



Essays on Financial Economics and Macroeconomics

Citation

Agarwal, Ruchir. 2012. Essays on Financial Economics and Macroeconomics. Doctoral dissertation, Harvard University.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:9295156>

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](#)

©2012 — *Ruchir Agarwal*

All rights reserved.

Essays on Financial Economics and Macroeconomics

ABSTRACT

The first chapter studies mass layoff decisions. Firms in the SP 500 often announce layoffs within days of one another, despite the fact that the average SP 500 constituent announces layoffs once every 5 years. By contrast, similar-sized privately-held firms do not behave in this way. This paper provides a theoretical model and empirical evidence illustrating that such clustering behavior is largely due to CEOs managing their reputation in financial markets. The model's predictions are tested using two novel datasets of layoff announcements and actual mass layoffs. I compare the layoff behavior of publicly-listed and privately-held firms to estimate the impact of reputation-based incentives on cyclical behavior of layoffs. I find that relative to private firms, public firms are twice as likely to conduct mass layoffs in a recession month. In addition, I find that the firms that cluster layoff announcements at high frequencies are also the ones that are more likely to engage in mass layoffs during recessions. My findings suggest that reputation management is an important driver of layoff policies both at daily frequencies and over the business cycle, and can have significant macroeconomic consequences.

In the second chapter I present a theory of the safe assets market and make three central points. First, the quantity of safe assets has a strong influence on equilibrium risk premium and households' willingness to hold risky assets. Second, the banking

system and its regulation largely determine the quantity of safe assets (money-like claims) available to households. Lastly, by regulating banks' safe asset creation, central bank policy influences risk premium even in a flexible-price world. I show that the optimal central banking policy involves managing risk in the economy, which sometimes calls for large interventions.

The third chapter studies the asset allocation decisions of investors and central banks. This chapter identifies the fundamental drivers for these decisions and determines whether their influence has been altered by the global financial crisis and subsequent low interest rate environment in advanced economies.

The fourth chapter analyzes the welfare losses of taxation in a simple dynamic moral hazard model under symmetric information.

Contents

Acknowledgement	vii
Chapter 1: Strategic Corporate Layoffs	1
1 Introduction	1
2 Statistical Evidence of Clustering	7
3 Model	10
3.1 Setup	12
3.2 Reputation and Updating Rules	14
3.3 Equilibrium Selection	17
3.4 Equilibrium Implications	19
3.5 Extending the Model: Impact of a Large Firm	24
3.6 Mapping the Model to the Data	28
4 Data Construction	29
4.1 Constructing Dataset of Layoff Announcements and Firm Characteristics	29
4.2 Confidential Microdata from the Mass Layoff Statistics Program of BLS	32
5 Business Cycle Frequency Tests	34
5.1. Comparing Public-Private Firms: Empirical Strategy	35
5.2 Comparing Public-Private Firms: Results	39
5.3 Possible Alternate Explanations for the Difference in Public vs. Private Layoff Behavior	45
5.4 Variation within Public Firms	49
6 Daily Frequency Tests	52
6.1 Is the Reputation Penalty of Layoff Announcements Lower after Layoffs by Other Firms?	52
6.2 Do Firms Announce Layoffs after other Large Firm Layoff Announcements?	57
6.3 Which Type of Firms Layoff after a Large Firm Layoff Announcement?	63
7 Linking High Frequency Results to Business Cycle Frequency Results	68

8 Alternate Explanations of Layoff Behavior	70
8.1 Common Shocks	72
8.2 Compassionate CEOs	73
8.3 Market Inattention	74
8.4 Learning from Other Managers	77
8.5 Alternate Mechanisms: Taking Stock	78
9 Conclusion	79
Chapter 2: Uncertainty and the Special Role of Safe Assets	82
1 Introduction	82
2 Safe Assets in General Equilibrium	86
2.1 Setup	86
2.2 Frictionless Benchmark	90
2.4 Why Self Regulation Fails?	97
2.5 Central Bank Policy	102
2.6 Central Bank Policy with Nominal Contracts	107
3 Central Bank Policy with Shadow Banks	109
3.1 A Simple Example of Monetary Policy with Shadow Banks	109
4 Quantitative Effects and Implications for Asset Prices	112
4.1 Basic Asset Pricing	113
4.2 The Effect of Money on Risk Premium in a Safe World .	116
4.3 The Effect of Money in a World with Disaster Risk . . .	117
5 Extension: Shadow Banking Shut-Down and Unconventional Policy	123
Chapter 3: What Drives Global Asset Allocation Across Countries?	128
1 Introduction	128
2 Literature Review	131
3 Data and Methodology	133
3.1 Data	133
3.2 Methodology	134
5 Empirical Results	140

6 Currency Composition of Central Banks	147
6.1 Background	147
6.2 Defining Foreign Exchange Reserves and Sovereign Wealth Fund	150
6.3 Rise in Reserves and Change in Composition	153
6.4 Results	157
6 Concluding Remarks	160
Chapter 4: Dynamic Incentives and Welfare Costs of Taxation	163
1 Introduction	163
2 Model with Endogenous Explicit Incentives	167
2.1 Setup	167
2.2 Equilibrium	169
2.3 Welfare	172
3 Pure Career Concerns	174
3.1 Setup	174
3.2 Equilibrium	175
3.3 Welfare	177
4 Conclusion	178
Appendix	187
Appendix A.1: Updating rule for the <i>terminate</i> policy	187
Appendix A.2: Updating rule for the <i>delay</i> policy	188
Appendix A.3: Updating rule for a mixed-strategy policy	190
Appendix A.4: Proof for proposition 1	191
Appendix A.5: Proof of proposition 2	192
Appendix A.6: Proof of corollaries 3-5.	192
.....	193
Appendix A.7: Dynamic Version of the Model with Exogenous Updating Rule	193

Acknowledgement

This dissertation is the endpoint of an intense journey of learning about economics and finance. Looking back at the past six years, one realizes how much less fruitful and enjoyable this journey could have been had I not received the help of so many people around me.

First and foremost, my advisors — N. Gregory Mankiw, Jeremy Stein, Lawrence Katz, Emmanuel Farhi, and Alp Simsek — provided different and complementary perspectives on this dissertation, and most importantly provided warm support throughout my graduate school experience.

I am most fortunate to have had Greg as my advisor since the very beginning of the Ph.D. program. He guided me through the process of becoming an economist with great patience. Even before the program began, he hired me to work for him as a research assistant. Subsequently, for three years, I had the pleasure of being a Teaching Fellow for his famous year-long introductory economics course, “EC 10.” Such interactions with Greg, in addition to the almost-weekly guidance on my research, served as the most invaluable learning force for me during my doctoral studies. Through Greg’s guidance not only did I develop a solid framework for Macroeconomics, but also learned to become a thoughtful economist. As I step into the policy world, my interactions with Greg will serve as a constant valuable source of guidance.

Given my interest in Financial Economics, the support and advice of Jeremy has proven fundamental during the past years. His deep knowledge of the frontiers of Corporate Finance and intricate insights about the various policy responses to the Global Financial Crisis inspired me to con-

duct research in these fields. I am really indebted to him not only for the many insightful suggestions and comments on my papers, but also for the support and guidance that he patiently gave me over the years. I also had the pleasure of being a Teaching Fellow for his undergraduate course on 'Financial Crisis.' Jeremy's fantastic lectures and my slow walks with him back to Littauer, is where I learned most of my Finance.

I was very lucky to have been influenced by both Greg and Jeremy immediately after they returned from the White House. Their penchant for bringing clarity to the complexities of the policy world has been strongly impressed upon me and will guide my work for years to come.

As a researcher interested in empirical questions, I am very fortunate to have Larry as my advisor and mentor. My deep appreciation for data and measurement stems directly from my interactions with Larry. I will very fondly remember my frequent visits to Larry's office and the walks around the NBER with him and Pika. I am also greatly thankful to him for supporting me tirelessly in the process of obtaining confidential data on Mass Layoff Statistics from the Bureau of Labor Statistics. Larry taught me to how to become a meticulous researcher while simultaneously maintaining an ambitious research agenda.

With great fortune, my arrival at Harvard as a graduate student coincided with Emmanuel's arrival as an assistant professor. Given my interest in both Macroeconomics and Finance, the guidance and advice of Emmanuel has proven fundamental during the past years. His deep knowledge of the two fields enabled me to get involved in research that intersects both fields. My conversations with him during office hours and in reading groups have been most helpful in my attempts to understand the causes

and consequences of the Global Financial Crisis. I am also greatly thankful to him for his support during the intense job market months, when he was often available to provide encouragement—both in person and on Skype. During these years Emmanuel has been an invaluable advisor and a great friend.

I have known Alp for several years—first as a close friend when he was a graduate student at MIT, and then as the brilliant young assistant professor at Harvard. More than anyone else, Alp helped me uncover the deep economics hidden in my theoretical models. His ability to simplify complex ideas to its core essence is unmatched, and I very much hope that some of his skills have rubbed off on me as I prepare to embark on my professional life. Also, I am very thankful to him for patiently reading and re-reading my papers. Much of the clarity in my dissertation stems from his invaluable insights and unyielding support.

I also would like to thank my other co-authors of the chapters in this dissertation, namely Julian Kolev, Serkan Arslanalp, Tao Sun, and Michael Powell. Their contributions in the respective chapters are as much theirs as they are mine. I also want to thank others who have provided helpful feedback on parts of the thesis: Daron Acemoglu, Philippe Aghion, Alberto Alesina, Robert Barro, Effi Benmelech, Shai Bernstein, John Campbell, Patrick Carey, Josh Coval, Dan Fetter, Ben Friedman, Joan Farre-Mensa, Ed Glaeser, Robin Greenwood, Paul Goldsmith-Pinkham, Elhanan Helpman, Rustam Ibragimov, Andrew Karolyi, Laura Kodres, Sam Hanson, Oliver Hart, David Laibson, Ryan Michaels, Sendhil Mullainathan, Erik Oppers, Andrei Shleifer, Kelly Shue, and Marty Weitzman. I am also very grateful to the members of the Mass Layoff Statistics program at the Bureau of La-

bor Statistics for hosting me as a visiting researcher, to Kevin Hallock for sharing his data and insights with me, and to Lisa Ursino for hosting me at NBER for six long years.

The Ph.D. would have been a completely different experience, both academically and personally, had I not benefited from conversations and support of so many of my classmates, students and friends — in the program and outside — and in particular Fatin Abbas, Coren Apicella, Jonathan Beauchamp, Syon Bhanot, Dan Cao, Eliana Carranza, Michael Chesher, Jeff Clemens, Tom Cunningham, Mikkel Davies, Salvatore Dell’Erba, Aditya Dhanrajani, Lana Dinic, Nese Dogusan, Marco Di Maggio, Michal Fabinger, Janling Fu, Andreas Fuster, Tim Ganser, Patrick Gaule, Eduardo Giacomazzi, Maria Giavazzi, Stefano Giglio, Ross Hazelett, Scott Hirst, Fernando Innechi, Luigi Iovino, Huseyin Kadikoy, Martin Kanz, Tim Khoury, Jan Klingelhofer, Jo Lim, Shih En Lu, Hiroaki Matsuura, David Mericle, David Molitor, Eduardo Morales, Arash Nekoei, Francisco Queiro, Amul Sathe, Tom Sampson, Roberta Scarpato, Tiago Severo, Danny Shoag, Kelly Shue, Holger Spamann, Andrea Stella, Ija Trapeznikova, Tom Vogl, Charles-Henri Weymuller, Andres Zahler, and Clara Zverina.

My students also deserve a special mention since they made my Ph.D. experience thoroughly exciting. Teaching economics at Harvard has been one of the most rewarding experiences for me, and along the way it allowed to meet some very talented and motivated undergraduates.

Finally, and most importantly, I am incredibly grateful to my parents Maya and Vinod, my sister Paridhi, and my grandparents Kishan Lal and Gita Devi, who have always provided their full support and advice, and though distant were always right there for me.

Chapter 1: Strategic Corporate

Layoffs

1 Introduction

Voluntary disclosures of bad news by firms are often immediately followed by similar disclosures by other firms. Clustering of bad news is observed in the release of negative earnings announcements, write-downs, or layoff announcements.¹ This paper focuses on layoffs and investigates the mechanisms that influence managers' decisions to cluster their layoff announcements. It also studies the aggregate implications of such decisions, which is relevant for welfare since the timing and quantity of layoffs is tightly linked to unemployment dynamics.

As a case study of clustering in layoff announcements, we consider the behavior of the top three US firms in the banking industry (Bank of America, J.P. Morgan and Citigroup) and the automobile manufacturing industry (G.M., Ford and Chrysler) around the 2001 recession. Figure 1.1 represents the timelines of layoff announcements for each of these firms from 2000 to 2003. Layoff announcements tend to be clustered within industry: in many cases we observe announcements within the same week. We also observe clustering of announcements across industries.² To further investi-

¹Acharya, DeMarzo, and Kremer (2011) present a model based on asymmetric information that predicts such behavior. Tse and Tucker (2010) provide evidence of clustering of bad news in the form of earnings warnings. Our paper is the first to document herding in layoff announcements.

²In the Fall of 2002, Chrysler and Ford announced layoffs on the same day. Similarly, in the Fall of 2003, J.P. Morgan and G.M. announced their layoffs in the same week, and days later Bank of America announced a layoff.

gate the degree of clustering, we apply standard temporal clustering tests to our full sample from 1970 to 2010, and find statistically significant evidence for excess clustering in the time series of layoff announcements. Further, we show that such clustering behavior is observed only in publicly-traded firms (“public” firms), and not in comparable privately-held firms (“private” firms). Motivated by these facts, the central question in this paper is: why do public firms engage in clustering of layoffs, while private firms do not; and what are the aggregate implications of such behavior?

We interpret the observed degree of clustering and the differences between public and private firms through a model based on asymmetric information between managers of firms and the financial market. The central mechanism of the model is as follows. The market perceives a layoff announcement as a negative signal about the manager’s ability. When aggregate business conditions are adverse, such as during recessions or industry-wide downturns, the market will attribute greater blame to external factors than to managerial ability. This generates incentives for managers to time their layoff announcements to occur during downturns, thereby minimizing the blame for the bad news. This key idea of our paper is a counterpart to the early paper by Gibbons and Katz (1991), who provide evidence that workers laid off on a discretionary basis are viewed less favorably by the market than are those losing jobs in plant closings. We invoke the same Gibbons-Katz mechanism to illustrate how the cyclical nature of layoffs is linked to the lower reputation penalty that managers face in recessions.

The model has two main cross-sectional predictions. First, if managers care more about their reputation (relative to the cash flows of the firm), then they are more likely to engage in layoffs during adverse states. Second, the

model also predicts that firms with managers who don't have a long-track record are more likely to engage in layoffs during adverse states. This is because the market's perception of their ability is more sensitive to new information. These predictions of our model apply to both business cycle and daily frequencies. The business cycle frequency results predict differential layoff strategies in recessions, while the daily frequency results are associated with differential propensity in clustering of layoff announcements. We test both mechanisms in turn using two novel datasets.

The first dataset consists of layoff announcements by the largest publicly-listed firms (Fortune 500 constituents) and largest privately-held firms (Forbes 100 constituents), collected from daily issues of the Wall Street Journal between 1970 and 2010. The second dataset, which we are the first researchers to access, consists of confidential microdata on actual mass layoffs from the Bureau of Labor Statistics under their Mass Layoff Statistics program. Since this program collects data from unemployment insurance (UI) claims, it allows us to observe the timing and the exact number of displacements arising from mass layoffs.³

To estimate the impact of reputation management on layoff propensity at the business cycle frequency, we focus on differences between public and private firm behavior. This analysis is motivated by certain fundamental differences between public and private firms, which make the public firm managers relatively more likely to manage their reputation in financial markets. First, since public firms sell shares to outside investors who are not involved in managing the firm, there exists separation between ownership and control. Therefore, since the managers of public firms do

³Because it looks at actual displacements, the BLS data is not subject to reporting bias.

not completely internalize the costs of adopting inefficient policies, they are more likely to prioritize reputation management over maximizing firm value. Second, owners of public firms typically have shorter investment horizons: the increased liquidity of public markets makes it easy for shareholders to sell their stock at the first sign of trouble rather than actively monitoring management. This leads to relatively myopic behavior among investors of public firms, which weakens incentives for effective corporate governance (Amar (1993)) and generates incentives for managers to engage in myopic reputation management (Stein (1989)). Third, managers of public firms are subject to takeover threats, which depend on the stock price of targeted firms. This can lead to managerial myopia in public firms in order to actively manage current stock prices (Stein (1988) and Edmans, Goldstein, and Jiang (2011)). Because reputation management focuses on boosting perceptions of short-term performance at the expense of long-run value, myopic managers are more likely to engage in the strategic behavior predicted by our model. Taken together these differences imply that if reputation management drives layoff behavior over the business cycle, the effect should show up as differences in behavior of comparable public and private firms. We emphasize reputation management in *financial markets*, because the managers of both public and private firms are likely to have similar motivations for reputation management among other constituents.⁴

Using a pairwise matching estimator based on size and three-digit industry, we find that the layoff propensity of public firms is twice as sensitive to recessions, relative to their matched private counterparts.⁵ In a range of

⁴Notable recent work that explores public-private differences are Asker, Farre-Mensa, and Ljungqvist (2011) and Davis, Haltiwanger, and Jarmin (2006)

⁵That is, in recession months, the propensity to layoff for public firms increases by roughly 5 percentage points, whereas that of private firms increases by roughly 2.5 per-

robustness tests, we show that these differences are not driven by public-private differences in lifecycle effects, leverage, workforce size, or on the criteria we use to match public and private firms. Within our sample of public firms, we find that firms predicted to be more strategic by our theoretical analysis, are also the ones more likely to engage in mass layoffs during recessions. Our results therefore suggest that reputation management is an important driver of the observed differences in the cyclicalities of layoffs between public and private firms.

Next, we test our model at the daily frequency, and find further support for its predictions. We show that a large firm announcement (i.e. the 20 largest firms based on past year's revenue) is associated with future layoffs by other Fortune 500 firms, but not with past layoffs.⁶ We find that this effect is twice as strong if the large firm is in the same industry as the follower firm. For our sample of privately held firms we find no such clustering behavior either before or after the large-firm layoff announcement. Moreover, when we compare the characteristics of firms that lay off in the five days after a large-firm announcement ("followers") to those that lay off in the five days before a large-firm announcement ("counterfactual followers"), we find that follower firms have a greater likelihood to be managed by short-tenured CEOs (i.e. with a tenure between 0 and 4 years) and to place greater reliance on equity-linked compensation for their CEOs. Consistent with our theoretical framework, these results suggest that reputation management is an important driver for the timing of layoff announcements at high frequencies.

centage points.

⁶Our results are related to the work of Gabaix (2011), who examines the role of large firms in explaining aggregate fluctuations.

Lastly, we establish a connection between high-frequency clustering and layoff behavior at business cycle frequencies. We find that the follower firms are roughly 3 percentage points more likely to engage in a mass layoff during a recession month, compared to counterfactual follower firms. This link between the daily frequency reputation management and the cyclical-ity of layoffs over the business cycle provides significant evidence that reputation concerns are an important driver of firms' layoff policies.

To rule out the role of alternate theories in driving our results, we conduct a series of additional tests. The key alternate explanations we consider are: common shocks, learning from other managers, compassionate CEOs, and market inattention. While each of these mechanisms may contribute to some of the patterns we observe in layoff behavior, no combination of these effects can explain the full range of our results. Taken together, the findings of this paper suggest that managerial behavior not only has costs for the individual firm, but also has significant aggregate implications at business cycle frequencies.

The remainder of the paper is structured as follows: Section 2 presents statistical tests for detecting excess clustering in layoff announcements. Section 3 presents the theoretical model. In Section 4 we describe the construction of our two main datasets. Section 5 presents the empirical methodology and business-cycle frequency tests of the model, while Section 6 presents daily-frequency tests. In Section 7 we link the daily-frequency results to the business-cycle frequency results. Section 8 discusses alternate explanations of layoff behavior, and Section 9 concludes.




Figure 1.1: Timeline of layoff announcements by top 3 firms in the banking and automotive industry between 2001 and 2003. A dot represents a layoff announcement on a given day.

2 Statistical Evidence of Clustering

As illustrated in the case study (Figure 1.1), we observe that layoff announcements are often clustered within days of each other. The size of these clusters ranges from within a day to over two weeks. Though this is suggestive evidence of excess clustering, it is not clear whether this observation represents a general trend which applies to other periods and other firms in the economy. In this section, we take a systematic approach to identify and characterize the nature of excess clustering in firms' layoff announcements.

Our approach uses a measure called the scan statistic, which is used to

detect unusual clusters in a sequence of events that occur over time. The approach is known as “moving window analysis” in the engineering literature (see Glaz, Naus, and Wallenstein (2001)). To see how it works, consider N events that occur on an unit interval. First, consider the number of occurrences in each window of size w . Then consider the maximum of these over all windows of size w in the unit interval. Under the null of a uniform distribution, the distribution of this maximum can be calculated. For a confidence level, e.g. $\alpha = 0.05$, we can then construct a critical value, c_α , such that $\Pr[S_w > c_\alpha] = \alpha$. Here S_w denotes the largest number of events to be found in any subinterval of $[0, 1]$ of length w , and is called the scan statistic. If the maximum observed local statistic, S_w , is larger than or equal to c_α , then we should reject the null hypothesis and infer existence of a local region with a statistically significant cluster.

The distribution of scan statistic described above is a function of two parameters: the size of the subwindow, w (relative to the size of the entire interval); and the number of events, N , which occur in the entire interval $[0, 1]$. We denote the p-value of this test as $\Pr[k; N, w]$. This p-value should be interpreted as follows: *under the null of N events independently drawn from the uniform distribution on $[0, 1]$, $\Pr[k; N, w]$ is the probability that we observe k or more events in any subwindow of size w .*⁷

The unit of time for our tests is business days. We conduct our tests for two different interval sizes: 60 business days (approximately one quar-

⁷Under the null hypothesis, the probability of observing a scan statistic, S_w , greater than k , can be characterized as a function of the two parameters:

$$\Pr[S_w \geq k] \equiv \Pr[k; N, w]$$

Exact estimates of this common probability exists for certain cases, and researchers have to rely on approximations for the other cases. See chapter 10 of Glaz, Naus, and Wallenstein (2001).

ter) and 20 business days (approximately one month). Our sample period begins in 1970 and ends in 2010, and we run separate tests for every non-overlapping window in this period of 41 years. Since the test has low power when N is small, we exclude months in which we observe fewer than 5 lay-off announcements.⁸ Also, we run the tests for two categories: all industries combined; and the manufacturing industry. This allows us to assess whether there is excess clustering at both the aggregate and the industry level.

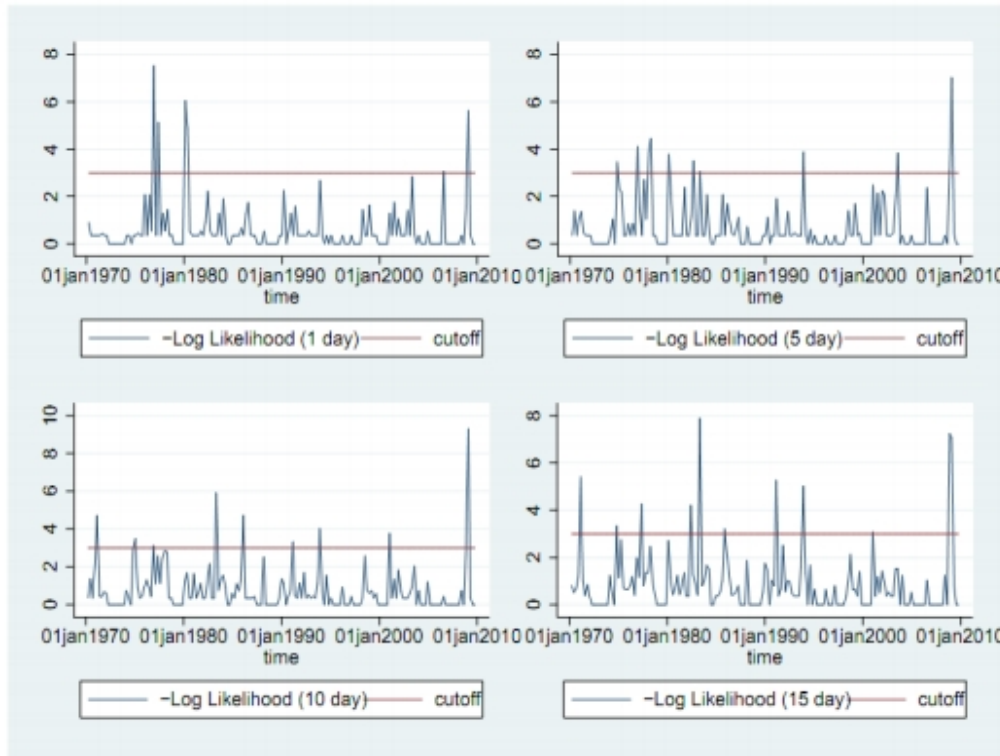


Figure 1.2: This graph reports the log p-values (multiplied by negative unity) of the sequence of tests for each non-overlapping 60 (business) day window to identify clustering at 1, 5, 10 or 15 day horizons when the null is that the layoff announcements are distributed uniformly over each 60 day window.

⁸Our results are do not change when we include all the months in our analysis.

Figure 1.2 plots the results of our analysis for each non-overlapping interval of 60 business-days. To facilitate viewing, we present the negative log of the p-value as our y-axis variable. Therefore, a higher value suggests that we can reject the null with greater confidence. Though this sequence of individual tests is suggestive of several episodes of clustering, we would like to combine the results from these different independent observations into a single statistical test. For this, we rely on Fisher’s method to combine the p-values from our tests into a single statistic, using the formula $X^2 = -2 \sum_{i=1}^k \log(p_i)$, where p_i is the p-value for the i-th independent test. This ‘combined p-value’ is reported in the last column of Table 1.1. The main conclusion of this analysis is that we can reject the null of no excess clustering for subwindows of 5 or more days using one-month intervals, and for subwindows as small as three days using quarterly intervals. Having established the existence of excess clustering, we proceed to offer a potential explanation for this phenomenon in Section 3, where we present our theoretical model.

3 Model

In this section, we present a reputation-based model of management layoff decisions, focusing on the tradeoff between firm profits and the perception of managerial talent. A similar model was presented in Rajan (1994), which studies the clustering of credit policies by banks. We focus on a three-period version of the model in the main text, incorporating fully rational Bayesian expectations and solving for the set of trembling-hand perfect equilibria. We discuss the implications of relaxing these assumptions and

Table 1.1: Fisher's Method for Combining Results from Independent Scan Statistics

This table reports our results for Fisher's Method, which is a method to combine results from the independent Scan Statistics we compute for each non-overlapping interval of 20 or 60 business-days between 1970 and 2010. The top half of the table conducts this analysis for all layoffs in our sample, while the bottom half restricts the sample to layoffs in the 3 digit NAICS industries. The null hypothesis is to assume that there is uniform layoff propensity for each 20 or 60 day window. The subwindow lists the length of window under consideration for excess clustering. The test statistic is a combined p-value of all individual tests, and is distributed with a chi-squared distribution. The degrees of freedom is simply twice the number of individual tests, which is given in the third column. The last column lists the 'combined p-value' from Fisher's method. If this p-value is below 0.05 then we can reject the null of no excess clustering for the given subwindow.

Window (days)	Sub-window (days)	# of months	Test Statistic (Fisher's Mehtod)	Degrees of Freedom	p-value (Fisher's Mehtod)
All Firms					
20	1	325	507.99	650	1.000
20	3	325	626.82	650	0.736
20	5	325	708.67	650	0.055
20	10	325	781.41	650	0.0003
60	1	109	202.26	218	0.771
60	3	109	255.70	218	0.041
60	5	109	254.52	218	0.045
60	10	109	268.86	218	0.011
60	15	109	288.03	218	0.001
3-digit Industry Firms					
20	1	252	374.27	504	1.000
20	3	252	437.20	504	0.986
20	5	252	551.12	504	0.072
20	10	252	631.32	504	0.0001
60	1	83	181.26	166	0.198
60	3	83	192.81	166	0.076
60	5	83	198.53	166	0.043
60	10	83	213.82	166	0.007
60	15	83	212.48	166	0.009

expand to multiple periods in the theoretical appendix.

3.1 Setup

Our model starts at date 0, and ends at date 2. There is no discounting between periods. There are two types of agents in this model: firm managers and the market.⁹ Managers care about the profits of their firm, and about their reputation with the market. The market takes no direct action in this model, but simply observes the actions of managers, and updates its priors according to Bayes' Rule.

Firm managers are the primary decision-making agents in this model. Each manager i is associated with a firm which, at date 0, begins a new project and hires one unit of labor to engage in production. There is continuum of manager types, which differ along the dimension of managerial talent, denoted by η_i . The only restriction on the distribution of talent, described by density $f(\eta_i)$, is that its support be within the unit interval, $[0,1]$. For convenience, we define the mean and variance of the distribution of talent to be μ and σ_η^2 , respectively.

After the project is undertaken, the aggregate economic state is realized at date 1.¹⁰ The aggregate state is denoted by $s \in S = \{N, A\}$: it can be adverse (A) with probability π , or normal (N) with probability $1 - \pi$.

The probability that a project succeeds depends on both the talent of

⁹There are many possible interpretations for the role of the market. One possibility is the population of equity market investors - this interpretation links our model to concerns about stock price responses to layoffs. Another option is the demand side of the market for managerial talent - this interpretation is more in line with the literature on career concerns (Holmström (1999)). We do not pin down a specific interpretation in order to allow for the broadest possible application of the model.

¹⁰We can interpret the aggregate economic state as either an economy-wide indicator, or a measure of the health of a particular sector.

the manager and the aggregate state, and is given by:

$$\theta_s^i = \eta_i \lambda_s \quad (1)$$

such that $\lambda_N = 1$, and $\lambda_A = 1 - \delta$. In the adverse state the probability of project success is $\eta_i(1 - \delta)$, and in normal states it is η_i . There is symmetric uncertainty about the aggregate state for both managers and the market throughout all time periods.¹¹

The manager privately observes the outcome of his firm's project at date 1. If the project was successful, there is no decision to make: the project continues into date 2, where it generates earnings of \$1, and then ends. If the project is not successful, the manager has to decide whether to terminate or continue the project. Termination involves firing the labor force hired at time 1, and is therefore fully observable. The firm's date 2 earnings are zero if it terminates an unsuccessful project at date 1. We label this approach as the "*terminate*" policy.

Instead of termination, the firm can hide the unsuccessful outcome of the project from the market by not laying off the labor force assigned to the project. If the manager adopts such a policy, he must pay the worker for one more period even though the worker will not be productive. We denote this cost as C . Relative to the *terminate* policy, this decision delays the end of the project by one period. We therefore label this approach as the "*delay*" policy.

We assume that adopting the *delay* policy is costly relative to the deci-

¹¹The results of this model do not rely on the particular functional form assumed here for the probability of success. The key comparative statics are identical if instead of a multiplicative function we assume an additive function: $\theta_s^i = \alpha \eta_i + (1 - \alpha) \lambda_s$, for $0 < \alpha < 1$.

sion to *terminate*. This is given by *Assumption 1*: $C > 0$. This assumption implies that in a first-best world, the *delay* policy should not be adopted.

Despite its inefficiency, managers have an incentive to adopt the *delay* policy because it is better for their reputation to hide an unsuccessful project outcome. In setting up the maximization problem for managers, let $D \in \{0, 1\}$ represent whether or not a given manager adopts the *delay* policy when his firm's project fails. We can then describe managers' preferences by the following utility function:

$$\max_D U_i = -DC + \gamma E_{mkt} [\eta_i | \hat{D}, \text{Layoffs}] \quad (2)$$

such that γ is the utility weight the management places on his reputation in the eyes of the market, and \hat{D} is the conjecture of the manager's strategy that the market uses to interpret the observation of layoffs or no layoffs.

3.2 Reputation and Updating Rules

The market's updating rule depends on 2 factors: a) its conjecture about the manager's strategy in addressing a failed project; and b) whether or not it observes layoffs. We begin our equilibrium analysis by focusing on four primary cases:

	No Layoffs	Layoffs
Conjecture: 'Terminate'	$E_{mkt} [\eta_i \hat{D} = 0, L = 0]$	$E_{mkt} [\eta_i \hat{D} = 0, L = 1]$
Conjecture: 'Delay'	$E_{mkt} [\eta_i \hat{D} = 1, L = 0]$	$E_{mkt} [\eta_i \hat{D} = 1, L = 1]$

When the market conjectures that the firm will adopt the *Terminate* policy

If the market believes the firm is going to adopt the *terminate* policy, i.e. $\hat{D} = 0$, then the firm will lay off the project's workers upon failure, and not lay off workers when the project succeeds. Therefore, not observing layoffs implies that the firm's project has succeeded. This makes the updating rule straightforward: the observation of layoffs or no layoffs is perfectly correlated with the outcome of the project.

Using Bayes Rule we can calculate the resulting posteriors as:

$$E_{mkt} [\eta_i | \hat{D} = 0, \text{Layoffs} = 0] = E_{mkt} [\eta_i | \text{Project Succeeds}] = \mu + \frac{\sigma_\eta^2}{\mu} \quad (3)$$

$$E_{mkt} [\eta_i | \hat{D} = 0, \text{Layoffs} = 1] = E_{mkt} [\eta_i | \text{Project Fails}] = \mu - \frac{\sigma_\eta^2 (1 - \pi\delta)}{1 - \mu(1 - \pi\delta)}$$

(Proof in Appendix A.1)

To interpret these results, note that the market's prior about managerial talent is given by μ . When the market observes no layoffs, they update their beliefs about managerial talent positively, which is reflected by the positive additive term in the first equation. Analogously, when the market observes a layoff, then they update their beliefs negatively, which is reflected by the negative additive term in the second equation. The weight of the additive terms depend positively on σ_η^2 since it measures how noisy was the market's prior was at the start of the period. The second equation also depends negatively on the probability of being in an adverse aggregate state (π). Therefore, when π is high, the reputation penalty of laying off is lower.

When the market conjectures that the firm will adopt the *Delay* policy

We get an analogous updating rule for the case in which the market believes the firm is going to adopt the *delay* policy, i.e. $\hat{D} = 1$. Under standard equilibrium assumptions, the outcome of layoffs under this policy will

never occur, and the Bayesian posterior would not be uniquely determined. We therefore introduce trembles and focus on trembling-hand-perfection as our equilibrium concept. Because successful projects continue automatically, they do not require any action from the manager and are not susceptible to trembles. By contrast, the decision to adopt the *delay* policy requires a direct action by the manager, who could tremble and choose to terminate the project instead. Therefore, whenever the firm engages in layoffs the market knows the project must have failed, even though this outcome will (almost) never be observed in equilibrium. The updating rule in this situation can be calculated as follows:

$$E_{mkt} [\eta_i | \hat{D} = 1, \text{Layoffs} = 1] = E_{mkt} [\eta_i | \text{Project Fails}] = \mu - \frac{\sigma_\eta^2 (1 - \pi\delta)}{1 - \mu (1 - \pi\delta)} \quad (4)$$

Once again when the market observes a layoff they update their beliefs about managerial talent negatively. By contrast, when the market observes no layoffs, they do not know whether the project failed or not. This is because the firm is expected to adopt the *delay* policy of no layoffs irrespective of project outcomes. Therefore, when the market observes no layoffs, they get no new information about the firm, and the updating rule is simply:

$$E_{mkt} [\eta_i | \hat{D} = 1, \text{Layoffs} = 0] = E_0 [\eta_i] = \mu \quad (5)$$

(Proof in Appendix A.2)

3.3 Equilibrium Selection

In equilibrium the market conjecture about the manager's policy must be correct, and hence $\hat{D} = D_i$.

To support the equilibrium where the manager always adopts the *terminate* policy, the following incentive compatibility condition must hold:

$$\gamma E_{mkt} [\eta_i | \hat{D} = 0, \text{Layoffs} = 1] \geq -C + \gamma E_{mkt} [\eta_i | \hat{D} = 0, \text{Layoffs} = 0] \quad (6)$$

By contrast, to support the equilibrium where the manager always adopts the *delay* policy, the IC constraint is:

$$\gamma E_{mkt} [\eta_i | \hat{D} = 1, \text{Layoffs} = 1] \leq -C + \gamma E_{mkt} [\eta_i | \hat{D} = 1, \text{Layoffs} = 0] \quad (7)$$

Using the Bayesian analysis in the previous section, the above constraints, respectively, simplify to:

$$C \geq \gamma \sigma_\eta^2 \left[\frac{1}{\mu} + \frac{1 - \pi\delta}{1 - \mu(1 - \pi\delta)} \right] \quad (8)$$

and

$$C \leq \gamma \sigma_\eta^2 \left[\frac{1 - \pi\delta}{1 - \mu(1 - \pi\delta)} \right] \quad (9)$$

From the above constraints, it is clear that for sufficiently high values of γ and σ_η^2 , managers will choose to always adopt the *delay* policy. At the same time, for sufficiently low values of these variables, managers will always choose to adopt the *terminate* policy. Having already characterized these equilibria, we now move to consider the intermediate set of parameter values, which support neither pure-strategy equilibrium.

For intermediate values of γ and σ_η^2 , the equilibrium reputation penalty under the *terminate* policy is so large that managers prefer the *delay* policy, and the equilibrium reputation penalty under the *delay* policy is so small that they prefer the *terminate* policy. This means that for these parameter values, there is no equilibrium in pure strategies. We therefore proceed to analyze a mixed-strategy equilibrium, where managers randomize between adopting the *terminate* and *delay* policies.

