Essays on Labor Markets in Developing Countries

Dissertation Advisor: Professor Sendhil Mullainathan Supreet Kaur Anand

Essays on Labor Markets in Developing Countries

Abstract

This dissertation consists of three empirical essays on distortions in labor market outcomes in developing countries.

Chapter 1 tests for downward nominal wage rigidity in markets for casual daily agricultural labor. It examines responses to rainfall shocks in 500 Indian districts from 1956-2008. First, nominal wages rise in response to positive shocks but do not fall during droughts. Second, after transitory positive shocks have dissipated, nominal wages do not fall back down. Third, inflation moderates these effects. Fourth, rigidities lower employment: landless laborers experience a 6% reduction in employment in the year after positive shocks. Fifth, consistent with separation failures, rationing leads to increased labor supply to small farms. New survey evidence suggests that agricultural workers and employers view nominal wage cuts as unfair and believe that they reduce effort.

Chapter 2 (with Michael Kremer and Sendhil Mullainathan) describes the results of a field experiment that tests for self-control problems in labor supply. First, we find that workers will choose dominated contracts—which pay less for every output level but have a steeper slope—to motivate themselves. Second, effort increases significantly as workers' (randomly assigned) payday gets closer. Third, the demand for dominated contracts (and their benefits) is concentrated amongst those with the highest payday effects. Finally, as workers gain experience, they appear to learn about their self control problems: the correlation between the payday effect and the demand for the dominated contract grows with experience. These results together suggest that self-control, in this context at least, meaningfully alters the firm's contracting problem.

Chapter 3 empirically examines the impact of multiple market failures on allocative efficiency in farm production in poor countries. In years when labor rationing is more likely in villages (due to wage rigidity), there is a 63% increase in sharecropped and leased land by small farmers.

This is consistent with the prediction that distortions from a failure in one market can be reduced by reallocating other factors of production. In areas with worse credit access, there is less land adjustment in response to labor rationing. These results provide evidence for separation failures resulting from multiple missing markets.

Table of Contents

Acknowledgements

I owe my development as a researcher to the faculty who provided guidance and encouragement during my time at Harvard. I am deeply grateful to my primary advisor, Sendhil Mullainathan, for teaching me how to think about research and instilling confidence in my own creativity. I could not have asked for a more inspiring mentor. Michael Kremer taught me to not shy away from big questions, but rather to find them in every research topic. Rohini Pande was unfailingly generous in providing guidance on my work and career development. Lawrence Katz was a model for excellence in research and made me a better researcher through his efforts and advice.

Many other faculty and students have influenced my journey through graduate school. Asim Khwaja, Mark Rosenzweig, and Dani Rodrik introduced me to the power of theory for understanding development questions during the MPA/ID program. I am grateful to Asim and Dani for supporting my decision to pursue a PhD, and have valued my interactions with them immensely. I have also benefited greatly from conversations with Edward Glaeser and David Laibson. I thank my friends and classmates for support, instruction, laughter, and commiseration—especially Melissa Adelman and Heather Schofield, who have weathered many days in the Littauer basement with me.

I am indebted to my family for fueling my aspirations with their encouragement, expectations, and love. My wonderful husband, Narinder, has been more generous of spirit and staunch in his support than I could ever deserve. His boundless optimism has kept me afloat through every up and down. I am humbled by the sacrifices of my parents, Devinder and Surinder, who signed up for a life of struggle to enable a future for their children. They placed the importance of education above all else—it is those values that are responsible for this accomplishment. I am grateful to my brothers, Harpreet and Prabhjot, for raising me as a child to become the adult I am today. They encouraged me to select a career focused on poverty and gave me the capacity to aspire. I would be lost without them. Finally, I thank the newest members of my family—Davinder, Gurjit, Simran, Raminder, Sarabjot, and Mallika—for the love, warmth, and patience with which they have supported my efforts in graduate school.

Introduction

Labor markets are central to economic activity in both poor and rich countries. They allocate the use of labor in production. They enable individuals to trade what is usually their most abundant resource—their labor—for income. As such, efficiency in these markets is intrinsically linked to firm productivity and consumer welfare in any context. In addition, two features market structure in developing countries augment the importance of labor markets in that context. First, given that the poor often have worse access to smoothing mechanisms—credit, insurance, and savings—selling labor is a way in which the poor cope with shocks (Kochar 1999; Jayachandran 2006). Second, a distinguishing feature of economic organization in developing countries is that there is often little distinction between the household and the firm (Banerjee and Duflo 2006). This creates a direct channel through which changes in household welfare may impact firm production decisions (Benjamin 1992).

This dissertation empirically documents two sources of labor market distortions: wage rigidity and self-control problems. It examines the implications of these distortions for employment, output, and contract structure.

Chapter 1 documents downward nominal wage rigidity by examining wage and employment responses to rainfall shocks in 500 Indian districts over the past 50 years. Such rigidities lower employment, with the biggest employment reductions for the poorest village residents—those with little or no land. As a result, rigidities reduce total labor use in production. When households with small landholdings are rationed out of the external labor market, they increase labor supply to their own farms. This is consistent with separation failures and is expected to create a wedge in the returns to labor on small and large farms. An implication of this is that gains from trade can potentially be realized by redistributing land from large to small farms until the marginal product of labor is equalized across farms.

Chapter 3 provides evidence that this does indeed occur. In years where labor rationing is more likely, small farmers increase the amount of land they lease and sharecrop from large farmers. This is consistent with the idea that a failure in one factor market can be mitigated by reallocating other factors of production. However, land requires substantial capital outlays (Shaban 1987). The acquisition of land by poor households requires the presence of well-functioning credit markets. As a result, in areas with worse credit access, there is less land adjustment in response to labor rationing. Thus, when there are multiple market failures (i.e. labor and credit), it is not necessarily possible to achieve allocative efficiency.

Chapter 2 (joint with Michael Kremer and Sendhil Mullainathan) documents a way in which labor outcomes can be distorted even when factor markets function perfectly. Using a year-long field experiment with data entry workers in India, it shows that self-control problems reduce labor supply, earnings, and output. Workers are willing to pay to overcome these problems: they select dominated contracts—which pay less for every output level but have a steeper slope—in order to motivate themselves. This highlights a way in which firms can use incentives to increase worker welfare and output, even in the presence of individual production and perfect information. This motivates an additional view of the role of the firm.

Collectively, the above results have implications for how labor productivity evolves through the process of development. To the extent that growth is accompanied by improvements in financial markets, this reduces the consequences of labor market failures. For example, farmers can use credit to reallocate land, mitigating the distortions from labor rationing (Chapter 3). Access to financial services provides a (potentially less distortionary) way for households to cope with shocks. Thus, an inability to supply as much labor as they would like—due to rationing or self-control problems—has potentially lower welfare consequences for the poor.

In addition, industrialization and growth generally involve substantial changes in labor market arrangements themselves, which have bearing on labor productivity. Notably, workers move from self-employment (in farming or cottage industry production) to employment in larger firms such as factories. This decouples the consumption and production decisions, potentially weakening the feedback loop through which household shocks directly affect production decisions (Chapters 1 and 3). Employment in firms can also mitigate self-control problems through the provision of incentives, increasing worker productivity (Chapter 2). The subsequent chapters examine wage rigidity and self-control problems in more detail and discuss their implications for the study of poverty.

1 Nominal Wage Rigidity in Village Labor Markets

1.1 Introduction

This article empirically examines downward nominal wage rigidity and its employment consequences in a developing country context. Under such rigidities, wages are expected to exhibit three features. First, wages are rigid—they do not adjust fully to productivity shocks. Second, the rigidity is asymmetric—adjustment is hindered particularly in the downward direction. Third, the rigidity applies to nominal wage reductions—real wage cuts are not impeded if they occur through inflation. Put simply, wages resist falling from their current values in nominal terms. Such rigidities, if present, can deepen the impact of recessions and heighten employment volatility. As a result, they have been a focus of much debate in a broad literature on unemployment and business cycle dynamics.[1] In addition, they could help explain apparent labor market imperfections in poor countries, such as labor rationing and differences in labor allocation on small and large farms.[2]

Establishing the presence of downward nominal wage rigidity has posed an empirical challenge. A literature in OECD countries finds evidence based primarily on examining distributions of wage changes (e.g., McLaughlin 1994; Kahn 1997; Dickens et al. 2006). While this approach yields compelling documentation, it is vulnerable to measurement error and requires limiting analysis to workers employed by the same firm in consecutive years.[3] Importantly, it also does not allow for analysis of the employment effects of rigidities. A more direct test would involve examining how wages react to changes in the marginal revenue product of labor, but shifters of this are typically

[1] For overviews, see Tobin (1972); Greenwald and Stiglitz (1987); Blanchard (1990); Clarida, Galì, and Gertler (1999); Akerlof (2002); and Galì (2002).

[2] See Rosenzweig (1988) and Behrman (1999) for reviews of the debate on labor market imperfections in this context.

[3] Akerlof, Dickens, and Perry (1996) and Card and Hyslop (1997) provide excellent discussions of measurement challenges associated with the histogram approach.

difficult to isolate.[4]

To overcome this challenge, I exploit a feature of agricultural production in developing countries. In this setting, local rainfall variation generates transitory labor demand shocks. I focus on markets for casual daily agricultural labor—a major source of employment in poor countries. To test for rigidity, I examine market-level wage and employment responses to rainfall shocks in 500 Indian districts from 1956 to 2008.

Wage responses are consistent with downward rigidities. First, wage adjustment is asymmetric. While nominal wages rise robustly in response to positive shocks, they do not fall during droughts on average. Second, transitory positive shocks cause a persistent increase in wages. When a positive shock in one year is followed by a non-positive shock in the following year, wages do not adjust back down—they remain higher than they would have been in the absence of the lagged positive shock.

Third, particularly consistent with *nominal* rigidity, inflation moderates these wage distortions. When inflation is higher, droughts are more likely to result in lower real wages. In addition, transitory positive shocks are less likely to have persistent wage effects. For example, when inflation is above 6%, positive shocks have no impact on future wages. Since local rainfall is uncorrelated with inflation levels, these tests have a causal interpretation.[5] The findings support the hypothesis that inflation "greases the wheels" of the labor market (see, e.g., Tobin 1972; Akerlof, Dickens, and Perry 1996; Card and Hyslop 1997).

When nominal rigidities bind—keeping wages above market clearing levels—this distorts employment. For a given non-positive shock in the current year, a transitory positive shock in the previous year raises current wages without affecting current productivity. This causes a 3% average drop in total worker-days spent in agriculture. This magnitude is equivalent to the employment decrease during a drought. There is heterogeneity in these effects: workers with less land are considerably more likely to be rationed out of the labor market. This is because landed households exhaust their own labor supply on their farms before hiring outside labor, whereas those with little or no

[4] Holzer and Montgomery (1993) perform analysis in this spirit. They examine correlations of wage and employment growth with sales growth, which they assume reflects demand shifts. They find that wage changes are asymmetric and are small compared to employment changes. Card (1990) uses a different approach in the context of unionized Canadian firms. When nominal wages are explicitly indexed to expected inflation, real wages do not adjust to inflation surprises. As a result, these firms adjust employment down (up) when inflation surprises raise (lower) real wages.

[5] Rainfall shocks are local to small geographic areas and do not affect national price levels—a fact I verify empirically in Section 1.4.

land must supply to other farms at the prevailing wage. As a result, landless laborers experience a 6% average drop in employment. Overall, these employment dynamics are consistent with boom and bust cycles in village economies. They also match observations from other contexts that labor markets exhibit relatively large employment volatility and small wage variation.

When workers face rationing in external employment, they increase labor supply to their own farms. Specifically, households in the bottom tercile of landholdings supply 7% more labor to their own land in the year after a positive shock than if the positive shock had not occurred. This is consistent with the prediction that in the presence of labor market failures, a household's labor supply decision will not be separable from its decision of how much labor to use on its farm (see, e.g.: Singh, Squire, and Strauss 1986; Benjamin 1992; Udry 1996). The fact that smaller farms tend to use more labor per acre and have higher yields per acre than larger farms has been widely documented in the development literature (e.g. Bardhan 1973). These results support the hypothesis that this relationship can be attributed to separation failures, which lead the marginal product of labor to be lower on smaller farms than larger ones.

Could the above findings be explained by factors other than nominal wage rigidity? For example, if positive shocks have persistent productivity effects, this could explain why wages remain high in the following year. However, in this case, employment should not fall in the next year. Alternately, inter-temporal substitution in labor could explain why positive shocks increase future wages and lower future employment. However, this is not consistent with the inflation results—labor supply shifters should not be affected by inflation. In addition, small farms respond to the decrease in external employment by supplying more intensively to their own farms—this is also inconsistent with the idea that the wage dynamics are driven by a decrease in labor supply in the year after a positive shock. Similar arguments imply that factors such as income effects, migration, or capital investment are also not driving the empirical findings. I argue the pattern of results is most consistent with nominal rigidities.

There is some evidence that wages are less rigid in areas where the costs of rigidity are likely to be higher. Certain crops are especially sensitive to the amount of labor hired—for example, they experience large output losses if not harvested immediately upon reaching maturity. In areas with such crops, price flexibility is particularly important because inefficient labor allocation will lead to especially large profit losses. Consistent with this, districts that grow more labor-sensitive crops are more likely to experience nominal wage cuts during droughts. In addition, while these districts are equally likely to raise wages in response to transitory positive shocks, such shocks are less likely to have persistent wage effects. These results provide suggestive evidence that rigidity is endogenous to local economic conditions.[6]

Having established the presence of nominal rigidity, I explore possible mechanisms using a survey I conducted in rural India. A growing body of evidence argues that nominal wage cuts are perceived as unfair, causing decreases in worker productivity.[7] Following Kahneman, Knetsch, and Thaler (1986), I presented 400 agricultural laborers and landed farmers in 34 villages with scenarios about wage setting behavior, and asked them to rate the behaviors as fair or unfair on a 4-point scale. The results suggest that wage cuts strongly violate fairness norms. For example, 62% of respondents thought it was unfair for an employer to cut wages after a surge in unemployment. To examine differences between nominal and real wages, I presented a scenario in which employers cut real wages by 5%, but varied whether this was achieved through a nominal wage cut or inflation. 64% of respondents thought it was unfair to cut nominal wages by 5% during a period of no inflation. In contrast, only 9% of respondents thought it was unfair to raise nominal wages by 5% during a period of 10% inflation. Respondents also displayed a strong belief that workers decrease effort when fairness norms are violated. Consistent with the crop heterogeneity results, in villages with more labor-sensitive crops, respondents were considerably less likely to view wage cuts as unfair. This suggests that fairness norms may form, at least in part, endogenously.

The results point to the relevance of nominal rigidities in a setting with few of the institutional constraints that have received prominence in the existing literature. For example, in villages, minimum wage legislation is largely ignored and formal unions are rare (Rosenzweig 1980; 1988). Wage contracts for casual laborers are typically of short duration (on the order of days) and can more easily reflect recent changes in market conditions. Observing rigidity in such a context is consistent with the potential importance of non-institutional forces such as fairness norms in labor markets.

[6] These findings are only suggestive since planting decisions are endogenous. The causality could also run in the opposite direction: farmers could be more likely to plant labor sensitive crops in areas where there are weaker norms for rigid wages.

[7] Individual responses to a range of scenarios suggest the relevance of nominal variables (Shafir, Diamond, and Tversky 1997). Employers express perceptions that nominal wage cuts damage worker morale, with potential consequences for labor productivity (Blinder and Choi 1990; Bewley 1999). Lab and field studies validate the survey evidence (Fehr, Kirchsteiger, and Riedl 1993; Fehr and Falk 1999; Gneezy and List 2006). See Fehr, Goette, and Zehnder (2009) for a broader discussion of the relevance of fairness preferences in labor markets.

The rest of the paper proceeds as follows. Section 1.2 presents a model of nominal wage rigidity. Section 1.3 lays out the empirical strategy that will be used to test the model's predictions. Section 1.4 presents the results. Section 1.5 evaluates whether explanations other than nominal rigidity can explain the results. Section 1.6 discusses mechanisms for nominal rigidities and presents survey evidence for the role of fairness norms in villages. Section 1.7 concludes.

1.2 Model

In this section, I model a small open economy with decentralized wage setting and exogenous product prices. Rigidities arise because workers view wage cuts below a nominal reference wage as unfair, and retaliate to such cuts by decreasing effort.[8] In the empirical work, the reference wage will be the average nominal wage in the market in the previous period. I use this framework to develop testable implications of fairness preferences on labor market outcomes.

1.2.1 Set-up

The labor force is comprised of a unit mass of potential workers. All workers are equally productive. They are indexed by parameter ϕ_i, which equals worker i's cost of supplying 1 unit of effective labor. This parameter is distributed uniformly over the interval $\left[0, \overline{\phi}\right]$. The worker's payoff from accepting a nominal wage offer of w equals the utility from consuming her real wage minus the disutility of working:

$$u\left(\frac{w}{p}\right) - \phi_i e \left\{1 + \frac{1-\lambda}{\lambda} I_{\{w<w_R\}}\right\},$$

where p is the price level. The disutility of work equals ϕ_i times the amount of effort, e, exerted by the worker. The term in brackets captures fairness preferences. Workers view working for a wage below an exogenous nominal reference wage, w_R, as unfair. The worker's disutility of work is scaled up by $\frac{1-\lambda}{\lambda} I_{\{w<w_R\}}$, where $I_{\{w<w_R\}}$ is an indicator for whether the wage is below w_R and $\lambda \in (0, 1]$. When $\lambda = 1$, the disutility of work is the same regardless of whether $w < w_R$. As λ decreases, working for a wage below the reference wage imposes larger costs.

A market-wide fairness norm governs workers' effort behavior. The worker usually exerts a

[8] In Section 1.6, I provide support for this modeling assumption using survey evidence on village fairness norms. I also discuss whether other micro-foundations for rigidity are consistent with the context of the study and empirical results.

standard amount of effort: $e = 1$. However, when she feels treated unfairly by the firm, she reduces her effort to exactly offset the disutility from the fairness violation:

$$e = \begin{cases} 1 & w \geq w_R \\ \lambda & w < w_R \end{cases}. \tag{1.1}$$

In the model, I take this fairness norm as exogenous.[9] More generally, it can be conceptualized as the reduced form for a strategy in a repeated game. Worker i's payoff from accepting wage offer w reduces to $u\left(\frac{w}{p}\right) - \phi_i$. I normalize the payoff from not working as 0. When all firms offer wage w, aggregate labor supply is given by:

$$L^S = \frac{1}{\phi} u\left(\frac{w}{p}\right). \tag{1.2}$$

There are J firms (indexed by j), where J is large so that each firm's wage contributes negligibly to the average market wage. Firm j's profits from hiring L_j workers at nominal wage w_j equals:

$$\pi_j = p\theta f\left(eL_j\right) - w_j L_j, \tag{1.3}$$

where θ is a non-negative stochastic productivity parameter whose realization is common to all firms and $f(\bullet)$ is a continuous, increasing, twice-differentiable concave function. Note that output depends on effective labor—the number of workers times the effort exerted by each worker.

1.2.2 Benchmark Case: No Rigidity

I begin by solving the benchmark case in which there are no fairness preferences, i.e., when $\lambda = 1$. In this case, $e = 1$ for all w. Firm j's profits are given by:

$$\pi_j = p\theta f\left(L_j\right) - w_j L_j. \tag{1.4}$$

[9] This is similar to the conceptualization of worker retaliation in Akerlof and Yellen (1990). They assume an exogenous effort rule according to which workers reduce effort in proportion to how far their wage falls below a perceived fair wage, and examine the implications of this in an economy with inelastically fixed labor supply.

I focus on the symmetric pure strategy Nash Equilibrium, in which all firms offer the same wage:[10]

$$w_j = w^* (\theta, p) \ \forall j,$$

where $w^* (\theta, p)$ will be used to denote the benchmark equilibrium wage level at θ and p. Since firms are identical, they all demand the same amount of labor, $L^* (\theta, p)$. The firm's first order condition is:

$$p\theta f' (L^*) = w^*. \tag{1.5}$$

This pins down the optimal choice of labor at w^*; since this condition is the same for all firms, all firms will demand the same amount of labor L^*. The market clearing condition is characterized by:

$$JL^* = \frac{1}{\phi} u \left(\frac{w^*}{p} \right). \tag{1.6}$$

This condition simply equates the amount of aggregate labor demand with aggregate labor supply at (w^*, L^*).

Proposition 1.1: Market clearing in benchmark case

If workers do not exhibit fairness preferences, the unique pure strategy symmetric Nash Equilibrium will satisfy conditions (1.5) and (1.6). The labor market will clear for all realizations of θ.

Proof: See Appendix A. ■

Note that equations (1.5) and (1.6) correspond exactly to the conditions in a competitive equilibrium. Combining these equations and taking the derivative of w^* with respect to θ gives $\frac{\partial w^*(\theta,p)}{\partial \theta} > 0$ for all values of θ. Consequently, any decrease in θ will lead to a reduction in the equilibrium wage level.

[10] In villages, it is common for employers to conform to a single prevailing wage for agricultural workers. Section 1.4 (Table 1.5) presents evidence in support of this. It is therefore reasonable in this setting to focus on pure strategy symmetric equilibria.

Corollary 1.1: Complete adjustment to negative shocks in benchmark case

If workers do not exhibit fairness preferences, the wage will adjust downward with a decrease in θ over all θ-values.

1.2.3 Downward Rigidity at the Reference Wage

I now turn to examine the implications of fairness preferences on labor market outcomes. Firm profits are given by expression (1.3). Note that for any (w_j, L_j) combination, profits are always weakly lower in the fairness case than the benchmark case.

In the symmetric pure strategy Nash equilibrium:

$$w_j = \overline{w}\left(w_R, \theta, p\right) \; \forall j,$$

where $\overline{w}\left(w_R, \theta, p\right)$ will be used to denote the equilibrium wage level corresponding to reference wage w_R, TFP θ, and price p in the fairness case. All firms demand the same amount of labor, $\overline{L}\left(w_R, \theta, p\right)$. For a given $\overline{w}$, this is pinned down by the firm's first order condition, which exhibits a discontinuity around w_R :

$$\overline{w} = \begin{cases} p\theta f'\left(\overline{L}\right) & \overline{w} \geq w_R \\ p\theta\lambda f'\left(\lambda\overline{L}\right) & \overline{w} < w_R \end{cases}. \tag{1.7}$$

When $\overline{w} \geq w_R$, this corresponds exactly to the first order condition in the benchmark case. However, when $\overline{w} < w_R$, retaliation by the firm's workers makes them less productive. I assume:

$$f'\left(\overline{L}\right) > \lambda f'\left(\lambda\overline{L}\right) \; for \; \lambda < 1. \tag{1.8}$$

Condition (1.8) implies that for a given wage level $\overline{w} < w_R$, firms demand less labor than in the benchmark case. This condition always holds, for example, under Cobb-Douglas production: $f\left(eL\right) = \left(eL\right)^{\alpha}$.

Define θ_R as the unique value of θ at which w_R is the equilibrium wage and the labor market clears. Specifically, θ_R is pinned down by the the following two conditions:[11]

$$\overline{w}\left(w_R, \theta_R, p\right) = w_R$$

$$J\overline{L}\left(w_R, \theta_R, p\right) = \frac{1}{\phi} u\left(\frac{w_R}{p}\right).$$

The next proposition establishes asymmetric adjustment in wages around θ_R.

Proposition 1.2: Asymmetric adjustment to shocks

In the unique pure strategy symmetric Nash equilibrium:

(i) There exists a $\theta'_R < \theta_R$ such that for all $\theta \in (\theta'_R, \theta_R)$:

$$\overline{w}\left(w_R, \theta, p\right) = w_R \quad \textit{and} \quad J\overline{L}\left(w_R, \theta, p\right) < \tfrac{1}{\phi} u\left(\tfrac{\overline{w}(w_R, \theta, p)}{p}\right) \quad .$$

In addition, $\lim_{\lambda \to 0} \theta'_R = 0$.

(ii) For $\theta \geq \theta_R$, the wage will correspond to the benchmark case and the labor market will clear:

$$\overline{w}\left(w_R, \theta, p\right) = w^*\left(\theta, p\right) \quad \textit{and} \quad J\overline{L}\left(w_R, \theta, p\right) = \tfrac{1}{\phi} u\left(\tfrac{\overline{w}(w_R, \theta, p)}{p}\right) .$$

Proof: See Appendix A. ■

[11] Note that these two conditions imply that θ_R is the value of θ for which w_R is the equilibrium wage in the benchmark case: $w^*\left(\theta_R, p\right) = w_R$.

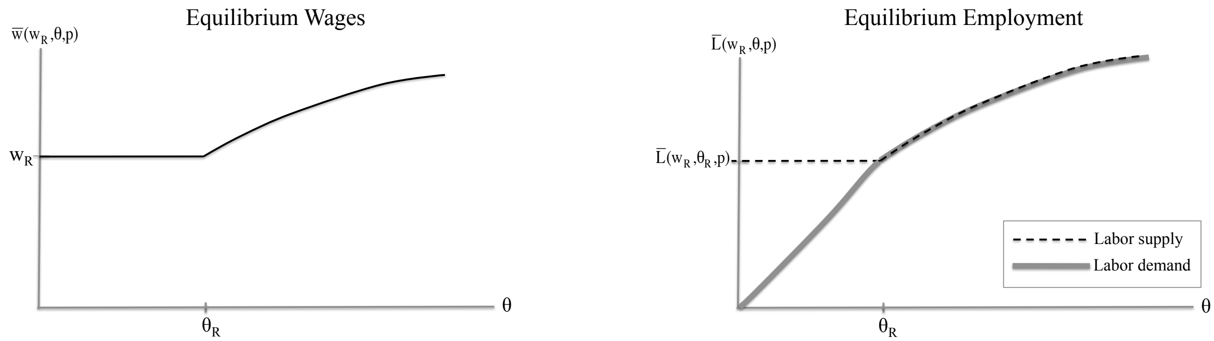

Figure 1.1: Asymmetric Adjustment to Shocks

Notes: This figure illustrates the relationship between wages and θ, and employment and θ (Proposition 1.2). The figure is drawn for the case of $\lambda \approx 0$.

For values of θ above θ_R, firms will increase wages smoothly as θ rises. However, for values of θ below θ_R, if firms cut nominal wages, they will suffer profit losses from decreases in worker effort. For sufficiently small decreases in θ below θ_R, it will be more profitable to maintain wages at w_R. Since the wage will remain the same, aggregate labor supply will remain the same. However, the firm's first order condition (1.7) implies labor demand will fall due to $\theta < \theta_R$, leading to excess supply in the market. Once θ falls to a sufficiently low level, below θ'_R, w_R can no longer be sustained as an equilibrium; equilibrium wages will be below w_R and the labor market will clear. Note that θ'_R will be lower for smaller values of λ: as λ approaches 0, firms will never find it profitable to lower wages below w_R. As a simple illustration, Figure 1.1 shows the relationship between realizations of θ and labor market outcomes for the case of $\lambda \approx 0$.

1.2.4 Impact of Increases in the Reference Wage

The above analysis implies that increases in the reference wage will expand the range of θ-values at which distortions occur. In addition, if λ is small so that wage cuts below the reference wage are rare, then reference wage increases will be particularly distortionary.[12]

Proposition 1.3: Distortions from reference wage increases

[12] In the empirical work, reference wage increases will arise from transitory positive shocks in the previous year—which raise the wage in the previous year and therefore lead to a higher reference wage in the current year.

Suppose the reference wage increases to $w_S > w_R$. For any $\theta < \theta_S$ and λ sufficiently small:

$$\overline{w}\left(w_S, \theta, p\right) > \overline{w}\left(w_R, \theta, p\right)$$
$$\overline{L}\left(w_S, \theta, p\right) < \overline{L}\left(w_R, \theta, p\right).$$

Proof: See Appendix A. ■

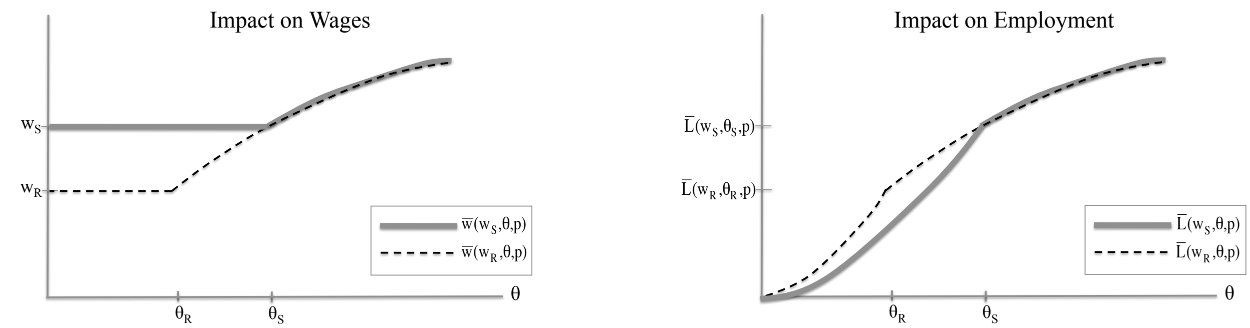

Figure 1.2: Distortions from Reference Wage Increases

Notes: This figure illustrates the impact of an increase in the reference wage on equilibrium wages and employment (Proposition 1.3). The figure is drawn for the case of λ≈0.

Since $\theta_R < \theta_S$, the wage distortions for $\theta \leq \theta_R$ will now be larger under w_S than they were under w_R. This will cause a particularly large excess supply of labor for $\theta \leq \theta_R$: labor demand will now be lower (leading to a drop in employment) while labor supply will actually be higher due to the wage increase. Figure 1.2 illustrates the impact of a reference wage increase on labor market outcomes for the case of $\lambda \approx 0$.

1.2.5 Impact of Inflation

In the benchmark case, prices are neutral. It is straightforward to verify from conditions (1.5) and (1.6):

$$\frac{\partial w^*\left(\theta, p\right)}{p} = \frac{w^*}{p}$$

$$\frac{\partial L^*(\theta, p)}{p} = 0.$$

Firms raise nominal wages to exactly offset the change in real wages from a price increase, keeping real wages constant and employment at the same level.

In contrast, when workers have fairness preferences over a nominal wage, inflation will no longer be neutral. When price levels rise, a given real wage level is associated with a higher nominal wage. As a result, for any w_R, a price increase means that the value of θ at which w_R is the market clearing nominal wage will now be lower. The rigidity will bind to the left of this lower θ value; this means distortions will affect a smaller portion of the θ-distribution.

Proposition 1.4: Inflation will mitigate distortions from nominal rigidity

For any fixed $\theta = \widetilde{\theta}$ and $p = \widetilde{p}$ such that:

$$\overline{w}\left(w_R, \widetilde{\theta}, \widetilde{p}\right) = w_R \text{ and } J\overline{L}\left(w_R, \widetilde{\theta}, \widetilde{p}\right) < \tfrac{1}{\phi} u\left(\tfrac{\overline{w}\left(w_R, \widetilde{\theta}, \tilde{p}\right)}{p}\right),$$

$\exists p' > \widetilde{p}$ such that $\forall p \geq p'$:

$$\overline{w}\left(w_R, \widetilde{\theta}, p\right) = w^*\left(\widetilde{\theta}, p\right) \text{ and } J\overline{L}\left(w_R, \widetilde{\theta}, p\right) = \tfrac{1}{\phi} u\left(\tfrac{\overline{w}\left(w_R, \widetilde{\theta}, p\right)}{p}\right)$$

Proof: See Appendix A. ■

For any fixed $\widetilde{\theta}$ at which the nominal rigidity binds (i.e. the wage is at the reference wage and there is excess supply), a sufficiently large increase in prices will lead to nominal wages rising above the reference wage. This will enable real wages to fall without incurring effort retaliation from workers. The wage at $\widetilde{\theta}$ will correspond to the benchmark case and the labor market will clear.

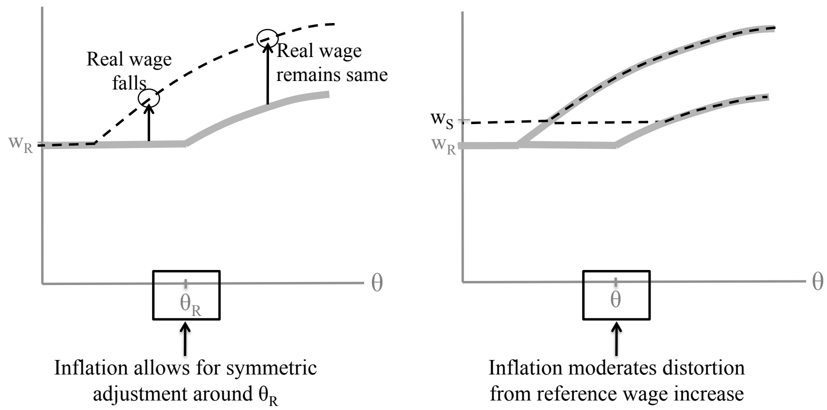

Figure 1.3: Inflation Mitigates Wage Distortions

Notes: This figure illustrates the impact of inflation in moderating wage distortions (Proposition 1.4). The graph on the left demonstrates the impact of an increase in price levels on equilibrium wages. The graph on the right illustrates the impact of inflation under an increase in the reference wage. The figure is drawn for the case of $\lambda \approx 0$.

Figure 1.3 illustrates how inflation moderates the distortions described in Propositions 1.2 and 1.3 for the case of $\lambda \approx 0$. By Proposition 2, at a fixed $\widetilde{p}$, there will be a θ^A around which adjustment will be asymmetric. For $p > \widetilde{p}$, the wage outcomes for θ-values just below θ^A will adjust so that: $w^* (\theta, p) = \overline{w} (w_R, \theta, p) < \overline{w} \left(w_R, \theta^A, p\right) = w^* \left(\theta^A, p\right)$. In other words, inflation will lead to symmetric adjustment around θ^A. In addition, inflation can also offset distortions from reference wage increases (Proposition 3). For any set of reference wages w_R and w_S, at a given $\widetilde{\theta}$, a sufficiently high p will cause nominal wages to be above w_S. Then, whether the reference wage is w_R or w_S will clearly have no impact on the equilibrium wage, which will correspond to the benchmark.

1.3 Empirical Strategy

1.3.1 Context: Rural Labor Markets in India

Agricultural production in India, as in most developing countries, is largely undertaken on smallholder farms. The median farm size is 1 acre, with considerable variation in landholdings.[13] The composition of farm employment varies and is often a mix of household and hired labor. Markets

[13] This, along with the remaining statistics in this sub-section, are from India's National Sample Survey data (1982-2008), described in Section 1.3.3 below.

for hired labor are active: most households buy and/or sell labor.[14] Workers are typically hired for standard tasks such as plowing, sowing, weeding, and harvesting.

The vast majority of hired labor is traded in decentralized markets for casual daily workers. For example, 97% of agricultural wage contracts are reported as casual wage contracts. In addition, 67% of landless rural workers report casual employment as their primary source of earnings.

There are few institutional constraints in these markets. Contracts are usually negotiated bilaterally between landowners and laborers in a decentralized manner; unions or other formal labor institutions are rare. Wage contracts are typically of short duration (on the order of days).[15] As a result, they can more easily reflect recent changes in market conditions and time worked is more flexible than in other contexts. Minimum wage policies are in practice ignored and there is little government intervention in the labor market (Rosenzweig 1980; 1988).

Agricultural production is heavily rainfall dependent and exhibits considerable seasonality in work intensity. The major rainfall episode is the yearly monsoon. The monsoon typically arrives between June-July in most parts of the country and marks the beginning of the agricultural year. For rice (the major crop in India) as well as some other crops, this is when field preparation and planting occur. The subsequent months involve various maintenance activities such as fertilizer application and weeding. Rice harvesting typically occurs between November and January, followed by post-harvest activities such as threshing and processing. March-May is the lean season. As discussed in Section 1.3.4, the monsoon affects labor demand in the various seasons over the agricultural year through impacts on planting levels, harvest yields, and the intensity of post-harvest tasks.

1.3.2 Empirical Tests

Testing the model's predictions requires identifying variation in both total factor productivity, θ, and in the nominal reference wage, w_R. I exploit rainfall variation to isolate shifters of both parameters in rural labor markets. A distinct labor market is defined as an Indian district (an administrative geographic unit). The empirical implementation will focus on discrete rainfall shocks: in each year, a labor market can experience a positive shock, no shock, or a negative shock (corre-

[14] See, for example, Table I of Rosenzweig (1980) for evidence from India, Tables I-III in Benjamin (1992) for evidence from Indonesia, and Bardhan (1997) for a broader discussion of the composition of agricultural employment.

[15] Of course, this does not rule out longer-term informal implicit contracts.

sponding to a high, medium, or low realization of θ, respectively). These shocks have no persistent productivity impacts: the value of θ in each year is determined solely by the rainfall shock in that year.[16] In addition, I assume the reference wage in a season equals the average market wage during that season in the previous year. A positive shock in the previous year would on average have raised wages in each season in the previous year, leading to a higher reference wage in each season this year.

This implies that examining the joint impact of lagged and current shocks on current wages can be used to test the model's predictions. Since there are 3 possible shocks in a given year, over every consecutive 2-year period, there are 9 possible realizations of shocks. This gives rise to the following estimating equation for wages:

$$\begin{aligned} w_{idst} \;=\; & \alpha_0^i + \alpha_0^{ii} S_{dt}^{\{-,0\}} + \alpha_1^i S_{dt}^{\{0,+\}} + \alpha_1^{ii} S_{dt}^{\{-,+\}} + \alpha_1^{iii} S_{dt}^{\{+,+\}} + \alpha_2^i S_{dt}^{\{0,-\}} + \alpha_2^{ii} S_{dt}^{\{-,-\}} \\ & + \alpha_3^i S_{dt}^{\{+,-\}} + \alpha_3^{ii} S_{dt}^{\{+,0\}} + \varphi \mathbf{X_{idst}} + \psi_{\mathbf{d}} + \eta_{\mathbf{t}} + \tau_{\mathbf{s}} + \varepsilon_{idst}, \end{aligned} \tag{1.9}$$

where w_{idst} is the nominal wage of worker i in district d in season s of year t; X_{idst} is a vector of controls; $\psi_{\mathbf{d}}$, $\eta_{\mathbf{t}}$, and $\tau_{\mathbf{s}}$ are vectors of district, year, and season fixed effects, respectively.

