\Delta \rightarrow 0^+} \min_{r \in \mathcal{P}(A)} \|p_{\Delta}^*(a|x) - r(a)\| = 0,$$

under the arbitrary norm $\|\cdot\|$. Moreover, it follows that

$$\sum_{x \in X} q(x) \sum_{a \in A} [\lim_{\Delta \rightarrow 0^+} \{p_{\Delta}^*(a|x)\} - \phi^*(a)] u(x, a) \geq 0.$$

Defining the set of ex-ante optimal actions as $A_+ \subseteq A$, it follows that

$$\lim_{\Delta \rightarrow 0^+} \min_{r \in \mathcal{P}(A_+)} \|p_{\Delta}^*(a|x) - r(a)\| = 0.$$

5.2.3 Proof of theorem 2.2

We begin by stating the following lemma, which is useful in the proof.

Lemma 5.3. *Define $r_{\Delta}^*(a) = \sum_{x \in X} q(x) p_{\Delta}^*(a|x)$. For all sub-sequences Δ_n such that $p_n^*(a|x) = p_{\Delta_n}^*(a|x)$ converges to $r^*(a) \in \mathcal{P}(A_+)$, and all $\alpha \in [0, \frac{1}{2})$,*

$$\lim_{\Delta \rightarrow 0^+} \Delta^{\alpha} \|p_n^*(a|x) - r_n^*(a)\| = 0.$$

Proof. See the appendix, section 5.2.4. □

Let $\Delta_m, m = \{1, 2, \dots\}$ denote a sequence such that $\lim_{m \rightarrow \infty} \Delta_m = 0$, and let $p_m^*(a|x) = p_{\Delta_m}^*(a|x)$ be the corresponding sequence of optimal policies. By the Bolzano-Weierstrauss theorem and the boundedness of $p_m^*(a|x)$, every sequence has a convergent sub-sequence. Denote that sequence $p_n^*(a|x)$, with associated time interval Δ_n . By lemma 5.3,

$$p_n^*(a|x) = r^*(a) + \phi_n^*(a) + \Delta_n^{\frac{1}{2}} \tau(a|x) + o(\Delta_n^{\frac{1}{2}}), \quad (5.5)$$

where $\sum_{x \in X} \tau(a|x) q(x) = 0$, $\sum_{a \in A} \tau(a|x) = 0$, $\sum_{a \in A} \phi_n^*(a) = 0$, and $\lim_{n \rightarrow \infty} \phi_n^*(a) = 0$.

$$\begin{aligned}
\max_{\{p_{\Delta}(a|x) \in \mathbb{R}_{|A|}^+\}, \kappa_n(x), \nu_n(a|x)} \quad & \Delta_n^{0.5} \sum_{x \in X} q(x) \sum_{a \in A} p(a|x) u(x, a) - C(\{p(a|x)\}) \\
& + \Delta_n^{0.5} \sum_{x \in X} q(x) \kappa_n(x) (1 - \sum_{a \in A} p(a|x)) \\
& + \Delta_n^{0.5} \sum_{x \in X} q(x) \sum_{a \in A} p(a|x) \nu_n(a|x)
\end{aligned}$$

The first order condition is

$$\frac{\partial C(\{p_n^*(a'|x')\})}{\partial p(a|x)} = q(x) \Delta_n^{\frac{1}{2}} [u(x, a) - \kappa_n(x) + \nu_n(a|x)].$$

Using the mean value theorem, for some $c \in [0, 1]$,

$$\frac{\partial C(\{p_n^*(a'|x')\})}{\partial p(a|x)} = \sum_{x' \in X} \sum_{a' \in A} [p_n^*(a'|x') - r^*(a')] \frac{\partial^2 C(\{cr^*(a'') + (1-c)p_n^*(a''|x'')\})}{\partial p(a|x) \partial p(a'|x')}.$$

Using equation (5.5) and the first order condition,

$$\begin{aligned}
\sum_{x' \in X} \sum_{a' \in A} [\phi_n^*(a') + \Delta_n^{\frac{1}{2}} \tau(a'|x') + o(\Delta_n^{\frac{1}{2}})] \frac{\partial^2 C(\{cr^*(a'') + (1-c)p_n^*(a''|x'')\})}{\partial p(a|x) \partial p(a'|x')} \\
= q(x) \Delta_n^{\frac{1}{2}} [u(x, a) - \kappa_n(x) + \nu_n(a|x)].
\end{aligned}$$

By the continuity of the second derivatives of $C(\cdot)$ (condition 4) and the convergence of $p_n^*(a''|x'')$ to $r^*(a)$, we have

$$\lim_{n \rightarrow \infty} \frac{\partial^2 C(\{cr^*(a'') + (1-c)p_n^*(a''|x'')\})}{\partial p(a|x) \partial p(a'|x')} = k(x, x') g_{a, a'}(r^*(a'')),$$

where $g_{a, a'}(r^*(a''))$ is the Fisher information matrix evaluated at r^* . Considering the limit,

$$\begin{aligned} \lim_{n \rightarrow \infty} \Delta_n^{-\frac{1}{2}} \sum_{x' \in X} \sum_{a' \in A} [\phi_n^*(a') + \Delta_n^{\frac{1}{2}} \tau(a'|x) + o(\Delta_n^{\frac{1}{2}})] \frac{\partial^2 C(\{cr^*(a'') + (1-c)p_n^*(a''|x'')\})}{\partial p(a|x) \partial p(a'|x')} = \\ \lim_{n \rightarrow \infty} \Delta_n^{-\frac{1}{2}} \sum_{x' \in X} \sum_{a' \in A} \phi_n^*(a') k(x, x') g_{a, a'}(r^*(a'')) + \\ \sum_{x' \in X} \sum_{a' \in A} \tau(a'|x') k(x, x') g_{a, a'}(r^*(a'')). \end{aligned}$$

By the fact that $\sum_{x' \in X} k(x, x') = 0$, we have

$$\sum_{x' \in X} \sum_{a' \in A} \tau(a'|x') k(x, x') g_{a, a'}(r^*(a'')) = \lim_{n \rightarrow \infty} q(x) [u(x, a) - \kappa_n(x) + v_n(a|x)].$$

It follows that $\tau(a|x) = 0$ if $r^*(a) = 0$. Multiplying the equation above by the inverse Fisher matrix,

$$\sum_{x' \in X} \tau(a|x') k(x, x') = \lim_{n \rightarrow \infty} q(x) r^*(a) [u(x, a) - \kappa_n(x) + v_n(a|x)].$$

By the assumption of convergence, $\lim_{n \rightarrow \infty} p_n^*(a|x) = r^*(a)$. Therefore, for any a such that $r^*(a) > 0$, $v_n(a|x) = 0$. Therefore,

$$\sum_{x' \in X} \tau(a|x') k(x, x') = \lim_{n \rightarrow \infty} q(x) r^*(a) [u(x, a) - \kappa_n(x)].$$

By the condition that $\sum_{a \in A} \tau(a|x) = 0$, we must have

$$\sum_{a \in A} \sum_{x' \in X} \tau(a|x') k(x, x') = 0 = \lim_{n \rightarrow \infty} q(x) \sum_{a \in A} r^*(a) [u(x, a) - \kappa_n(x)].$$

Therefore,

$$\lim_{n \rightarrow \infty} \kappa_n(x) = \sum_{a \in A} r^*(a) u(x, a).$$

We rewrite the FOC as

$$\sum_{x' \in X} \tau(a|x')k(x, x') = q(x)r^*(a)[u(x, a) - \sum_{a' \in A} r^*(a')u(x, a')].$$

The only vectors in the null space of $k(x, x')$ are proportional to $\mathbf{1}$, the vector of ones. It follows that this equation has no solution if

$$\sum_{x \in X} q(x)r^*(a)[u(x, a) - \sum_{a' \in A} r^*(a')u(x, a')] \neq 0.$$

However, $r^* \in \mathcal{P}(A_+)$, and therefore $\sum_{x \in X} q(x)u(x, a) = \sum_{x \in X} q(x)u(x, a')$ for all a, a' such that $r(a) > 0$ and $r(a') > 0$. Therefore, a solution exists.