As in the previous section, we begin our characterization of the mixed-strategy equilibrium by focusing on the market's posterior following an observation of either layoffs or no layoffs. In this case, instead of a binary conjecture about the policy of the manager, we move to a continuous conjecture $\hat{D} \in (0,1)$ which corresponds to the probability with which the market expects the manager to choose the *delay* policy, conditional on project failure. In such a setting, we calculate the market's posteriors as:

$$E_{mkt} [\eta_i | \hat{D}, \text{Layoffs} = 0] = \mu + \frac{(1 - \hat{D})(1 - \pi\delta)\sigma_\eta^2}{\hat{D} + (1 - \hat{D})(1 - \pi\delta)\mu} \quad (10)$$

and

$$E_{mkt} [\eta_i | \hat{D}, \text{Layoffs} = 1] = \mu - \frac{\sigma_\eta^2(1 - \pi\delta)}{1 - \mu(1 - \pi\delta)}$$

(Proof in Appendix A.3)

Note that the above posteriors match up with pure-strategy beliefs when we take the limits as $\hat{D} \rightarrow 0$ and $\hat{D} \rightarrow 1$ for the cases of the *terminate* and *delay* policies, respectively.

To complete the characterization of the mixed-strategy equilibrium, we need the manager to be indifferent between the two strategies available to him. Using the same IC constraint framework as in the pure-strategy case,

we need:

$$\gamma E_{mkt} [\eta_i | \hat{D}, \text{Layoffs} = 1] = -C + \gamma E_{mkt} [\eta_i | \hat{D}, \text{Layoffs} = 0] \quad (11)$$

Using the posterior market beliefs for the mixed-strategy case, the above constraint simplifies to:

$$C = \gamma \sigma_\eta^2 \left[\frac{(1 - \hat{D})(1 - \pi\delta)}{\hat{D} + (1 - \hat{D})(1 - \pi\delta)\mu} + \frac{1 - \pi\delta}{1 - \mu(1 - \pi\delta)} \right] \quad (12)$$

The equation above allows us to solve for the manager's randomization probabilities in the mixed-strategy equilibrium. Note that because the RHS is monotonically decreasing in \hat{D} , there is a unique set of mixing probabilities that supports equilibrium play for any given set of parameter values. Further, because the limits of the expression match up to the pure-strategy equilibria described in the previous section, there is a continuous progression from always choosing the *terminate* policy, through a mix of both options, and finally to always choosing the *delay* policy, as the product of γ and σ_η^2 increases from zero.

3.4 Equilibrium Implications

The results of the previous section describe conditions under which the firm will undertake the *delay* policy, despite the inefficient reduction in earnings that result from it. To gain an insight into these central results, Figure 1.3 plots the equilibrium policies of managers based on their values of γ (degree of reputational concerns) and σ_η^2 (variance of market's prior about firm). Managers who adopt the *delay* policy will lie above and to the right of mixed-strategy region. In this region, managers and their firms will

delay project termination and avoid layoffs, despite their project failing.

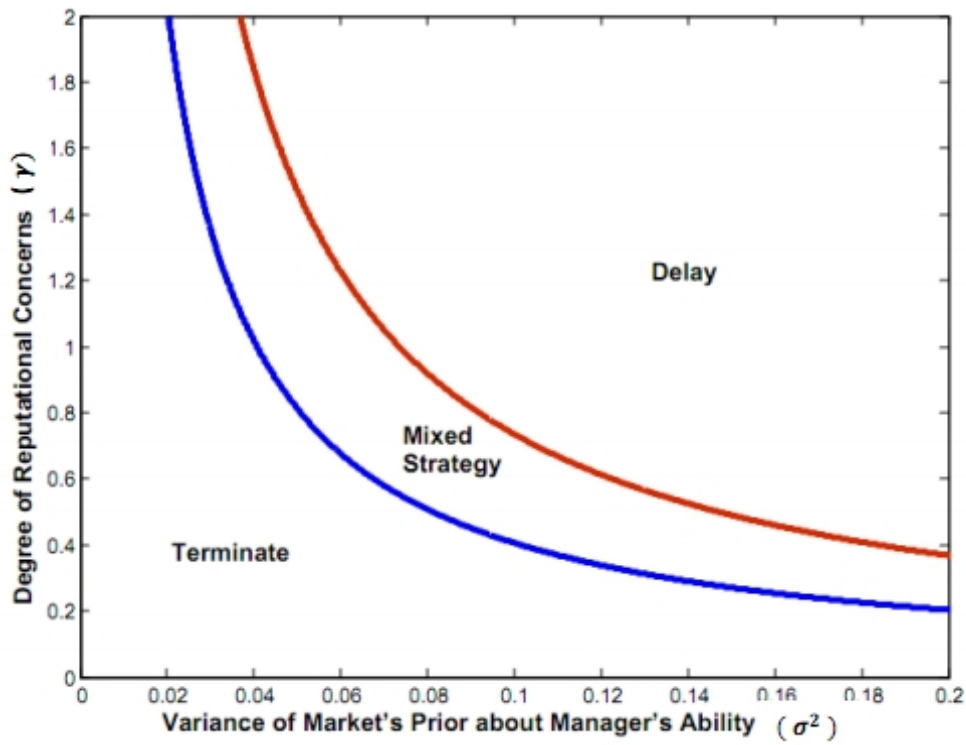


Figure 1.3: Parameter Space and Equilibrium Layoff Policy

Based on this analysis, we can conclude that the firm has an incentive to undertake the *delay* policy when:

- the manager places a high weight on reputation (as measured by γ)
- layoffs are particularly informative about the manager's ability, due to significant uncertainty in the market's prior beliefs (i.e. a high value of σ_η^2).

The above implications have direct links to observable variables in empirical corporate finance. A high value of γ is likely to be associated with firms that incentivize their management with high-powered, market-based

compensation packages. As for the informativeness of layoffs, a high value of σ_η^2 (a sufficient statistic for the signal-to-noise ratio in our model) is likely to be associated with firms that have a new management, precisely because the market will have less information about them, and any action taken by them will be relatively more informative.

An interesting implication of the model is the effect of changing beliefs about the aggregate state S , where the expectation of an adverse state is measured by π . Figure 1.4 plots the same boundaries as Figure 1.3, and adds another set of boundaries to demonstrate the effect of the market's perception about aggregate state becoming pessimistic. Assuming that this perception is justified, there will be a direct effect of fewer successful projects in an adverse economic environment. Because of this, the market is less likely to attribute the negative signal of a layoff to the manager's level of talent, and consequently the reputational concern associated with layoffs diminishes. This is illustrated by the rightward shift of the boundaries in Figure 1.4. As a result, the parameter region for which firms will adopt the *delay* policy shrinks. More firms now will choose to announce layoffs if their project fails.




Figure 1.4: Comparative Statics — when belief about aggregate state worsens.

In addition to the direct effect of adverse economic conditions leading to higher rates of project failure, our model also predicts a shift-in-equilibrium effect: conditional on project failure, a larger fraction of managers will choose to terminate their projects and engage in layoffs during these economic downturns.¹² In the pure-strategy regions of the parameter space, there is no shift in equilibrium because managers are effectively at a corner solution. Those who strictly prefer the *delay* policy will continue

¹²In the multi-period model discussed in the appendix, we show that firms with failed projects are likely to continue them until the next economic downturn, effectively saving their layoffs until they can implement them without suffering the normal reputational penalty.

to have a layoff rate of precisely zero, while those who strictly prefer the *terminate* policy will have a layoff rate equal to their project failure rate: $1 - \eta(1 - \pi\delta)$, or $1 - \mu(1 - \pi\delta)$ on average. The interesting case is that of managers in the mixed-strategy region. We describe their layoff rate in the following proposition:¹³

PROPOSITION 1. For managers in the mixed-strategy region of the parameter space, an increase in the expected probability of a downturn leads to a strictly higher rate of layoffs. Specifically:

$$\frac{\partial \Pr[\text{Layoffs} \mid \pi, \gamma, \sigma_\eta^2; C, \delta]}{\partial \pi} = \frac{\delta \gamma \sigma_\eta^2}{C [1 - \mu(1 - \pi\delta)]^2} > 0$$

PROOF. See appendix A.4. ■

From this result, it follows that managers who care most about reputation (i.e. have a high value of γ) and who have a short track record (i.e. a high value of σ_η^2) are most likely to be affected by the adverse shift in market's perception of the aggregate state. These insights are summarized in the following two corollaries:

COROLLARY 2. If the market's belief about a firm's management is less precise, then the manager is more likely to announce layoffs in downturns. That is, $\frac{\partial^2 \Pr[\text{Layoff} \mid \pi, \gamma, \sigma_\eta^2; C, \delta]}{\partial \pi \partial \sigma_\eta^2} \geq 0$.

COROLLARY 3. If the manager's utility function puts more weight on his reputation, then the manager is more likely to announce layoffs in downturns. That is, $\frac{\partial^2 \Pr[\text{Layoff} \mid \pi, \gamma, \sigma_\eta^2; C, \delta]}{\partial \pi \partial \gamma} \geq 0$.

¹³While Proposition 1 and the two corollaries below and the rest of the analysis focus on changes in the probability of experiencing an adverse aggregate state, similar results obtain when considering an increase in the severity of the adverse state, represented by the magnitude of δ .

3.5 Extending the Model: Impact of a Large Firm

In this section we extend the model presented above by introducing a large firm into the economy. The notion of large here is that the firm's performance contains information about the aggregate state of the industry. By contrast, the performance of small firms is heavily influenced by conditions in their local market, so the ability to obtain information about the aggregate state of the industry from the performance of a small firm is assumed to be negligible. As a result, the large firm's layoffs decision will influence the other firms in the industry through the information it provides about the aggregate state about the industry. As shown in Propositions 1 and 2, the market's beliefs about the aggregate state have a strong influence on firms' layoff decisions. With the addition of a large firm, the model therefore generates strategic interaction between firms.

In formulating this extension of the model, we move from a general-purpose metric of beliefs about the aggregate state, π , to a firm-specific metric represented by π_i . Here it's more appropriate to interpret π_i as the probability of an adverse economic state in the industry or local market of small firm i . Specifically, it combines the outcome of the aggregate state s_{agg} with firm-specific conditions described by ε_i :

$$\pi_i = f(1 \cdot (s_{agg} = A) + \varepsilon_i)$$

and we restrict $f(\cdot)$ to be a monotonically increasing function with a support equal to the interval $[0,1]$.¹⁴ For simplicity, we assume that the project outcome for the large firm is directly dependent on s_{agg} as before, with the

¹⁴While we do not specify a functional form, the most widely-used options include the logistic function and the probit, or normal quantile function.

prior probability of an adverse aggregate state measured by π_0 . For the small firm, an adverse aggregate state increases the chances that the firm-specific economic state will also be adverse, but does not predict this perfectly. We maintain the same framework of observability as before, where prior distributions are common knowledge, but only the layoff decision is observed by the market.

The market updates the smaller firms' reputation in two steps. First, it updates its prior on the realization of s_{agg} , using the large firm's layoff decision. This, in turn, leads to an updated belief about π_i for the small firm, in turn impacting the reputational penalty the small firm would face if it announced layoffs of its own.

Using Bayes' Rule, the process of updating expectations about s_{agg} using the layoff decision of the large firm is straightforward. Letting π_0 be the prior expectation of the adverse state, the posterior expectation conditional on observing layoffs by the large firm, $\pi_1 \equiv \Pr[s_{agg} = A | \text{Large Firm Layoff} = 1]$ is given by:

$$\pi_1 = \pi_0 \left(\frac{1 - \eta(1 - \delta)}{1 - \eta(1 - \pi_0\delta)} \right) > \pi_0 \quad (13)$$

Note that the above updating rule is true for all possible strategies employed by the large firm, as long as we allow for trembles in the case where the large firm would like to always choose the *delay* policy. Moreover, while the market does not know the value of η for the large firm, taking expectations over any prior distribution leads to the same conclusion: $\pi_1 > \pi_0$. Thus, whenever the market observes layoffs by the large firm, its posterior beliefs imply that there is a higher chance that the aggregate state is adverse. This, in turn, increases the likelihood that the firm-specific economic state s_i will be adverse for the small firm. As a result, layoffs by the large firm lead

to an increase in π_i as the market prepares to observe the action of the small firm.

Combining this result with the analysis in Figures 1.3 and 1.4, it follows that following a layoff by the large firm, the small firm will be more likely to choose the *terminate* policy. Intuitively, the layoff by the large firm means the market will be more willing to attribute poor performance to an adverse state rather than a lack of managerial talent, making further layoffs more likely. In effect, our model predicts a clustering or "safety in numbers" effect, where some firms will strategically announce layoffs close to the announcements of other firms, in groups in order to minimize the reputational costs they incur. In particular, we expect that firms whose characteristics normally push them toward the *delay* policy will be followers in such situations, announcing layoffs in the wake of firms whose characteristics push them toward the *terminate* policy. The following proposition summarizes this insight:

PROPOSITION 4. Firms tend to cluster layoff announcements after layoffs by a large firm. That is,

$$\frac{\Delta \Pr[\text{Layoffs}_i=1 \mid \pi_i, \gamma, \sigma_\eta^2, C, \delta]}{\Delta(\text{Large Firm Layoff})} \geq 0.$$

PROOF. See appendix A.5. ■

The mechanism described above can also occur in response to economic news that signals a deterioration of firm performance. Consequently, we expect that adverse aggregate news, correlated with real firm performance, will also trigger clustering of layoff announcements. This gives us the following corollary:

COROLLARY 5. Firms cluster layoff announcements after negative macro-

economic news. The strength of the effect depends on the predictive power of the negative news with respect to firm performance. Specifically,

$$\frac{\Delta \Pr [\text{Layoffs}_i = 1 \mid \pi_i, \gamma, \sigma_\eta^2, C, \delta]}{\Delta (\text{Adverse Macro Event})} \geq 0$$

and

$$\partial \left[\frac{\Delta \Pr [\text{Layoffs}_i = 1 \mid \pi_i, \gamma, \sigma_\eta^2, C, \delta]}{\Delta (\text{Adverse Macro Event})} \right] / \partial \left[\frac{\Delta p_i}{\Delta (\text{Adverse Macro Event})} \right] \geq 0$$

PROOF. See appendix A.6. ■

The key message of this analysis is that after a large firm layoff (or release of negative macroeconomics news), perceived probability of industry downturn (π_i) increases. This in turn leads to an increase in the layoff propensity of other firms. Since other firms' layoff propensity increases simultaneously, they all tend to lay off at the same time, leading to clustering. Therefore clustering in this model is not driven by the desire of firms to lay off close to other firms. Instead, it is driven by an aggregate shock (layoff of a large firm in the same industry, or other types of common bad news).

Extending the analysis further, we turn to the types of firms which are more likely to cluster their layoffs in response to shocks such as layoffs by large firms and adverse macroeconomic shocks. Similar to the results of proposition 1 and 2, we find with strong reputation-based incentives and shorter track records are most likely to engage in layoff clustering. The following two corollaries summarizes these results:

COROLLARY 6. If there is significant uncertainty about the manager's talent, then he is more likely to cluster layoff announcements after layoffs by

a large firm: $\frac{\partial}{\partial \sigma_\eta^2} \left[\frac{\Delta \Pr[\text{Layoffs}_i=1 \mid \pi_i, \gamma, \sigma_\eta^2, C, \delta]}{\Delta(\text{Large Firm Layoff})} \right] \geq 0.$

PROOF. See appendix A.6. ■

COROLLARY 7. If the manager's utility function puts more weight on his reputation then he is more likely to cluster layoff announcements after leader layoff. That is $\frac{\partial}{\partial \gamma} \left[\frac{\Delta \Pr[\text{Layoffs}_i=1 \mid \pi_i, \gamma, \sigma_\eta^2, C, \delta]}{\Delta(\text{Large Firm Layoff})} \right] \geq 0.$

PROOF. See appendix A.6. ■

The propositions and corollaries in Section 3 summarize the testable predictions of the model. We now move to Section 4, which discusses how we link the parameters of the model to measurable attributes of firms and managers. With respect to the empirical tests, proposition 2 and its associated corollaries deal with leader-follower behavior, which we test using high-frequency data over short time horizons. By contrast, we test propositions 1 and its corollaries using lower-frequency data over the course of the business cycle.

3.6 Mapping the Model to the Data

The model presented in this section is a static three-period model. Therefore, in order to test the model's predictions we need to specify the appropriate time horizon. In principle, the time horizon depends on the persistence of beliefs about the economic state, and the corresponding persistence of reduced reputational costs to layoffs. Thus, to guide our empirical tests we choose the appropriate time horizon for our tests based on the frequency at which market's belief about the aggregate state of the economy changes. As summarized in the propositions above, the key comparative statics involve change in manager's behavior after the release of *adverse aggregate*

news. Guided by this principle, our empirical tests are conducted at two frequencies: business cycle frequencies and daily frequencies. The timescale of business cycles is a natural candidate because market participants are much more pessimistic about the aggregate states during recessions relative to booms. Similarly, testing the predictions at daily frequencies is informative, since release of unexpected bad news by a leading firm in an industry often drastically changes market's beliefs about the state of the industry in a matter of hours. The next section describes how we construct our datasets, and then turns to the empirical strategies and results for both the business cycle frequency tests (Section 5), and the daily frequency tests (Section 6).

4 Data Construction

4.1 Constructing Dataset of Layoff Announcements and Firm Characteristics

The data for this study are based on two sets of firms: large publicly-listed firms, and large privately-held firms. The publicly-listed firms are the population of firms in the annual Fortune 500 from 1970 to 2010. Analogously, the privately-held firms are the population of firms in the Forbes annual list of largest 100 privately held firms ("Forbes 100") from 1985 to 2010. To minimize selection bias, we restrict the sample in any given year to the subset of firms that are contemporaneously constituents of the Fortune 500 or Forbes 100 in that year.¹⁵ Over the relevant range of years, we track 1013 different publicly-listed firms and 436 privately-held firms at an

¹⁵However, conducting our empirical analyses on the entire sample of firms that were *ever* in the Fortune 500 or Forbes 100 does not alter our key results.

annual level. With this framework, we track announcements of layoffs by these firms in the Wall Street Journal, the definitive source of news for large US-based firms. For the publicly-held firms data from 1970 to 2006 comes from Kevin Hallock.¹⁶ Using the same methodology as Hallock (2009), we extended this dataset to 2010¹⁷, and independently constructed a dataset of layoff announcements for the private firms from 1985 to 2010 (see data appendix for more details).

From the Wall Street Journal announcement dataset, we focus not only on the number of layoffs by a particular firm in a given year, but also track the total number of workers laid off. We then match our firms to four of COMPUSTAT's datasets: Prices, Dividends, and Earnings; Fundamentals Annual, Fundamentals Quarterly, and ExecuComp. From the Prices, Earnings, and Dividends dataset, we obtain a firm's NAICS code, as well as information on its annual earnings and its equity: shares outstanding, market and book values, and dividends. From Fundamentals Annual, we obtain firm employment numbers and information from balance sheets and income statements: measures of debt, revenues, income, and capital expenditures. From Fundamentals Quarterly, we obtain date of earnings announcements, which serves as an important control variable in some of our empirical tests.

The data from ExecuComp is limited by the fact that it starts in 1992; however, it provides valuable information on the tenure and compensation

¹⁶For related research using this data see Billger and Hallock (2005) and Farber and Hallock (2009). Also, Hallock (2009) provides an interesting discussion of other aspects of this dataset which we do not explore.

¹⁷There are several approaches to conducting searches on historical news database. In consultation with Kevin Hallock we narrowed the search criteria to three different methods. Using the three criteria we re-constructed the dataset for the publicly-held firms for three random years in the period 1970 and 2006. To ensure consistency we settled on the search criteria that yielded the maximum amount of overlap between the two datasets.

of the CEOs of firms in our sample. We supplement this dataset with information from the Forbes CEO Compensation list of the largest 500 firms from 1970 to 1991. This allows us to construct measures of CEO tenure over several decades for a large subset of firms. This is critical for some of our empirical tests involving CEO tenure, as it allows us to include firm-fixed effects to examine within-firm variation.

In addition to these firm-specific measures, we also obtain sector-level data from the BLS Current Employment Statistics National Survey covering employment levels and number of hours worked, and measures of value-added from the National Income and Product Accounts of BEA, decomposed by NAICS major industry groups. We also obtain daily stock market returns from the CRSP database for the entire sample period, 1970 to 2010.¹⁸

With these data, we first construct a range of standard control variables in order to cover a wide range of standard predictors of firm behavior. Specifically, the following variables are constructed based on firms' annual earnings reports covering the year prior to the layoff announcements being analyzed. We begin with the standard measures investors use to categorize companies into groups: firm size and value vs. growth. For the former, we include both the traditional market capitalization measure, as well as a measure of total firm value which combines equity market capitalization with the firm's long-term debt obligations. For the latter, we use both the ratio of equity book value to equity market value, as well as the earnings to price ratio for the firm's stock. In addition to these, we include a measure of financial leverage, equal to the ratio of the value of long-term debt obligations to the sum those obligations and the firm's equity. We also construct a

¹⁸The data appendix goes into more detail about our methodology and procedures for constructing and merging the different datasets.

measure of firm maturity as measured by years since initial public offering (IPO) date.

To test the propositions outlined in the theory section we construct two different datasets. In the first dataset each firm is tracked annually (Annual dataset), and in the second each firm is tracked every business day (Business Day dataset). Out of 5569 layoff announcements we only find two to be announced in the Weekend edition of the Wall Street Journal. Consequently, the Business Day Histories is at the business day level rather than the calendar day level.¹⁹

4.2 Confidential Microdata from the Mass Layoff Statistics Program of BLS

The Mass Layoff Statistics program (MLS) of the Bureau of Labor Statistics (BLS) is a Federal-State cooperative statistical effort which uses a standardized, automated approach to identify, describe, and track the effects of major job cutbacks, using data from each State's unemployment insurance database. Establishments which have at least 50 initial claims for unemployment insurance (UI) filed against them during a consecutive 5-week period are contacted by State agencies to determine whether the claimants are facing separations of at least 31 days duration, and, if so, information is obtained on the total number of separations, the reasons for these separations, and recall expectations. Establishments are identified according to industry classification and location, and unemployment insurance claimants are identified by demographic characteristics including age, race, sex, ethnic group, and place of residence. The data is collected at a monthly frequency

¹⁹These two layoff announcements were recoded as occurring on the following Monday. Our results are identical when we drop these two observations.

starting in April of 1995. We end our sample in December 2010.

According to the MLS definitions, a mass layoff occurs when at least 50 initial claims are filed against an establishment during a consecutive 5-week period. An extended mass layoff occurs when at least 50 initial claims are filed against an establishment during a consecutive 5-week period and at least 50 workers have been separated from jobs for more than 30 days. Since extended mass layoffs are a better measure of layoffs that lead to more permanent job dislocation (greater than 30 days), we focus on this measure in our analysis.

Our focus on the subset of establishments employing 50 or more workers means that, according to the 2003 data, 4.6 percent of all covered employers and 56.7 percent of covered employment are in the program's scope.²⁰ This measure has been quite stable over time: more than two decades ago, 5 percent of employers and 61 percent of total employment were reported in establishments with 50 or more workers (Brown (2004)).

The Bureau of Labor Statistics keeps the identity of companies that engage in mass layoffs confidential. Under the auspices of the onsite researcher program of BLS, we were able to access the confidential microdata, which allowed us to extend our empirical analysis to actual mass layoffs, in addition to the layoff announcement observations based on the Wall Street Journal data. Five state employment offices, however, rejected our proposal to access the confidential data citing state legislation that disallows them to share the identity of establishments even for research purposes. Nevertheless, the researchers at the BLS estimate that the confidential data that was

²⁰The large difference in percentages reflects the strongly right-skewed distribution of employer size, where a relatively small fraction of establishments provide a majority of jobs.

accessible to us covered more than 85% of all the mass layoff events they track. Having access to actual mass layoffs data allows us to examine the degree to which strategic behavior by firms can lead to actual changes in the labor market outcomes.

5 Business Cycle Frequency Tests

One of the main predictions of the model presented in Section 3 is that firms with managers who care most strongly about reputation are more likely to engage in layoffs during downturns (proposition 2 and 3). In an ideal experiment, we would estimate the magnitude of this effect using two identical firms, such that the manager of one firm has incentives to manage reputation while the other does not. In the absence of such an experiment, we exploit differences in the incentives faced by publicly-listed firms (“public” firms) and privately-held firms (“private” firms). Public firms differ from private firms along three major margins, all of which make their managers more likely to manage their reputation in financial markets, relative to a similar private firm. First, since public firms sell shares to outside investors who are not involved in managing the firm, there exists separation between ownership and control. This may lead to agency problems if managers’ interests diverge from those of their investors (Jensen and Meckling (1976)). Second, owners of public firms typically have shorter horizons, since liquidity makes it easy for shareholders to sell their stock at the first sign of trouble rather than actively monitoring management. This relative myopic behavior of investors, weakens incentives for effective corporate governance (Amar (1993)), and generates incentives for managers to be my-

opic in their reputation management (Stein (1989)). Third, managers of public firms are subject to takeover threats, which are, in part, dependent on the stock price of targeted firms. This can lead to managerial myopia in public firms in order to actively manage current stock prices (Stein (1988) and Edmans, Goldstein, and Jiang (2011)).²¹

If reputation management drives layoff behavior over the business cycle, the first-order effect should show up when comparing differences in behavior of *similar* public and private firms. Here we emphasize reputation management in *financial markets*, since the managers of public and private firms are likely to have similar motivations for reputation management among other constituents. The next section describes the empirical strategy and data samples we use for our tests.

5.1. Comparing Public-Private Firms: Empirical Strategy

The analysis of this section is based on the confidential microdata collected at a monthly frequency from April 1995 to December 2010 by the Mass Layoff Statistics program of the Bureau of Labor Statistics. The dataset includes firms that ever engaged in a layoff during the sample period. Using this data we create three different samples for our study.

5.1.1 Full Sample. — The construction of our full sample for this portion of our analysis begins with all public firms that are Fortune 500 constituents between 1985 and 2010; and all private firms that are Forbes 100 constituents

²¹Private firms, in contrast, are often owner-managed and even when not, are both illiquid and typically have highly concentrated ownership, which encourages their owners to monitor management more closely. Indeed, evidence from the Federal Reserve's 2003 Survey of Small Business Finances (SSBF) shows that 94.1% of the larger private firms in the survey have fewer than ten shareholders (most have fewer than three), and 83.2% are managed by the controlling shareholder. As a result, agency problems are likely to be greater among public firms than among private ones.

between 1985 and 2010. We then match these firms to the microdata we accessed from the MLS database, resulting in a total of 478 public firms, and 135 private firms tracked over the period covered by MLS, namely, 1995-2010²². We call this sample the “full sample.” Table 1.2 reports the characteristics of both the private and public firms in this sample. Over the 1995-2010 time period, the public firms tend to be larger than the private firms in terms of both revenue and number of employees. Also, the baseline lay-off propensity of public firms is about 1.8 percentage points greater than the private firms, although we find no difference in the number of workers laid off by both these firms in a given mass layoff event. In an ideal world, we would like to compare the investment behavior of two otherwise identical firms that differ only in their listing status. To get closer to this ideal we need to find pairs of public and private firms that are observably similar to each other. One convenient way to do this is through matching, which is what we turn to next.

5.1.2 Matching Sample. — Since size is an important observable difference between the public and private firms in our sample, we match on size (revenue) in addition to matching on industry. This procedure closely follows the methodology of Asker, Farre-Mensa, and Ljungqvist (2011), who conduct a similar matching between public and private firms to investigate differences in investment sensitivities. Matching on size implies that our matched sample consists of the bottom half of public firms in the Fortune 500, which correspond to the size of all private firms in the Forbes 100. (see

²²The number of firms we track for this analysis is reduced by two factors: First, five states did not allow us to access their mass layoffs information. Second, not all firms in our broad sample engaged in layoffs between 1995 and 2010. This sample differs from our other results in that it considers non-contemporaneous constituents. The vast majority of our results are unchanged when restrict the sample to contemporaneous constituents of the two lists between 1995 to 2010, although this substantially reduces the sample size.

TABLE 1.2: Public - Private Firm Characteristics

This table presents descriptive statistics for the full samples of public and private firms, for a size-and-industry matched sample, and for a sample of public and private firms that were successfully or unsuccessfully targeted by leveraged buyouts (LBO), over a period from 1995 and 2010. The table reports mean, median, and standard deviation of the key variables used in our empirical analysis that compares public-listed and privately-held firms. We also report pairwise differences in means, with *** and ** indicating a difference that is significant in a t-test at the 1% and 5% level, respectively.

	FULL SAMPLE			MATCHED SAMPLE			BUYOUT SAMPLE		
	Public Firms	Private Firms	Difference in Means	Public Firms	Private Firms	Difference in Means	Public Firms	Private Firms	Difference in Means
Revenue (in 2005 USD)									
Mean	15143.29	8067.53	7064.76***	7162.07	6992.29	169.78	9218.31	8609.35	608.97
Median	6035.37	4406.83		5075.34	4903.13		4871.72	5258.40	
St. Dev	30540.05	10817.58		5122.21	4951.14		14895.95	6995.246	
Employees									
Mean	46803.01	36624.15	10178.86***	29341.66	29614.29	-272.63	43101.70	46120.20	-3018.50
Median	21586	17649		19117	15000		21000	28000	
St. Dev	98809.90	57237.13		20212.44	26967.78		68204.04	44141.21	
Layoff Propensity									
Mean	0.0721	0.0541	0.0180***	0.0584	0.0558	0.0027	0.0689	0.0698	-0.0009
St. Dev	0.2587	0.2263		0.2345	0.2295		0.2533	0.2550	
Number Laid off									
Mean	337.74	357.88	-20.14	281.59	248.82	32.77	284.32	409.57	-125.25**
Median	174	200		160	150		166	191	
St. Dev	953.98	490.70		521.45	418.56		446.33	615.03	
Share Laid Off									
Mean	0.0134	0.0165	-0.0031	0.0138	0.0169	-0.0031	0.0134	0.0168	-0.0034
Median	0.0035	0.0060		0.0047	0.0056		0.0043	0.0056	
St. Dev	0.0819	0.0356		0.0132	0.0442		0.0686	0.0280	
No. of Firms	478	135		206	74		42	23	
No. of Observations	70665	10632		10470	10470		5502	1803	

Table 1.2 for a comparison).

In the language of the matching literature surveyed in Imbens and Wooldridge (2009), we use a nearest-neighbor match adapted to a panel setting. Starting in fiscal year 1985, for each public firm, we find the private firm that is closest in size and that operates in the same three-digit NAICS industry, requiring that the ratio of their total revenue (TR) is less than 2 (i.e., $\max(TR_{public}, TR_{private}) / \min(TR_{public}, TR_{private}) < 2$). If no match is found, we discard that observation and look for a new match for that firm in the following year. Once a match is formed, it is kept in subsequent years to ensure the panel structure of the data remains intact. If a matching firm exits the panel, a new match is spliced in. Because we match with replacement, to maximize the match rate and match quality, the matched sample contains 206 public firms and 74 private firms. Our results are not sensitive to matching without replacement, although this substantially reduces the sample size. The standard errors are appropriately clustered to account for the resampling of private firms. The middle three columns in Table 1.2 compare the characteristics of the matched sample, and allows us to assess how good this match is. Since we match on size as measured by revenue, it is not surprising to find no statistical difference in the average revenue of public and private firms in our matched subsample. We find almost no difference in the average number of employees between public and private firms, and no difference in average layoff propensity or severity.

5.1.3 Leverage Buyout Sample. — Next, we create an alternate subsample based on leverage buyout (LBO) attempts. From the full sample we only keep private firms that were once public and went private through a LBO after 1985. We obtain this data from the Forbes annual survey of largest

private firms in the US. As for the public firms, we track firms after 1985, and only include the public firms that were targeted by an unsuccessful LBO attempt.²³

In Table 1.2 we report the observed differences between public and private firms in this subsample. We again find no significant differences in revenue, number of employees, layoff propensity, and the share of employees laid off between public firms (resulting from unsuccessful LBO attempts) and private firms (resulting from successful LBO attempts). We do find a difference in the average number of workers laid off, but the difference between medians is much smaller.

5.2 Comparing Public-Private Firms: Results

The main results of the analysis are reported in Table 1.3. We estimate the same set of two regressions using the three different samples described above. In each set the first regression has an indicator for mass layoff as the dependent variable, while the second regression uses the share of employees laid off (conditional on a layoff). All regressions include controls for the log of the previous year's revenues. Additionally, the regressions with layoff indicator as the dependent variable controls for the previous year's employee size. We are unable to control for other firm characteristics since the Forbes dataset only reports these two variables for private firms.

5.2.1 Full Sample Results. — In the first set of specifications, (1) and (2), we estimate the regressions on the full sample, with no firm fixed effects

²³If the withdrawal of the LBO action is random, differences between successful LBOs and withdrawn LBOs will allow us to identify the differences in layoff propensity of public firms relative to private firms. However, the withdrawal of LBOs are not always random. We run several tests to examine the observable differences between the two sets of firms. We find no systematic differences in firm characteristics. Still, unobservable differences may exist, which we are unable to control for.

or time fixed effects, but with 4-digit industry fixed effects. In addition we control for seasonality by including calendar month fixed effects, and for the overall time trend using a quadratic function. Using the BLS microdata for the period 1995 to 2010 at a monthly frequency, we find that private firms are roughly 2.01 percentage points more likely to layoff in a recession month. Compared to them, the public firms' propensity to lay off workers in recession months is 2.47 percentage points *greater*, indicating that the layoffs of public firms are more than twice as sensitive to recessions as those of private firms. The relatively modest response of private firms to recessions also shows up in the share of employees they lay off in recession years. Conditional on a layoff, private firms exhibit no difference in the share of employees they layoff in recession months and non-recession months. Contrastingly, the share of workers laid off by public firms goes up by 0.28 percentage points (conditional on a layoff) in a recession month when compared to the share laid off by private firms.

In the next set of specifications, (3) and (4), we conduct the same analysis but with firm fixed effects instead of industry fixed effects, and with month fixed effects instead of the controls for seasonality and time trends. We are still able to identify the impact of the private firm indicator, since this regression takes advantage of public-to-private and private-to-public transitions. During our sample years of 1995-2010, we have 38 such transitions. The recession indicator in these regressions is not identified due to presence of month fixed effects. Under these specifications we find very similar results: a firm is more likely to announce layoffs in a recession if it is public. Similarly, when compared to private firms, public firms lay off more workers in recession years, although this coefficient is imprecisely estimated.

Table 1.3: Public - Private Comparison over the Business Cycle

This table analyzes differences in actual layoff behavior over the business cycle between public and private firms. The unit of observation is a the firm level tracked monthly between April 1995 and December 2010. In this table, a layoff constitutes a mass layoff as defined by the Bureau of Labor Statistics under its Mass Layoff Statistics Program (see section 4.2 for further discussion). In this table we report four sets of regressions, each with two different dependent variables. The first dependent variable is layoff indicator, which takes a value of one if the firm engages in a mass layoff in a given month. The second dependent variable is the number of workers laid off as a share of previous year's employees. The first two set of regressions are estimated over the entire sample, which include firms that were in the Fortune 500 (public firms) or Forbes 100 (private firms) between 1985 and 2010. The only difference between the first and second set of regressions is that the first set includes 4-digit industry fixed effects instead of firm fixed effects, and a quadratic time trend and controls for seasonality (calendar-month fixed effects) instead of month fixed effects. On the other hand the second set includes firm fixed effects and month fixed effects, instead of industry fixed effects or controls for time trend. In the second set of regressions we are able to identify the effect of the private firm indicator variable since we have 38 cases of firms that transitioned from private-to-public or vice-versa. In the last two set of regressions, (5)-(8), we restrict the sample to matched public-private pairs (see section 5.1.2 for the methodology) based on size (revenue) and 3-digit NAICS industry. These specifications rely on comparing each public firm to its matched private counterpart since we include matched-pair fixed effects. All regressions include controls for previous year's log revenue and its interaction with the recession dummy, and previous year's number of employees and its interaction with the recession dummy (except when the dependent variable is share laid off). Since the month fixed effects absorb all variation over time, the effect of a recession indicator is not separately identified in the specifications that include month fixed effects. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, * and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Dependent Variable	Full Sample			Matched Sample				
	Layoff Indicator (1)	Share Laid Off (2)	Layoff Indicator (3)	Share Laid Off (4)	Layoff Indicator (5)	Share Laid Off (6)	Layoff Indicator (7)	Share Laid Off (8)
Public Indicator	0.0051 (0.0085)	0.0020 (0.0013)	-0.0029 (0.0127)	0.0029 (0.0017)	-0.0064 (0.0090)	-0.0069*** (0.0025)	-0.0044 (0.0063)	-0.0060 (0.0037)
Recession Indicator	0.0201** (0.0087)	-0.0011 (0.0015)			0.0240** (0.0161)	0.0008 (0.0014)		
Public x Recession	0.0247*** (0.0097)	0.0028* (0.0015)	0.0259*** (0.0090)	0.0025 (0.0018)	0.0320** (0.0136)	0.0040* (0.0022)	0.0307** (0.0142)	0.0012 (0.0028)
Mean	0.0696	0.0137	0.0696	0.0137	0.0569	0.0155	0.0569	0.0155
Std. Dev	0.2544	0.0785	0.2544	0.0785	0.2318	0.0267	0.2318	0.0267
Industry Fixed Effect (4-digit NAICS)	✓	✓						
Quadratic Time Trend & seasonality controls	✓	✓			✓	✓	✓	✓
Month Fixed Effects			✓	✓				
Firm Fixed Effects			✓	✓				
Matched-Pair Fixed Effects					✓	✓	✓	✓
Log Employees and Recession Interaction	✓		✓		✓		✓	
Log Revenue and Recession Interaction	✓	✓	✓	✓	✓	✓	✓	✓
Observations	81414	5587	81414	5587	20592	1167	20592	1167

5.2.2 Matched Sample Results. — Our results using the matching estimator are reported in columns (5)-(8). These four specifications are a counterpart to the first four specifications discussed above. In addition to using our smaller matched sample, the key difference is in the control structure: we include a matched-pair fixed effect instead of the industry- and firm-level fixed effects in the previous specifications. Therefore, the identification in these regressions is based off within-pair variation, where each pair consists of one public and one private firm within the same subindustry matched on size. The matching estimator results are in line with the results above: the layoff propensity of public firms is more than twice as sensitive to recessions as that of private firms, and the same is true for the share of employees laid off conditional on a layoff. Specifically, the propensity of private firms to engage in a mass layoff increases by 2.4 percentage points in recession months, while public firms experience an increase that is 3.2 percentage points *greater* than the effect we see in private firms. We also find that public firms are 0.64 percentage points less likely to lay off workers outside recession months when compared to private firms, although this result is not precisely estimated. We find similar results for the fraction of employees laid off conditional on a layoff. Outside recession months public firms lay off a relatively smaller fraction of employees compared to private firms, whereas the opposite holds in recession months. These coefficients, however, are not precisely estimated and cannot be interpreted as definitive results.