Each of the remaining 8 covariates is an indicator for the realization of a particular sequence of shocks. The indicators take the form $S_{dt}^{\{i,j\}}$, where i denotes district d's shock in year $t-1$ and j denotes the district's shock in year t. The i and j take the values $-$, 0, and $+$, which correspond to the realization of a negative shock, no shock, and a positive shock, respectively. Each indicator equals 1 if that particular sequence of shocks was realized and equals 0 otherwise. For example $S_{dt}^{\{+,0\}}$, equals 1 if district d had a positive shock last year and no shock this year, and equals 0 otherwise. The sequence $S_{dt}^{\{0,0\}}$, which is the case when the district experienced no shock last year and no shock this year, is omitted and serves as the reference case. Shocks are drawn from an iid distribution each year and are uncorrelated with the residual error, ε_{idst}. Thus, each of the coefficients on the indicator functions in equation (1.9) represents the reduced form average effect of that particular sequence of shocks on year t wages relative to $S_{dt}^{\{0,0\}}$.

Proposition 1.2 of the model predicts asymmetric adjustment in wages.

[16] This is a standard assumption in prior work that exploits rainfall shocks to investigate a range of outcomes in India (e.g., Paxson 1992; Rosenzweig and Wolpin 1993; Townsend 1994; Jayachandran 2006). In Section 1.5, I use the results to rule out persistent productivity impacts of shocks.

Prediction 1.1: Wage distortions: negative shocks

If there is no rigidity, wages will fall in response to negative shocks. In the presence of downward rigidities, wages may not fall in response to negative shocks.

I test this prediction by checking whether wages fall in response to contemporaneous negative shocks—the sequences $S_{dt}^{\{0,-\}}$, $S_{dt}^{\{-,-\}}$, and $S_{dt}^{\{+,-\}}$—relative to the reference case of $S_{dt}^{\{0,0\}}$. Note that in the presence of downward rigidities, $S_{dt}^{\{+,-\}}$ may cause wages to be even higher than the reference case, as discussed in Prediction 1.2 below.

Test 1.1:

$H_0 : \alpha_2^i < 0$, $\alpha_2^{ii} < 0$, and $\alpha_3^i < 0$.

$H_1 : \alpha_2^i = \alpha_2^{ii} = 0$ and $\alpha_3^i \geq 0$ under sufficiently severe downward rigidities.

Under the null of no rigidities, only current shocks should predict wages: for $S_{dt}^{\{i,j\}}$, the sign of the coefficient should be determined solely by j. However, under rigidities, lagged shocks can matter through their impact on the current reference wage (Proposition 1.3 of the model).

Prediction 1.2: Wage distortions: lagged positive shocks

If there is no rigidity, lagged positive shocks will have no impact on current wages. In the presence of downward rigidities, when a positive shock last year is followed by a non-positive shock this year, this may lead to higher current wages than if the lagged positive shock had not occurred.

I test this by checking whether the sequences $S_{dt}^{\{+,-\}}$ and $S_{dt}^{\{+,0\}}$ raise wages relative to the reference case of $S_{dt}^{\{0,0\}}$.

Test 1.2:

$H_0 : \alpha_3^i < 0$ and $\alpha_3^{ii} = 0$.

$H_1 : \alpha_3^i > 0$ and $\alpha_3^{ii} > 0$ under sufficiently severe downward rigidities.

Note that Tests 1.1 and 1.2 only have power under sufficiently large costs of worker retaliation (i.e. sufficiently small λ). If this cost is sufficiently low, firms will cut wages in response to decreases in θ and this estimation strategy may fail to produce a rejection of the null even though rigidities exist.

The remaining sequences of shocks in model (1.9) do not distinguish downward nominal rigidity from the benchmark case—they will have the same effects on wages in both cases. Specifically, when a negative shock is followed by no shock, wages should be the same as in the reference case: $\alpha_0^{ii} = 0$. In addition, contemporaneous positive shocks will increase wages relative to the reference case: $\alpha_1^i > 0$, $\alpha_1^{ii} > 0$, and $\alpha_1^{iii} > 0$. These sequences are included for completeness in the estimating equations.

To simplify the empirical analysis, the below alternate specification combines sequences of shocks into groups with common predictions:

$$\begin{aligned} w_{idst} \;=\; & \beta_0 \,+\, \beta_1 \left[S_{dt}^{\{0,+\}} + S_{dt}^{\{-,+\}} + S_{dt}^{\{+,+\}}\right] + \beta_2 \left[S_{dt}^{\{0,-\}} + S_{dt}^{\{-,-\}}\right] + \beta_3^i S_{dt}^{\{+,-\}} + \beta_3^{ii} S_{dt}^{\{+,0\}} \\ & + \varphi \mathbf{X_{idst}} + \psi_\mathbf{d} + \eta_\mathbf{t} + \tau_\mathbf{s} + \varepsilon_{idst}. \end{aligned} \tag{1.10}$$

The omitted category in model (1.10) is the sequences $S_{dt}^{\{0,0\}}$ and $S_{dt}^{\{-,0\}}$. Each term in brackets constitutes a new indicator function: $\left[S_{dt}^{\{0,+\}} + S_{dt}^{\{-,+\}} + S_{dt}^{\{+,+\}}\right]$ equals 1 if district d experienced a contemporaneous positive shock in year t and equals 0 otherwise, and $\left[S_{dt}^{\{0,-\}} + S_{dt}^{\{-,-\}}\right]$ equals 1 if the district had a non-positive shock last year followed by a negative shock this year. All other covariates are the same as in (1.9). Tests 1.1 and 1.2 imply that under sufficiently severe rigidities, $\beta_2 = 0$, $\beta_3^i > 0$, and $\beta_3^{ii} > 0$.

In the presence of rigidities, inflation will enable symmetric wage adjustment and moderate the effects of reference wage increases (Proposition 1.4 of the model):

Prediction 1.3: Impact of inflation on wage distortions

In the absence of rigidities, inflation will not alter the impact of shocks. In the presence of rigidities, when inflation is higher, wages will be more likely to be lower during negative shocks. In addition, when inflation is higher, lagged positive shocks will be less likely to raise current wages.

The following regression model adds interactions of each of the shock categories with the inflation rate:

$$\begin{aligned} w_{idst} \;=\; & \gamma_0 \,+\, \gamma_1 \left[S_{dt}^{\{0,+\}} + S_{dt}^{\{-,+\}} + S_{dt}^{\{+,+\}}\right] \,+\, \delta_1 I_{dt} \left[S_{dt}^{\{0,+\}} + S_{dt}^{\{-,+\}} + S_{dt}^{\{+,+\}}\right] \\ & + \gamma_2 \left[S_{dt}^{\{0,-\}} + S_{dt}^{\{-,-\}}\right] \,+\, \delta_2 I_{dt} \left[S_{dt}^{\{0,-\}} + S_{dt}^{\{-,-\}}\right] \\ & + \gamma_3^{i} S_{dt}^{\{+,-\}} \,+\, \gamma_3^{ii} S_{dt}^{\{+,0\}} \,+\, \delta_3^{i} I_{dt} S_{dt}^{\{+,-\}} \,+\, \delta_3^{ii} I_{dt} S_{dt}^{\{+,0\}} \\ & + \varphi \mathbf{X_{idst}} + \psi_{\mathbf{d}} \,+\, \eta_{\mathbf{t}} + \tau_{\mathbf{s}} + \varepsilon_{idst}, \end{aligned} \tag{1.11}$$

where I_{dt} is the percentage change in price levels in district d between years $t-1$ and t.

Test 1.3:

$H_0 : \delta_1 = \delta_2 = \delta_3^{i} = \delta_3^{ii} = 0$

$H_1 : \delta_1 = 0$, while $\delta_2 < 0$, $\delta_3^{i} < 0$, and $\delta_3^{ii} < 0$.

In the reference case, employers will raise nominal wages by the inflation rate to keep real wages constant. In cases where rigidities cause wage distortions—contemporaneous negative shocks and lagged positive shocks—employers can simply not adjust nominal wages upward, thereby achieving real wage reductions. As a result, when inflation is higher, nominal wages will be more likely to be lower in these cases relative to the reference case.

For a given non-positive shock in the current year, a transitory positive shock in the previous year raises current wages without affecting the current value of θ. This wage distortion should generate a distortion on employment:

Prediction 1.4: Employment distortions: lagged positive shocks

If there is no rigidity, lagged positive shocks will have no impact on current employment. In the presence of downward rigidities, when a positive shock last year is followed by a non-positive shock this year, this will lead to lower employment than if the lagged positive shock had not occurred.

The following model allows for tests for the impact of shocks on employment:

$$\begin{aligned} e_{idst} \;=\;& \rho_0 \,+\, \rho_1\left[S_{dt}^{\{0,+\}} + \rho_{dt}^{\{-,+\}} + \rho_{dt}^{\{+,+\}}\right] +\, \rho_2\left[S_{dt}^{\{0,-\}} + S_{dt}^{\{-,-\}}\right] \,+\, \rho_3^{i} S_{dt}^{\{+,-\}} \,+\, \rho_3^{ii} S_{dt}^{\{+,0\}} \\ & + \varphi\mathbf{X_{idst}} + \psi_{\mathbf{d}} \,+\, \eta_{\mathbf{t}} + \tau_{\mathbf{s}} + \varepsilon_{idst}, \end{aligned} \tag{1.12}$$

where e_{idst} is the employment level of worker i in district d in season s of year t, and all other covariates are the same as in model (1.10). As tests of Prediction 4, sequences $S_{dt}^{\{0,0\}}$ and $S_{dt}^{\{-,0\}}$ (which are the omitted category) serve as counterfactuals for $S_{dt}^{\{+,0\}}$—the value of θ is the same in the current year, the only difference is whether there is a wage distortion from a lagged positive shock. Similarly, $S_{dt}^{\{0,-\}}$ and $S_{dt}^{\{-,-\}}$ serve as counterfactuals for $S_{dt}^{\{+,-\}}$—while employment should fall under all these sequences, the fall should be relatively more severe for $S_{dt}^{\{+,-\}}$ due to the added wage distortion from the lagged positive shock.[17]

Test 1.4:

$H_0 : \rho_3^{i} = \rho_2 < 0$ and $\rho_3^{ii} = 0$

$H_1 : \rho_3^{i} < \rho_2 < 0$ and $\rho_3^{ii} < 0$.

In villages, those who own land have the option to exhaust their own labor supply on their farms before hiring external non-household labor. In contrast, the landless must sell their labor externally

[17] The model predicts employment distortions from all contemporaneous negative shocks. However, testing for distortions under the sequences $S_{dt}^{\{0,-\}}$ and $S_{dt}^{\{0,-\}}$ requires a counterfactual benchmark of how much employment would have fallen if wages were flexible. However, there is no clear benchmark for this; constructing one would require imposing assumptions about the parameters of the production function and labor supply elasticity. Consequently, I focus in the employment tests on the effect of lag positive shocks, which have clear counterfactuals with clean qualitative predictions.

to other farms at the prevailing wage. As a result, when nominal rigidities bind, those with less land will be most likely to be rationed out of the labor market.

Prediction 1.4A: Employment distortions will be more severe for those with less land

Those with less land will be relatively more likely to suffer employment losses after lagged positive shocks.

This prediction is readily tested by adding interactions of land size with the shock categories in regression model (1.12) and checking if the distortions from $S_{dt}^{\{+,-\}}$ and $S_{dt}^{\{+,0\}}$ are higher for those with less land.

1.3.3 Data

Wage and employment data for over 500 Indian districts during the years 1956-2008 is constructed using two primary datasets.

The first source is the rural sample of the Employment/Unemployment rounds of the Indian National Sample Survey, a nationally representative survey of over 500 Indian districts.[18] Households in each district are sampled on a rolling basis over the agricultural year (July to June). The agricultural years 1982, 1983, 1987, 1993, 1999, 2004, 2005, and 2007 are covered. The survey elicits daily employment and wage information for each household member over the 7 days preceding the interview. I compute the daily agricultural wage as paid earnings for casual agricultural work divided by days worked.[19] I measure agricultural employment as the total percentage of the interview

[18] A district is an administrative unit in India, with an average of 17 districts per state. Like counties in the US, districts vary greatly in size. On average, a district has approximately 2 million total residents.

[19] Agricultural work is identified in the questionnaire as work activity corresponding to one of 6 possible agricultural operations such as plowing, sowing, weeding, etc. The wage data is restricted to observations in which a worker was paid for work performed; these do not include imputed wages for self-employment. I use total wage earnings: cash plus in-kind wages. 93% of wage observations in the sample have some cash component. Given potential measurement error in the valuation of in-kind wages, as a robustness check, I have also performed the analysis using log cash wages as the dependent variable in the wage regressions. The results are similar.

reference period during which a worker was employed in agricultural activities (own farm work plus hired out labor). This employment variable is constructed for members of the agricultural labor force (i.e. individuals who report agriculture as their primary or subsidiary source of employment).

The second source is the World Bank Agriculture and Climate dataset, which provides yearly panel data on 228 Indian districts in 13 states over the agricultural years 1956-1987.[20] The unit of observation is a district-year. The reported wage variable equals the mean daily wage for a male ploughman in the district-year.[21] Data on 20 crops, including acres planted and yields, is also included.

Rainfall data is taken from *Terrestrial Precipitation: 1900-2008 Gridded Monthly Time Series* (version 2.01), constructed by the Center for Climatic Research, University of Delaware. Rainfall estimates are constructed for 0.5 by 0.5 degree latitude-longitude grids by interpolating from 20 nearby weather stations. I match the geographic center of each district to the nearest latitude-longitude node in the rain data. These district coordinates are included in the World Bank data; for the NSS data, I have obtained them using district boundaries from the Indian census. The measure of interest is rainfall in the first month when the monsoon typically arrives in a district, which ranges from May to July for each district. Rainfall shock definitions, discussed below, are constructed as deviations from the district's usual rainfall in the sample. Rainfall distributions are computed for each district separately for each dataset: they are based on the years 1956-1987 for the World Bank wage data and the years 1982-2007 for the NSS data.

The national inflation rate is constructed from CPI indices from the government publication *Agricultural Prices in India.* Inflation in year t is defined as the average change in monthly inflation from July of year t-1 to June of year t. For 1965-1987, this is computed as the mean inflation level across all states using the state CPI for Agricultural Workers. For 1956-1962, it is computed from the national Working Class Cost of Living Index, since agricultural CPI numbers are not available for these earlier years. There is no inflation data available for 4 of the years in the World Bank

[20] The dataset includes data on 271 districts. I limit analysis to 228 agricultural districts, which I define as the districts whose mean percentage of land area planted with rice in the sample is at least 1%. Since rice is the dominant crop in India, districts that do not grow any rice are unlikely to engage in substantial agricultural activity. Performing the analysis below with all 271 districts gives similar results, with slightly larger standard errors.

[21] This information was collected from sampled villages within each district. A knowledgeable person in each village, such as a school teacher or village official, was asked the prevailing wage rate in the village. In years when the data for a male ploughman are not available, wages for a general male agricultural laborer are used instead.

data: 1960, 1963, 1964, and 1975.

1.3.4 Definition of Shocks

Figure 1.4 shows the non-parametric relationship between rainfall levels and 3 outcomes: crop yields, agricultural employment, and agricultural wages. The yield and wage graphs use observations from the World Bank data. The employment results are from the NSS data since the World Bank dataset does not contain employment information. I regress each dependent variable on controls (including year and district fixed effects) and rainfall decile dummies.[22] Each decile dummy is an indicator for whether the district's rainfall in the first month of the monsoon that year fell within the given decile of the district's rain distribution. The graphs plot the coefficients for the decile dummies.

Figure 1.4 indicates that on average, high rainfall levels are associated with increased crop yields, higher agricultural employment, and wage increases. Note that even rainfall in the uppermost decile is a positive productivity shock. In contrast, low rainfall levels (droughts) are associated with lower average yields and employment; there is weak evidence that wages are lower. In the wage regression shown in Panel C, the F-test for joint significance of the ninth and tenth decile coefficients has a p-value of 0.066, while the F-test for the first and second deciles has a p-value of 0.357.

Panels A and B are consistent with the presumption that high (low) rainfall levels constitute positive (negative) shocks to the marginal product of labor and increase (decrease) labor demand. I create discrete categories for positive and negative shocks to reflect the non-linear effects of rainfall on productivity and to increase statistical power. A positive demand shock is defined as rainfall above the eightieth percentile for the district; a drought as rainfall below the twentieth percentile; and no shock as rainfall between the twentieth and eightieth percentiles.[23] Jayachandran (2006) also uses rainfall to identify labor demand shifts and employs the same percentile cutoffs in defining shocks. Table 1.1 provides summary statistics on rainfall shocks in the sample.

[22] District identifiers are not available for the first three rounds of the NSS data. For these years, the smallest geographic identifier is the region—there are on average 2.6 regions per state in the NSS data, and a region is comprised of 8 districts on average. As a result, for all regressions using the NSS dataset, the geographic fixed effects are region fixed effects for the first three rounds and district fixed effects for the remaining rounds. This is equivalent to using two pooled panels with separate fixed effects for analysis. Using a common set of region fixed effects for all rounds gives similar (though slightly less precise) results in the regressions.

[23] Although the cut-offs are symmetric, this does not presume that the magnitude of shocks from the upper and lower tails of the rainfall distribution is symmetric.

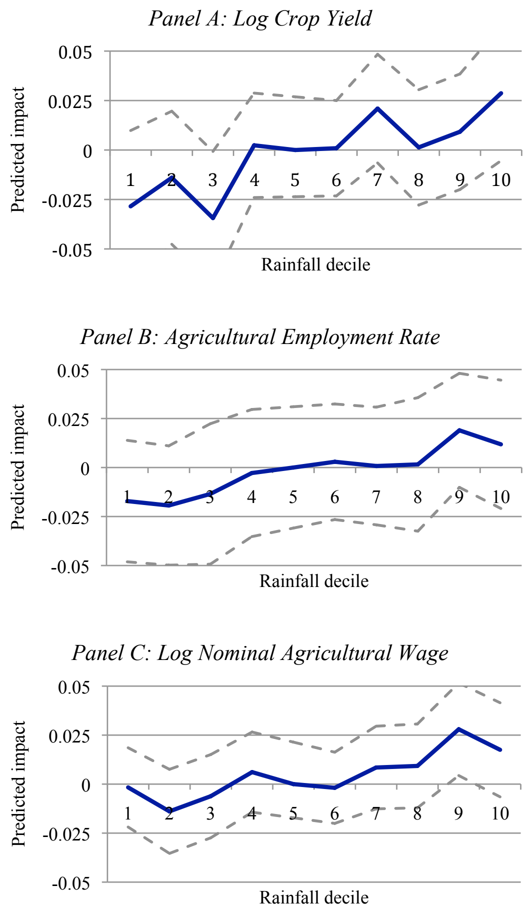

Figure 1.4: Impact of Rainfall on Agricultural Outcomes

Notes:
1. This figure shows the impact of rainfall on 3 outcome measures—log agricultural yields (from the World Bank data), the employment rate for workers in the agricultural labor force (from the NSS data), and the log nominal daily agricultural wage (in the World Bank data).
2. The panels plot coefficients from a regression of each outcome on dummies for each decile of the rainfall distribution and district and year fixed effects. Each decile dummy equals 1 if the district's rainfall in the current year fell within the given decile of the district's usual rainfall distribution and equals 0 otherwise. The 5th decile is the omitted category in each regression. The coefficients on the decile dummies are shown, along with 95% confidence intervals. The confidence interval for the 5th decile is computed by averaging the confidence intervals for the 4th and 6th deciles.
3. Standard errors are corrected to allow for clustering by region-year.

Rainfall is serially uncorrelated. I verify this in Appendix Table B.1. To allow for the possibility that shocks across districts may be correlated within a given year, standard errors are clustered by region-year in all regressions, using the region definitions provided in the NSS data.[24]

[24] Appendix Table B.1 provides some evidence for negative serial correlation in rainfall. Clustering standard errors by region makes minor difference in the results, and slightly improves precision in some cases. To be conservative, I cluster by region-year.

Table 1:1: Summary Statistics

Variable	Mean	Standard Deviation	Observations: District-years	Observations: Individual-years	Source
Rainfall shocks					
% Positive Shocks (1956-1987)	0.222	0.416	7,296	--	Univ of Delaware
% Droughts (1956-1987)	0.150	0.357	7,296	--	Univ of Delaware
% Positive Shocks (1982-2008)	0.178	0.383	3,419	--	Univ of Delaware
% Droughts (1982-2008)	0.186	0.389	3,419	--	Univ of Delaware
Wage and employment variables					
Log nominal wage (1956-1987)	1.200	0.815	7,296	--	World Bank
Log nominal wage (1982-2008)	3.279	0.890	--	154,578	Natl Sample Survey
Agricultural employment rate	0.494	0.484	--	1,003,431	Natl Sample Survey
Other measures					
Log crop yields index	0.237	0.271	7,296	--	World Bank
Acres per adult in household	0.776	8.260	--	1,003,431	Natl Sample Survey
Inflation rate	0.074	0.079	6,384	--	Agricultural Prices in India, Consumer Price Indices

Notes:

1. This table presents summary statistics for variables used in the analysis. Means and standard deviations are presented for each variable.
2. % Positive shocks and % Droughts gives the percentage of district-years in the data in which there was a positive rainfall shock or drought, respectively. A positive shock is defined as rainfall in the first month of the monsoon above the 80th percentile of the district's usual distribution and a drought is defined as rainfall below the 20th percentile of the district's usual distribution.
3. Log nominal agricultural wage is the log of the mean nominal daily wage for a male ploughman in the World Bank data during a district-year, and the log of the nominal daily wage for agricultural activities for an individual in the National Sample Survey data.
4. Log crop yields index is defined as the log of a composite index measure of the yields variable. The index is a weighted mean of yields of all 20 crops for which yields data are available, where the yield has first been normalized by the mean yield of that crop in the district. Weights are the mean percentage of landarea planted with a given crop in a district.
5. The national inflation rate is defined as the mean change in the CPI in the past agricultural year. For 1956-1962, this is computed from the national Working Class Living Index. For 1965-1987, this is computed by taking the average of the state CPI for Agricultural Workers across all states in the World Bank sample.

While shocks occur at the start of the agricultural year, the empirical approach assumes their impact persists over the entire year. Appendix Table B.2 examines differential impacts over calendar quarters. The variation in employment levels across quarters attests to the substantial seasonality in agriculture. However, I cannot reject that the impact of shocks on wages and employment is the same across quarters. As a result, in the analysis that follows, I pool observations within each agricultural year and examine mean impacts of shocks over the entire year.

1.4 Results

1.4.1 Distributions of Wage Changes

Before moving to the main empirical tests, I examine the distribution of wage changes for evidence of wage stickiness. Figure 1.5 displays histograms of year-to-year percentage wage changes in the World Bank panel. Panel A shows the distribution of nominal wage changes. The figure shows a bunching of mass to the right of nominal zero, with a discontinuous drop to the left of zero. 17 percent of observations are zero nominal changes. Since the district wage data is computed by averaging wages from sampled villages, this likely underestimates the percentage of zero changes in the underlying micro-data. In an economy experiencing a continuous distribution of shocks (from rainfall or other events) to the marginal product of labor, one would not expect a large discrete jump at zero in the absence of nominal rigidities (Kahn 1994; McLaughlin 1994). Consequently, this figure provides prima facie evidence for nominal rigidity. However, an important concern with this approach is measurement error in reported wages. If wages are reported in round increments (while actual wages vary continuously) or there is recall bias in reporting, this would make observing nominal zero changes more likely.

Panel B displays the distribution of real wage changes, using the local state CPI to compute real wages. Only 0.07% of observations are zero real wage changes. In addition, the mass is distributed fairly smoothly to the left and right of zero. There is little evidence of real wage rigidity.

Panels C-D examine whether real wage cuts are more likely when inflation is higher. I define high inflation years state inflation above 6% (slightly below the sample median). Both panels use observations in which real wage cuts should be especially likely. The histograms in Panel C limit observations to district-years with contemporaneous droughts. Only 29% of observations are real wage cuts in low inflation years, contrasted with 64% of observations in high inflation years.[25] The histograms in Panel D limit observations to district-years in which the district experienced a positive shock in the previous year, which would have caused an increase in wages in the previous year. These

[25] Of course, not all districts would be expected to cut wages since rainfall shocks are not the only determinants of labor productivity. Indian agriculture has gone through periods of strong national and localized growth—for example, from the adoption of green revolution technologies or infrastructure investments. Rainfall shocks are uncorrelated with these developments. Real wage increases are therefore expected even in the presence of negative rainfall shocks.

districts would on average have experienced a productivity decrease in the current year: real wages should be likely to fall. Again, the histograms show that real wage cuts are considerably more likely in high inflation years (67% of observations) than in low inflation years (30% of observations).

1.4.2 Tests for Wage Distortions

Table 1.2 tests for wage distortions from rigidities. The dependent variable is the log of the nominal daily wage for agricultural work. Columns (1)-(3) show results from the World Bank district data, covering the years 1956-1987. Columns (4)-(6) shows results from the NSS individual data, covering the years 1982-2008. Columns (1) and (4) provide estimates of regression model (1.9).[26] The results are qualitatively similar in both columns. As expected under both rigidity and flexible wage models, the coefficient on the sequence $S_{dt}^{\{-,0\}}$ (row 2) is indistinguishable from zero and contemporaneous positive shocks (rows 3-5) raise wages. For example, a zero shock last year followed by a positive shock this year increases wages by approximately 2.1% in the World Bank data and 4.5% in the NSS data.

The coefficients on $S_{dt}^{\{0,-\}}$, $S_{dt}^{\{-,-\}}$, and $S_{dt}^{\{+,-\}}$ (rows 6-8) estimate the impact of contemporaneous droughts (Test 1.1). Consistent with downward rigidity, there is little evidence of wage decreases under droughts in both datasets. While the $S_{dt}^{\{0,-\}}$ and $S_{dt}^{\{-,-\}}$ coefficients have a negative sign, they are generally small in magnitude and I cannot reject they are zero; the $S_{dt}^{\{+,-\}}$ coefficient is actually positive.

Finally, the coefficients on $S_{dt}^{\{+,-\}}$ and $S_{dt}^{\{+,0\}}$ (rows 8-9) test for effects of lag positive shocks (Test 1.2). In the World Bank data, a positive shock followed by no shock raises wages by 2.1% (significant at the 5% level). Even when a positive shock is followed by a drought, wages are 3.8% higher than the reference case (significant at the 10% level). These results bear out in the NSS data as well: this year's wages are about 2.6% and 11.5% higher on average when a positive shock last year is followed by a zero shock or drought, respectively, this year.[27]

[26] There is a small change in the specification for the World Bank data. Since the unit of observation is a district-year, the dependent variable is w_{dt}, the log nominal wage in district d in year t, and there are no individual-level controls or season fixed effects.

[27] This coefficient of 0.115, which measures the mean impact of a positive shock followed by a negative shock in the NSS data, is surprisingly large. However, this seems to be a result of sampling variation in the data. One cannot reject, for example, that this coefficient is the same as the measured impact of a negative shock followed by a positive shock in the NSS data.

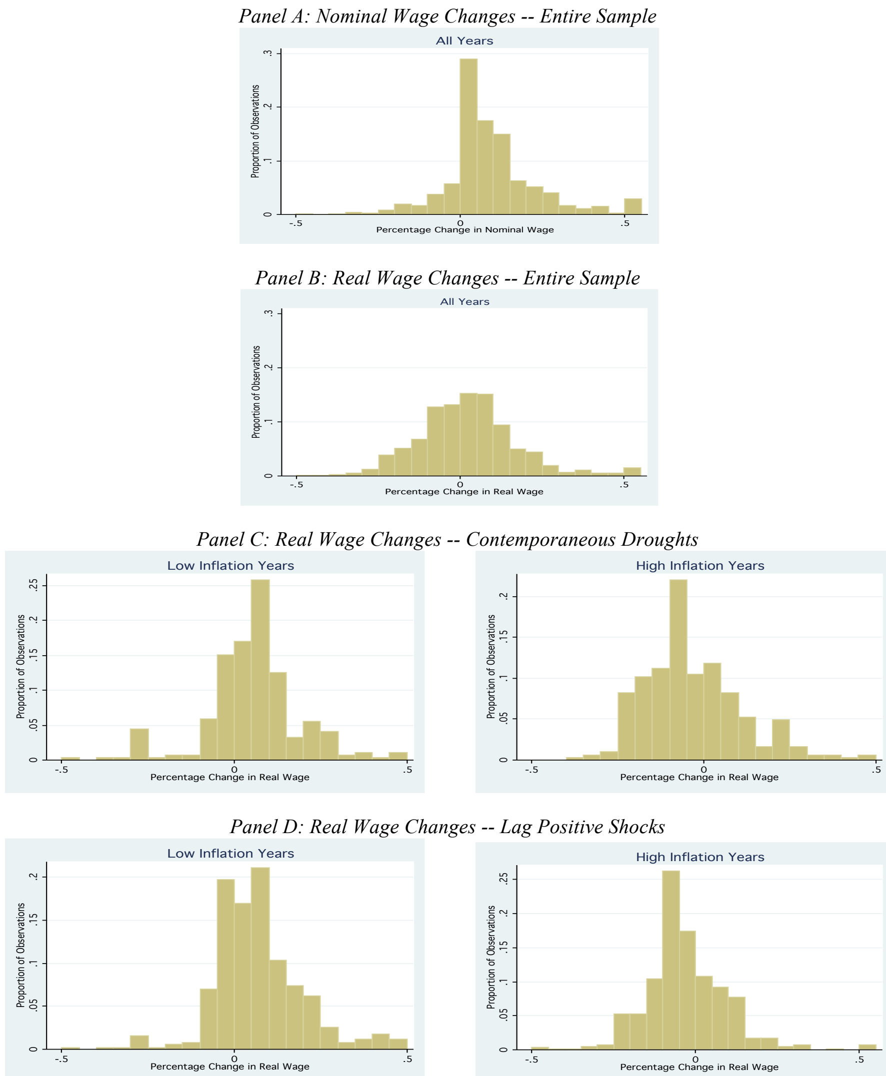

Figure 1.5: Distributions of Wage Changes

Notes:
1. This figure shows distributions of year-to-year percentage changes in agricultural wages in the World Bank dataset.
2. The unit of observation is a district-year. Real wages are computed as the nominal wage divided by the state CPI for agricultural workers. Low (high) inflation years are defined as years in which the state inflation rate was below (above) the sample median of 6%. Droughts and positive shocks are defined as rainfall in the first month of the monsoon below the 20th percentile and above the 80th percentile, respectively, of the district's usual rainfall distribution in that month.
3. Panel A shows nominal wage changes in the entire sample (8,401 district-years). The observations in Panels B-D are drawn from the 4,642 district-years for which state CPI data is available. Panel B shows real wage changes in this sample. Panel C uses observations in which a district experienced a drought in that year. Panel D uses observations in which a district experienced a positive shock in the previous agricultural year. Panels C and D show distributions separately for low and high inflation years.

Table 1.2: Effect of Shocks on Equilibrium Wages

Dependent Variable: Log Nominal Daily Agricultural Wage

			Source: World Bank Data (1956-1987)			Source: NSS Data (1982-2008)		
			(1)	(2)	(3)	(4)	(5)	(6)
	Shock$_{t-1}$	*Shock*$_t$						
1	Zero	Zero	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted
2	Drought	Zero	0.003	Omitted	Omitted	0.002	Omitted	Omitted
			(0.011)			(0.015)		
3	Zero	Positive	0.021			0.045		
			(0.010)**			(0.012)***		
4	Drought	Positive	0.064	0.026	0.026	0.079	0.052	0.052
			(0.019)***	(0.009)***	(0.009)***	(0.028)***	(0.011)***	(0.011)***
5	Positive	Positive	0.014			0.066		
			(0.016)			(0.023)***		
6	Zero	Drought	-0.006			0.006		
			(0.013)			(0.016)		
				-0.010	-0.010		-0.003	-0.002
				(0.011)	(0.011)		(0.013)	(0.013)
7	Drought	Drought	-0.015			-0.025		
			(0.018)			(0.028)		
8	Positive	Drought	0.038	0.037		0.115	0.114	
			(0.021)*	(0.020)*		(0.018)***	(0.019)***	
					0.024			0.056
					(0.010)**			(0.013)***
9	Positive	Zero	0.021	0.021		0.026	0.025	
			(0.010)**	(0.010)**		(0.014)*	(0.015)*	
District and year FE?			Yes	Yes	Yes	Yes	Yes	Yes
Additional controls?			No	No	No	Yes	Yes	Yes
Obs: district-years			7,296	7,296	7,296	--	--	--
Obs: individual-years			--	--	--	154,476	154,476	154,476
Dependent var mean			1.197	1.197	1.197	3.261	3.261	3.261

Notes:

1.This table tests for the impacts of sequences of shocks on the agricultural wage.

2. The dependent variable is the log of the nominal daily agricultural wage.

3. The shocks are defined as drought, zero, or positive, and correspond to rainfall below the 20th percentile, between the 20th-80th percentiles, and above the 80th percentile, respectively. The covariates are indicators that equal 1 if a given sequence of shocks was realized and zero otherwise. The sequences are presented as the shock in the previous year and the shock in the current year.

4. Columns (1) and (3) omit the sequence {Zero, Zero} and include separate dummies for each of the remaining 8 combinations of shocks. The remaining columns group shocks into categories with similar predictions. Columns (2) and (4) also omit the sequence {Drought, Zero}; combine rows 3-5 into one indicator function for whether the district experienced a contemporaneous positive shock; and combine rows 6-7 into an indicator function for whether the district had a zero shock or drought last year followed by a contemporaneous drought. Columns (3) and (6) repeat this specification, but also combine rows 8-9 into one indicator for whether the district had a positive shock last year followed by a drought or zero shock this year.

5. Each regression also contains year and district fixed effects. Regressions (4)-(6) from the NSS data also include fixed effects for calendar quarters of the year and a dummy for gender.

6. Standard errors are corrected to allow for clustering by region-year.

Columns (2) and (5) of Table 1.2 repeat this analysis for the simpler specification in regression model (1.10). The results are similar to the previous columns. Finally, columns (3) and (6) repeat this specification, but also collapse the sequences used for the lag positive shock tests ($S_{dt}^{\{+,-\}}$ and $S_{dt}^{\{+,0\}}$; rows 8-9) into one cell. The results in these columns indicate that in both datasets, wages are on average the same this year regardless of whether the positive shock occurred last year or this year.[28]

For simplicity, the main specification focuses on shocks in the current year and previous year only. Appendix Table B.4 examines the duration of persistence of shocks. In the World Bank data—positive shocks raise nominal wages for up to 5 years. In the NSS data, they do not significantly impact wages for more than 1 future year on average. This is consistent with higher levels of real agricultural growth in India during the NSS data years. As expected, droughts have no persistent effects in either dataset. Note that focusing on only last year's shocks in the main specification makes a rejection in Test 1.2 (lag positive shocks) less likely—the main specification therefore enables simplicity without biasing the results towards finding rigidity.

1.4.3 Impact of Inflation on Wage Distortions

To test whether inflation moderates the wage distortions documented above, I use the World Bank data since it covers 32 years, providing substantial variation in inflation. Column (1) of Table 1.3 shows estimates of model (1.10) for the restricted sample for which inflation data is available for comparison purposes. The regressions in columns (2)-(3) add interactions of each of the shock categories with measures of the national inflation rate. In column (2), the measure is the continuous inflation rate–this corresponds to the specification in model (1.11). In column (3), the inflation measure is an indicator that equals 1 if the inflation rate is above 6% (slightly below the sample median) and equals 0 otherwise.

[28] In Appendix Table B.3, I use an alternate specification to test for the impact of shocks on wages in both datasets. Instead of the 9 discrete cells, I include dummies for positive shocks and droughts in current and previous periods, along with a full set of interactions between current and lagged shocks. The model offers 2 sets of predictions under the null of no rigidity. First, contemporaneous droughts should lead to wage decreases. As in Table 1.2, there is no support for this. Second, lag shocks should not predict current wages. The F-test p-values reported at the bottom of the table test this restriction for covariates involving lag positive shocks and also for covariates involving any lag shocks—these tests are significant at the 5% level or less in each case.