Let $k^+(x, x')$ denote the pseudo-inverse of $k(x, x')$. Any function of the form

$$\tau(a|x') = \sum_{x \in X} k^+(x, x')q(x)r^*(a)[u(x, a) - \sum_{a' \in A} r^*(a')u(x, a')] + \lambda(a)$$

is a solution. However, by the condition that $\sum_{x \in X} \tau(a|x)q(x) = 0$, we have

$$\sum_{x' \in X} \sum_{x \in X} q(x')k^+(x, x')q(x)r^*(a)[u(x, a) - \sum_{a' \in A} r^*(a')u(x, a')] + \lambda(a) = 0.$$

Therefore,

$$\begin{aligned} \tau(a|x') &= \sum_{x \in X} k^+(x, x')q(x)r^*(a)[u(x, a) - \sum_{a' \in A} r^*(a')u(x, a')] - \\ &\quad \sum_{x' \in X} \sum_{x \in X} q(x')k^+(x, x')q(x)r^*(a)[u(x, a) - \sum_{a' \in A} r^*(a')u(x, a')]. \end{aligned}$$

Using theorem 2.1 to do a Taylor expansion around the policy function $p_n^*(a|x) = r^*(a)$, for some $c \in [0, 1]$,

$$C(\{p_n^*(a''|x'')\}) = \frac{1}{2} \sum_{x \in X} \sum_{a \in A} \sum_{x' \in X} \sum_{a' \in A} [p_n^*(a'|x') - r^*(a')][p_n^*(a|x) - r^*(a)] \cdot \frac{\partial^2 C(\{c r^*(a'') + (1-c)p_n^*(a''|x'')\})}{\partial p(a|x) \partial p(a'|x')}.$$

By similar arguments to the ones used earlier,

$$\lim_{n \rightarrow \infty} \Delta_n^{-1} C(\{p_n^*(a''|x'')\}) = \frac{1}{2} \sum_{x \in X} \sum_{a \in A} \sum_{x' \in X} \sum_{a' \in A} \tau(a|x) \tau(a'|x') k(x, x') g_{a, a'}(r^*(a'')).$$

The terms in $\tau(a|x)$ associated with $\lambda(a)$ become zero, and

$$\lim_{n \rightarrow \infty} \Delta_n^{-1} C(\{p_n^*(a''|x'')\}) = \frac{1}{2} \sum_{a \in A} \sum_{x \in X} \sum_{x' \in X} q(x') k^+(x, x') q(x) r^*(a) \cdot [u(x, a) - \sum_{a' \in A} r^*(a') u(x, a')] [u(x', a) - \sum_{a' \in A} r^*(a') u(x', a')].$$

Returning to the original problem, we have

$$p_n^*(a|x) \in \arg \max_{\{p_n(a|x) \in \mathcal{P}(A)\}} \Delta_n^{\frac{1}{2}} \sum_{x \in X} q(x) \sum_{a \in A} (r^*(a) + \phi_n(a)) u(x, a) + \Delta_n^{\frac{1}{2}} \sum_{x \in X} q(x) \sum_{a \in A} \tau(a|x) u(x, a) - C(\{p(a|x)\}) + o(\Delta_n).$$

Using the limit for the cost function,

$$\begin{aligned}
p_n^*(a|x) \in \arg \max_{\{p_n(a|x) \in \mathcal{P}(A)\}} & \Delta_n^{\frac{1}{2}} \sum_{x \in X} \sum_{a \in A} q(x) r^*(a) u(x, a) + \\
\frac{1}{2} \Delta_n \sum_{a \in A} \sum_{x \in X} \sum_{x' \in X} & q(x) k^+(x', x) q(x') r^*(a) [u(x', a) - \sum_{a' \in A} r^*(a') u(x', a')] u(x, a) - \\
& \Delta_n \sum_{x \in X} q(x) \sum_{a \in A} \lambda(a) u(x, a) + \\
& \Delta_n^{\frac{1}{2}} \sum_{x \in X} q(x) \sum_{a \in A} \phi_n(a) u(x, a) + o(\Delta),
\end{aligned}$$

where the solution for $\lambda(a)$ is defined above. Note that $\lambda(a) > 0$ if and only if $r^*(a) > 0$. The choice of $\phi_n(a)$ is restricted by the requirement that $p_n(a|x) \geq 0$ to be positive if $r^*(a) = 0$. It follows any $\phi(a)$ such that $\phi(a) \neq 0$ if and only if $a \in A_+$ cannot change the total utility. Moreover, the effects of $\lambda(a)$ on the optimal policy are subsumed by $\phi_n(a)$ — it is without loss of generality to assume that $\phi_n(a) = 0$.

$$\begin{aligned}
p_n^*(a|x) \in \arg \max_{\{p_n(a|x) \in \mathcal{P}(A)\}} & \Delta_n^{\frac{1}{2}} \sum_{x \in X} \sum_{a \in A} q(x) r^*(a) u(x, a) + \\
\frac{1}{2} \Delta_n \sum_{a \in A} \sum_{x \in X} \sum_{x' \in X} & q(x) k^+(x', x) q(x') r^*(a) [u(x', a) - \sum_{a' \in A} r^*(a') u(x', a')] u(x, a) \\
& + o(\Delta). \tag{5.6}
\end{aligned}$$

Because $r^*(a) \in \mathcal{P}(A_+)$, the first term is invariant to the choice of $r^*(a)$. Therefore,

$$r^*(a) \in \arg \max_{r(a) \in \mathcal{P}(A_+)} \sum_{a \in A} m(a, a) r(a) - \sum_{a \in A} \sum_{a' \in A} m(a, a') r(a) r(a'),$$

where

$$\begin{aligned} m(a, a') &= \sum_{x \in X} \sum_{x' \in X} q(x) k^+(x', x) [q(x') u(x', a) - \sum_{x'' \in X} q(x'') u(x'', a)] [u(x, a') \sum_{x'' \in X} q(x'') u(x'', a')] \\ &= \sum_{x \in X} \sum_{x' \in X} q(x) k^+(x', x) q(x') u(x', a) u(x, a'). \end{aligned}$$

By the assumption that $u(x, a)$ has full row rank, $m(a, a')$ is positive definite. It follows that $r^*(a)$ is unique.

We have proved that every convergent sub-sequence of $p_m^*(a|x)$ converges to a unique $r^*(a)$. It follows that $p_m^*(a|x)$ converges to $r^*(a)$, and therefore

$$\lim_{\Delta \rightarrow 0^+} p_{\Delta}^*(a|x) = r^*(a).$$

5.2.4 Proof of lemma 5.3

Let $p_n^*(a|x) = p_{\Delta_n}^*(a|x)$ denote the convergent sub-sequence. By definition,

$$\lim_{n \rightarrow 0^+} p_n^*(a|x) = r^*(a) = \lim_{n \rightarrow 0^+} r_n^*(a)$$

for some $r^* \in \mathcal{P}(A_+)$. We can Taylor-expand the cost function around r^* , to find, for some $c \in [0, 1]$, that

$$C(\{p_n^*(a''|x'')\}) = \sum_{x \in X} \sum_{a \in A} \sum_{x' \in X} \sum_{a' \in A} [p_n^*(a'|x') - r^*(a')] [p_n^*(a|x) - r^*(a)] \frac{\partial^2 C(\{c r^*(a'') + (1-c) p_n^*(a''|x'')\})}{\partial p(a|x) \partial p(a'|x')}.$$

By the continuity of the second derivatives of $C(\cdot)$ (condition 4) and the convergence of $p_n^*(a''|x'')$ to $r^*(a)$, we have

$$\lim_{n \rightarrow \infty} \frac{\partial^2 C(\{c r^*(a'') + (1-c) p_n^*(a''|x'')\})}{\partial p(a|x) \partial p(a'|x')} = k(x, x') g_{a, a'}(r^*(a'')),$$

where $g_{a,a'}(r^*(a''))$ is the Fisher information matrix evaluated at r^* .

By the inequality derived earlier, we must have

$$\Delta_n^{-\alpha} \sum_{x \in X} q(x) \sum_{a \in A} [p_n^*(a|x) - r^*(a)] u(x, a) \geq \Delta_n^{-\alpha - \frac{1}{2}} C(\{p_n^*(a|x)\}).$$

Using the Taylor expansion, it follows that

$$\begin{aligned} \lim_{n \rightarrow \infty} \Delta_n^{-\alpha} \sum_{x \in X} q(x) \sum_{a \in A} [p_n^*(a|x) - r^*(a)] u(x, a) \geq \\ \Delta_n^{-\alpha - \frac{1}{2}} \sum_{x \in X} \sum_{a \in A} \sum_{x' \in X} \sum_{a' \in A} k(x, x') [p_n^*(a'|x') - r^*(a')] [p_n^*(a|x) - r^*(a)] g_{a,a'}(r^*(a'')). \end{aligned}$$

Suppose that

$$\lim_{n \rightarrow \infty} \Delta_n^{-\alpha} (p_n^*(a|x) - r^*(a)) = \omega(a|x),$$

for some $\omega(a|x)$. We can rewrite the inequality as

$$\begin{aligned} \sum_{x \in X} q(x) \sum_{a \in A} \omega(a|x) u(x, a) \geq \\ \lim_{n \rightarrow \infty} \Delta_n^{\alpha - \frac{1}{2}} \sum_{x \in X} \sum_{a \in A} \sum_{x' \in X} \sum_{a' \in A} k(x, x') \omega(a|x) \omega(a'|x') g_{a,a'}(r^*(a'')). \end{aligned}$$

For this inequality to be satisfied, it must be the case that either $\alpha \geq \frac{1}{2}$ or $\omega(a|x)$ is uninformative ($\omega(a|x) = \phi(a)$ for all $x \in X$). In the latter case,

$$\lim_{n \rightarrow \infty} \Delta_n^{-\alpha} (r_n^*(a) - r^*(a)) = \omega(a|x).$$

Therefore, for all convergent sub-sequences $p_n^*(a|x)$,

$$\lim_{\Delta \rightarrow 0^+} \Delta^\alpha \|p_n^*(a|x) - r_n^*(a)\| = 0.$$

5.2.5 Proof of lemma 2.2

The sequence problem can be written

$$V_{\Delta}(t, w_t) = \max_{\{p_{s,\Delta}(a|x)\}_{x \in X, a \in A, s \in [t,1]}} E_t[\bar{V}(w_1) - \sum_{j=0}^{(1-t)\Delta^{-1}-1} C(\{p_{t+\Delta j, \Delta}(a|x)\})],$$

subject to $w_{t+\Delta} = w_t + \mu\Delta + \Delta^{0.5}f(a, x)$. By the envelope theorem,

$$V_{\Delta, w}(t, w_t) = E_t[\bar{V}_w(w_1)],$$

and by the Lipschitz continuity of $\bar{V}(w_1)$, $V_{\Delta, w}(t, w_t) \in [0, K]$ for some constant K . That is, $V_{\Delta}(t, w_t)$ is Lipschitz continuous.