The general message from the matched sample results is that the public firms are more cyclical in their layoff policies compared to their matched private counterparts. These results suggest that public firms may be car-

rying excess capacity and waiting for longer periods between layoffs when compared to similar private firms. To investigate this further, we use the same matching methodology to estimate both the total number of workers laid off and the median duration between layoffs over the course of a full business cycle. We present these results in Table 1.4. In examining the total number of workers laid off, we use a peak-to-peak identification, starting from October of 2000 and ending in July of 2007. We report our findings in column (1): public firms tend to lay off approximately 30% more workers over the course of the entire business cycle, though this result is not precisely estimated. When we consider the median duration between layoffs, we use a slightly different time period. We begin our sample in April of 2002, approximately six months after the trough of the 2001 recession, and end in December of 2009. The motivation for this is to observe firm behavior in the period after they are most likely to have adjusted their labor force to their desired optimal level: most layoffs occur between the start and the trough of a recession, and the firms in our sample would have had ample opportunity to adjust their labor force. Using within-pair variation, we find evidence for a difference in layoff timing between public and private firms, and report these results in column (2) of Table 1.4. We find that the average duration between layoffs for a sample of firms that engage in layoffs during the 2002-10 period is 7.76 months. Using within-pair variation, we find that the duration between layoffs for public firms is roughly 0.65 months greater than their matched private counterpart. While this effect is not precisely estimated, it is consistent with the view that public firms may wait longer to announce layoffs.

5.2.3 Leveraged Buyout Targets Sample Results. — The last set of results

Table 1.4: Public - Private Comparison over the Business Cycle (part II)

This table reports characteristics of the matched sample and examines differences in public and private firm behavior over the business cycle with respect to their layoff policies. The first specification (1) computes the total number of workers laid off by a firm over an entire business cycle (peak-to-peak from 6 months prior to the peak of the 2001 recession (October 2000) to 6 months prior to the peak of the 2008 recession (July 2007)). Correspondingly each firm has one observation in this sample. We control for matched pair fixed effects, and correspondingly the coefficient on the public indicator variable relies on a comparison of a public firm with its matched private counterpart. In specification (2) using the same data structure we study the median distance mass layoffs for each firm over an entire business cycle. For this specification we rely on a trough-to-trough period, since the period right after a recession is a more natural starting point for this analysis. Once again each firm has one observation in the sample. We again control for matched pair fixed effects so as to rely on within pair variation. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Sample	Matched Sample	
Dependent Variable	workers laid off over business cycle (1)	Median Months Since Last Layoff (2)
Public Indicator	0.3039 (0.4251)	0.0656* (0.0363)
Mean	2.5745	7.76
Std. Dev	3.1077	6.69
Matched-Pair Fixed Effects	✓	✓
Observations	236	260

are estimated using a fixed effects specification with the firms in our LBO sample, and we report these results in Table 1.5. We find that after controlling for 4-digit industry category, log of revenues and log of employees, the layoff propensity of public firms increases by 6.04 percentage points in recession months when compared to private firms in the sample. Outside of recession months, however, we find that public firms are slightly less likely to engage in layoffs. Examining the share of workers laid off, we once again we find that public firms are likely to lay off a larger fraction of their workforce in a recession month when compared to a private firm, but the effect is not statistically significant. Overall, these results are consistent with those from our other samples: public firms are much more cyclical when compared to private firms.

5.3 Possible Alternate Explanations for the Difference in Public vs. Private Layoff Behavior

This section explores the plausibility of explanations other than reputation management for the difference in layoff propensity between public and private firms. The key concern is that when we compare public and private firms, there are unobservable differences between unrelated to reputation management which may be driving the results presented above. Among the set of possible sources of unobservable heterogeneity between public and private firms, financial leverage and lifecycle effects are central and may directly alter the layoff behavior of firms independent of any reputation-management behavior of managers. Private firms tend to have greater degree of financial leverage compared to public firms, and also tend to be younger than public firms. Since we do not observe these characteristics for

Table 1.5: Public - Private Comparison over the Business Cycle (part III)

This table analyzes differences in actual layoff behavior over the business cycle between public and private firms. The unit of observation is a the firm level tracked montly between April 1995 and December 2010. In this table, a layoff constitutes a mass layoff as defined by the Bureau of Labor Statistics under its Mass Layoff Statistics Program (see section 4.2 for further discussion). The set of regressions is estimated on a subsample of firms that were targets of a leveraged buyout (LBO). Among the targets the public firms are those for whom the LBO offer was withdrawn, and the private firms are those for whom the buyout offer was successful. The regressions include month fixed effects, and controls for previous year's log revenue and its interaction with the recession dummy, and previous year's number of employees and its interaction with the recession dummy (except when the dependent variable is share laid off). Since the month fixed effects absorb all variation over time, the effect of a recession indicator is not separately identified in these specifications. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Sample	LBO Targets	
	Layoff Indicator (1)	Share Laid Off (2)
Public Indicator	-0.0079 (0.0261)	-0.0088 (0.0272)
Public x Recession	0.0604** (0.0258)	0.0181 (0.0176)
Mean	0.0692	0.0142
Std. Dev	0.2550	0.0610
Industry Fixed Effect (4-digit NAICS)	✓	✓
Month Fixed Effects	✓	✓
Log Employees and Recession Interaction	✓	
Log Revenue and Recession Interaction	✓	✓
Observations	7116	470

our sample of private firms we cannot control for them in our regressions. In order to assess the importance of these characteristics we instead compare our sample of private firms to the most levered public firms, and to the youngest public firms. If these characteristics are drivers of layoff policies, we would expect that in recessions, the layoff behavior of high-leverage public firms and young public firms will be quite similar to that of private firms. We investigate these alternate hypotheses in Table 1.6.

The dependent variable in all the regressions in Table 1.6 is the layoff indicator. Column 1 restricts the sample of public firms to ‘young’ firms (those whose time-since-IPO in their first year in our panel is less than the median time-since-IPO of all public firms in a given calendar year), while column 2 restricts the sample of public firms to ‘old’ firms. Similarly, Column 3 restricts the sample of public firms to ‘high leverage’ firms (those whose leverage ratio exceeds the median leverage ratio of all public firms in a given calendar year), while column 4 restricts the sample of public firms to ‘low leverage’ firms. In the last specification, we restrict the sample to public firms only, and estimate the interactions between the recession indicator and each of the two characteristics above: the log of years since IPO, and the leverage ratio. Each regression includes month fixed effects, the log of previous year’s employees and its interaction with recession indicator, and the log of the previous year’s revenue and its interaction with the recession indicator.

We find that the younger public firms are much more likely to lay off in a recession month compared to older public firms (specification (5)). We also find no significant effect of leverage on the sensitivity of layoff propensity to recessions. These results are consistent with specifications (1)-(4): the

Table 1.6: Assessing Alternate Explanations for Public - Private Differences in Cyclicity of Mass Layoffs

This table explores the plausibility of alternate explanations for the difference in layoff propensity between public and private firms. The dependent variable in all the regressions is the layoff indicator. Column 1 restricts the sample of public firms to 'young' firms (those whose time-since-IPO in their first year in our panel exceeds the median time-since-IPO of all public firms in the same calendar year), while column 2 restricts the sample of public firms to 'old' firms. Similarly, Column 3 restricts the sample of public firms to 'high leverage' firms (those whose leverage ratio exceeds the median leverage ratio of all public firms in the same calendar year), while column 4 restricts the sample of public firms to 'low leverage' firms. In the last specification, we restrict the sample to public firms only, and estimate the interaction between the recession indicator and, both, log of years since IPO and leverage ratio. Each regression includes month fixed effects, log of previous year's employees and its interaction with recession indicator, and log of the previous year's revenue and its interaction with the recession indicator. The first four specifications include industry fixed effects, whereas the last one includes firm fixed effects. Since the month fixed effects absorb all variation over time, the effect of a recession indicator is not separately identified in these specifications. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Sample	Young Firms	Old Firms	High Leverage Firms	Low Leverage Firms	Public Only
Dependent Variable	Layoff Indicator (1)	Layoff Indicator (2)	Layoff Indicator (3)	Layoff Indicator (4)	Layoff Indicator (5)
Public Indicator	0.0095 (0.0154)	0.0055 (0.0231)	0.0009 (0.0139)	0.0071 (0.0196)	
Public x Recession	0.0478*** (0.0127)	0.0133 (0.0099)	0.0187* (0.0108)	0.0250** (0.0108)	
Leverage Ratio					0.0721*** (0.0137)
Log(1+Years since IPO)					0.0005 (0.0136)
Recession X Lev. Ratio					0.0259 (0.0190)
Recession X Log Age					-0.0341*** (0.0116)
Mean	0.0567	0.0785	0.0602	0.0759	0.0718
Std. Dev	0.2314	0.2689	0.2379	0.2639	0.2582
Firm Fixed Effects					✓
Industry Fixed Effect (4-digit NAICS)	✓	✓	✓	✓	
Month Fixed Effect	✓	✓	✓	✓	✓
Log Emp. and Interaction with Rec. Indicator	✓	✓	✓	✓	✓
Log Rev. and Interaction with Rec. Indicator	✓	✓	✓	✓	✓
Observations	44805	47019	45078	46746	69621

difference between public and private firms is strongest when comparing against young public firms, and is relatively consistent when comparing against high- and low-leverage public firms. The key implication of the results in this table is that the observed differences in layoff propensity in recessions between public and private firms are not being driven by unobservable differences in either leverage or lifecycle effects.

Though these results rule out two key possible alternate explanations, there may be other forms of unobserved variation between public and private firms. In order to investigate this further, we refined our matching criteria to match on size and 4-digit subindustry (instead of 3-digit subindustry). This reduces our sample size by approximately 50% , but we have enough observations to conduct similar analysis as reported in Table 1.3. The results of this analysis is reported in Table 1.7. These results based on the four digit industry level replicate what we find in Table 1.3, and the magnitudes of the coefficients in this table line up with the analysis using a matching criteria based on the three digit industry level. Therefore, our results is robust to changes in the matching criteria we use. Moreover, using a more stringent matching criteria (i.e. at the four digit industry level) allows us to mitigate unobservable differences between public and private firms.

5.4 Variation within Public Firms

So far, our analysis has been based on comparing public and private firms. In this section we look for differences in layoff behavior within our sample of public firms. In seeking to identify managers that are more likely to engage in reputation management and time layoffs strategically, we rely on the following two measures. First, we use an indicator variable called

Table 1.7: Public - Private Comparison over the Business Cycle (Robustness Tests)

This table is a robustness test for results presented in Table 2, in which the matching algorithm relies on 4-digit industry instead of 3-digit industry. The table exploits within-firm or within-industry variation to analyze differences in actual layoff behavior over the business cycle between public and private firms. The unit of observation is a the firm level tracked montly between April 1995 and December 2010. In this table, a layoff constitutes a mass layoff as defined by the Bureau of Labor Statistics under its Mass Layoff Statistics Program (see section 4.2 for further discussion). In this table we have two different dependent variables. The first dependent variable is layoff indicator, which takes a value of one if the firm engages in a mass layoff in a given month. The second dependent variable is the number of workers laid off as a share of previous year's employees. The sample includes matched public-private pairs (see section 5.1.2 for the methodology) based on size (revenue) and 4-digit NAICS industry. These specifications rely on comparing each public firm to its matched private counterpart since we include matched-pair fixed effects. All regressions include controls for previous year's log revenue and its interaction with the recession dummy, and previous year's number of employees and its interaction with the recession dummy (except when the dependent variable is share laid off). Since the month fixed effects absorb all variation over time, the effect of a recession indicator is not separately identified. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Sample	Matched Sample	
	Layoff Indicator (1)	Share Laid Off (2)
Public Indicator	-0.0194** (0.0090)	-0.0117*** (0.0041)
Recession Indicator		
Public x Recession	0.0312* (0.0185)	0.0115** (0.0048)
Mean	0.0569	0.0155
Std. Dev	0.2318	0.0267
Quadratic Time Trend & seasonality controls		
Month Fixed Effects	✓	✓
Matched-Pair Fixed Effects	✓	✓
Log Employees and Recession Interaction	✓	
Log Revenue and Recession Interaction	✓	✓
Observations	10920	656

short-tenured CEO, which takes a value of 1 if the CEO's tenure with the firm is four years or less.²⁴ Second, we measure a firm's past equity-linked compensation share as the share of total CEO compensation that comes from equity-linked instruments over the past 5 years. We are interested in the impact of these variables on actual layoff behavior over the business cycle, and we report the results in Table 1.8. We use the same set of firm-level controls as in Table 1.6, and also include month fixed effects and firm fixed effects. The sample includes all contemporaneous constituents of the Fortune 500 that ever engaged in a layoff. In specification (1) we find that short-tenured CEOs are roughly 1.44 percentage points more likely to engage in a layoff in a recession month, relative to firms with longer-tenured CEOs.²⁵ Similarly, in specification (2), we find that a one-standard-deviation increase in the share of equity-linked compensation (controlling for the level of total compensation) is associated with an increase of roughly 0.7 percentage points in layoff propensity during a recession month.²⁶

In specification (3) we restrict the sample to all firms that have option share above the median level for a given year. Similarly, in specification (4) we restrict the sample to all firms that have CEO tenure below the median level for a given year. We estimate the same regressions as in specifications

²⁴Some CEOs may choose to lay off workers immediately after they are hired, so as to start with a 'clean sheet.' Therefore, we also run the same test with a separate indicator variable for CEOs with a tenure between 0 and 1 year. Our results remain unchanged.

²⁵One might expect that firms which bring in new CEOs are also aiming to maintain (or even reduce) their labor force, rather than looking to expand aggressively. If true, this could lead to the result that firms with younger CEOs also engage in relatively more layoffs during recessions: firms looking to expand aggressively can respond by cutting back on hiring rather than announcing layoffs. To evaluate this alternative interpretation, we examine the relationship between new CEOs and employment growth at their firms. We find no effect, indicating that our layoff results are not coming from reduced hiring by new CEOs.

²⁶The equity-linked compensation findings are consistent with results in the earnings management literature with respect to CEO incentives (see Bergstresser and Philippon (2006))

(1) and (2), and find that our results get stronger: both short-tenured CEOs and equity-linked compensation are linked to a greater incidence of layoffs during recessions. These results suggest that within the set of public firms, the CEOs who are most likely to care strongly about reputation are also the ones who are most likely to engage in layoffs during recessions. This suggests that reputation management by CEOs plays a significant role in determining the cyclical nature of their firms' layoff policies.

6 Daily Frequency Tests

In this section we turn to the daily frequency tests of our model's predictions. First we establish that the reputation penalty is lower if a firm announces a layoff right after other large firms in the economy announce layoffs (Section 6.1). We evaluate the strength of the response to this reduction in reputation penalty when we test Proposition 4, which predicts that firms will cluster layoffs after layoff announcements by large firms. Finally, we test the differential sensitivity results in Corollaries 6 and 7, which predict that strategically-motivated managers (those with high γ and σ_{η}^2) will be more likely to cluster their layoff announcements.

6.1 Is the Reputation Penalty of Layoff Announcements Lower after Layoffs by Other Firms?

Though the manager may be managing his reputation with several constituents—the stock market, the board of directors, employees of the firm—the analysis in this section focuses on financial market reputation, which we can test this using daily stock returns. This focus on stock-market-based

Table 1.8: Actual Layoff Propensity over the Business Cycle (Public Firms), 1995-2010

This table exploits within industry variation to analyze the impact of the equity-linked executive compensation and short tenure of CEO on layoff propensity over the business cycle. The unit of observation is at the firm level tracked monthly between April 1995 and December 2010. In this table, a layoff constitutes a mass layoff as defined by the Bureau of Labor Statistics under its Mass Layoff Statistics Program (see section 4.2 for further discussion). The sample for the first two specifications includes all contemporaneous constituents of Fortune 500 that ever announced a layoff. In specification (3) we restrict the sample to all firms that are above the median with respect to option share of compensation in the previous year. Similarly, in specification (4) we restrict the sample to all firms above the median with respect to CEO tenure in the previous year. The dependent variable is an indicator variable, which takes a value of one when a firm announces a layoff in a given month, and zero otherwise. In each regression specification the main explanatory variables are the recession indicator, key variable, and the interaction of the key variable with the recession indicator. The recession indicator takes a value of one in months classified as recession months by the NBER Business Cycle Dating Committee. The first key variable (specification (1) and (3)) is a measure of short-tenured CEO, which is an indicator variable that takes a value of one if the tenure of a given CEO at the firm is between 0 and 4 years. The second key variable (specification (2) and (4)) is a measure of equity-linked compensation, which is the Black-Scholes past five year average of the value of stock-option grants a CEO receives as a share of the average total compensation. All the specifications include month fixed effects, industry fixed effects, and firm-level controls and the interaction of each of these controls with the recession indicator. The firm level controls include Earnings-Price ratio, Book-to-Market Ratio, Market Capitalization, Leverage Ratio, Years since IPO, which are all obtained from COMPUSTAT. Since the month fixed effects absorb all variation over time, the effect of a recession indicator is not separately identified in these specifications. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Sample	Full Sample		High Option Share Firms	Short-tenured CEO Firms
Dependent Variable	Layoff Indicator (1)	Layoff Indicator (2)	Layoff Indicator (3)	Layoff Indicator (4)
KV = Short CEO Tenure Indicator	0.0032 (0.0050)		0.0025 (0.0062)	
KV = Avg. Equity-Linked Compensation Share		0.0072 (0.0071)		0.0164 (0.0091)
Key Variable (KV) x Recession	0.0144* (0.0079)	0.0261** (0.0126)	0.0255** (0.0123)	0.0352** (0.0168)
Mean	0.0787	0.0790	0.0812	0.0831
Std. Dev	0.2693	0.2697	0.2732	0.2761
Industry Fixed Effects	✓	✓	✓	✓
Month Fixed Effects	✓	✓	✓	✓
Firm Controls and Recession Interaction	✓	✓	✓	✓
Observations	28746	27882	16818	13410

measures is driven by ease of testability, and should not be seen as a narrow interpretation of the reputation management mechanism in our model. Several studies have documented a negative stock market reaction to layoff announcements (Farber and Hallock (2009), Hallock (2009)). In this section we are interested in whether this negative penalty is lower if a firm announces a layoff within days of a layoff announcement by a large firm. Our empirical strategy follows the conventional event-study approach.

Using a sample of the Fortune 500 constituents from 1970 to 2010, we calculate cumulative excess returns using return data from the Center for Research in Security Prices (CRSP) at the University of Chicago. The excess return is the part of the movement in the stock return of a company that is not correlated with overall market movement in stock returns and presumably reflects unexpected firm-specific factors. To do this we run a first-stage regression where the daily stock return for company i on day t , denoted by R_{it} , is regressed on the value-weighted return of the market, R_{mt} :

$$R_{it} = \alpha_i + \beta_i R_{mt} + \eta_{it}$$

Next, for days around the event, the daily abnormal (or excess) returns is calculated as follows:

$$ER_{it} = R_{it} - (\hat{\alpha}_i + \hat{\beta}_i R_{mt})$$

where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are estimates of the previous regression. The first stage regression is run for a period in the past, which in this study ranged for a period of one year ending 30 days before the event. We rely on the average cumulative excess return over a five-day window—two days before, the day

of, and two days after the event for each of the 41 years from 1970 to 2010. Changes to this window length has no material effect on the results. The results are listed in Table 1.9 Panel A. Specification (1) restricts the sample to all layoff announcements that occur within the three days following a layoff by the largest 20 public firms in the economy as measured by the previous year's revenue. Specification (2) considers the complementary case, in which the sample includes layoff announcements that occur on all other days.

We find that the cumulative excess return around a layoff announcement that occurs within the three days following a top 20 firm announcement is -0.38 percentage points. On the other hand, if a layoff occurs on any other day, the cumulative excess return is more than twice as large: -0.87 percentage points. This suggests that in the context of financial markets, the reputation penalty of layoffs is lower immediately after negative signals about the state of the economy.

In Panel B of Table 1.9 we further examine the nature of stock market penalty of layoff announcements. Our model predicts a mitigated reputation penalty for the first-mover when their layoff is followed by layoffs by others in the industry. This is because the initial negative penalty should partially be undone when the market realizes other firms in the industry also are laying off. We test this in column (1) of Panel B. For a top-20 firm that laid off within the last three days, we find a three-day cumulative abnormal return of 0.33 percentage points when other firms in the industry engage in a layoff. Consistent with the results of Panel A, this result also confirms the mitigated penalty prediction of our model.

In column (2) we estimate the stock market reaction of other firms in

Table 1.9: Stock Penalty of Layoffs: Cumulative Abnormal Returns (from day -2 to day +2)

This table reports the cumulative abnormal returns (CAR) around a layoff announcement. The sample includes daily observations from 1970 to 2010. In Panel A, the baseline CAR is reported in specification (1) and (3) against the constant term of the regression. The main coefficient of interest is associated with the indicator variable reported in the first line, which takes a value of one if the firm under observation announces a layoff within three days after a layoff announcement by any of the largest 20 firms in the economy (as measured by previous year's revenue). In specifications (3) and (4) we also include an indicator variable for whether there was a layoff announcement by the same set of largest 20 firms within the three days following the layoff announcement by the firm under observation. Specification (2) and (4) also includes year fixed effects and industry fixed effects. Inclusion of these fixed effects implies that we cannot estimate the constant term in these specifications. In Panel B, we estimate the cross-effects of a layoff announcement on other firms. Column (1) reports the CAR of a top-20 firm that laid off within the last three days when another firm in the industry engages in a layoff. This is effectively the reverse effect on the industry leader when other firms follow up with a layoff. Column (2) reports the CAR of other firms in the industry when a top-20 firm lays off. Heteroskedasticity-consistent standard errors clustered at the industry level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Panel A

Dependent Variable	Cumulative	Cumulative	Cumulative	Cumulative
	Abnormal Return	Abnormal Return	Abnormal Return	Abnormal Return
	(1)	(2)	(3)	(4)
Layoff by Top 20 Firm in Previous 3 days = 1	0.0049*** (.0017)	0.0052*** (.0019)	0.0042** (.0020)	0.0044** (.0021)
Layoff by Top 20 Firm in Next 3 days = 1			-0.0025 (.0029)	-0.0027 (.0029)
Constant (Baseline CAR)	-0.0087*** (0.0013)		-0.0081*** (0.0016)	
Year Fixed Effect		✓		✓
Industry Fixed Effect		✓		✓
Observations	4796	4796	4796	4796

Panel B

Dependent Variable	CAR of Industry Leader that Laid off Within Last 3 Days when Follower Firms Lay off	CAR of Other Firms in Industry when Industry Leader Lay off
	(1)	(2)
Constant (CAR)	0.0033*** (.0011)	-0.0003 (.0007)
Observations	212	8286

the industry when a top-20 firm (as measured by previous year’s revenue) engages in a layoff. The model predicts a negative effect since a layoff by a large firm changes market perception about industry condition. We find a small negative but insignificant effect (-0.03 percentage points). This result indicates that the change in perceptions regarding industry conditions is largely offset by a countervailing effect, likely stemming from the dynamics of product market competition. Specifically, when one firm in an industry does poorly, other firms may benefit due to reduced competition in the product market. The results suggest that this second effect is offsetting the primary effect predicted by our model (i.e. adversely changing the belief about industry condition).

Lastly, in Table 1.10 we find that the first-mover penalty to persist even after one month. We refer you to table 1.10 for further discussion of this result. In the next section, we proceed to test whether firms respond to these incentives by timing their layoff announcements to occur immediately after a layoff announcement by top-20 firms.

6.2 Do Firms Announce Layoffs after other Large Firm Layoff Announcements?

6.2.1 Empirical Strategy. — To assess whether firms engage in a ‘leader-follower’ behavior with respect to layoff announcements, our estimation strategy relies on a dynamic regression model with lagged and future effects. For a firm i at time t , the regression specification is

$$Layoff_{it} = \alpha_i + \sum_{j=-p}^p \beta_j MacroEvent_{g,t-j} + X'_{it}\phi + Y'_t\omega + \varepsilon_{it} \quad (14)$$

Table 1.10: Stock Penalty of Layoffs: One-Month Cumulative Abnormal Returns

This table reports the one-month cumulative abnormal returns (CAR) after a layoff announcement. The sample includes daily observations from 1970 to 2010. In Column (1) the sample includes all firms that announce layoffs and are not classified as industry leaders (based on prior-year revenue). This column estimates the cumulative 1-month return of being a follower layoff firm (i.e. layoff within 1-week after a leader layoff) versus being a counterfactual-follower layoff firm (i.e. layoff within 1-week before a leader layoff). In column (2) the sample includes all firms that are industry leaders (i.e. firms that are classified as industry leaders based on prior-year revenue). This column estimates the cumulative 1-month return after a leader firm layoff for firms that have follower firms layoff within the next week versus those who do not using an indicator variable. Heteroskedasticity-consistent standard errors clustered at the industry level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Dependent Variable	Cumulative One Month Return of Firms that Lay off	Cumulative One Month Return of Firms that Lay off
Sample	Firms that Announce Layoffs (Excluding Industry Leaders)	Industry Leader Firms
	(1)	(2)
Laid off within 1-week after Leader Layoff = 1	0.0039 (.0088)	
Laid off within 1-week before Leader Layoff =1	-0.0168 (.0113)	
Industry leader with Follower Layoffs		0.0229*** (.0068)
Constant (Baseline CAR)	0.0040 (.0029)	0.0023 (.0022)
Observations	3125	1882

The central variables are $Layoff_{it}$, an indicator variable which takes a value of one when firm i announces layoff on business day t , and $MacroEvent_{g,t-j}$, which is an indicator variable that takes a value of one if there was a negative macroeconomic news released on date $t - j$ relevant for a firm in industry g . We also include firm-level controls, denoted by X'_{it} to control for firm-level heterogeneity. We begin our analysis by focusing on the layoff announcements of public firms. For public firms, these controls include total revenue, number of employees, years since IPO, book-to-market ratio, earnings-price ratio, leverage ratio, and days since last earnings announcement.²⁷ For private firms, we include the two control variables that we are able to observe: total revenues and number of employees. In addition to concerns about firm heterogeneity, we also need to account for the possibility that firms may be more likely to announce layoffs on certain days of the week (e.g. Friday) or in certain months. To address these concerns, the vector Y_t includes year fixed effects, month fixed effects, and day-of-week fixed effects.

We consider three different types of macro events: layoff by a top 20 firm (as measured by previous year's revenue), layoff by a firm in the top 20 which shares the same 1-digit NAICS code as firm i , and, as a placebo test, unexpected negative news announcements that are not directly linked to economic performance. We conduct this analysis separately for public firms and private firms. In the first measure, which we label "Leader Layoffs," we restrict our group of large firms to the largest 20 firms as measured by previous year's revenue.²⁸ This ensures that the layoff announcements

²⁷Controlling for earnings announcement date is important since firms may be clustering layoff announcements around earnings announcement date, and we may observe clustering merely because different firms have earnings announcement dates close to each other.

²⁸Press and analyst coverage of publicly-listed firms is highly skewed, with the largest

of these firms correspond closely with the notion of ‘negative macroeconomic news’ in our model. This measure takes the form of an indicator variable equal to one on the business day the WSJ reports any such layoff announcement. Our second measure, which we label “Industry Leader Layoffs,” is similar in structure. It takes a value of one whenever the first measure takes a value of one *and* the large firm is in the same 1-digit NAICS industry as firm *i*. The final measure, which we label “News Media Events”, is based on the ‘biggest news stories’ as measured by press coverage. Using a survey of major news events from USA today (2007), we construct a list of events for the period 1970 to 2010. Events are selected to be negative in nature (e.g. the Sept. 11 attacks), and to occur over a span of a day or less (i.e. news events such as the Afghanistan invasion of 2001 are excluded). As above, this measure is an indicator variable that takes the value of one on days when these events occur.²⁹

Our model predicts a strong asymmetry in layoff behavior before and after a large-firm layoff announcement. This is in sharp contrast to the common shock hypothesis, which predicts layoff announcements by smaller firms both before and after the leader layoff. The event study framework also enables us to tackle the issue of reverse causality: a potential concern is that smaller firms may drive large firm layoff announcements. There-

firms receiving an inordinately high degree of coverage compared to slightly less large firms (Fang and Peress (2009)). Our results are almost identical when we use a threshold of 5, 10, 25 or 30 for classification of large firms instead of the 20 largest firms measured by previous year’s revenue.

²⁹The purpose of considering these three different measures of macro events is to evaluate the predictions of the model in Proposition 4 and Corollary 5. As we will discuss in the alternate hypotheses section below, juxtaposing the results of these three regressions will allow us to exclude several alternate hypotheses that are otherwise consistent with some of our leader-follower results. Future terms are included in the dynamic regression model for falsification purposes. Most importantly, it allows us to tackle one of the main alternate hypotheses of common shocks.

fore, the coefficient on future events will enable us to establish whether this mechanism is at play. We report results for a lag length of $p = 5$, but note that our results are almost identical when we choose $p = 10$ or $p = 15$. To ensure comparability across all the regressions the sample of public firms is restricted to all contemporaneous constituents of the annual Fortune 500 list excluding the top 20 firms. Similarly, our sample of private firms include all contemporaneous constituents of the annual Forbes 100 list.

6.2.2 Results: High Frequency Announcement Behavior of Firms. — Figure 1.5 plots the sequence of β_j estimates from the event study specifications (14), along with point-wise 95% confidence intervals using standard errors clustered at the firm level. Panels A and B report the response to all “Leader Layoffs” by public and private firms, respectively. In Panel A, we see that the sequence of β_j estimates is roughly flat and close to zero before the leader layoff announcement ($j < 0$), and then jumps discretely at $t = 0$, and thereafter decreases gradually over the next 5 business days. Thus, a large firm layoff announcement is associated with future layoffs by other large public (Fortune 500) firms, but not with past layoffs. The magnitudes of these β_j should be compared to the unconditional average daily layoff announcement propensity of 0.0008. In Panel B, by contrast, we find no similar response to “Leader Layoffs” by private firms. Specifically, we find no evidence of clustering either before or after a top 20 layoff announcement. This is consistent with our business cycle frequency results, in which the public firms exhibited much greater propensity to engage in actual layoffs in recession months compared to the private firms.

In Panel C, we return to public firms and investigate the response to “Industry Leader Layoffs.” We again find the pattern of no effect prior to

the event date, followed by a strong jump and gradual decline at $t = 0$. Notably, the magnitude of the within-industry response of public firms to a layoff by a large firm is almost twice that of the economy-wide response described in Panel A. In Panel D, we turn our focus to the response of public firms to “News Media Events.” In contrast to events based on large-firm layoffs, a large negative media event does not predict future or past layoff announcements by the Fortune 500 firms. This suggests that firms are not attempting to time layoff announcements during periods when investors may be distracted by non-economic events.

Taken together, these results suggest that the Fortune 500 firms are more likely to time their layoff announcements in the days immediately after negative economic news is released, such as the aftermath of a layoff announcement by a very large firm. The strength of the ‘follower’ behavior is stronger when the negative news is more related to the firm’s own productivity; and the effect is absent after large negative news that are non-economic in nature. Therefore, these results are more in line with the leader-follower mechanism being driven by an informational channel, rather than a ‘hiding behind the headlines’ channel. Moreover, the asymmetric response of the firms before and after the large layoff announcements offer evidence against the ‘common shock’ hypothesis, which would predict a more symmetric response in layoff announcement responses of the Fortune 500 firms.

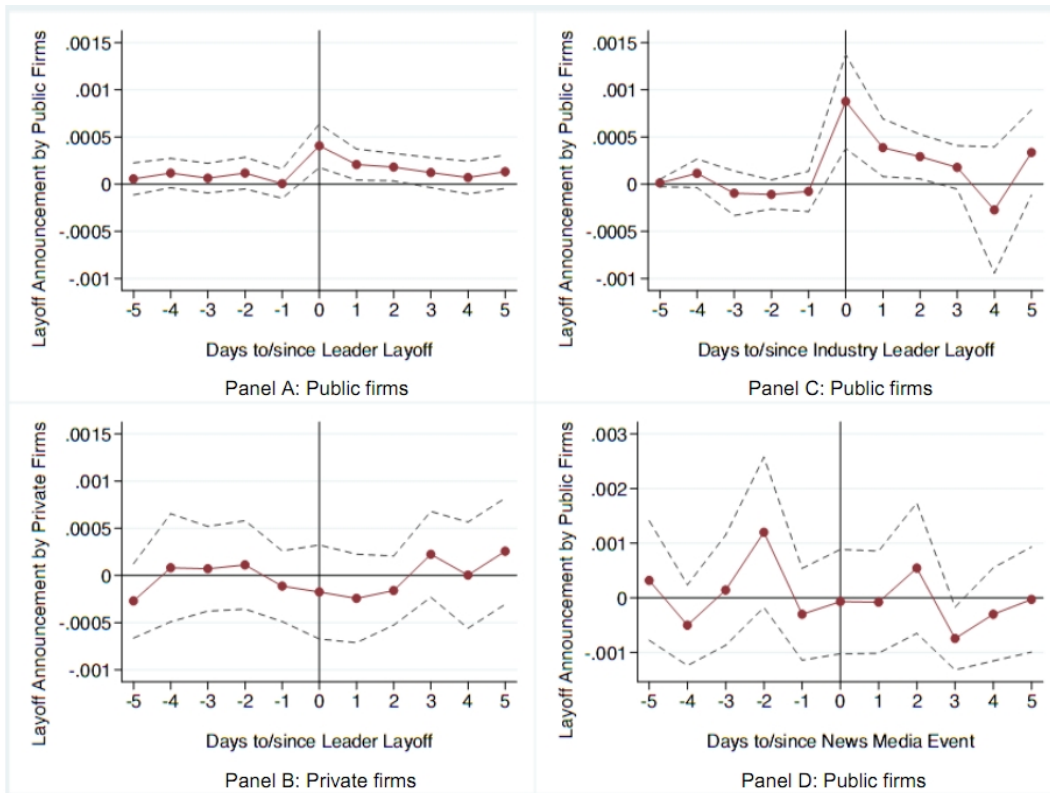


Figure 1.5: These graphs report the response of layoff propensities to “Leader Layoffs” — layoff announcements by the largest 20 firms in the economy, based on previous year’s revenue. We plot the coefficients of dynamic regressions which predict daily layoff propensities over the 11-day event-time periods surrounding large-firm layoff announcements.

6.3 Which Type of Firms Layoff after a Large Firm Layoff Announcement?

6.3.1 Empirical Strategy. — To identify the characteristics of firms that announce layoffs immediately after a large firm layoff announcement, we construct an annual dataset in which the unit of observation is at the firm level for every year. With this framework, we aim to test the model’s predictions on the type of firms which are most likely to behave strategically.

Specifically, corollaries 6 and 7 predict that managers for whom the market lacks strong priors, and who care most about reputation, will be most likely to engage in strategic layoff timing. These results rely on public firms since we do not observe clustering behavior for private firms. The basic regression specification is:

$$\begin{aligned}
 y_{it} &= \alpha_i + Z'_{it}\phi + V'_t\omega + \varepsilon_{it} \\
 z_{it} &= \alpha_i + Z'_{it}\phi + V'_t\omega + \varepsilon_{it}
 \end{aligned}
 \tag{15}$$

In these regressions, y_{it} is an indicator variable which takes a value of one when a firm is a ‘follower’ firm and zero otherwise; analogously, z_{it} is an indicator variable which takes a value of one when a firm is a ‘counterfactual follower’ firm and zero otherwise. A follower firm is defined as any firm that announces a layoff in a 5-day window following a large layoff (including the day of the large layoff). Similarly, a counterfactual follower firm is a firm that announces a layoff in a 5-day window prior to a large layoff (excluding the day of the large layoff). The first set of regressions (specification (1) and (2)) is based on all firms in the dataset. This effectively treats all other Fortune 500 constituents as the control group. The second pair of regressions (specification (3) and (4)) restricts the sample to all firms that announce a layoff in year t . In effect, the control group in these specifications is the set of firms which also announced layoffs, but were outside the ten-day window which identifies firms as followers or counterfactual followers. In the last column (specification (5)), we restrict the sample to all firms that announce a layoff within 5 days (before or after) a large firm layoff announcement. In this last specification, the control group is simply

the counterfactual followers: firms that announced layoffs in the five days *prior* to a large-firm layoff.

In all five specifications, the vectors Z_{it} and V_t include the same control variables and explanatory variables used in Table 1.8 and Figure 1.5: total revenue, number of employees, years since IPO, book-to-market ratio, earnings-price ratio, leverage ratio, and days since last earnings announcement. We also include a full set of time-based (annual) fixed effects.

In a closely-related test, we examine whether clustering at these short horizons is driven by information. Specifically, we test whether firms change their clustering behavior after they start being covered by financial analysts. We estimate the same set of five specifications as described above, and include an indicator variable for analyst coverage as an explanatory variable. Using I/B/E/S sell-side analyst recommendations for U.S. stocks from 1993 to 2010, we construct an aggregate analyst coverage indicator variable. I/B/E/S codes recommendations from 1 (strong buy) to 5 (sell). We first restrict our sample to all firms that appear at least once in the I/B/E/S database. Next, we create an indicator variable, 'Past 3 years coverage,' which takes a value of 1 if an analyst covered by the I/B/E/S dataset made a recommendation in the previous three years. We report the results on determinants of follower behavior in Tables 1.11 and 1.12, with the former focusing on compensation structure and CEO tenure, while the latter focuses on analyst coverage.