Table 3: Impact of Inflation on Wage Distortions

Dependent Variable: Log Nominal Daily Agricultural Wage

	Definition of Lag Positive Shocks:	*Positive shock in previous year*			*At least one positive shock in last 3 years*		
	Measure of Inflation	--	Inflation rate	Inflation>6%	--	Inflation rate	Inflation>6%
		(1)	(2)	(3)	(4)	(5)	(6)
	$\{Shock_{t-1}$=Drought or Zero$\}$; $\{Shock_t$=Zero$\}$	Omitted	Omitted	Omitted	Omitted	Omitted	Omitted
1	$\{Shock_{t-1}$=Drought, Zero, or Positive$\}$; $\{Shock_t$=Positive$\}$	0.017	0.020	0.021	0.024	0.035	0.034
		(0.009)*	(0.011)*	(0.011)*	(0.010)**	(0.011)***	(0.012)***
2	$\{Shock_{t-1}$=Drought, Zero, or Positive$\}$; $\{Shock_t$=Positive$\}$ x		-0.031	-0.006		-0.127	-0.017
	Inflation measure		(0.101)	(0.017)		(0.109)	(0.018)
3	$\{Shock_{t-1}$=Drought or Zero$\}$; $\{Shock_t$=Drought$\}$	-0.020	0.000	0.001	-0.038	0.014	-0.004
		(0.011)*	(0.014)	(0.016)	(0.015)***	(0.018)	(0.020)
4	$\{Shock_{t-1}$=Drought or Zero$\}$; $\{Shock_t$=Drought$\}$ x		-0.254	-0.036		-0.577	-0.056
	Inflation measure		(0.156)	(0.023)		(0.188)***	(0.028)**
5	$\{Shock_{t-1}$=Positive$\}$; $\{Shock_t$=Drought$\}$	0.019	0.039	0.061	0.028	0.040	0.057
		(0.020)	(0.034)	(0.033)*	(0.016)*	(0.024)*	(0.027)**
6	$\{Shock_{t-1}$=Positive$\}$; $\{Shock_t$=Drought$\}$ x		-0.218	-0.080		-0.154	-0.057
	Inflation measure		(0.251)	(0.040)**		(0.198)	(0.034)*
7	$\{Shock_{t-1}$=Positive$\}$; $\{Shock_t$=Zero$\}$	0.015	0.044	0.042	0.029	0.049	0.052
		(0.011)	(0.016)***	(0.018)**	(0.008)***	(0.012)***	(0.014)***
8	$\{Shock_{t-1}$=Positive$\}$; $\{Shock_t$=Zero$\}$ x		-0.336	-0.047		-0.248	-0.040
	Inflation measure		(0.128)***	(0.021)**		(0.096)***	(0.017)**
	Year and district fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
	F-test p-value: Coefficient 3 + Coefficient 4 = 0	--	--	0.027**	--	--	0.002***
	F-test p-value: Coefficient 5 + Coefficient 6 = 0	--	--	0.400	--	--	0.990
	F-test p-value: Coefficient 7 + Coefficient 8 = 0	--	--	0.678	--	--	0.207

Notes:

1. This table tests whether inflation mitigates distortions from rigidity. Observations are from the 6,384 district-years in the World Bank data for which inflation data is available (years 1956-87 except 1960, 1963-64, and 1975). The dependent variable is the log of the nominal daily agricultural wage; the dependent variable mean is 1.27.
2. Shocks are defined as drought, zero, or positive, and correspond to rainfall below the 20th percentile, between the 20th-80th percentiles, and above the 80th percentile, respectively. Covariates 1,3,5, and 7 are of the form $\{Shock_{t-1}$=X$\}$;$\{Shock_t$=Y$\}$; they are indicators that equal 1 if the district experienced shock X in the previous year and shock Y in the current year, and equal 0 otherwise. The other covariates are interactions of these sequences of shocks with measures of inflation.
3. In Columns (1)-(3), $\{Shock_{t-1}$=Positive$\}$ refers to having a positive shock in the previous year; in columns (4)-(6), this refers to having at least one positive shock in the last 3 years. In all columns, $\{Shock_{t-1}$=Drought or Zero$\}$ refers to having a drought or zero shock in the previous year.
4. Columns (1) and (4) show OLS regressions of the dependent variable on the sequences of shocks. Columns (2) and (5) add interactions of each shock category with the national inflation rate. Columns (3) and (6) add interactions of each shock category with an indicator for whether the national inflation rate was above 6%.
5. Each regression contains year and district fixed effects. Standard errors are corrected to allow for clustering by region-year.

Consistent with Test 1.3, contemporaneous droughts and lag positive shocks are less likely to cause wage distortions when inflation is higher. For example, the results in column (3) indicate that when a non-positive shock is followed by a drought, wages are the same as the reference cell on average when the inflation rate is below 6% (row 3). In contrast, when inflation is above 6%, wages are 3.6% lower than the reference cell (row 4). The F-test for whether, under high inflation, wages are the same during droughts as the reference cell has a p-value of 0.027 (reported at the bottom of the table). Thus, wages are indeed lower under droughts when inflation is sufficiently high. Similarly, when inflation is low, lag positive shocks increase current nominal wages (rows 5 and 7). When inflation is high, lag positive shocks do not cause persistent effects on future wages (the interactions in rows 6 and 8 are negative). For example, when a positive shock is followed by no shock, I cannot reject that wages are the same as the reference cell when inflation is above 6% (p-value 0.678).

The regressions in columns (4)-(6) repeat this analysis, with one change in the definition of shocks. To exploit the fact that positive shocks persist over many years in the World Bank data, I define a lag positive shock as at least one positive shock anytime in the past 3 years.[29] The remaining shock definitions remain the same. This yields qualitatively similar results to the regressions in columns (1)-(3), but increases precision.

In Appendix Table B.5, I rule out two sets of potential concerns. The first is that rainfall shocks may influence the inflation rate. Columns (1)-(2) show regressions of the national inflation rate on the shock categories (as defined in column (1) of Table 1.3). There is little correlation between shocks and inflation—the coefficients on contemporaneous droughts and lag positive shocks (rows 3, 5, and 7) are especially small and insignificant. As a further check, column (3) shows a regression of the log nominal wage on the shock categories and an interaction with inflation, where the inflation rate has been computed as the mean inflation rate across all states except the district's own state. This is a useful robustness check since a district's local rainfall is especially unlikely to be correlated with inflation in other states. The results are similar to those in Table 1.3, though are less significant since state-level inflation data is available for a limited number of years. The second concern is that there are co-trends in inflation and the impact of rainfall shocks. For example, if inflation and the

[29] The results are similar if other definitions for lag positive shocks are used instead, such as at least one shock in the past 2 years or 4 years.

adoption of irrigation (which makes crops less reliant on rainfall) both trend upward over time, this could create a spurious correlation. I check for such co-trends by interacting the shocks with a linear time trend in column (4) and a dummy for whether the year is after 1970 (the sample mid-point and the beginning of India's green revolution) in column (5). The interaction coefficients in both columns are extremely small and insignificant, indicating that the inflation results are not driven by co-trends.

1.4.4 Tests for Employment Distortions

Figure 1.6 compares kernel density estimates of mean employment in district-years with and without lagged positive shocks in the NSS data. The observations are limited to district-years in which there was no contemporaneous positive shock. Consistent with Prediction 1.4, lagged positive shocks cause the employment distribution to shift to the left. This provides initial evidence that downward rigidity reduces aggregate employment.

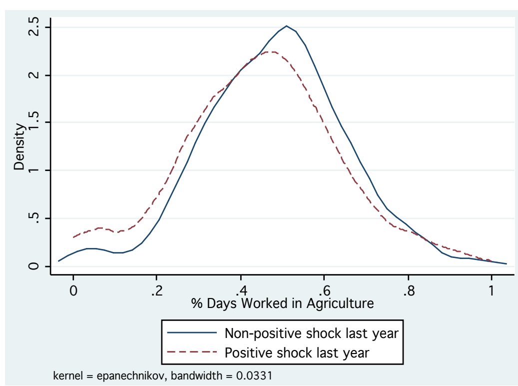

Figure 1.6: Employment Distortions from Rigidity – Impact of Lag Positive Shocks

Notes:
1. This figure displays the impact of lag positive shocks on agricultural employment in the NSS data.
2. The outcome variable is the mean of the percentage of days in the interview reference period in which agricultural workers were employed in agricultural work (own farm work plus hired out work) in each district-year.
3. A positive (non-positive) shock in a district-year is defined as rainfall above (below) the 80th percentile of the district's rainfall distribution.
4. The solid line plots kernel density estimates of the outcome variable for district-years in which there was a non-positive shock in the previous year and a non-positive shock in the current year. The dashed line plots kernel density estimates for district-years in which there was a positive shock in the previous year and a non-positive shock in the current year.
5. The estimates use the Epanechnikov kernel function. The bandwidth minimizes the mean integrated squared error assuming a Gaussian distribution and kernel.

Table 1.4: Effect of Shocks on Employment

Dependent variable: Total worker-days in agriculture

	Sample	Full Sample (1)	Full Sample (2)	Lean Season Excluded (3)
	Panel A: Average Impact of Lag Positive Shocks			
	Lag positive shock	-0.111	-0.217	-0.220
		(0.046)**	(0.049)***	(0.052)***
	Lag positive shock x Acres per adult in household		0.141	0.139
			(0.029)***	(0.028)***
	Panel B: Full Specification			
	{$Shock_{t-1}$=Drought or Zero; $Shock_t$=Zero}	Omitted	Omitted	Omitted
1	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive}	0.078	0.078	0.104
		(0.047)*	(0.047)*	(0.051)**
2	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive} x Acres per adult in household		-0.006	-0.005
			(0.005)	(0.004)
3	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Negative}	0.116	-0.112	-0.095
		(0.049**	(0.050)**	(0.051)*
4	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought} x Acres per adult in household		-0.001	-0.001
			(0.015)	(0.013)
5	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought}	-0.244	-0.352	-0.365
		(0.076)***	(0.073)***	(0.074)***
6	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought} x Acres per adult in household		0.123	0.135
			(0.037)***	(0.038)***
7	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero}	-0.107	-0.213	-0.205
		(0.058)*	(0.064)***	(0.072)***
8	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero} x Acres per adult in household		0.151	0.132
			(0.040)***	(0.038)***
9	Acres per adult in household	0.047	0.007	0.037
		(0.015)***	(0.002)***	(0.014)***
10	(Acres per adult in household)^2	$-1.04x10^{-5}$	$-1.50x10^{-6}$	$-7.73x10^{-6}$
		$(3.37x10^{-6})$***	(5.40x10-7)***	$(3.02x10^{-6})$**
	F-test p-value: Coefficient 3 = Coefficient 5	0.117	0.002***	0.001***
	Observations: individual-years	1,003,030	1,003,030	755,347
	Dependent variable mean	3.48	3.48	3.62

Notes:

1. This table tests for employment impacts of shocks. The dependent variable is the number of worker-days in the last 7 days in which the worker was employed in agricultural work (own farm work plus hired out work). Columns (1)-(2) include all observations for agricultural workers (workers whose primary or subsidiary work activity is agriculture). Column (3) excludes observations for the lean quarter (April-June).
2. Shocks are defined as drought, zero, or positive, and correspond to rainfall below the 20th percentile, between the 20th-80th percentiles, and above the 80th percentile, respectively.
3. Panel A shows regressions of the dependent variable on an indicator for whether the district experienced a positive shock in the previous year, and an interaction of this indicator with acres per adult in the household.
4. In Panel B, covariates 1,3,5, and 7 are of the form {$Shock_{t-1}$=X};{$Shock_t$=Y}; they are indicators that equal 1 if the district experienced shock X in the previous year and shock Y in the current year, and equal 0 otherwise. The other covariates are interactions of these sequences of shocks with acres per adult in the household.
5. Each regression also contains year fixed effects, district fixed effects, fixed effects for calendar quarters of the year, a gender dummy, and a quadratic function of acres per adult in the household.
6. Standard errors are corrected to allow for clustering by region-year.

Table 1.4 provides statistical tests of employment distortions and quantifies their magnitude. The dependent variable is the number of worker-days in the last 7 days in which the worker was employed in agricultural work (own farm work plus hired out work). Panel A begins by examining the mean impact of lagged positive shocks on employment. Agricultural laborers and farmers experience an average decrease in employment of 0.111 days per week if their district experienced a positive shock in the previous year (relative to no positive shock in the previous year). This constitutes a 3% decrease in agricultural activity. Column (2) adds an interaction with a measure of landholding: acres per adult in the household. In the year after a positive shock, landless laborers experience a 6% decrease in employment (significant at the 1% level). In contrast, those with land are less likely to face rationing. Column (3) repeats the analysis in column (2) but excludes observations from the lean quarter (April-June), when there is limited agricultural activity; the results are quite similar.

Panel B examines employment effects using the full specification. Column (1) provides estimates of regression model (1.12). Contemporaneous positive shocks (row 1) raise average employment by 0.078 days per week, or 2.2 percent. Contemporaneous droughts—which did not lead to wage cuts—do decrease employment. The $S_{dt}^{\{0,-\}}$ and $S_{dt}^{\{-,-\}}$ sequences (row 3) reduce employment by 0.116 days; this constitutes a 3.3% reduction and is significant at the 5% level.

The coefficients in row 5 and 7 provide tests of Prediction 1.4. When a drought is preceded by a positive shock, employment drops by about 0.25 days per week (row 5). This magnitude is twice as large as the decrease that occurs when a drought is not preceded by a lag positive shock (row 3). In addition, when a lag positive shock is followed by no shock (row 7), the average worker experiences a drop in employment of about 0.107 days (or 3%) relative to the reference cell.[30]

The regression in column (2) add interactions of acres per adult with each of the shock categories. The results conform to Prediction 1.4A. When a positive shock last year is followed by a drought in the current year (sequence $S_{dt}^{\{+,-\}}$, row 5), landless laborers are predicted to experience an employment decrease of 0.352 days per week; this corresponds to 10% of the mean employment level (significant at 1%). This magnitude is significantly larger than the 0.112 day decrease that results from a drought that wasn't preceded by a positive shock last year (row 3). The F-test for

[30] Appendix Table B.6 repeats this analysis, showing the impact of each of the 9 sequences of shocks on employment separately.

equality of the two coefficients has a p-value of 0.002. Note that difference in these coefficients is 0.240—about twice as large as the magnitude of the baseline employment effect of a drought. As expected, these employment decreases are less severe for those with more land: each additional acre of land per household adult is associated with an increase in employment of 0.123 days per week. Similarly, when a positive shock last year is followed by no shock in the current year ($S_{dt}^{\{+,0\}}$, row 7), the employment of landless agricultural workers is 0.213 days lower than if there hadn't been a positive shock in the previous year. It constitutes a 7% employment reduction and is significant at the 5% level. Again, this effect is about twice as large as the decrease in employment under a drought. Also in this case, landholdings mitigate these adverse employment effects. Finally, column (3) excludes observations from the lean quarter; the results are quite similar to column (2), and slightly stronger.[31]

In Appendix Table B.7, I investigate a potential concern with the interpretation of the employment results: the possibility that rainfall shocks alter the composition of the agricultural labor force. In the presence of compositional effects, the employment variable will not accurately estimate changes in aggregate employment levels. For example, if lag positive shocks cause in-migration, increasing the number of agricultural workers, the percentage of days worked by each worker could decrease even if the aggregate number of worker-days has gone up. Appendix Table B.7 investigates two ways in which shocks could create compositional changes—through migration into the village and by altering the probability that respondents identify agriculture as their occupation. There is little evidence that lag positive shocks influence either of these outcomes.[32]

[31] While landholding is an important determinant of worker-days spent in agriculture, it does not impact the wage received by workers. When the log nominal daily agricultural wage is regressed on the covariates in the regression in Column (2) of Table 1.4, the coefficients on the landholding controls and interactions terms are all extremely small in magnitude (between 0.00-0.002) and insignificant. These results are consistent with the presence of a prevailing market wage, which is the same for all agricultural workers who sell their labor externally on the market.

[32] There is evidence that individuals are less likely to migrate into the village during contemporaneous droughts—in the main specification (Panel B), migration falls by 0.1%. However, the fact that the labor force is relatively smaller during droughts is unlikely to be the reason wages don't fall during these shocks. As a simple calibration, since the mean employment rate is 0.498, this can explain only a 0.001*0.498 = 0.000498 percentage point change in the number of worker-days.

1.4.5 Separation Failures Test: Compositional Effects on Household Employment

When employment is rationed, the household's labor supply decision will no longer be separable from its decision of how much labor to use on its farm. Households with less land, who cannot find external employment when rigidities bind, will supply more intensively to their own farms. Table 1.5 provides a test of this prediction. It decomposes total household agricultural employment into worker-days in the external labor market (as a paid agricultural labor) and worker-days on the household's own farm.

Panel A begins by examining the average impact of lagged positive shocks by landholding. Households are defined as having small, medium, and large landholdings, corresponding to the the lower, middle, and upper terciles of the sample distribution of acres per adult in the household, respectively. The sample is limited to agricultural households with positive landholding. The dependent variable in column (1) is the total number of days spent by household members in external employment (as a hired agricultural laborer on someone else's farm). Consistent with the results in Table 1.4, households with small landholdings face reductions in external employment after lagged positive shocks, while households with medium and large landholdings do not.

Column (2) provides the key test of the separation failure prediction. The dependent variable is the number of days spent by household members on their own farm. In the year after a positive shock, households with small landholdings—who are rationed out of the external market—increase labor supply on their own farms by half a day a week on average. This is a 7% increase relative to the mean, and is about the same as how much these households increase own-farm production during a contemporaneous positive shock (see Panel B below). In contrast, own-farm labor supply does not change after lagged positive shocks for medium-landholding households and actually decreases for large landowners (perhaps due to decreased supervision time in the field since less external labor is being hired). Column (3) shows the sum of off-farm and own-farm employment (the same dependent variable as in Table 1.4). Because households supplement decreased external employment with increases in own-farm work, there is little aggregate movement in total household employment after lagged positive shocks.

Table 1.5: Separation - Compositional Effects on Employment

	Dependent Variable	*Worker-days as agric laborer* (1)	*Worker-days on own farm* (2)	*Total worker-days in agric* (3)
	Panel A: Average Impact of Lag Positive Shocks			
	Lag positive shock	-0.797	0.502	-0.230
		(0.283)***	(0.228)**	(0.269)
	Lag positive shock x Medium landholding	0.641	-0.523	0.118
		(0.270)**	(0.240)**	(0.269)
	Lag positive shock x Large landholding	0.870	-1.216	-0.346
		(0.321)***	(0.288)***	(0.314)
	Panel B: Full Specification			
1	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive}	0.446	0.599	1.045
		(0.303)	(0.299)**	(0.289)***
2	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive} x Medium landholding	-0.392	-0.170	-0.562
		(0.315)	(0.336)	(0.283)**
3	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive} x Large landholding	-0.580	-0.238	-0.818
		(0.339)*	(0.392)	(0.369)**
4	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought}	-1.506	-0.014	-1.521
		(0.427)***	(0.269)	(0.401)***
5	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought} x Medium landholding	1.192	-0.126	1.066
		(0.439)***	(0.324)	(0.420)**
6	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought} x Large landholding	1.556	0.118	1.673
		(0.448)***	(0.367)	(0.430)***
7	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought}	-0.913	0.322	-0.592
		(0.392)**	(0.326)	(0.401)
8	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought} x Medium landholding	0.812	-0.274	0.538
		(0.439)*	(0.381)	(0.392)
9	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought} x Large landholding	0.938	-1.047	-0.109
		(0.472)**	(0.439)**	(0.495)
10	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero}	-1.234	0.701	-0.533
		(0.444)***	(0.313)**	(0.426)
11	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero} x Medium landholding	0.896	-0.676	0.220
		(0.379)**	(0.333)**	(0.401)
12	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero} x Large landholding	1.328	-1.402	-0.073
		(0.499)***	(0.388)***	(0.491)
Dependent variable mean		3.31	7.06	10.38

Notes:

1. This table tests if household labor supply is consistent with separation failures. Dependent variables are number of worker-days employed as an agricultural laborer (col 1), in own-farm work (col 2), and total labor supply (own farm work plus hired out) (col 3) by the household in the past 7 days. The sample comprises agricultural households with land and excludes lean quarter observations (April-June); it contains 203,073 household-year observations.
2. Shocks are defined as drought, zero, or positive, and correspond to rainfall below the 20th percentile, between the 20th-80th percentiles, and above the 80th percentile, respectively. Medium and large landholding are indicators for whether acres per adult in the household is in the second or third tercile of the sample distribution, respectively.
3. Panel A shows regressions of the dependent variable on an indicator for whether the district experienced a positive shock in the previous year, and an interaction of this indicator with landholding terciles.
4. In Panel B, covariates 1,3,5, and 7 are of the form {$Shock_{t-1}$=X};{$Shock_t$=Y}; they are indicators that equal 1 if the district experienced shock X in the previous year and shock Y in the current year, and equal 0 otherwise. The other covariates are interactions of these sequences of shocks with landholding terciles.
5. Each regression contains year, district, and calendar quarter fixed effects, quadratic functions of the number of males and females in the household. Standard errors are corrected to allow for clustering by region-year.

Panel B repeats this analysis using the full specification. Rows 7 and 10 provide the coefficients of primary interest. When a positive shock is followed by no shock (row 10), households with small landholdings experience an average decrease in external employment of 1.23 days per week. In these years, they increase their supply of labor to their own farm by 0.701 days, or 10% of mean own-farm labor supply. The impacts on households with medium and large landholdings (rows 11-12) are similar to the pattern shown in Panel A.

In addition, when a positive shock is followed by a drought (row 7), households with small farms do not decrease own-farm labor supply, despite the negative productivity shock–the coefficient is positive (though insignificant). However, I cannot reject that this coefficient is equal in magnitude to the effect of a non-positive shock followed by a drought (row 4).

As a whole, these results provide evidence that households respond to rationing by increasing labor supply on their own farms. However, a full test of whether rationing leads the marginal product of labor on small farms to be lower than that on large farms requires farm-level data on total labor inputs. Farms in the bottom tercile of the landholding distribution are quite small and unlikely to hire much labor, so own-farm employment is likely highly correlated with total farm labor use. However, if these farms do hire some external labor, then some of the increase in own-farm supply may be offsetting decreases in labor hired by the farm. Farm-level labor use data is needed for a more complete understanding of how rationing affects the allocation of labor across farms.

1.4.6 Heterogeneity in Wage Rigidity: Crop Variation

Districts exhibit substantial heterogeneity in the extent of rigidity. To test for heterogeneity, in the World Bank panel data, I regress the log nominal wage on the three main categories of shocks (contemporaneous positive shocks, contemporaneous droughts preceded by a non-positive shock, and lag positive shocks followed by a non-positive shock), year fixed effects, district fixed effects, and an interaction of each of the district dummies with the contemporaneous droughts indicator. This is the same specification as in column (3) of Table 1.2 plus the interaction terms. The coefficient on each interaction term provides an estimate of that district's mean wage change to a drought in the sample (relative to the omitted district). If the effect of droughts is the same across all districts, then the coefficients on the interactions should be 0. The F-test of joint significance of the interaction coefficients has a p-value of 0.000, indicating heterogeneity in the extent to which districts respond to droughts. Repeating this analysis by instead interacting each district dummy with the lag positive shock indicator also suggests heterogeneity in the extent to which lag positive shocks influence future wages (the F-test of joint significance of the interaction coefficients has a p-value of 0.000).

Districts in India differ substantially in crops grown. The World Bank dataset contains data on 20 crops, including the percentage of land area in each district-year planted with each crop. Five of these twenty crops—soybeans, sesame, rapseed/mustard, sunflowers, and sugarcane—are extremely sensitive to the amount of labor hired during harvest.[33] For example, if the first three are not harvested quickly upon reaching maturity, their pods burst, spilling their seeds onto the ground and leading to large output losses.[34] In these areas, price flexibility is particularly important because inefficient labor allocation will lead to especially large profit losses.

I investigate whether rigidities are lower in areas where the costs of rigidity are likely to be higher due to crop characteristics. Specifically, I test whether districts with a greater percentage of land area planted with labor-sensitive crops are more responsive to shocks. The crop sensitivity measures were constructed as follows. For each of the five crops, the percentage of land-area planted with the crop in each district-year was regressed on year fixed effects to remove national time trends. The residuals for each of the five regressions were then summed to give the total adjusted percentage of land planted with these crops in each district-year.

[33] These crops were identified in the following manner. A researcher with a background in agricultural extension work in India compiled a timeline of work activities for each of the 20 crops, along with identifying which activities were particularly important for output. This information was based on consultations with an expert at an agricultural research university and numerous field interviews with farmers. The researcher identified these 5 crops as extremely sensitive to timely labor inputs for the reasons listed below. He did not know how this information would be used and did not have access to the World Bank dataset.

[34] Similarly, sunflower seeds will fall to the ground when they become over-ripe. A bigger practical concern, however, is that birds are relentless in eating the seeds as soon as they reach maturity. This poses such a large threat that farmers in richer countries like the US cut sunflowers early and let the seeds ripen indoors, or cover each sunflower head with protective covering to protect it from birds. These practices are not often followed in India, where farming is less capital intensive. Harvesting sunflowers quickly is therefore important for output levels. The constraint on sugarcane is institutional. Each sugarcane mill in India is assigned a command area; all growers within that area are required by law to sell their crop to that mill. To manage supply chains, mills assign farmers a harvest date on which they are allowed to sell their output to the mill; output is not accepted on other dates. Therefore, farmers must ensure their crop is prepared for delivery to the mill on their assigned date.

Table 1.6: Heterogeneity in Wage Rigidity - Sensitivity of Crops to Labor Inputs

Dependent Variable: Log Nominal Wage

		Average over Last 5 Years		Average over Sample	
	District Crop Sensitivity Measure	*% Land with Sensitive Crops*	*% Land with Sensitive Crops is > Median*	*% Land with Sensitive Crops*	*% Land with Sensitive Crops is > Median*
		(1)	(2)	(3)	(4)
	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Zero}	Omitted	Omitted	Omitted	Omitted
1	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive}	0.027	0.020	0.025	0.028
		(0.009)***	(0.010)**	(0.009)***	(0.009)***
2	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive} x District crop sensitivity measure	-0.061	0.014	-0.130	-0.005
		(0.137)	(0.015)	(0.146)	(0.014)
3	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought}	-0.009	0.016	-0.010	0.010
		(0.011)	(0.015)	(0.011)	(0.013)
4	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought} x District crop sensitivity measure	-0.425	-0.045	-0.263	-0.039
		(0.200)**	(0.020)**	(0.184)	(0.018)**
5	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought}	0.044	0.069	0.036	0.033
		(0.021)**	(0.032)**	(0.020)*	(0.021)
6	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought} x District crop sensitivity measure	-0.351	-0.042	-0.148	0.008
		(0.325)	(0.036)	(0.293)	(0.033)
7	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero}	0.022	0.017	0.020	0.028
		(0.010)**	(0.011)	(0.010)*	(0.010)***
8	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero} x District crop sensitivity measure	-0.142	0.010	-0.206	-0.015
		(0.160)	(0.018)	(0.166)	(0.017)
9	District crop sensitivity measure	-0.375	-0.028	--	--
		(0.205)*	(0.012)**		
	Observations: district-years	6,840	6,840	7,296	7,296
	F-test p-value: Coefficient 3 + Coefficient 4 + Coefficient 9 = 0	--	0.001***	--	0.061*
	F-test p-value: Coefficient 5 + Coefficient 6 + Coefficient 9 = 0	--	0.970	--	0.175
	F-test p-value: Coefficient 7 + Coefficient 8 + Coefficient 9 = 0	--	0.962	--	0.398

Notes:

1. This table tests whether districts with crops that are more sensitive to labor inputs have less rigid wages. Observations are from the World Bank dataset. The dependent variable is the log of the district's mean nominal daily wage.
2. Shocks are defined as drought, zero, or positive, and correspond to rainfall below the 20th, between the 20th-80th, and above the 80th percentiles, respectively. Covariates 1, 3, 5, and 7 are of the form {$Shock_{t-1}$=X};{$Shock_t$=Y}; they are indicators that equal 1 if the district experienced shock X in the previous year and shock Y in the current year, and equal 0 otherwise. The other covariates are interactions of these sequences of shocks with district crop sensitivity measures.
3. Crop sensitivity measures capture the % of landarea planted with crops that are highly sensitive to the amount of labor hired (soybeans, sesame, rapseed/mustard, sunflowers, and sugarcane). These measures were constructed as follows. For each of the 5 crops, the % of land planted with the crop in each district-year was regressed on year fixed effects to remove national trends. The residuals for the 5 regressions were then summed for each district-year. Columns (1)-(2) use the district's mean of this sum over the last 5 years. Columns (3)-(4) use the district's mean in the sample as a whole. In both cases, the percentage values from the current year and previous year are excluded when computing means. In columns (1) and (3), the measure is the continuous mean percentage. In columns (2) and (4), it is a binary indicator that equals 1 if the mean percentage is above the sample median and equals 0 otherwise.
4. Each regression contains district and year fixed effects. Standard errors are corrected to allow clustering by region-year.
5. The last 3 rows report p-values of F-tests for sums of coefficients as indicated. For the regressions in columns (3)-(4), coefficient 9 is absorbed by the district fixed effects; this coefficient is therefore not included in the sums in column (4).

Table 1.6 shows results from OLS regressions of the log nominal wage on indicators for each shock category, interactions of each shock category with a measure of the district's crop sensitivity, and year and district fixed effects. The measure in column (1) is the running average of the adjusted percentage over the past 5 years in each district. The current year and previous year are omitted when computing means, to preclude any possibility that current or lag rainfall shocks influence the crop sensitivity measure (via effects on planting decisions). The regression in column (2) shows interactions with a binary version of this measure, which equals 1 if the district's running average over the last 5 years is above the sample median and 0 otherwise.

The results indicate that districts with more labor-sensitive crops are substantially more likely to cut wages in response to droughts (row 4). For example, during droughts, wages are about 4.5% less in districts in the upper half of the distribution in terms of land area planted. The F-test for whether, in these high crop sensitivity districts, wages are the same during droughts as the reference cell has a p-value of 0.001 (reported at the bottom of the table). In addition, I cannot reject that lag positive shocks have no persistent wage effects in these districts.

The regressions in columns (3)-(4) repeat this analysis, but the crop sensitivity measure is constructed by averaging the adjusted percentage over the entire sample for each district. As before, current year and previous year are omitted when computing means. This reflects the time-invariant proportion of land planted with sensitive crops in a district. Column (3) uses the continuous percentage measure, while column (4) uses a binary indicator that equals 1 if the percentage is above the sample median and 0 otherwise. The results are similar to those in columns (1)-(2).

One potential concern with the interpretation of these results is that rainfall shocks may impact the marginal product of labor more in high sensitivity districts—the results could stem from different productivity effects of shocks. Alternately, if workers in these different types of districts have different labor supply elasticities, this may alter the equilibrium wage response. However, wages in high and low sensitivity districts are equally responsive to contemporaneous positive shocks (rows 1-2). But the effects of positive shocks only carry over to future wages in low sensitivity districts. This should not be the case if the results are due to differences in productivity effects or labor supply elasticities.

As a whole, these results are consistent with less rigidity in areas where crops are more sensitive to labor inputs. However, this evidence is only suggestive since crop choice is endogenous. It may be more profitable for farmers to grow labor-sensitive crops in less rigid areas, which could influence their planting decisions.

1.5 Alternate Explanations

Could the results be explained by reasons other than downward nominal wage rigidity? In this section, I discuss potential alternate explanations. I focus on three sets of competing explanations: alternate models of equilibrium unemployment, the possibility that rainfall shocks have persistent effects in future periods, and measurement error.

Efficiency wage models with micro-foundations that do not involve nominal rigidities—such as moral hazard, screening, labor turnover, or nutrition—also predict that wages may remain above market clearing levels in equilibrium. However, these models do not predict rigid wages—they generally predict that wages will decrease when labor productivity declines. For example, none of these models can account for why wages would rise under a positive shock but then not come back down to their prior level once the shock has dissipated, or why this should be influenced by inflation. Similar arguments apply to search friction models that do not incorporate some degree of nominal rigidity.

The second set of concerns is that rainfall shocks are not actually transitory, but have future effects on the economy through channels other than wage rigidity. For example, one potential confound arises if positive shocks have persistent productivity effects in future years. If this is the case, then lag positive shocks could raise future wages because they positively impact the marginal product of labor in the next year. However, in this case, employment should also be higher in the following year, whereas the results indicate substantial employment decreases.

Alternately, inter-temporal substitution of labor could cause an increase in future wages and decrease in employment. However, such labor supply shifters are difficult to reconcile with the inflation results: it is not clear why any of these should be more likely when inflation is lower. These explanations are also inconsistent with the heterogeneity in the employment results by landholding. Reductions in external employment are especially likely for those with less land, and these households respond by *increasing* labor supply on their own farms. It is therefore unlikely that the wage dynamics are driven by a decrease in labor supply in the year after a positive shock. Similar arguments imply that factors such as income effects, migration, or capital investment are also not driving the empirical findings. In addition, such explanations do not account for why wages do not fall in response to droughts.

A third concern is that measurement error—driven for example, by rounding error—could make wages appear more sticky than they are. As discussed above, this is especially a concern when using histograms of wage changes to document rigidity. However, the empirical tests in this paper are less subject to biases from such errors. Specifically, measurement error will only confound the above results if it is correlated in a very particular way with the random rainfall shocks: the error must be more likely in years with negative shocks and in years after positive shocks than in other years. It is not clear why respondents should be differentially more likely to round wages in these two special cases. In addition, if the observed wage persistence associated with lag positive shocks was simply due to reporting errors, we should not observe real employment effects or variation caused by inflation.

1.6 Mechanisms: Survey Evidence on Fairness Norms

The presence of nominal rigidities in markets for casual daily labor is perhaps especially surprising given the lack of institutional constraints in these markets. This suggests that non-institutional mechanisms discussed in the literature—such fairness norms against wage cuts—may play a role in maintaining rigid wages. To explore the relevance of fairness considerations, I conducted a survey in 34 villages in the Indian states of Orissa and Madhya Pradesh. 396 respondents (196 agricultural laborers and 200 landed farmers) were interviewed. Following Kahneman, Knetsch, and Thaler (1986), I presented workers with scenarios about wage setting behavior and asked them to rate the actions described as "Very fair", "Fair", "Unfair", or "Very unfair". Table 1.7 presents the text of these scenarios and reports the percentage of respondents who viewed each scenario as "Very unfair" or "Unfair". Any given respondent was asked only half the questions to prevent the survey from becoming tedious. Some questions involve paired scenarios, which alter the text of the scenario slightly (questions 1A/1B, 3A/3C, and 9A/9B in Table 1.7); for these questions, each respondent was asked only one version of the scenario.

Table 1.7: Fairness Norms in Rural Labor Markets

		Proportion Selecting Unfair			
		Full Sample	Orissa	Madhya Pradesh	Difference
	Panel A: Acceptability of Wage Reductions				
1	A farmer hires a laborer to weed his land for 1 day at a wage of Rs. 120. There is a local factory that pays Rs. 100 per day. One month later, the factory shuts down and many people in the area become unemployed.				
	A) … After this, the farmer decides to do a second weeding and hires the same laborer as before at a wage of Rs. 100.	0.62	0.74	0.49	0.24***
	B) … After this, the farmer decides to do a second weeding and hires one of the newly unemployed laborers at a wage of Rs. 100.	0.55	0.72	0.36	0.36***
2	A farmer usually pays laborers Rs. 120 per day. His son becomes sick and the medical bills are very expensive. He lowers the wage to Rs. 110 per day.	0.79	0.79	0.79	0.00
	Panel B: Money Illusion				
3	Last year, the prevailing wage in a village was Rs. 100 per day. This year, the rains were very bad and so crop yields will be lower than usual.				
	A) … There has been no change in the cost of food and clothing. Farmers decrease this year's wage rate from Rs. 100 to Rs. 95 per day.	0.64	0.81	0.45	0.36***
	B) …. The price of food and clothing has increased so that what used to cost Rs. 100 now costs Rs. 105. Farmers keep this year's wage rate at Rs. 100.	0.38	0.42	0.32	0.10
	C) … The price of food and clothing has increased, so that what used to cost Rs. 100 now costs Rs. 110. Farmers increase this year's wage rate from Rs. 100 to Rs. 105.	0.09	0.08	0.09	0.00
4	A farmer usually pays Rs. 100 per day plus food. There is not much work in the area and many are looking for work. He stops providing food but continues to pay Rs. 100.	0.29	0.24	0.34	-0.11*
	Panel C: Market Clearing Mechanisms				
5	A farmer needs to hire a laborer to plough his land. There is not much work in the area, and 5 laborers want the job. The farmer asks each of them to state the lowest wage at which they are willing to work, and then hires the laborer who stated the lowest wage.	0.61	0.75	0.45	0.29***
6	A farmer needs to hire a laborer to plough his land. The prevailing rate is Rs. 120. The farmer knows there is a laborer who needs money for a family expense and is having difficulty finding work. The farmer offers the job to that laborer at Rs. 110 per day.	0.53	0.65	0.38	0.27***
7	It is harvest time and all farmers pay laborers Rs. 120 per day. One large farmer decides to harvest some of his land immediately and needs to hire 10 laborers. To find enough laborers, he pays them Rs. 150 per day for one week. In the following weeks, he decides to harvest the rest of his land, and re-hires 5 of the laborers at Rs. 120 per day.	0.63	0.81	0.43	0.38***
8	There are 20 landowners in a village. The prevailing wage during plowing is Rs. 120. 10 landowners want to attract extra laborers, and they increase the wage they pay to Rs. 130. The other 10 landowners don't need much labor and maintain the wage at Rs. 120.	0.45	0.60	0.27	0.33***

Table 1.7 (Continued)

Panel D: Fairness Norms and Effort			
9 A farmer needs a laborer to weed his land. The prevailing wage is Rs. 120. There isn't much work in the area and many want the job. A laborer named Balu has family expenses for which he needs money. The farmer knows of Balu's situation, and so he offers him the job at: A) Rs. 120 B) Rs. 100. Given his need for money, Balu accepts the job. How carefully will he do the weeding?			
	More carefully than usual	With the normal amount of care	Less carefully than usual
A) Rs. 120	0.55	0.44	0.01
B) Rs. 100	0.06	0.54	0.40

Notes:
1. This table presents survey evidence on fairness norms. Respondents were asked to rate each scenario as "Very fair", "Fair", "Unfair", or "Very Unfair". The percentage of respondents that selected "Unfair" or "Very Unfair" is shown. Each respondent received half the scenarios. For the paired scenarios (1A/1B, 3A/3C, and 9A/9B), each respondent was asked 1 scenario in each pair.
2. The sample is comprised of 396 respondents (196 casual agricultural laborers and 200 landowning farmers) from 34 villages in the Indian states of Orissa and Madhya Pradesh. All respondents were males aged 20-80. Interviews were conducted July-August 2011.

The states of Orissa and Madhya Pradesh were chosen because they differ greatly in the type of crops grown and other area characteristics. Orissa is poorer with a greater emphasis on staple crops; rice is the dominant crop in the areas surveyed. In contrast, Madhya Pradesh is more affluent; a large portion of the districts in which surveys were conducted are dominated by soybean farming. As discussed in Section 1.4.6, soybean output is substantially more sensitive to the amount of labor hired than rice. The crop heterogeneity results suggest that if fairness norms affect wage setting behavior, norms may be weaker in areas where costs of rigidity are likely to be higher. To check for suggestive evidence along these lines, Table 1.7 also reports responses to each scenario separately for each state along with a test for whether whether the differences are significant.

Panel A establishes baseline norms relating to wage cuts in 2 sets of situations. Question 1 presents a scenario in which a farmer pays a worker Rs. 120 for a task, and then cuts the wage for future work after a factory closure increases local unemployment. 62% of respondents believed it was unfair for the farmer to rehire his old employee at a lower wage, and 55% felt it was unfair for the farmer to hire one of the newly unemployed workers at a lower wage. Note that respondents in Orissa were about 30 percentage points more likely to denote these actions as unfair than respondents in Madhya Pradesh; the t-tests for equality of the means is significant at the 1% level. In Question 2, 79% of respondents indicated that it was unfair for a farmer who was facing personal financial distress to cut the wage of his workers. These perceptions were the same in both states.