We can write the agent's single-period maximization problem as

$$V_{\Delta}(t, w_t) = \max_{\{p_{t,\Delta}(a|x)\} \in \mathcal{P}(A)} -C(\{p_{t,\Delta}(a|x)\}) + \sum_{x \in X} q(x) \sum_{a \in A} p_{t,\Delta}(a|x) V_{\Delta}(t + \Delta, w_t + \mu(a)\Delta + \Delta^{0.5}f(a, x)).$$

By the optimality of $p_{t,\Delta}^*(a|x)$, it must be the case that

$$\sum_{x \in X} q(x) \sum_{a \in A} (p_{t,\Delta}^*(a|x) - r_{t,\Delta}^*(a)) V_{\Delta}(t + \Delta, w_t + \mu(a)\Delta + \Delta^{0.5}f(a, x)) \geq C(\{p_{t,\Delta}^*(a|x)\}),$$

where $r_{t,\Delta}^*(a) = \sum_{x \in X} q(x) p_{t,\Delta}^*(a|x)$. By the Lipschitz continuity of $V_{\Delta}(t, w_t)$, this can be rewritten as

$$K \sum_{x \in X} q(x) \sum_{a \in A} (p_{t,\Delta}^*(a|x) - r_{t,\Delta}^*(a)) |\Delta^{0.5}f(a, x)| \geq C(\{p_{t,\Delta}^*(a|x)\}).$$

Define

$$\bar{K} = K \sum_{x \in X} q(x) [(\max_{a \in A} f(a, x)) - (\min_{a \in A} f(a, x))].$$

We can write this expression as

$$+\bar{K}\Delta^{\frac{1}{2}} \geq C(\{p_{t,\Delta}^*(a|x)\}).$$

It follows that the optimal policy becomes uninformative as $\Delta \rightarrow 0^+$.

The remainder of the proof is identical to the proof of theorem 2.2 (section 5.2.3 in the appendix), with the substitution

$$\begin{aligned} u(x, a) &= \Delta^{-\frac{1}{2}} V_{\Delta, w}(t + \Delta, w_t + \mu\Delta + c(a, x)\Delta^{\frac{1}{2}} f(a, x)) f(a, x) \\ &= V_{\Delta, w}(t, w_t) f(a, x) + o(1). \end{aligned}$$

5.2.6 Proof of lemma 2.3

Under the optimal policy,

$$p_{t,\Delta}^*(a|x) = r_t^*(a) + \phi_\Delta(a) + \Delta^{0.5} V_{\Delta, w}(t, w_t) r_t^*(a) \sum_{x' \in X} q(x') k^+(x, x') [f(a, x') - \sum_{a' \in A} r_t^*(a') f(a', x')] + o(\Delta^{\frac{1}{2}})$$

By definition,

$$E_t[w_{t+\Delta} - w_t] = \mu(r_t)\Delta + \sum_{x \in X} \sum_{a \in A} q(x) p_{t,\Delta}^*(a, x) \Delta^{0.5} f(a, x).$$

Therefore,

$$\begin{aligned} E_t[w_{t+\Delta} - w_t] &= \mu(r_t)\Delta + \Delta^{0.5} \sum_{x \in X} \sum_{a \in A} q(x) r_t^*(a) f(a, x) + \\ &\Delta V_{\Delta, w}(t, w_t) \sum_{x \in X} \sum_{a \in A} \sum_{x' \in X} r_t^*(a) q(x) q(x') k^+(x, x') [f(a, x') - \sum_{a' \in A} r_t^*(a') f(a', x')] + \\ &\Delta^{0.5} \sum_{x \in X} \sum_{a \in A} q(x) \phi_n(a) f(a, x) + o(\Delta). \end{aligned}$$

By assumption, $\sum_{x \in X} \sum_{a \in A} q(x)r(a)f(a,x) = 0$ for all $r(a) \in \mathcal{P}(A_+)$, and therefore

$$E[w_{t+\Delta} - w_t] = \mu(r_t)\Delta + \Delta V_{\Delta,w}(t, w_t)\sigma_{c,t}^2 + o(\Delta).$$

Similarly,

$$\begin{aligned} \text{Var}_t[w_{t+\Delta} - w_t - \mu(r_t)\Delta] &= \Delta \sum_{x \in X} \sum_{a \in A} q(x)p_t(a,x)f(a,x)^2 - \\ &\quad (\Delta V_w(t, w_t)\sigma_{c,t}^2)^2 + o(\Delta^2) \\ &= \Delta \sum_{x \in X} \sum_{a \in A} q(x)r_t^*(a)f(a,x)^2 + o(\Delta). \end{aligned}$$

5.2.7 Proof of lemma 2.4

This proof essentially follows the proof for the theorem 2.2. The problem is

$$V_{\Delta}(t, w_t) = \max_{\{p_{t,\Delta}(a|x) \in \mathcal{P}(A)\}} -C(\{p_{t,\Delta}(a|x)\}) + \sum_{x \in X} q(x) \sum_{a \in A} p_{t,\Delta}(a|x)V_{\Delta}(t + \Delta, w_t + \mu(a)\Delta + \Delta^{0.5}f(a,x)).$$

Using the results of lemma 2.2 and theorem 2.2, we can expand the cost to order Δ as

$$\begin{aligned} C(\{p_{t,\Delta}^*(a|x)\}) &= \frac{1}{2}\Delta V_{\Delta,w}(t, w_t)^2 \sum_{a \in A} \sum_{x' \in X} \sum_{x \in X} r_t^*(a)q(x)q(x')k^+(x, x')[f(a, x') - \sum_{a' \in A} r_t^*(a')f(a', x')] \cdot \\ &\quad [f(a, x) - \sum_{a' \in A} r_t^*(a')f(a', x)] + o(\Delta). \end{aligned}$$

Defining

$$m(a, a') = \sum_{x \in X} \sum_{x' \in X} q(x)k^+(x', x)q(x')f(x', a)f(x, a'),$$

and

$$\sigma_c^2(r_t) = \sum_{a \in A} r_t(a)[m(a, a) - \sum_{a' \in A} r_t(a')m(a, a')],$$

we have

$$C(\{p_{t,\Delta}^*(a|x)\}) = \frac{1}{2}\Delta V_{\Delta,w}(t, w_t)^2 \sigma_c^2(r_t^*).$$

We can rewrite the Bellman equation, assuming twice-differentiability for the value function, using a Taylor expansion for V_Δ , as

$$\begin{aligned} V_\Delta(t, w_t) &= \max_{r_t \in \mathcal{P}(A)} -\frac{1}{2}\Delta V_{\Delta,w}(t, w_t)^2 \sigma_c^2(r_t) + \\ &V_\Delta(t, w_t) + \Delta V_{\Delta,t}(t, w_t) + \Delta V_{\Delta,w}(t, w_t) \sum_{a \in A} \mu(a) r_t(a) + \\ &\Delta V_{\Delta,w}(t, w_t)^2 \sigma_c^2(r_t) + \\ &\Delta \frac{1}{2} V_{\Delta,ww}(t, w_t) \sum_{a \in A} \sum_{x \in X} q(x) r(a) f(a, x)^2 + o(\Delta). \end{aligned}$$

This simplifies to

$$0 = \max_{r_t \in \mathcal{P}(A)} \frac{1}{2}\Delta V_{\Delta,w}(t, w_t)^2 \sigma_c^2(r_t) + \Delta V_{\Delta,t}(t, w_t) + \Delta V_{\Delta,w}(t, w_t) \mu(r_t) + \Delta \frac{1}{2} V_{\Delta,ww}(t, w_t) \sigma^2(r_t) + o(\Delta).$$

The optimal $r_t(a)$ (which is unique by the arguments in section 5.2.3) solves

$$r_t^* = \arg \max_{r_t \in \mathcal{P}(A)} V_{\Delta,w}(t, w_t) \mu(r_t) + \frac{1}{2} V_{\Delta,w}(t, w_t)^2 \sigma_c^2(r_t) + \frac{1}{2} V_{\Delta,ww}(t, w_t) \sigma^2(r_t).$$

Under the optimal choice of the variable e_t , the Bellman equation in the lemma is equivalent to the one derived here. The evolution of the wealth process follows from lemma 2.3.