6.3.2 Results. — The key independent variables in Table 1.11 are the same ones we used in the business cycle frequency results of public firms in Table 1.8: an indicator variable for short-tenured CEO and the average of the share of CEO compensation that derives from equity-linked compensation

over the past 5 years.

The results in this panel suggest that firms with a higher degree of equity-linked compensation are more likely to be followers; by contrast, this variable has no predictive power for counterfactual followers. Similarly, firms with short-tenured CEOs are more likely to be followers, while the same quality has a weak negative effect on the likelihood of being a counterfactual follower. These results are robust to the estimation using the full sample (specifications (1) and (2)), or the more restricted samples in specifications (3)-(5), although the coefficients are less precisely estimated due to reduced sample size.³⁰ Overall, we find that the same characteristics which predict layoff cyclicalities on a business-cycle level also predict strategic behavior over shorter horizons: firms with short-tenured CEOs and with significant equity-linked compensation are much more likely to act as follower firms, but not as counterfactual followers. This suggests an important role of reputation management in driving the high degree of observed ‘leader-follower’ behavior in layoff announcements.

To further explore the role of asymmetric information and reputation management, Table 1.12 reports our results from the analyst coverage regressions. The specifications here are identical to those in Table 1.11, except here the key explanatory variable is the indicator variable of analyst coverage. The results are broadly similar: when a firm is being covered by analysts, it is much less likely to announce layoffs after a large firm layoff. At the same time, the analyst indicator has no predictive power for counterfactual follower firms. We interpret these results as supporting the idea of

³⁰To ensure that the short-tenured CEO result is not being driven by the desire of newly-appointed CEOs to start with a ‘clean sheet’ we run the same regressions with a separate indicator variable for CEOs with tenure between 0 and 1 year. The results of this analysis are almost identical to those discussed here.

Table 1.11: Follower Characteristics, 1970-2010

In this table we report our results about the characteristics of public firms that layoff before and after the largest 20 firms in the economy as measured by previous year's revenue. Correspondingly the sample is restricted to the Fortune 500 constituents that are not the largest 20 firms as measured by revenue. The unit of observation is at the firm level observed over every business day between 1970 and 2010. There are two indicator variables that serve as dependent variables in the analysis. The follower indicator takes a value of one if in a given year the firm announces the layoff in a 5-day window following a large firm layoff (including the day of the large layoff). Similarly, the counterfactual follower indicator takes a value of one if in a given year the firm announces a layoff in a 5-day window prior to a large layoff (excluding the day of the large layoff). The first pair of regressions ((1) and (2)) are estimated using the entire sample (i.e. all Fortune 500 firms not in the top 20). The second pair of regressions ((3) and (4)) are estimated using only the sample of firms that laid off in a given year. By restricting the sample to this pool of firms, we are effectively relying on variation within firms that have relatively high propensity to layoff. In the last specification (5), we restrict the sample even further and include only firms that laid off within 5 days before or after a large firm announcement. Since in this case the two dependent variables are perfectly negatively correlated, we just need to estimate one regressions (instead of the pair). All the specifications include year fixed effects and firm level controls. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Dependent Variable	Counterfactual		Follower		
	Follower Indicator (1)	Follower Indicator (2)	Follower Indicator (3)	Counterfactual Follower Indicator (4)	Follower Indicator (5)
Sample	All Firms	All Firms	All Firms that Laid Off in Given Year	All Firms that Laid Off in Given Year	All Firms that Laid off within 5 days before/after large firm layoff
Avg. Equity-linked Compensation Share	0.0741*** (0.0186)	0.0186 (0.0165)	0.1850 (0.1546)	-0.1086 (0.1415)	0.4689* (0.2833)
Total Compensation	-0.0014*** (0.0004)	0.0004 (0.0011)	-0.0043 (0.0029)	0.0110*** (0.0025)	-0.0073* (0.0038)
Ceo Tenure = 0 - 4 years	0.0229*** (0.0075)	-0.0042 (0.0056)	0.1458** (0.0583)	-0.0945* (0.0496)	0.2534*** (0.0928)
Mean	0.0292	0.01763	0.2764	0.1645	0.6238
Std. Dev	0.1685	0.1316	0.4482	0.3715	0.4869
Year Fixed Effects	✓	✓	✓	✓	✓
Firm Level Controls	✓	✓	✓	✓	✓
Observations	2151	2151	228	228	101

asymmetric-information based factors driving the leader-follower behavior.

7 Linking High Frequency Results to Business Cycle Frequency Results

The high frequency results in the previous section offer strong evidence for strategic behavior, but on their own, they do not indicate the existence of significant welfare-relevant effects. In order to assess this, we investigate the impact of being a *follower* firm on changes in layoff propensity over the business cycle. Being a high-frequency follower in layoff announcements is, in principle, a better measure of active reputation management than other measures we have used (such as short-tenured CEO and equity-linked compensation). This is because we can rely on direct observations of reputation management behavior rather than predictions of such behavior. The results in this section seek to establish a connection between our high frequency results and the business-cycle-frequency results described in Section 5.

7.1 Empirical Strategy. — For this analysis, we first identify firms that have been ‘follower’ firms in the past five years, using the same methodology as in the previous section. This is our basis for our measure ‘Past 5 year follower,’ which takes the value of one for a specific firm in a given calendar month if, at any point over the prior five years, that firm has engaged in a layoff announcement within the five days following a layoff announcement by a large (i.e. top-20) firm. Analogously, we create a measure of ‘Past 5 year counterfactual followers,’ which takes a value of one for a specific firm in a given calendar month if, at any point over the prior five years, that firm has engaged in a layoff announcement within the five days prior to (and not

Table 1.12: Follower Characteristics, 1970-2010 (part II)

In this table we conduct the same analysis as in Table 8, except now our key dependent variable is a measure of analyst coverage. The sample is still restricted to the Fortune 500 constituents that are not the largest 20 firms as measured by revenue. The unit of observation is at the firm level observed over every business day between 1970 and 2010. There are two indicator variables that serve as dependent variables in the analysis. The follower indicator takes a value of one if in a given year the firm announces the layoff in a 5-day window following a large firm layoff (including the day of the large layoff). Similarly, the counterfactual follower indicator takes a value of one if in a given year the firm announces a layoff in a 5-day window prior to a large layoff (excluding the day of the large layoff). The first pair of regressions ((1) and (2)) are estimated using the entire sample (i.e. all Fortune 500 firms not in the top 20). The second pair of regressions ((3) and (4)) are estimated using only the sample of firms that laid off in a given year. By restricting the sample to this pool of firms, we are effectively relying on variation within firms that have relatively high propensity to layoff. In the last specification (5), we restrict the sample even further and include only firms that laid off within 5 days before or after a large firm announcement. Since in this case the two dependent variables are perfectly negatively correlated, we just need to estimate one regressions (instead of the pair). All the specifications include year fixed effects and firm level controls. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Dependent Variable	Follower Indicator (1)	Couterfactual Follower Indicator (2)	Follower Indicator (3)	Couterfactual Follower Indicator (4)	Follower Indicator (5)
Sample	All Firms	All Firms	All Firms that Laid Off in Given Year	All Firms that Laid Off in Given Year	All Firms that Laid off within 5 days before/after large firm layoff
Analyst Coverage in Past 3 years	-0.0254*** (0.0093)	-0.0068 (0.0060)	-0.2061** (0.0898)	-0.0476 (0.0777)	-0.1222 (0.1935)
Mean	0.0361	0.0246	0.2464	0.1836	0.5940
Std. Dev	0.1864	0.1550	0.4439	0.3874	0.4916
Firm Fixed Effects	✓	✓	✓	✓	
Year Fixed Effects	✓	✓	✓	✓	✓
Firm Level Controls	✓	✓	✓	✓	✓
Observations	7185	7185	964	964	436

after) a layoff announcement by a large firm.

7.2 Results. — The results of this analysis are reported in Table 1.13. We find that outside of recession months the propensity of follower and counterfactual follower firms are statistically indistinguishable. However, within recession months, a firm that has been a follower firm in the previous five years is roughly 3.14 percentage points more likely to engage in a mass layoff. By contrast, we find the impact of recessions on layoff propensity is not statistically different from zero for counterfactual follower firms. These results suggest that the firms we identify as ‘strategic’ over the very short horizons of our daily frequency data are also the firms that are more likely to have cyclical layoff policies over the course of the business cycle.

8 Alternate Explanations of Layoff Behavior

The results in Sections 5 and 6 are consistent with the predictions of the model of strategic corporate layoffs; however, several alternative theories also predict the temporal clustering of both mass layoffs and announcements of mass layoffs. In this section, we discuss whether these alternate theories are partially or fully consistent with the broad set of results presented so far. The leading alternative theories focus on common shocks, compassionate CEOs, management learning from other firms, and market inattention. Sections 8.1-8.4 discuss these alternative mechanisms and explore whether their implications match the empirical results. Notably, we do not seek to reject these alternate explanations; rather, we argue that none of them can explain the full range of results in the previous sections. We see this as strong support for the conclusion that reputation management in

Table 1.13: Linking Daily Frequency Layoff Behavior to Business Cycle Outcomes, 1995-2010

This table examines whether firms that engage in high frequency clustering of layoff announcements also are the ones that are more likely to layoff in recessions. The unit of observation is at the firm level tracked monthly between April 1995 and December 2010. In this table, a layoff constitutes a mass layoff as defined by the Bureau of Labor Statistics under its Mass Layoff Statistics Program (see section 4.2 for further discussion). The dependent variable is an indicator variable, which takes a value of one when a firm announces a layoff in a given month, and zero otherwise. In each regression specification the main explanatory variable is the interaction of the key variable with the recession indicator. The recession indicator takes a value of one in months classified as recession months by the NBER Business Cycle Dating Committee. The first key variable (specification (1)) is a follower indicator, which takes a value of one if in the past five years a given firm has announced a layoff in a 5-day window following a large firm layoff (including the day of the large layoff). Similarly, the counterfactual follower indicator (specification (2)) takes a value of one if in the past five years a given firm has announced a layoff in a 5-day window prior a large firm layoff excluding the day of the large layoff). All the specifications include month fixed effects, firm fixed effects, and firm level controls and the interaction of each of these controls with the recession indicator. The firm level controls include Earnings-Price ratio, Book-to-Market Ratio, Market Capitalization, Leverage Ratio, Years since IPO, which are all obtained from COMPUSTAT. Since the month fixed effects absorb all variation over time, the effect of a recession indicator is not separately identified in these specifications. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Sample		
Dependent Variable	Layoff Indicator (1)	Layoff Indicator (2)
Past 5 year Follower Indicator	0.0135*** (0.0035)	
Past 5 year Counterfactual Follower Indicator		0.0120*** (0.0036)
Recession × Past 5 year Follower Indicator	0.0314*** (0.0085)	
Recession × Past 5 year Counterfactual Follower Indicator		0.0041 (0.0095)
Mean	0.0722	0.0722
Std. Dev	0.2587	0.2587
Firm Fixed Effects	✓	✓
Month Fixed Effects	✓	✓
Firm Controls and Recession Interaction	✓	✓
Observations	70170	70170

the context of financial markets plays a significant role in determining the layoff behavior of large firms.

8.1 Common Shocks

The first alternative explanation for the patterns in Sections 5 and 6 is rooted in common shocks. If an aggregate shock hits a large subset of firms simultaneously, this will lead them to announce layoffs within a short period of time. A simple model with such common shocks will generate temporal clustering of layoff announcements, and also that of actual layoffs. A more sophisticated model of common shocks may generate excess sensitivity of firms with certain characteristics to these common shocks. If these firm characteristics are correlated with the structure of executive compensation or CEO tenure, we would not only observe temporal clustering of layoffs, but also that firms with greater equity-linked compensation or short-tenured CEOs are more sensitive to the common shocks.

While common shocks are certainly part of the story, the results in Figure 1.5 provide suggestive evidence against the common shock theory at a daily frequency. If either the crude or the more-sophisticated version of the common shock theory were true, we should observe responses in layoff announcements both before and after layoff announcements by large firms, both in the aggregate and within industries. However, the results demonstrate a stark asymmetry in the dynamic response of layoff announcements of Fortune 500 firms. In the business days leading up to a large firm announcement, the response of layoff announcements of the other firms is flat and close to zero. By contrast, it jumps up on the day of the large firm announcement, and gradually returns to zero in the next 4-5 business days.

Additionally, in the results of Table 1.11 we find that firms that announce layoffs within the five days following a large firm announcement are more likely to have greater equity-linked compensation and short-tenured CEOs. Conversely, no such association was found among the firms that lay off in the five days prior to a large firm's layoff announcement.

Over the longer horizons explored in our business cycle frequency data, our matching estimator results (based on size and four-digit industry) suggest a differential sensitivity of public firms' layoff behavior in response to recessions, compared to the behavior of matched private firms. Neither version of the common shock theory offers a clear prediction on the differential sensitivity of public and private firms, indicating that common shocks alone cannot explain the full range of our results.

8.2 Compassionate CEOs

Another mechanism that may generate the observed layoff behavior is that CEOs are compassionate and care about their labor force. This interpretation is compatible with the analysis in our model, but shifts the context of reputation from financial markets to the firm's employees. The degree of reputation concerns would therefore reflect factors such as altruistic motives or the CEO having strong ties to the labor force. Such motivations may lead the CEO to be biased against engaging in layoffs, and delay their layoffs until absolutely necessary. Consequently, CEOs with such characteristics will appear to pursue a cyclical layoff policy, or announce layoff announcements after other large firms have announced a layoff. Moreover, it is quite reasonable to expect that there may be strong correlations between a CEO's level of compassion for his employees, and his tenure and compen-

sation structure.

Testing the general validity of this mechanism is difficult because the level of compassion of a CEO may manifest itself in a number of different ways. We propose the following narrow test to identify the strength of such a mechanism in explaining our observed empirical patterns. If the CEO has spent many years at the firm (greater than 10 or 15 years) before being appointed its CEO (“home-grown CEOs”), then he is more likely to be compassionate. Conversely, CEOs that are externally-recruited or did not spend many years at the firm before being appointed to head it are less likely to be compassionate towards their labor force. We therefore evaluate the viability of this mechanism by testing its ability to explain the pattern of strategic firms having a greater sensitivity of layoff propensity to recessions.

We report the results of these tests in Table 1.14. The results suggest that home-grown CEOs who have been with the firm for many years are less likely to announce layoffs in general (although the point-estimates are imprecisely estimated). However, we find no effect of being home-grown on the cyclical nature of layoff announcements. Therefore, we conclude that though compassionate CEOs may alter the firm’s baseline layoff propensity, we find no evidence for this mechanism affecting layoff behavior differently over the business cycle. Thus, we cannot appeal to this mechanism in explaining the cyclical layoff behavior of the strategic firms in our analysis.

8.3 Market Inattention

Another alternate mechanism that can lead to ‘leader-follower’ behavior and potentially cyclical layoff policies is that firms are relying on market inattention to hide behind bad news. There are two versions of this alter-

Table 1.14: Compassionate CEOs and Layoff Announcements over the Business Cycle, 1970-2010

This table analyzes the impact of compassionate CEOs on layoff announcement propensity over the business cycle. The unit of observation is at the firm level tracked yearly between 1970 and 2010. The sample includes all contemporaneous constituents of Fortune 500 that ever announced a layoff. The dependent variable is an indicator variable, which takes a value of one when a firm announces a layoff in a given year, and zero otherwise. Our first measure of compassionate CEO is an indicator variable which takes a value of one, if the tenure of a given CEO at the firm before he was appointed CEO is greater than 5 years. This is used in specification (1) and (2). For specification (3) and (4) we change the cutoff from 5 years to greater than 15 years. In specifications (1) and (3) all firms in the sample are used for estimation, whereas in (2) and (4) we restrict the sample to firms that had no externally-appointed CEOs in the given year (i.e. only "home-grown CEOs). All the specifications include year-fixed effects, firm-fixed effects, and firm-level controls and the interaction of each of these controls with the recession indicator. The firm level controls include Earnings-Price ratio, Book-to-Market Ratio, Market Capitalization, Leverage Ratio, Years since IPO, which are all obtained from COMPUSTAT. Note that since we include year-fixed effects, the recession indicator cannot be separately identified in these regressions. Heteroskedasticity-consistent standard errors clustered at the firm level are shown underneath the coefficient estimates in all columns. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Dependent Variable	Layoff Indicator (1)	Layoff Indicator (2)	Layoff Indicator (3)	Layoff Indicator (4)
Sample	All Firms	Only Firms with "home-grown" CEOs	All Firms	Only Firms with "home-grown" CEOs
Recession x Years at Firm before CEO > 5 years	-0.0160 (0.0372)	0.0058 (0.0435)		
Years at Firm before CEO > 5 years	-0.0011 (0.0554)	-0.0194 (0.0588)		
Recession x Years at Firm before CEO > 15 years			-0.0129 (0.0413)	0.0048 (0.0437)
Years at Firm before CEO > 15 years			-0.1005* (0.0530)	-0.1240** (0.0573)
Mean	0.1438	0.1438	0.1438	0.1438
Standard Deviation	.3509	.3509	.3509	.3509
Firm Fixed Effects	✓	✓	✓	✓
Year Fixed Effects	✓	✓	✓	✓
Firm Level Controls and Recession Interactions	✓	✓	✓	✓
Observations	2295	1694	2295	1694

nate theory. The first relies on the market's underreaction to information caused by limited attention of market participants (Dellavigna and Pollet (2009)). In this version, there are certain days, e.g. Fridays, on which the market participants pay less attention to news, and therefore firms choose to release negative information on such days so as to reduce the adverse reputational effect. Thus, we would expect to observe firms clustering their layoff announcements around certain dates, but this would not be driven by interactions between firms. Rather, it would be the direct result of firms responding to common external drivers of market inattention.

The second version of the market inattention theory is that firms are following large firm layoff announcements because it allows their news article to be pushed to the back pages of the newspapers (or analogously gain less prominence in televisions news or other media). If market participants have some information processing cost, they are less likely to chance upon this negative news, allowing the firms to release negative news in a relatively 'concealed' manner.

The high-frequency event study results presented in Figure 1.5 already controls for day of week and calendar month to control for predictable market inattention. In addition, the regressions also control for whether the daily observation occurs within a week (before or after) the firm's scheduled earnings announcement date. Correspondingly, the leader-follower behavior observed in Panels A and C of Figure 1.5 stems from mechanisms other than those suggested by the first version the market inattention theory. As for the second version, we refer to the results of Panel D. If firms were trying to hide behind other negative news, we should find the same mechanism to hold after days of major negative non-economic news (e.g.

Hurricane Katrina). The results presented in the last panel, however, illustrate that there is no systematic change in layoff announcement behavior either before or after such major negative non-economic news. In addition, we separately conduct an analysis of the page of the Wall Street Journal on which each layoff announcement was originally reported. We find no significant difference in the placement of layoff announcement coverage on days of layoff announcements by large firms, when compared to other days. While it is possible that market inattention may play a role in determining the timing of layoff announcements, we are unable to find any evidence for this in our analysis, and cannot appeal to this mechanism to explain our results at either the daily or the business cycle frequency.

8.4 Learning from Other Managers

A final alternate mechanism that may lead to 'leader-follower' behavior is that managers are uncertain about the state of the aggregate economy, and they are waiting to receive a signal from the actions of the largest firms in the economy. By virtue of being larger, the managers of the largest firms may have better information about the aggregate state. Consequently, managers learn about the aggregate state from the performance of large firms, and respond to layoff announcements by large firms with layoffs of their own. This mechanism would also predict that short-tenured CEOs would be more likely to react to the announcements of large firms, as they are more likely to be inexperienced, and thus more reliant on learning from other managers.

Despite its intuitive appeal, this mechanism cannot account for the differences we observe between public and private firm layoff behavior. If

the ‘learning-from-others’ theory is the primary driver of layoff policy, we should observe little difference between public and private firms, particularly when matching on size and industry. We would expect all firms within an industry to face similar market conditions regardless of their ownership status, so the optimal response to learning about changes in market conditions should be identical across the two groups.

Moreover, the stock penalty results we find in Table 1.9 fit squarely with our model, but cannot be explained by this alternate mechanism. In particular, the reputation-based mechanism has strong predictions about the cross-effects of one firm’s action on the reputation penalty of other firms in the same industry (e.g. mitigated reputation penalty for layoffs followed by a large-firm layoff). Our empirical findings confirm these predictions. By contrast, the learning-based mechanism has limited predictions with respect to such cross-effects. In particular, mechanisms based on learning should be largely predictable, and therefore subsequent layoffs should have no observable (abnormal) cross effects on stock prices. In sum, to explain the full range of our empirical findings we need to appeal to reputation-based mechanism such as the one described by our model.

8.5 Alternate Mechanisms: Taking Stock

The key conclusion of this section is *not* that the alternate mechanisms discussed here play *no* role in the high-frequency clustering of layoff announcements or the cyclicity of layoffs at the business-cycle level. Instead, we conclude that the main results of this paper cannot be fully explained solely by these alternate mechanisms. Instead, we argue that reputation concerns in the context of financial markets represent the most salient ex-

planation for the patterns we observe.

9 Conclusion

In this paper we document that there is excess clustering of layoff announcements at weekly and daily horizons. We interpret this clustering of announcements within a theoretical framework in which managers delay layoffs during good economic states to avoid damaging the market's perception of their ability. We test the implications of our model at both daily and business-cycle frequencies using two novel datasets. Using a pairwise matching estimator based on size and three-digit industry, we find that the layoff propensity of public firms is twice as sensitive to recessions, relative to their matched private counterparts.

In a range of robustness tests we show that these differences are not being driven by public-private differences in lifecycle effects, leverage, workforce size, or on our matching criteria. Within our sample of public firms, we find that firms predicted to be more strategic by our theoretical analysis, are also the ones more likely to engage in mass layoffs during recessions. Our results therefore suggest that reputation management is an important driver of the observed differences in the cyclicity of layoffs between public and private firms.

At the daily frequency, we also find significant support for our model. We show that a large firm announcement (i.e. the 20 largest firms based on past year's revenue) is associated with future layoffs by other Fortune 500 firms, but not with past layoffs. We find that this effect is twice as strong if the large firm is in the same industry as the follower firm. For our sam-

ple of privately held firms we find no such clustering behavior either before or after the large-firm layoff announcement. Moreover, when we compare the characteristics of firms that lay off in the five days after a large-firm announcement (“followers”) to those that lay off in the five days before a large-firm announcement (“counterfactual followers”), we find that follower firms have a greater likelihood to be managed by short-tenured CEOs (i.e. with a tenure between 0 and 4 years) and to place greater reliance on equity-linked compensation for their CEOs. Consistent with our theoretical framework, these results suggest that reputation management is an important driver for the timing of layoff announcements at high frequencies.

Lastly, we establish a connection between high-frequency clustering and layoff behavior at business cycle frequencies. We find that the follower firms are roughly three percentage points more likely to engage in a mass layoff during a recession month, compared to counterfactual follower firms. This link between the daily frequency reputation management and the cyclicity of layoffs over the business cycle provides significant evidence that reputation concerns are an important driver of firms’ layoff policies. Taken together, the findings of this paper suggest that managerial behavior not only has costs for the individual firm, but also has significant aggregate implications at the business cycle frequencies.

Taken together, the findings of this paper indicate that reputation management in financial markets may strongly impact the real decisions of firms, particularly with respect to labor decisions. At the firm level, this mechanism can lead to significant deviations from optimal policy by delaying the termination of unprofitable projects. In addition, this mechanism also impacts industry dynamics by influencing the timing of layoffs by other firms

in the same industry.

This paper suggests that understanding how distortions in managerial behavior impact industry dynamics can lead to a number of implications and directions for future research. First, to understand industry dynamics one should place added emphasis on shocks to large players such as Walmart, Ford, or Nokia. For example, in our analysis, layoffs by one of these firms may trigger clustering of layoffs by other firms in the economy. Second, this paper highlights a novel mechanism through which corporate governance can not only influence the financial decisions, but also the real decisions of firms and industries. Third, our research suggests that the timing of other forms of corporate disclosure, such as dividend cuts and writedowns, may also be subject to clustering, and have implications for determining firm value and performance. Future papers can take advantage of the empirical methodologies used here to understand the mechanisms driving disclosure dynamics.

Chapter 2: Uncertainty and the Special Role of Safe Assets

1 Introduction

Money serves three functions: medium of exchange, unit of account, and safe store of value. This paper is exclusively concerned with the last function: safety. The central premise of the paper is that the quantity of money-like instruments that provide safe storage between periods, crucially determines households' appetite for holding risky assets. That is, the risk premium demanded for holding risky assets will be lower when there is a greater quantity of money-like claims in households' portfolio. Using this novel connection between money and risk premia, I characterize how central bank policy affects the real economy. The key insight is that by regulating the quantity of federally-insured deposits available to households, central banks directly control the quantity of safe assets available to households. In such an economy, the role of central banking authority becomes stronger when private agents' (such as banks) ability to produce safe assets weakens. This weakening may stem from agency frictions, e.g. incentives for banks to cheat by selling more safe securities than they actually own. I present a model that highlights why the banking sector fails to self-regulate itself and consequently it is socially optimal for the central bank authority to play an active role in regulating the banking sector. In particular, I show that these interventions will mirror standard policy tools in modern economies: deposit-insurance and reserve requirements. Using these insights, this pa-

per attempts to provide a theory of central banking in a flexible-price world with a strong emphasis on the mechanics and regulation of safe asset creation (by traditional banks and shadow banks).

The next few paragraphs discuss the mechanics of my model to facilitate understanding of the novel features of the model. The model in the paper presents an economy in which there are two types of agents, risk-averse households and risk-neutral banks. These agents have access to projects (or assets) with returns that can be decomposed into a safe and a risky component. Since the households are risk-averse they will prefer to hold safer assets instead of the original projects which have a risky component. The risk-neutral banks recognize this shortage of safety, and sells money-like safe instruments to households to offer them insurance. How can the banking sector create these safe instruments? The simple answer is capital structure. The banks will find it optimal to buy the original projects (assets side) and then tranche and sell securities that are backed by the safe component of these assets (liability side). This tranching technology essentially enables the banks to create 'inside money' and increase the degree of risk-sharing in the economy. The banks benefit from selling these safe securities, since in turn they can create more assets with the proceeds of the sale of the tranced safe security. This 'safety transformation' role of banking effectively makes the household more insured, thereby lowering the risk premium households demand for holding risky securities.

In such a setup, the fundamental friction I introduce is the inclination of banks to create too many safe securities. This effectively captures one of the most perennial problems of the banking sector. Historically, the unregulated banking sector has had strong incentives to create too many demand

deposits that are riskless (nominally), than the safety of their assets permit. Such incentives were forcefully highlighted during the financial crisis of 2007-09 in which the shadow banking sector bundled together substantial tail risk in their super-senior AAA-rated securities. This friction lowers the degree of risk-sharing in the economy relative to the optimal frictionless benchmark, leading to an increase in risk premium.

In presence of such frictions, I show that the optimal policy is for the central bank authority to offer deposit insurance to households for the safe securities issued by the banks. In addition, to ensure that this deposit insurance program is self-financing (i.e. to limit the burden to the banking sector), the central bank must regulate the quantity of safe securities banks are permitted to issue. This policy ensures that the central bank will be able to tax the banks to finance the deposit insurance in all states of the world. Thus, simultaneously offering deposit insurance, and limiting the quantity of safe securities banks issue enables the central bank to restore the optimal frictionless benchmark. The immediate question is what instruments do central banks have to regulate the quantity of safe securities? I show that either of the two following instruments can be used to achieve this: 1) a permit system in the form of reserve requirements, or 2) a limit on the degree of leverage banks undertake.

This gives us a theory of how central banks in practice may influence risk premia. When the central bank drains reserves from the system, they force intermediation activity to take place outside the regulated banking sector. Correspondingly, the shadow banking sector expands. The shadow banks, however, are unregulated, and therefore are unable to commit to sell completely safe securities without bundling together unsafe tranches. This

effectively leads to a rise in the risk premia as the degree of risk-sharing in the economy falls. Thus, the main idea here is that private money creation by unregulated financial players is inefficient, and by draining reserves the central bank redirects households to these unregulated players, away from the safety provided by federally-backed demand deposits.

The logic of safe assets also highlights that inflation uncertainty can be very costly from a welfare perspective. When there is substantial inflation uncertainty, nominal debt instruments that were otherwise safe (e.g. money) cease to be safe. This in turn can increase risk premia in the economy. This may potentially explain the large bond risk premia observed during episodes of high inflation uncertainty. Thus, if the active management of reserve requirements leads to a rise in inflation uncertainty, the central bank authority may face a tradeoff between keeping inflation uncertainty low, or having a larger quantity of safe securities be created through the regulated banking sector.

In the last theory section, I discuss how counterparty risk can lead to a shutdown of the shadow banking sector, and how the central bank can restore optimal allocations by buying up risky assets in exchange for safe claims. Using their power of taxation, the central bank can credibly commit to pay back the safe claims. I discuss how the effect of such a policy depends crucially on who bears the tax burden. If the risk-averse households are likely to bear the tax burden, then such a policy is unlikely to reduce risk premia in the economy. On the other hand, if the burden of taxation is levied on risk-neutral shadow bankers, then this can bring about a Pareto improvement, thereby reducing risk premia.

I conclude the paper by presenting some measurements of quantity of

safe assets created by the banking sector that are held outside the banking sector. One such measure is M3, which is one of the broadest measures of the stock of money in circulation. Since the Federal Reserve stopped publishing this measure in 2006, I reconstruct this measure for the period 2006-10. Incidentally, just after the Fed stopped publishing this measure, there was an unprecedented growth in M3 between 2006-08. Most of this growth can be attributed to the money market mutual funds that were central to financing the operations of the shadow banks by purchasing asset-backed commercial paper and entering repurchase agreements. Analogously, there was a sharp contraction in this broad-money measure in 2008 after Lehman's failure, which has continued into 2010. The magnitude of this contraction (roughly 15% between 2008 and 2010) mirrors the contraction of measures of broadly-defined money after the Great Depression (roughly 18% between 1929 and 1932), which continued into the mid 1930s. I also present other evidence on the operation of the traditional and shadow banking sectors during the recent financial crisis. The central message of this evidence section will be that to understand central banking in the modern economy, it is essential to take a broad view of money, which includes private money creation by the shadow banking sector.

2 Safe Assets in General Equilibrium

2.1 Setup

2.1.1 Preliminaries. — Consider an economy with two periods and a single consumption good. Agents have endowments in period 1 but they only consume in period 2. There is a risky asset which is supplied elastically in

period 1, and which pays off in period 2. In particular, one unit of the asset pays

$$\bar{s} + v,$$

where $\bar{s} \in (0, 1)$ is constant (the safe portion) and v is a nonnegative random variable with expected value equal to $1 - \bar{s}$ (so that the asset's expected payoff is equal to 1, which is a normalization that simplifies the analysis). Here \bar{s} captures the fraction of the asset payoff that is safe and v captures the uncertainty in the asset payoff. Note that so far this is a real economy, and therefore considerations such as inflation uncertainty that may make \bar{s} also risky are suppressed. Such cases will be discussed in subsequent sections.

There are two types of agents. There is a measure one of households who start with w^H units of endowment in period 1, and who are risk averse with utility function $u(\cdot)$, which is continuous, strictly-concave and monotone. There is also a measure one of banks who start with w^B units of endowment in period 1, and who are risk neutral. The banks are in perfect competition and provide insurance to households. They can potentially invest in the risky asset (asset-side) and issue a relatively less risky security (liability-side), which they then sell to the households.

Here it is worth being clear about the interpretation of banks in the model. In the real world, banks do not directly invest in physical projects but rather lend to firms who in turn do the project selection. Abstracting away from this extra layer is equivalent to assuming that there is no contracting friction between firms and banks. The defining feature of banks here is that they engage in tranching (i.e. issuance of structured senior securities) similar to the models in Gorton and Pennacchi (1990), DeMarzo and

Duffie (1999), DeMarzo (2005)³¹. Also, here the assumption of risk neutrality on the part of banks (instead of relatively low risk aversion) is made for convenience. You may interpret this difference in risk preferences arising from a case in which the owners of banks are rich individuals, whereas the household sector is comprised of relatively poorer households.

Formally, let T denote the set of feasible contracts (tranches) which the bank might issue against one unit of investment (which includes also the original asset, $\bar{s} + v$). More specifically, each $\tilde{t} \in T$ is a random variable that describes the tranche's payoff in each state of the world. Many of the securities in T will not be issued in equilibrium (in fact only one will be issued). Nonetheless, each $\tilde{t} \in T$ has an equilibrium price, which we denote by $p(\tilde{t})$. Here it is implicitly assumed that if households want to buy the original asset it must go through the banking sector. Also, the original risky asset will serve as our reference security, and the price of all other securities will be expressed relative to the price of this security, $p(\bar{s} + v)$. For all subsequent exposition, note that the use of tilde is merely to represent choice variables in the respective agent's maximization problem. Correspondingly, whenever the choice variable is presented without tilde, it represents the agent's optimal choice.

2.1.2 Households. — Let $z(\tilde{t})$ denote the households' holding of tranche \tilde{t} in period 1, and let c^H denote the random variable that represents the households' consumption in period 2. In equilibrium, households will only hold a single tranche (as noted above). However, households' portfolio choice enables us to price the securities in T . In particular, the households' problem

³¹For realism I could easily introduce the role of pooling risk (i.e. diversification) in addition to tranching, by introducing idiosyncratic shocks. But since it does not add much to the analysis I have suppressed this feature.

is given by:

$$\begin{aligned}
& \max_{\{z(\tilde{t})\}_{\tilde{t} \in \{T\}}, c^H} E \left[u \left(c^H \right) \right] & (16) \\
\text{s.t.} \quad & c^H = \int_{\tilde{t} \in T} z(\tilde{t}) \tilde{t} d\tilde{t} \\
& \int_{\tilde{t}} z(\tilde{t}) p(\tilde{t}) d\tilde{t} \leq w^H.
\end{aligned}$$

Using the households' first order condition, the price $p(\tilde{t})$ of a tranche \tilde{t} must satisfy

$$\frac{p(\tilde{t})}{p(\bar{s} + v)} \leq \frac{E[u'(c^H) \tilde{t}]}{E[u'(c^H) (\bar{s} + v)]} \quad (17)$$

This condition will hold with equality if there is an interior solution, i.e. households hold positive quantities of $(\bar{s} + v)$. Note also that, when the prices satisfy the first order condition (17), then an allocation $\{z(\tilde{t})\}_{\tilde{t} \in \{T\}}$ is optimal for the households iff it satisfies the households' budget constraint in (16) with equality.

2.1.3 Banks. — In principle, the bank could issue different types of tranches for different units of investment. However, in our setting, it is sufficient to focus attention on allocations in which banks issue the same tranche, t , for each unit of investment³². Denoting the banks' investment level by k , its problem can be written as:

$$\begin{aligned}
& \max_{\tilde{k}, \tilde{t} \in T} E \left[\tilde{k} (\bar{s} + v - \tilde{t}) \right] & (18) \\
\text{s.t.} \quad & \tilde{k} \leq w^B + \tilde{k} p(\tilde{t}).
\end{aligned}$$

³²This is because there is no heterogeneity among the households, and thus they choose identical tranches. Allowing for heterogeneity in risk aversion will introduce the possibility of banks issuing multiple tranches, but such considerations are not addressed here.

That is, banks choose their investment level, k , and the tranche, t , to maximize expected utility (taking equilibrium prices as given).

For the market for tranches to clear, consumers must hold the tranche that is offered by banks. That is, they should choose $z(t) = k$ and $z(\tilde{t}) = 0$ for any $\tilde{t} \neq t$. Note, given that prices satisfy the inequality in (17), this allocation is optimal for consumers if it satisfies their budget constraint with equality, which can be written as:

$$kp(t) = w^H. \quad (19)$$

We are now in a position to define the equilibrium.

2.1.4 Equilibrium. —

DEFINITION 1 *An equilibrium in this economy is a collection of allocations, (t, k, c^H, c^B) , and prices $\{p(\tilde{t})\}_{\tilde{t} \in T}$ such that:*

1. *Consumers choose their allocation optimally, that is, prices satisfy (17) and the budget constraint (19) holds.*
2. *Banks optimize, that is, (k, t) solves problem (18).*

2.2 Frictionless Benchmark

First suppose there are no contracting or agency frictions and banks can issue any security subject to feasibility constraints. In this case, the set T can be represented as follows (because all possible securities can be reproduced using such tranches):

$$T = \{\tilde{s} + \tilde{\eta}v \mid \tilde{s} \in [0, \bar{s}] \text{ and } \tilde{\eta} \in [0, 1]\}.$$

Hence, banks' choice of contract is reduced to the choice of two weights, $\tilde{s}, \tilde{\eta}$. To see this note that this linear contract spans the contingent-state space, since it allows the banks to offer a contract contingent on each realization of v (and \bar{s} which is a constant).

Under certain parametric conditions (which will be specified below), we conjecture an equilibrium in which banks choose

$$s \in [0, \bar{s}] \text{ and } \eta = 0, \quad (20)$$

That is, banks choose $t = s$, where $s \in [0, \bar{s}]$ will be determined below. In this conjectured competitive equilibrium, the zero profit condition will apply, and thus for any security t we must have

$$p(t) = E[t] = s$$

(where we have used the assumption that the expected value of the risky asset is 1).

Plugging these prices back into the inequality on prices imposed by household's first order condition (17) we know that these prices satisfy the household's problem. Hence, the conjecture in (20) is optimal for households.

Plugging in these prices, banks' problem (18) can be written as:

$$\begin{aligned} & \max_{k, \tilde{s}, \tilde{\eta}} E [k (\bar{s} - \tilde{s} + (1 - \tilde{\eta}) v)] \\ \text{s.t.} \quad & k \leq w^B + kE [\tilde{s} + \tilde{\eta}v]. \end{aligned}$$

Note that in equilibrium the banks do not profit from intermediation.