Panel B explores the extent to which fairness norms are anchored on the nominal wage rather than the real wage. Question 3 investigates whether respondents are less likely to view real wage cuts as unfair if they do not involve nominal cuts. Respondents were told that last year the prevailing wage was Rs. 100 and that this year real wages are cut by about 5% because a drought will lower yields. However, the 5% real wage cut is presented in three different ways. When the cut consists of a 5% nominal decrease in a period of no inflation, 64% of respondents think it unfair. When the cut consists of no change in the nominal wage during a period of 5% inflation, the percentage viewing it as unfair drops to 38%. Finally, when the cut results from a 5% nominal increase in a period of 10% inflation, only 9% of respondents viewed it as unfair. This pattern is strongly consistent with the idea that workers are averse to nominal (and not necessarily real) wage cuts. Such questions produce comparable responses in other contexts like the US and Canada (Kahneman, Knetsch, and Thaler 1986; Shafir, Diamond, and Tversky 1997). Note that again there is a substantial difference between the 2 states. 81% of Orissa respondents think the nominal wage cut is unfair, whereas 45% of Madhya Pradesh respondents deem it as unfair. In contrast, when real wage cuts do not involve a decrease in the nominal wage, responses from the two states are similar.

Question 4 provides further evidence for the relevance of the nominal wage. When a farmer who pays workers a nominal wage plus food reduces real wages by eliminating food, only 24% of respondents viewed this as unfair. This is sharply lower than the reactions to nominal wage cuts of comparable magnitude in Panel A.[35]

Panel C demonstrates that several wage setting behaviors that are associated with market clearing are at odds with expressed fairness norms. For example, 61% of respondents felt it would be unfair if, during a period of high unemployment, a farmer asks workers for their reservation wage and then offers a job to the worker with the lowest reservation wage (Question 5). Question 7 presents a scenario in which a farmer raises the wage during a period of high labor demand to attract enough workers, and then lowers it again in later weeks when demand is lower. 63% viewed such behavior as unfair. As above, these behaviors violate norms in Orissa much more so than in Madhya Pradesh.

Finally, Panel D investigates whether respondents think worker effort depends on fairness perceptions. Question 9 presents a scenario in which a farmer offers a job to a worker in financial distress. In one version of the question, the farmer offers the prevailing wage rate; this would up-

[35] The value of the food, expressed in the vernacular as high quality food during lunch and other bonuses, exceeds Rs. 10. The magnitude of the real wage cut in question 4 is therefore comparable or greater than the cuts in Panel A.

hold fairness norms and possibly also show benevolence given the laborer's distress. In another version, the farmer sets the wage below the prevailing wage; this strongly violates fairness norms (see Question 6). Among respondents who were told that the wage was set at the prevailing wage, 55% percent believed the worker would exert more effort than usual and only 1% believed he would exert less effort than usual. In sharp contrast, when told the wage was below the prevailing rate, only 6% believed the worker would exert extra effort while 40% believed the worker would exert less effort than usual. This indicates a belief that worker effort responds to violations of fairness norms. Responses to this question were not substantially different in the two states.

As further evidence along these lines, Table 1.8 tabulates responses from survey questions about respondents' views about their own behavior or those of their fellow villagers. For example, when laborers were asked whether they offer to work at a wage below the prevailing rate when they have difficulty finding work (Question 2), only 31% said yes while 47% said no. As before, the differences between Orissa and Madhya Pradesh are stark—76% of Orissa laborers said no while only 14% of Madhya Pradesh laborers said no. Question 3 presents a more extreme scenario—whether the worker would accept a wage cut if he had faced prolonged unemployment and was in urgent need of money. Only 38% of Orissa laborers said yes, while 79% of Madhya Pradesh laborers said yes.

Responses by landowning farmers are consistent with these views. The overwhelming majority state that they have not themselves ever hired a laborer at a wage below the prevailing wage (Question 4).[36] In addition, when farmers are asked if a worker in their village would accept a wage cut if he had faced prolonged unemployment and was in urgent need of money, only 39% say yes. This number is considerably larger for Madhya Pradesh (67%) than Orissa (14%).

Of course, responses to hypothetical scenarios may not reflect the actual actions people take when the stakes are real. However, given the strength of the pattern of results, this evidence lends support to the view that fairness norms are a plausible way in which rigid wages are maintained in village labor markets. In addition, the stark difference in results between Orissa and Madhya Pradesh is consistent with the findings in Table 1.6 that areas with crops that are more sensitive to the amount of labor hired have more flexible wages. They suggest that labor market norms may

[36] For this question, concerns that farmers may not truthfully answer a question about their past hiring behavior are warranted. However, whether the answers are truthful reports of past behavior or are driven by a desire by respondents to show that they conform to norms, at the very least, the results speak to the strength of the norms against wage cuts.

form, at least in part, endogenously in response to local conditions.

Table 1.8: Survey Responses to Employment Scenarios

			Proportion of Responses		
				State Breakup	
			Full Sample	Orissa	Madhya Pradesh
	Panel A: Laborers (N=196)				
1	If a laborer was willing to accept work at a rate lower than the prevailing wage, would he be more likely to obtain work from farmers in the village?	Yes	0.61	0.53	0.70
		Maybe	0.20	0.19	0.22
		No	0.19	0.28	0.09
2	When you have difficulty finding work at the prevailing wage, do you offer to work at a lower wage?	Yes	0.31	0.08	0.58
		Sometimes	0.22	0.16	0.28
		No	0.47	0.76	0.14
3	Suppose the prevailing wage is Rs. 100 per day. You have been unemployed for a long time and are in urgent need of money. If a farmer offers you Rs. 95 for one day of work, would you accept the job?	Yes	0.58	0.38	0.79
		Maybe	0.24	0.38	0.09
		No	0.18	0.23	0.12
	Panel B: Landowners (N=200)				
4	In non-peak periods, have you ever hired a laborer for agricultural work at a wage below the prevailing wage?	Yes	0.05	0.03	0.09
		No	0.95	0.97	0.91
5	Suppose the prevailing non-peak wage rate is Rs. 100. There is a laborer in your village who has been unemployed for a long time and is in urgent need of money. If a farmer offers him Rs. 95 for one day of work, would the laborer accept the job?	Yes	0.39	0.14	0.67
		Maybe	0.25	0.38	0.09
		No	0.37	0.48	0.24
6	Suppose you need to hire a laborer to work during the non-peak period. The prevailing wage is Rs. 100. There is a laborer who would accept the job at Rs. 95 because of money problems. What wage rate would you offer him?	Rs. 95	0.40	0.27	0.54
		Rs. 100	0.60	0.73	0.46

Notes:
1. This table tabulates responses of agricultural workers and employers to survey questions.
2. The sample is comprised of 396 respondents (196 casual agricultural laborers and 200 landowning farmers) from 34 villages in the Indian states of Orissa and Madhya Pradesh. All respondents were males aged 20-80. Interviews were conducted July-August 2011.

1.7 Conclusion

In addition to their broad implications for unemployment and business cycle dynamics, the presence of nominal rigidities in village labor markets has particular relevance for the study of developing country labor markets.

Such rigidities give rise to an additional route through which production volatility (e.g., rainfall shocks) can have adverse consequences for the poor. One focus in the development literature has been that shocks cause shifts in the production frontier, leading to volatility in income; this affects welfare because the poor have limited ability to smooth income across periods. In the presence of wage rigidity, volatility has an additional implication: production may often not be at the frontier because labor markets do not adjust to optimize fully in each period. As implied by the employment results, this means rigidities may lower the levels and increase the volatility of output and income—they may compound the adverse consequences of production volatility.

The fact that those with less land respond to rationing by increasing production on their own farms provides another channel through which rigidities impact efficiency and output. Specifically, it suggests the presence of separation failures in rural labor markets. This is consistent, for example, with the widely documented fact that smaller farms tend to use more labor per acre and have higher yields per acre. This suggests that the distribution of landholdings in poor countries does not have only distributional consequences—it can impact the allocation of labor use in production, and through it, aggregate output.

Finally, the survey results suggest that fairness norms against wage cuts are strong, but they also differ substantially across areas. It is unclear whether such fairness preferences are inherent features of utility or whether they arise endogenously—for example, in response to worker demand for wage stability. The implicit insurance literature has discussed this as a potential source of wage rigidities (see, e.g., Rosen 1985). Insurance demand may be especially relevant given the low income levels in poor countries. In decentralized markets where it is difficult to contract on real wages and explicit contracts are difficult to enforce, fairness norms around nominal wages could be a way to maintain stable real wages. However, it is unclear why workers should be willing to accept employment losses in exchange for wage stability. Ultimately, identifying the cause of nominal rigidities requires better understanding of these factors. Further exploration of fairness norms in labor markets and the underlying mechanisms that give rise to them is a promising direction for future research.

2 Self-Control at Work

(With Michael Kremer and Sendhil Mullainathan)

2.1 Introduction

Agency theory emphasizes a tension between workers and firms: workers do not work as hard as firms would like (Holmstrom 1979; Grossman and Hart 1983). Wages do not reflect the full benefits of work—the employer provides some insurance—and so workers are imperfectly incentivized. Introspection suggests another problem at work: self-control.[1] Looking to the future, agents would like to work hard. Acting in the moment, though, they would rather relax. This raises a new tension: workers do not work as hard as they themselves would like.

This tension on the worker side changes the logic of contract design. A simple example illustrates how. Suppose a firm earns revenue by entering data and faces a penalty C for not completing by a certain date. It hires a worker to enter this data at some wage w and gives her a penalty c for not meeting the deadline. The penalty creates risk for the worker: the data could turn out to be complex to enter so that even at her best effort the worker may not be able to meet the deadline. This increased risk requires compensation. Self-control changes this simple logic. A worker with self-control problems may see benefits to an increased penalty in tomorrow's contract (assuming she is sophisticated as in O'Donoghue and Rabin 1999). An increased penalty will motivate her to work harder tomorrow, which she values today. As a result, she may not need compensation. This generates a striking prediction. Sophisticated workers with self-control problems may prefer a dominated contract, one that pays less for every output realization. In effect, the incentive scheme can be an implicit commitment device (Laibson 1997; Ashraf et. al. 2006; Gine et. al. 2010).

[1] Frederick et al (2002) and DellaVigna 2009 review the self-control literature. Prominent models include Laibson (1996), O'Donoghue and Rabin (1999; 2001), and Fudenberg and Levine (2006). Bernheim, Ray and Yeltekin (2011) and Banerjee and Mullainathan (2009) examine it in the development context. Gul and Pesendorfer (2001; 2004) provide a different account of the demand for commitment.

This observation separates work from several other self-control domains that have been studied (e.g. savings and smoking). An independent agent—the employer—has both the means and motives to reduce worker self control. Sharp monetary incentives or firing for tardiness and inadequate performance can ameliorate self-control problems. This could alter our understanding of work arrangements. Employers can now actually increase productivity beyond what workers could achieve as full residual claimants.[2] This argument makes two key assumptions: self-control at work is quantitatively important and workers are sufficiently sophisticated as to value the implicit self-control benefits of stronger incentives.

We performed a field experiment to test these assumptions.[3] In the experiment, data entry workers are paid weekly according to a piece rate that depends on the number of accurate fields entered each day. Workers were randomized into two conditions. Workers in the Control condition face the standard piece rate of w. Workers in the Choice condition are allowed to choose a target T for the day: if they meet the target they receive the standard piece rate w; if they fail, they receive half the piece rate, w/2. Like the penalty in the fictional example, the target increases penalty without increasing reward. In this sense, a zero target contract dominates every positive target contract. Workers were allowed to choose their targets either the evening before (for the next day) or that morning; they were randomly assigned into these Morning and Evening choice conditions. All contract assignments were randomly assigned daily for each worker.[4] To measure the impact of time horizon, workers were also randomized into different payday groups: all were paid weekly but the exact day of payment varied. These randomizations were at the worker level and paydays once assigned were fixed. Our experiment takes place over 13-months in Mysore, India. Workers were typically high school graduates for whom employment in the firm was their primary source of

[2] Clark (1994) makes this case for the rise of the factory during the industrial revolution. O'Donoghue and Rabin (1999; 2006) formalize how firms use deadlines to motivate a procrastinator. They also produce interesting implications for screening which we examine only briefly here. DellaVigna and Malmendier (2004) also study contract design in a different context. Kaur et al (2010) discuss the work context. Firms could also increase output by solving the free riding problem in team production (Cheung 1969; Alchian and Demsetz 1972).

[3] Ariely and Wertenbroch (2002) and Burger et. al. (2008) provide evidence of self control in work. Though quite interesting, both papers involved student populations and smaller stakes. Shearer (2004), Gneezy and List (2006), Bandiera et al (2007), Fehr and Goette (2008), and Hossain and List (2009) use field experiments to study other features of worker psychology.

[4] In addition to the control and two choice conditions above, we also randomize workers into one of three dominated contracts without choice. While not useful for understanding self-control, these provide useful calibration benefits. This is reported in greater detail below. We also randomized workers seats so that every 1 to 3 weeks they moved to different seats. This allows us to estimate peer effects as discussed in Kaur, Kremer and Mullainathan (2010, 2012).

earnings.

We find three main results. First, workers work harder as the payday gets closer. They earn 8% more on paydays than at the beginning of the weekly pay cycle. This is not concentrated on the payday: production rises smoothly through pay cycle.[5] Second, dominated contracts are chosen on 35% of the worker-days when they are offered (among workers that are present). The ability to choose a dominated contract increases productivity: the Choice condition shows a Local Average Treatment Effect of 6%. Third, these two effects are related. We find significant heterogeneity in the payday effect and this heterogeneity predicts the impact of Choice. Workers with above mean payday effects are 49% more likely to choose the dominated contracts, and show a 20% treatment effect of Choice on output. The workers with the biggest self-control problems appear to value the dominated contracts the most. We argue below that while any one of these findings may be explained in some other way (e.g. workers trying to signal ability to employers), only the self-control interpretation fits these three facts together.

We calibrate these effects against two benchmarks. We compare their magnitudes to a simple OLS estimate of the returns to education: the payday effect is comparable to a little more than a 1-year increase in education and the ability to choose the dominated contract increases output at the same level roughly of two-thirds of a year of education. We also calibrate a simple model to estimate the implied time inconsistency—the extent to which the discount rate changes. This requires mapping out the cost of effort curve, which we do using an exogenous change in the overall piece rate that was implemented after the treatments ended. The payday effect suggests a discount rate of roughly 5% per day. The Choice treatment suggests that the difference in discounting of benefits between the self that chooses the contract and the one that works appears to be at least 18%. For the workers with above average payday effects, the impact of choice rises to roughly 2.5 years of education or a time inconsistency between the chooser and doer of 64%.

Two other findings stand out. First, workers seem to learn about their self-control problems. Early on, many workers experiment with dominated contracts when in the Choice condition. As they gain experience, workers diverge: some choose positive targets more while others choose zero targets more regularly. The workers who increase demand are also those with the highest payday

[5] To avoid selection bias due to selective attendance, we report production results with absent workers coded as zero production. Workers are told the next day's contract assignment before leaving for the day.

effects, suggesting accurate learning. The treatment effect of Choice on productivity also increases with worker experience: as workers sort better, the benefits of choice increase. In contrast, the effect of paydays neither declines nor increases with experience. These results suggest that workers learn about the extent of their self-control problems but do not necessarily learn away these problems (possibly suggesting the availability of external commitment devices to help with self control at work is low).

Second, inconsistent with the simplest self-control models, we fail to find a difference between Morning and Evening choice: targets chosen for the next day are the same as for the same day. Ex post analysis of the data suggests a possible reason: workers may face uncertainty the evening before that will be realized when they arrive at work. Variability in computer speed and time they can reach work (e.g. due to uncertain buses) may offset a greater desire for self control the evening before. In fact, targets are higher the evening before when these measures of uncertainty are low and this reverses when uncertainty is high. Self-control models that allowed for such uncertainty that is revealed over time could easily explain these results but we ourselves did not make such a prediction prior to running the experiment.[6]

Finally, we find mixed evidence of heterogeneous treatment effects. Workers with above average productivity were 40% more likely to demand dominated contracts but interestingly benefited less from their provision. This hints at the possibility that there may be variability in the extent of sophistication. At the end of the experiment, we conducted surveys to elicit subjective measures of self-control at work and in other domains (such as smoking).[7] As a whole, these measures, however, showed little predictive power, either because of a lack of statistical power or because these abstract measures are inherently noisy.

2.2 Model

We present a simple principal-agent model that incorporates worker self-control. Since our empirical work will focus on the *demand* for contracts, we focus only on worker utility under different contracts: we do not explicitly derive the optimal contract here.

[6] We were motivated to do this analysis when we noticed a significant drop in demand for the targets during days of slow network speed. This led us to see whether this uncertainty mediated the evening-morning effect.

[7] For example, agreement with statements like, "Some days I don't work as hard as I would like."

An agent exerts unobservable continuous effort e to produce stochastic binary output, y. Output equals 1 with probability $p(e)$ and equals 0 otherwise. Output is perfectly observable, and the agent is paid in period T as a function of it. Write L and H for the pay in the low and high states and $\Delta = H - L$. Agents exert effort in period 1, output is realized and pay is given in period T, and incentive contracts are signed in period 0.

Agents discount a payoff at horizon τ (i.e. τ periods in the future) by $d(t)$ where $d(t)$ is decreasing in t, $d(t) \leq 1$ and $d(0) = 1$. For the time consistent case, we write $d(t) = \delta^t$ where δ equals the daily discount factor and t is measured in days. For time inconsistent agents, we assume that the impatience for a delay of s periods $\frac{d(t+s)}{d(t)}$ is decreasing in t for any fixed s.[8] These assumptions generate a discount function that matches what is used to model hyperbolic discounting: people are particularly impatient for receiving a payment today versus tomorrow, less impatient for tomorrow versus day after tomorrow and more generally their impatience falls with the horizon. For the empirical work, we further implicitly assume that $d(1)$ is sufficiently less than 1. The agent is risk averse, with concave utility, $u(\bullet)$, over income and a cost of effort $c(e)$ which is convex in e.[9]

We will write utility as $U_{0,1}^{\{C,I\}}$ to indicate the utility for either a time consistent (C) or time inconsistent (I) agent at time 0 or 1. Based on our assumptions:

$$\begin{aligned}
U_1^C &= \delta^T \left[p(e)u(H) + (1 - p(e))u(L)\right] - c(e) \\
U_0^C &= \delta^{T+1} \left[p(e)u(H) + (1 - p(e))u(L)\right] - \delta c(e) \\
U_1^I &= d(T)([p(e)u(H) + (1 - p(e))u(L)] - c(e) \\
U_0^I &= d(T+1)([p(e)u(H) + (1 - p(e))u(L)] - d(1)c(e)
\end{aligned}$$

Note that $U_0^C = \delta U_1^C$ but $U_0^I \neq d(1)U_1^C$.

[8] A hyperbolic discount factor $d(\tau) = (1+\alpha\tau)^{-\gamma/\alpha}$, where α captures deviations from exponential discounting (Lowenstein and Prelec 1992), will satisfy this property. Quasi-hyperbolic discounting (see Laibson 1997) or more broadly, present bias (see Benhabib et al 2007) satisfies these assumptions with one possible caveat. A quasi-hyperbolic function has $d(\tau) = \beta\delta^\tau$ so that $\frac{d(1)}{d(0)} = \beta$ and $\frac{d(\tau+1)}{d(\tau)} = \delta$ for $\tau > 0$. So the rate of time preference is strictly decreasing between $\tau = 0$ and $\tau = 1$, and is the same between other periods. The conflict is purely about now versus later with no conflict in future periods. For our purposes we are interested in conflict even in future horizons: for example we will look not just at a *payday* effect but also affects as the payday gets closer.

[9] In writing this utility function we are assuming that the agent only consumes pay when it is given. If agents have access to perfect credit markets and suffer from no other psychological biases (e.g. mental accounting) then consumption utility would not depend on the actually date of pay.

Optimal effort for the time consistent denoted by is given by the first order condition:

$$\delta^T \left[p'(e_C^*)u'(H) + (1 - p'(e_C^*))u'(L)\right] - c'(e_C^*).$$

For the time inconsistent, desired effort differs between time 0 and 1, is denoted by e_0^* and e_1^*. Since time 1 chooses effort the realized effort is given by the first order condition:

$$d(T)(\left[p'(e_1^*)u'(H) + (1 - p'(e_1^*))u'(L)\right] - c'(e_1^*).$$

The difference in these formulas trivially illustrates the payday effect. As T goes down (a closer payday), for a time consistent agent, since $\delta \approx 1$ there should be no noticeable changes in e_C^*.[10] Since we are assuming that d(1) < 1, e_1^* will decrease in T.

Prediction 2.1: Timing of Compensation

As the lag between effort and compensation decreases, a time inconsistent agent will supply greater effort. In contrast, there will be no noticeable changes in effort provision if the agent is time consistent.

Period 0's desired effort level is given by the first order condition:

$$d(T+1)(\left[p'(e_0^*)u'(H) + (1 - p'(e_0^*))u'(L)\right] - c'(e_0^*).$$

Since period 0 weighs the benefits of effort relative to the costs more heavily than period 1 $\left(\frac{d(T+1)}{d(T)} > \frac{d(T)}{1}\right)$, period 0 desires more effort than is provided: $e_0^* > e_1^*$. This is the heart of the time inconsistency problem.

This inconsistency changes period 0's "demand" for different contracts. Suppose we change the

[10] When presenting results, we return to this assumption and explicitly calibrate the exponential discount rate that is implied by our empirical results.

payment in the low state L. For the time consistent worker at time *0*:

$$\begin{aligned}
\frac{\partial U_0^C}{\partial L} &= \delta^{T+1}\left[1-p(e_C^*)\right]u'(L)+\delta\frac{\partial e_C^*}{\partial L}\left(\delta^T\left[p'(e_C^*)u'(H)+(1-p'(e_C^*))\,u'(L)\right]-c'(e_C^*)\right) \\
&= \delta^{T+1}\left[1-p(e_C^*)\right]u'(L) \\
&> 0
\end{aligned}.$$

This is intuitive: more pay in the low state generates an income effect. The disincentive effect can be disregarded because of the envelope theorem—the agent was already equating marginal cost and benefit of effort. Importantly, period 1 who chooses the effort has the same first order condition as period 0 who is valuing the contract. As a result, a decrease in L lowers utility and an increase raises utility.

A time inconsistent worker has a different perspective. As far as she is concerned, the marginal costs and benefits of effort, as she weighs them, have not been equalized by period 1. So the impact of a change in L on incentives must also be considered:

$$\frac{\partial U_0^I}{\partial L}=d(T+1)\left[1-p(e_1^*)\right]u'(L)+\delta\frac{\partial e_1^*}{\partial L}\left(d(T+1)\left[p'(e_1^*)u'(H)+\left(1-p'(e_C^*)\right)u'(L)\right]-d(1)c'(e_C^*)\right).$$

As before the first term is positive: change L and income changes. The second term—the incentive effect—is negative. This is easiest to see for a reduction in L: lowering pay in the low state increases effort. Since 0 wants more effort and this incentive effect raises utility for 0 and offsets the income effect. Thus, if the agent's self-control problem is sufficiently severe, in period 0 utility may rise when L falls. This means that even holding pay constant in the high state (H) she may be happier with a contract that pays her less in the low state. Thus a *dominated* contract—one that pays less in some states and no more in all states—can improve her utility. This leads to the next prediction.

Prediction 2.2: Demand for Dominated Contracts

A (sophisticated) time inconsistent agent may prefer a dominated contract that increases the marginal returns to effort. In contrast, a time consistent agent would never prefer such a contract. Providing a (sophisticated) time inconsistent agent the option to select a dominated contract will increase effort, output, and earnings.

Note that a direct implication of our model is that there will be a correlation between our first two sets of predictions. Time consistent workers will not be affected by the timing of compensation, nor will they select dominated contracts. In contrast, sophisticated time inconsistent workers would select a dominated contract:[11]

Prediction 2.3: Correlation between Compensation Effect and Demand for Dominated Contracts

An agent that is affected by the timing of compensation will be more likely to select and benefit from a dominated contract.

When will time inconsistent sophisticates be most likely to prefer the dominated contract? Trivially in our model period 0 could choose a dominated contract, whereas period 1 would not. If we expand our model to allow for $s+1$ periods prior to effort the term in front of the incentive effect $\frac{\partial e_1^*}{\partial L}$ would become $\frac{d(T+s+1)}{d(T+s)} - d(1)$ and would therefore be increasing in s. In other words, demand for the dominated contract would increase with the horizon between the choice and effort period. Inconsistent individuals are more likely to choose a dominated contract for next week than for tomorrow.[12]

Prediction 2.4: Horizon of Choice and Demand for Dominated Contracts

A time inconsistent agent will be more likely to prefer a dominated contract farther in advance of the effort period.

[11] Note that paydays should lead to production increases among all time inconsistent workers, whereas only time inconsistent sophisticates should demand commitment. In a model that incorporates naiveté, the naifs would attenuate the correlation towards zero.

[12] Our stylized model predicts that workers will only demand targets in advance of the effort period. However, the continuity of time means that even in the morning of the workday, workers may value targets—for example, because the morning self wants the afternoon self to work hard. Thus, our prediction is simply that there will be *less* demand for the dominated contract in the morning than the evening, not that there will be none. In addition, in a quasi-hyperbolic model, this prediction would be more trivial: there will be demand for commitment for tomorrow but no demand for commitment today. The horizon of commitment would not otherwise since the model is exponential except for today versus tomorrow.

An implication of this observation is that the principal can use the incentive scheme to increase productivity and utility of the period 0 self. Consider the case of a risk neutral agent where $u(w)$ is linear. In this case the optimal incentive scheme for a time consistent agent is $L = 0$ and $H = 1$, that is, the agent owns the output.

From before, we know that for a time inconsistent agent, less effort will be realized under this scheme than 0 would like. Following the previous notation let $e*$ be the effort when the agent owns the output. Suppose the firm perturbs this scheme by x so that:

$$\begin{array}{r} L = -\frac{x}{p(e_1^*)} \\ H = -\frac{x}{1-p(e_1^*)} \end{array}.$$

Since the probability of the low state is this merely increases incentives without producing an income effect (for constant effort). The incentive effect on period 1 is positive. As a result, we can see that this scheme generates positive net benefits to both firm and the worker at time 0. The worker at time 0 would prefer it and the firm would weakly prefer it (this can be made a strong preference by letting the firm share in some of the increased revenues). This leads to the following implication:

Prediction 2.5: Full Incentives

A time consistent agent has maximum productivity when she owns the output. A time inconsistent agent, however, will produce more with an employer providing a different incentive scheme than the agent simply owning the output.

In traditional agency theory, the firm provides insurance. Here we can see the firm can increase productivity as well.

We do not directly test this implication but we see it indirectly when we examine nonlinear (dominated) incentive contracts and examine impacts on productivity and workers' willingness to choose them. The derivation of this prediction implicitly shows another prediction. Time inconsistency generates incentive schemes that give super-normal incentives: workers will be incentivized at a rate that is larger than the actual impact of effort on output.[13]

2.3 Experiment Design

2.3.1 Experimental Context

We test these predictions within an Indian data entry firm in the city of Mysore, located in one of the country's major data entry hubs. In this firm, workers use data entry software to type information from scanned images into fields on their computer screen. To control for quality, we measured accuracy using dual entry of data, with manual checks by separate quality control staff when there were discrepancies. These are standard practices in the data entry industry. Workers were paid piece rates for production as a function of how many accurate fields they entered every day. The specific piece rate schedule depended on the contract assignment (see below) but all contracts were functions of accurate fields entered. At all times, the screen showed the worker total fields entered so far that day.

In this context, the physical production function itself is completely individualistic. This means that incentive schemes do not need to be concerned about production externalities.[14] Moreover, the incentives to work came primarily from the piece rate contracts. There were no penalties for being late or leaving early. There were also few reputational concerns or potential for promotion to drive effort since workers were hired for a fixed duration job as is common in this industry. Of course there may be some residual career concerns and we discuss how this would affect our findings in Section 2.5.

Employees were recruited through the standard procedures used by the firm with which we worked—from the pool of resumes submitted by walk-ins to the firm and solicitations via posters and announcements in surrounding villages. Applicants were required to have completed tenth grade education and be at least eighteen years of age. Employees were hired in order of application. Upon

[13] Both these predictions are seen in O'Donoghue and Rabin (1996). The first is implicit in their Propositions 2 and 3 where the employer is able to increase probability of task completion. The second is a direct consequence of employers in their model willing to give sharper penalties for missing a deadline than project value.

[14] This does not rule out social externalities in production. See Kaur et al. (2010; 2012) for an analysis of those in this context.

joining the firm, workers received about 2 weeks of training. This included technical instruction on the data entry software, the production task, and other aspects of computer usage. They were also trained on the two types of incentive contracts and the four contract treatments. During the initial part of training, workers were paid a flat stipend of Rs. 100/day while they learned the task. Trainees then worked under assignment to the control contract. At the end of the training period, they were assigned to the dominated contract for two days under the low and medium targets, respectively. This gave them the opportunity to observe their production under both types of incentive schemes before beginning contract randomizations.[15]

2.3.2 Treatments

To test Prediction 2.1, we randomized employees into three payday groups—Tuesday, Thursday, and Saturday. One-third of workers were assigned to each group at the beginning of the study, and these assignments determined which day of the week each worker received her full weekly pay. For example, on Tuesday evening of each week, employees in the Tuesday payday group were paid for work completed since the previous Wednesday. Randomly assigning paydays removes other reasons that specific days might impact effort. For example, workers might work less hard on Fridays since they'd like to enjoy their Friday evening out with friends. Alternatively they might work harder on Monday after a weekend's rest. In this design, the same day is a payday for some (randomly chosen) workers and not for others. As a result, we can identify the effect of aligning compensation with effort by comparing production on paydays with production on non-paydays.[16]

To test Prediction 2.2, the demand for dominated incentive schemes, we focus on two types of contracts. The first is a linear "control" contract that paid a piece rate wage of w for each unit of production. The second is a nonlinear "dominated" contract that imposed a production target. Under this latter contract, workers received the piece rate of w if they met the target, but only received w/2 for each entered field if they fell short of the target. As shown in Figure 2.1, the control contract dominates the treatment contract in earnings—for any given production level,

[15] The training period for some workers (particularly those that were the first to joined the project) lasted longer than 2 weeks. However, the structure of the training remained the same, regardless of duration.

[16] In a quasi-hyperbolic model, where time horizon is defined by a day, one can make a sharper prediction: production will only be higher on the payday itself. Of course if the time horizon is longer than a day or if discounting is hyperbolic (and not quasi-hyperbolic) this sharp prediction will fail.

earnings are always weakly higher under the control contract.

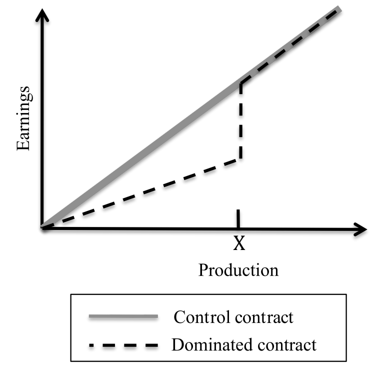

Figure 2.1: Incentive Contracts

Notes: This figure displays the two types of incentive contracts offered to workers. The linear "control" contract paid a piece rate wage of w for each accurate field entered. The nonlinear "dominated" contract imposed a production target, X; workers were paid w for each accurate field if they met the target, but only received w/2 for each field if they fell short of the target. Thus, earnings are equivalent under both contracts for output levels above X. However, if a worker fails to achieve X, earnings are substantially less under the dominated contract.

Every day one quarter of employees were placed in Control and simply received the linear contract. One quarter were placed in Target and were assigned to the dominated contract, with an exogenously chosen production target. The imposed target was selected from three target levels—low, medium, or high.[17] One quarter of workers were assigned to Evening choice, in which they chose their own target the evening before (for the next day). They could always choose a target of 0 (and many did), which is the equivalent of choosing the simple linear contract. One quarter were placed in Morning choice and chose their targets in the morning when they arrived to work. To make the information workers have similar across these conditions, all workers were told their treatment assignment for each day the evening before.

We randomized contract treatments daily at the individual level. We used a balanced design, where every worker received each of the four contract treatments in random order exactly 25 percent of the time over every 8-day or 12-day work period.[18] This ensured that each of the four treatment

[17] For about half the randomization period (mostly in the first half of the study), the Target Assignment treatment consisted of assignment to "low" or "medium" targets only. Assignment to the "high" target was added later, as worker production levels increased. These target levels are explained in greater detail in Appendix C.

[18] During the period when the Target Assignment treatment consisted of assignment to only low or medium targets, randomizations were on an 8-day cycle: 2 Control; 2 Target Assignment (1 low and 1 medium target); 2 Evening Choice; and 2 Morning Choice. During the period when the Target Assignment treatment consisted of assignment to low, medium, and high targets, randomizations were on a 12-day cycle: 3 Control; 3 Target Assignment (1 low, 1 medium, and 1 high target); 3 Evening Choice; and 3 Morning Choice. The proportional weight on each of the four treatments therefore remained 25% at all times.

cells had an equal number of observations, both within each worker and across the entire sample. The vector of assignments was independent across workers. As an example, Appendix Table D.1 displays the contract assignments for 5 workers in the sample over a 24-day period. Daily variation enables us to generate a large number of data points within the study period. This will give us sufficient power for a richer set of analyses, such as looking at trends in behavior over time.[19]

Appendix Table D.2 shows that the 8,673 observations in the Analysis Sample are spread evenly across the treatment cells due to the balanced randomization design; the minor differences in observations are caused by worker turnover—vacancy time until worker replacement, variation in start day from first payday, and the random order of contract assignments.

In addition to the above predictions, our design provides the opportunity to test whether heterogeneity in treatment effects is predicted by correlates of self-control that have been posed elsewhere in the literature. These include preference reversals in estimated discount rates; workers' own assessments of their self-control problems; failed attempts at quitting addictive behaviors; and measures of ability such as productivity, education and IQ. Appendix C provides additional details about the experiment context and protocols.

2.4 Results

2.4.1 Summary Statistics

Panel A of Table 2.1 displays participants' characteristics. Most workers were males (74%). 63 workers reported their age on resumes or elsewhere in their application. These workers ranged in age from 19 to 38 years, with a mean of 24. We collected information on educational attainment and experience during baseline surveys administered to 101 of the 111 workers.[20] Employees had 13 years of education on average. The majority had taken a computer course and had an email address prior to joining the firm.

[19] A drawback of daily randomization is the potential for bias from inter-temporal substitution in effort. We empirically test for this concern in Section 2.4.3. In addition, note that daily randomization reduces the likelihood of potential Hawthorne effects. It is unlikely that workers would persistently alter their behavior each day for a year in response to knowledge of their treatment status.

[20] In this and other information presented in Table 2, some of the employees that were hired in later stages of the project were not surveyed because of clerical oversight.

Table 2.1: Worker Characteristics and Survey Responses

	Mean	Standard Deviation	10th pctile	90th pctile	Obs
	(1)	(2)	(3)	(4)	(5)
A. Worker Characteristics					
Proportion female	0.26	0.44	--	--	111
Age	24	4	20	29	63
Years of education	13	2	11	15	101
Completed high school	0.84	0.37	--	--	101
Used computer prior to joining firm	0.67	0.47	--	--	101
Taken computer training course	0.70	0.46	--	--	101
Had email address prior to joining firm	0.60	0.49	--	--	101
B. Performance on Tests Administered During Experiment					
Contracts comprehension quiz: percentage score	93	13	80	100	79
Raven's Matrix score	35	12	17	49	107
Digit Span composite score	27	6	20	36	109
IQ composite score (Raven's Matrix plus Digit Span)	62	15	40	84	106
C. Discount Rate Measurement					
Proportion of times worker chose smaller immediate reward	0.31	0.28	0	0.67	58
Proportion of times worker displayed preference reversal	0.17	0.23	0	0.67	58
D. Endline Survey: Self-Reported Measures of Self-Control Problems					
Worker agreed or agreed strongly with the statement:					
"Some days I don't work as hard as I would like to."	0.76	0.43	--	--	70
"At the end of the day, I get tempted to leave work earlier than I would like."	0.40	0.49	--	--	70
"I wish I had better attendance at work."	0.86	0.35	--	--	70
"It would be good if there were rules against being absent because it would help me come to work more often."	0.73	0.45	--	--	70
Self-control index: mean of responses to all 9 self-control questions (1=disagree strongly; 5=agree strongly)	3.43	0.55	3.11	4.17	70
Worker has tried to quit an addictive behavior and failed	0.12	0.33	--	--	51
Factor analysis: self-control factor	0.00	0.86	-1.37	0.72	70
E. Endline Survey: Self-Reported Measures of External Constraints					
Worker agreed or agreed strongly with the statement:					
"If I miss one bus or train, the next one I can take is much later."	0.61	0.49	--	--	70
"I do not have much flexibility in how late I can stay in the office because I have to leave at a certain time."	0.56	0.50	--	--	70
Constraints index: mean of responses to all 4 constraints questions (1=disagree strongly; 5=agree strongly)	3.61	0.85	2.25	4.50	70

Notes: This table presents summary statistics for the 111 workers that participated in the study. Panel A presents statistics on worker characteristics, gathered from a baseline survey. Panel B provides information on tests that were administered to workers: a quiz that tested their comprehension of the contract treatments; their performance on the Raven's Matrix IQ test; their total score on the Digit Span test (administered forwards and backwards in each English and the local language); and the sum of their Raven's Matrix and Digit Span score. Panel C describes worker behavior during a discount rate exercise in which they traded 3 sets of off cash awards (Rs. 20 vs. Rs. 24; Rs. 50 vs. Rs. 57; and Rs. 100 vs. Rs. 110) under 2 different horizons: short horizon (the smaller amount today vs. the larger amount in 3 days) and long horizon (the smaller amount in 14 days vs. the larger amount in 17 days). Panel C reports statistics on the proportion of times the worker choose the smaller immediate reward out of the 6 questions, and the number of times the worker showed preference reversal (chose the smaller immediate reward in the short horizon but showed patience by choosing the larger reward in the long horizon). Panel D summarizes responses by workers to questions during the endline survey that asked them to agree or disagree with statements relating to their self-control behavior. It also reports summary statistics for the Self-Control Factor, which is determined using a Factor Analysis on all the endline survey questions. Panel E provides details of responses to 4 endline survey questions that asked about external constraints.