5.2.8 Proof of lemma 5.4

Define the norm $\|\cdot\|$ such that

$$\|\{p(s|x) - r(s)\}\|^2 = \left(\sum_{x \in X} \sum_{a \in A} q(x) (p(s|x) - r(s)) |f(a, x)| \right)^2.$$

By the results in section 5.2.5, we have

$$\frac{1}{2}\lambda(1-\lambda)m\left(\sum_{x \in X} \sum_{a \in A} q(x)(p_{t,\Delta}^*(a|x) - r_{t,\Delta}^*(a))|f(a,x)|\right)^2 \leq \lambda C(\{p(s|x)\}).$$

By the convexity of $C(\cdot)$,

$$\langle p_{t,\Delta}^*(a|x) - r_{t,\Delta}^*(a), \frac{\partial C(\{p_{t,\Delta}^*(a''|x'')\})}{\partial p(a|x)} \rangle \geq C(\{p(s|x)\}).$$

By the Cauchy-Schwarz inequality,

$$\|\{p(s|x) - r(s)\}\| \cdot \left\| \frac{\partial C(\{p_{t,\Delta}^*(a''|x'')\})}{\partial p(a|x)} \right\| \geq C(\{p(s|x)\}).$$

Therefore,

$$\frac{1}{2}(1-\lambda)m\left(\sum_{x \in X} \sum_{a \in A} q(x)(p_{t,\Delta}^*(a|x) - r_{t,\Delta}^*(a))|f(a,x)|\right) \leq \left\| \frac{\partial C(\{p_{t,\Delta}^*(a''|x'')\})}{\partial p(a|x)} \right\|.$$

By the first-order condition,

$$\frac{\partial C(\{p_{t,\Delta}^*(a''|x'')\})}{\partial p(a|x)} = q(x)[V_\Delta(t + \Delta, w_t + \mu(a)\Delta + \Delta^{\frac{1}{2}}f(a,x)) - V_\Delta(t + \Delta, w_t)].$$

By the Lipschitz continuity of V_Δ ,

$$\left| \frac{\partial C(\{p_{t,\Delta}^*(a''|x'')\})}{\partial p(a|x)} \right| \leq q(x)K|\Delta\mu(a) + \Delta^{\frac{1}{2}}f(a,x)|.$$

Therefore, for some $\bar{\Delta} > 0$, there exists a \bar{C} such that, for all $\Delta < \bar{\Delta}$,

$$\left\| \frac{\partial C(\{p_{t,\Delta}^*(a''|x'')\})}{\partial p(a|x)} \right\| \leq \bar{C}\Delta^{\frac{1}{2}},$$

and the result follows.

5.2.9 Proof of theorem 2.3

This proof employs the following lemma.

Lemma 5.4. *Under the optimal policy $p_{t,\Delta}^*(a|x)$, there exists a $\bar{\Delta} > 0$ and constant \bar{C} such that*

$$\bar{C}\Delta^{\frac{1}{2}} \geq \sum_{a \in A} \sum_{x \in X} q(x) [p_{t,\Delta}^*(a|x) - \sum_{x' \in x} q(x') p_{t,\Delta}^*(a|x')] |f(a,x)|$$

for all $t \in [0, 1]$, w_t , and $\Delta < \bar{\Delta}$.

Proof. See appendix, section 5.2.8. □

Consider the sequence problem:

$$V_{\Delta}(t, w_t) = E_t[\bar{V}(w_1) - \sum_{j=0}^{(1-t)\Delta^{-1}-1} C(\{p_{t+\Delta j, \Delta}^*(a|x)\})].$$

Define the function $G_{\Delta}(t) = \Delta^{-1}C(\{p_{t+\Delta j, \Delta}^*(a|x)\})$ for all $t \in [j\Delta, (j+1)\Delta)$ (that is, $G_{\Delta}(t)$ is the simple process representing the cumulative information costs). The random variable representing the remaining cumulative information costs is

$$\sum_{j=0}^{(1-t)\Delta^{-1}-1} C(\{p_{t+\Delta j, \Delta}^*(a|x)\}) = \int_0^1 G_{\Delta}(s) ds.$$

By lemma 5.4 and section 5.2.5, it follows that $C(\{p_{t+\Delta j, \Delta}^*(a|x)\}) \leq \Delta(\bar{K}_1 + \bar{C}\bar{K}_2)$. By the dominated convergence theorem,

$$\lim_{\Delta \rightarrow 0^+} \sum_{j=0}^{(1-t)\Delta^{-1}-1} C(\{p_{t+\Delta j, \Delta}^*(a|x)\}) = \lim_{\Delta \rightarrow 0^+} \int_0^1 G_{\Delta}(s) ds = \int_0^1 (\lim_{\Delta \rightarrow 0^+} G_{\Delta}(s)) ds.$$

By the results of section 5.2.7,

$$\lim_{\Delta \rightarrow 0^+} G_{\Delta}(t) = \frac{1}{2} V_w(t, w_t)^2 \sigma_{c,t}^2.$$

Therefore,

$$\lim_{\Delta \rightarrow 0^+} V(t, w_t) = E_t[\bar{V}(w_1) - \frac{1}{2} \int_t^1 V_w(s, w_s)^2 \sigma_{c,s}^2 ds].$$

Next, we consider the evolution of the wealth process. Define the variable

$$\mu_{t,\Delta} = \sum_{a \in A} \sum_{x \in X} p_{t,\Delta}^*(a|x) q(x) f(a, x),$$

and the variable

$$\sigma_{t,\Delta}^2 = \sum_{a \in A} \sum_{x \in X} p_{t,\Delta}^*(a|x) q(x) f(a, x)^2 - \mu_{t,\Delta}^2.$$

Define the random variable $\varepsilon_{t+\Delta,\Delta} = \sigma_{t,\Delta}^{-1}(f(a, x) - \mu_{t,\Delta})$. Note that $\varepsilon_{t+\Delta,\Delta}$ is a mean-zero, unit variance random variable. The wealth process is

$$w_t = w_0 + \sum_{j=0}^{t\Delta^{-1}-1} (\mu_{\Delta j,\Delta} + \mu_{\Delta j,\Delta} \Delta^{0.5}) + \sum_{j=0}^{t\Delta^{-1}-1} \Delta^{0.5} \sigma_{\Delta j,\Delta} \varepsilon_{(j+1)\Delta,\Delta}.$$

Defining

$$B_t = \Delta^{0.5} \sum_{k=0}^{[t\Delta^{-1}]} \varepsilon_{(k+1)\Delta} + \Delta^{0.5} (t - \Delta[[t\Delta^{-1}]]) \varepsilon_{\Delta[[t\Delta^{-1}]]+2\Delta},$$

where $[[t\Delta^{-1}]]$ denotes the largest integer less than or equal to $t\Delta^{-1}$, then by the invariance principle (Karatzas and Shreve (1991), chapter 2, theorem 4.20), we have that B_t converges in distribution to a standard Brownian motion on the canonical probability space.

First, we show that the mean converges. By section 5.2.8, it follows that $|\mu_{\Delta j,\Delta}| \leq \bar{C}\Delta$. By the dominated convergence theorem,

$$\sum_{j=0}^{t\Delta^{-1}-1} \mu_{\Delta j,\Delta} \Delta^{0.5} = \Delta^{0.5} \sum_{j=0}^{t\Delta^{-1}-1} \lim_{\Delta \rightarrow 0^+} \mu_{\Delta j,\Delta} = \int_0^t V_w(s, w_s) \sigma_{c,s}^2 ds,$$

where the last line follows from lemma 2.3. Next, we show that the random term converges. We have (see lemma 2.3)

$$\sum_{j=0}^{t\Delta^{-1}-1} \Delta^{0.5} \sigma_{\Delta j, \Delta} \varepsilon_{(j+1)\Delta, \Delta} = \sum_{j=0}^{t\Delta^{-1}-1} \Delta^{0.5} \sigma_{\Delta j} \varepsilon_{(j+1)\Delta, \Delta} + \sum_{j=0}^{t\Delta^{-1}-1} \varepsilon_{(j+1)\Delta, \Delta} \Delta^{0.5} (\sigma_{\Delta j, \Delta} - \sigma_{\Delta j}).$$

Define $\bar{\sigma}^2 = \sum_{x \in X} q(x) \max_{a \in A} f(a, x)^2$. The variance of the remainder term is

$$V_0 \left[\sum_{j=0}^{t\Delta^{-1}-1} \varepsilon_{(j+1)\Delta, \Delta} \Delta^{0.5} (\sigma_{\Delta j, \Delta} - \sigma_{\Delta j}) \right] = \sum_{j=0}^{t\Delta^{-1}-1} \Delta (\sigma_{\Delta j, \Delta} - \sigma_{\Delta j})^2.$$