Therefore, the banks are indifferent between any $\tilde{s}, \tilde{\eta} \in [0, \bar{s}] \times [0, 1]$.

Hence, the conjecture in (20) is also optimal for banks. In the conjectured allocation, banks' investment is then given by:

$$k = \frac{w^B}{1 - p(s)}.$$

The only remaining condition to check is the households' budget constraint (16), which can be written as:

$$\frac{w^B}{1 - p(s)} p(s) = w^H.$$

By rewriting we get $p(s) = w^H / (w^B + w^H)$. Due to the zero profit condition, we have $p(s) = s$, and substituting this gives us a unique solution for s :

$$s = \frac{w^H}{w^B + w^H}$$

The conjectured allocation is indeed an equilibrium if and only if

$$\bar{s} \geq s \tag{21}$$

That is, the conjectured equilibrium exists if the security has a sufficiently large safe portion and/or if the banks' endowment (capital) is sufficiently large. To recap, the results for the frictionless benchmark are summarized in the following proposition.

PROPOSITION 1 Under the parametric condition (21), there is an equilibrium in which banks invest in the original asset and issue a safe security, $s \in [0, \bar{s}]$. The households are fully insured since they buy money-like debt claims from the

banks, and the banks retain the risky tranches on their balance sheet. The equilibrium allocations are given by

$$\begin{aligned}
 s &= \frac{w^H}{w^B + w^H} \\
 k &= w^B + w^H \\
 c^B &= (w^B + w^H) [\bar{s} + v] - w^H \\
 c^H &= w^H
 \end{aligned}$$

2.3 Equilibrium with Agency Frictions and Complexity

Let us next introduce the key friction and analyze the resulting equilibrium, with two assumptions. Let λ be a positive but small constant (i.e. closer to zero than to one).

Assumption (A1). For each $\tilde{s} \in [0, \bar{s}]$, households are unable to tell the difference between the security \tilde{s} and the security $(1 - \lambda) \tilde{s} + \lambda \tilde{s} \frac{v}{E[v]}$.

Note that the corresponding security has two features: (i) it has the same expected payoff as the original security \tilde{s} , (ii) it is riskier than the original security \tilde{s} . In words, this assumption posits that, households are unable to tell the difference between a completely safe security and a security that is slightly risky. This can be motivated by complexity of the tranced securities. We also make the following assumption, which generates an agency friction in our setup.

Assumption (A2). After contracting with households, a bank that chooses $t = \tilde{s} \in [0, \bar{s}]$ is unable to commit to not replace the security.

Denote bank's profits as ω . Then from the bank's first order condition, taking prices as given we get $\partial E[\omega] / \partial s \geq 0$, as long as $p(1) \geq 1$. This condition always holds since households are risk-averse and are willing to

pay a premium for safe securities. Thus, the banks always have incentive to replace the security (because the replaced security allows the banks to sell more safe securities). Off the equilibrium path, when households don't take this replacement into account, this allows banks to create $\lambda\bar{s}$ units of additional securities that have face value \bar{s} , but is effectively backed by payoff $(1 - \lambda)\bar{s} + \lambda\bar{s}\frac{v}{E[v]}$. Hence, the key implication of these two assumptions is that banks have incentives to create too many safe claims, which effectively reduces the safety of each claim they issue.

However, households will take this into account when they price the security. Hence, these assumptions effectively reduces the set of securities to:

$$T^f = \left\{ (1 - \lambda)\bar{s} + \tilde{\eta}v \mid \bar{s} \in [0, \bar{s}] \text{ and } \tilde{\eta} \geq \frac{\lambda\bar{s}}{1 - \bar{s}} \right\}.$$

After replacing T with T^f , the rest of the equilibrium definition is identical. We next conjecture an equilibrium in which the replacement constraint binds, that is, banks choose:

$$t = (1 - \lambda)s + \frac{\lambda s}{1 - \bar{s}}v = \left[(1 - \lambda) + \frac{\lambda}{1 - \bar{s}}v \right] s \quad (22)$$

Note that the expected payoff of this security is s , since $E[v] = 1 - \bar{s}$. But note also that this security is risky, unlike the frictionless benchmark. Hence, households' consumption is given by, $c^H = kt$, which is riskier than the frictionless benchmark. This implies that the risk premium in this economy may be greater than the frictionless benchmark since expected marginal utility will be higher.

Once again, due to the zero profit condition we have

$$p(t) = E[t] = s$$

(where we have used the assumption that the expected value of the risky asset is 1). These prices satisfy the inequality (17) and therefore, the conjecture (22) is optimal for households.

Plugging in these prices, banks' problem (18) can be written as:

$$\begin{aligned} & \max_{k, \tilde{s}, \tilde{\eta}} E[k(\tilde{s} - \tilde{s} + (1 - \tilde{\eta})v)] \\ \text{s.t.} \quad & k \leq w^B + kE[\tilde{s} + \tilde{\eta}v]. \end{aligned}$$

Note that the banks are again indifferent between any $\tilde{s}, \tilde{\eta} \in [0, \bar{s}] \times [0, 1]$. This is because in equilibrium the banks do not profit from intermediation. Hence, the conjecture in (22) is also optimal for banks. In the conjectured allocation, banks' investment is then given by:

$$k = \frac{w^B}{1 - p(t)}.$$

The only remaining condition to check is the households' budget constraint (16), which can be written as:

$$\frac{w^B}{1 - p(t)} p(t) = w^H.$$

By rewriting we get $p(t) = w^H / (w^B + w^H)$. Using the zero profit

condition this solves uniquely for s :

$$s = \frac{w^H}{w^B + w^H}$$

Again the conjectured allocation is indeed an equilibrium if and only if

$$\bar{s} \geq s \tag{23}$$

These results are summarized in the following proposition.

PROPOSITION 2 *Under the parametric condition (21), and assumptions (A1), (A2), there is an equilibrium in which banks invest in the original asset and issue a quasi-safe security, $s + \eta v$. The households are not fully insured since the banks are unable to commit to not replace the safe security. The equilibrium allocations are given by*

$$\begin{aligned} s &= \frac{w^H}{w^B + w^H} \\ k &= w^B + w^H \\ t &= (1 - \lambda) s + \frac{\lambda s}{1 - \bar{s}} v \\ c^B &= (w^B + w^H) [\bar{s} + v] - w^H \left[(1 - \lambda) \cdot 1 + (\lambda) \frac{v}{E[v]} \right] \\ c^H &= w^H \left[(1 - \lambda) \cdot 1 + (\lambda) \frac{v}{E[v]} \right] \end{aligned}$$

The key implication of this proposition is that when all banks issue quasi-safe securities (instead of fully safe), households are less insured and consequently the risk premium in the economy goes up.

To gain more intuition, suppose we start from the allocation of the economy with frictions and suppose we remove the frictions. Then, a bank at

the margin has an incentive to offer a safer security (because consumers are willing to pay an even greater premium for that security). Hence, the above equilibrium unravels. But when all banks do that, consumers become more insured and the risk premium diminishes. In the resulting (frictionless) equilibrium, the consumers are better off.

The main message of this analysis is that agency frictions may hinder the degree of risk-sharing between risk-neutral banks and risk-averse households. In particular, the banks will always have incentives to bundle together some risk with the safe securities, which in equilibrium will make households less insured. This agency friction that prevents optimal risk sharing creates room for regulation to restore the optimal allocation. The next two sections discuss the role of regulation.

2.4 Why Self Regulation Fails?

This subsection discusses why the banking sector as a whole may not want to regulate themselves, even if they had access to a technology that will completely eradicate the incentive to cheat among *all* banks. To see this we need to modify our setup in a way that allows banks to make some profits so as to cleanly understand their incentives to self-regulate. The way I do this here is to introduce an interim period of trading. The characteristics of the interim period are summarized in the following definition.

DEFINITION 2 *The interim period has the following properties:*

1. *It occurs after period 1 contracts are written, but before period 2 uncertainty is realized.*
2. *With probability α , households receive an endowment of w_1 units of original*

assets that pay $(\bar{s} + v)$ in period 2.

- 3. Banks can make a take-it-or-leave-it offer to sell tranches to households in exchange for the endowments households receive.*

The rationale behind the third property is that the quantities of securities cannot be adjusted in the interim period (since they are fixed in period 1). This gives the banks a temporary market power since supply is inelastic in the interim period. In addition I impose the following assumption, which makes the analysis simpler.

Assumption (A3). *Households cannot write period 1 contracts contingent on endowments that they may receive in the interim period.*

This assumption leaves household's actions in period 1 unaltered. To see this note that households have no incentives to hoard resources to trade in the interim period, since they can always get a better deal in period 1 when the banking sector is competitive. Also due to assumption (A3), they cannot buy additional securities based on the interim period endowments. This gives us the following lemma, which simplifies exposition.

LEMMA 1 *Under the definition of interim period, assumption (A3), and when w_I is sufficiently small :*

- 1. Households' period 1 problem and the interim problem are independent.*
- 2. Banks' period 1 problem and the interim problem are independent.*

PROOF. (see Appendix). ■

Note here the assumption w_I being sufficiently small is needed to ensure that the banks have sufficiently large quantity of safe assets on their

balance sheet to sell to households in the interim period. It will turn out that the parametric restriction on w_I will be $w_I \leq (w^B + w^H) \bar{s} - w^H$.

Since the banks are able to make a take-it-or-leave-it offer, they will find it optimal to set prices such that households are indifferent between trading with them (which will allow the banks to extract all the surplus from the trade). This implies that the prices in the interim period will be pinned down by the household's first order condition. Let's denote the interim prices as $p_I(\tilde{t}_I)$ for tranche \tilde{t}_I .

2.4.1 Households. — Based on the lemma, we can write household's interim period problem as:

$$\begin{aligned} & \max_{\{z(\tilde{t}_I)\}_{\tilde{t}_I \in T}, c^H} E \left[u \left(c^H \right) \right] \\ \text{s.t.} \quad & c^H = kt + \int_{\tilde{t}_I \in T} z(\tilde{t}_I) \tilde{t}_I d\tilde{t} \\ & \int_{\tilde{t}_I} z(\tilde{t}_I) p_I(\tilde{t}_I) d\tilde{t}_I \leq w_I p_I(\bar{s} + v) \end{aligned}$$

Under the assumption of take-it-or-leave-it contracts, the price $p_I(\tilde{t}_I)$ of a tranche \tilde{t}_I must satisfy

$$\frac{p_I(\tilde{t}_I)}{p_I(\bar{s} + v)} = \frac{E \left[u' \left(c^H \right) \tilde{t}_I \right]}{E \left[u' \left(c^H \right) (\bar{s} + v) \right]} \quad (24)$$

2.4.2 Banks. — The banks take these prices as given, and decide which securities to trade in exchange for the w_I units of assets, each of which will pay $\bar{s} + v$ in period 2. The banks' problem is then simply given by:

$$\max_{\tilde{t}_I \in T} E \left[(\bar{s} + v) - \frac{p_I(\bar{s} + v)}{p_I(\tilde{t}_I)} \tilde{t}_I \right]$$

That is, banks choose a tranche t_I that they exchange with households for one unit of the original asset that pays $(\bar{s} + v)$.

For the market of tranches to clear, consumers must hold the tranche that is offered by banks. That is, they should choose $z(t_I) p_I(t_I) = w_I p_I(\bar{s} + v)$, and $z(\tilde{t}_I) = 0$ for any $\tilde{t}_I \neq t_I$. Note that given the prices satisfy household's first order condition, this allocation is optimal for consumers if it satisfies their budget constraint with equality.

2.4.3 Solution. — Since the prices are pinned down by household's problem (due to the take-it-or-leave-it offer), households are indifferent between any tranche \tilde{t}_I . Therefore, the optimal tranche is pinned down by banks' first order condition taking household's problem as given. Choosing \tilde{t}_I is equivalent to choosing $(\tilde{s}_I, \tilde{\eta}_I)$. The first order condition, with respect to $\tilde{\eta}$, at $\tilde{\eta} = 0$, is negative:

$$\left. \frac{\partial E[(\bar{s} + v) - \{p_I(\bar{s} + v) / p_I(\tilde{t}_I)\} \tilde{t}_I]}{\partial \tilde{\eta}} \right|_{\tilde{\eta}=0} \leq 0$$

Therefore, the optimal tranche is $\tilde{t}_I = s_I$, $s_I \in [0, \bar{s}]$.

In equilibrium the budget constraint of households hold with equality, which pins down total resources transferred from banks to households in period 2, in exchange for w_I units of the original asset:

$$s_I z(s_I) = w_I \frac{p_I(\bar{s} + v)}{p_I(1)}$$

Therefore, the solution to this problem is that households give up their endowment that pays off $w_I(\bar{s} + v)$ units in period 2, in exchange for securities that pays off $w_I p_I(\bar{s} + v) / p_I(1)$ units in period 2.

2.4.4 Frictionless Case. — In the frictionless benchmark, the households

are perfectly insured, and therefore their first order condition (24) gives us that the relative prices $p_I(\bar{s} + v) / p_I(1) = 1$. Therefore total resources households receive from trading with banks in the interim period is:

$$s_I z(s_I) = w_I$$

This indeed is an equilibrium as long as the banks have enough safe resources to sell to households in the interim period: $w_I \leq (w^B + w^H)(\bar{s} - s)$. Replacing s , we get

$$w_I \leq (w^B + w^H)\bar{s} - w^H$$

2.4.5 Case with Frictions. — In the case with frictions as specified in (A1),(A2), the relative prices $p_I(\bar{s} + v) / p_I(1) < 1$. This is because household consumption is risky, and therefore they are willing to pay a premium to get additional riskless securities in the interim period. This implies that the total resources households receive is less than w_I , i.e. less than the frictionless benchmark:

$$s_I z(s_I) < w_I$$

Equivalently, the banks get greater profits in the case of frictions than no frictions. This implies that the banking sector as a whole will be unwilling to regulate themselves in order to eliminate the frictions. The following proposition summarizes these results.

PROPOSITION 3 *Under the conditions given in lemma 1, and in presence of frictions given in A1-A2:*

1. *An individual bank has incentives to eliminate the agency problem (A2) in period 1.*

2. *The banking sector as a whole has no such incentives to eliminate the agency problem (A2) in period 1.*

The main message of this proposition is that the banking sector as a whole is more profitable due to the agency frictions. This is because the agency friction acts as a coordination device to sell less insurance to households, which in turn makes the premium for safe assets higher in the interim period. This is somewhat counter-intuitive since agency frictions usually makes the agent worse-off, but here due to the general equilibrium effect arising from risk-sharing we get the opposite result. Since the banking sector as a whole does not have incentives to collectively create a self-regulation body, this creates room for public regulation of the banking sector, which I turn to next.

2.5 Central Bank Policy

2.5.1 Implementing the frictionless benchmark. — The results of the friction-based model showed that banks have incentives to bundle together risk with safe securities, which renders them quasi-safe. This effectively is a market failure. This section shows how the central banking authority can regulate the banking sector to restore the frictionless benchmark. The advantage they will have over the private agents is that they enjoy the power of taxation, which gives them a broader set of instruments. It will be shown that the central bank can implement the frictionless benchmark with two instruments: deposit insurance and a constraint on bank leverage. The key insights of this section will be close to the insights presented in Stein (2010).

The role of deposit insurance provided to households will be to ensure that households have access to a completely safe security. Recall that this

safe security belongs to the set T , but does not belong to the restricted set with agency frictions, T^f . In this case, even when banks have incentives to replace a fraction of the safe security with risky tranches, the backing of the central bank will ensure that the households effectively get a safe security. The obvious question is how will the central bank finance this deposit insurance?

In principle, to restore the frictionless benchmark we want the banks to bear the burden of cheating so as to leave household's consumption unaffected. More specifically, since the resources that are required to finance the deposit insurance will be a random variable, it will not be optimal for the risk-averse households to bear the burden of financing this. If the households were to be taxed to finance the deposit insurance, it would make their consumption riskier, thereby moving us away from the frictionless benchmark.

Consequently, the tax burden will be optimally levied on the banks. When banks have promised s to households, then in the worst case scenario (i.e. when $v = 0$) they have only $(1 - \lambda)s$ units to give the households, and the central bank has to pay λs to meet the deposit insurance claim. That is, the maximum payout of the deposit insurance is λs units. To make this a self-financing program, the central bank should be able to tax the banks λs ex-post (in period 2 after the realization of uncertainty). Consequently, the central bank must ensure that the banks must hold at least λs units of the safe tranche ex-ante (in period 1). This gives us the following restriction on

s :

$$\begin{aligned}\bar{s} - s &\geq \lambda s \\ \text{or } s &\leq \frac{\bar{s}}{1 + \lambda}\end{aligned}$$

Therefore, when the central bank can limit the quantity of safe securities the banks can issue, they will have no incentive to replace the security, since they the central bank will be able to tax them in all states of the world to recover the replacement. The key principle here is that the central bank must constrain the banks from issuing more safe securities than the safe assets they have on their books. Recall that in the frictionless benchmark $s = w^H / (w^B + w^H)$. Plugging this into the above inequality we find that the frictionless benchmark will be restored if

$$\frac{w^H}{w^B + w^H} \leq \frac{\bar{s}}{1 + \lambda} \quad (25)$$

i.e., if the constraint needed to be imposed by the central bank is not too restrictive. This gives us the following proposition.

PROPOSITION 4 *When (25) holds, the central banking authority can restore the frictionless benchmark by using two tools:*

1. *Offer deposit insurance to households against all safe tranches $t = \tilde{s} \in [0, \bar{s}]$ issued by the banks.*
2. *Constrain bank leverage such that the set of tranches the banks can issue is restricted to $T^C = \{\tilde{s} + \tilde{\eta}v \mid \tilde{s} \in [0, \frac{\bar{s}}{1+\lambda}] \text{ and } \tilde{\eta} \in [0, 1]\}$.*

This implementation of central bank authority to make the deposit-insurance program self-financing may look like a capital adequacy ratio (or

a leverage ratio), which regulatory authorities often impose on commercial banks. The central bank can impose this by essentially imposing a tax $\tau(s)$ on every unit of investment made by the banks such that

$$\tau(s) = \begin{cases} 0 & \text{if } s \leq \frac{\bar{s}}{1+\lambda} \\ \infty & \text{if } s > \frac{\bar{s}}{1+\lambda} \end{cases}$$

Alternatively, the central bank can implement the same allocation using permits to issue safe securities. The real-world analogue of these permits are reserves (that satisfy reserve requirements). The Federal Reserve's actively manages the creation of safe securities by participating in the Fed Funds market by draining and creating reserves. At a 10% reserve ratio, each permit (reserve) effectively allows the banks to create 10 units of safe assets, and therefore the central bank can issue $(1/10) \frac{\bar{s}}{1+\lambda}$ permits to implement the self-financing deposit-insurance program that mimics the frictionless benchmark.

This discussion highlights an important role of the central bank. The central bank policy matters in the economy because it has a special monopoly power: it controls the quantity of safe (deposit-insured) securities the banking sector can create. Consequently, when the central bank creates additional reserves, it permits the banking sector to create more safe assets, thereby lowering risk premia in the economy.

2.5.2 Limits of Central Bank Authority. — What if the parameters of the model are such that the inequality (25) does not hold? In this case, the only way the central bank can make the deposit-insurance program self-financing (i.e. not tax the households) is by constraining the safe security issuance to be at a lower level than in the frictionless benchmark. In this case

the central bank cannot restore the frictionless benchmark (first best) using the instruments discussed above. If the central bank allows the banks to create as many safe securities as they want and it provides deposit insurance on all these securities, then there exist states in which it has to tax the household to finance the deposit-insurance program. In particular, the household will have to pay a tax in period 2 of the size $V = \max \{0, k [(\bar{s} + v) - t]\}$, or plugging in the appropriate values:

$$V = \max \left\{ 0, (\bar{s} + v) (w^H + w^B) - \left((1 - \lambda) \cdot 1 + (\lambda) \frac{v}{E[v]} \right) w^H \right\} \quad (26)$$

Note that this tax is a random variable, and will effectively increase the variance of household consumption. In fact, when the inequality (25) does not hold, any increase in the permits issued by the central bank will have no effect on risk premium. This is because the central banks will be taxing households to pay households, so the net effect is a wash.

In practice, how might the government levy this tax on households to raise the additional financing it needs for the deposit insurance program? The government might explicitly implement this tax by increasing the level of income taxes. The other possibility for the government is to induce an inflation which reduces the real value of the nominal government debt outstanding. Consequently, in a nominal economy when the government offers excessive deposit insurance, by increasing the degree of permits in the system, it might be associated with an increase in inflation uncertainty. These considerations are considered more explicitly below.

2.6 Central Bank Policy with Nominal Contracts

This section briefly considers the case in which contracts are denominated relative to the value of some government debt (e.g. government money). The key idea here will be to capture the general equilibrium effects of inflation uncertainty.

Consider the case in which households have an endowment of M units of debt issued by the government, such that each unit will pay P^M units of consumption goods in period 2. For simplicity assume that the household sector holds this debt (i.e. I shut down trading between banks and households with respect to government debt). The government finances this debt by taxing households in period 2. The government imposes a fixed period 2 tax, \overline{Tax} , and after paying off V (as specified in (26)) pays the residual to the debt holders. The payoff to the holder of each unit of this debt in period 2 is net tax revenue divided by total quantities outstanding:

$$P^M = \frac{\overline{Tax} - V}{M}$$

Then each unit of government debt M , can buy $1/P^M$ units of period 2 consumption. This allows us to define the price level as the price of consumption goods relative to the value of government debt:

$$price\ level = \frac{1}{P^M} = \frac{M}{\overline{Tax} - V}$$

This looks like the quantity equation of money, except instead of output, here we have net tax revenues. Note that if V is constant, then the price level will be constant. Equivalently, any variation in V will translate into

variation in the price level. Therefore when V is uncertain, the price level in period 2 is also uncertain. Also, any increase in M leads to a proportional increase in the price level.

In this case, if (25) does not hold, and the central bank insures all safe securities, V will be a nonnegative random variable. This is because under this condition there exist states in which the bank has no resources which the government can tax, and therefore has to turn to the household sector to raise finances. This effectively makes the period 2 price level uncertain from the perspective of period 1.

Now consider the case in which the contracts between households and banks are nominally-denominated. That is, the only tranches that banks can sell households will have the following payoff in units of period 2 consumption goods:

$$\tilde{t} \cdot \text{price level} = \tilde{t} \cdot \frac{M}{\text{Tax} - V}$$

Then any uncertainty in the price level, translates into uncertainty in payoff of the tranches sold by the banks to households. This will therefore make even the safest tranches issued by the bank unsafe, thereby limiting the degree of risk sharing in the economy. In this model, with nominally-denominated contracts it will be optimal for the central bank to limit offering deposit insurance (by limiting how many safe securities they insure) so as to minimize price-level uncertainty in the economy.

This then predicts that for a given quantity of nominally-denominated safe assets, an increase in inflation uncertainty leads to an increase in macroeconomic risk premia for the same reasons laid out above. This then gives us a prediction about the interaction of nominal-denominated safe asset assets, inflation uncertainty and observed risk premia in asset prices.

3 Central Bank Policy with Shadow Banks

3.1 A Simple Example of Monetary Policy with Shadow Banks

In reality, the central bank only regulates a fraction of the banks in the economy. A large quantities of money-like instruments are created in the shadow banking sector, which are not directly regulated by the central bank. We can create a simple analogue of our model to fit this feature of the modern banking system.

Let's continue to consider the model with frictions summarized in A1-A2, and in addition consider the following constraint on central bank policy.

Assumption (A4). *The government has regulatory authority only over a fraction ρ of the banks. Label the regulated banks as 'traditional banks,' and the unregulated banks as 'shadow banks.'*

This assumption essentially implies that for a fraction $(1 - \rho)$ of the banks operating in the shadow banking sector, the government has no ability to offer deposit insurance or regulate the quantity of safe securities issued. To map to reality, you may think of the securities issued by these banks as asset backed commercial paper, which households hold via money market funds.

The only difference this has on the solution is that now there will be two types of banks. The shadow banks, will freely choose the tranche they sell, and will continue to choose $t = (1 - \lambda) s^{sh} + \frac{\lambda s^{sh}}{1 - \bar{s}} v$, exactly as in the previous section (here I use s^{sh} to denote the safe component offered by the shadow banks). The traditional banks on the other hand will be subject to the constraint imposed by the central bank. Since they will be regulated (exactly as in the previous section) they will be able to commit to sell safe

securities. Suppose the restriction imposed by the central bank binds on them, then they will choose $t = s^{tr} < s$, where s^{tr} is the restriction imposed by the central bank. Then household's budget constraint (16) is given by

$$w^H = p(s^{tr})z(s^{tr}) + p(t)z(t) \quad (27)$$

where s^{tr} is the safe tranches with deposit insurance purchased from traditional banks, and $t = (1 - \lambda)s^{sh} + \frac{\lambda s^{sh}}{1 - s^{tr}}v$ are quasi-safe tranches purchased from the shadow banks. Plugging in the appropriate values³³ into the budget constraint (27) uniquely solves for s^{sh} :

$$s^{sh} = \frac{w^H - \rho \left(\frac{w^B s^{tr}}{1 - s^{tr}} \right)}{w^H - \rho \left(\frac{w^B s^{tr}}{1 - s^{tr}} \right) + (1 - \rho) w^B}$$

It can be checked that $\partial s^{sh} / \partial s^{tr} < 0$. That is, when the central bank reduces permits by lowering s^{tr} , it must be the case that revenue of the traditional banks (given by $p(s^{tr})z(s^{tr})$) shrinks, and that of the securitized banks (given by $p(t)z(t)$) expands. This effectively implies that now households have to hold more quasi-safe securities and less safe securities, which in turn makes their consumption riskier. Hence, a contraction by the central bank leads to an increase in risk premia. In fact, in this case, due to the symmetry, a decrease in s^{tr} that leads to a reduction of x units of total safe securities from the traditional banking sector, corresponds to a one-for-one increase of x units of total quasi-safe securities by the shadow banking

³³In equilibrium $z(s^{tr})$, $z(t)$ have to equal the securities created by the banks, and the banks make zero profits. This implies $p(s^{tr}) = s^{tr}$, and $p(t) = s^{sh}$, and correspondingly, $z(s^{tr}) = \rho w^B / (1 - s^{tr})$, and $z(t) = (1 - \rho) w^B / (1 - s^{sh})$.

sector. This is simply saying that

$$w^H = \{p(s^{tr})z(s^{tr}) - x\} + \{p(t)z(t) + x\}$$

But since the shadow banks create quasi-safe securities there is a loss of λx quantities of total safe tranches from the household's portfolio. Thus, by reducing permits for creating safe security from the hands of the traditional banking sector, the central bank effectively reduces the total quantities of safe securities available to the households, since the shadow banking sector is ineffective in producing completely safe securities. This leads to a reduction of money-like securities in the economy, thereby leading a rise in the risk premia. Figures 2.1 and 2.2 present a numerical example.

Securitized Banks	
Assets	Liabilities
<ul style="list-style-type: none"> • Projects $w/2+40$ 	<ul style="list-style-type: none"> • Wealth $w/2$ • Debt (t) 40

Traditional Banks	
Assets	Liabilities
<ul style="list-style-type: none"> • Projects $w/2+60$ • (Permits 6) 	<ul style="list-style-type: none"> • Wealth $w/2$ • Debt (s) 60

Household	
Net Assets	
<ul style="list-style-type: none"> • s securities 60 • t securities 40 	

Figure 2.1: Time 0 — Balance Sheets before contracting reserves.




Figure 2.2: Time 1 — Balance sheets after reserve contraction.

4 Quantitative Effects and Implications for Asset Prices

So far we did not make any stringent functional form assumptions, and therefore the results were mostly qualitative. In this section the goal is to quantitatively assess the effect of safe assets supply on equilibrium risk premium. To do that I begin with the basic asset pricing setup.

4.1 Basic Asset Pricing

The model in this paper effectively takes consumption-savings decision as given, and focuses on the portfolio choice problem. The structure of my model, therefore, is similar to the Lucas-tree environment, which takes consumption endowments as given, and solves for asset prices. The novel feature here is that given total resources, the quantity of safe assets that end up in household portfolio will depend on frictions and regulation of the banking sector. In what follows I will use the same approach as in the Lucas-fruit-tree framework, and examine how changing the nature of frictions and regulatory response from the government affects asset prices. This will allow us to identify conditions under which the mechanisms described in this paper are quantitatively relevant.

Typically households have access to other assets other than just bank securities. The setup so far had abstracted away from such considerations. Here I introduce this feature by expanding household's endowment to include L units of assets with payoff $(\bar{s}_{out} + v_{out})$, such that \bar{s}_{out} is a constant. For simplicity assume that $v_{out} = \left(\frac{1-\bar{s}_{out}}{1-\bar{s}}\right) v$, such that v_{out} is a nonnegative random variable with $E[v_{out}] = 1 - \bar{s}_{out}$. To retain simplicity I restrict the set of contracts that banks can write to be the same (T) as in previous sections. This is not a trivial assumption since I am preventing the banks from writing derivative contracts with households that are contingent on the payoff of these L units of assets. In other words, these L units cannot be pledged. However, allowing such contracts will not substantially alter the main insights of this section, and hence are precluded. Under these assumptions,

we can represent household's period 2 consumption as

$$\begin{aligned}
 c^H &= \underbrace{L(\bar{s}_{out} + v_{out})}_{\text{Outside Endowment}} + \underbrace{K \left[(1 - \psi) + \psi \frac{v}{E[v]} \right] s}_{\text{Bank Securities}} \\
 &= [L + Ks] \left\{ (\phi) \cdot 1 + (1 - \phi) \frac{v}{E[v]} \right\}
 \end{aligned}$$

such that ψ is the fraction of securities that is effectively replaced (by unregulated banking), and $\phi \equiv \frac{L\bar{s}_{out} + K(1-\psi)s}{L+Ks}$ represents fraction of household consumption that is safe. Here asset prices will correspond to the case in which household's first order condition will hold with equality. Therefore the asset prices derived here will correspond to the interim period trading, rather than period 1 trading (since in period 1 households were at a corner solution). Normalizing period 1 marginal utility of consumption as 1, we can represent the stochastic discount factor as

$$\bar{M} = \beta u'(c^H)$$

such that $0 \leq \beta \leq 1$ is the time-discount factor. This allows us to represent the price of any asset with payoff d , as

$$p(d) = E[\bar{M}d]$$

Using this fundamental pricing equation we can represent the price of any security in this economy. The price of a money-like claim which pays off 1 unit for sure in the next period is given by

$$p(1) = E[\bar{M}]$$

while the gross one-period return on the risk-free asset R^f is $1/p(1)$.

Analogously, a one-period equity is a hypothetical asset that pays next period's consumption. The price of this asset is

$$p(c^H) = E[\bar{M}c^H]$$

while the gross one-period return is $R^e = c^H/p(c^H)$.

Let's define the equity premium as $ep \equiv \ln[R^e/R^f]$, and plugging in appropriate values we get

$$ep = \ln E[c^H] + \ln E[u'(c^H)] - \ln E[u'(c^H)c^H]$$

and the log riskfree rate is

$$rf = \rho - \ln E[u'(c^H)]$$

such that $\rho \equiv -\ln(\beta)$.

The standard convention in the asset-pricing literature is to use power utility with coefficient of relative risk aversion γ . Define $X \equiv \ln[c^H]$, and we get

$$ep = \ln E[\exp(X)] + \ln E[\exp(-\gamma X)] - \ln E[\exp((1-\gamma)X)]$$

$$rf = \rho - \ln E[\exp(-\gamma X)]$$

It can be checked that the equity premium rises, and the risk-free rate falls when ϕ falls. The factors that lead ϕ to fall are:

1. The size of replacement friction in the banking sector rises (ψ increases)

2. The share of safety in aggregate assets falls (i.e. when \bar{s} or \bar{s}_{out} falls)
3. The aggregate endowment become riskier (i.e. v becomes riskier)

The general principle here is that whenever aggregate non-diversifiable uncertainty increases, there is a flight-to-safety to money-like instruments. Consequently, for a given supply of such instruments, asset market clearing leads to an increase in the premium households are offered to hold non-safe instruments (i.e. a rise in the risk premium).

A similar effect happens when the central bank drains reserves (as in the shadow banking section above). This forces households to rely on the unregulated banking sector, thereby making consumption riskier, leading to a rise in equity premium, and a decline in the riskfree rate. The following subsections presents a calibration of the model to assess the magnitude of these effects.

4.2 The Effect of Money on Risk Premium in a Safe World

For what follows the key parameter I will work with is ϕ , which measures the fraction of consumption that is safe. First, I consider a parameterization of the model that corresponds to the baseline textbook case—when v is lognormally distributed. In this case, as ϕ converges to zero, we get the ubiquitous result that under power utility the equity premium equals γ times the variance of v : i.e. $ep \rightarrow \gamma\sigma^2$ as $\phi \rightarrow 0$. The other extreme here is when $\phi \rightarrow 1$, in which case $ep \rightarrow 0$, since household consumption is completely safe. Reasonable values of γ range between 1 and 3, and the standard deviation of aggregate consumption growth in the US has been between 1-2%. The risky component of consumption will have higher variance. As

a benchmark I consider this to be more than double that of consumption, such that $\sigma = .04$. Consequently, by setting $\gamma = 2$ and $\sigma = .04$, we get the equity premium ranges between 0 and .0032 (i.e. .32%) as a function of ϕ . This effectively implies that in the extreme case of making household consumption completely risky from completely safe, will increase the equity premium by at most .32%. This effect compared to actual observable values of equity premia (upwards of 6%) is relatively small.

The main result here is that the quantitative effect of safe assets and the sensitivity of risk premium with respect to ϕ is relatively small in an environment with longnormally distributed risky assets. This, however, changes drastically when we consider the possibility of catastrophe risk, which I turn to next.

4.3 The Effect of Money in a World with Disaster Risk

4.3.1 Baseline. — Recent work by Weitzman (2007) and Barro (2006) has highlighted the importance of tail risk in explaining large observable risk premia. The tail risk effectively captures the possibility of large contractions by assigning such episodes non-negligible probabilities. In Weitzman's framework, such tail-thickening arises from subjective uncertainty about the scale parameters of the distribution of log consumption. Analogous to these recent papers, this section will work with cases in which the distribution of $\log(v)$ is thick-tailed. The key difference here will be that unlike in previous papers, I assume that only a fraction of consumption is subject to such disaster risk. This will allow me to characterize how changing the fraction of safe assets will influence risk premia. It will turn out that in a framework with thick-tails, risk premium will be very sensitive to

changes in ϕ .

Recall household consumption can be represented as a constant times a weighted average of unity and a random variable with mean one³⁴:

$$c^H = [L + Ks] \left\{ (\phi) \cdot 1 + (1 - \phi) \frac{v}{E[v]} \right\}$$

Let the distribution of $\log(v)$ be given by Student's t distribution. The left tail of log consumption is governed by the fraction of consumption arising from the safe assets, and the properties of the tails of the risky tranche v . The following basic proposition establishes a crucial insight for the rest of the results that follow.

PROPOSITION 5 *Under constant relative risk aversion, $E[u(c^H)] \rightarrow u(L + Ks)$ as $\phi \rightarrow 1$, and $E[u(c^H)] \rightarrow -\infty$ as $\phi \rightarrow 0$.*

PROOF. See appendix. ■

This proposition essentially implies that expected utility converges to a constant when the share of safe assets is large in household portfolio, and diverges to negative infinity when share of safe assets become negligible. A related result about the menacing effects of thick-tails on utility functions that represent decreasing absolute risk aversion was independently pointed out by Geweke (2001) and Weitzman (2007). The key intuition behind this result is that under power utility, $u(0)$ is a drastically perverse state, which households would like to avoid at all costs (since the utility at zero consumption diverges to negative infinity). Correspondingly, whenever log consumption is distributed with thick-tails, the probability weight around zero consumption is non-negligible, thereby inducing expected utility to di-

³⁴The results of this section is not going to change much if instead we considered a case in which the payoff of the safe assets were lognormally distributed instead of being constant.

verge to negative infinity.

When there is a sufficiently large quantity of safe assets in the portfolio, the total portfolio is no longer exposed to the menacing effects of the tails. This is because mixing sufficiently large quantities of assets distributed with constant returns (or thin-tailed returns) makes the distribution of the total portfolio returns thin-tailed. This effect is captured in Figure 2.3, which plots the expected utility of this mixed portfolio as a function of the share of safe assets, ϕ . The benchmark line which is the horizontal line is expected utility when $\phi = 1$. There are two key points to note from this figure:

1) when share of safe assets (ϕ) is small, expected utility and marginal utility are very sensitive to changes in the share of safe assets in household's portfolio.

2) when share of safe assets is sufficiently large (say greater than 20 percent) household's expected utility is not going to be very sensitive to changes in higher moments of log consumption.




Figure 2.3: Expected Utility as a function of Fraction of Safe Assets, ϕ . The dashed line represents expected utility when the entire portfolio is riskless.

The following corollary summarizes the asset pricing implications of this result.

COROLLARY 5.1 $ep \rightarrow 0$ as $\phi \rightarrow 1$, and $ep \rightarrow \infty$ as $\phi \rightarrow 0$.

To understand this qualitatively, note that ep is an inverse-function of the supply of safe assets as measured by ϕ in this specification. This is because an increase in supply of safe assets decreases the marginal utility in disaster states, thereby lowering the value of the riskfree asset which pays off in these disaster states. This makes the difference between the two securities smaller, thereby lowering the equity-risk premia. This difference widens whenever there is greater mass in the left tail of the unconditional distribution of log consumption growth. This happens when the fraction of safe assets (ϕ) contracts. The opposite happens when ϕ expands. This generates an inverse relationship between money supply and the equity risk premia. Appendix (...) discusses the mathematical underpinnings of the quantitative effect. Those who are familiar with the use of moment-generating functions in asset pricing may find this appendix useful. Below I present an example which illustrates what is driving the quantitative effect behind this corollary.