The key outcomes of interest are worker output and demand for dominated contracts. As defined above, output is measured as the number of accurate fields entered in a day. We have 2 measures of demand for dominated contracts under the Choice treatments—an indicator for whether a positive target was chosen and the target level chosen.

Table 2.2: Summary Statistics of Outcome Measures

	Analysis Sample	*Full Payday Sample*
	(1)	(2)
Attendance	0.88	0.88
	(0.33)	(0.32)
Production	5337	5665
	(3404)	(3651)
Production conditional on attendance	6094	6433
	(2935)	(3193)
Indicator for whether positive target was selected under Choice treatment	0.35	--
	(0.48)	
Target level selected under Choice treatment	974	--
	(1502)	
Number of workers in sample	102	111
Number of observations in sample	8,423	11,744

Notes: This table shows summary statistics for four outcome measures: attendance rates, production conditional on attendance (where production is measured as the number of accurate fields entered by a worker in a day), take-up rates for the commitment contract (where take-up is defined as whether a worker selected a positive target when assigned to Choice), and the target levels selected by workers when assigned to Choice (this includes observations where the worker selected a target of 0). These statistics are summarized for each of 2 samples—the 11-month period during which payday treatments were run; the 8-month period during which both the contract and payday treatments were run. The table presents means for each measure and standard deviations are shown in parentheses.

Table 2.2 reports summary statistics for measures of these outcomes. Column (1) provides means and standard deviations for the 8-month period during which the contract randomizations were run. This constitutes the main Analysis Sample—when both contract and payday treatments occurred simultaneously—and is comprised of 102 workers and 8,423 observations. Attendance was 88 percent and mean production conditional on attendance was 6,094 accurate fields per day in the Analysis sample. This amounts to mean daily earnings of Rs. 183 (conditional on attendance). Column (2) reports statistics for the entire 11-month period during which the payday randomizations were run. This constitutes the Full Payday Sample and is comprised of 111 workers and 11,744 observations. While mean attendance was the same in the full payday period, mean production was somewhat higher (6,433 fields). This difference stems from the fact that the payday period ran for 3 additional months, and therefore reflects production increases over time by workers. For consistency of analysis, we will use the Analysis Sample throughout the empirical analysis that follows. To demonstrate that restricting analysis to this sample does not impact the payday results, all tables pertaining to the payday treatments also contain a column showing regression estimates for the Full Payday Sample.

2.4.2 Payday Effects on Production (Prediction 2.1)

We now turn to our first prediction: productivity should increase as the payday increases. We estimate an OLS model of the form:

$$Y_{i,t} = \alpha_0 + \alpha_1 Payday_{i,t} + \alpha_\mathbf{2}\mathbf{W_i} + \alpha_\mathbf{3}\mathbf{D_t} + \alpha_\mathbf{4}\mathbf{S_{i,t}} + \mu_{i,t}. \tag{2.1}$$

$Payday_{i,t}$ is an indicator for whether worker i had a payday on date t. W_i, D_t, and $S_{i,t}$ are vectors of worker, date, and seat assignment dummies, respectively, and are included to increase the precision of the payday coefficient estimate. $\mu_{i,t}$ is the residual error. $Y_{i,t}$ is the number of accurate fields entered by worker i on date t. When workers are absent, this variable is coded as 0.[21] Random assignment of paydays ensures that: $E[Payday_{i,t}\mu_{i,t}] = 0$. α_1 is therefore interpretable as the causal impact of paydays on production.

Columns (1)-(4) of Table 2.3 provide estimates of the payday effect on output in the Analysis Sample. Column (1) estimates the specification in regression model (2.1). It shows that workers produce 215 fields more on average on paydays than non-paydays (significant at the 1% level). There is persistent serial correlation in output, which we control for in column (2), and continue to find a positive and significant effect of the payday. Average output is roughly 5,300 fields. These coefficients (and the rest in the Table) suggest a treatment effect of 2.6% to 10% increased production depending on the specification and whether we are comparing paydays to all days or to the day furthest from the payday.

[21] Simply dropping absentee observations from the sample would produce selection and could bias estimates. In this context, the zero assignment has an economic interpretation: it also corresponds to the workers' earnings that day.

Table 2.3: Treatment Effect of Paydays on Worker Production

	Analysis Sample								*Full Payday Sample*	
	Production	*Production*	*Production*	*Production*	*Attendance*	*Attendance*	*Attendance*	*Attendance*	*Production*	*Attendance*
	OLS	OLS	OLS	OLS	Linear Probability Model	Linear Probability Model	Linear Probability Model	Probit Marginal Effects	OLS	Linear Probability Model
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Payday	215	140	225	428	0.048	0.055	0.077	0.065	414	0.072
	(70)***	(63)**	(68)***	(94)***	(0.009)***	(0.010)***	(0.013)***	(0.009)***	(85)***	(0.011)***
1 day before payday			353	539		0.035	0.053	0.043	469	0.047
			(76)***	(95)***		(0.011)***	(0.013)***	(0.010)***	(85)***	(0.011)***
2 days before payday			207	417		0.016	0.037	0.028	394	0.041
			(94)**	(113)***		(0.014)	(0.016)**	(0.013)*	(103)***	(0.014)***
3 days before payday				374			0.026	0.019	412	0.036
				(112)***			(0.017)	(0.012)	(101)***	(0.014)***
4 days before payday				332			0.047	0.039	254	0.036
				(123)***			(0.017)***	(0.012)***	(105)**	(0.014)***
5 days before payday				176			0.023	0.019	115	0.017
				(119)			(0.017)	(0.013)	(107)	(0.014)
Production on previous workday		0.355	0.353	0.355					0.349	
		(0.016)***	(0.016)***	(0.016)***					(0.013)***	
Production from two workdays ago		0.135	0.136	0.137					0.139	
		(0.015)***	(0.015)***	(0.015)***					(0.012)***	
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8423	8423	8423	8423	8423	8423	8423	7920	11744	11744
R2	0.50	0.59	0.59	0.59	0.11	0.11	0.11	0.14	0.60	0.11
Dependent var mean	5337	5337	5337	5337	0.88	0.88	0.88	0.87	5665	0.88

Notes: This table reports the effects of paydays and distance from paydays on worker production and attendance. Production is defined as the number of accurate fields completed by a worker in a day, and equals zero on days workers are absent. Attendance is an indicator variable that takes the value of 1 if a worker is present and 0 otherwise. Payday is an indicator variable for whether the day was the worker's assigned payday. 1 day before payday to 5 days before payday are indicators for whether the current day is 1 day to 5 days away, respectively, from her assigned payday. Production on previous workday and Production from two workdays ago are lag production controls that measure worker i's production on workdates t-1 and t-2, respectively. Columns (1)-(8) show results from the Analysis Sample—the 8-month period during which both payday and contract randomizations were run. Columns (9)-(10) report results from the Payday Sample—the full 11-month period during which payday randomizations were run. All columns except column (8) show estimates from OLS regressions. Column (8) gives results from a probit regression where the dependent variable is attendance. Estimated discrete changes for each dummy variable are reported. All regressions include fixed effects for each date in the sample, each worker in the sample, and each computer seating assignment. Robust standard errors are reported in parentheses.

In columns (3) and (4) we examine in detail the dynamics of the payday cycle. Instead of a single dummy for the payday, we include dummies for each of the days leading up to the payday:

$$Y_{i,t} = \beta_0 + \sum_k \beta_1^k Payday_{i,t}^{-k} + \beta_\mathbf{2}\mathbf{W_i} + \beta_\mathbf{3}\mathbf{D_t} + \beta_\mathbf{4}\mathbf{S_{i,t}} + \beta_5 Y_{i,t-1} + \beta_6 Y_{i,t-2} + \mu_{i,t}, \qquad (2.2)$$

where $Payday_{i,t}^{-k}$ are dummies indicating that the payday is k days away and $Y_{i,t-1}$ and $Y_{i,t-2}$ are lag production controls.[22] Figure 2.2 graphs the regression coefficients from column (4) of Table 5. The days that immediately follow a payday, and are therefore furthest from the next payday, are when employees are least productive. Production then rises steadily through the pay cycle.[23]

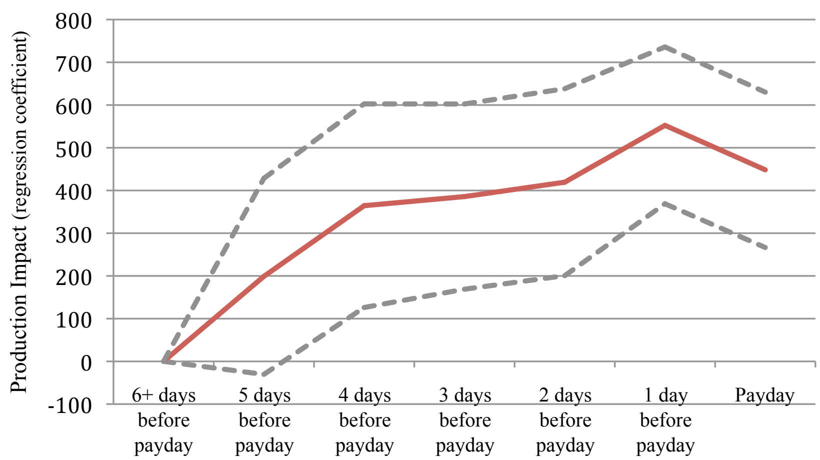

Figure 2.2: Production over the Pay Cycle

Notes: This figure graphs the coefficients and confidence intervals from a regression of production on 6 binary indicators that capture distance from a worker's next payday (payday, 1 day before payday, 2 days before payday, etc). The regression includes controls for lagged production as well worker, date, and seat assignment fixed effects. Confidence intervals are based on 1 Production Impact (regression coefficient) robust standard errors. Note these coefficients correspond to those shown in column (4) of Table 5.

[22] We will look at earnings as the primary outcome. In Panel A of Appendix Table D.3 we show how our results decompose into hours and productivity conditional on hours worked.

[23] Figure 2.2 shows that production steadily increases as the payday approaches. It does not, however, pin down the shape of this increase—one could fit a linear, convex, or concave curve through the confidence intervals. Some time inconsistency models predict convexity. At the extreme, β-δ and dual-self models where time periods are defined as days would predict that all the effect should be concentrated on the payday itself; our results would seem to refute this special case. However, this is not necessarily true under alternate ways of conceptualizing time periods or the horizon of β and δ. More generally, predictions in any model will be sensitive to how the discount function is defined and how time periods are specified. We are interested in testing core predictions that are common across self-control models. We therefore focus here on the qualitative prediction that production should be higher closer to paydays—a result that bears out strongly in the data.

Since our output measures includes both attendance and productivity while at work, we separately examine the impact on attendance using both the Analysis (columns 5 to 8) and the Full Sample (9 and 10) and varying the inclusion of controls. All are linear probability models except for column 8, which uses a probit estimator. Unsurprisingly, people are more likely to show up to work on the payday (by 4.8 percentage points); if nothing else, the benefit of collecting one's pay should increase attendance. In columns (6) and (7), however, we see this "check collection" motive is an incomplete explanation: attendance increases steadily in the days before as well. This is consistent with our model since the benefits of showing up at work will appear larger when the payday is closer. This effect is magnified by the fact that the workers know they will also earn more because they work harder once at work.

These payday effects might imply that paying people daily will increase total output. But in what we have shown some (or all) of the increased effort on paydays may reflect a substitution of effort from other days. Such substitution will happen whenever the cost of effort is not separable across days; for example, workers may be tired from prior days' work effort or may be less motivated if they had high income recently. In Section (2.4.3) below, we test directly for these substitution effects. Our results suggest that in fact effort across days is independent. While this hints that more frequent payments would increase outputs (that the payday effects reflect net new output), one should keep in mind that we have no direct experimental data on payment frequency. Even if this were true, there might be several reasons for not paying workers daily. In our particular case, for example, there are costs of making more frequent payment that are substantive. Cash management poses significant costs to the employer. On a different note, even if more frequent payment solves the work self control problem, it may exacerbate the consumption self control problem. Specifically infrequent payment may be an implicit savings commitment device.

Calibration of Effect Sizes

How can we understand these effect sizes? To calibrate these effects, throughout the paper we will compare the production increases to two other values. The first calibration is straightforward. An OLS estimate in our data finds a return to an additional year of education of 501 fields, or 9%

of mean production. Calibrated in this way, these results suggest that the timing of pay alone has the equivalent of 29% to 112% of the impact of a year of education.

A more theoretically insightful calibration would be to use our payday coefficients to calibrate the implied discount rate. Let e_P be the effort on paydays and e_N be the effort on non-paydays; P refers to the payday time index and N is the index in time for the non-payday. Write π'(e) to be the marginal return to the worker of a unit of effort. In the model above π'(e) = p'(e)[u'(H) – u'(L)]. Since the subjects face continuous output and a linear pay scheme, the original model does not match this context. Since we have a linear pay scheme for the experiment we can model π'(e) as u'(•)k where k is the coefficient on the linear pay scheme.

From the first order condition above we can write:

$$\begin{array}{c} d(P)\pi'(e_P) = c'(e_P) \\ d(N)\pi'(e_N) = c'(e_N) \end{array}.$$

Assuming a linear approximation for utility of money, for small changes we can write:

$$[d(P) - d(N)]\,\pi'(e_P) \approx c'(e_P) - c'(e_N).$$

Our data only gives us the output impact of the payday. This formula suggests that converting the output impact into an estimate of the discount rate ($d(P)$) requires knowledge about the marginal cost of effort. Producing this requires another treatment. The easiest way to do this would be to use variation in the linear pay scheme. For this reason, at the end of the study (after contract randomizations were finished), we randomized workers into two piece-rate wages: Rs. 3 (their usual wage) and Rs. 4 per 100 accurate fields. Each worker received each piece rate 5 times over a 10-day period in random order, with approximately half the workers in the office assigned to each wage within any given day. If we write these two pay schemes as $\pi_1(e)$ and $\pi_2(e)$ we can write:

$$\begin{array}{rcl} d(N)\pi_1'(e_1) & = & c'(e_1) \\ d(N)\pi_2'(e_2) & = & c'(e_2) \\ \Rightarrow d(N)\left[\pi_2'(e_2) - \pi_1'(e_1)\right] & = & c'(e_2) - c'(e_1) \end{array}.$$

We find that the higher piece rate increases productivity by 11%. Thus a 33% increase in wages increases output by 11% or an elasticity of 0.33. By assuming a constant elasticity, we can use the change in pay impact to calibrate the output change due to the payday (or any intervention). Assuming e_2 is chosen in this way we can write:

$$\begin{array}{rrcl} & [d(P)-d(N)]\,\pi'(e_P) & \approx & d(N)\,[\pi_2'(e_2)-\pi_1'(e_1)] \\ \Rightarrow & \frac{d(P)-d(N)}{d(N)} & \approx & \frac{\pi_2'(e_2)-\pi_1'(e_1)}{\pi_1'(e_1)} \end{array}.$$

This tells us that the percentage change in discount rate can be estimated by the percentage change in slope needed to get the same output effect as the payday scheme. In Table 2.3, the production difference between paydays and the beginning of the pay cycle is 428 fields, or 8%. Thus, the production increase on paydays is roughly comparable to the impact of raising the piece rate by 24% between the beginning and end of the pay week. This suggests that the discount factor changes 24% between non-paydays and paydays. If we use the full estimation in Figure 2.2, this suggests that relative to day 0 (the day after payday), the discount factor rises by approximately 5% each day over the pay week cycle.

As a whole, these results are consistent with Prediction 2.1 above. Might they be driven by other factors as well? First, perhaps the results reflect natural impatience as alluded to earlier. Since payments are made weekly, the maximum gap between paydays and non-paydays is about 6 days and the average gap is 3 days. Could a reasonable discount factor produce the increase in production of 8% on paydays? Recall that for a time consistent individual, the marginal cost of effort equals δ^T times the marginal benefit of effort. We know from the calibration above, that the payday increase is equivalent to a 24% increase in the piece rate or the marginal benefit of effort. This suggests that $\delta^3 \approx 1/(1.24) \approx .81$. This implies a daily discount rate of 7 percent, which is highly implausible. This calculus also rules out variants of impatience, such as credit constraints. Second, suppose that there are transaction costs to simply showing up for work. Workers may then appear to have higher output on paydays simply because they are more likely to attend. We cannot address this problem by simply examining productivity conditional on attendance (as opposed to total output) on paydays since those who attend may be a selected sample.[24] The results in column (6), however, provide direct evidence that the 'show up to collect the paycheck effect' does not drive our results: we find increased productivity in the days before the payday. These pre-payday effects could not be explained by a transaction cost of attendance. Third, perhaps the payday effect is not about higher production on the payday but lower production on the day after. Perhaps people take the day after off. Again the dynamics belie this explanation: we find an increase in the days leading up to the payday and the payday effect is neither concentrated on the payday nor the day after. Finally, and perhaps most importantly, suppose that consumption and work effort are complements. People might work harder on paydays because they look forward to consumption that day. A night out with friends might be more enjoyable after a hard day at work. Once again, it is hard to see how this would explain the pre-payday effects; one would require the complementarities with consumption to depend (in a decreasing way) on prior days' work effort. Finally, rather than looking to rule out explanations of the payday effect in isolation, one should examine how they fit the full set of facts, a point we return to in Section 2.5.[25]

Finally, testing Prediction 2.3 involves correlating heterogeneity in the payday effect with demand for dominated contracts. For this test to be valid, there must be heterogeneity in payday effects. We test this directly by interacting a worker fixed effect with the payday dummy:

$$Y_{i,t} = \phi_0 + \phi_1 Payday_{i,t} + \phi_{\mathbf{2}}\mathbf{W_i} + \phi_{\mathbf{3}} Payday_{i,t} * \mathbf{W_i} + \phi_{\mathbf{4}}\mathbf{D_t} + \phi_{\mathbf{5}}\mathbf{S_{i,t}} + \phi_6 Y_{i,t-1} + \phi_7 Y_{i,t-2} + \mu_{i,t}. \quad (2.3)$$

$Payday_{i,t} * W_i$ is a vector of interactions between the vector of worker dummies and the payday dummy. All other variables are as defined above. An F-test for joint significance of the coefficients

[24] In Appendix Table D.4, we show regression results where the dependent variable is production conditional on attendance. However, these results are difficult to interpret. Since attendance is higher closer to paydays due to the self-control benefits of paydays, the composition of workers in attendance is correlated with distance from payday. For example, when workers face negative productivity shocks like sickness, they may be more likely to come into work on paydays than non-paydays (i.e. the higher cost of effort is more likely to be justified when benefits of effort are more immediate). Alternately, if low ability workers also have greater self-control problems, the mean ability of the worker pool will be lower on paydays. Such selection problems could lead average production, conditional on attendance to be lower on paydays than non-paydays because different groups of workers are being compared across days. This also applies to columns (5)-(8) of Appendix Table D.3, where we condition outcomes on attendance.

[25] We find another piece of evidence consistent with our model. In India, festivals involve large expenditures by households (Banerjee and Duflo 2006). Under convex effort costs, time consistent workers should not show large production spikes in the days leading up to festivals (which are perfectly foreseeable); time inconsistent workers, however, would be expected to show such spikes. Indeed, we find that average production increases by 15% in the week prior to major festivals (significant at 1%).

ϕ_3 has a p-value < 0.001, suggesting significant heterogeneity in the payday effect.

2.4.3 Demand for and Treatment Effects of Dominated Contracts (Prediction 2.2)

For the contract treatments, we first analyze take-up of the dominated contracts. First, we focus on days when the worker was present both the day before and the day of the Choice assignment. Absent workers would not have a chance to choose for that day in Morning Choice or the day after in Evening choice. Measured in this way 35% of the Analysis Sample chooses a positive target (see Appendix Table D.5). Since this constitutes a selected sample, we also analyze take-up for all 4,193 Choice observations. We define target choice to be 0 if a worker was absent the day of Choice assignment. This is sensible if we think that not showing up to work indicates a preference for a 0 production level. For consistency, we also define target choice to be 0 if a worker was absent the day before Choice assignment, and therefore did not receive notice of assignment as per protocols and could not select a target if assigned to Evening Choice. These conventions provide a lower bound on the level of demand for dominated contracts by workers. Under this definition, the take-up rate across observations is 28%, as is the mean of the workers' take-up rates.

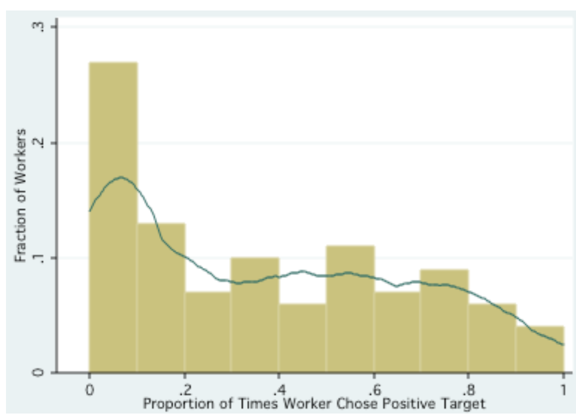

Figure 2.3: Take-up of Dominated Contracts - Distribution of Worker Means

Notes: This figure shows the distribution of take-up rates of the dominated contract. A worker's take-up rate is the proportion of times the worker selected a positive target when assigned to Choice (and present the day before and day of Choice assignment). The distribution is shown for the 101 workers in the Analysis Sample that were assigned to Choice at least once.

Figure 2.3 plots a histogram and kernel density estimate of worker take-up rates. The figure reveals substantial variation in demand for targets. Some workers (16% of the sample) always chose a target of 0. The bottom quarter of the distribution chose positive targets less than 10% of the time. The top quarter of the distribution chose positive targets at least 60% of the time.

To test whether these targets are binding, we examine their impact on production by estimating:

$$Y_{i,t} = \lambda_0 + \lambda_1 Choice_{i,t} + \lambda_2 Target_{i,t} + \lambda_3 W_i + \lambda_4 D_t + \lambda_5 S_{i,t} + \lambda_6 Y_{i,t-1} + \lambda_7 Y_{i,t-2} + \mu_{i,t} \quad (2.4)$$

$Choice_{i,t}$ is an indicator for whether worker i was assigned to one of the Choice treatments on day t; $Target_{i,t}$ is an indicator for the Target Assignment treatment. As before, $Y_{i,t}$ measures production, W_i, D_t, and $S_{i,t}$ are vectors of worker, date, and seating assignment dummies, respectively, and $Y_{i,t-1}$ and $Y_{i,t-2}$ are controls for lagged production. Due to random assignment, $E[Choice_{i,t}\mu_{i,t}] = E[Target_{i,t}\mu_{i,t}] = 0$. The key coefficient of interest is λ_1. It represents the Intent to Treat estimate of giving workers the option to take-up the dominated contract.

In columns (1)-(4) of Table 2.4, we estimate variants of the above regression model for the Analysis Sample. Column (1) shows Choice increased production by 111 fields on average (2% of mean production). This effect is significant at the 10% level. Being assigned to the low target did not significantly increase production in the sample overall. In contrast, assignment to the medium and high targets led to average production increases of 213 fields (4% of mean production) and 335 fields (6% of mean production) respectively. These estimates are significant at the 5% level.

Column (2) separately estimates Evening and Morning choice: Evening choice increases output by 150 fields (3% of mean production, significant at the 5% level); Morning Choice increases output by 73 fields, but is not significant. Columns (3) and (4) limit the analysis to those observations in which workers were assigned to the Control or Choice treatments—the Target Assignment observations are excluded—and show similar results.

Columns (5)-(8) examine the impact of contract assignments on attendance using a linear probability model. Unlike the payday treatment, the contract treatments do not appear to impact whether employees show up to work on average.

The Choice treatments increase production by about 2%, implying a local average treatment effect of approximately 6%. Using the piece rate treatment calibration logic from above, we can lower bound the time inconsistency at 18%—that is, the difference in discounting of benefits between the self that chooses the contract and the one that works to be at least 18%. This is a lower bound since the dominated contract may not achieve optimal effort (from the perspective of today's self).

Table 2.4: Treatment Effects of Contract Assignment on Worker Production

	Dependent variable: Production				*Dependent variable: Attendance*			
Observations	All obs	All obs	Control & Choice	Control & Choice	All obs	All obs	Control & Choice	Control & Choice
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Assignment to choice	111		120		0.007		0.007	
	(59)*		(59)**		(0.009)		(0.009)	
Assignment to evening choice		150		156		0.01		0.014
		(69)**		(69)**		(0.01)		(0.010)
Assignment to morning choice		73		84		-0.00		0.001
		(69)		(69)		(0.01)		(0.010)
Assignment to low target	3	3			-0.002	-0.00		
	(90)	(90)			(0.013)	(0.01)		
Assignment to medium target	213	213			-0.006	-0.01		
	(91)**	(91)**			(0.013)	(0.01)		
Assignment to high target	335	334			-0.005	-0.01		
	(150)**	(150)**			(0.019)	(0.02)		
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lag production controls	Yes	Yes	Yes	Yes	No	No	No	No
Observations	8423	8423	6310	6310	8423	8423	6310	6310
R2	0.59	0.59	0.60	0.60	0.15	0.15	0.11	0.11
Dependent variable mean	5337	5337	5311	5311	0.88	0.88	0.88	0.88
Proportion choosing a positive target (cond'l on attendance)	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35
Proportion choosing a positive target (target=0 when absent)	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28

Notes: This table estimates the treatments effects of the contracts on production. The dependent variable in columns (1)-(4) is production. Production is defined as the number of accurate fields completed by the worker in a day, and equals zero on days workers are absent. The dependent variable in columns (5)-(8) is attendance. Attendance is a binary variable that takes the value of 1 when workers are present and 0 otherwise. The regressions in columns (1)-(2) and (5)-(6) include all observations in the Analysis sample. The regressions in columns (3)-(4) and (7)-(8) include observations in which workers were assigned to the Control contract or to one of the Choice treatments (Evening choice or Morning choice). Each column shows the estimates from an OLS regression of the dependent variable on indicators for contract treatment. Assignment to choice is an indicator that equals 1 if the worker was assigned to one of the Choice treatments (Evening Choice or Morning Choice) and equals 0 otherwise. All regressions include worker, date, and seat fixed effects. In addition, columns (1)-(4) also include controls for lagged production. Robust standard errors are reported in parentheses.

The bottom 2 rows present the proportion of times workers chose a positive target when assigned to a Choice treatment. The first row limits the estimate to observations in which a worker was present the day of and day before Choice assignment. The second row includes all Choice observations and codes target choice as 0 if the worker was absent the day of or day before Choice assignment. Standard deviations are reported in parentheses.

Our rich production data allows us in principle to better understand the nature of the targets. First, we calibrate how large the targets are. The average target chosen is 974 fields. To interpret this, we compute the probability that a worker would have failed to meet the chosen target *if her output distribution matched the distribution under control*: this counterfactual represents what the period 1 self would produce if not tied to a positive target by the period 0 self. Specifically, for observations where workers were in attendance, we estimate a regression of production on worker, date, and computer fixed effects; lag production controls; payday distance dummies; contract assignment dummies; and log experience. For each of the 1,168 observations in which a worker was assigned to Choice, selected a positive target, and was present, we predict the worker's production under the control contract on that day using the estimates from the above regression. To this predicted value, we add the worker's vector of residuals from the above regression to arrive at a vector of potential production values, which we fit to a lognormal distribution. Evaluating the CDF of this distribution at the chosen target level provides an estimate of the probability that the worker would have missed her chosen target under the control contract.

Row 1 of Appendix Table D.6 displays the mean of worker averages for this statistic. Workers' average probability that they would have missed their selected targets if they had been assigned to the control contract is 9.1% in the sample as a whole. For about 60% of workers, the mean target miss probability was 6% or less. The top quarter of the distribution selected targets associated with 16% or higher probability of incurring the penalty. In Row 2 of the table, we report the proportion of times workers actually failed to reach their targets under Choice (conditional on choosing a positive target). The mean worker missed her chosen targets 2.6% of the time. Columns (2)-(3) of the table display these statistics separately for high and low payday impact workers. They indicate that the workers that are most affected by paydays also select considerably more aggressive targets. For high impact workers, the average probability of missing selected targets under control is 11.8%—this is 60% higher than the average miss probability of low impact workers. In addition, while low impact workers rarely actually fail to reach their chosen targets, high impact workers miss them over 5% of the time on average.

Interpreting the aggressiveness of these targets is difficult. When deciding on a target level, time inconsistent workers would weigh the motivational benefits to their future selves against the probability of incurring these costs. The penalty for missing one's target is substantial: half of one's piece rate earnings for the day. If shocks generate uncertainty in output (see Section 2.4.5), then choosing overly aggressive targets can be extremely costly—either due to the financial penalty, or from having to achieve the target even on days when the cost of effort turns out to be high. One might think that the model would give guidance as to the magnitude of targets we might expect. But the marginal cost of effort is unobservable, making it hard to make these predictions. For example, if the penalty is large enough, ex post workers will always stretch to meet the goal. But ex ante some of this stretching is inefficient.

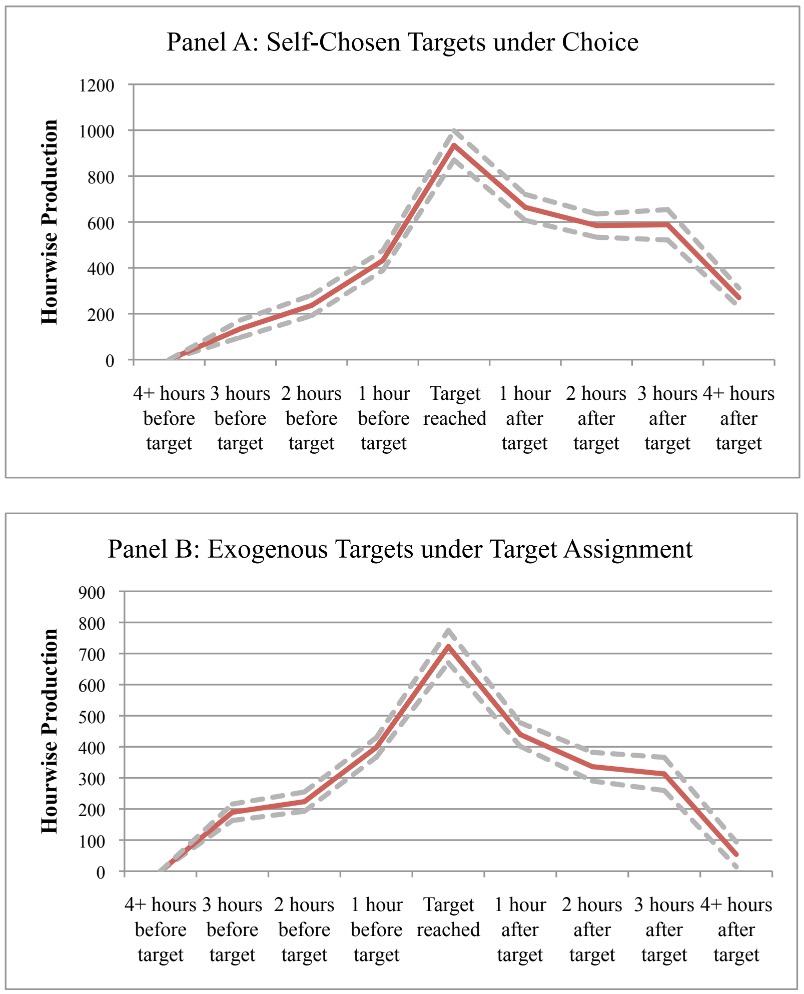

Figure 2.4: Production Behavior Around Target Thresholds

Notes: This figure graphs the coefficients and associated 95% confidence intervals from regressions of hour-wise production (production by worker i on date t in calendar hour h) on binary indicators of distance from when a worker achieved her target within a day, and controls for worker, date, and seat assignment fixed effects. Standard errors are corrected to allow for clustering by worker-day. Panel A reports estimates from days when workers selected positive targets under Assignment to Choice and were present. Panel B reports estimates from days when workers were assigned to exogenous targets under Target Assignment and were present. The coefficients plotted in Panels A and B correspond to those shown in columns (2) and (5), respectively, of Appendix Table D.7.

One might also imagine that the targets would imply a bunching of production around the target level. In fact, we find no such bunching. This is also true, however, for the exogenously assigned targets in the Target treatments. To understand this behavior, we use the hourly data (where hours are defined as calendar hours such as 8-9 am, 9-10 am, etc) to paint a qualitative picture of how production changes as workers approach (and surpass) their targets. Of course since targets are set endogenously, these numbers cannot be interpreted causally but are informative nonetheless. In column (1) of Appendix Table D.7, we regress hourly production on a dummy that equals 1 for the calendar hour in which the worker hit her target that day and 0 otherwise, along with worker, date, and computer fixed effects.[26] In columns (2) and (3), we add a series of dummies for the hours immediately before and after the target was hit; the omitted category is 4 or more hours before the target was reached. These coefficients are graphed in Panel A of Figure 2.4. In all specifications, we see that output increases as workers reach the target and then falls off. We can reject that production in the hour when the target was reached equals production in the hour after it was reached at the 1% significance level.[27] Interestingly, the fall off, though significant, is not huge. Workers continue to work even after reaching the target. The target (self-imposed or externally imposed) appears to serve as a way to get workers working early in the day, a rhythm they keep up even after they hit the target.

As discussed in the payday results above, our interpretation of treatment effects estimates relies on the assumption that labor supply is separable across periods. In the case of the contract treatments, if hard work under targets increases the cost of effort on future days, this would change our interpretation of the estimates in Table 2.4. To test for this concern, we exploit the random ordering of contract treatment assignment. If effort costs are not independent across periods, then today's

[26] Production is defined as 0 in hours when workers did not work. In cases where a worker did not manage to reach her target during the day, the hour when she reached it is coded at 7 pm (which is when the office closed for the day). Note that there are compositional issues in the sample of which the distance from the target dummies are estimated. For example, workers who reach their targets at 5 pm are never observed 3 hours after the target is reached—the value of the 3 hours after target dummy will always be 0 in these cases. This raises selection problems that make it difficult to interpret the coefficients on the distance dummies. The comparison in which we are most interested is between the hour when the target is reached and the hour right after. We observe workers in the hour after the target is reached in 98.7% of cases; so while the selection problem undermines the validity of this comparison, it is unlikely to drive the difference in the coefficients. In Appendix Figure D.3, we graph the proportion of worker-days observed in each distance category.

[27] As a benchmark, we repeat this analysis for the Assignment to Target observations in columns (4)-(6) of Appendix Table D.7 and Panel B of Figure 2.4. We find similar patterns in how workers behave around the point at which they hit the targets we exogenously impose on them.

production should be lower if the worker was assigned to a high effort treatment last period (such as choice or assignment to the high target). In Appendix Table D.8, we explore various specifications in line with this approach and find no evidence that there is dependence in effort between periods. For example, in column (1), we regress production on dummies for yesterday's contract assignment. We cannot reject that being assigned to choice or a target (relative to being assigned to control) has no impact on the next day's production. In fact, the coefficient estimates in all specifications are usually positive (though insignificant). In contrast, if there were intertemporal substitution, these coefficients should be negative. The positive (though insignificant) coefficient could itself be of interest. It could reflect some type of habit-formation in which working becomes easier if you have been working recently.

Another potential concern arises from the block randomization design of the contract treatments. Since workers are assigned to each treatment a fixed number of times within each 8- or 12-day period, treatment assignment on a given day is correlated with the probability of future treatments within each block. For example, conditional on receiving Choice today, a worker is less likely to receive Choice (and more likely to receive the Control contract) tomorrow. This could induce a mechanical correlation that affects what is being captured by the coefficients on the treatment assignment dummies—a concern that would not arise under independent randomization. We test for this concern in Appendix Table D.9. For each observation in the Analysis sample, we compute the probabilities of receiving each contract assignment in that worker-day; these probabilities are determined by the worker's previous assignments in that randomization block. We then directly control for these probabilities in a regression of production on the contract dummies. The results indicate that the assignment probabilities have little predictive power and their inclusion has little impact on the estimated treatment effects. This is confirmed by an F-test of joint significance of the probability controls—the test p-value is 0.45.

Finally, as we did for the payday treatment effects, we check for heterogeneity in treatment effects of Choice. Using only Control and Choice observations, we regress production on: a dummy for assignment to Choice; worker fixed effects; interactions of each worker fixed effect with the Choice dummy; and date fixed effects, computer fixed effects, and lag production controls. The p-value of the F-test of joint significance of the interaction coefficients is 0.003. We interpret this as evidence for heterogeneity in the individual treatment effects of Choice.

2.4.4 Correlation between Payday and Contract Effects (Prediction 2.3)

The payday and contract results each support the predictions of time inconsistency models. As noted above, we also see substantial heterogeneity among workers in the payday and contract effects. We now explore this heterogeneity by checking whether the two sets of effects are correlated—whether those that are most affected by the payday treatments are also those that select and derive the greatest benefits from the dominated contracts.

We define the payday impact measure as: [(mean production on paydays)−(mean production on non-paydays)]/(mean production in sample). We compute this differential for each worker in the Analysis Sample, using only those observations in which the worker was assigned to the Control contract. We then define a worker as having a high payday impact if her differential is above the mean differential in the sample.[28]

In Table 2.5, we test whether workers with higher payday impacts are more likely to demand the dominated contract. We use two outcome variables to measure take-up: the target level chosen (which includes targets choices of 0) and a binary indicator that equals 1 if the worker chose a positive target. Following the conventions described in Section 2.4.3, in this and future regressions, we define both take-up dependent variables as 0 on days that the worker was absent the day before or day of assignment to Choice. In each column, we regress the dependent variable on the high payday differential dummy and controls. On average, workers that are more affected by paydays select a target that is 353 fields higher and are 13.8 percentage points more likely to select a positive target. These coefficients correspond to a striking 47% and 49% of the mean target level and take-up rate, respectively, and are both significant at the 1% level.