Note that $(\sigma_{\Delta j, \Delta} - \sigma_{\Delta j})^2 \leq \bar{\sigma}^2$. By the dominated convergence theorem, it follows that

$$\lim_{\Delta \rightarrow 0^+} \sum_{j=0}^{t\Delta^{-1}-1} \Delta (\sigma_{\Delta j, \Delta} - \sigma_{\Delta j})^2 = 0.$$

Therefore, $\sum_{j=0}^{t\Delta^{-1}-1} \Delta^{0.5} \sigma_{\Delta j, \Delta} \varepsilon_{(j+1)\Delta, \Delta}$ converges in an L^2 sense to

$$\lim_{\Delta \rightarrow 0^+} \sum_{j=0}^{t\Delta^{-1}-1} \Delta^{0.5} \sigma_{\Delta j} \varepsilon_{(j+1)\Delta, \Delta}.$$

By the definition of the Ito integral, and the square integrability and adaptedness of σ_t , it follows that

$$\lim_{\Delta \rightarrow 0^+} \sum_{j=0}^{t\Delta^{-1}-1} \Delta^{0.5} \sigma_{\Delta j, \Delta} \varepsilon_{(j+1)\Delta, \Delta} = \int_0^t \sigma_s dB_s.$$

Putting all of this together, w_t converges in distribution to

$$w_t = w_0 + \int_0^t (\mu_s + V_w(s, w_s) \sigma_{c, s}^2) ds + \int_0^t \sigma_s dB_s.$$

It follows that

$$dw_t = \mu_s dt + V_w(t, w_t) \sigma_{c,t}^2 dt + \sigma_t dB_t.$$

The HJB equation for the value function, under the optimal policy, is

$$\begin{aligned} 0 = & -\frac{1}{2} V_w(t, w_t)^2 \sigma_{c,t}^2 dt + V_t(t, w_t) dt + \\ & V_w(t, w_t) (\mu_s + V_w(t, w_t) \sigma_{c,t}^2) + \\ & \frac{1}{2} V_{ww}(t, w_t) \sigma_t^2 dt. \end{aligned}$$

Introducing the choice variable e_t , we can derive this HJB equation from the following equation:

$$\begin{aligned} 0 = & \max_{e_t} -\frac{1}{2} \sigma_{c,t}^2 e_t^2 dt + V_t(t, w_t) dt + \\ & V_w(t, w_t) (\mu_s + e_t \sigma_{c,t}^2) + \\ & \frac{1}{2} V_{ww}(t, w_t) \sigma_t^2 dt \end{aligned}$$

Under the optimal choice of e_t , $e_t^* = V_w(t, w_t)$, and the evolution of the wealth process is

$$dw_t = \mu_s dt + e_t \sigma_{c,t}^2 dt + \sigma_t dB_t.$$

Introducing the choice variable r_t , we can rewrite this as

$$\begin{aligned} 0 = & \max_{e_t, r_t \in \mathcal{P}(A)} -\frac{1}{2} \sigma_c^2(r_t) e_t^2 dt + V_t(t, w_t) dt + \\ & V_w(t, w_t) (\mu(r_t) + e_t \sigma_c^2(r_t)) + \\ & \frac{1}{2} V_{ww}(t, w_t) \sigma^2(r_t) dt. \end{aligned}$$

It follows that r_t will maximize

$$V_w(t, w_t)\mu(r_t) + \frac{1}{2}V_w(t, w_t)^2\sigma_c^2(r_t) + \frac{1}{2}V_{ww}(t, w_t)\sigma^2(r_t),$$

and therefore be equivalent to r_t^* .

5.3 Appendix: The Costs of Sovereign Default

5.3.1 GMM and Returns-IV

In this section, we present results for the GMM and Returns-IV estimators discussed in the text. Given the problematic behavior of these estimators under the null hypothesis that $\alpha = 0$, we cannot interpret our bootstrapped confidence intervals as providing correct coverage for the t-tests and J-tests conducted in this section. We therefore have removed all the of asterisks from the tables, although we list the standard errors and confidence intervals generated by our procedure.

For our GMM confidence intervals, we use the moment-recentering procedure discussed by Horowitz (2001). We also employ this bootstrap strategy to estimate the 95% confidence interval for the over-identification test (J statistic). The 95% confidence interval for the J-statistic is based on the 95th percentile of the sampling distribution, and the associated test is one-sided. Currently, we run our GMM procedure on abnormal returns/CDS changes, treating them as known. The GMM estimator is a two-step GMM estimator.

Table 5.2: Returns-IV

	(1)	(2)	(3)	(4)
	Index	Banks	Non-Financial	Real Estate
ΔD	-61.56***	-68.53***	-46.66*	176.7
Robust SE	(15.50)	(21.09)	(20.05)	(506.9)
95% CI	[-99.3,-13.5]	[-104.4,-29.5]	[-84.8,9.0]	[-3.5e+3,472.2]
Observations	413	413	413	413
1st Stage F-Stat	81.99	71.37	41.32	0.237

	(5)	(6)	(7)
	FX (ADR)	FX (On.)	FX (Official)
ΔD	55.59***	-9.886	499.5
Robust SE	(12.23)	(11.61)	(1,058)
95% CI	[13.1,81.1]	[-43.8,34.9]	[-1.3e+5,721.9]
Observations	368	413	413
1st Stage F-Stat			

Notes: This table reports the results for the variance-based estimator estimated as the ratio of $\lambda\alpha^2$ to $\lambda\alpha$. This estimator is called the “Returns-IV” estimator because it depends on the excess variance of the ADR return on event days. The column headings denote the outcome variable. Index is the MSCI Argentina Index, Banks is our equally weighted index of Argentine bank ADRs, Industrial is our equally weighted index of Argentine industrial ADRs, and REIT is our equally weighted index of Argentine real estate holding companies. FX (ADR) is the ARS/USD exchange rate derived from the ratio of ADR prices (in USD) to the price of the underlying equity (in ARS). FX (On.) is the ARS/USD exchange rate offered by onshore currency dealers. FX (Official) is the exchange rate set by the Argentine government. The coefficient on ΔD is the effect on the percentage returns of an increase in the 5-year risk-neutral default probability from 0% to 100%, implied by the Argentine CDS curve. Standard errors and confidence intervals are computed using the stratified bootstrap procedure described in the text. The underlying data is based on the two-day event windows and non-events described in the text.

Table 5.3: GMM

	(1)	(2)	(3)	(4)
	Index	Banks	Non-Financial	Real Estate
α (ΔD)	-54.45**	-60.09***	-30.95	-0.466
Robust SE	(14.05)	(14.16)	(16.67)	(14.58)
95% CI	[-99.4,-15.5]	[-99.1,-28.6]	[-79.0,33.9]	[-64.5,37.7]
λ	30.04	30.55	26.42	33.05
Robust SE	(13.21)	(13.64)	(12.54)	(13.02)
95% CI	[-143.8,57.4]	[-41.8,56.2]	[-264.1,56.7]	[-497.5,57.5]
Observations	413	413	413	413
J-Stat	0.180	0.280	0.596	0.0207
J-Stat-CI	[0,6.4]	[0,4.4]	[0,4.9]	[0,8.8]

	(5)	(6)	(7)
	FX (ADR)	FX (On.)	FX (Official)
α (ΔD)	40.82	17.87**	-0.166
Robust SE	(21.24)	(4.218)	(0.851)
95% CI	[-80.6,129.4]	[0.9,29.1]	[-2.3,1.5]
λ	21.39	50.85**	33.60
Robust SE	(10.44)	(11.66)	(14.17)
95% CI	[-166.3,52.2]	[3.2,69.3]	[-4.4e+03,48.9]
Observations	368	413	413
J-Stat	1.160	5.521**	0.872
J-Stat-CI	[0,15.5]	[0,2.8]	[0,25.0]

Notes: The GMM estimates are based on a two-step estimator, run once for each outcome variable. The column headings denote the outcome variable. Index is the MSCI Argentina Index, Banks is our equally weighted index of Argentine bank ADRs, Non-Financial is our equally weighted index of Argentine non-financial ADRs, and Real Estate is our equally weighted index of Argentine real estate holding companies. FX (ADR) is the ARS/USD exchange rate derived from the ratio of ADR prices (in USD) to the price of the underlying equity (in ARS). FX (On.) is the ARS/USD exchange rate offered by onshore currency dealers. FX (Official) is the exchange rate set by the Argentine government. The parameter α is the effect on the percentage returns of an increase in the probability of default, from 0% to 100%. λ is proportional to the difference in the variance of the default probability shocks during event and non-event windows. Standard errors and confidence intervals are computed using the stratified bootstrap procedure described in the text, with moment-recentering. J-Stat is an over-identification test of the validity of the assumptions described in Rigobon and Sack (2004). The underlying data is based on the two-day event windows and non-events described in the text.