4.3.2 Pricing of Tail Risk. — Introducing safe assets in household portfolio effectively truncates the left tails of consumption. This positive effect of truncation on risk premia, competes with the negative effect of thick-tails on risk premia. Note we will achieve this tail-thinning effect also when we consider lognormally distributed assets (as opposed to completely riskless assets). The tension between these two effects is what generates the large sensitivity of risk premium with respect to safe assets. To see this cleanly here I present a simple case that will highlight how tail risk gets priced, and

under what conditions households will be less bothered by tail properties of traded assets.

Consider two assets A and B which are identical, except that asset A pays off 1 unit surely in disaster states, and asset B pays off zero in disaster states. You may think of asset B as a MBS with some tail risk, and asset A as a riskfree security. The payoffs of asset $i \in [A, B]$ is denoted by d^i , and is given by

$$d^A = \begin{cases} 1 & \text{if } v > \underline{v} \\ 1 & \text{if } v \leq \underline{v} \end{cases}$$

and

$$d^B = \begin{cases} 1 & \text{if } v > \underline{v} \\ 0 & \text{if } v \leq \underline{v} \end{cases}$$

Define $q \equiv \Pr[v \leq \underline{v}]$. Then the price of each asset is given by

$$\begin{aligned} p(d^A) &= (1 - q) E[\bar{M} | v > \underline{v}] + q E[\bar{M} | v \leq \underline{v}] \\ p(d^B) &= (1 - q) E[\bar{M} | v > \underline{v}] \end{aligned}$$

The spread of these two assets, $\ln [p(d^A) / p(d^B)]$, is given by

$$\ln \left[\frac{p(d^A)}{p(d^B)} \right] = \ln \left[1 + \left(\frac{q}{1 - q} \right) \frac{E[\bar{M} | v \leq \underline{v}]}{E[\bar{M} | v > \underline{v}]} \right] \quad (28)$$

The question we are after is ‘when do agents in the economy start caring about tail risk of any asset being traded in the economy?’ In the context of this example, this question translates into asking ‘when does the difference in these two assets matter?’ If the spread in the valuation of these two assets $\ln [P^A / P^B]$ is ‘almost’ zero, then these two assets will appear roughly similar to the agents. From examining (28), it is evident that the

difference between the two will matter whenever conditional expectation of marginal utility is high, or whenever q is high. However, since we are focusing on disaster states we can restrict attention to the cases in which q is very low – to be precise less than .005. Then whenever expected marginal utility in disaster states is not much larger than expected marginal utility in non-disaster states the two assets will be valued almost identically (i.e. when $E[\bar{M} | v \leq \underline{v}]$ is not too large relative to $E[\bar{M} | v > \underline{v}]$). On the other hand whenever $E[\bar{M} | v \leq \underline{v}]$ is large, the valuation of the two assets diverge. Under our assumption of power utility with $\log(v)$ distributed with a student-t distribution, $E[\bar{M} | v \leq \underline{v}]$ diverges to positive infinity whenever the fraction of safe assets go to zero (such that the distribution of \log consumption converges to a student-t). Having sufficiently large fractions of safe assets ensures that the expected utility is finite and small. This is summarized in the following corollary.

COROLLARY 5.2 The tail properties of any asset traded has vanishingly negligible effect in the pricing of assets as $\phi \rightarrow 1$.

This sensitivity can be seen in Figure 2.4. This result may give an insight into why we witnessed large spikes in spreads of relatively safe securities (e.g. LIBOR-OIS) between August 2007 and January 2009. During these episodes we had a massive contraction in safe securities, such as asset-backed commercial paper, and repo, which according to the logic of this model may lead to an increase in bond risk premia between two almost identical securities: one safe, and the other quasi-safe.

The general message of the calibration subsections is that the quantitative effect of safe assets on risk premium is relatively small in an environment with no disaster risk. But as soon as we introduce disaster risk, small

changes in quantity of safe assets can have a large effect on equilibrium risk premium.




Figure 2.4: Bond Risk Premia as a function of Fraction of Safe Assets (ϕ).

5 Extension: Shadow Banking Shut-Down and Unconventional Policy

The discussion thus far was restricted to how monetary policy works during ‘normal’ times. We are, however, also interested in understanding the role of central bank policy in ‘abnormal times’. To do so we need to first define what ‘abnormal times’ exactly means. The most interesting case is when the shadow banking sector (or in general the unregulated banking sector) shuts down. In reality, various micro-mechanisms may lead to such

a shut down. For the purposes of this paper I consider one fundamental micro-friction to capture such states: counterparty risk in distress states. The following assumption summarizes this friction:

Assumption (A5). *After the realization of outcomes, if $v \leq \underline{v}$ (i.e. in distress states), a bank that chooses $t = \tilde{s} \in [0, \bar{s}]$ fails to pay with probability w .*

This is the classic agency problem of post-contractual opportunism, in which self-regulation breaks down when the banks themselves lose most of their capital and enter distress states. The real-world analogue in terms of safe securities issued by ABS conduits is when the sponsoring banks get distressed and refuse to assume rollover risk or even credit risk of the conduits. Since the traditional banks have deposit insurance, the mechanics will be similar with or without counterparty risk. Therefore for simplicity I assume that only shadow banks are subject to counterparty risk.

In this case, it can be checked under our parametric assumptions that there exists a $0 < w < 1$ and \underline{v} such that the household's willingness to pay for any security t issued by the shadow banks will be below s . That is, for all t ,

$$p(t) < s$$

Then the shadow banks will make negative profits if they issue any security, and consequently will chose to shut down. Thus, possibility of default in disaster states can undo the shadow banks ability to sell insurance to households. The fundamental insight here is that under this 'safety transformation' view of banking, the key role banking sector plays is to provide insurance to households. Consequently, any possibility of the banks failing in disasters states (systemic risk) can undo their advantage and render them unprofitable. Hence, in such an economy, households will transact

with shadow banks only when they are sure that the banks will payout in high marginal utility states.

When the shadow banking sectors shut down, what can the central bank do? Consider one policy: the government enters a swap agreement with the traditional banks in which they buy the claims to the payoff of the risky investment project, $\bar{s} + v$. In turn the banks receive a riskfree claim from the central bank. How can the central bank create riskfree claims? Suppose the government exchanges one unit of investment payoff ($\bar{s} + v$) for f riskless units, and buys G quantities of such exchange-claims. The risk-neutral banks will be willing to enter this trade as long as $f > E(\bar{s} + v) = \bar{s}$. Then in period $t + 1$, whenever $\bar{s} + v \geq f$, the government just takes the proceeds from the payoff of each exchange-claim they own to pay the bank Gf units. On the other hand, when $\bar{s} + v < f$, the government has to rely on its power of taxation to pay the difference: $G[f - (\bar{s} + v)]$.

Can this policy affect the risk premia and investment in the economy? The answer to this question depends on who the government will tax. Let's first consider the case in which the government taxes the households. In this case it can be shown that the government policy will have no effect or will increase risk premia (but not decrease it). This is because the government will force households to hold more risk (in the form of taxation) than they want.

What if the government taxes a group other than households in the states of the world they have a shortfall of magnitude $f - (\bar{s} + v)$? The only group in this economy other than households who will have riskless quantity of resources in time $t + 1$ is the bankers who own the securitized banks, and are currently shut down. In the state of shut down, they are using their

wealth $(1 - \rho) w^H$ to invest in projects, and each unit gives them \bar{s} units of certain resources. In this case of taxing the shadow bankers, the government's asset swap has a direct impact on asset prices, risk premia, and investment. In particular this asset swap makes the traditional banks' portfolio safer, which in turn allows them to issue more riskless claims to the households, which reduces the risk premia in the economy. The government can restore optimal allocation as long as $fG = (1 - \rho) w^B s^{sh} / (1 - s^{sh})$. The government brings about this Pareto improvement by using its power of taxation and enabling the banking sector to credibly expand its creation of safe assets, which satiates households desire to risklessly transfer resources intertemporally.

What does this asset swap look like in practice? This effectively mirrors an asset purchase by the Federal Reserve, which they pay for by printing reserves under the auspices of a 'quantitative easing' program. The analysis above yields an important insight about such programs. The central bank, or more generally the government, has a special power of taxation which allows them to create safe assets. However, the only way this can have a real positive effect on the economy is when they transfer the burden of taxation in disaster states onto some agents who are not the marginal investors in asset markets, and are willing to bear the risk of taxation (shadow banks in the example above).

The government's action of quantitative easing can be thought of as a case in which the government forces a disaster insurance contract between two groups of agents. If the central bank incurs losses on its asset purchase they will impose a tax on a group of agents (either an inflation tax or one directly levied through the Treasury). Consequently, this group of agents

is effectively selling disaster insurance to those who rush to deposit money in commercial banks, which in turn increase their holdings of the reserves issued by the central bank. Equivalently, the U.S. government's doubling of Treasury Bills debt in September 2008 (from \$1 trillion to \$2 trillion) and using the proceeds to finance various funding facilities (either directly or through the Federal Reserve) can be thought of as one such forced-trade of disaster insurance.

The behavior of agents after Lehman's failure mirrors the one predicted by this model. In the week of September 15th, 2008, after Primary Reserve Fund 'broke the buck,' there was a massive run on prime money market mutual funds, and in general on anything that was not deemed safe. Households increased their holdings of various instruments offered by commercial banks as long as they were federally insured. Similarly, there was a flow of money into money market mutual funds that only invested in Treasury and agency debt. These commercial banks and government MMMFs in turn were looking for safe havens to park this newly-found influx of cash. The government satiated this appetite by issuing reserves that paid interest, and substantially increasing the quantity of Treasury debt outstanding.

Lastly, by following the logic of the model it also becomes clear we will observe 'monetization' of central bank's losses (i.e. an inflation tax) only if the risky assets they hold have bad realizations.

Chapter 3: What Drives Global Asset Allocation Across Countries?

1 Introduction

Asset allocation decisions of global investors lie at the heart of financial flows between markets, currencies and countries. The paper takes as its point of departure the asset allocation decision of the individual investor. This sets it apart from much of the existing literature, which focuses on investment flows from a macroeconomic point of view, and derives most of its analysis from balance of payments data (Forbes and Warnock (2011); IMF (2011b); IMF (2011c)). In this paper's more integrated view, changes in risk and return preferences of individual investors are the fundamental driver of asset allocation over time and, consequently, financial flows into and out of markets, currencies, and countries.

This paper aims to understand recent trends in global asset allocation across countries and their determinants. First we focus on the unleveraged, real-money investors, including individuals, public and private pension funds, insurance companies, as sizable sources of underlying capital flows. Then we turn to the decisions of central banks with respect to their reserve allocation. While the overwhelming majority of financial assets is owned and managed by private investors, sovereign investors have grown to become important players in international capital markets. Sovereign wealth funds (SWFs) hold some \$4 trillion in assets, while international reserves amount to \$10 trillion. Their combined assets amount to about a quarter

of the assets under management of private institutional investors. Their increasing size makes sovereigns important investors in international capital markets.

An extensive literature links asset allocation to an investor's objectives and the risk and return characteristics of individual assets. It is assumed that investors behave predictably when such characteristics change; when the expected return of an asset increases without changes in its riskiness, investors are expected to want to hold more of that asset. Similarly, when an asset becomes more risky (because its return is more variable or the risk of default is higher), investors would want to hold less of it, unless the asset offers a higher return.

The global financial crisis has raised the possibility that some of the parameters in these relationships may have changed, including investor's objectives themselves. Anecdotal evidence abounds and can sometimes seem contradictory. For example, investors, spooked by the financial turmoil, are said to have become much more sensitive to risk, in particular to events with small probability but large adverse effects ("tail event"). They are now seeking more protection against such events. Similarly, after disruptions in some markets during the height of the financial turmoil, investors are said to be much more focused on liquidity risk. These structural changes interact with cyclical factors: despite increased sensitivity to risk, the low-interest rate environment may push some investors, especially those with the need to earn a certain minimum return to match expected payouts on their liabilities, to take on more risk in alternative assets and less liquid markets to increase returns on their assets.

In this context, the paper focuses on the following questions:

- What are the trends in global asset allocation since 2005, and what are their determinants? Do trends and determinants differ by country or region?
- Have the financial crisis, the sovereign debt crisis in Europe, and low interest rates in advanced economies fundamentally altered investment decisions, perhaps pressing long-term investors toward riskier investment to augment their poor returns in advanced economies?
- Are there growing risks for a reversal of investment flows to emerging economies, and if so, how would that affect capital flows? In the longer term, is financial stability compromised as a result of these developments?

The analysis shows that global asset allocation is driven most strongly by growth prospects and risks in the recipient countries, while interest rate differentials between countries play a lesser role. The analysis does not, however, imply that capital flows in general do not respond to interest rate differentials, since other components, including investment flows of short-term leveraged investors (such as those from the carry trade)—which this paper does not examine—might still be affected by changes in interest rates.

Beyond these long-term trends, the empirical results indicate that asset-allocation strategies of real money investors have changed since the onset of the global financial crisis. Most importantly, investors are more risk conscious, including regarding the risks associated with liquidity and sovereign credit. Also, the structural trend of investing in emerging market assets has accelerated following the crisis; and with many first-time investors taking advantage of the relatively better economic performance of these

countries, the risk of a reversal cannot be discounted if fundamentals (such as growth prospects or country or global risk) change. For larger shocks, the impact of such reversals could be of the same magnitude as the pull-back in flows experienced during the financial crisis.

With respect to reserve allocation of central banks we find a stronger role for interest rates. Reserve managers appear to respond to U.S. interest rates: increases in the U.S. dollar interest-rate are associated with a rebalancing away from euro and towards the U.S. dollar.

This paper proceeds as follows. Section 2 briefly reviews the literature. Section 3 discusses data and methodology. Section 4 provides stylized facts and Section 5 summarizes empirical results. Section 6 concludes with implications of our findings.

2 Literature Review

This section discusses the relationship of our analysis to existing literature on portfolio investment flows and global asset allocation.

Several studies have looked into the relationship between international portfolio flows and local market returns: Froot, O'Connell, and Seasholes (2001), Bekaert, Harvey, and Lumsdaine (2002), and Froot and Ramadorai (2008) document significant effects of portfolio investment flows on local market equity and bond returns.

Other studies have looked at the reverse relationship between local market returns and subsequent international flows. This literature finds some evidence of "performance chasing" behavior. Such behavior is documented at both the domestic level (see Grinblatt, Titman, and Wermers

(1995), Grinblatt and Keloharju (2000)) and the international level. In the international context, Brennan and Cao (1997) find that portfolio flows are associated with returns on national market indices. They attribute this to informational differences between foreign and domestic investors. For emerging markets they find that in addition to contemporaneous returns, lagged local market returns are a strong predictor of portfolio flows. Froot and Teo (2004) attribute such behavior to behavioral factors and style investing.

A separate literature looks at the international allocation of domestic investors. This literature finds large home bias in portfolio holdings. French and Poterba (1991) was an early paper to document too little cross-border diversification, despite the large portfolio diversification benefits. Recent papers continue to find home bias in portfolio allocation at both the individual and institutional level. Hau and Rey (2008) find that there is large heterogeneity in home bias across mutual funds. They also document that a representative fund usually invests in a limited number of countries, although fund size is positively related to number of countries in which funds invest.

This paper adds another important dimension to this literature, focusing on the link between asset allocation decision of institutional investors and international portfolio flows. This paper also relies on a relatively high frequency public dataset. Most of the other studies use either (quarterly) balance of payment data or propriety data maintained by large custodian banks. In our framework, we focus on the international portfolio flows that result from the choices made by institutional investment managers. This approach is desirable since a substantial body of already theoretical literature deals with the optimal portfolio choice at the individual or institutional

level. Correspondingly, we first write down a simple model of optimal portfolio choice that is motivated by the recent literature in portfolio choice. The model then guides our empirical investigation about the determinants of international portfolio flows.

3 Data and Methodology

3.1 Data

For our empirical investigation, we use a dataset compiled by Emerging Portfolio Fund Research (EPFR). EPFR provides global fund flows and asset allocation data from some 20,000 equity funds and 10,000 bond funds with \$14 trillion in total assets. The investors are a mix of retail and institutional investors; EPFR estimates that 70 percent of assets are institutional, mainly from pension funds and insurance companies. It covers funds registered in most major developed market jurisdictions and offshore domiciles. EPFR samples a subset of funds to give insights into the destination countries for equity and bond investments. Data at the monthly frequency are used below, covering the period from January 2005 to May 2011. EPFR has widened its coverage of fund flows over time, which may raise data consistency issues; the period of study was chosen to minimize these concerns.

Capital flows in and out of countries may include other types of investments, such as bank loans or FDI, that are not tracked by the EPFR dataset. However, since we are mainly interested in the decisions of long-term portfolio investors, the EPFR data suits our purposes better than the standard balance of payments data that includes these other types of investment. Also, for some countries, especially emerging markets which are

traditionally underweighted in portfolios of international investors, flows in and out of bond and equity funds may to a considerable extent capture the corresponding cross-border flows.³⁵

3.2 Methodology

Using the EPFR data, this section addresses the following questions: First, what global and domestic factors have driven the asset allocation of international bond and equity fund investors? Second, has their investment behavior changed fundamentally after the global financial crisis? To capture the truly global picture, a panel regression is estimated covering 50 advanced and emerging market economies for which we have complete and consistent data. The regressions are run separately for equity funds and bond funds, and are estimated for the whole sample and for five geographic groupings separately.³⁶

The theoretical setup is based on an optimal portfolio problem with constant relative risk aversion (CRRA) utility. This allows the portfolio decision to be derived from microeconomic foundations. The CRRA function has a constant parameter of risk aversion, which implies that risk aversion does not change with investor's wealth. Furthermore, the CRRA utility function is time separable, which means that total utility is a simple sum of utilities in each separate period.

³⁵Specifically, reducing their underweighting in international capital market indices may lead to increased portfolio flows into emerging markets, with corresponding capital inflows.

³⁶The regressions are run on flow data, since the stock data are generally nonstationary. The dependent variables are defined for each country as the valuation-adjusted flows into equity and bond funds in the country, divided by the stock at the beginning of the month. All variables are used at a monthly frequency. For variables of higher frequency, the end-of-month value is used. All regressions include country-fixed effects to account for any country specific factors not identified by the other explanatory variables. Dropping country-fixed effects does not alter the signs or statistical significance of the results.

Using a CRRA utility function, the portfolio decision can be represented as follows:

$$\max E_t \left[\sum_{i=0}^{\infty} \delta^i U(C_{t+i}) \right] = E_t \left[\sum_{i=0}^{\infty} \delta^i \frac{C_{t+i}^{1-\gamma} - 1}{1-\gamma} \right] \quad (29)$$

where $U(\cdot)$ is the CRRA utility function, C_{t+i} is consumption at time $t+i$, γ is the coefficient of relative risk aversion, δ is a discount factor, and $E_t[\cdot]$ is the expectations operator taking into account all information up through period t .

The intertemporal budget constraint of the investor is given by

$$W_{t+1} = (1 + R_{p,t+1}) (W_t - C_t) \quad (30)$$

where $R_{p,t+1}$ is portfolio return between period t and $t+1$, and W_{t+1} is wealth in period $t+1$.

Regarding the portfolio, suppose the investor can choose from N risky assets and one risk-free asset. \mathbf{R}_{t+1} is a vector of risky returns with N elements. It has a mean vector $E_t \mathbf{R}_{t+1}$ and a variance-covariance matrix $\mathbf{\Sigma}_{t+1}$. $\boldsymbol{\alpha}_t$ is a vector of allocations to the risky asset. The riskless asset has return R_{t+1} from time t to $t+1$.

The portfolio manager's decision, therefore, is to optimally choose $\boldsymbol{\alpha}_t$ to maximize his utility (1) subject to his budget constraint (2).

Unfortunately, a closed-form solution to this investment problem does not exist. However, based on a linearized approximation, the following

solution can be derived (see Campbell and Viceira (2002) for details):

$$\begin{aligned}\alpha_t &= \frac{1}{\gamma} \Sigma_{t+1}^{-1} \left(E_t \mathbf{r}_{t+1} - r_{f,t+1} + \sigma_t^2 / 2 \right) + \left(1 - \frac{1}{\gamma} \right) \Sigma_{t+1}^{-1} \sigma_{ht} \quad (31) \\ \sigma_{ht} &= -\text{Cov}_t \left[\mathbf{r}_{t+1}, (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{p,t+1+j} \right]\end{aligned}$$

Where σ_t^2 is the portfolio variance, σ_{ht} is the vector of covariances of each risky asset return with revisions in expected future portfolio returns, and ρ is a parameter of the linearization. When the consumption-wealth ratio is constant, ρ can be interpreted as the ratio of reinvested wealth to total wealth.

Equation (31) shows that the optimal weight for each asset in the portfolio is a function of two terms: (i) the asset's risk premium, based on its excess return, variance, and covariance with other assets; and (ii) the asset's covariance with revisions in expected portfolio returns, i.e. its hedge against future declines in portfolio returns (intertemporal hedging component).

The weights placed on these two terms are proportional to the investor's risk tolerance ($1/\gamma$). At one extreme, when the investor has a risk tolerance of 1 (or "log utility"), asset weights are determined solely by the risk premium of each asset. At the other extreme, when the investor has a risk tolerance of zero, asset weights are only a function of the intertemporal hedge provided by each asset.

This result, therefore, predicts that an investor will choose to allocate more of his portfolio to a given asset i when:

- it is expected to generate high excess returns, that is, the i th term in $E_t \mathbf{r}_{t+1} - r_{f,t+1}$ is high;

- it has low variance, that is, the i th diagonal term in Σ is low;
- it has low covariance with other assets, that is, the nondiagonal terms in Σ are low; and
- it offers a hedge against future declines in portfolio returns, that is, σ_{ht} is high

Furthermore, when risk-aversion γ increases, an investor will shift his portfolio toward less risky assets, or more precisely, to assets that offer a better hedge against future declines in portfolio returns. Therefore, in periods of elevated risk aversion, investors will move out of risky bonds and equities to “risk-free” instruments.

On the basis of this model, the following factors are used in the regression analysis to explain global asset allocation:

- Return factors: (i) policy rate differentials of countries relative to the simple G-4 average; and (ii) the one-year-ahead GDP growth forecast from Consensus Economics.
- Volatility factors: these represent the variance of returns as measured by (i) the volatility of host country expected inflation; (ii) the volatility of GDP growth; and (iii) the volatility of the exchange rates.
- Risk tolerance: perceptions of risk are (i) country risk, as proxied by the measure of country risk compiled by the International Country Risk Group; and (ii) global risk, as proxied by the Chicago Board Options Exchange Volatility Index (VIX).

- Other variables of interest: (i) an IMF measure of capital controls (both on inflows and outflows)³⁷, (ii) the covariance between country returns and world portfolio return (to capture the diversification effect), (iii) the covariance between country returns and changes in world portfolio return (to capture intertemporal hedging demand), and (iv) dummies to account for any structural changes in investor behavior that may have occurred after the global financial crisis.³⁸

The variables that are used as a proxy for the various determinants above are listed in Table 3.1.

Model determinant	Equities		Bonds	
	Empirical equivalent	Proxy in regressions	Empirical equivalent	Proxy in regressions
Expected returns	Capital gains Dividends Country risk	Real GDP growth Real GDP growth Country risk	Interest income Default/credit risk Country risk	Risk-free interest rate Real GDP growth Country risk
Variance	Stock market volatility	Real GDP growth volatility	Inflation risk	Inflation volatility
Covariance (diversification effect)	Covariance with world portfolio returns	Covariance of MSCI country equity returns with MSCI world returns	Covariance with world portfolio returns	Covariance of country bond returns with world portfolio returns
Intertemporal hedge	Covariance with changes in world portfolio returns	Covariance of MSCI country equity returns with changes in MSCI world returns	Covariance with changes in world portfolio returns	Covariance of country bond returns with changes in world portfolio returns

Table 3.1 Determinants of Equity and Bond Flows.

The regions we examine are based on the Morgan Stanley Capital International (MSCI) regional classification and are as follows:

- Asia-Pacific (excluding Australia, Japan, New Zealand): China, Hong Kong SAR, India, Indonesia, Korea, Malaysia, Philippines, Singapore, Thailand (9)

³⁷The model employs a six-month lagged capital control measure, for two reasons. First, capital control measures are expected to take effect with a time lag. Second, large flows could in fact prompt the imposition of capital controls, forcing an opposite (positive) sign as reflected in this type of the regression; the lagged capital control variable addresses this concern of reverse causality.

³⁸Two crisis dummies are included, one for the period between June 2007 and August 2008 (global credit crunch) and one for the period after September 2008 (Lehman Brothers bankruptcy).

- Eastern Europe, Middle East and Africa (EMEA): Bulgaria, Croatia, Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Romania, Russia, Slovakia, Slovenia, Turkey, Egypt, Nigeria, Saudi Arabia, South Africa (17)
- Latin America: Argentina, Brazil, Chile, Colombia, Mexico, Venezuela (6)
- G7 : United States, Canada, Japan, France, Germany, Italy, UK (7)
- Non-G7: Austria, Belgium, Denmark, Finland, Greece, Ireland, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, Israel, Australia, New Zealand (15)

As a preliminary analysis, we examine the relationship between equity flows and GDP. Figure 3.1 shows the positive relationship between equity flows and GDP growth since March 2009. An interesting case is China, which was forecast to have high GDP growth compared to peers, but received relatively lower equity inflows. This may reflect the existence of capital controls in China. A flip side of the story is Turkey, which has attracted more equity flows, despite a lower projected GDP growth rate.

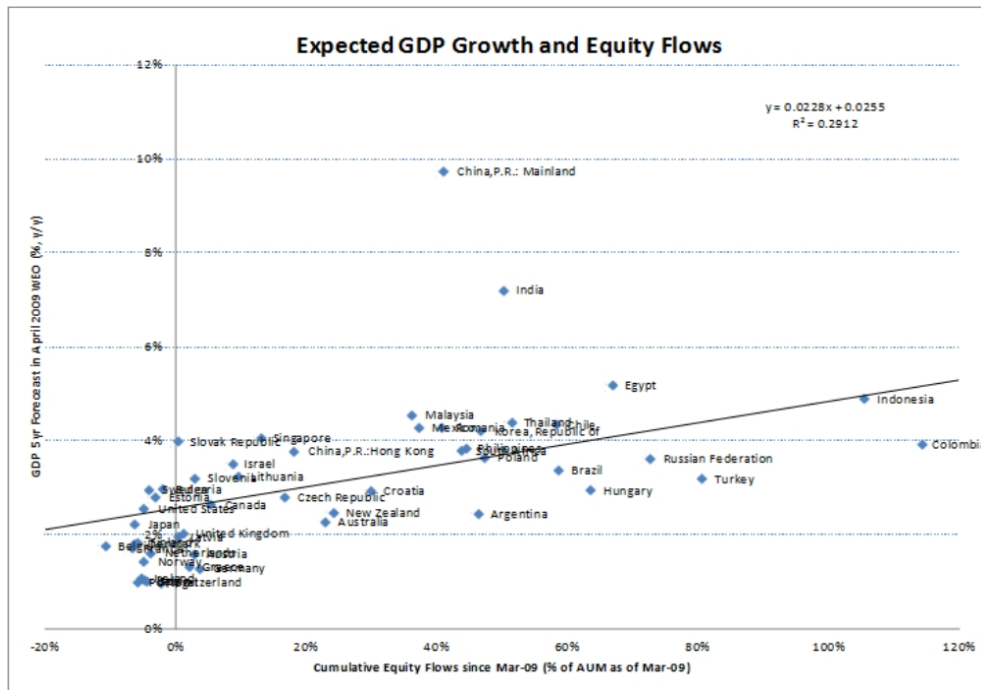


Figure 3.1: Expected GDP Growth and Equity Flows

5 Empirical Results

The analysis yields the following main results about the drivers of flows into equity and bond funds (Tables 3.2 provides detailed results):

- Interest rate differentials in most cases have no effect on flows into equity and bond funds. These flows generally do not respond to policy rate differentials in a statistically significant way. These results are generally invariant to using policy rate differentials relative to the G-4 (as used in the baseline regression), nominal policy rates, nominal or real long-term interest rates (for countries where long-term rates are available), nominal or real long-term interest rate differentials relative to the G-4, and lagged policy rate differentials.³⁹ The implications of

³⁹Because policymakers may use policy rates to dampen undesirable capital flows

this finding are discussed further below.

- Improving GDP growth prospects in general positively affect flows. Globally, an increase in the forecast GDP growth rate in the investment destination country leads to an increase in bond and equity investments. GDP growth is important for equity investors because higher GDP would lead to higher corporate earnings growth, making equities more attractive. It could also affect bond investors if higher GDP growth reduces credit risk, making bond investments more attractive.
- A rise in country risk generally reduces flows. The regression analysis confirms that, in many cases, an increase in country risk in emerging markets reduces their attractiveness for equity and bond investors. The effect is not statistically significant in advanced economies, perhaps partly because these showed little variation in country risk until recently.
- A rise in global risk generally reduces flows. Globally and for all regions, an increase in global risk (proxied by the VIX variable) discourages flows into equities and bonds.
- Lower return covariance generally leads to increased flows. In many cases, lower covariance of a country's equities and bonds leads to higher flows into these investments. This is as expected, since an asset that tends to have low covariance to other assets in the portfolio reduces the risk of the overall portfolio.

(which may partly flow into bond and equity investments), the regression may suffer from an "endogeneity" problem. To get around this issue, as noted, a regression was run with lagged policy rate differentials. Expected changes in foreign exchange rates (proxied by the forward less the spot rate) are not included in the regression because any expected change would be captured by the interest rate differential through covered interest parity.

- Higher uncertainty tends to reduce flows. Uncertainty about future exchange rates and GDP growth, measured by changes in the volatility of exchange rates and GDP forecast, are found in general to reduce flows into equities and bonds.
- Capital control measures show only weak effects. Capital control measures negatively affect bond flows on a global scale but not in most of the regressions for emerging markets. This weak finding may result in part because such controls are usually placed on money market and exchange rate instruments and not on longer-term equity and bond investments, where the interests of real-money investors lie; this is consistent with findings in other IMF studies (IMF (2010), in particular). Also, there is evidence that controls tend to lose effectiveness as market participants find ways to circumvent them, which occurs as long as the return on the controlled transaction exceeds the cost of circumvention.
- The crisis appears to have had an enduring effect on investor behavior. We find structural breaks in investor behavior after the global financial crisis. After the initial stage of the crisis (June 2007 to August 2008), there was a general slowdown in both equity and bond flows to all regions. However, after the second stage (beginning in September 2008), there was an increase in equity flows to Latin America (although there was no effect on Asian equity investments). There is for now no firm evidence that these effects have faded.⁴⁰

⁴⁰Specifically, the explanatory power of the crisis dummy variables do not improve significantly if it is terminated before the end of the sample, suggesting that the alteration during the crisis continues through the end of the sample.

Table 3.2. Panel Regression Results

	World		Asia		Latin America		Europe Middle East and Africa		G7 countries		Non-G7 advanced countries	
	equity	bond	equity	bond	equity	bond	equity	bond	equity	bond	equity	bond
Policy rate differential (host-G4 average)	-0.042 (-0.597)	-0.145 (-1.542)	-0.362** (-2.391)	-0.499 (-1.073)	-0.068 (-0.840)	-0.003 (-0.056)	0.081 (0.457)	-0.636*** (-3.246)	0.053 (0.152)	-0.519 (-1.120)	0.412*** (3.807)	0.410 (1.454)
GDP growth forecast	0.418*** (4.389)	0.775*** (5.557)	0.274** (2.426)	0.168 (0.790)	0.506** (2.509)	0.564*** (5.189)	0.483** (2.364)	1.265*** (5.479)	0.359** (2.054)	0.452*** (3.418)	0.256*** (3.376)	0.604*** (2.777)
Inflation volatility	-0.095 (-0.478)	-0.634** (-2.145)	-0.534** (-2.015)	-0.883** (-2.058)	0.293 (1.441)	0.147 (0.218)	0.654 (1.065)	-0.548* (-1.666)	-0.327* (-1.890)	-1.955** (-2.441)	-0.436** (-2.501)	-1.439* (-2.175)
GDP growth volatility	-1.908*** (-4.958)	-4.654*** (-8.389)	-2.835*** (-3.110)	-7.134*** (-8.558)	-2.997*** (-5.029)	-5.597*** (-7.783)	-2.423*** (-2.983)	-3.405*** (-3.356)	0.171 (0.419)	-3.410*** (-4.256)	-1.029*** (-3.287)	-3.463*** (-5.481)
Exchange rate volatility	-0.369*** (-3.630)	-0.835*** (-3.026)	-0.485*** (-2.165)	-1.161*** (-4.641)	-0.268 (-0.877)	-0.080 (-0.135)	-0.556*** (-1.177)	-1.104*** (-3.864)	-0.361*** (-2.805)	-0.850** (-2.422)	-0.334*** (-5.757)	-1.063*** (-9.843)
Return Covariance (cross-country)	-0.054 (-1.484)	-0.462*** (-6.532)	0.025 (0.721)	-0.656*** (-9.542)	0.235*** (4.034)	-0.422*** (-10.262)	-0.129*** (-3.506)	-0.275*** (-2.989)	0.002 (0.036)	-0.644*** (-5.005)	0.019 (0.599)	-0.574*** (-13.405)
Return Covariance (intertemporal)	-0.002 (-0.028)	0.094 (0.784)	0.001 (0.037)	0.198** (3.791)	-0.456*** (-6.069)	-0.289*** (-3.367)	0.120* (1.673)	0.049 (1.033)	-0.136* (-1.867)	0.077 (1.167)	-0.173*** (-6.011)	-0.030 (-0.690)
Country risk	-0.052* (-1.833)	-0.147*** (-2.996)	-0.157*** (-2.296)	-0.351*** (-4.957)	-0.066 (-0.963)	-0.282 (-1.594)	-0.080 (-0.963)	-0.151*** (-2.981)	-0.019 (-0.304)	-0.058 (-0.796)	-0.023 (-0.836)	0.005 (0.051)
VIX Index	-0.047*** (-8.368)	-0.071*** (-7.792)	-0.084*** (-10.264)	-0.073*** (-8.984)	-0.096*** (-6.437)	-0.079*** (-4.972)	-0.044*** (-3.892)	-0.117*** (-5.074)	-0.022*** (-9.109)	-0.047*** (-6.170)	-0.024*** (-6.977)	-0.051*** (-4.224)
Capital control index	-1.154 (-1.125)	-3.897*** (-2.625)	-7.776*** (-3.342)	-1.605 (-0.561)	0.364 (0.176)	5.839** (1.988)	0.981 (0.517)	-3.070 (-0.951)	-9.775* (-1.902)	-11.491*** (-2.915)	-0.902 (-1.113)	-8.481*** (-5.989)
Crisis dummy1	-0.723*** (-7.155)	-2.388*** (-16.002)	-0.516** (-2.375)	-2.216*** (-8.895)	0.205 (0.927)	-2.758*** (-6.505)	-0.617*** (-3.102)	-2.536*** (-6.603)	-0.848*** (-4.152)	-1.602*** (-4.147)	-1.023*** (-13.087)	-2.309*** (-9.220)
Crisis dummy2	-0.513* (-1.914)	-2.308*** (-6.899)	0.327 (0.843)	-3.597*** (-7.363)	1.033*** (5.464)	-4.119*** (-6.103)	-1.285 (-1.104)	-3.632*** (-5.847)	-0.823*** (-3.295)	-0.203 (-0.391)	-0.741*** (-8.102)	-0.906** (-2.079)
Time trend	0.000 (0.052)	0.071*** (10.724)	-0.026*** (-3.253)	0.105*** (10.226)	-0.034*** (-4.048)	0.088** (6.367)	0.033 (1.575)	0.069*** (9.091)	0.014** (2.361)	0.033*** (3.651)	0.003 (1.227)	0.043*** (5.639)
Constant	0.932 (1.240)	-7.071*** (-6.913)	8.375*** (4.193)	-11.554*** (-4.710)	4.979*** (3.203)	-12.243*** (-9.145)	-4.228 (-1.542)	-11.027*** (-5.524)	-0.227 (-0.392)	-2.464** (-1.950)	0.352 (1.385)	-3.551*** (-3.795)
Number of countries	50	50	9	9	6	6	13	13	7	7	15	15
Number of observations	2,966	2,845	504	504	395	397	527	461	490	485	1,050	998
R-squared (within)	0.160	0.472	0.271	0.510	0.147	0.490	0.214	0.615	0.295	0.526	0.400	0.475
R-squared (between)	0.047	0.003	0.493	0.288	0.068	0.003	0.370	0.173	0.073	0.351	0.121	0.210
R-squared (overall)	0.083	0.303	0.004	0.477	0.145	0.418	0.212	0.502	0.054	0.293	0.388	0.428

Source: Authors' estimates.

Note: The table presents panel fixed-effects regressions on factors affecting equity and bond flows to advanced and emerging market economies between January 2005 and May 2011. The results are presented for the whole sample as well as for five separate regions. Dependent variables are monthly equity and bond flows as a proportion of assets dedicated to each country at the beginning of the month. The policy rate differential is the difference between the policy rate in host country and the simple average policy rate for G4. GDP growth forecast is the change in the one-year-forward GDP forecast for host country, provided by Consensus Economics. Inflation volatility, GDP growth volatility, exchange rate volatility are the standard deviation of inflation, GDP growth, and exchange rate forecasts over the past year. Country risk is a measure of country risk from International Country Risk Group (ICRG). Risk is the VIX index as a measure of global risk. Return covariance 1 is a measure of the covariance of country returns with the world portfolio return (cross-country correlation factor). Return covariance 2 is a measure of the covariance of country returns with changes in the world portfolio return (intertemporal correlation factor). Control variables are capital control index, crisis dummy 1, and crisis dummy 2. Capital control index is the 6 month lagged capital control index produced by the Monetary and Capital Markets Department at the International Monetary Fund. Crisis dummy 1 represents the period after June 2007 (global credit crunch). Crisis dummy 2 represents the period after September 2008 (Lehman's collapse). All independent variables, except for control variables, are in first-differences. A time trend is included. T-statistics are in parenthesis. ***, **, and * denote statistical significance at the 1 percent, 5 percent, and 10 percent level of confidence based on robust standard errors.