In Table 2.6, we explore whether workers with high payday impacts derive more benefit from the contracts. In column (1), we provide the estimates of the average treatment effects of Assignment

[28] In the results presented above, we also see production increases in the days leading up to paydays. We use the payday-nonpayday difference for simplicity. In addition, since we can only compute this statistic for workers that were assigned to the Control contract on both paydays and non-paydays during their employment, it cannot be computed for some workers that were in the sample for shorter periods of time. This reduces our sample for this analysis from 8,423 to 8,240 observations.

to Choice and Assignment to a Target for reference. In column (2), we estimate:

$$
\begin{aligned}
Y_{i,t} = \quad & \theta_0 + \theta_1 Choice_{i,t} + \theta_2 Choice_{i,t} * HighImpact_i + \theta_3 Target_{i,t} + \theta_4 Target_{i,t} * HighImpact_i \\
& + \theta_{\mathbf{5}} \mathbf{W_i} + \theta_{\mathbf{6}} \mathbf{D_t} + \theta_{\mathbf{7}} \mathbf{S_{i,t}} + \theta_8 Y_{i,t-1} + \theta_9 Y_{i,t-2} + \mu_{i,t}
\end{aligned}
\tag{2.5}
$$

$HighImpact_i$ is the indicator for whether worker i has an above average payday differential. We are interested in the coefficients on the interactions—θ_2 and θ_4 . If the workers that are most affected by paydays are also those that benefit the most from the dominated contracts, then these coefficients will be positive.

Table 2.5
Heterogeneity in Take-up of Dominated Contracts:
Correlation with Payday Impact

Dependent variable	*Target level chosen*	*Positive target indicator*
	(1)	(2)
High payday production impact	353	0.138
	(129)***	(0.044)***
Seat fixed effects	Yes	Yes
Date fixed effects	Yes	Yes
Lag production controls	Yes	Yes
Observations	4098	4098
R2	0.22	0.20
Dependent variable mean	759	0.28

Notes: This table tests whether workers that are most affected by paydays are also more likely to demand the dominated contract. The dependent variables are two measures for take-up of the dominated contract. The first variable is the continuous measure of the target level selected by workers when assigned to Choice. The second variable is a binary indicator for whether the worker selected a target above zero when assigned to a Choice treatment. Both dependent variables are defined as 0 if a worker was absent the day before or day of Choice assignment. Each column shows results from an OLS regression of the dependent variable on an indicator for whether a worker had a high payday impact, computer fixed effects, and date fixed effects. High payday production impact is an indicator that equals 1 if the worker's payday impact—defined as the difference between mean production on paydays and non-paydays under Control divided by mean production under assignment to Control—is above the sample average. Standard errors allow for clustering by worker and are reported in parentheses.

Table 2.6: Heterogeneity in Contract Treatment Effects - Correlation with Payday Impact

Dependent variable	*Production*	*Production*	*Production*	*Attendance*	*Attendance*	*Attendance*
	(1)	(2)	(3)	(4)	(5)	(6)
Assignment to choice	118	-69	-146	0.007	-0.016	-0.028
	(60)*	(74)	(84)*	(0.009)	(0.010)	(0.011)**
Assignment to choice *		482	735		0.058	0.091
High payday production impact		(126)***	(144)***		(0.019)***	(0.022)***
Assignment to choice *			401			0.064
Payday			(179)**			(0.024)***
Assignment to choice * Payday *			-1314			-0.178
High payday production impact			(288)***			(0.041)***
Assignment to a target	153	-35	-48	-0.003	-0.019	-0.024
	(71)**	(86)	(96)	(0.010)	(0.012)*	(0.013)*
Assignment to a target *		483	673		0.042	0.066
High payday production impact		(148)***	(168)***		(0.022)*	(0.025)***
Assignment to a target *			68			0.026
Payday			(219)			(0.029)
Assignment to target * Payday *			-972			-0.120
High payday production impact			(348)***			(0.049)***
Payday			-183			-0.009
			(153)			(0.021)
High payday impact *			1178			0.164
Payday			(234)***			(0.032)***
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Lag production controls	Yes	Yes	Yes	No	No	No
Observations	8240	8240	8240	8240	8240	8240
R2	0.60	0.59	0.59	0.11	0.11	0.11
Dependent variable mean	5355	5355	5355	0.875	0.875	0.875

Notes: This table reports estimates of how heterogeneity in treatment effects of the contracts correlates with effects of paydays. Columns (1)-(3) report results from OLS regressions in which the dependent variable is production. Production is defined as the number of accurate fields entered in a day and equals zero on days workers are absent. Columns (4)-(6) report results from OLS Linear Probability Models in which the dependent variable is a binary indicator for attendance. Columns (1) and (4) show regressions of the dependent variable on indicators for contract treatment. Columns (2) and (4) add interactions of each of the contract treatment indicators with an indicator for High payday production impact. High payday production impact is an indicator that equals 1 if the worker's payday impact—defined as the difference between mean production on paydays and non-paydays under Control divided by mean production under assignment to Control—is above the sample average. For each of the contract treatments, columns (3) and (6) add a triple interactions of the contract treatment indicator with the High payday impact indicator and an indicator for whether the current day was the worker's assigned payday (along with the pair wise double interactions between these variables). All regressions include worker, date, and seat fixed effects. In addition, columns (1)-(3) include controls for lagged production. Robust standard errors are reported in parentheses.

The results indicate that this is indeed the case. The average treatment effects of Choice and Target Assignment for workers with low payday impacts are statistically indistinguishable from 0. In contrast, compared to the production of low impact workers, high impact workers produce about 480 fields more under Choice and Target Assignment on average. These coefficients are significant at the 1% level and their magnitudes correspond to 9% of mean production. Using our estimate of the return to education, providing high payday impact workers with simply the option to select targets leads to production increases comparable to a one-year increase in education. In addition, using our benchmark production elasticity of 0.33 (see Section 2.4.2), the local average treatment effect for high impact workers is comparable to a 64% increase in the piece rate wage.

In column (3), we explore how these treatment effects vary on paydays versus non-paydays. The payday impact statistic measures the extent to which a worker is affected by aligning the compensation period with the effort period. It constitutes an imperfect proxy for the level of a worker's self-control problem. On non-paydays, when the effort and compensation period are not aligned, time inconsistency is likely to create greater distortions on effort for workers with greater self-control problems. We therefore see that it is these workers that benefit the most from the provision of targets on non-paydays. Specifically, on non-paydays, high payday impact workers produce 735 fields (14% of mean production) more under Choice than Control. At the same time—allowing for heterogeneity in the extent to which both types of treatments help a worker overcome her time inconsistency problem—the high difference workers are the ones that are most helped by paydays. They therefore have less need for dominated contracts to solve the self-control problem on paydays than those workers for whom paydays don't produce large benefits. As a result, we see that on paydays the Choice treatment is relatively more beneficial for workers with a low payday impact. The estimated coefficients on Target assignment in column (3) tell a similar story.

In columns (4)-(6), we repeat this analysis for attendance as the dependent variable. While the average treatment effects on attendance are indistinguishable from zero, we see in column (5) that this masks substantial heterogeneity. The Choice treatment does not affect attendance of low payday impact workers, and increases the attendance of high impact workers by 5.8 percentage points. This effect is significant at 1%. If workers face a self-control problem in not just how hard they work in the office, but also in the decision to show up to work (as implied by the payday results), then this effect is consistent with a model of time inconsistency with sophisticated agents. There is a sizable fixed cost of attendance—for example, up to a 2-hour commute in each direction. Workers that are sophisticated enough to pick targets are also sophisticated enough to know that in the absence of a target, they will be tempted to exert low effort. Consequently, they're more likely to go in when they can select targets, because they know their earnings on those days will justify paying the attendance fixed cost. This is consistent with the results in column (6), which match the trends in column (3). Choice boosts attendance for high impact workers especially on non-paydays, whereas on paydays the impact of Choice is relatively lower for high impact workers than low impact ones.[29]

In Appendix Table D.11, we perform the analogous exercise for payday treatment effects—we examine whether those most impacted by the contract treatments have higher increases in effort on paydays. We find that treatment effects under Choice and Target Assignment are highly predictive of payday spikes. However, take-up of dominated contracts is not predictive of payday effects.[30]

2.4.5 Morning and Evening Choice (Prediction 2.4)

In Table 2.7, we examine the impact of evening choice on take-up of dominated contracts. The OLS specification is:

$$Takeup_{i,t} = \delta_0 + \delta_1 Eve_{i,t} + \delta_{\mathbf{2}}\mathbf{W_i} + \delta_{\mathbf{3}}\mathbf{D_t} + \delta_{\mathbf{4}}\mathbf{S_{i,t}} + \mu_{i,t}, \tag{2.6}$$

where $Eve_{i,t}$ is an indicator that equals 1 if worker i was assigned to Evening Choice on date t and equals 0 if the worker was assigned to Morning Choice. $Takeup_{i,t}$ measures take-up of the dominated contract by worker i on date t. As before, we use two measures of take-up—the target level chosen and a binary indicator for whether a positive target was chosen. Column (1) shows no difference in take up of dominated contracts the evening before.

[29] Is the large production effect of Choice on high difference workers driven completely by the impact on attendance? In Appendix Table D.10, we estimate the average treatment effects of Choice on production and attendance as 395 fields and 4.4 percentage points, respectively, for high payday difference workers. For these workers, mean production conditional on attendance is 5581. As a simple calibration, 5581*0.044 = 245 < 395. This implies that the entire effect of Choice on Production for these workers is not driven by attendance increases. In column (3), we regress production conditional on attendance on the contract treatment dummies. While the coefficients are positive and significant, as discussed above, they are difficult to interpret since attendance is an endogenous outcome.

[30] A different interaction between dominated contracts and paydays may be interesting: is there greater or lower demand for dominated contracts on paydays? One intuition suggests that self-control problems may be lower on paydays but this need not be the case: this depends on whether the payday as a motivator is a substitute or a complement for the dominated contract as a motivator. In our data we find no difference in demand for dominated contracts on paydays versus non-paydays on average

Table 2.7: Demand for the Dominated Contract - Impact of Uncertainty

			Definition of High Uncertainty Indicator							
			Worker is assigned to high uncertainty computer				Worker's arrival time is sensitive to bus/train schedules		Worker has above average score on Constraints Index	
Dependent Variable	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Evening choice	-18	-0.002	-26	-0.003	168	0.066	87	0.027	126	0.041
	(37)	(0.012)	(32)	(0.010)	(72)**	(0.022)***	(46)*	(0.016)	(65)*	(0.021)*
High uncertainty indicator			-134	-0.013	-20	0.027	282	0.104	329	0.099
			(63)**	(0.016)	(82)	(0.022)	(206)	(0.062)	(193)*	(0.061)
Evening choice *					-230	-0.082	-253	-0.070	-233	-0.070
High uncertainty indicator					(78)***	(0.024)***	(97)**	(0.029)**	(89)**	(0.029)**
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No
Seat fixed effects	Yes	Yes	No	No	No	No	Yes	Yes	Yes	Yes
Observations	4193	4193	4193	4193	4193	4193	3106	3106	3106	3106
R2	0.34	0.33	0.33	0.32	0.33	0.32	0.13	0.12	0.13	0.12
Dependent variable mean	767	0.28	767	0.28	767	0.28	803	0.3	803	0.3

Notes: This table tests whether the horizon of choice and uncertainty impact demand for the dominated contract. The observations used are the Assignment to Choice treatment observations in the Analysis Sample. Two outcome measures of demand are used. Target level chosen is the continuous measure of the target level selected by workers when assigned to Choice. Positive target indicator is a binary indicator for whether the worker selected a target above zero when assigned to a Choice treatment. Both dependent variables are defined as 0 if a worker was absent the day before or day of Choice assignment.

Evening choice is a dummy that equals 1 if the worker was assigned to Evening Choice and 0 if the worker was assigned to Morning Choice. In columns (3)-(6), the high uncertainty indicator equals 1 if the worker was assigned to a computer that was highly sensitive to office network speed, and equals 0 otherwise. In columns (7)-(8), the high constraint indicator measures workers' response to a question in the endline survey that asked them to agree or disagree with the statement: "The bus/train schedules really impact whether I can get to work on time because if I miss one bus or train, the next one I can take is much later," and takes a value of 1 if the worker's response was "Agree Strongly" and 0 otherwise. In columns (9)-(10), the high constraint indicator equals 1 if a workers constraint index score—computed by averaging her responses to the 4 questions on constraints in the endline survey—is above the mean score in the sample. Standard errors are reported in parentheses. Robust standard errors are reported in columns (1)-(2). Standard errors are corrected to allow for clustering by computer in columns (3)-(6) and by worker in columns (7)-(10).

Why might this prediction fail? Ex post analysis and qualitative work suggest a hypothesis. In choosing the night before, workers may fear uncertainty that may affect productivity the next day—uncertainty which may be realized after showing up to work, making workers less likely to commit to a target the day before.

An important source of this volatility is network speed fluctuations.[31] These fluctuations affected the rate at which workers could send data entered from an image to the central server and retrieve the next image for entry. The wait time between images could range from one second to over five minutes. Some computers in the office were more sensitive to network fluctuations than others. We asked the office management staff to consult workers in identifying the set of computers that were perceived as more sensitive to network slowdowns. Management did not know the list would be used for this purpose.

We then tested whether the computers identified as more uncertain are indeed more sensitive to overall network fluctuations. In Appendix Table D.12, we show that these computers are more sensitive and illustrate this in Appendix Figure D.5. When workers arrive to the office in the morning, they receive new information on the network speed and can use this to inform their target choice. This information is especially valuable for workers on bad computers, since network shocks will greatly impact their productivity.

To test this, we estimate:

$$Takeup_{i,t} = \gamma_0 + \gamma_1 Eve_{i,t} + \gamma_2 BadComputer_{i,t} + \gamma_3 Eve_{i,t} * BadComputer_{i,t} + \gamma_4 \mathbf{W_i} + \gamma_5 \mathbf{D_t} + \mu_{i,t} \tag{2.7}$$

in columns (3)-(6) of Table 2.7. Columns (3)-(4) provide some evidence that workers assigned to bad computers are less likely to demand targets in the sample as a whole. In columns (5)-(6), we estimate the above OLS regression model. Our predictions hold strongly in these results. When assigned to a good computer, selected targets are 168 fields higher on average in the evening than the morning. However, when assigned to bad computers, selected targets are 82 fields lower on average in the evening than the morning. These are sizable magnitudes (equivalent to 22% and 11%, respectively, of the mean target levels chosen by workers in the sample overall). The results are similar if we use our binary measure of demand. Workers on good computers are 6.6 percentage

[31] Appendix Figure D.4 displays substantial fluctuations in day-to-day output in the office.

points more likely to pick a positive target in the evening than the morning. In contrast, those on bad computers are more likely to demand targets in the morning than the evening.

A second source of uncertainty faced by workers stemmed from external constraints on time. For example, as discussed in Appendix C, workers that lived in more remote areas faced long and uncertain commute times. These impacted morning arrival time and therefore how much the worker could produce in a day. In addition, some workers had duties or binding constraints on time outside the office. This made it more difficult to absorb production shocks. For example, if the network unexpectedly slowed down, it would have been harder for workers with external constraints to stay late in the office to ensure their targets were met. Much of the uncertainty from these sources was resolved by the morning of the workday—by then employees knew their arrival time at the office and would have had a better sense of duties at home for that day.

Thus, we expect workers with greater external constraints to be relatively more likely to demand the dominated contract in the morning. To test this prediction, we estimate:

$$Takeup_{i,t} = \varphi_0 + \varphi_1 Eve_{i,t} + \varphi_2 Constraint_{i,t} + \varphi_3 Eve_{i,t} * Constraint_{i,t} + \varphi_{\mathbf{4}} \mathbf{D_t} + \varphi_{\mathbf{5}} \mathbf{S_{i,t}} + \mu_{i,t}. \quad (2.8)$$

$Constraint_i$ is a measure of the external constraints faced by worker i. In columns (7)-(10) of Table 2.7, we present results for two binary measures of the constraint variable. In columns (7)-(8), the variable measures workers' response to a question in the end-line survey that asked them to agree or disagree with the statement: "The bus/train schedules really impact whether I can get to work on time because if I miss one bus or train, the next one I can take is much later." The high constraint indicator takes a value of 1 if the worker's response was "Agree Strongly" and 0 otherwise. The results indicate that workers with more uncertain commute times select targets more often under Morning Choice than Evening Choice, and the opposite is true for workers with less uncertain commute times. The end-line survey asked four questions related to external constraints. For the analysis in Columns (9)-(10), we compute a Constraint Index for each worker by averaging his or her answers to the four questions. The high constraint indicator equals 1 if the worker's constraint index score was above the sample mean score and equals 0 otherwise. The results in columns (9) and (10) are similar to those in columns (7) and (8), respectively.

2.4.6 Learning over Time

As workers gain experience, do they learn about the value of the dominated contracts or perhaps find other ways around their self-control problems? We examine how workers' choices and treatment effects evolve with their experience. We define $experience_{i,t}$ as the number of workdays worker i has been in the Analysis Sample on date t.[32]

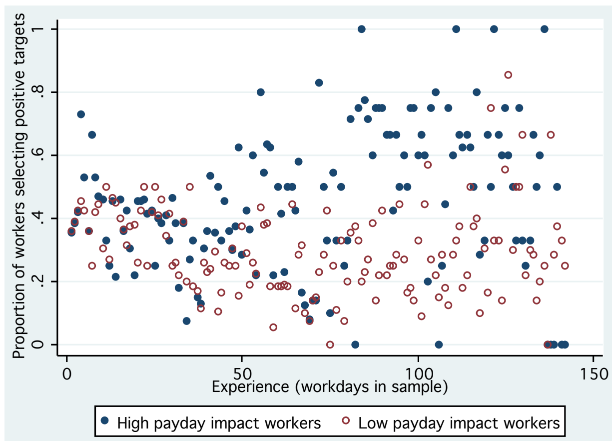

Figure 2.5: How Demand for Dominated Contracts Changes with Experience

Notes: This figure shows how demand for the dominated contract evolves with worker experience. The x-axis measures worker experience, defined as the number of workdays the worker has been in the sample. The y-axis measures the proportion of workers that chose a positive target under assignment to Choice for each value of the experience variable. The closed and open circles show the value of this statistic separately for high payday impact workers and low payday impact workers, respectively. High impact workers are those for whom the mean payday production impact—the difference in production on paydays and non-paydays under control, divided by overall mean production under control—is higher than the sample average payday impact. The proportion of times positive targets were chosen is computed using Choice observations in the Analysis Sample in which the worker was present the day before and day of Choice assignment.

In Figure 2.5, we explore how demand for the dominated contracts evolved with worker experience for high and low payday impact workers. For each value of the experience variable, we compute the proportion of high impact workers that choose a positive target under assignment to Choice (and were present the day of and day before Choice assignment); these values are plotted in closed circles. The open circles plot the value of this statistic for the low impact workers. The figure shows little initial difference in mean take-up rates between high and low impact workers. However, as workers gain experience, we see a divergence. Over time, those that have the largest self-control

[32] Recall that days during which workers are in training are not included in the Analysis Sample. As a result, $experience_{i,t} = 1$ on worker i's first day of contract randomizations. Note also that the experience variable suffers from selective attrition.

problems (as measured by our payday difference proxy) end up demanding the dominated contract at substantially higher rates than the workers that do not have large payday impacts.[33]

We explore these trends more formally in Panel A of Table 2.8. In columns (1) and (2), we regress each of our measures of take-up on the log of the experience variable and our standard controls. In both columns, we cannot reject that mean demand for the dominated contract does not change over time in the sample as a whole. In columns (3) and (4), we add the high payday impact indicator and an interaction of high payday impact with log experience. The results are consistent with the trends in Figure 6. As low impact workers gain experience, they decrease their take-up of the targets; a 1% increase in experience is associated with about a 0.066 percentage point decrease in take-up (significant at 5%). In contrast, we cannot reject that the demand among high impact workers stays constant over time. The F-tests of whether the log experience coefficient and interaction coefficient sum to zero are insignificant (with p-values of 0.493 and 0.895 in columns (3) and (4), respectively). As a result, at higher values of experience, the high payday impact workers exhibit substantially higher demand on average for the dominated contracts than the low impact workers.

These results are consistent with a story in which workers initially try the dominated contracts. Over time, they continually receive opportunities to observe their production under targets—both through Target assignment and potentially also when on Choice assignment. Those workers for whom the targets do not yield utility benefits stop selecting the dominated contract. In contrast, the workers with large self-control problems see that the targets are helpful, and continue to select them. Consistent with Test 3 of our model, this latter group of workers correlates with the group that is most affected by paydays.

Next, in Panel B of Table 2.8, we test whether the treatment effects on production persist over time. For reference purposes, in column (1), we first regress production on: log experience; dummies for Choice assignment, Target assignment, and Payday; and our vector of standard controls. As before, we define experience as the number of workdays the employee has been in the Analysis Sample. Not surprisingly, we see that production increases strongly with experience. The remaining results in column (1) are consistent with those presented in earlier tables.

[33] The figure also shows that there is variation in the level of day-to-day take-up within each group of workers over time. This is not surprising since the composition of workers assigned to Choice changes each day with the contract randomizations. In addition, day-to-day shocks (such as network speed fluctuations) impact take-up.

Table 2.8: Trends Over Time

	(1)	(2)	(3)	(4)
Panel A: Trends in Demand for the Dominated Contract				
Dependent variable	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*
Log experience	-102	-0.025	-128	-0.066
	(96)	(0.032)	(88)	(0.030)**
High payday production impact			-337	-0.087
			(277)	(0.100)
High payday production impact * Log experience			189	0.062
			(76)**	(0.026)**
Observations	4098	4098	4098	4098
Dependent variable mean	759	0.28	759	0.28
F-test p-value: (Log experience) + (High payday difference * Log experience) = 0			0.493	0.895
Panel B: Persistence in Treatment Effects Over Time (Dependent variable: Production)				
Sample	Analysis	Analysis	Analysis	Full Payday
Experience measure	*Log experience*	*Log experience*	*> Two months experience*	*Log experience*
Experience measure	257	176	459	314
	(66)***	(73)**	(148)***	(54)***
Assignment to choice	109	-269	27	
	(59)*	(174)	(76)	
Assignment to choice * Experience measure		104	144	
		(51)**	(114)	
Assignment to a target	145	-264	42	
	(70)**	(217)	(93)	
Assignment to a target * Experience measure		113	174	
		(62)*	(135)	
Payday	140	131	225	110
	(63)**	(179)	(82)***	(169)
Payday * Experience measure		3	-142	11
		(51)	(116)	(45)
Observations	8423	8423	8423	11744
Dependent variable mean	5337	5337	5337	5665

Notes: Panel A examines trends in demand for targets. It uses observations from the Analysis Sample in which the worker was assigned to a Choice treatment. The dependent variable in columns (1) and (3) is the target level selected, and is a binary indicator for whether the worker selected a positive target in columns (2) and (4). Both variables are defined as 0 if a worker was absent the day before or day of Choice assignment. Log experience is the log of the number of workdays an employee has been in the Analysis sample. High payday impact is a binary indicator that equals 1 if the worker's payday impact—the difference between mean production on paydays and non-paydays under Control divided by mean production under assignment to Control—is above the sample average. All regressions include seat and date fixed effects and lagged production controls. Regressions (1)-(2) also include worker fixed effects. Standard errors are clustered by worker.

Panel B tests whether treatments effects on production persist over time. All columns show results from OLS regressions in which the dependent variable is production—the number of accurate fields completed by a worker in a day, and equals zero if the worker is absent. Regressions (1)-(3) use observations from the Analysis Sample; Regression (4) uses observations from the Full Payday Sample. Log experience is as defined above. More than 2 months in sample is a binary indicator for whether the worker has been in the sample for more than 2 calendar months. All regressions include lagged production controls and worker, date, and seat fixed effects. Robust standard errors are reported in parentheses.

In column (2), we add interactions of log experience with each of the treatment variables of interest: Choice, Target assignment, and Payday. We are interested in the coefficients on the interactions. If treatment effects diminish over time—for example, once the novelty of the treatments wears off—then these coefficients will be negative. Instead, the results in column (2) reveal positive interaction coefficients. The interaction of log experience with Choice assignment is positive and significant at the 5% level. This is consistent with the findings in Panel A, which indicate that the workers that derive the largest benefits from the dominated contracts are the ones that are most likely to select them over time. In addition, the interaction with Target assignment is also positive and significant at the 10% level. The coefficient on the payday interaction is essentially 0, indicating that the payday effect is constant over time on average.

In column (3), we repeat this exercise using a different measure of experience: a binary indicator for whether the worker has been in the Analysis Sample for more than two calendar months. The coefficients Choice and Target interaction coefficients are positive (and insignificant). The interaction on the payday coefficient is now negative, but insignificant. In column (4), we check for persistence in the payday effect in the full 11-month payday sample and again find no change in effects over time. Together, columns (2) to (4) provide compelling evidence that the treatment effects of the contract treatments and paydays persist over time, and some evidence that the effects of Choice increase over time.

Overall, we see that workers select and derive steady benefits from the dominated contracts throughout the experiment. Similarly, the production increases on paydays persist week after week. Given the long horizon of the study, our results imply that time inconsistency is a perpetual problem in the workplace. They lend credence to our view that many workplace features can plausibly be interpreted as arrangements that seek to solve self-control problems.

2.4.7 Heterogeneity in Treatment Effects

How well do survey measures predict self-control? In Table 2.9, we use external measures to explore interpersonal differences. Each column conducts this analysis for a different potential correlate; the correlate for each column is specified at the top of that column. In Panel A, we show coefficient estimates from a regression of target level chosen on the correlate and controls. In Panel B, we regress the binary take-up indicator on the correlate and controls. In Panel C, we report estimates from a regression of production on: the correlate; dummies for Choice and Payday; interactions of each of these dummies with the correlate; and controls.

In columns (1)-(3), we look for evidence on whether more able workers are differentially affected by self-control. The correlate in column (1) is a measure of worker productivity: whether the worker's mean production under assignment to the control contract is above the sample mean. In column (2), the correlate is years of education. Both these correlates positively predict dominated contract demand—for example, high productivity workers are 40% more likely to choose positive targets under Choice. Interestingly, these measures do not positively predict treatment effects of choice, potentially indicating higher levels of sophistication among higher ability workers. In columns (3), we look at a measure of IQ—the sum of the worker's scores on the Raven's Matrix and Digit Span tests. IQ does not predict take-up or contract treatment effects. In addition, none of the ability correlates predicts treatment effects of paydays.

The literature on psychology and economics has proposed a range of correlates of self-control problems. In the remaining columns, we examine the predictive power of some of these correlates, collected through the end-line surveys. In columns (4)-(6), we look at measures of self-control problems based on self-reports by workers. The correlate in column (4) is the Self-control Factor, obtained from a factor analysis on the end-line survey data. In column (5), we construct a Self-Control Index from the end-line survey responses by averaging each worker's responses to the 9 self-control questions in the end-line survey. Both the Self-control Factor and Self-control index values have been de-meaned in the analysis. In column (6), we use self-reports of addictive behaviors by male workers. In this column, the correlate equals 1 if the worker said he had tried to quit drinking, smoking, or chewing tobacco and failed, and equals 0 otherwise. Each of these three columns shows similar results. These three correlates from the end-line surveys positively predict demand for the dominated contract, and also positively predict treatment effects of the contracts. However, among these, only the coefficients on the Self-control Factor are generally significant. None of the correlates predicts the payday effect.

Table 2.9: Heterogeneity in Treatment Effects - Correlates of Ability and Self-Control

	Correlate of Ability			*Correlate of Self-control Problems*				
	High productivity worker	*Years of education*	*IQ test index score*	*Self-control factor*	*Self-control index*	*Addictive behaviors dummy*	*Proportion of impatient responses*	*Proportion of preference reversals*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Dependent Variable: Target Level Chosen (Choice observations)								
Correlate	472	120	-1	124	148	352	230	533
	(159)***	(43)***	(5)	(74)*	(140)	(239)	(342)	(573)
Observations	4187	4056	4089	3106	3106	2245	2454	2454
R2	0.12	0.23	0.22	0.25	0.25	0.29	0.27	0.28
Panel B: Dependent Variable: Positive Target Indicator (Choice observations)								
Correlate	0.116	0.029	-0	0.056	0.057	0.139	0.070	0.143
	(0.049)**	(0.015)*	(0)	(0.025)**	(0.046)	(0.082)	(0.115)	(0.189)
Observations	4187	4056	4089	3106	3106	2245	2454	2470
R2	0.11	0.19	0.19	0.22	0.21	0.26	0.23	0.23
Panel C: Dependent Variable: Production (Choice and Control observations)								
Correlate	237	127	13	-255	-209	-247	-1254	-826
	(93)**	(65)*	(6)**	(103)**	(148)	(395)	(318)***	(461)*
Assignment to choice	1428	147	150	91	91	38	113	116
	(147)***	(73)**	(73)**	(82)	(83)	(114)	(83)*	(83)*
Assignment to choice *	-258	51	-1	167	215	429	706	765
Correlate	(154)*	(45)	(5)	(92)**	(146)	(263)	(305)**	(447)*
Payday	241	153	181	156	156	115	135	131
	(74)***	(74)**	(71)**	(87)*	(87)*	(119)	(106)	(107)
Payday *	-148	27	-0	31	77	58	-53	-234
Correlate	(121)	(40)	(4)	(87)	(138)	(246)	(308)	(544)
Observations	6304	6101	6149	4674	4674	3376	3701	3701
R2	0.57	0.56	0.55	0.57	0.57	0.55	0.59	0.58

Notes: This table tests whether correlates of ability and self-control explain heterogeneity in results. Columns (1)-(3) use measures of ability. The correlate in column (1) is an indicator for whether the worker's mean production under assignment to the control contract is above the sample mean. The correlates in columns (2) and (3) are, respectively, years of education and composite IQ score, defined as the sum of the worker's score on the Raven's Matrix and Digit Span tests. Both education and IQ have been de-meaned. The correlates in columns (4)-(8) are various proxies for self-control. The column (4) correlate is the Self-control Factor, obtained from a principal factors analysis on the endline survey data. The column (5) correlate is a Self-Control Index, obtained by averaging each worker's responses to the 9 self-control questions in the endline survey. Both the Self-control Factor and Self-control index have been de-meaned. The correlate in column (6) is computed for male workers; it equals 1 if the worker said he has tried to quit drinking, smoking, or chewing tobacco and failed, and equals 0 otherwise. The correlates in columns (7)-(8) are computed from the discount rate exercise, in which workers traded off cash rewards between different time horizons. The column (7) correlate measures the proportion of times the worker chose the smaller immediate reward instead of the larger delayed reward. The column (8) correlate measures preference reversals—the proportion of times a worker chose the larger immediate reward in the short horizon, but then chose the smaller delayed reward when choosing among the same amounts in the long horizon.

The dependent variable in Panel A is the mean target level chosen by workers under Choice. The dependent variable in Panel B is a binary indicator for whether the worker selected a positive target under Choice. Both these dependent variables are defined as 0 if a worker was absent the day before or day of Choice assignment. Both Panels A and B report estimates from an OLS regression of the dependent variable on the specified correlate and date and computer fixed effects; columns (2)-(8) also contain lag production controls. The dependent variable in Panel C is production—defined as the number of accurate fields entered in a day, and equal to zero on days workers are absent. Panel C reports estimates from a regression of production on: the specified correlate; Choice and Payday dummies; interactions of each of these dummies with the correlate; date and seat fixed effects; and controls for lagged production. In each regression, standard errors are corrected to allow for clustering by worker. Note that observations change between columns because not all workers provided education information or took the IQ tests, and because the endline survey and discount rate exercise were administered at the end of the project.

In columns (7)-(8), we look at outcomes from the discount rate exercise, in which we asked workers to trade off cash rewards between different time horizons—a standard way of testing for self-control problems in the literature. In column (7), we look at impatience. Our measure of impatience is defined as the proportion of times in the 6 questions the worker chose the smaller immediate reward rather than the larger delayed reward. In column (8), we look at preference reversals. Here, our self-control correlate is defined as the proportion of times a worker chose the larger immediate reward in the short horizon, but then displayed patience when choosing between the same amounts in the long horizon. As in the case of the Self-control Index and Addictive behaviors, these correlates do not appear to predict demand for the dominated contracts—the coefficients in Panels A and B are positive but significant. Also as before, we see in Panel C that workers with greater self-control problems (as measured by these correlates) are less productive on average. We also see evidence that these workers benefit more from the contract treatments. For example, the coefficient on the interaction between proportion of impatient responses and Choice assignment is 706 and is significant at the 5% level. It indicates that workers that show impatience in all 6 questions produce 706 more fields (13% of mean production) under Choice than Control. As before, the correlates do not seem to predict payday behavior.

Overall, we find some support that proxies of self-control posed in the literature correlate with behavior under the contract treatments. In contrast, none of these proxies is correlated with the payday effects. Thus, while these survey-based proxies have some predictive power, we find that the strongest predictor of effects under each set of treatments is workers' behavior under the other set of treatments (Section 2.4.4).

2.5 Alternative Explanations

The results are largely consistent with a self-control agency model. Could they be explained without appealing to self-control? We will argue that while any one result could be explained by other factors, it is hard to fit an alternative theory to the full pattern of results: the production increases on paydays; demand for dominated contracts and treatment effects of Choice; and the correlation between the payday and contract effects.

First, could workers be choosing dominated contracts because they are confused? The experiment was designed to minimize this possibility. During the training period, all workers were subjected to targets. At the end of training, we administered a quiz that tested their comprehension of the contracts. The mean score was 93%, indicating that workers understood the contracts.[34] Even after training, given the length of the experiment, workers receive (in both assignment and choice treatments) a great deal of experience with these contracts. As a result, the qualitative exit interviews with workers suggest they were well aware that the dominated contracts are dominated.

Second, could workers be choosing dominated contracts to signal ability to employers? It is not clear that demand for the dominated contract actually should serve as a positive signal. Since the employer observes production directly, there is no reason to believe a worker that can achieve high production under the control contract should not appear more impressive than one that needs a dominated contract to increase output. Still to mitigate this, we advertised the job as a one-time employment opportunity. Of course that might not have worked fully. The stumbling block for this—as with all the other explanations—is how it would explain both the payday effect and the correlation of dominated contract choice with the payday effect.

Third, a different psychological explanation could be that targets serve as non-monetary impact of targets. It need not be, as we have modeled them, that targets are merely monetary motivators. It is possible that the targets generate intrinsic motivation: the desire to hit the target alone may motivate workers. With data such as ours, of course, one cannot separate intrinsic from extrinsic motivation generated by the target. However, without time inconsistency it is unclear how this would explain the payday findings or the correlation between payday and contract effects. As a result, while our data cannot rule out non-monetary motivations provided by the target, it does suggest that time inconsistency is needed in this case as well.

Finally, while it may not explain choice of dominated contracts, could income targeting explain the payday effects? If workers target a fixed weekly income level, then small amounts of impatience or the realization of shocks could lead workers to backload effort closer to the payday. We can test for such behavior directly in the data. Income targeting implies a sharp decrease in marginal utility for income levels above the weekly target (see Camerer et. al. 1997). Two pieces of evidence

[34] In Appendix Table D.13, we test whether quiz score is correlated with demand for the dominated contracts. If workers mistakenly chose dominated contracts because they did not understand the contract treatments, then we would expect quiz score to be negatively correlated with take-up. Instead, quiz performance positively predicts take-up, although the coefficients in columns (1)-(4) are insignificant. In addition, as noted above, education strongly predicts take-up of the dominated contract.

suggest this is not happening in our data. First, we saw in our test for intertemporal substitution, exogenous production increases caused by target assignment do not lead to production decreases on subsequent days (see Appendix Table D.8). Second, a targeting model delivers an even finer testable prediction: an unexpected production increase today will lead to a larger reduction in tomorrow's effort if the worker is closer to her payday, because there are fewer subsequent days over which the adjustment needs to be made. We test for this in Appendix Table D.14. Under income targeting, the interactions in columns (2) and (3) should be negative. However, the interactions are positive (though largely insignificant). [35]

2.6 Conclusion

These results are interesting because they may help us explain workplace arrangements. Agency theory understands workplace arrangements—the existence of bosses, worker discipline and other pejorative mechanisms (even physical punishment Chwe 1990)—in one of two ways. The first view is that the firm exists to provide insurance. This insurance creates moral hazard. These workplace arrangements exist to mitigate that moral hazard. The second view is summarized in a story of Steven Cheung (1983): "On a boat trip up China's Yangtze River in the 19th Century, a titled English woman complained to her host of the cruelty to the oarsmen. One burly coolie stood over the rowers with a whip, making sure there were no laggards. Her host explained that the boat was jointly owned by the oarsmen, and that they hired the man responsible for flogging." Joint production necessitates the need for monitoring (Alchian and Demsetz 1972).

Our results suggest a different way to understand a diverse host of workplace arrangements. Might certain contract features (such as nonlinear contracts) be thought of as partly reflecting self-control benefits? Might discipline at the workplace or workplace rules be thought of as demand

[35] We also test an additional prediction of income targeting. Since the impact of day-to-day shocks is adjusted within the payweek to arrive at the weekly target, the variance in production among payweeks should be less than the variance in production among weeks defined according to some other arbitrary cycle, such as calendar weeks. To check this, we compare production across workers' payweek cycles with production across 4 artificial weekly cycles, created by shifting forward days from the worker's actual pay cycle. For example, for a worker assigned to the Saturday pay group, her true pay week is from each Monday to Saturday. The 4 artificial cycles for this worker would be from Tuesdays to Mondays, Wednesdays to Tuesdays, Thursdays to Wednesdays, and Fridays to Thursdays. For each worker, we then compute the standard deviation of weekly production across her actual payweeks and across each of the 4 associated artificial weekly cycles. On average, the standard deviation of weekly production for actual payweeks is 1838. The mean standard deviations for each of the 4 artificial cycles are lower than this, ranging from 1731 to 1809. Overall, all 5 standard deviation estimates are close to each other (within 5% or less from the payweek mean). This provides further support against weekly income targeting.

for features to help workers avoid the temptation to shirk? Might the organization of production itself, such as the presence of a boss and task division with deadlines serve to mitigate self-control problems? Clark (1994) advances this interpretation of the industrial revolution for example. One of the major changes in the organization of production in economic history has been the transition from the putting out system (under which workers were paid piece rates according to work performed and could choose production levels and work hours themselves) to the more rigid workplace system that is the norm today (with features like assembly lines, production minimums, rigid work hours, and hefty punishments for even momentary lapses in behavior). One interpretation is that increases in capital since the industrial revolution place a premium on increasing labor productivity (Clark 1994). Finding ways of reducing worker self-control problems could be one response to this problem. Under this view, the rise of the factory and its associated disciplinary infrastructure was in part an attempt to solve self-control problems.