Table 5.4: Regressions for Brazil and Mexico

	(1)	(2)	(3)	(4)
	Brazil CDS	Brazil Index	Mexico CDS	Mexico Index
OLS ΔD	53.08***	-12.21***	42.75***	-5.978**
Robust SE	(8.632)	(3.823)	(7.861)	(3.150)
95% CI	[30.53,76.38]	[-18.40,-6.50]	[20.43,67.46]	[-11.18,-0.98]
Event IV ΔD	23.16**	-3.035	6.330	0.634
Robust SE	(15.22)	(6.592)	(13.98)	(5.426)
95% CI	[1.14,53.00]	[-15.04,9.30]	[-25.62,48.25]	[-4.14,7.93]
CDS-IV ΔD	20.68	-1.098	1.728	1.669
Robust SE	(16.44)	(7.075)	(15.09)	(5.812)
95% CI	[-4.28,54.10]	[-9.84,10.18]	[-33.83,48.53]	[-3.22,8.73]

Notes: This table reports the results for the OLS, IV-style event study, and CDS-IV estimators of the effect of changes in the risk-neutral default probability (ΔD) on the 5-year CDS spreads and stock market indices of Brazil and Mexico. The coefficient on ΔD is the effect on the percentage returns (of stocks) and change in the 5-year CDS spread (in bps) of an increase in the 5-year risk-neutral default probability from 0% to 100%, implied by the Argentine CDS curve. Standard errors and confidence intervals are computed using the stratified bootstrap procedure described in the text. The underlying data is based on the two-day event windows and non-events described in the text. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.3.2 Mexico and Brazil

5.3.3 Risk-Neutral Default Probabilities

We convert CDS spreads into risk-neutral default probabilities to provide a clearer sense of the magnitude of the estimated coefficients. We emphasize that we work with risk-neutral probabilities and do not attempt to convert them to physical probabilities. Pan and Singleton (2008) and Longstaff et al. (2011a) impose additional structure to estimate the physical default probabilities.

We begin with data from Markit on CDS par spreads. The par spread is the coupon payment that a buyer of CDS protections pays to the seller of the contract such that the CDS contract has zero cost at initiation. Because the seller of a CDS insures the buyer of a CDS against credit losses throughout the duration of the contract, pricing the contract involves calculating the term structure of credit risk on the bond.

The market standard for pricing CDS is a reduced form model that models time-varying credit risk as a time-varying hazard rate of default.¹²⁶ The simplest version of such a model would be to

¹²⁶White (2013) provides a very thorough discussion of the ISDA standard model.

assume that throughout the life of CDS contract there is a constant default hazard rate λ . In this simple case, we can convert the par spread S_T for a contract with maturity T to the hazard rate,

$$\lambda = \frac{S_T}{1 - Recov.} \quad (5.7)$$

where $Recov.$ is the recovery rate, which is assumed to be known and constant. Once this hazard rate λ is calculated, we can calculate the probability that a bond defaults before time $t = T$ as

$$Pr(Def < T) = 1 - \exp(-\lambda T).$$

For example, a 1-year CDS with zero recovery and a par spread of 100% would imply a hazard rate of 1. This means that half the time the bond would fully default and the seller would fully compensate the buyer, and half the time the underlying bond would not default and the seller would earn an annual interest rate of 100%, breaking even on average.

If there were only one tenor of CDS observed in the market, this constant hazard rate calculation would be all that is feasible. However, with the multiple tenors we do not need to restrict the hazard function λ to be constant throughout the duration of the CDS. Our dataset includes quotes at the 6 month, 1 year, 2 year, 3 year, 4 year, 5 year, 7 year, 10 year, 15 year, 20 year and 30 year tenors. We follow the ISDA standard and construct the risk-free yield curve using the Libor deposit rate for the 6 month tenor and interest rate swap rate for all longer maturities. This data can be downloaded from FRED, the online database of the Federal Reserve Bank of St. Louis.

We implement the standard market model in Matlab using *cdsbootstrap*. The assumption of the standard model is that the time varying hazard rate $\lambda(t)$ is constant between all of the nodes of the CDS curve. This mean that we begin by using 5.7 to calculate λ_{6m} , the hazard rate between the initiation of the contract and its expiration 6 months later. We use the recovery rate that Markit calculates by polling the reporting dealers as our assumed recovery upon default. This recovery rate is assumed to be the same for all tenors. Having calculated λ_{6M} , we can calculate λ_{1Y} , the hazard rate between 6 months and 1 year of the contract consistent with the observed par spread,

then λ_{2Y} between 1 and 2 years, and so on up the curve. Having calculated these hazard rates, we can then compute the probability of a default during the life of each contract as:

$$\begin{aligned}
 Pr(D \leq 6M) &= 1 - \exp\left(-\lambda_{6M} \cdot \left(\frac{1}{2}\right)\right) \\
 Pr(D \leq 1Y) &= 1 - \exp\left(-\lambda_{6M} \cdot \left(\frac{1}{2}\right) - \lambda_{1Y} \cdot \left(\frac{1}{2}\right)\right) \\
 &\vdots \\
 Pr(D \leq 5Y) &= 1 - \exp\left(-\lambda_{6M} \cdot \left(\frac{1}{2}\right) - \lambda_{1Y} \cdot \left(\frac{1}{2}\right) - \lambda_{2Y} - \lambda_{3Y} - \lambda_{4Y} - \lambda_{5Y}\right)
 \end{aligned}$$

We perform this bootstrapping for our 11 tenors for the full sample period (January 3, 2011 to July 30, 2014).

5.3.4 Econometric Model

The model we use is

$$\begin{aligned}
 \Delta D_t &= \mu_d + \omega_D^T X_t + \gamma^T r_t + \beta_D F_t + \varepsilon_t \\
 r_t &= \mu + \Omega X_t + \alpha \Delta D_t + \beta F_t + \eta_t,
 \end{aligned}$$

where r_t is a vector of returns, ΔD_t is the change in the default probability, X_t is a set of global factors (S&P 500, etc...), F_t is an unobserved factor, and ε_t is the idiosyncratic default probability shock, and η_t is a vector of return shocks that do not directly affect the probability of default. Through some algebra, we show that this is equivalent to the systems described in equations 3.1 and 3.2, used in most of our analysis, and equations 3.8 and 3.9 used in the cross-sectional analysis.

We begin by separating the equation governing the vector of returns r_t into the return of asset i , $r_{i,t}$, which is the asset of interest, and the returns of some other assets, denoted $r_{-i,t}$. We separate the various coefficient vectors and matrices, $\mu, \Omega, \alpha, \beta, \gamma$, and shocks η_t , into versions for asset i ,

μ_i, ω_i^T , etc..., and versions for the other assets, μ_{-i}, Ω_{-i} , etc... This system can be written as

$$\begin{aligned}\Delta D_t &= \mu_d + \omega_D^T X_t + \gamma_i^T r_{i,t} + \gamma_{-i}^T r_{-i,t} + \beta_D F_t + \varepsilon_t \\ r_{i,t} &= \mu_i + \omega_i^T X_t + \alpha_i \Delta D_t + \beta_i F_t + \eta_{i,t} \\ r_{-i,t} &= \mu_{-i} + \Omega_{-i} X_t + \alpha_{-i} \Delta D_t + \beta_{-i} F_t + \eta_{-i,t}.\end{aligned}$$

Most of our analysis considers only a single asset, $r_{i,t}$, and the default probably change ΔD_t . Substituting the returns $r_{-i,t}$ into the ΔD_t equation,

$$\begin{aligned}\Delta D_t &= \frac{\mu_d + \gamma_{-i}^T \mu_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} + \frac{\omega_D^T + \beta_{-i}^T \Omega_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} X_t + \frac{\gamma_i^T r_{i,t}}{1 - \gamma_{-i}^T \alpha_{-i}} + \\ &\quad \frac{\beta_D + \gamma_{-i}^T \beta_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} F_t + \frac{1}{1 - \gamma_{-i}^T \alpha_{-i}} (\gamma_{-i}^T \eta_{-i,t} + \varepsilon_t) \\ r_{i,t} &= \mu_i + \omega_i^T X_t + \alpha_i \Delta D_t + \beta_i F_t + \eta_{i,t}.\end{aligned}$$

This system, for the two assets, is equivalent to the one in equations 3.1 and 3.2, except that it has two shocks, $\gamma_{-i}^T \eta_{-i,t}$ and ε_t , that directly affect ΔD_t without affecting $r_{i,t}$, and includes constants and observable controls X_t . Neither of these differences substantially alter the identification assumptions or analysis. The event study and Rigobon (2003) approach both identify the coefficient α_i , under their identifying assumptions, which is the coefficient of interest.