The empirical results show that investors' asset allocation behavior changed at the time of the crisis. The dummies included in the regressions to capture the effects of the crisis show that globally, and for most regions separately, investors changed their behavior toward equities and bonds in a way not captured by the regular drivers (that is, the other independent variables in the regression). This "crisis effect" began, first, at the onset of the crisis, in mid-2007, and continued around the time of the Lehman Brothers bankruptcy, in September 2008. These were statistically significant changes in behavior.

A useful metric is the Z-score, which relates the size of the change in asset allocation at the time of the crisis to shocks that would normally have been experienced before the crisis. The Z-score is the size of the change in allocation implied by the dummy coefficient, minus the pre-crisis mean, divided by the pre-crisis standard deviation. Note that the Z-score is meaningless if the dummy is not statistically significant, as in such cases there was no statistically significant change at all in asset allocation at the time of the crisis. Under the assumption of a normal distribution for shocks to investment flows, a Z score of about 2 indicates that the shock would be classified as among the 5 percent most severe.

The Z-scores indicate that the crisis effect was quite large for bonds and advanced economy equities (Table 3.3). For bonds, the Z-score was in many cases close to, or exceeded 2, so that the outflows from bond funds during the crisis were among the 5 percent most severe compared to the pre-crisis period. For equities, there is a distinction between emerging markets and advanced markets. In emerging markets, even though the coefficients for the first dummy (June 2007–August 2008) were generally significantly

negative, the effects were small (i.e., in line with usual volatility in the pre-crisis period). In addition, the coefficients on the second crisis dummy (after September 2008) were not significantly different from zero, except for Latin America, where the coefficient was positive and significant. In these cases, the low Z scores imply that investors in emerging market equities continued during and after the crisis to let themselves be guided by the established drivers of asset allocation. Not so in advanced markets, where the “crisis” effect on equity funds was large, with Z scores around 2, meaning that the crisis-induced outflows from equity funds in advanced markets were among the 5 percent most severe compared to the pre-crisis period.

Table 3.3. Economic Significant of Crisis Indicators

	World		Asia		Latin America		Europe, Middle East and Africa		G-7 Countries		Non G-7 Advanced Countries	
	Equities	Bonds	Equities	Bonds	Equities	Bonds	Equities	Bonds	Equities	Bonds	Equities	Bonds
Crisis Indicator 1	-0.723*** (-7.155)	-2.388*** (-16.002)	-0.516** (-2.375)	-2.216*** (-8.895)	0.205 (0.927)	-2.758*** (-6.505)	-0.617*** (-3.102)	-2.536*** (-6.603)	-0.848*** (-4.152)	-1.602*** (-4.147)	-1.023*** (-13.087)	-2.309*** (-9.220)
Crisis Indicator 2	-0.513* (-1.914)	-2.308*** (-6.899)	0.327 (0.843)	-3.597*** (-7.363)	1.033*** (5.464)	-4.119*** (-6.103)	-1.285 (-1.104)	-3.632*** (-5.847)	-0.823*** (-3.295)	-0.203 (-0.391)	-0.741*** (-8.102)	-0.906** (-2.079)
Pre-crisis Mean	0.61	1.03	0.83	1.16	0.95	1.65	0.45	1.32	0.38	0.34	0.58	0.65
Pre-crisis Std. Deviation	1.71	1.74	1.41	1.72	2.43	1.34	2.52	2.36	0.61	1.14	0.67	1.13
Z-score for crisis period 1	-0.78	-1.97	-0.96	-1.96	...	-3.30	-0.43	-1.64	-2.00	-1.70	-2.40	-2.63
Z-score for crisis period 2	...	-1.92	...	-2.76	0.04	-4.32	...	-2.10	-1.96	...	-1.98	-1.38

Source: Authors' estimates.

Notes: In the first two columns the table re-reports the regression coefficients and standard errors for the two crisis indicator variables from the regression analysis. The next two rows report the mean and standard deviation of the growth of equity and bond flows estimated for the period before June 2007 (the pre-crisis period). The last two rows report the z-scores of both equity and bond flows using the mean and standard deviations reported in the previous two rows. The z-score is defined as the regression coefficient minus the pre-crisis mean, divided by the pre-crisis standard deviation, and is one metric for gauging the size of the regression coefficients on the crisis indicator variables.

"..." indicates that the original estimate is not statistically significant, so no Z-score is calculated.

6 Currency Composition of Central Banks

6.1 Background

While the overwhelming majority of financial assets is owned and managed by private investors, sovereign investors have grown to become important players in international capital markets. Sovereign wealth funds (SWFs) hold some \$4 trillion in assets, while international reserves amount to \$10 trillion. Their combined assets amount to about a quarter of the assets under management of private institutional investors.⁴¹ Their increasing size makes sovereigns important investors in international capital markets.

Management of international reserves are distinct from sovereign wealth funds in that they are explicitly held for balance of payments or monetary policy purposes, and as a result, the objectives of reserve managers may depart from pure return maximization. The asset allocation and management of reserves can be different from other types of asset management: it is traditionally driven by safety, liquidity, and return, in that order (IMF (2011d)). The requirement that reserves are available at short notice and at low cost to meet balance-of-payments needs and financial stability objectives leads to an allocation that is traditionally dominated by short-term government bonds issued by only a few countries.

However, global foreign exchange reserve holdings (excluding gold) have grown so fast in recent years that their size for many countries now exceeds the amount needed for balance of payments and monetary purposes. After having expanded five-fold between 2000 and 2008, reserve levels saw a brief decline during the global financial crisis, but rebounded quickly and

⁴¹Using the IMF's definition of foreign exchange reserves and sovereign wealth funds; see Section 6.2.

accumulation has resumed. Today's reserve levels in emerging and developing economies well exceed levels traditionally considered adequate IMF (2011d)).

This means that an increasing share of reserves could be available for potential investment in less liquid and longer-term risk assets. A new IMF estimate puts core reserves needed for balance of payments purposes in emerging market economies at \$3.0-4.4 trillion, leaving \$1.0-2.3 trillion available to be invested beyond the traditional mandate of reserve managers, i.e. more like SWFs.⁴² Some central banks have facilitated this by splitting their reserves into a "liquidity tranche" and an "investment tranche," with the latter aiming to generate a higher return over the long run (Borio and others, 2008). In the aggregate, however, these investment tranches are, to date, quite small, and government bonds remain the dominant asset class in reserves.

Before the financial crisis, concerns about high opportunity costs of holding large reserves and a low interest rate environment pushed central banks in the direction of expanding the investment tranche. Other factors that contributed to this trend include the shrinking supply of government debt (in the late 1990s, early 2000s) and the "learning by investing" argument for non-traditional asset classes. In the recent financial crisis, however, reserve managers interrupted their trend towards diversification of the investment tranche, and central banks rapidly exited unsecured bank deposits. The proportion invested in these deposits dropped rapidly from its peak in July 2007 (17.2 percent of reserves including gold at market prices) to less than 5 percent in June 2010. Several surveys, in particular

⁴²Metric for reserve adequacy developed in IMF (2011d); the suggested adequacy range is 100-150 percent of the metric, leading to the ranges given here.

those conducted annually by Central Banking Publications, and other studies (Pihlman and Van der Hoorn (2010)) confirm (qualitatively) that reserve managers' risk aversion increased and that reserve managers participated in the global flight to quality and liquidity.

Now as we emerge from the financial crisis and the European debt crisis, the reserve managers are faced with another low-interest rate environment. Looking ahead, some reserve managers are rethinking their asset allocation strategies (IMF (2011a)).

This raises an important set of questions: Will the reserve managers deploy 'excess' reserves to buy riskier securities to seek higher yields? Will they continue the pre-crisis trend of increasing the diversification of their investment tranche? The answer to this question has important ramifications for global asset markets, since the size of reserve accumulation makes the reserve managers an important player in a market where the private asset allocators have become more risk averse and less willing to hold risky assets. Also, given the unprecedented size of reserves held by central banks, the opportunity cost of holding most of the reserve portfolio in securities that offer close-to-zero yield, there is reason to believe that reserve managers may consider such a shift.

A lack of detailed data on asset allocation by reserve managers makes it difficult to investigate their investment behavior and answer this question. Few countries publish details on the composition of their international reserves by currency (i.e., the destination country for reserve investments), asset class, or maturity. In addition, transaction data for the buying and selling of foreign exchange by central banks is also generally not publicly available.

However, the IMF collects data on the size and currency composition of foreign exchange reserves of member countries in its Currency Composition of Official Foreign Exchange Reserves (COFER) database that could be used for a limited empirical analysis. These data are not published, and are not available for all countries. Still, they may be used for a limited empirical investigation of reserve management when we consider that the currency composition of reserves is equivalent to the country destination of investment by reserve managers. While the choice of reserve currency is subject to additional considerations of balance of payments need, in principle, reserve managers have a choice in which currency to hold their reserves, as most reserve currencies have deep and liquid exchange markets and can quickly be converted into a different currency if needed.

Given the data, we can answer the following question: do reserve managers also respond to the incentives of private investors, such as risk and return? Reserve managers are not expected to behave as fully return maximizing investors for their core reserves. Still, they may be more responsive to risk and return incentives at the margin, i.e. for reserves that exceed the core—their “investment tranche.” If we find reserve managers responding to such incentives, then the investment incentives that drive private investors may also induce reserve managers to be a potential source of longer-term risk capital provision by sovereigns, in addition to SWFs.

6.2 Defining Foreign Exchange Reserves and Sovereign Wealth Fund

6.2.1 Foreign Exchange Reserves. — The IMF’s primary definition of reserves is contained in Chapter VI of its Balance of Payments and International Investment Position Manual, sixth edition (2009): “Reserve assets are

those external assets that are readily available to and controlled by monetary authorities for meeting balance of payments financing needs, for intervention in exchange markets to affect the currency exchange rate, and for other related purposes (such as maintaining confidence in the currency and the economy, and serving as a basis for foreign borrowing).”

The IMF defines reserve assets further by stating that “reserve assets must be both denominated and settled in foreign currency” (paragraph 6.71); that “reserve assets must be denominated and settled in convertible foreign currencies” (paragraph 6.72); and that “reserve assets, other than gold bullion, must be claims on nonresidents.” (paragraph 6.65). It should be noted that there are not many restrictions on the asset classes that can be used for reserve asset investments. The main constraints concern liquidity (“readily available”) and they must constitute claims on “nonresidents” in “convertible foreign currencies.”

6.2.2 Sovereign Wealth Funds. — SWFs are defined as follows: “SWFs are defined as special purpose investment funds or arrangements, owned by the general government. Created by the general government for macroeconomic purposes, SWFs hold, manage, or administer assets to achieve financial objectives, and employ a set of investment strategies which include investing in foreign financial assets. The SWFs are commonly established out of balance of payments surpluses, official foreign currency operations, the proceeds of privatizations, fiscal surpluses, and/or receipts resulting from commodity exports.”

This definition excludes, inter alia, foreign currency reserve assets held by monetary authorities for the traditional balance of payments or monetary policy purposes, operations of state-owned enterprises in the tradi-

tional sense, government-employee pension funds, or assets managed for the benefit of individuals.

Three key elements define an SWF:

- **Ownership:** SWFs are owned by the general government, which includes both central government and sub-national governments.⁴³
- **Investments:** The investment strategies include investments in foreign financial assets, so it excludes those funds that solely invest in domestic assets.
- **Purposes and Objectives:** Established by the general government for macroeconomic purposes, SWFs are created to invest government funds to achieve financial objectives, and (may) have liabilities that are only broadly defined, thus allowing SWFs to employ a wide range of investment strategies with a medium- to long-term timescale. SWFs are created to serve a different objective than, for example, reserve portfolios held only for traditional balance of payments purposes. While SWFs may include reserve assets, the intention is not to regard all reserve assets as SWFs.⁴⁴

Furthermore, the reference in the definition that SWFs are “commonly established out of balance of payments surpluses, official foreign currency operations, the proceeds of privatizations, fiscal surpluses, and/or receipts

⁴³Note that the use of the word arrangements as an alternative to funds allows for a flexible interpretation of the legal arrangement through which the assets can be invested. SWFs vary in their institutional arrangements, and the way they are recorded in the macroeconomic accounts may differ depending on their individual circumstances. See also the IMF’s Government Finance Statistics Manual, 2001.

⁴⁴Likewise, the intention is not to exclude all assets on the books of central banks: SWFs can be on the books of central banks if they also are held for purposes other than balance of payments purposes (e.g., as intergenerational wealth transfer).

resulting from commodity exports” reflects both the traditional background to the creation of SWFs—the revenues received from mineral wealth—and the more recent approach of transferring “excess reserves.”

It should be noted that reserve assets and assets held by an SWF can overlap. Reserve assets can be held within an SWF. This can only occur, though, when the SWF is permitted to transact in such assets only on terms specified by the monetary authorities or only with their express approval.” (see Balance of Payments and International Investment Position Manual, sixth edition (2009), paragraph 6.67).

6.3 Rise in Reserves and Change in Composition

Global foreign exchange reserves (excluding gold) have grown five-fold since end-2000, and now represent approximately 6.2 percent of global debt and equity markets’ capitalization (end-2009). The growth of reserves has been concentrated in emerging and developing economies, which have accumulated reserves at a rapid pace in the aftermath of the Asian financial crisis of the late 1990s. Following a brief decline during the recent global financial crisis, reserve levels rebounded quickly and reserve accumulation has resumed (Figure 3.2). Today’s reserve levels in emerging and developing economies well exceed levels traditionally considered adequate (IMF (2011d)). These reserves are mostly managed by central banks, although in some countries this is the responsibility of the ministry of finance.

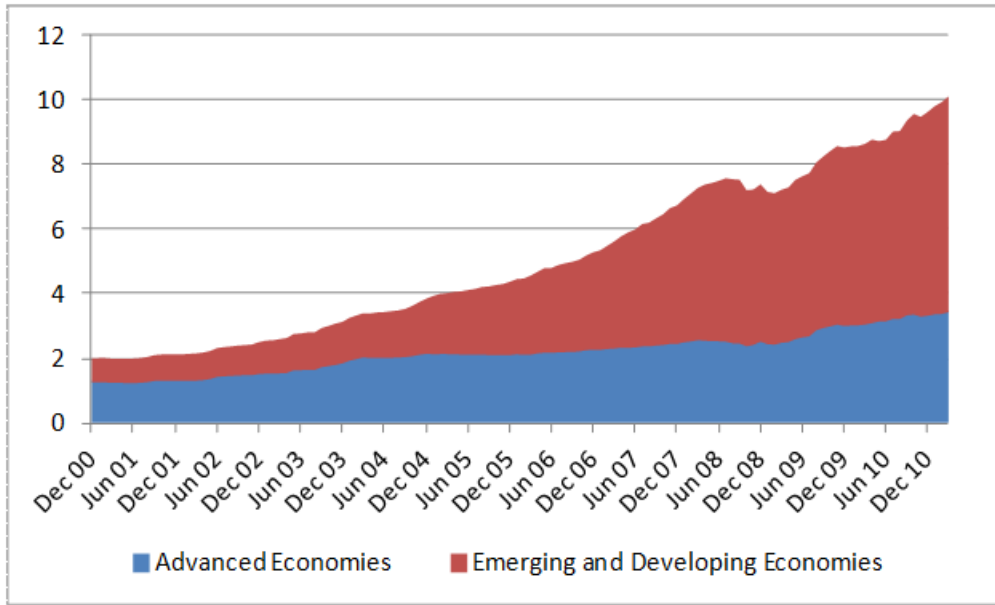


Figure 3.2: Foreign Exchange Reserves Excluding Gold (in trillion of U.S. dollars). Source: IMF, International Financial Statistics (IFS).




Figure 3.3: Instrument Composition of Official Reserves Including Gold. Source: IMF.

The focus on safety and liquidity of reserves returned at the start of the global financial crisis, when central banks rapidly exited unsecured bank deposits (Figure 3.3). The proportion invested in unsecured bank deposits dropped rapidly from its peak in July 2007 (17.2 percent of reserves including gold at market prices) to less than 5 percent since June 2010, most likely as a result of the increased perceived risk associated with this asset class. At the same time, a previous trend of selling gold has been reversed: the total holdings of gold among reporting central banks peaked at 25,353 tons in October 2000 and dropped to 22,599 tons in March 2009. Since then, central banks have purchased 367 tons of gold on a net basis. Given the recent increase in the market price of gold, the proportion of reserves held in gold has increased, despite the large increase in total reserves.

The currency composition of reserves changed with the introduction of the euro, but has been fairly stable in recent years (Figure 3.4). The currency composition is closely related to the objectives for holding reserves. Countries hold currencies in line with the composition of their short-term external debt and/or import basket, which tend to change only slowly over time. There may be an over-weighting of the U.S. dollar as the most liquid and most widely-used currency in foreign exchange markets. The liquidity of the U.S. dollar makes it therefore the preferred currency for foreign exchange interventions.

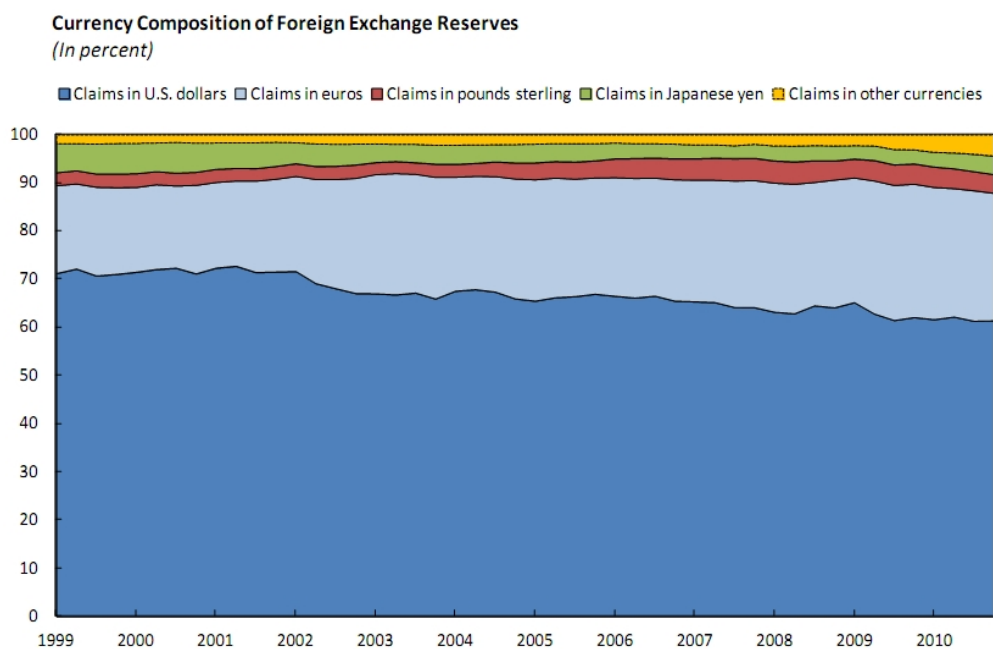


Figure 3.4: Currency Composition of Foreign Exchange Reserves.

The stability of the currency composition in the face of large swings in exchange rates indicates that central banks rebalance the currency composition of their reserves (see Lim (2007)). Figure 3.5 plots the annual change in the relative share of euros versus U.S. dollars at constant exchange rates against the annual change in the euro/dollar exchange rate. There is a clear negative correlation indicating strong rebalancing effects: whenever the euro appreciates, central banks sell euros against dollars and vice versa, thereby reducing volatility in the foreign exchange markets. The negative correlation persists in recent years; rebalancing strategies appear to have been unaffected by the crisis.

Beyond these longer-term strategic asset allocation objectives of reserve managers, do they also respond to the incentives of private investors, such as risk and return? Since monetary authorities do not necessarily maximize returns, their sensitivity to interest rate changes may be fairly low. On




Figure 3.5: Changes in Euro Holdings Versus Changes in the Euro/Dollar Exchange Rate (in percent). Source: IMF, Currency Composition of Foreign Exchange Reserves (COFER).

the other hand, foreign reserves are primarily invested in fixed-income securities, with the bulk of these investments likely in short-duration liquid instruments. Thus, it is plausible that reserve managers may be sensitive to interest rate changes, analogous to carry traders in the private market.

6.4 Results

In what follows, this question is examined using the COFER database. The database contains country-level currency composition data from the 1960s to the present. The investigation uses quarterly data from 1999 to 2011 for 102 countries that include a number of the variables we used for the private mutual fund data above, in addition to variables to measure the conventional objectives of reserve managers, including debt to GDP ratios

and export and import propensities.⁴⁵ The four dependent variables used in this case are the shares of total reserves allocated to the four major currencies, the US dollar, euros, pound Sterling, and Japanese, which constitute more than 90 percent of total reserve holdings for most of the countries in our sample.

⁴⁵Since the portfolio choice of reserve managers is determined simultaneously, we jointly estimate our system of regression equations using the seemingly unrelated regressions model.

Table 3.4 Regression Results for the Currency Composition of Reserves

	U.S. dollar share	Euro share	Pound sterling share	Yen share
U.S. policy rate	0.0048***	-0.0029**	-0.0008	0.0003
Euro policy rate	-0.0016	0.0026	-0.0003	-0.0011
U.K. policy rate	-0.0036	0.0018	0.0011	0.0009
Japan policy rate	0.0146	-0.0112	-0.0015	0.0014
Euro-U.S. FX volatility	0.0109***	-0.0061**	-0.0008	0.0009
U.K.-U.S. FX volatility	-0.0058	0.0015	0.0016	0.0002
Japan-U.S. FX volatility	0.0020	-0.0002	-0.0008	0.0007
U.S. GDP forecasts	-0.0006	-0.0012	0.0001	0.0003
Euro GDP forecasts	0.0010	-0.0029	0.0002	-0.0002
U.K. GDP forecasts	-0.0012	0.0021	0.0007	0.0001
Japan GDP forecasts	0.0011	-0.0002	-0.0003	0.0002
Crisis indicator 1	-0.0158**	0.0013	0.0031	0.0029
Crisis indicator 2	-0.0046	-0.0012	-0.0003	0.0043**

Source: IMF staff estimates.

Notes: The table presents results of a system of regression equations estimated using seemingly unrelated regressions. The dependent variables are shares of foreign reserves allocated to the four major reserve currencies. The omitted category is 'other' currencies, and the shares of the five categories add up to one. Data for the dependent variable is from the COFER statistical database, at a quarterly basis from Q1 1999 to Q1 2011 for 102 countries. The policy rate variables measure the short-term policy rate for the four major currencies. The FX volatility is computed using the exchange rate volatility of each country (with U.S. dollar as base currency) over a rolling period of one year. GDP forecasts are mean forecasts of one-year GDP growth acquired from Consensus Forecasts. Crisis indicator 1 represents the period between June 2007 and August 2008 (global credit crunch). Crisis indicator 2 represents the period after September 2008 (Lehman's collapse). The regression also controls for total government debt to GDP ratio, real GDP per capita, import share of GDP, export share of GDP, and foreign exchange regimes. ***, ** & * denote statistical significance at the 1%, 5% and 10% level of confidence based on robust standard errors.

The key results of the analysis are as follows (Table 3.4):

- Reserve managers appear to respond to U.S. interest rates: as shown in the first row of the table below, increases in the U.S. dollar interest-rate are associated with a rebalancing away from euro and towards the U.S. dollar.
- An increase in the volatility of the euro/dollar exchange rate tends to favor the U.S. dollar as a reserve currency, also at the expense of the euro.
- The shares of the other two main reserve currencies, the pound and the yen, appear not to be affected by interest rates or exchange-rate volatility.
- Economic growth differentials (which are found to be important for private asset allocation (IMF (2011a)) appear not to matter for the currency composition of international reserves.
- At the start of the global credit crunch in the summer of 2007, there was a drop in the share of U.S. dollars in international reserves. This may have been associated with central banks providing dollar liquidity support to domestic banks.

6 Concluding Remarks

With respect to institutional investors, the above findings show the main “pull” and “push” factors for these investors’ asset allocations. The main “pull” factor is the long-term growth prospects in destination countries, which may be diminished to some extent by rising country risk. The

main “push” factor is the risk appetite of global investors. These factors are robust over the period studied (2005–11).⁴⁶

The most notable of the above findings is that interest rate differentials do not significantly affect real money investor flows. Neither bond nor equity flows respond to changes in interest rate differentials, globally or for any of the regions. This result is not fully in line with previous findings (see, for example, IMF (2011b)).⁴⁷ A few of the possible explanations are the following:⁴⁸

- The result applies only to real-money flows in and out of bond and equity investment funds. Short-term flows, usually seen as more interest-sensitive, are less likely to be invested through these funds; leveraged flows (including from the carry trade), which are not captured in these data, may still respond to differentials in policy rates and other interest rates.
- The EPFR data include bond funds that hold bonds with a wide range of maturities, which respond differently to changes in rates at different points along the yield curve. Therefore, the effect of short-term rates on bond flows, presumably concentrated on short-term bonds, is

⁴⁶These push and pull factors are also found to be important according to the IMF Survey on Global Asset Allocation (IMF (2011a), Annex 2.2).

⁴⁷Although Forbes and Warnock (2011) also found weak evidence for the effect of global interest rates on gross capital flows using balance of payments data.

⁴⁸One possible explanation was not borne out in the data. Countries with high interest rate differentials may carry risks of large and sudden devaluations (the “peso problem”). There may therefore be a heterogeneous impact of policy rate differentials on bond flows that may increase the standard error of the estimated coefficient, rendering it insignificant. To try to solve this potential problem, the regression was rerun including an interaction term defined as the product of the policy rate differential and the country risk. Whereas the interest rate differential was positively associated with bond flows when the interaction term is included for the global sample, the results in the regional regressions were unchanged, with bond flows not significantly positively responsive to interest rate differentials.

obscured by possible differing (and perhaps opposing) effects on long-term bonds. The converse appears also to be true, as using long-term rates in the regressions does not change the results. Thus, whereas different interest rates along the yield curve may affect flows into bonds of different maturities, their effect on total flows into bonds of all maturities is not statistically significant in these data.

The finding of this study that interest rate differentials do not affect bond and equity flows should not be extended to capital flows in general, for two reasons: First, flows in and out of bond and equity investments may come out of domestic funds, and to the extent that they do, they would not directly affect capital flows. Second, as noted, capital flows may be dominated by other types of investments, including flows from leveraged investors (such as the carry trade), which this analysis does not cover.

By contrast, we find a stronger role for interest rates with respect to reserve allocation of central banks. Reserve managers appear to respond to U.S. interest rates: increases in the U.S. dollar interest-rate are associated with a rebalancing away from euro and towards the U.S. dollar. We also find that at the start of the global credit crunch in the summer of 2007, there was a drop in the share of U.S. dollars in international reserves. This may have been associated with central banks providing dollar liquidity support to domestic banks.

Chapter 4: Dynamic Incentives and Welfare Costs of Taxation

1 Introduction

Executive compensation in financial institutions has become a prime target for criticism and regulation, in response to the turmoil of the financial crisis that began in 2007. Much of this criticism is fueled by the enormous growth in executive compensation in the past two decades. Since the inception of the crisis, the US government and several other Western nations have implemented various bailout plans, totalling several trillion dollars, to prop up the financial system from failing. This expanded government role in financial institutions has allowed public policy participants to prominently voice their criticism of excessively high executive compensation. Consequently, several countries, like UK and France, have passed legislation to tax executive compensation at financial institutions at very high rates. Other countries are considering broader measures of taxing executives. Such measures stem mostly from the notion that top executives are being grossly overpaid, and the perception that there is a growing wedge between compensation and actual value-added.

This paper informs this debate by analyzing the welfare losses of taxation in a simple dynamic moral hazard model under symmetric information. We analyze how the principal-agent relationship is affected by taxation, and the consequent welfare implications. We show that the strength of incentives is central to understand the welfare costs of taxation. When ex-

explicit incentives are chosen optimally in a compensation contract, then the deadweight loss of taxation is proportional to the degree of high-powered incentives. That is, the welfare cost of taxation will be higher when the agent is incentivized to exert effort using very high-powered incentives. This is because the elasticity of taxable income is proportional to the power of incentives. That is, when the power of incentives is high, an increase in tax rate leads to a relatively greater reduction in taxable income. Hence, the deadweight loss is higher when the power of incentives is higher.

Critics often argue, however, that the explicit incentives are chosen exogenously (or suboptimally). In particular, due to weak corporate governance, the compensation contracts may be such that executive compensation is not highly linked to performance (Bebchuk and Fried (2006)). In such a case, the relation between explicit incentives and deadweight loss need not hold. Instead the career concerns of executives become a key determinant of the deadweight loss of taxation. This is because under such contracts, the incentives to provide effort, in general, will arise from both the exogenously-determined explicit incentives, and career concerns. Under career concerns the incentives are such that the executive's current effort will improve the market's perception of his talent, resulting in higher compensation in the future. Hence, under such a setting, the deadweight loss of a tax in period t will depend on the sensitivity of market's perception of the executive's talent in period t to his current effort.

This paper is related to several literatures in the economics of organizations, corporate finance and public finance. Relatively few papers exist that introduce taxation in a principal-agent setup. From a contract theoretic perspective the results of this exercise is useful in understanding the effect of

taxation on the principal-agent relationship and the consequent changes in the incentive structure. A huge literature in corporate finance deals with executive compensation, and a few papers have explored the impact of taxes on executive compensation. Holte (2007) is a recent paper that explores this linkage. Using a model of career concerns, he argues that lowering of top income tax rates over the past decades, strengthened career incentives, which in turn led to better talent identification and hence greater top income inequality. Katuscak (2005) develops a theoretical model in which an increase in the marginal tax rate decreases the equilibrium level of managerial effort and the after-tax pay-to-performance sensitivity. He subsequently tests this model and finds some empirical support for it. Hall and Liebman (2000) explore the implications of tax changes over the past decades. They argue that the tax changes cannot explain the dramatic increase in the share of executive compensation paid through stock options. Related to this, recent work by Frydman and Molloy (2011) use a sample of top executives in large firms from 1946 to 2005, and find little response in the mix of executive compensation to changes in taxes.

This paper differs from this existing literature on several margins. Primarily, this paper focuses on deriving the welfare implications of taxing executive compensation, which is governed by a principal-agent relationships between the executives and the corporate board. It is crucial to understand the welfare costs imposed by such the new legislations introduced in several countries that impose substantially high taxes on executives. The paper also derives sufficient statistics for welfare costs of taxation under principal-agent relationships. This allows for potential empirical investigation of such welfare costs. The next section details the model.

- Our model also has implications for optimal taxation in the presence of endogenous effort choices. We rely on a model of symmetric uncertainty (between an employee, an employer, and the market). While understanding the effects of taxation in models of asymmetric information is important, such a model captures some important aspects of the labor market. We feel that this model is especially informative about settings in which a worker is recently promoted to or hired into a job with significantly different demands than his previous job.
- It has been a challenge to quantify the dynamic efficiency loss of taxation in the public finance literature. While in general, it is impossible to know exactly what these effects are, the contracting literature proposes several models of dynamic incentives. To the extent that one believes that such a model captures the important determinants of compensation practices, it is possible to compute the welfare impacts of taxation within the context of such a model. Further, for the models we consider in this paper, the welfare loss depends only on a single economic parameter [Chetty (2009)]. In the case of the mixed-incentives model of Gibbons and Murphy (1992), this parameter is also a measure of the quality of the contracting environment. In the pure career-concerns model of Holmström (1999), this parameter is a measure of how informative the market perceives a worker's current performance is about his ability. It turns out that in each of these models, given only data on the cross-sectional distribution of wages within a cohort in a particular career, these parameters can be extracted from the evolution of cross-sectional inequality, potentially allowing for empirical analysis.

- Frydman and Molloy (2011) find that there are little changes in observable compensation practices following an increase in tax rates. In our model, performance pay is independent of the tax rate. However, taxable income is not, to the extent that performance pay is an important determination of effort. To the extent that career concerns are important, it is not the current tax rate that is important in determining the before-tax wages, but rather the market's beliefs about the worker's beliefs about future tax rates. Thus, our model explains why a change in the tax rate need not result in a change in the observed performance pay component, nor even in a change in taxable income.

In the following two sections we derive simple expressions for the dead-weight loss of taxation under two cases of the model: one in which the strength of explicit incentives is determined endogenously as above, and one in which the strength of incentives are exogenously fixed to zero. The latter case is the pure career-concerns model analyzed by Holmström (1999).

2 Model with Endogenous Explicit Incentives

2.1 Setup

Our setup closely follows the framework of Gibbons and Murphy (1992) who analyze the optimal provision of explicit performance-based incentives when workers are motivated by career concerns (Holmström (1999), Dewatripont, Jewitt, and Tirole (1999)). In order to simplify the analysis, we assume that workers are risk neutral, which necessitates a multitask (Holmstrom and Milgrom (1991), Baker (1992), Baker (2002)) model of the agency problem as opposed to the standard risk-incentives tradeoff of Holmstrom

(1979) and others. In the appendix, we analyze a similar model with risk averse agents. For clarity, we will focus on a two period problem, though extending the model to T periods is possible under restrictive assumptions about the distribution of noise terms.

There are a continuum of principals and a single agent who, at time t , chooses an effort vector $a_t \in \mathbb{R}_+^N$, where $N \geq 2$ at cost $c(a_t) = \frac{1}{2} \sum_{\ell=1}^N a_{\ell t}^2$. Output, which is unobservable by all parties and non-contractible, is given by

$$y_t = \sum_{i=1}^N f_i a_{it} + \eta + \varepsilon_t,$$

where η is the agent's ability and is symmetrically unobservable by the agent and all the principals. While principals care about y_t , neither they nor the market can observe it in the time frame the agent works for them. However, there is an observable performance measure which is commonly observed by all principals and is verifiable by a court, given by

$$p_t = \sum_{i=1}^N g_i a_{it} + \eta + \varepsilon_t.$$

Throughout, we will assume that $\|f\| = \|g\| = 1$, which is akin to assuming that the marginal impact of effort on the performance measure and on the output of the firm is of similar magnitude as the impact of ability. This is not without loss of generality, but we will indicate later where it can be relaxed. We assume that the principal can write enforceable contracts contingent on the performance measure but not on the actual output, and it can be shown that under risk neutrality, linear contracts of the form $w_t = s_t + b_t p_t$ are optimal. These wages are taxed at rate τ_t , which is exogenously imposed by the government. Agents care about the discounted sum of their

flow wages minus effort costs

$$U(\{w_t\}, \{a_t\}) = \sum_{t=1}^2 \delta^{t-1} (w_t - c(a_t)),$$

where $\delta \leq 1$ is the discount factor.

The timing is as follows. For each period t , there are four stages. In the first stage, a continuum of principals offers contracts (s_t, b_t) to the agent. In the second stage, the agent chooses which contract (of any) to accept. In stage three, if the agent has accepted the contract, he chooses effort vector a_t at private cost $c(a_t)$. Finally, in the fourth period, the performance measure p_t is realized, and the agent receives wage $w_t = (1 - \tau_t)(s_t + b_t p_t)$.

Ability is unobserved by all, but the principals and the agent have a common prior $\eta \sim N(m_1, h_1^{-1})$, where h_1 is the precision of the prior. We also assume that ε_t is distributed according to $N(0, h_\varepsilon^{-1})$. If we let $\varphi = \frac{h_\varepsilon}{h_1 + h_\varepsilon}$ denote the signal-to-noise ratio, then a principal who observes p_1 and conjectures that the agent chose effort vector \hat{a}_1 believes $\eta | p_1 \sim N(m_2, (h_\varepsilon + h_1)^{-1})$, where $m_2 = (1 - \varphi)m_1 + \varphi(p_1 - g \cdot \hat{a}_1)$ is a weighted average of two estimates of the agent's ability, the prior mean and the performance measure in excess of expected effort.

2.2 Equilibrium

Letting H^{t-1} denote the history of performance measures observed up to time t , perfect competition among principals yields a zero profit condition $E[y_t | H^{t-1}] = E[w_t | H^{t-1}]$ in which expected wages are paid according to the expected marginal product of the worker. Substituting y_t and w_t

into these expressions yields

$$w_t = (f - b_t g) \cdot \hat{a}_t + (1 - b_t) m_t + b_t p_t,$$

so that wages in period t are a weighted average of the market's perception of the agent's ability based on his past performance (m_t) and on the measure of his current period performance (p_t) plus an additional term which is constant with respect to the agent's effort choice in each period. Of these wages, the agent receives $(1 - \tau_t) w_t$.

To solve this problem, we work backwards, as in Gibbons-Murphy. At time $t = 2$, taking b_2 as given, the agent chooses an effort vector

$$\max_{a_2} (1 - \tau_2) E_2 [w_2 | a_1, p_1] - c(a_2),$$

which yields

$$a_{\ell 2}^* = (1 - \tau_2) b_2 g_\ell \text{ for } \ell = 1, \dots, N.$$

In particular, in the second period, absent explicit incentives based on the performance measure (i.e. if $b_2 = 0$), the agent would put in no effort. At the beginning of $t = 2$, principals compete to offer (s_2, b_2) to the agent, which ensures that (s_2, b_2) will be chosen to maximize the agent's equilibrium expected utility, or

$$\max_{b_2} (1 - \tau_2) (m_2 + f \cdot a_2^*(b_2)) - c(a_2^*).$$

Some simple computation shows that $b_2^* = \frac{\sum_{i=1}^N f_i g_i}{\sum_{i=1}^N g_i^2} = \rho_{fg}$, where ρ_{fg} is the correlation coefficient between the vectors f and g . Note that b_2^* is

independent of the tax rate.

Next, from the perspective of the first period, taking $b_2^*, a_2^*(b_2)$, and b_1 as given, the agent chooses effort to

$$\max_{a_1} E_1 [(1 - \tau_1) w_1 - c(a_1) + \delta ((1 - \tau_2) w_2 - c(a_2^*(b_2)))]$$

which yields

$$\begin{aligned} a_{\ell 1}^* &= (1 - \tau_1) \left(b_1 + \delta \frac{1 - \tau_2}{1 - \tau_1} (1 - b_2) \varphi \right) g_\ell \\ &= (1 - \tau_1) B_1 g_\ell \text{ for } \ell = 1, \dots, N \end{aligned}$$

The first term in B_1 is the standard pay-for-performance component, and the second captures the career-concerns component of incentives, which incentivizes a worker to exert effort in order to increase the market's perception of his ability and hence his second-period expected wages.