Indeed, the process of development—which entails movements from agriculture to manufacturing, or from cottage industry to factory work—may increase labor productivity not just through technological innovation, but also because the work arrangements associated with these advances mitigate workers' self-control problems. (See Clark 1994 and Kaur et al. 2010 for further exposition of this view).

We may even need to enrich how we conceptualize the production function. Take the basic prediction that the self-control problem increases as the returns to work are further in the future. This mechanism suggests a new look at a variety of naturally occurring production function differences. In agriculture in developing countries, should we view productivity in long-horizon crops differently from productivity in short-horizon crops? Might farmers choose shorter horizon crops because effort distortions are smaller when effort and compensation are more closely aligned? Might the move from farm work to formal sector work with regular pay have self-control productivity benefits?

These arguments are, of course, speculative. However, given that we find strong evidence that self-control problems distort worker effort at economically meaningful magnitudes, a closer exploration of these possibilities is warranted in future research.

3 Multiple Market Failures and Allocative Efficiency in Villages

3.1 Introduction

A key distinguishing feature of poor countries is that, in contrast to rich counties, there is often little distinction between the household unit and the firm. In rural areas, agricultural production is largely undertaken on small family farms. In urban areas, large portions of the poor are self-employed in small family businesses (Banerjee and Duflo 2006).

If markets are complete, this organizing feature of production will have little impact on aggregate output: initial endowments will not affect how factors are allocated in production. However, in the presence of market failures, this feature can have potentially severe implications. Households with smaller endowments (such as smaller farms) will use productive factors (such as labor and capital) differently than households with larger endowments. This will lead to allocative inefficiency: the marginal product of inputs will not be equalized across "firms". This, in turn, will reduce aggregate output.

This paper empirically examines how multiple market failures interact to create allocative inefficiency in production. Using panel data on households in Indian villages, I focus on three markets that are central to agricultural production—labor, land, and credit.

I begin by documenting failures in the labor market. Following the approach in Chapter 1, I use rainfall shocks to demonstrate the presence of rigidity in wages for casual agricultural laborers. Specifically, positive rainfall shocks increase labor demand, leading to wage increases. Once these shocks have dissipated, however, wages do not return to their prior levels—they remain high in subsequent years. Chapter 1 shows that when a positive shock in one year is followed by a non-positive shock in the subsequent year, this leads to labor rationing. Thus, lagged rainfall shocks can be used to predict when the ration in the labor market is more likely to bind.

Due to separation failures, labor rationing may lead small farms to use labor more intensively than large farms (Singh, Squire, and Strauss 1986; Bardhan and Udry 1999). This implies that gains from trade can potentially be realized by redistributing land from large to small farms until the marginal product of labor is equalized across farms. In villages, land transfers are typically achieved through hiring (i.e. leasing or sharecropping) land. Such hiring arrangements are common (accounting for 20% of cultivated land) and their prevalence varies considerably from year to year.

In years when labor market rigidities are more likely to bind, there is a substantial increase in land transfers from large to small farms. Specifically, when a lagged positive shock is followed by a non-positive shock this year, there is a 63% increase in land hiring by households with below-average initial landholdings. In contrast, there is no increase in land hiring when a lagged positive shock is followed by a positive shock in the current year. This indicates that these landholding changes are not driven by wealth effects of lagged positive shocks, but stem from a response to labor market failures. These results are consistent with the prediction that failure in one market alone is not enough to create distortions—transactions in other markets can be used to restore allocative efficiency.

I then turn to examine the implications of multiple market failures. Hiring land requires substantial capital outlays (Shaban 1987). Given that the delay between the main harvest and the next major planting season is typically 5 months or longer, farmers cannot necessarily self-finance these outlays from prior harvest profits (Duflo, Kremer, and Robinson 2010). Thus, the acquisition of land by poor households requires the presence of well-functioning credit markets.

Consistent with this, land markets are relatively more responsive to labor market frictions in areas with better financial services. In years when labor rigidities bind, a one standard deviation increase in banking access leads to a 30% increase in land hiring by small farmers. These results hold for three separate measures of banking access (taken from Jayachandran 2006). The banking measures do not differentially predict land hiring responses to other sequences of rainfall shocks. This implies that banking is not simply capturing, for example, regional differences in the extent to which rainfall constitutes a labor demand shock. These results are consistent with the prediction that when there are multiple market failures (i.e. labor and credit), it will not necessarily be possible to achieve allocative efficiency.

These findings relate most closely to the empirical literature on missing markets and separation failures. Most of this literature has focused on testing the implications of separation failures—for example, by checking whether the number of adults in a household predicts labor use on its farm (see Rosenzweig 1988 and Behrman 1999 for reviews). This is primarily because it is difficult to identify variation in market failures. By using rainfall variation to isolate changes in rationing of hired labor, I directly test for the impact of labor market failures and explicitly examine interactions with other markets.

The results also contribute to the literature on land markets in developing countries. First, variation in take-up of sharecropping arrangements is typically understood through, for example, interpersonal differences in risk aversion. The above findings indicate that market-level conditions are also important determinants of these arrangements. Second, sharecropping has been viewed as providing reduced effort incentives (relative to land ownership), with potentially negative consequences for output (Shaban 1987). The results in this paper indicate that sharecropping can potentially increase the return to labor inputs because it involves land transfers from households where the marginal product of labor is higher to those where it is lower. In addition, if the margin of labor adjustment on sharecropped land is from being a hired laborer on someone else's land to farming one's own sharecropped land, the net incentive effects on effort are potentially ambiguous. Consequently, viewing sharecropping as a response to labor market failures provides additional ways to understand land arrangements in developing countries.

The rest of the paper proceeds as follows. Section 3.2 outlines the conceptual prediction and presents the empirical tests. Section 3.3 describes the data and defines rainfall shocks. The results are presented in Section 3.4, and Section 3.5 concludes.

3.2 Conceptual Predictions and Empirical Tests

Following Chapter 1, sequences of lagged and current rainfall shocks can be used to identify periods when labor rationing is more likely in village labor markets. Specifically, a positive rainfall shock in the current year—which constitutes a positive transitory labor demand shock—leads to an increase in the agricultural wage in the current year. If there is a non-positive shock in the subsequent year, under rigidities, the wage will not adjust back downwards.[1] This result from Chapter 1 can be verified using the following simplified regression model:

$$w_{ihvt} = \alpha_1 S_{vt}^{\{+,0\}} + \alpha_2 S_{vt}^{\{+,+\}} + \alpha_3 S_{vt}^{\{0,+\}} + \varphi \mathbf{X_{ihvt}} + \eta_{\mathbf{t}} + \psi_{\mathbf{hv}} + \varepsilon_{ihvt}, \tag{3.1}$$

where w_{idst} is the log of the nominal wage of worker i in household h in village v of year t; X_{ihvt} is a vector of controls; $\eta_{\mathbf{t}}$ is a vector of year fixed effects; and ψ_{hv} represents fixed effects for each village-landclass. Landclass is the category of landownership into which the household belongs; each household is categorized into one of 4 landclass groups—landless laborers and the 3 terciles of the village's landholding distribution.

Each of the remaining 3 covariates is an indicator for the realization of a particular sequence of shocks. The indicators take the form $S_{vt}^{\{i,j\}}$, where i denotes village v's shock in year $t-1$ and j denotes the village's shock in year t. The i and j take the values 0 or +, which correspond to the realization of a non-positive shock (referred to as "no shock") and a positive shock, respectively. Each indicator equals 1 if that particular sequence of shocks was realized and equals 0 otherwise. For example $S_{vt}^{\{+,0\}}$, equals 1 if village v had a positive shock last year and no shock this year, and equals 0 otherwise. The sequence $S_{vt}^{\{0,0\}}$, which is the case when the village experienced no shock last year and no shock this year, is omitted and serves as the reference case. Shocks are drawn from an iid distribution each year and are uncorrelated with the residual error, ε_{idst}. Thus, each of the coefficients on the indicator functions in equation (3.1) represents the reduced form average effect of that particular sequence of shocks on year t wages relative to $S_{dt}^{\{0,0\}}$.[2]

Under the null of no rigidities, only current shocks should predict wages: for $S_{dt}^{\{i,j\}}$, the sign of the coefficient should be determined solely by j. However, under rigidities, a lagged positive shock will have persist wage effects: it will positively predict current wages:

Prediction 3.1: Wage distortions from lagged positive shocks:

[1] As in Chapter 1, this strategy presumes that these shocks have no persistent productivity impacts—they last only for 1 agricultural year.

[2] Note that, unlike in Chapter 1, this specification does not distinguish between zero and negative shocks—it groups them into one category. This substantially simplifies the specification without compromising the ability to conduct the relevant tests. In addition, this simpler specification allows for analysis using data with more limited geographic and timeseries variation than that used in Chapter 1, since it requires observing 4 rather than 9 categories of shocks over time for each labor market. This is helpful given the data used for analysis (described below).

$\alpha_1 > 0$ *under sufficiently severe downward rigidities*

Note that regardless of the presence of rigidities, contemporaneous positive shocks should increase wages relative to the reference case: $\alpha_2 > 0$ and $\alpha_3 > 0$.

Since wages are higher under $S_{dt}^{\{+,0\}}$ than $S_{dt}^{\{0,0\}}$—while labor demand is the same under both sequences—farmers will demand less hired labor under $S_{dt}^{\{+,0\}}$ than $S_{dt}^{\{0,0\}}$. In the absence of any other changes in factor allocation, households with larger landholdings will demand less hired labor, while those with small landholdings (who are net suppliers of hired labor) will face labor rationing. In response, this latter group is likely to increase labor supply to its own farms until the marginal return from labor use on the family farm equals the marginal disutility from foregone leisure. As a result, labor will be inefficiently allocated: the marginal product of labor on large farms will be higher than that on small farms. This is the classic separation failure result (e.g. Benjamin 1992). This will result in lower aggregate production.

However, allocative efficiency can potentially be regained (or at least increased) by redistributing other factors of production. Specifically, in the absence of frictions in the land market, land can be transferred from large farms to small ones until the marginal product of labor is equalized across farms in the village. Such transfers will raise aggregate output and therefore constitute potential gains from trade [3]. In villages, land sales are extremely uncommon, due to substantial frictions in obtaining deed changes and a general unwillingness among farmers to permanently sell land. Land transfers are therefore typically achieved through leasing or sharecropping arrangements.[4] The following model tests whether labor market rigidities lead to land transfers from large to small households:

[3] If farm-specific human capital is an important input in production, or if there are substantial economies of scale in production, reallocation of land in this way will not necessarily increase aggregate output. These predictions assume that such considerations are sufficiently small relative to the importance of the level of physical labor used in production.

[4] Given incentive differences, sharecropping households may supply fewer labor inputs per acre than leasing or owning households (Shaban 1987). However, when land is transferred, on the margin, the change is from a small farmer working as a hired laborer on someone else's land to working on his own sharecropped plot of land. Given potential agency issues with hired labor, the net impact on effective effort supplied is unclear. More generally, if the incentive distortions from sharecropping are severe, then the below tests will fail to find an effect on land transfers even in the presence of labor market rigidities.

$$
\begin{aligned}
L_{hvt} \quad = \quad & \beta_1^i S_{vt}^{\{+,0\}} + \beta_1^{ii} large_{hv} * S_{vt}^{\{+,0\}} + \beta_2^i S_{vt}^{\{+,+\}} + \beta_2^{ii} large_{hv} * S_{vt}^{\{+,+\}} \\
& + \beta_3^i S_{vt}^{\{+,0\}} + \beta_3^{ii} large_{hv} * S_{vt}^{\{+,0\}} + \varphi \mathbf{X_{ihvt}} + \eta_\mathbf{t} + \psi_\mathbf{hv} + \varepsilon_{ihvt}
\end{aligned} \quad . \tag{3.2}
$$

L_{hvt} is the net amount of land that is sharecropped or leased by household h in village v in year t. $large_{hv}$ is an indicator that equals 1 if the household's initial landholding (at the beginning of the study period) is above the sample average and equals 0 otherwise. Note that $large_{hv}$ only appears in the interaction terms because its level effect is absorbed by the landclass-village fixed effects, ψ_{hv}. All other covariates are the same as above.

Prediction 3.2: Land re-allocation from large to small farms:

$\beta_1^i > 0$ and $0 < \beta_1^{ii} = -2\beta_1^i$

A concern with this test is that land increases among small farmers may be due to a wealth effect. Specifically, the lagged positive shock could allow credit constrained small farmers to finance land acquisition in the next year. However, this case, land re-allocation should occur under both $S_{dt}^{\{+,0\}}$ and $S_{dt}^{\{+,+\}}$. In contrast, if reallocation is due to the labor market failure, it should be higher under $S_{dt}^{\{+,0\}}$ than $S_{dt}^{\{+,+\}}$.

Auxiliary Prediction 3.2A: Distinguishing labor market failures from wealth effects:

$\beta_1^i > \beta_2^i$

The acquisition and operation of leased or sharecropped land requires substantial capital outlays by farmers: up front rental payments (in the case of lease contracts) and working capital costs for seeds, fertilizer, pesticides, irrigation, labor, and animal inputs (Shaban 1987). Harvest profits from Kharif (the main agricultural season) are realized in November-January, while substantial field operations for the next Kharif season do not begin until May-July. Given this gap, farmers are often unable to retain profits from the previous year to purchase inputs for the next year (Duflo, Kremer, Robinson). As a result, they are heavily reliant on credit to finance farm operations.

This suggests that credit access will be an important determinant of whether small farmers acquire land in response to the labor market failure. In the absence of credit access, land reallocation will be difficult to initiate (in the case of lease arrangements) and will not lead to increases in aggregate output (due to an inability of small farmers to pay for variable inputs). Thus, gains from trade in land will not be realized. This is consistent with the prediction that when there is a failure in more than one market (in this case, labor and credit), then it will not necessarily be possible to reallocate factors of production to achieve efficiency.

The following model adds a triple interaction of landholding and credit access with each shock category:

$$
\begin{aligned}
L_{hvt} = \quad & \beta_1^i S_{vt}^{\{+,0\}} + \beta_1^{ii} credit_v * S_{vt}^{\{+,0\}} + \beta_1^{iii} large_{hv} * S_{vt}^{\{+,0\}} + \beta_1^{iv} credit_v * large_{hv} * S_{vt}^{\{+,0\}} \\
& + \beta_2^i S_{vt}^{\{+,+\}} + \beta_2^{ii} credit_v * S_{vt}^{\{+,+\}} + \beta_2^{iii} large_{hv} * S_{vt}^{\{+,+\}} + \beta_2^{iv} credit_v * large_{hv} * S_{vt}^{\{+,+\}} \\
& + \beta_3^i S_{vt}^{\{0,+\}} + \beta_3^{ii} credit_v * S_{vt}^{\{0,+\}} + \beta_3^{iii} large_{hv} * S_{vt}^{\{0,+\}} + \beta_3^{iv} credit_v * large_{hv} * S_{vt}^{\{0,+\}} \\
& + \varphi \mathbf{X_{ihvt}} + \eta_{\mathbf{t}} + \psi_{\mathbf{hv}} + \varepsilon_{ihvt}
\end{aligned} \tag{3.3}
$$

The primary coefficients of interest are those in the first row: β_1^i to β_1^{iv}. Land reallocation should be especially likely when credit access is higher:

Prediction 3.3: The role of credit access in enabling land adjustments:

$\beta_1^{ii} > 0$ and $\beta_1^{iv} < 0$

Note that while the level of credit access is endogenous and will be correlated with the error term, rainfall shocks are iid within each village. As a result, $E[credit_v S_{vt}^{\{+,0\}} * \varepsilon_{ihvt}] = 0$ and $E[credit_v large_{hv} S_{vt}^{\{+,0\}} * \varepsilon_{ihvt}] = 0$. Thus, β_1^{ii} *and* β_1^{iv} capture causal effects.

However, the interpretation of these estimates is subject to a concern: rainfall may be a more important productivity shock in areas with higher credit access. For example, areas with better credit availability may grow crops that are more sensitive to rainfall variation. This differential productivity effect could give rise to the results in Prediction 3.3 even if credit has no direct effect on land hiring. Two auxiliary tests can be used to rule out this concern. First, if credit is a proxy for the importance of rainfall as a labor market shock, then the interactions of credit with $S_{vt}^{\{+,+\}}$ and $S_{vt}^{\{0,+\}}$ should also be significant (and move in the same direction as the interaction with $S_{vt}^{\{+,0\}}$). Second, the addition of triple interactions of $S_{vt}^{\{+,0\}}$ with the landholding indicator and agricultural measures (e.g., crops grown, proportion of irrigated land) should reduce the predictive power of the credit interactions.

Auxiliary Predictions 3.3A-B: Distinguishing credit access from differing effects of rain shocks across villages:

3A) $\beta_1^{ii} > \beta_2^{ii}$ and $\beta_1^{ii} > \beta_3^{ii}$

3B) When triple interactions of $S_{vt}^{\{+,0\}}$ with landholding and agricultural indicators are added to model (3.3), Prediction 3 will continue to hold.

3.3 Data and Definitions

Wage and land data are taken from the ICRISAT Village Level Studies. This is a panel dataset covering ten Indian villages in five districts in the states of Andhra Pradesh, Gujarat, Madhya Pradesh, and Maharashtra. It spans the years 1975-1984. Three villages are covered for all ten years, and the remaining villages were added in later years, for a total of 62 village-year observations. Forty households were randomly selected in each village: ten landless households, and ten households sampled from each tercile of the village's landholding distribution. The wage variable is constructed from Schedule K, which reports wages for casual agricultural employment. Household-level land data is constructed from Schedule Y, which includes plot-level details of each household's farm operations, including land ownership status. The land panel has been balanced so that any household (including landless households) which has owned or operated land in any year in the village is included in each year, with landholding defined as 0 when no land was used/operated. Demographic variables are taken from Schedule C, which captures household characteristics.

Table 3.1: Summary Statistics

	Mean	Std Dev	Unit of observation	Obser-vations	Source
Wages and Landholding					
Log nominal daily wage	1.140	0.549	Individual-months	4,406	ICRISAT
Hired land (acres)	0.872	3.287	Household-years	1,726	ICRISAT
Within-HH standard deviation in hired land (acres)	1.148	2.379	Households	336	ICRISAT
Financial Access					
Bank credit per capita (1981)	100	43	District	5	1981 Indian Census
Bank deposits per capita (1981)	154	73	District	5	1981 Indian Census
Bank branches (1975-1984)	85	40	District-years	37	Reserve Bank of India
% of Observations in Each Rainshock Category					
{*t-1*=Shock, *t*=No shock}	0.290	--	Village-years	62	Indian Meteorological Dept
{*t-1*=Shock, *t*=Shock}	0.210	--	Village-years	62	Indian Meteorological Dept
{*t-1*=No shock, *t*=Shock}	0.226	--	Village-years	62	Indian Meteorological Dept
{*t-1*=No shock, *t*=No shock}	0.274	--	Village-years	62	Indian Meteorological Dept

Notes:
1. This table presents means and standard deviations for variables used in the analysis.
2. Hired land is computed as sharecropped land plus hired land in a village-year.
3. Potential rainfall sequences are displayed in the form {t-1=X, t=Y}. X corresponds to last year's rainfall shock and Y corresponds to the rainfall shock in the current year. Shock is defined as rainfall in the first month of the monsoon above the 60th percentile of the village's usual distribution in that month; no shock is defined as rainfall below the 60th percentile in that month. The percentage of village-years for which each particular sequence of shocks was realized is reported.

To proxy for the level of financial services in an area, three district-level measures of banking access are used. The first two are average credit per capita and deposits per capita in the village's district, taken from the 1981 Indian Census. The third measure is the number of bank branches per capita, taken from the Reserve Bank of India. This latter variable changes from year to year, while the first two are constant across the sample. These are the same three measures of financial access used by Jayachandran (2006).

Rainfall data is taken from the IMD Gridded Daily Rainfall dataset, constructed by the Indian Meteorological Department. Rainfall estimates are provided for 1-by-1 degree latitude-longitude grids by interpolating from 1,803 regional weather stations located across the country. The coordinates of each village (provided in the ICRISAT data) are matched to their nearest node in the rain data. As in Chapter 1, the measure of interest is rainfall in the first month when the monsoon could arrive in a village, which ranges from May to June for the ICRISAT villages. Measuring rainfall in this month proxies for both high monsoon rainfall levels and early monsoon arrival—both of which constitute positive shocks. Rainfall distributions are computed for each village separately, using the years 1950 (the first available year in the IMD dataset) to 1984. A positive rainfall shock is defined as rainfall above the 60th percentile of the village's 1950-1984 rainfall distribution for the relevant month. A non-positive shock, or no shock, is defined as rainfall below the 60th percentile of the village's rainfall distribution.[5]

Table 3.1 provides summary statistics for the key variables used in the analysis.

3.4 Results

3.4.1 Wage Distortions (Prediction 3.1)

Table 3.2 verifies the finding in Chapter 1 that lagged positive shocks positively predict current wages, indicating the presence of wage rigidity in the ICRISAT data. Column (1) shows a regression of the log of the daily nominal wage on the set of shock sequences, and year and village fixed effects. Column (2) shows the specification in regression model (3.1), with the full set of standard controls. The results are similar in both columns. Consistent with Prediction 3.1, a village that received a positive shock last year and no shock this year has about 10% higher wages than a village that received no shock in both years. One cannot reject that current wages are the same regardless of whether the shock occurred in the previous year or current year—the F-test for equality of the coefficients on the three shock categories has a p-value of 0.542.

Since the subsequent analysis will focus on differences among small and large landowners, Column (3) tests whether the wage effects of shocks vary with landholding status. The regression adds interactions of each shock sequence with an indicator for whether the household has above average landholding. Each interaction is insignificant, and the F-test for joint significance of the interaction terms has a p-value of 0.430. Thus, to the extent that shocks affect labor market opportunities for large and small farmers, these effects do not operate through differences in the wage.

[5] Note that this is different than the cut-offs used in Chapter 1, in which a positive shock is defined as rainfall above the 80th percentile, a negative shock is defined as rainfall below the 20th percentile, and no shock is defined as rainfall between the 20-80th percentiles. Defining positive shocks in this manner is not feasible in the ICRISAT data: given the relatively small geographic and timeseries variation, rainfall above the 80th percentile is seldom observed, and two consecutive years of rainfall above the 80th percentile are quite rare.

Table 3.2: Rigidity in Agricultural Wages

Dependent variable: Log nominal daily wage

		(1)	(2)	(3)
	{*t-1*=No shock; *t*=No shock}	(Omitted)	(Omitted)	(Omitted)
1	**{*t-1*=Shock; *t*=No shock}**	**0.127**	**0.102**	**0.095**
		(0.039)* **	**(0.036)*** **	**(0.039)** **
2	{*t-1*=Shock; *t*=No shock} x Large landholding			0.047
				(0.053)
3	{*t-1*=Shock; *t*=Shock}	0.193	0.160	0.157
		(0.086)**	(0.079)**	(0.079)*
4	{*t-1*=Shock; *t*=Shock} x Large landholding			0.014
				(0.054)
5	{*t-1*=No shock; *t*=Shock}	0.153	0.131	0.135
		(0.056)***	(0.047)***	(0.046)***
6	{*t-1*=No shock; *t*=Shock} x Large landholding			-0.021
				(0.041)
Year fixed effects?		Yes	Yes	Yes
Village fixed effects?		Yes	No	No
Village-landclass fixed effects?		No	Yes	Yes
Demographic controls?		No	Yes	Yes
Observations		4,406	4,406	4,406
Dependent variable mean		1.140	1.140	1.140
F-test p-value: Coeff 1 = Coeff 3 = Coeff 5		0.559	0.542	0.449
F-test p-value: Coeff 2 = Coeff 4 = Coeff 6 = 0		--	--	0.430

Notes:
1. This table tests for rigidity in agricultural wages. The dependent variable is the log of the nominal daily agricultural wage for casual labor.
2. Rainfall shocks are defined as Shock (corresponding to rainfall above the 60th percentile) and No shock (corresponding to rainfall below the 60th percentile). The shock covariates are of the form {t-1=X, t=Y}; they are indicators that equal 1 if the village experienced shock X in the previous year and shock Y in the current year.
3. Large landholding is an indicator that equals 1 if the household's landclass category (as defined in the ICRISAT data) is below the sample average, and equals 0 otherwise.
4. Demographic controls consist of a gender dummy, fixed effects for the household's caste ranking (as defined in the ICRISAT data), and fixed effects for the education category of the household head.
5. Standard errors are corrected to allow for clustering by district-year.

3.4.2 Land Reallocation in Response to Labor Market Rigidities (Prediction 3.2)

There are active markets for sharecropping and fixed-rent leasing in the ICRISAT villages. 25% of households are involved in sharecropping or leasing, and these arrangements account for about 20% of total cultivated area in the ICRISAT sample (Shaban 1987). As reported in Table 3.1, the mean amount of hired land (defined as sharecropped or leased land) in a household-year is 0.872

acres. In addition, there is considerable variation in the amount of land hired by a given household from year to year. The within-household standard deviation in hired land is 1.148 acres—132% of the mean hired landholding. This suggests that households actively adjust their farm size through changes in hired land.

Table 3.3: Land Adjustments in Response to Labor Market Failures

Dependent variable: Net hired land (acres)

		(1)	(2)	(3)
1	**{*t-1*=Shock; *t*=No shock}**	**0.547**	**0.518**	**0.446**
		(0.267)**	**(0.279)***	**(0.313)**
2	**{*t-1*=Shock; *t*=No shock} x Large landholding**	**-0.768**	**-0.838**	**-0.807**
		(0.273)***	**(0.294)*****	**(0.317)****
3	{*t-1*=Shock; *t*=Shock}		-0.189	-0.318
			(0.251)	(0.363)
4	{*t-1*=Shock; *t*=Shock} x Large landholding		-0.286	-0.251
			(0.246)	(0.275)
5	{*t-1*=No shock; *t*=Shock}			-0.149
				(0.256)
6	{*t-1*=No shock; *t*=Shock} x Large landholding			0.063
				(0.299)
Year fixed effects?		Yes	Yes	Yes
Village-landclass fixed effects?		Yes	Yes	Yes
Observations: household-years		1,726	1,726	1,726
Mean sharecropped or leased landholding (acres)		0.872	0.872	0.872
F-test p-value: Coeff 2 = Coeff 4		--	0.035**	0.035**

Notes:

1. This table tests whether households alter landholdings in years when labor market rationing is more likely. The dependent variable is net hired land: (land sharecropped or leased in) minus (land sharecropped or leased out) in each household-year.
2. Rainfall shocks are defined as Shock (corresponding to rainfall above the 60th percentile) and No shock (corresponding to rainfall below the 60th percentile). The shock covariates are of the form {t-1=X, t=Y}; they are indicators that equal 1 if the village experienced shock X in the previous year and shock Y in the current year.
3. Large landholding is an indicator that equals 1 if the household's landclass category (as defined in the ICRISAT data) is below the sample average, and equals 0 otherwise.
4. Each regression includes year fixed effects, village-landclass fixed effects, a gender dummy, fixed effects for the household's caste ranking (as defined in the ICRISAT data), and fixed effects for the education category of the household head.
5. Standard errors are corrected to allow for clustering by district-year.

Table 3.3 examines whether variation in hired land is explained by labor market rigidities. The dependent variable is net hired land, defined as (total land sharecropped in or leased in) minus (total land sharecropped out or leased out). Column (1) presents a regression of net hired land on $S_{vt}^{\{+,0\}}$, an interaction for whether the household's initial landholding (at the start of the study in

1974) is above the sample average, and standard controls. Consistent with Prediction 3.2, in years where labor market rigidities are more likely to bind, small landholders hire 0.547 acres of additional land on average—this corresponds to a striking 63% of the mean amount of hired land. In contrast, large landowners decrease their landholding by hiring out land. The estimated decrease in hired out land is 0.221 acres, but one cannot reject that Coefficient2 = 2*Coefficient1 (F-test p-value: 0.429).

Column (2) adds the sequence $S_{vt}^{\{+,+\}}$, and Column (3) adds all sequences for the full specification shown in regression model (3.2). These columns allow for tests of Auxiliary Prediction 3.2A. If the increase in hired land is due to wealth effects from the lagged positive shock, then hired land should increase under $S_{vt}^{\{+,+\}}$ (Coefficient 3) as well. However, in both columns, small farmers do not increase their hired land when a lagged positive shock is followed by a positive shock in the current year—in fact, the coefficients are negative, though insignificant. The F-test for the equality of the coefficients on $S_{vt}^{\{+,0\}}$ and $S_{vt}^{\{+,+\}}$ has a p-value of 0.035. Thus, small farmers increase their landholding only in years when lagged positive shocks are likely to cause labor market rigidities to bind.

3.4.3 Impact of Credit Access on Land Adjustments (Prediction 3.3)

Table 3.4 examines whether land adjustments are more likely in areas with better access to financial services. Columns (1)-(3) present results from a regression of net hired land on a triple interaction of $S_{vt}^{\{+,0\}}$ with the large landholding indicator and a measure of banking access. Each of the three columns uses a different measure of banking access: credit per capita, deposits per capita, and banks per capita, respectively. Each measure has been normalized to have mean 0 and standard deviation 1.

As before, on average, small landholders increase their landholding by about half an acre in years when labor rigidities are more likely to bind (Coefficient 1). This increase is larger in areas with more financial access (Coefficient 2). For example, a one standard deviation increase in credit per capita leads to a 0.258 acre relative increase in hired land. This corresponds to 30% of the mean amount of hired land. Similarly, the amount of land hired out by large landowners is greater in areas with better financial access (Coefficient 4). These results are consistent with Prediction 3.3.

Column (4) repeats the specification in Column (1), but without village-landclass fixed effects so that the main coefficients on banking access and landholding can be viewed.

Columns (5) provides a test of Auxiliary Prediction 3.3A. It shows estimates for regression model (3.3), with triple interactions included for all three shock sequences. Increases in credit access are not associated with increases in hired land among small farmers under $S_{vt}^{\{+,+\}}$ or $S_{vt}^{\{0,+\}}$. In fact, both coefficients are negative. In addition, as a whole, banking access does not predict effects of these shocks on hired land; the F-test for joint significance of Coefficients 6, 8, 10, and 12 has a p-value of 0.243. In contrast, as in the earlier regressions, banking access continues to predict the effects of $S_{vt}^{\{+,0\}}$ on hired land; the F-test for joint significance of Coefficients 2 and 4 has a p-value of 0.027.

To test Auxiliary Prediction 3.3B, a series of crop controls to the basic specification shown in Column (1). Four crop indicators capture each major category of crops defined in the ICRISAT data: cereals (e.g. rice and maize), pulses (e.g. lentils), oilseeds (e.g. groundnuts), and fiber crops (e.g. cotton). For the first year that each household appears in the sample, a crop indicator was constructed that equals 1 if the household grew the given category or crops and 0 otherwise.[6] There is substantial variation in which crops are grown in each village; this variation is captured by these crop indicators. A triple interaction between the banking measure, the landholding dummy, and each of these 4 crop indicators was added to the specification in Column (1). The results from this regression are displayed in Column (6). If variation in banking is simply proxying for differences in crops across villages, then the inclusion of these additional controls should render Coefficients 2 and 4 insignificant. However, there is little change in these coefficients relative to Column (1) and they continue to have predictive power; the F-test for joint significance of Coefficients 2 and 4 has a p-value of 0.045. The results are similar if triple interactions with other controls, such as the proportion of land with irrigation, are included. Together, Columns (5)-(6) provide evidence that the banking measures are not just reflecting differences in the importance of rainfall as a labor market shock across areas.

[6] These indicators were based on only the first year in the sample (rather than re-defined for each year) since crop planting decisions are likely endogenous to rainfall conditions. Since rain shocks are iid in each year, initial planting decisions will be uncorrelated with future shocks.

Table 3.4: Heterogeneity in Land Adjustments: Banking Access

Dependent variable: Net hired land (acres)

	Measure of Banking	*Credit per capita* (1)	*Deposits per capita* (2)	*Number of Banks* (3)	*Credit per capita* (4)	*Credit per capita* (5)	*Credit per capita* (6)
1	{t-1=Shock; t=No shock}	0.502 (0.117)**	0.497 (0.123)**	0.514 (0.219)*	0.421 (0.098)**	0.433 (0.170)*	0.483 (0.161)*
2	**{*t-1*=Shock; *t*=No shock} x Banking measure**	**0.258 (0.036)*** **	**0.252 (0.042)*** **	**0.311 (0.329)**	**0.197 (0.060)** **	**0.124 (0.084)**	**0.256 (0.107)* **
3	{*t-1*=Shock; *t*=No shock} x Large landholding	-0.588 (0.245)*	-0.589 (0.155)**	-0.753 (0.382)	-0.557 (0.239)	-0.500 (0.264)	-1.034 (0.387)*
4	**{*t-1*=Shock; *t*=No shock} x Banking measure x Large landholding**	**-0.617 (0.098)*** **	**-0.661 (0.036)*** **	**-0.266 (0.433)**	**-0.619 (0.109)*** **	**-0.680 (0.139)** **	**-0.980 (0.217)** **
5	{*t-1*=Shock; *t*=Shock}					-0.316 (0.287)	
6	{*t-1*=Shock; *t*=Shock} x Banking measure					-0.441 (0.114)**	
7	{*t-1*=Shock; *t*=Shock} x Large landholding					0.016 (0.203)	
8	{*t-1*=Shock; *t*=Shock} x Banking measure x Large landholding					-0.022 (0.108)	
9	{*t-1*=No shock; *t*=Shock}					-0.088 (0.293)	
10	{t-1=No shock; t=Shock} x Banking measure					-0.153 (0.135)	
11	{t-1=No shock; t=Shock} x Large landholding					0.232 (0.214)	
12	{t-1=No shock; t=Shock} x Banking measure x Large landholding					-0.223 (0.123)	
	Banking measure				0.668 (0.036)***		
	Large landholding				0.208 (0.243)		
	Banking measure x Large landholding				-0.123 (0.163)		
	Village-landclass fixed effects?	Yes	Yes	Yes	No	Yes	Yes
	Year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
	Crop triple interactions?	No	No	No	No	No	Yes
	F-test p-value: crop interactions	--	--	--	--	--	0.191
	Observations: household-years	1,726	1,726	1,726	1,726	1,726	1,726
	Mean hired landholding (acres)	0.872	0.872	0.872	0.872	0.872	0.872

Notes:

1. This table tests whether land adjustments depend on credit access. The dependent variable is net hired land: (land sharecropped or leased in) minus (land sharecropped or leased out) in each household-year.
2. Rainfall shocks are defined as Shock (rainfall above the 60th percentile) and No shock (rainfall below the 60th percentile). The shock covariates are of the form {t-1=X, t=Y}; they are binary indicators that equal 1 if the village experienced shock X in the previous year and shock Y in the current year.
3. Banking measure is one of three measures of banking access. The specific measure used in each regression is reported at the top of the table. Each measure has been normalized to have mean 0 and standard deviation 1.
4. Large landholding is an indicator that equals 1 if the household's landclass catgory is below the sample average.
5. Each regression includes fixed effects for year, village-landclass, household caste category, and education category of the household head, and a gender dummy. Standard errors are corrected to allow for clustering by district-year.

3.5 Conclusion

This chapter empirically examines land responses to labor market frictions. It uses rainfall shocks to demonstrate the presence of wage rigidity and identify years in which labor rationing is more likely to occur in a village. Households respond to labor rationing by reallocating land from large to small farms. However, their ability to make land adjustments is hindered in areas with poor credit access—highlighting a way in which multiple market failures distort allocative efficiency in production. These results provide support for the presence of separation failures in village economies. They are consistent, for example, with the view that labor rationing leads to more intensive use of labor inputs by small farms compared to large ones.

The findings have important policy relevance. They suggest that correcting failures in one market can have indirect benefits on other market distortions as well. Specifically, improvements in credit access will improve farmers' ability to mitigate distortions from labor market failures.

Appendices

Appendix A: Chapter 1 Model Proofs

This appendix presents proofs of the model propositions in Chapter 1 Section 1.2. Before proceeding, it is useful to specify an allocation mechanism by which workers are matched to firms. This is needed to formalize the impact of off-equilibium deviations on firm profits. I assume all firms simultaneously post a wage. Firms satisfy labor demand in descending order of posted wages. If multiple firms post the same wage, those firms proceed in random order. This ensures that the firms offering the highest wage receive priority in hiring. For simplicity, I assume each firm hires the available workers with the lowest ϕ values that are willing to work for it. This maximizes gains from trade in the narrow sense that for a given wage offer, those workers that would benefit the most from employment (the lowest ϕ workers) are the ones that get the job.

Proof of Proposition 1.1: Market Clearing in Benchmark Case

First, I show that the market clearing condition must hold in the benchmark case.

(i) Suppose there is excess labor supply: $JL^* < \frac{1}{\phi}u\left(\frac{w^*}{p}\right)$. Then firm j can cut its wage to some $w^* - \epsilon$ and still hire L^* workers. To see this, define δ as the slack in the market: $\delta = JL^* - \frac{1}{\phi}u\left(\frac{w^*}{p}\right)$. At wage $w_j = w^* - \epsilon$, by the allocation mechanism defined above, the supply of workers available to j equals the mass of workers that would be willing to work for j minus the mass of workers employed by the other (higher-wage) firms:

$$L_j^{Avail} \quad = \quad max\left\{\frac{1}{\phi}u\left(\frac{w^*-\epsilon}{p}\right) - (J-1)L^*\ ,\ 0\right\}$$

Firm j can cut wages by ϵ and still hire L^* workers as long as ϵ satisfies the following condition:

$$\begin{aligned} & L^* \leq \tfrac{1}{\phi}u\left(\tfrac{w^*-\epsilon}{p}\right) - (J-1)L^* \\ \implies & \tfrac{1}{J}\left[\tfrac{1}{\phi}u\left(\tfrac{w^*}{p}\right) - \delta\right] \leq \tfrac{1}{\phi}u\left(\tfrac{w^*-\epsilon}{p}\right) - \tfrac{J-1}{J}\left[\tfrac{1}{\phi}u\left(\tfrac{w^*}{p}\right) - \delta\right] \\ \implies & \tfrac{1}{\phi}u\left(\tfrac{w^*}{p}\right) - \delta \leq \tfrac{1}{\phi}u\left(\tfrac{w^*-\epsilon}{p}\right). \end{aligned}$$

Such a wage cut will strictly decrease j's wage bill while holding revenue constant, thereby strictly increasing profits. Thus, there cannot be excess labor supply.