Next, we discuss a version of this system with the market return. Let the market return be a weighted version of the return vector, $r_{m,t} = w^T r_t$. Separating the vectorized version of the system into four equations,

$$\begin{aligned}\Delta D_t &= \mu_d + \omega_D^T X_t + \gamma_i^T r_{i,t} + \gamma_{-i}^T r_{-i,t} + \beta_D F_t + \varepsilon_t \\ r_{i,t} &= \mu_i + \omega_i^T X_t + \alpha_i \Delta D_t + \beta_i F_t + \eta_{i,t} \\ r_{-i,t} &= \mu_{-i} + \Omega_{-i} X_t + \alpha_{-i} \Delta D_t + \beta_{-i} F_t + \eta_{-i,t} \\ r_{m,t} &= \mu_m + \omega_m^T X_t + \alpha_m \Delta D_t + F_t + w^T \eta_t,\end{aligned}$$

where $\mu_m = w^T \mu$, $\omega_m^T = w^T \Omega$, and so on. We have assumed that $w^T \beta = 1$, which is a normalization. Substituting out $r_{-i,t}$,

$$\begin{aligned}\Delta D_t &= \frac{\mu_d + \gamma_{-i}^T \mu_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} + \frac{\omega_D^T + \beta_{-i}^T \Omega_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} X_t + \frac{\gamma_i^T r_{i,t}}{1 - \gamma_{-i}^T \alpha_{-i}} + \\ &\quad \frac{\beta_D + \gamma_{-i}^T \beta_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} F_t + \frac{1}{1 - \gamma_{-i}^T \alpha_{-i}} (\gamma_{-i}^T \eta_{-i,t} + \varepsilon_t) \\ r_{i,t} &= \mu_i + \omega_i^T X_t + \alpha_i \Delta D_t + \beta_i F_t + \eta_{i,t} \\ r_{m,t} &= \mu_m + \omega_m^T X_t + \alpha_m \Delta D_t + F_t + w^T \eta_t,\end{aligned}$$

as above. Next, we solve for F_t using the market return equation:

$$F_t = r_{m,t} - \mu_m - \omega_m^T X_t - \alpha_m \Delta D_t - w^T \eta_t.$$

Plugging this into our system of equations,

$$\begin{aligned}(1 + \alpha_m \frac{\beta_D + \gamma_{-i}^T \beta_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}}) \Delta D_t &= (\frac{\mu_d + \gamma_{-i}^T \mu_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} - \frac{\beta_D + \gamma_{-i}^T \beta_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} \mu_m) \\ &\quad + (\frac{\omega_D^T + \beta_{-i}^T \Omega_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} - \frac{\beta_D + \gamma_{-i}^T \beta_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} \omega_m^T) X_t + \\ &\quad \frac{\gamma_i^T r_{i,t}}{1 - \gamma_{-i}^T \alpha_{-i}} + \frac{\beta_D + \gamma_{-i}^T \beta_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} r_{m,t} + \\ &\quad (\frac{\gamma_{-i}^T}{1 - \gamma_{-i}^T \alpha_{-i}} - \frac{\beta_D + \gamma_{-i}^T \beta_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} w_{-i}^T) \eta_{-i,t} + \\ &\quad \frac{\beta_D + \gamma_{-i}^T \beta_{-i}}{1 - \gamma_{-i}^T \alpha_{-i}} w_i \eta_{i,t} + \frac{1}{1 - \gamma_{-i}^T \alpha_{-i}} \varepsilon_t \\ r_{i,t} &= (\mu_i - \beta_i \mu_m) + (\omega_i^T - \beta_i \omega_m^T) X_t + (\alpha_i - \beta_i \alpha_m) \Delta D_t \\ &\quad + \beta_i r_{m,t} + (1 - w_i \beta_i) \eta_{i,t} + w_{-i}^T \eta_{-i,t}.\end{aligned}$$

This system is equivalent to the one presented in equations 3.8 and 3.9, except that there are multiple common factors ($\eta_{i,t}$ and $\eta_{-i,t}$) and no idiosyncratic return shocks. The event study

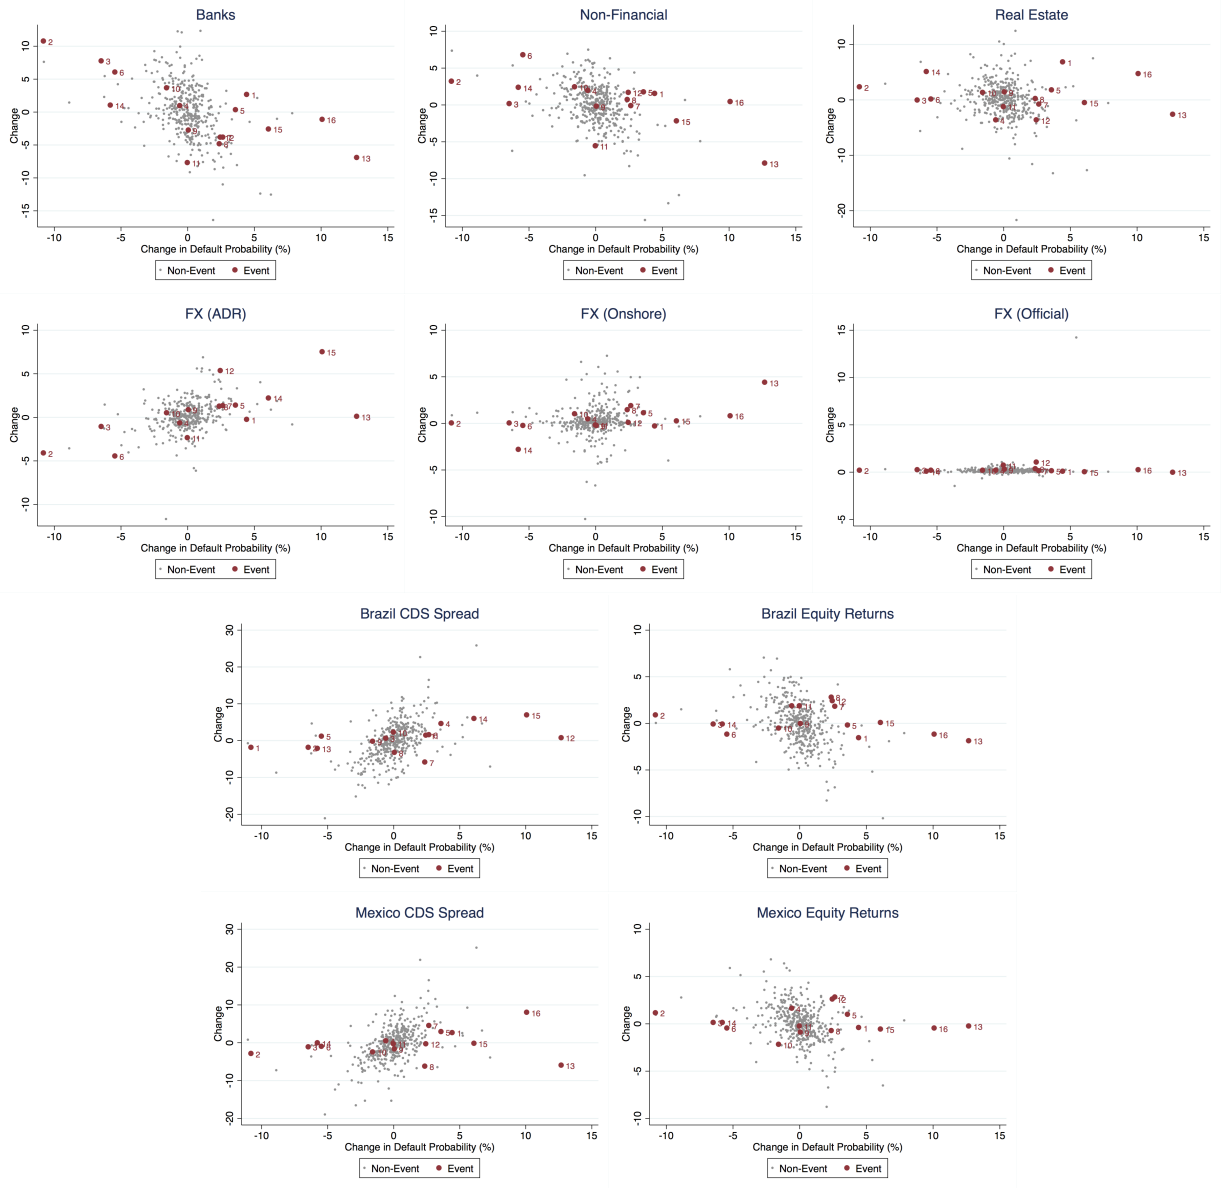
and Rigobon (2003) approach both identify the coefficient $(\alpha_i - \beta_i \alpha_m)$, under their identifying assumptions, which is the coefficient of interest.

5.3.5 Figures and Tables

Table 5.5: Events and Non-Events

Two-Day Window End	Event Type	Description	PDF Time (EST) and Link	News Time (EST) and Link	Comments
7Dec11	Excluded	Original ruling by Judge Griesa with regards to Pari Passu clause.	7Dec11, 12:55pm Decision	Missing	There was very little contemporaneous news coverage, and we are unable to determine when the ruling became public. The first story we found about the ruling is based on an article in “La Nacion” published on 5Mar12.
23Feb12	Excluded	Order by Judge Griesa requiring “ratable payment.”	Missing Order	Missing	See above.
05Mar12	Excluded	Stay granted by Judge Griesa, pending appeal.	Missing Stay	05Mar12, 7:11am Bloomberg	See above.
26Oct12	Excluded	Appeals court upholds Judge Griesa’s ruling that the Pari Passu clause requires equal treatment of restructured bondholders and holdouts.	25Oct12, 12:43pm Decision	26Oct12, 2:14pm Bloomberg	The appeals court releases opinions during the middle of the day. Unfortunately, the closing marks on this day are questionable, given the impending impact of “Superstorm Sandy.”