Finally, principals compete at the beginning of $t = 1$ to offer an equilibrium expected utility-maximizing contract (s_1, b_1) , which solves

$$\max_{b_1} E_1 [(1 - \tau_1) w_1^* - c(a_1^*) + \delta ((1 - \tau_2) w_2^* - c(a_2^*))].$$

The solution to this problem is for the principal to choose explicit incentives in order to keep total incentives constant across periods, because effort costs are convex. That is, $B_1^* = \rho_{fg}$ and hence

$$b_1^* = B_1^* - \delta \frac{1 - E_1[\tau_2 | \tau_1]}{1 - \tau_1} (1 - b_2^*) \varphi,$$

where $E_1[\tau_2 | \tau_1]$ is the agent's beliefs about the future tax rate from the

perspective of the first period. Note that if he believes that tax rates follow a martingale, then b_1^* is independent of both the first- and second-period tax rates. This leads to the following proposition.

PROPOSITION 8. In this model, if the agent believes that tax rates follow a martingale, then both b_1^* and b_2^* are independent of tax rates.

This proposition helps us understand why we might not expect endogenous performance pay to vary with the tax rates, a finding confirmed in much of the literature on CEO compensation. (See, for example, Frydman and Molloy (2011).)

2.3 Welfare

In order to compute the welfare loss of taxation, we follow Chetty (2009). The social surplus from the perspective of period 1 is the expected discounted net utility plus tax revenues, assuming these tax revenues will be redistributed back to the agent in a lump-sum fashion. That is,

$$W_1 = E_1 \left[\sum_{t=1}^2 \delta^{t-1} ((1 - \tau_t) w_t^* - c(a_t^*)) + \sum_{t=1}^2 \delta^{t-1} \tau_t w_t^* \right].$$

Since a_t^* is chosen optimally, the envelope theorem gives us

$$\frac{dW_1}{d\tau_t} = \delta^{t-1} \tau_t \frac{\partial E_1 [w_t^*]}{\partial \tau_t} = -\delta^{t-1} \tau_t \rho_{fg}^2.$$

That is, the marginal increase in deadweight loss as a result of an increase in taxes is proportional to the correlation between the f and g vectors. To the extent that we believe explicit incentives are chosen optimally and competitively, then the dynamic inefficiencies resulting from an increase in the

tax rate depend critically on the quality of the contracting environment. The loss or gain in efficiency from a discrete change in the proposed tax system from $\tau' = (\tau'_1, \tau'_2)$ to $\tau'' = (\tau''_1, \tau''_2)$ is then $W(\tau'') - W(\tau') = -\rho_{fg}^2 \int_{\tau'_1}^{\tau''_1} \tau d\tau - \delta \rho_{fg}^2 \int_{\tau'_2}^{\tau''_2} \tau d\tau$

In order to quantify this, one would need a measure of the quality of the contracting environment. While it may be possible to compute this directly from performance reports and accounting records of individual firms, it may be possible to get at this number using less direct means. Suppose there are several agents who are all drawn from the same commonly known distribution of ability $N(m_1, h_1^{-1})$ and who all begin working at the same time. It turns out that ρ_{fg} can be backed out from the growth rate in income variance between the two periods. In the appendix, we show that if we let $I_1 = \text{Var}(w_1^*)$ and $I_2 = \text{Var}(w_2^*)$, then

$$\rho_{fg}^2 = 1 - \frac{2h_1^2 I_1 (I_2 - I_1)}{1 - \sqrt{1 - 4h_1^2 I_1 (I_2 - I_1)}}$$

$$\varphi = \frac{1}{2h_1 I_1} - \frac{1}{2h_1 I_1} \sqrt{1 - 4h_1^2 I_1 (I_2 - I_1)}.$$

Thus, given knowledge of the prior distribution h_1 and measures of income inequality within a given cohort in the experience cycle over two periods, we can back out the structural parameters of interest. In a three-period model, inequality in the three periods can be used to also back out h_1 . In a model with more than three periods, the structural parameters are over-identified.

The intuition here is the following.

3 Pure Career Concerns

3.1 Setup

In the previous section we derived expressions for the deadweight loss of taxation under the case in which the strength of explicit incentives is determined endogenously. By contrast, in this section the strength of incentives are exogenously fixed to zero. This case is the pure career-concerns model analyzed by Holmström (1999).

There is a single agent and a continuum of principals. Time is indexed by $t = 1, \dots, T$, and in each period, the agent chooses an effort level $a_t \in \mathbb{R}_+^N$, where $N \geq 2$, at cost $c(a_t) = \frac{1}{2}a_t \cdot a_t$. Effort generates output for the principal for whom the agent works at period t according to

$$y_t = f \cdot a_t + \eta + \varepsilon_t,$$

where η is the agent's innate ability, $f \in \mathbb{R}_+^N$ is a vector of weights, and ε_t is an error term. The agent and all the principals are uninformed about the agent's ability. Throughout, we will assume $\eta \sim N(m_1, h_1^{-1})$, where h_1 is the precision (ex ante uncertainty) of the distribution of ability in the population, and $\varepsilon_t \sim N(0, h_\varepsilon^{-1})$, where h_ε is the precision of the error term. Throughout, we will assume that output is commonly observed by the agent as well as all principals.

The timing of the game is follows. In each period t , there are four stages. In the first stage, each principal offers the agent a wage w_t^P . After observing the offers from each principal, the agent chooses which to accept, and receives after-tax wage $w_t^A = (1 - \tau_t) w_t^P$. If the agent has ac-

cepted, he chooses an effort vector. Output is then publicly observed. Define $H^t = (y_1, \dots, y_{t-1})$ to be the public history of output.

Given a sequence of wage payments $\{w_t^A\}_{t=1}^T$ and a sequence of effort choices $\{a_t\}_{t=1}^T$, the agent's preferences are given by

$$U = \sum_{t=1}^T \delta^{t-1} \left((1 - \tau_t) w_t^A - c(a_t) \right),$$

and the principal that employs the agent in period t receives expected profits

$$\pi_t = E[y_t | \mu_t] - w_t^P,$$

where μ_t is the (public) belief about the agent's ability, which depends on the history of realized output prior to t . If $\tau_t = 0$ for each t and $N = 1$, this model would be identical to Holmstrom's career concerns model.

3.2 Equilibrium

DEFINITION 9. A Perfect Bayesian Equilibrium with competition is a sequence of effort choices $\{a_t^*\}_{t=1}^T$, a sequence of public beliefs $\mu_t^*(H^t)$, a sequence of wage functions $w_t^{P*}(\mu_t)$ and $w_t^{A*}(\mu_t)$ such that

1. Given μ_t^* and a_t^* , wages are determined by the zero-profit condition
2. Given $w_t^{P*}(\mu_t(H^t))$, the agent optimally chooses $\{a_t^*\}_{t=1}^T$
3. $\mu_t(H^t)$ is determined by Bayes's Rule

It is important to note that the optimal sequence of effort choices will be history-independent in this model, because output is additively separable in ability and effort. This in turn implies that the marginal returns to

effort are independent of ability, and therefore the agent does not have an additional direct motive for signalling about past effort through current effort choices. We now solve for the equilibrium. Given a history H^t , by the normal updating rule, the public beliefs about the agent's type are given by

$$\eta | H^t \sim N(m_t, h_t^{-1}),$$

where

$$\begin{aligned} m_t &= \frac{h_1}{h_t} m_1 + \frac{h_\varepsilon}{h_t} \sum_{s=1}^{t-1} (y_s - f \cdot a_s^*) \\ h_t &= h_1 + (t-1) h_\varepsilon. \end{aligned}$$

The wages are given by $w_t^{P*} = E[y_t | \mu_t^*, a_t^*] = f \cdot a_t^* + m_t$ and thus at t , the agent's problem is to

$$\max_{a_t} \sum_{s=t}^T \delta^{s-t} (1 - \tau_s) E[w_s^P | H^t] - c(a_t),$$

which gives us, for $i = 1, \dots, N$,

$$\begin{aligned} a_{it}^* &= \left(\sum_{s=t+1}^T \delta^{s-t} (1 - \tau_s) \frac{\partial E[w_s^P | H^t]}{\partial a_{it}} \right) \\ &= \left(\sum_{s=t+1}^T \delta^{s-t} \frac{1 - \tau_s h_\varepsilon}{1 - \tau_t h_s} \right) f_i \equiv B_t f_i. \end{aligned}$$

We can think of B_t as capturing the total incentive strength in period t that is derived from career concerns.

3.3 Welfare

Define ex ante expected equilibrium discounted welfare as

$$W_1 = E_1 \left[\sum_{t=1}^T \delta^{t-1} \left((1 - \tau_t) w_t^{P*} - c(a_t^*) \right) + \sum_{t=1}^T \delta^{t-1} \tau_t w_t^{P*} \right].$$

Fix $\tau_t = \tau$ and consider an increase in τ_k . By the envelope theorem, we have

$$\frac{dW_1}{d\tau_k} = \sum_{t=1}^T \tau_t \delta^{t-1} \frac{\partial E_1 [w_t^{P*}]}{\partial \tau_k}.$$

It is easy to show that

$$\delta^{-(k-1)} \frac{dW_1}{d\tau_k} = -\tau \frac{h_\varepsilon}{h_k} (k-1),$$

where $\varphi = \frac{h_\varepsilon}{h_\varepsilon + h_1}$ is the first-period signal-to-noise ratio. Welfare losses due to taxation are proportional to the tax level τ . They are increasing in φ and k .

A sufficient statistic for computing the welfare loss of taxation in this model, then, is φ . This is a theoretical object that does not necessarily have an observable real-world counterpart. However, given panel data on wages, and viewing the world through the lens of this model, we can back out φ as a function of the growth-rate of wage dispersion for a given cohort (conditional on observables). That is, since

$$I_t = \text{Var} \left(w_t^{P*} \right) = \text{Var} (m_t) = \frac{1}{h_1} - \frac{1}{h_t},$$

given data on the pre-tax wage distribution for at least three periods, a researcher can back out h_1 and h_ε . Since $h_k = h_1 + (k-1)h_\varepsilon$, then, h_1 and h_ε

are sufficient statistics for estimating the welfare losses of increasing taxes k periods from now.

4 Conclusion

This paper informs this debate by analyzing the welfare losses of taxation in a simple dynamic moral hazard model under symmetric information. We analyze how the principal-agent relationship is affected by taxation, and the consequent welfare implications. We show that the strength of incentives is central to understand the welfare costs of taxation. When explicit incentives are chosen optimally in a compensation contract, then the deadweight loss of taxation is proportional to the degree of high-powered incentives. That is, the welfare cost of taxation will be higher when the agent is incentivized to exert effort using very high-powered incentives. This is because the elasticity of taxable income is proportional to the power of incentives. That is, when the power of incentives is high, an increase in tax rate leads to a relatively greater reduction in taxable income. Hence, the deadweight loss is higher when the power of incentives is higher.

Critics often argue, however, that the explicit incentives are chosen exogenously (or suboptimally). In particular, due to weak corporate governance, the compensation contracts may be such that executive compensation is not highly linked to performance (Bebchuk and Fried (2006)). In such cases, the relation between explicit incentives and deadweight loss need not hold. Instead the career concerns of executives become a key determinant of the deadweight loss of taxation. This is because under such contracts, the incentives to provide effort, in general, will arise from both the exogenously-

determined explicit incentives, and career concerns. Under career concerns the incentives are such that the executive's current effort will improve the market's perception of his talent, resulting in higher compensation in the future. Hence, under such a setting, the deadweight loss of a tax in period t will depend on the sensitivity of market's perception of the executive's talent in period t to his current effort.

References

- ACHARYA, V. V., P. M. DEMARZO, AND I. KREMER (2011): "Endogenous Information Flows and the Clustering of Announcements," *American Economic Review*, forthcoming.
- AMAR, B. (1993): "The hidden costs of stock market liquidity," *Journal of Financial Economics*, 34(1), 31 – 51.
- ASKER, J. W., J. FARRE-MENSA, AND A. LJUNGQVIST (2011): "Comparing the Investment Behavior of Public and Private Firms," *SSRN eLibrary*, forthcoming.
- BAKER, G. (2002): "Distortion and Risk in Optimal Incentive Contracts," *Journal of Human Resources*, 37(4), 728–751.
- BAKER, G. P. (1992): "Incentive Contracts and Performance Measurement," *Journal of Political Economy*, 100(3), 598–614.
- BARRO, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *The Quarterly Journal of Economics*, 121(3), 823–866.
- BEBCHUK, L. A., AND J. FRIED (2006): *Pay without performance: The unfulfilled promise of executive compensation*. Harvard University Press.
- BEKAERT, G., C. R. HARVEY, AND R. L. LUMSDAINE (2002): "The dynamics of emerging market equity flows," *Journal of International Money and Finance*, 21(3), 295–350.
- BERGSTRESSER, D., AND T. PHILIPPON (2006): "CEO incentives and earnings management," *Journal of Financial Economics*, 80(3), 511–529.

- BILLGER, S. M., AND K. F. HALLOCK (2005): "Mass Layoffs and CEO Turnover," *Industrial Relations*, Vol. 44, No. 3, pp. 463-489, July 2005.
- BRENNAN, M. J., AND H. H. CAO (1997): "International Portfolio Investment Flows," *Journal of Finance*, 52(5), 1851-80.
- BROWN, S. (2004): "Mass Layoff Statistics Data in the United States and Domestic and Overseas Relocation," *Bureau of Labor Statistics papers*.
- CAMPBELL, J. Y., AND L. M. VICEIRA (2002): *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, no. 9780198296942 in OUP Catalogue. Oxford University Press.
- CHETTY, R. (2009): "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance," *American Economic Journal: Economic Policy*, 1(2), 31-52.
- DAVIS, S. J., J. HALTIWANGER, AND R. JARMIN (2006): "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," *NBER Macroeconomics Annual*, 21, pp. 107-179.
- DELLAVIGNA, S., AND J. M. POLLET (2009): "Investor Inattention and Friday Earnings Announcements," *Journal of Finance*, 64(2), 709-749.
- DEMARZO, P., AND D. DUFFIE (1999): "A Liquidity-Based Model of Security Design," *Econometrica*, 67(1), 65-100.
- DEMARZO, P. M. (2005): "The Pooling and Tranching of Securities: A Model of Informed Intermediation," *Review of Financial Studies*, 18(1), 1-35.

- DEWATRIPONT, M., I. JEWITT, AND J. TIROLE (1999): "The Economics of Career Concerns, Part I: Comparing Information Structures," *Review of Economic Studies*, 66(1), 183–98.
- EDMANS, A., I. GOLDSTEIN, AND W. JIANG (2011): "The Real Effects of Financial Markets: The Impact of Prices on Takeovers," *Journal of Finance*.
- FANG, L., AND J. PERESS (2009): "Media Coverage and the Cross-section of Stock Returns," *The Journal of Finance*, 64(5), 2023–2052.
- FARBER, H. S., AND K. F. HALLOCK (2009): "The changing relationship between job loss announcements and stock prices: 1970 to 1999," *Labour Economics*, 16(1), 1 – 11.
- FORBES, K. J., AND F. E. WARNOCK (2011): "Capital Flow Waves: Surges, Stops, Flight, and Retrenchment," *NBER Working Papers*.
- FRENCH, K. R., AND J. M. POTERBA (1991): "Investor Diversification and International Equity Markets," *American Economic Review*, 81(2), 222–26.
- FROOT, K. A., P. G. J. O'CONNELL, AND M. S. SEASHOLES (2001): "The portfolio flows of international investors," *Journal of Financial Economics*, 59(2), 151–193.
- FROOT, K. A., AND T. RAMADORAI (2008): "Institutional Portfolio Flows and International Investments," *Review of Financial Studies*, 21(2), 937–971.
- FROOT, K. A., AND M. TEO (2004): "Equity Style Returns and Institutional Investor Flows," NBER Working Papers 10355, National Bureau of Economic Research, Inc.

- FRYDMAN, C., AND R. S. MOLLOY (2011): "Does Tax Policy Affect Executive Compensation? Evidence from Postwar Tax Reforms," NBER Working Papers 16812, National Bureau of Economic Research, Inc.
- GABAIX, X. (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79(3), 733–772.
- GEWEKE, J. (2001): "A note on some limitations of CRRA utility," *Economics Letters*, 71(3), 341–345.
- GIBBONS, R., AND L. F. KATZ (1991): "Layoffs and Lemons," *Journal of Labor Economics*, 9(4), 351–80.
- GIBBONS, R., AND K. J. MURPHY (1992): "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy*, 100(3), 468–505.
- GLAZ, J., J. NAUS, AND S. WALLENSTEIN (2001): *Scan statistics*, Springer series in statistics. Springer.
- GORTON, G., AND G. PENNACCHI (1990): "Financial Intermediaries and Liquidity Creation," *Journal of Finance*, 45(1), 49–71.
- GRINBLATT, M., AND M. KELOHARJU (2000): "The investment behavior and performance of various investor types: a study of Finland's unique data set," *Journal of Financial Economics*, 55(1), 43–67.
- GRINBLATT, M., S. TITMAN, AND R. WERMERS (1995): "Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior," *American Economic Review*, 85(5), 1088–1105.

- HALL, B. J., AND J. B. LIEBMAN (2000): "The Taxation of Executive Compensation," in *Tax Policy and the Economy, Volume 14*, NBER Chapters, pp. 1–44. National Bureau of Economic Research, Inc.
- HALLOCK, K. F. (2009): "Job Loss and the Fraying of the Implicit Employment Contract," *Journal of Economic Perspectives*, 23(4), 69–93.
- HAU, H., AND H. REY (2008): "Home Bias at the Fund Level," *American Economic Review*, 98(2), 333–38.
- HOLMSTRÖM, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–182.
- HOLMSTROM, B. (1979): "Moral Hazard and Observability," *Bell Journal of Economics*, 10(1), 74–91.
- HOLMSTROM, B., AND P. MILGROM (1991): "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics and Organization*, 7(0), 24–52.
- HOLTE, M. (2007): "Taxation, Career Concerns and CEO Pay," *Working Paper*.
- IMBENS, G. W., AND J. M. WOOLDRIDGE (2009): "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature*, 47(1), 5–86.
- IMF (2010): "Global Financial Stability Report, Chapter IV "Global Liquidity Expansion: Effects on Receiving Economies and Policy Response Options"," *World Economic and Financial Surveys*.

- (2011a): “Global Financial Stability Report, Chapter II Long-Term Investors and Their Asset Allocation: Where are they now?,” *World Economic and Financial Surveys*.
- (2011b): “International Capital Flows: Reliable Or Fickle?” *World Economic Outlook*,” *World Economic and Financial Surveys*.
- (2011c): “Recent Experiences in Managing Capital Inflows Cross-Cutting Themes and Possible Policy Framework,” *World Economic and Financial Surveys*.
- (2011d): “Assessing Reserve Adequacy,” *IMF Policy Papers*.
- JENSEN, M. C., AND W. H. MECKLING (1976): “Theory of the firm: Managerial behavior, agency costs and ownership structure,” *Journal of Financial Economics*, 3(4), 305 – 360.
- KATUSCAK, P. (2005): “The impact of personal income taxation on executive compensation,” *The Proceedings from the 98th Annual Conference of the National Tax Association*, 55.
- LIM, E.-G. (2007): “Do Reserve Portfolios Respond to Exchange Rate Changes Using a Portfolio Rebalancing Strategy? An Econometric Study Using COFER Data,” IMF Working Papers 07/293, International Monetary Fund.
- PIHLMAN, J., AND H. VAN DER HOORN (2010): “Procyclicality in Central Bank Reserve Management: Evidence from the Crisis,” *SSRN eLibrary*.
- RAJAN, R. G. (1994): “Why Bank Credit Policies Fluctuate: A Theory and Some Evidence,” *The Quarterly Journal of Economics*, 109(2), 399–441.

- STEIN, J. C. (1988): "Takeover Threats and Managerial Myopia," *Journal of Political Economy*, 96(1), 61–80.
- (1989): "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *The Quarterly Journal of Economics*, 104(4), 655–69.
- STEIN, J. C. (2010): "Monetary Policy as Financial-Stability Regulation," *NBER Working Paper Series*, (16883).
- TSE, S., AND J. TUCKER (2010): "Within-industry timing of earnings warnings: do managers herd?," *Review of Accounting Studies*, 15, 879–914.
- WEITZMAN, M. L. (2007): "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, 97(4), 1102–1130.

Appendix

Appendix A.1: Updating rule for the *terminate* policy

Given the market's conjecture that the firm has adopted the *efficient* policy, $\hat{D} = 0$, the conditional distribution of η given that the project failed is given by Bayes' Rule:

$$\begin{aligned}
 f(\eta|\text{Proj. Fails}) &= \frac{f(\eta) \Pr(\text{Proj. Fails}|\eta)}{\int_{\hat{\eta}} f(\hat{\eta}) \Pr(\text{Proj. Fails}|\hat{\eta}) d\hat{\eta}} \\
 &= \frac{f(\eta) [\pi(1 - \eta(1 - \delta)) + (1 - \pi)(1 - \eta)]}{\pi(1 - \mu(1 - \delta)) + (1 - \pi)(1 - \mu)} \\
 &= \frac{f(\eta) [1 - \eta(1 - \pi\delta)]}{1 - \mu(1 - \pi\delta)}
 \end{aligned}$$

Therefore, we can calculate the conditional expectation of η as:

$$\begin{aligned}
 E_{mkt}[\eta|\text{Proj. Fails}] &= \int_{\eta} \eta f(\eta|\text{Proj. Fails}) d\eta \\
 &= \frac{1}{1 - \mu(1 - \pi\delta)} \int_{\eta} \eta [1 - \eta(1 - \pi\delta)] f(\eta) d\eta \\
 &= \frac{\mu - E[\eta^2](1 - \pi\delta)}{1 - \mu(1 - \pi\delta)} \\
 &= \frac{\mu - (\mu^2 + \sigma_{\eta}^2)(1 - \pi\delta)}{1 - \mu(1 - \pi\delta)} \\
 &= \mu - \frac{\sigma_{\eta}^2(1 - \pi\delta)}{1 - \mu(1 - \pi\delta)}
 \end{aligned}$$

Similarly, the conditional distribution of η given that the project suc-

ceeded is given by the Bayes Rule:

$$\begin{aligned}
f(\eta|\text{Proj. Succeeds}) &= \frac{f(\eta) \Pr(\text{Proj. Succeeds}|\eta)}{\int_{\hat{\eta}} f(\hat{\eta}) \Pr(\text{Proj. Succeeds}|\hat{\eta}) d\hat{\eta}} \\
&= \frac{f(\eta) [\pi\eta(1-\delta) + (1-\pi)\eta]}{\pi\mu(1-\delta) + (1-\pi)\mu} \\
&= \frac{f(\eta) [\eta(1-\pi\delta)]}{\mu(1-\pi\delta)} \\
&= f(\eta) \frac{\eta}{\mu}
\end{aligned}$$

Therefore, conditional expectation is

$$\begin{aligned}
E_{mkt} [\eta|\text{Proj. Succeeds}] &= \int_{\eta} \eta \cdot f(\eta|\text{Proj. Succeeds}) d\eta \\
&= \int_{\eta} \frac{\eta^2}{\mu} f(\eta) d\eta \\
&= \frac{E[\eta^2]}{\mu} \\
&= \frac{\sigma_{\eta}^2 + \mu^2}{\mu} \\
&= \mu + \frac{\sigma_{\eta}^2}{\mu}
\end{aligned}$$

Appendix A.2: Updating rule for the *delay* policy

Given the market's conjecture that the firm has adopted the defensive policy, $\hat{D} = 1$, the outcome of observing layoffs would be off the equilibrium path and Bayes' Rule would not apply. However, as specified in the main text, introducing trembles to the model allows for a positive probability of observing layoffs. Further, because the manager only has a decision point when the project fails, the observation of layoffs in this setting is

equivalent to observing project failure. Thus, the application of Bayes' Rule is identical to the case of the *terminate* policy:

$$\begin{aligned} f(\eta|\text{Proj. Fails}) &= \frac{f(\eta) \Pr(\text{Proj. Fails}|\eta)}{\int_{\hat{\eta}} f(\hat{\eta}) \Pr(\text{Proj. Fails}|\hat{\eta}) d\hat{\eta}} \\ &= \frac{f(\eta) [1 - \eta(1 - \pi\delta)]}{1 - \mu(1 - \pi\delta)} \end{aligned}$$

Again, we calculate the conditional expectation of η as:

$$\begin{aligned} E_{mkt}[\eta|\text{Proj. Fails}] &= \int_{\eta} \eta f(\eta|\text{Proj. Fails}) d\eta \\ &= \mu - \frac{\sigma_{\eta}^2(1 - \pi\delta)}{1 - \mu(1 - \pi\delta)} \end{aligned}$$

For the case of no layoffs, because the market expects firms to never announce a layoff when $\hat{D} = 1$, the observation of no layoffs should not lead to any updating. This is exactly what we find:

$$\begin{aligned} f(\eta|\text{No Layoffs}, \hat{D} = 1) &= \frac{f(\eta) \Pr(\text{No Layoffs}|\eta)}{\int_{\hat{\eta}} f(\hat{\eta}) \Pr(\text{No Layoffs}|\hat{\eta}) d\hat{\eta}} \\ &= \frac{f(\eta) [\pi(\eta(1 - \delta) + (1 - \eta(1 - \delta))) + (1 - \pi)(\eta + (1 - \eta))]}{\pi(\mu(1 - \delta) + (1 - \mu(1 - \delta))) + (1 - \pi)(\mu + (1 - \mu))} \\ &= f(\eta) \end{aligned}$$

By inspection, the posterior distribution is identical to the prior, so the market's expectation of talent remains at μ .

Appendix A.3: Updating rule for a mixed-strategy policy

For the case where the market observes layoffs, the signal is once again equivalent to observing project failure. Thus, the calculation proceeds as in the previous sections, and yields the posterior expectation:

$$\begin{aligned} E_{mkt} [\eta | \text{Proj. Fails}] &= \int_{\eta} \eta f(\eta | \text{Proj. Fails}) d\eta \\ &= \mu - \frac{\sigma_{\eta}^2 (1 - \pi\delta)}{1 - \mu(1 - \pi\delta)} \end{aligned}$$

For the case of no layoffs, the market expects firms to announce “no layoffs” with probability \hat{D} conditional on project failure, and with probability 1 conditional on project success. The updating rule proceeds as follows:

$$\begin{aligned} f(\eta | \text{No Layoffs}, \hat{D}) &= \frac{f(\eta) \Pr(\text{No Layoffs} | \eta, \hat{D})}{\int_{\hat{\eta}} f(\hat{\eta}) \Pr(\text{No Layoffs} | \hat{\eta}, \hat{D}) d\hat{\eta}} \\ &= \frac{f(\eta) [\pi(\eta(1 - \delta) + \hat{D}(1 - \eta(1 - \delta))) + (1 - \pi)(\eta + \hat{D}(1 - \eta))]}{\pi(\mu(1 - \delta) + \hat{D}(1 - \mu(1 - \delta))) + (1 - \pi)(\mu + \hat{D}(1 - \mu))} \\ &= \frac{f(\eta) [\eta(1 - \pi\delta) + \hat{D}(1 - \eta(1 - \pi\delta))]}{\mu(1 - \pi\delta) + \hat{D}(1 - \mu(1 - \pi\delta))} \end{aligned}$$

Therefore, conditional expectation is

$$\begin{aligned}
E_{mkt} [\eta | \text{No Layoffs}, \hat{D}] &= \int_{\eta} \eta f(\eta | \text{No Layoffs}, \hat{D}) d\eta \\
&= \frac{\int_{\eta} \eta [\eta(1 - \pi\delta) + \hat{D}(1 - \eta(1 - \pi\delta))] f(\eta) d\eta}{\mu(1 - \pi\delta) + \hat{D}(1 - \mu(1 - \pi\delta))} \\
&= \frac{E[\eta^2](1 - \pi\delta) + \hat{D}(\mu - E[\eta^2])(1 - \pi\delta)}{\mu(1 - \pi\delta) + \hat{D}(1 - \mu(1 - \pi\delta))} \\
&= \frac{(\mu^2 + \sigma_{\eta}^2)(1 - \pi\delta) + \hat{D}(\mu - (\mu^2 + \sigma_{\eta}^2)(1 - \pi\delta))}{\mu(1 - \pi\delta) + \hat{D}(1 - \mu(1 - \pi\delta))} \\
&= \mu + \frac{\sigma_{\eta}^2(1 - \pi\delta) + \hat{D}(-\sigma_{\eta}^2)(1 - \pi\delta)}{\mu(1 - \pi\delta) + \hat{D}(1 - \mu(1 - \pi\delta))} \\
&= \mu + \frac{(1 - \hat{D})(1 - \pi\delta)\sigma_{\eta}^2}{\hat{D} + (1 - \hat{D})(1 - \pi\delta)\mu}
\end{aligned}$$

Appendix A.4: Proof for proposition 1

Taking the total derivative of equation 12 with respect to both \hat{D} and π , we find:

$$\frac{\partial \hat{D}}{\partial \pi} \left[(1 - \mu(1 - \pi\delta)) \frac{C}{\gamma \sigma_e t a^2} \right] = -\delta \left[\frac{1}{[1 - \mu(1 - \pi\delta)]^2} - \mu(1 - \hat{D}) \frac{C}{\gamma \sigma_e t a^2} \right]$$

Next, we know that the probability of observing a layoff is $(1 - \hat{D})(1 - \eta(1 - \pi\delta))$, or $(1 - \hat{D})(1 - \mu(1 - \pi\delta))$ in expectation. Taking the total derivative with respect to π yields:

$$\frac{\partial \Pr[\text{Layoffs}]}{\partial \pi} = -\frac{\partial \hat{D}}{\partial \pi} (1 - \mu(1 - \pi\delta)) + (1 - \hat{D})\mu\delta$$

Combining the two results above and simplifying, we find:

$$\frac{\partial \Pr[\text{Layoffs}]}{\partial \pi} = \frac{\delta \gamma \sigma_{\eta}^2}{C [1 - \mu(1 - \pi\delta)]^2} > 0$$

Appendix A.5: Proof of proposition 2

We can rewrite $\Delta \Pr [\text{Layoffs}] / \Delta (\text{Large Firm Layoff})$ as

$$\frac{\Delta \Pr [\text{Layoffs}]}{\Delta (\text{Large Firm Layoff})} = \frac{\Delta \Pr [\text{Layoffs}]}{\Delta (\pi)} \frac{\Delta \pi}{\Delta (\text{Large Firm Layoff})}$$

From Proposition 1 we know $\frac{\Delta \Pr[\text{Layoffs}]}{\Delta(\pi)} > 0$. Also from equation (13) we know that

$\Delta \pi / \Delta (\text{Large Firm Layoff}) > 0$. These two inequalities imply

$$\frac{\Delta \Pr [\text{Layoffs}]}{\Delta (\text{Large Firm Layoff})} > 0.$$

Appendix A.6: Proof of corollaries 3-5.

Following from appendix A.5, we can write the following derivative as:

$$\partial \left[\frac{\Delta \Pr [\text{Layoffs}]}{\Delta (\text{Macro Event})} \right] / \partial \left[\frac{\Delta \pi}{\Delta (\text{Macro Event})} \right] = \frac{\Delta \Pr [\text{Layoffs}]}{\Delta (\pi)}$$

From appendix 4, we know the right hand side is positive. Therefore,

$$\partial \left[\frac{\Delta \Pr [\text{Layoffs}]}{\Delta (\text{Macro Event})} \right] / \partial \left[\frac{\Delta \pi}{\Delta (\text{Macro Event})} \right] > 0$$

Similarly from appendix 4, we know

$\partial \left[\frac{\Delta \text{Pr}[\text{Layoffs}]}{\Delta(\pi)} \right] / \partial \sigma_\eta^2 > 0$ and $\partial \left[\frac{\Delta \text{Pr}[\text{Layoffs}]}{\Delta(\pi)} \right] / \partial \gamma > 0$. Therefore we get:

$$\begin{aligned}
& \frac{\partial}{\partial \sigma_\eta^2} \left[\frac{\Delta \text{Pr}[\text{Layoffs}]}{\Delta(\text{Large Firm Layoff})} \right] \\
&= \frac{\Delta \pi}{\Delta(\text{Large Firm Layoff})} \partial \left[\frac{\Delta \text{Pr}[\text{Layoffs}]}{\Delta(\pi)} \right] / \partial \sigma_\eta^2 > 0 \\
& \frac{\partial}{\partial \gamma} \left[\frac{\Delta \text{Pr}[\text{Layoffs}]}{\Delta(\text{Large Firm Layoff})} \right] \\
&= \frac{\Delta \pi}{\Delta(\text{Large Firm Layoff})} \partial \left[\frac{\Delta \text{Pr}[\text{Layoffs}]}{\Delta(\pi)} \right] / \partial \gamma > 0
\end{aligned}$$

Appendix A.7: Dynamic Version of the Model with Exogenous Updating Rule

In this appendix we present a simple dynamic extension of our model based on McCall's (1970) model of intertemporal job search. At date 0 there is a continuum of managers who evaluate their investment opportunities and hire a worker to engage in production. Thereafter in each period t , the project has a positive chance of failure, which depends positively on both the manager's quality given by η_i , and the time-varying aggregate state, given by $\lambda_{s,t}$. The states can be normal (N) or adverse (A), such that such that $\lambda_{N,t} = 1$, and $\lambda_{A,t} = \delta$. After observing the realization of the project on each date, each firm has to decide whether to continue or terminate production. Given the parameters of this model, it will always be optimal for firms to continue if their projects do not fail. Therefore, the key decision has to be made after the project fails.

The main cost of announcing a layoff is that the market's belief about

the manager type is adversely affected when they observe a layoff. This is what we call *reputation cost*. In each period nature draws a reputation cost $z(\pi)$ for the manager. For simplicity, we take this reputation cost to be exogenously given. The market's belief that the aggregate is in a normal state at time t is given by π_t . The key assumption about the reputation cost is that $z'(\pi) > 0$, i.e. when the market thinks the aggregate state is more likely to be normal, then their posterior beliefs about manager talent will be more pessimistic after they observe a layoff. In each period there is a random draw of π , which maps 1-to-1 into a reputation cost. Let the distribution of reputation cost be given by $F(Z) = \Pr[z \leq Z]$, with $F(0) = 0$, $F(B) = 1$ for $B < \infty$. The manager has the option of not engaging in a layoff, in which case he pays c in this period and waits until next period for another draw of reputation cost from F . The per-period cost c is the net loss the firm bears every period by keeping the worker at a failed project for an additional period.

Let y_t be the manager's payoff in period t . To be consistent with the pre-existing search models it is convenient to characterize the payoff as $y_t = -\gamma z(\pi)$ if the manager with a failed project decides to layoff when the reputation cost is $z(\pi)$, and $y_t = -c$ when the manager decides to delay layoff. Here γ again measures the degree of reputational concern a manager has. The managers devise a strategy to maximize $E \sum_{t=0}^{\infty} \beta^t y_t$, such that $0 < \beta < 1$ is the discount factor.

Let $v(z(\pi))$ be the expected value of $\sum_{t=0}^{\infty} \beta^t y_t$ for an optimally-behaving manager who faces a reputation cost of $z(\pi)$, and is deciding whether to layoff or not. In this model we assume no recall. The value function $v(z(\pi))$

satisfies the Bellman equation

$$v(z(\pi)) = \max \left\{ -\frac{\gamma z(\pi)}{1-\beta}, -c + \beta \int v(z(\pi')) dF(\pi') \right\}$$

There exists a threshold reputation $z(\bar{\pi})$, such that if the manager is facing a reputation cost $z(\pi) \leq z(\bar{\pi})$, he should layoff, and delay otherwise. Solving for the threshold reputation, we can characterize his strategy as

$$\gamma z(\bar{\pi}) - c = \frac{\beta\gamma}{1-\beta} \int_0^{\bar{\pi}} [z(\pi) - z(\bar{\pi})] dF(\pi')$$

Further rewriting and by applying integration by parts we can characterize the reputation cost threshold as

$$\gamma z(\bar{\pi}) - c = \beta(\gamma E[z(\pi)] - c) + \beta\gamma \int_{\bar{\pi}}^{\infty} F(\pi') d\pi'$$

Thus, in this economy, managers whose project fails, will delay announcing layoffs until he faces a sufficiently high market-wide belief of being in an adverse aggregate state (i.e. a low value of π). This effectively generates periods of no or little layoffs, and large number of layoffs when market's belief about the aggregate state is adverse with high likelihood. In this model the firms need not all layoff in the same period. Their decision rule will depend on their degree of reputational concerns, γ . For a large class of functional forms for $z(\cdot)$ it can be shown that the threshold $\bar{\pi}$ is a decreasing function of γ . This suggests that when managerial reputational concerns rises (i.e. high γ), their threshold for waiting becomes more restrictive (i.e. lower $\bar{\pi}$), as these managers are waiting to engage in layoffs in periods when the market's belief puts a very high probability on the

aggregate state being adverse. As an example consider the case in which F is a uniform distribution with support $[0, B]$. Additionally assume that $z(x) = \phi x$. Under this example, we can characterize the threshold as

$$\bar{\pi} = \frac{\beta \{B + \phi E[\pi]\}}{\phi + \beta} + \frac{(1 - \beta) c}{\gamma (\phi + \beta)}$$

It is clear that $\partial \bar{\pi} / \partial \gamma < 0$, $\partial \bar{\pi} / \partial c > 0$, and $\partial \bar{\pi} / \partial B > 0$, while the effect of ϕ is ambiguous. This suggests that a greater reputational concern makes the threshold more restrictive, while a larger cost c (which is a per-period loss made by the firm because of following the inefficient policy) and a higher variance of beliefs as measured by B , leads to a less restrictive threshold.