Figure 5.3: Change in Default Probability and other Financial Variables on Event and Non-Event Days



Notes: This figure plots the change in the risk-neutral probability of default and returns on all indices and exchange rates, as well as Mexican and Brazilian equities and CDS, on event and non-event days. Each event and non-event day is a two-day event or non-event as described in the text. The numbers next to each maroon dot references each event-day in the table below Figure 3.1. The procedure for classifying events and non-events is described in the text.

Table 5.5: Events and Non-Events (Continued)

Two-Day Window End	Event Type	Description	PDF Time (EST) and Link	News Time (EST) and Link	Comments
23Nov12	Excluded	Judge Griesa removes the stay on his order that Argentina immediately pay the holdouts, if they also pay the exchange bondholders.	Missing Order	22Nov12, 5:33am. Business News Americas	Nov 22 was Thanksgiving in the United States, and all CDS marks on that date and the morning of the 23rd appear to be the same as on the 21st. The opinion was filed by Judge Griesa on the night of the 21st, but was embargoed until the 23rd. On the 22nd, the Argentine market fell a lot, but bounced back on the 23rd. We cannot observe this in the ADR data, so we exclude this event.
27Nov12	Open-to-Open, 26Nov12 to 27Nov12	Judge Griesa denies the exchange bondholders request for a stay. The bondholders immediately appealed.	26Nov12, 3:43pm Denial	27Nov12, 5:00am. New York Post	The denial occurred on the 26th, and both the government of Argentina and the exchange bondholders immediately appealed. We compare the open on the 27th to the open on the 26th. The 26th is an Argentine holiday, so the ADR Blue Rate is missing (for the open-to-open, but not the two day window).
29Nov12	Close-to-Open, 28Nov12 to 29Nov12	Appeals court grants emergency stay of Judge Griesa's order.	28Nov12, 5:04pm Stay	29Nov12, 8:24am. Bloomberg	
05Dec12	Open-to-Close, 04Dec12	Appeals court denies request of holdouts to force Argentina to post security against the payments owed.	04Dec12, 1:15pm. Denial	04Dec12, 1:46pm. Bloomberg	

Table 5.5: Events and Non-Events (Continued)

Two-Day Window End	Event Type	Description	PDF Time (EST) and Link	News Time (EST) and Link	Comments
07Dec12	Close-to-Close, 05Dec12 to 06Dec12	Appeals court allows the Bank of New York (custodian of the exchange bonds) and the Euro bondholders to appear as interested parties.	05Dec12, 10:13pm. Order	06Dec12, 11:47am Bloomberg	
11Jan13	Close-to-Open, 10Jan13 to 11Jan13	Appeals court denies certification for exchange bondholders to appeal to NY state court for interpretation on Pari Passu clause.	10Jan13, 4:10pm Order	11Jan13, 12:01am Bloomberg	The ruling was written immediately after the closes on the 10th.
28Feb13	Excluded	Appeals court denies request for en-banc hearing of appeal.	28Feb13, 3:27pm. Decision	Missing Shearman	The denial occurred at the beginning of a hearing, during which lawyers for both sides argued various issues. Lawyers from Argentina may have changed their arguments in response to expectations about the Argentine economy, violating the exclusion restriction.
04Mar13	Open-to-Open, 01Mar13 to 04Mar13	Appeals court asked Argentina for a payment formula.	01Mar13, 11:49am. Order	01Mar13, 4:46pm Financial Times	The FT story is the earliest we could find. Most other coverage is from the following day (a Saturday).

Table 5.5: Events and Non-Events (Continued)

Two-Day Window End	Event Type	Description	PDF Time (EST) and Link	News Time (EST) and Link	Comments
04Mar13	Open-to-Open, 01Mar13 to 04Mar13	Appeals court asked Argentina for a payment formula.	01Mar13, 11:49am. Order	01Mar13, 4:46pm Financial Times	The FT story is the earliest we could find. Most other coverage is from the following day (a Saturday).
27Mar13	Open-to-Open, 27Mar13 to 26Mar13	Appeals court denies Argentina's request for en-banc rehearing.	26Mar13, 11:58am Order	26Mar13, 2:35pm Bloomberg	The Bloomberg story specifically mentions a 374bp increase in the 5yr CDS spread, which does not appear in our data until after the NY close at 3:30pm. We use the one day window to ensure we are capturing the event.
01Apr13	Non-Event (neither event or excluded)	Argentina files payment plan. Offer roughly 1/6 of Judge Griesa ordered.	N/A	30Mar13, 12:05pm Bloomberg	Argentina filed just before midnight on 28Mar13. Actions by Argentina are endogenous. This neither an event nor excluded.
22Apr13	Non-Event (neither event or excluded)	Holdouts reject Argentina's payment plan.	19Apr13, 5:20pm Reply	20Apr13, 12:01am Bloomberg	Holdouts reject Argentina's payment plan. Also conceivably endogenous. The rejection was filed after business hours on Friday, 19Apr13. This is also neither an event nor excluded.
26Aug13	Close-to-Close, 22Aug13 to 23Aug13	Appeals court upholds Griesa's decision.	22Aug13, 4:21pm Decision	23Aug13, 4:02pm Bloomberg	The appeals court announces decisions during the business day. The modification date of the PDF is 10:17am. However, because the

Table 5.5: Events and Non-Events (Continued)

Two-Day Window End	Event Type	Description	PDF Time (EST) and Link	News Time (EST) and Link	Comments
11Sep13	Non-Event	Supreme court schedules hearing to consider Argentina's appeal.	Missing Docket Info .	Bloomberg	The supreme court distributed case materials related to Argentina's petition. We were advised that this is routine and not "news," so we do not count is as a ruling.
26Sep13	Excluded	Holdouts had petitioned Griesa to consider the Argentine central bank liable for the defaulted debt. Argentina motioned to dismiss, and Griesa rejected Argentina's motion.	Missing		We were not able to find Griesa's ruling, so we exclude this event.
04Oct13	Open-to-Open, 03Oct12 to 04Oct13	Griesa bars Argentina from swapping the exchange bonds into Argentine-law bonds.	03Oct13, 2:43pm. Order	03Oct13, 6:27pm. Bloomberg	
08Oct13	Open-to-Close, 07Oct13	Supreme court denies Argentina's first petition.	N/A Order	07Oct13, 11:45am SCOTUS Blog	The stock market opens (9:30am EST) before the Supreme court issues decisions (9:30am or 10:00am EST).
19Nov13	Open-to-Open, 18Nov13 to 19Nov13	Appeals court denies Argentina's request for an en-banc hearing.	18Nov13, 11:04am Denial	19Nov13, 12:01am Bloomberg	The modification time of the orders is 4:53pm.

Table 5.5: Events and Non-Events (Continued)

Two-Day Window End	Event Type	Description	PDF Time (EST) and Link	News Time (EST) and Link	Comments
13Jan14	Open-to-Close, 10Jan14	Supreme court agrees to hear Argentina case.	10Jan14, 2:42pm Order	10Jan14, 2:48pm SCOTUS Blog	The supreme court usually announces orders at 10am. The document was likely posted afterward.
17Jun14	Open-to-Close, 16Jun14	Supreme court denies Argentina's second petition.	See Text. Denial	See Text. SCOTUS Blog	The text discusses this event in detail.
23Jun14	Open-to-Open, 20Jun14 to 23Jun14	Griesa prohibits debt swap of exchange bonds to Argentine law bonds.	20Jun14, 2:17pm Order		20Jun14 is an Argentine holiday, so the local stocks are missing. This event is excluded from our ADR analysis because of the two-day windows (it overlaps with the event below).
24Jun14	Open-to-Open, 23Jun14 to 24Jun14	Griesa appoints special master to oversee negotiations.	23Jun14, 12:36pm Order	23Jun14, 7:35pm Bloomberg	The modification date for the order is 1:05pm. This event is excluded from our local stock analysis because of the two-day event windows (see the event above).
27Jun14	Open-to-Close, 26Jun14	Griesa rejects Argentina's application for a stay, pending negotiations.	26Jun14, 11:40am Order	26Jun14, 2:05pm Bloomberg	
30Jun14	Non-Event	Argentina misses a coupon payment			
29Jul14	Open-to-Open, 28Jul14 to 29Jul14	Griesa allows Citi to pay Repsol bonds for one month.	28Jul14, 3:51pm Order	28Jul14, 12:01am Bloomberg	The modification time on the order is 5:07. This event almost certainly occurred post-close, but we use the one day window to be safe.
30Jul14		The 30-day grace period for the missed payment expires.		Bloomberg	We end our dataset on 29Jul14.