(ii) Suppose there is excess labor demand: $JL^* > \frac{1}{\phi}u\left(\frac{w^*}{p}\right)$. This implies that each firm is hiring strictly less labor than demanded by its first order condition. If firm j raises its wage infinitesimally above w^* to $w^*+\epsilon$, it will be able to fully satisfy its labor demand by the allocation mechanism. In what follows, denote $L_j^{FOC}(w_j)$ as j's labor demand under wage w_j (this is determined by j's first order condition, (1.5)). This upward wage deviation will be profitable if profits from $w^*+\epsilon$ are higher than profits from w^*, i.e. if the following inequality holds:

$$\theta p f\left(L_j^{FOC}(w^*+\epsilon)\right) - (w^*+\epsilon)\, L_j^{FOC}(w^*+\epsilon) > \theta p f\left(\frac{1}{J\phi}u\left(\frac{w^*}{p}\right)\right) - w^*\frac{1}{J\phi}u\left(\frac{w^*}{p}\right).$$

Note that:

$$\begin{aligned} & \underset{\epsilon\to 0}{lim} \quad \theta p f\left(L_j^{FOC}(w^*+\epsilon)\right) - (w^*+\epsilon)\, L_j^{FOC}(w^*+\epsilon) \\ & = \theta p f\left(L_j^{FOC}(w^*)\right) - w^* L_j^{FOC}(w^*) \\ & > \theta p f\left(\tfrac{1}{J\phi}u\left(\tfrac{w^*}{p}\right)\right) - w^*\tfrac{1}{J\phi}u\left(\tfrac{w^*}{p}\right). \end{aligned}$$

The equality on the second line follows from the continuity of the first order condition and continuity of $f(\bullet)$. The inequality on the third line is due to the fact that at w^*, $L_j^{FOC}(w^*)$ maximizes profits. This implies that there exists some $\bar{\epsilon} > 0$ such that for all $\epsilon < \bar{\epsilon}$, profits from deviating to $w^*+\epsilon$ will be higher than maintaining wages at w^*.

Next, I show that no firm will deviate from the w^* pinned down by conditions (1.5) and (1.6).

(i) Suppose firm j raises its wage to some $w_j = w^* + \epsilon$. It follows from the first order condition, (1.5), that the firm will demand labor $L_j^{FOC} < L^*$. However, it could have hired L_j^{FOC} workers under wage w^*, with a lower wage bill and higher profits. This deviation cannot be profitable.

(ii) Suppose firm j lowers its wage to some $w_j = w^* - \epsilon$. The supply of workers available to j equals the mass of workers that would be willing to work for j minus the mass of workers employed by the other (higher-wage) firms:

$$\begin{aligned} L_j^{Avail} &= max\left\{\tfrac{1}{\overline{\phi}}u\left(\tfrac{w^*-\epsilon}{p}\right) - (J-1)L^*\ ,\ 0\right\} \\ &= max\left\{\tfrac{1}{\overline{\phi}}u\left(\tfrac{w^*-\epsilon}{p}\right) - \tfrac{J-1}{J\overline{\phi}}u\left(\tfrac{w^*}{p}\right)\ ,\ 0\right\}. \end{aligned}$$

Note that at $w^* - \epsilon$, $L_j^{Avail} < L^* < L_j^{FOC}$ by the above and the first order condition. This deviation will not be profitable iff $\pi_j\left(w^*, L^*\right) - \pi_j\left(w^* - \epsilon, L_j^{Avail}\right) \geq 0$.

(a) If $L_j^{Avail} = 0$, then $\pi_j\left(w^* - \epsilon, L_j^{Avail}\right) = 0$ and profits are trivially weakly higher from maintaining w^*.

(b) If $L_j^{Avail} > 0$, then profits from maintaining w^* will be higher for J sufficiently large. First, rewrite:

$$\begin{aligned} \pi_j\left(w^*, L^*\right) - \pi_j\left(w^* - \epsilon, L_j^{Avail}\right) &= p\theta\left[f\left(L^*\right) - f\left(L_j^{Avail}\right)\right] - \tfrac{\epsilon}{J\overline{\phi}}u\left(\tfrac{w^*}{p}\right) \\ &= F(J) - \tfrac{\epsilon}{J\overline{\phi}}u\left(\tfrac{w^*}{p}\right), \end{aligned}$$

where $F(J)$ is defined as the difference in output from L^* and L_j^{Avail}. Note that:

$$\frac{\partial}{\partial J}F(J) = \frac{1}{J^2\overline{\phi}}u\left(\frac{w^*}{p}\right)p\theta\left[f'\left(L_j^{Avail}\right) - f'\left(L^*\right)\right] > 0$$

by the concavity of $f(\bullet)$. Next, implicitly define $\widetilde{J}$as:

$$F(1) = \frac{\epsilon}{\widetilde{J}\overline{\phi}}u\left(\frac{w^*}{p}\right).$$

Cutting wages to $w^* - \epsilon$ will not be a profitable deviation for any J such that

$F(J) - \frac{\epsilon}{J\overline{\phi}} u\left(\frac{w^*}{p}\right) > 0$. The following shows this will hold for any $J \geq \widetilde{J}$. For any positive number X:

$$\begin{aligned} F(\widetilde{J}+X) &> F(\widetilde{J}) && (since\ \tfrac{\partial}{\partial J}F(J) > 0) \\ &> F(1) && (since\ \tfrac{\partial}{\partial J}F(J) > 0) \\ &= \frac{\epsilon}{\widetilde{J}\overline{\phi}} u\left(\frac{w^*}{p}\right) && (by\, definition\, of\ \widetilde{J}) \\ &> \frac{\epsilon}{(\widetilde{J}+X)\overline{\phi}} u\left(\frac{w^*}{p}\right). \end{aligned}$$

Thus for J sufficiently large, profits from maintaining w^* will be higher than from deviating to $w^* - \epsilon$. This is consistent with the assumption stated in the model that J is arbitrarily large. ■

Proof of Proposition 1.2: Downward rigidity at the reference wage

I prove each of the two parts of Proposition 1.2 in turn.

1) For $\theta \in (\theta'_R, \theta_R)$:

Define $\theta'_R = \frac{w_R}{pf'\left(\frac{1}{(J-1)\overline{\phi}} u\left(\frac{\lambda w_R}{p}\right)\right)}$. For $\theta \in (\theta'_R, \theta_R)$, no firm will deviate from wage offer w_R:

(a) Suppose firm j deviates by raising the wage to $w_j > w_R$. It follows from the first order condition, (1.7), that the firm will demand labor $L_j^{FOC} < \overline{L}$. However, it could have hired L_j^{FOC} workers under wage w_R, with a lower wage bill and higher profits. This deviation cannot be profitable.

(b) Suppose firm j deviates by lowering the wage to $w_j \in (\lambda w_R, w_R)$. By the firm's first order condition (1.7), j's labor demand will increase, but the supply of labor available to j will decrease to some L_j^{Avail}: $0 < L_j^{Avail} < \overline{L}(w_R, \theta, p)$. Then:

$$\begin{aligned}
\pi_j\left(w_j, L_j^{Avail}\right) &= p\theta f\left(\lambda L_j^{Avail}\right) - w_j L_j^{Avail} \\
&< p\theta f\left(\lambda L_j^{Avail}\right) - w_R\left(\lambda L_j^{Avail}\right) \qquad (since\, w_j > w_R\lambda) \\
&< p\theta f\left(\overline{L}(w_R, \theta, p)\right) - w_R\overline{L}(w_R, \theta, p) \quad (by\, FOC\, at\, w_R) \\
&= \pi_j\left(w_R, \overline{L}(w_R, \theta, p)\right).
\end{aligned}$$

This deviation is not profitable.

(c) Suppose firm j deviates by lowering the wage to $w_j \leq \lambda w_R$. Since $\theta > \theta'_R$, the definition of θ'_R above implies:

$$\overline{L}(w_R, \theta, p) > \frac{1}{(J-1)\overline{\phi}} u\left(\frac{\lambda w_R}{p}\right).$$

As a result, the supply of labor available to j is:

$$\begin{aligned}
L_j^{Avail} &= max\left\{\tfrac{1}{\phi}u\left(\tfrac{w_j}{p}\right) - (J-1)\bar{L}\,,\, 0\right\} \\
&\leq max\left\{\tfrac{1}{\phi}u\left(\tfrac{\lambda w_R}{p}\right) - (J-1)\bar{L}\,,\, 0\right\} \quad (since\, w_j \leq w_R\lambda) \\
&= 0 \qquad (by\, the\, expression\, for\, \overline{L}\, above).
\end{aligned}$$

The profits from cutting to $w_j \leq \lambda w_R$ are therefore 0. This deviation is not profitable.

The first order condition (1.7) implies that for $\theta \in (\theta'_R, \theta_R)$, $\overline{L}(w_R, \theta, p) < \overline{L}(w_R, \theta_R, p)$. This is because the wage remains fixed at w_R, while $\theta < \theta_R$, and $f(\bullet)$ is concave. Since by the definition of θ_R, $J\overline{L}(w_R, \theta_R, p) = \frac{1}{\phi}u\left(\frac{w_R}{p}\right)$, this implies that for $\theta \in (\theta'_R, \theta_R)$, $J\overline{L}(w_R, \theta, p) < \frac{1}{\phi}u\left(\frac{w_R}{p}\right)$. Thus, there will be excess labor supply in the market.

Finally, note that $\lim\limits_{\lambda \to 0} \theta'_R = \lim\limits_{\lambda \to 0} \frac{w_R}{pf'\left(\frac{1}{(J-1)\overline{\phi}}u\left(\frac{\lambda w_R}{p}\right)\right)} = 0$.

2) For $\theta \geq \theta_R$:

The definition of θ_R and Proposition 1 imply: $\overline{w}(w_R, \theta_R, p) = w^*(\theta_R, p) = w_R$. Since $\frac{\partial w^*(\theta,p)}{\partial\theta} > 0$ for all θ, $w^*(\theta_R, p) \geq w_R$ for $\theta \geq \theta_R$. The below arguments show that for $\theta \geq \theta_R$, no firm will want to deviate from $\overline{w}(w_R, \theta, p) = w^*(\theta, p)$:

(a) Suppose firm j raises its wage to some $w_j = \overline{w}(w_R, \theta, p) + \epsilon > w_R$. Since $w_j > w_R$, j's first order condition (1.7) coincides with the benchmark case. This deviation cannot be profitable by the same logic as part (i) of the proof of Proposition 1 above.

(b) Suppose firm j lowers its wage to some $w_j = \overline{w}(w_R, \theta, p) - \epsilon \geq w_R$. (Note that this implies $\theta > \theta_R$). The firm's choice of labor demand at w_j is given by first order condition (1.7). This deviation cannot be profitable by the same logic as part (ii) of the proof of Proposition 1 above.

(c) Suppose firm j lowers its wage to some $w_j = \overline{w}(w_R, \theta, p) - \epsilon < w_R$. Define $L_j^{FOC,\lambda}$ implicitly as: $p\theta\lambda f'\left(\lambda L_j^{FOC,\lambda}\right) = w_j$. In addition, define $L_j^{FOC,B}$ implicitly as: $p\theta f'\left(L_j^{FOC,B}\right) = w_j$. Note that by condition (1.8), $L_j^{FOC,\lambda} < L_j^{FOC,B}$. At w_j, j's optimal labor demand will correspond to $L_j^{FOC,\lambda}$. There are 2 possibilities:

1) If $L_j^{FOC,\lambda} > L_j^{Avail}$, then the amount of labor hired by the firm will correspond to L_j^{Avail} (the available labor supply). Then:

$$\begin{aligned}
\pi_j\left(w_j, L_j^{Avail}\right) &= p\theta f\left(\lambda L_j^{Avail}\right) - w_j L_j^{Avail} \\
&\leq p\theta f\left(L_j^{Avail}\right) - w_j L_j^{Avail} \qquad (since\ \lambda < 1) \\
&< p\theta f(L^*) - w^* L^* \qquad (by\ Proposition\ 1.1\ proof) \\
&= p\theta f\left(\overline{L}\right) - \overline{w}\,\overline{L} \\
&= \pi_j\left(\overline{w}, \overline{L}\right)
\end{aligned}$$

2) If $L_j^{FOC,\lambda} \leq L_j^{Avail}$, then the amount of labor hired by the firm will correspond to $L_j^{FOC,\lambda}$. Then:

$$\begin{aligned}
\pi_j\left(w_j, L_j^{FOC,\lambda}\right) &= p\theta f\left(\lambda L_j^{FOC,\lambda}\right) - w_j L_j^{FOC,\lambda} \\
&< p\theta f\left(L_j^{FOC,\lambda}\right) - w_j L_j^{FOC,\lambda} \qquad (since\ \lambda < 1) \\
&< p\theta f\left(L_j^{FOC,B}\right) - w_j L_j^{FOC,B} \qquad (by\ FOC\ condn\ (1.5)) \\
&< p\theta f(L^*) - w^* L^* \qquad (by\ Proposition\ 1.1\ proof) \\
&= p\theta f\left(\overline{L}\right) - \overline{w}\,\overline{L} \\
&= \pi_j\left(\overline{w}, \overline{L}\right)
\end{aligned}$$

Thus, such a downward deviation cannot be profitable.

Since $\overline{w}(w_R, \theta_R, p) = w^*(\theta_R, p)$ for $\theta \geq \theta_R$, this implies $\overline{L}(w_R, \theta_R, p) = L^*(\theta_R, p)$ because labor demand under the first order conditions (1.5) and (1.7) coincides for $w \geq w_R$. As a result, condition (1.6) implies $J\overline{L}(w_R, \theta, p) = \frac{1}{\phi} u\left(\frac{\overline{w}(w_R, \theta, p)}{p}\right)$ for $\theta \geq \theta_R$. ■

Proof of Proposition 1.3: Distortions from reference wage increases

Since, from Proposition 1.2, $\frac{\partial \theta'_S}{\partial \lambda} > 0$ and $\underset{\lambda \to 0}{lim}\ \theta'_S = 0$, for λ sufficiently small, it follows that $\overline{w}(w_S, \theta, p) = w_s$ for $\theta \leq \theta_S$.

First note that for $\theta \in (\theta_R, \theta_S)$:

$$\begin{aligned} \overline{w}(w_R, \theta, p) &= w^*(\theta, p) && by\, Proposition\, 1.2 \\ &< w^*(\theta_S, p) && by\, Corollary\, 1.1 \\ &= w_S && by\, definition\, of\, \theta_S \end{aligned} .$$

In addition, for $\theta \leq \theta_R$, $\overline{w}(w_R, \theta, p) \leq w_R < w_S$, where the first inequality follows from Proposition 1.2. Together, the above imply that $\overline{w}(w_R, \theta, p) < w_s$ for $\theta < \theta_S$.

Since Proposition 1.3 assumes $\overline{w}(w_S, \theta, p) = w_S$ for $\theta < \theta_S$, this implies: $\overline{w}(w_R, \theta, p) < w_S = \overline{w}(w_S, \theta, p)$ for $\theta < \theta_S$. Then, $\overline{L}(w_S, \theta, p) < \overline{L}(w_R, \theta, p)\ for\ \theta < \theta_S$ by the firm's first order condition (1.7). ■

Proof of Proposition 1.4: Inflation will mitigate distortions from nominal rigidity

Suppose that $\overline{w}\left(w_R, \widetilde{\theta}, \widetilde{p}\right) = w_R\ and\ J\overline{L}\left(w_R, \widetilde{\theta}, \widetilde{p}\right) < \frac{1}{\phi} u\left(\frac{\overline{w}\left(w_R, \widetilde{\theta}, \widetilde{p}\right)}{p}\right)$. As the price level rises above $\widetilde{p}$, holding the wage fixed at w_R, the first order condition (1.7) implies that labor demand will rise, while (1.2) implies that labor supply will fall. There will be a $p' > \widetilde{p}$ at which aggregate labor demand will be exactly equal to aggregate supply. This p' is pinned down by the following condition:

$$p'\widetilde{\theta} f'\left(\frac{1}{J\phi} u\left(\frac{w_R}{p'}\right)\right) = w_R.$$

Note that at p' and $\widetilde{\theta}$, w_R is the market clearing wage. This implies that: $\overline{w}\left(w_R, \widetilde{\theta}, p'\right) = w^*\left(\widetilde{\theta}, p'\right) = w_R$. In addition, for any $p'' \geq p'$:

$$\begin{aligned} \overline{w}\left(w_R, \widetilde{\theta}, p'\right) &= w_R && by\, definition\, of\, p'. \\ &= w^*\left(\widetilde{\theta}, p'\right) \\ &\leq w^*\left(\widetilde{\theta}, p''\right) && since\, \frac{\partial w^*}{\partial p} > 0 \\ &= \overline{w}\left(w_R, \widetilde{\theta}, p''\right) && by\, Proposition\, 1.2\, since\, w^*\left(\widetilde{\theta}, p''\right) \geq w_R \end{aligned} .$$

Thus, $\forall p \geq p'$, $\overline{w}\left(w_R, \widetilde{\theta}, p\right) = w^*\left(\widetilde{\theta}, p\right)$. In addition, this implies $\overline{L}\left(w_R, \widetilde{\theta}, p\right) = L^*\left(\widetilde{\theta}, p\right)$ since $\overline{w}\left(w_R, \widetilde{\theta}, p\right) \geq w_R$ and also implies market clearing by Proposition 1.2. ■

Appendix B: Chapter 1 Supplementary Tables

Table B.1: Test for Serial Correlation in Rainfall

Dependent variable: District's rainfall deviation in the current year

	Sample			
	World Bank data districts (1956 - 1987)		NSS data districts (1982 - 2008)	
	(1)	(2)	(3)	(4)
District's rainfall deviation in the previous year	-0.048	-0.056	-0.018	-0.016
	(0.035)	(0.033)*	(0.033)	(0.030)
District and year fixed effects?	No	Yes	No	Yes
Observations: district-years	8,672	8,672	15,392	15,392

Notes:

1. This table tests for serial correlation in rainfall. The unit of observation in each regression is a district-year. Regressions (1)-(2) perform the analysis for the districts in the World Bank dataset for rainfall over the years 1956-1987. Regressions (3)-(4) perform the analysis for the districts in the NSS dataset for rainfall over the years 1986-2007.
2. The dependent variable is a district's rainfall deviation, which equals the rainfall level in inches in the first month of the monsoon minus the district's mean rainfall level in that month in the sample.
3. Each column shows results of an OLS regression of the dependent variable on the district's rainfall deviation in the previous year. The regressions in columns (2) and (4) also include year fixed effects and district fixed effects.
4. Standard errors in each regression are corrected to allow for clustering by geographic region, as defined in the NSS data.

Table B.2: Test for Differential Impact of Rainfall Shocks by Season

	Dependent Variable	
	Log nominal wage	% Days Worked in Agriculture
	(1)	(2)
Positive shock	0.046	0.026
	(0.014)***	(0.012)**
Positive shock x	-0.015	-0.013
Harvest quarter (October-December)	(0.017)	(0.012)
Positive shock x	-0.023	-0.008
Post-harvest quarter (January-March)	(0.022)	(0.015)
Positive shock x	0.021	-0.027
Lean quarter (April-June)	(0.017)	(0.019)
Drought	0.023	-0.023
	(0.018)	(0.011)**
Drought x	0.023	-0.004
Harvest quarter (October-December)	(0.022)	(0.011)
Drought x	-0.018	0.015
Post-harvest quarter (January-March)	(0.022)	(0.014)
Drought x	-0.003	0.003
Lean quarter (April-June)	(0.023)	(0.018)
Harvest quarter (October-December)	0.035	0.010
	(0.010)***	(0.006)
Post-harvest quarter (January-March)	0.060	-0.054
	(0.011)***	(0.008)***
Lean quarter (April-June)	0.086	-0.097
	(0.011)***	(0.011)***
F-test p-value: Joint significance of interaction terms	0.118	0.546
Year and district fixed effects?	Yes	Yes
Obs: individual-years	154,476	1,002,005
Dependent var mean	3.244	0.483

Notes:

1.This table tests whether rainfall shocks have differential effects by season over the agricultural year.

2. Observations are from the NSS data. The dependent variable in Column (1) is the log of the nominal daily agricultural wage. The dependent variable in Column (2) is the percentage of days over the interview reference period that the worker was employed in agricultural activities.

3. Positive shock is an indicator that equals 1 if the district experienced rainfall in the first month of the monsoon above the 80th percentile and equals 0 otherwise. Drought is an indicator that equals 1 if the district experienced rainfall below the 20th percentile and equals 0 otherwise.

4. Each column shows results from an OLS regression of the dependent variable on rainfall shocks, dummies for quarter of the year, and interactions of each shock with quarters. The monsoon quarter (July-September) is omitted. Each regression contains year and district fixed effects and a dummy for gender. Regression (2) also contains a quadratic function of acres per adult in the household.

5. Standard errors are corrected to allow for clustering by region-year.

Table B.3: Specification Check - Effect of Shocks on Equilibrium Wages

Dependent Variable: Log Nominal Agricultural Wage

		Source: World Bank Data (1956-1987)			Source: NSS Data (1982-2008)		
		(1)	(2)	(3)	(4)	(5)	(6)
1	Positive shock	0.021	0.022	0.021	0.041	0.042	0.041
		(0.009)**	(0.009)**	(0.010)**	(0.010)***	(0.010)***	(0.012)***
2	Drought	-0.003	-0.003	-0.006	0.026	0.018	0.004
		(0.011)	(0.011)	(0.012)	(0.012)**	(0.011)	(0.016)
3	Lag positive shock		0.019	0.021		0.044	0.021
			(0.009)**	(0.010)**		(0.011)***	(0.014)
4	Lag drought		0.009	0.003		-0.002	-0.000
			(0.009)	(0.011)		(0.013)	(0.015)
5	Positive shock x Lag positive shock			-0.028			0.006
				(0.018)			(0.030)
6	Drought x Lag drought			-0.011			-0.029
				(0.022)			(0.034)
7	Positive shock x Lag drought			0.039			0.017
				(0.020)*			(0.031)
8	Drought x Lag positive shock			0.023			0.086
				(0.021)			(0.026)***
	F-test p-value: joint signif of Coeff 3, 5, & 8			0.042**			0.000***
	F-test p-value: joint significance of Coeff 3-8			0.020**			0.000**
	District and year FE?	Yes	Yes	Yes	Yes	Yes	Yes
	Additional controls?	No	No	No	Yes	Yes	Yes
	Obs: district-years	7,296	7,296	7,296	--	--	--
	Obs: individual-years	--	--	--	154,476	154,476	154,476
	Dependent var mean	1.197	1.197	1.197	3.261	3.261	3.261

Notes:

1.This table tests for the impacts of sequences of shocks on the agricultural wage.

2. The dependent variable is the log of the nominal daily agricultural wage.

3. Positive shock is an indicator that equals 1 if the district experienced rainfall above the 80th percentile in the current year and equals 0 otherwise. Drought is an indicator that equals 1 if the district experienced rainfall below the 20th percentile in the current year and equals 0 otherwise. Lag positive shock (Lag drought) is an indicator that equals 1 if the district experienced a positive shock (drought) in the previous year and equals 0 otherwise.

4. Each regression contains year and district fixed effects. Regressions (4)-(6) from the NSS data also include fixed effects for calendar quarters of the year and a dummy for gender.

5. Standard errors are corrected to allow for clustering by region-year.

Table B.4: Persistence of Lagged Shocks

Dependent Variable: Log Nominal Daily Agricultural Wage

	Source: World Bank Data (1956-1987)			*Source: NSS Data (1982-2008)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Positive shock this year	0.021	0.022	0.022	0.041	0.042	0.041
	(0.009)**	(0.009)**	(0.009)***	(0.010)***	(0.010)***	(0.009)***
Positive shock 1 year ago		0.019	0.020		0.044	0.042
		(0.009)**	(0.009)**		(0.011)***	(0.011)***
Positive shock 2 years ago			0.030			0.007
			(0.009)***			(0.013)
Positive shock 3 years ago			0.028			-0.012
			(0.010)***			(0.011)
Drought this year	-0.003	-0.003	-0.005	0.026	0.018	0.019
	(0.011)	(0.011)	(0.010)	(0.012)**	(0.011)	(0.011)*
Drought 1 year ago		0.009	0.009		-0.002	-0.001
		(0.009)	(0.009)		(0.013)	(0.013)
Drought 2 years ago			0.006			-0.009
			(0.009)			(0.013)
Drought 3 year ago			0.008			0.004
			(0.009)			(0.014)
District and year FE?	Yes	Yes	Yes	Yes	Yes	Yes
Obs: district-years	7,296	7,296	7,296	--	--	--
Obs: individual-years	--	--	--	154,476	154,476	154,476
Dependent var mean	1.197	1.197	1.197	3.261	3.261	3.261

Notes:

1.This table tests for the persistence of lag shocks on the agricultural wage.

2. The dependent variable is the log of the nominal daily agricultural wage. Columns (1)-(4) use observations from the World Bank data. Columns (5)-(7) use observations from the NSS data.

3. Positive shock is an indicator that equals 1 if the district experienced rainfall above the 80th percentile and equals 0 otherwise. Drought is an indicator that equals 1 if the district experienced rainfall below the 20th percentile and equals 0 otherwise. Each covariate is an indicator for whether the district experienced a positive shock or drought in the current or in a previous year, as described in the table.

4. Each regression contains year and district fixed effects. Regressions (4)-(6) from the NSS data also include fixed effects for calendar quarters of the year and a dummy for gender.

5. Standard errors are corrected to allow for clustering by region-year.

Table B.5: Inflation Results - Robustness Checks

	Dependent variable	Inflation rate		Log nominal wage		
	Interaction Term in Regressions			*Inflation in other states*	*Linear year trend*	*Post-1970 year dummy*
		(1)	(2)	(3)	(4)	(5)
	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Zero}	Omitted	Omitted	Omitted	Omitted	Omitted
1	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive}	-0.023	-0.023	0.014	0.017	0.023
		(0.011)*	(0.012)*	(0.009)	(0.010)*	(0.016)
2	{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive} x Interaction term			-0.024	-0.000	-0.010
				(0.123)	(0.001)	(0.022)
3	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought}	0.006	0.007	0.004	-0.020	-0.021
		(0.018)	(0.019)	(0.012)	(0.013)	(0.018)
4	{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought} x Interaction term			-0.280	0.000	0.002
				(0.160)*	(0.001)	(0.026)
5	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought}	0.003	0.004	0.025	0.019	0.001
		(0.014)	(0.015)	(0.025)	(0.023)	(0.029)
6	{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought} x Interaction term			-0.240	-0.000	0.028
				(0.216)	(0.002)	(0.039)
7	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero}	0.005	0.005	0.032	0.014	0.009
		(0.010)	(0.010)	(0.017)*	(0.011)	(0.016)
8	{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero} x Interaction term			-0.247	0.001	0.010
				(0.120)*	(0.001)	(0.024)
District fixed effects?		No	Yes	Yes	Yes	Yes
Year fixed effects?		No	No	Yes	Yes	Yes
Observations: district-years		6,384	6,384	5,016	6,384	6,384
Dependent variable mean		1.27	1.27	1.58	1.27	1.27

Notes:
1. This table provides robustness checks for the inflation results. Observations are from the World Bank data. The sample in column (3) is comprised of the years for which state inflation data is available (the years 1965-87, except 1975). The sample in all other columns is comprised of the years for which national inflation data is available (the years 1956-87 except 1960, 1963-64, and 1975).
2. Shocks are drought, zero, or positive, and correspond to rainfall below the 20th, between the 20th-80th, and above the 80th percentile, respectively. Covariates 1, 3, 5, and 7 are indicators of the form {$Shock_{t-1}$=X};{$Shock_t$=Y}; they equal 1 if the district experienced shock X in the previous year and shock Y in the current year, and equal 0 otherwise. The other covariates are interactions of the shock sequences with various controls.
3. Columns (1)-(2) show OLS regressions of the national inflation rate on shock sequences. Regression (2) also includes district fixed effects.
4. Columns (3)-(5) show OLS regressions of the log nominal daily agricultural wage on shock sequences and interactions. The interaction term in regression (3) is the average inflation rate across all states except the district's own state. The interaction term in regression (5) is a linear trend--the de-meaned calendar year. The interaction term in column (6) is an indicator that equals one if the year is after 1970.
5. Standard errors in all regressions are corrected to allow for clustering by year.

Table B.6: Effect of Shocks on Employment

Dependent Variable: Agricultural Employment Rate

			(1)	(2)
	Shock$_{t-1}$	*Shock$_t$*		
1	Zero	Zero	Omitted	Omitted
2	Drought	Zero	0.011 (0.008)	Omitted
3	Zero	Positive	0.016 (0.009)*	
4	Drought	Positive	-0.008 (0.016)	0.011 (0.007)*
5	Positive	Positive	0.030 (0.014)**	
6	Zero	Drought	-0.016 (0.009)*	-0.017 (0.007)**
7	Drought	Drought	-0.006 (0.014)	
8	Positive	Drought	-0.031 (0.011)***	-0.034 (0.011)***
9	Positive	Zero	-0.012 (0.008)	-0.015 (0.008)*
District and year FE?			Yes	Yes
Additional controls?			Yes	Yes
Observations: individual-years			1,002,005	1,002,005
Dependent variable mean			0.483	0.483

Notes:

1. This table tests for the impacts of sequences of shocks on the employment rate. Observations are from the NSS data. The dependent variable is the % of days in the past week in which the worker was employed in agricultural work (own farm work plus hired out work).
2. Shocks are defined as drought, zero, or positive, and correspond to rainfall below the 20th, between the 20th-80th, and above the 80th percentile, respectively. The covariates are indicators that equal 1 if a given sequence of shocks was realized and zero otherwise. The sequences are presented as the shock in the previous year and the shock in the current year.
3. Column (1) omits the sequence {Zero, Zero} and includes separate dummies for each of the remaining 8 combinations of shocks. The remaining columns group shocks into categories with similar predictions. Column (2) also omits the sequence {Drought, Zero}; combines rows 3-5 into one indicator function for whether the district experienced a contemporaneous positive shock; and combines rows 6-7 into an indicator function for whether the district had a zero shock or drought last year followed by a contemporaneous drought.
4. Each regression also contains year fixed effects, district fixed effects, fixed effects for calendar quarters of the year, a gender dummy, and a quadratic function of acres per adult in the household.
5. Standard errors are corrected to allow for clustering by region-year.

Table B.7: Tests for Compositional Effects of Rainfall on Agricultural Labor Force

Dependent Variable	*Individual Migrated into Village*			*Individual is in Agricultural Labor Force*	
	All village residents	Residents with positive agricultural employment	Members of agricultural labor force	All village residents	Residents with positive agricultural employment
	(1)	(2)	(3)	(4)	(5)
Panel A: Lag Shock Dummies					
$Shock_{t-1}$=Drought	-0.0063	-0.0069	-0.0091	0.000089	0.000408
	(0.0030)**	(0.0035)*	(0.0030)***	(0.003023)	(0.000360)
$Shock_{t-1}$=Positive	0.0042	0.0016	0.0036	-0.002853	0.000013
	(0.0039)	(0.0042)	(0.0044)	(0.002913)	(0.000326)
Panel B: Main Specification					
{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Zero}	Omitted	Omitted	Omitted	Omitted	Omitted
{$Shock_{t-1}$=Drought, Zero, or Positive}; {$Shock_t$=Positive}	-0.0086	-0.0122	-0.0124	-0.002173	0.000064
	(0.0045)*	(0.0058)*	(0.0058)*	(0.004099)	(0.000376)
{$Shock_{t-1}$=Drought or Zero}; {$Shock_t$=Drought}	0.0007	-0.0110	-0.0010	0.000860	0.000007
	(0.0041)	(0.003)***	(0.0035)***	(0.003687)	(0.000386)
{$Shock_{t-1}$=Positive}; {$Shock_t$=Drought}	-0.0013	-0.0087	-0.0090	0.002217	-0.001012
	(0.0077)	(0.0062)	(0.0070)	(0.004710)	(0.000711)
{$Shock_{t-1}$=Positive}; {$Shock_t$=Zero}	0.0059	-0.0016	0.0018	-0.006377	0.000480
	(0.0036)*	(0.0039)	(0.0040)	(0.003218)**	(0.000366)
Observations: individual-years	973,572	278,640	381,055	2,366,290	687,913
Dependent variable mean	0.114	0.120	0.135	0.405	0.994

Notes:

1. This table tests whether rainfall shocks lead to compositional changes in the agricultural labor force. Observations are from the NSS data.
2. The dependent variable in regressions (1)-(3) is a binary indicator that equals 1 if the individual is a migrant into the village. The dependent variable in regressions (4)-(5) is a binary indicator that equals 1 if agriculture is the respondent's primary or subsidiary occupation.
3. Panel A shows results from Probit regressions of the dependent variable on an indicator for whether the district had a positive shock in the previous year. Panel B shows results from Probit regressions of the dependent variable on categories of sequences of shocks. Estimated discrete changes are reported for each covariate. Each regression also contains year fixed effects, district fixed effects, and a gender control.
4. Columns (1)-(3) use data from NSS rounds in which migration questions were asked (agricultural years 1982, 1983, 1987, and 1999). Columns (4)-(5) use data from all rounds. Estimates are reported for different subgroups of the sample as indicated. All village residents includes all respondents interviewed in the rural sample. Residents with positive agricultural employment are those respondents who report being employed in agricultural work for at least half a day in the past week. Members of the agricultural labor force are respondents who indicate that agricultural work is their primary or subsidiary occupation.
5. Standard errors are corrected to allow for clustering by region-year.

Appendix C: Chapter 2 Study Details (Context and Protocols)

Production Task

Workers entered information from scanned images into fields on their screen (see Appendix Figure D.2). Once a worker finished entering data from an image, the software automatically sent the data to a central server and fetched the next image. This meant workers could not select the images on which they worked. Output was measured as the number of accurate fields entered. The data entry software displayed both the total and accurate number of fields entered so far that day (with about a 15 minute delay), so employees always had real time information on their own output.

Workers faced some uncertainty in production due to shocks. Two types of shocks are particularly relevant in our context. First, the office experienced network speed fluctuations that impacted productivity. Some computers were more sensitive to these fluctuations than others. As a result, workers were randomly assigned to seats in the office and these assignments changed every 1-3 weeks. Second, many employees commuted from surrounding villages using buses and trains, with some traveling up to two hours in each direction. Those from more remote locations faced increased uncertainty in morning arrival times and therefore production.

Paydays

Workers received their wages in cash on their assigned payday. Once they finished work for the day, they reported to the office manager, who computed and paid out their earnings for the previous week (including that day). If employees were absent on their payday, they could collect their owed earnings when they returned to work at no penalty.

Contract Treatments

Workers were paid piece rates based on output. The control contract paid Rs. 0.03 for each accurate field entered, regardless of production amount. The dominated contract paid Rs. 0.03 per accurate field if the worker met the day's production target, and Rs. 0.015 per accurate field otherwise.

Under the Assignment to a Target treatment, workers were assigned to low, medium, and high targets. These were set at 3,000, 4,000, and 5,000 accurate fields, respectively. In the first month of randomizations, these corresponded to the 30th, 50th, and 70th percentiles, respectively, of worker production under the control contract. During the last month of contract randomizations, we changed these levels to 4,000, 5,000, and 6,000 accurate fields to correspond to increases in worker production over time.

Before leaving work each day, employees were required to report to an office staff member in a separate area of the office. At that time, they were told their contract assignment for the next day. For example, employees were informed of Wednesday's assignment on Tuesday evening. If the assignment was Evening Choice, they also selected their target for Wednesday at that time. If the assignment was Morning Choice, then they selected their target upon arriving in the office Wednesday morning. This exchange was confidential and took place away from other workers.

Office Structure and Timeline

The office was open each day from 8:45 am to 6:30 pm, five to six days per week except holidays. Employees could choose when they worked, except for two 15-minute periods each day when work activity was halted for "server maintenance". In accordance with the norms of the firm with which we worked, employees were given tea in an outside area at 11 am and 3:30 pm each day. Workers could select the length of their tea breaks and lunch breaks. They were also free to check email, play computer games, or leave the office at any time.

The project ran for 15 months. During the first 2 months, the management staff established protocols, recruited subjects, and trained the new hires. After this, the contract and payday randomizations ran for 4 months. There was then a 2-month break while the office underwent changes to the data entry software and task. During this time, workers were generally not paid the standard piece rates and there were no contract randomizations. The contract and payday randomizations then resumed for another 4 months. In the final 3 months of the project, we ran end-line activities and surveys. We did not randomize workers into the four contract treatments during this time, but we continued to adhere to the payday assignments. Thus, the contract treatments ran for an approximate total of 8 months and the payday treatments for 11 months.

Sample construction

The office held 64 data entry operators at a time. Due to employee turnover, 111 workers participated in the experiment. When an employee quit, the management staff hired a replacement from a database of persons that had submitted applications for the job. As in the initial recruitment, workers were hired in order of application date. The gender composition of employees was kept fixed—if a female quit, the worker hired to replace her was a female. Each new worker "inherited" all the assignments of his or her predecessor—payday group, vector of contract assignments, and seating assignment. The payday and seat assignments took effect immediately. New hires began their scheduled contract assignments after completing training.

Appendix D: Chapter 2 Supplementary Figures and Tables

		Contract Assignment (Assignment changes daily)			
		Control Contract (0.25)	*Assigned to Target* (0.25)	*Evening Choice* (0.25)	*Morning Choice* (0.25)
Payday Assignment (Assigned once in beginning of study)	*Tuesday Payday* (0.33)				
	Thursday Payday (0.33)				
	Saturday Payday (0.33)				

Figure D.1: Experiment Design

Notes: This chart provides an overview of the treatment design. One-third of workers were assigned to each of the three payday groups. This assignment was done once for each worker, when the worker joined the firm, and remained fixed for the duration of the project. Workers were orthogonally assigned to each of the four contract treatments exactly 25% of the time. The assignments changed daily.

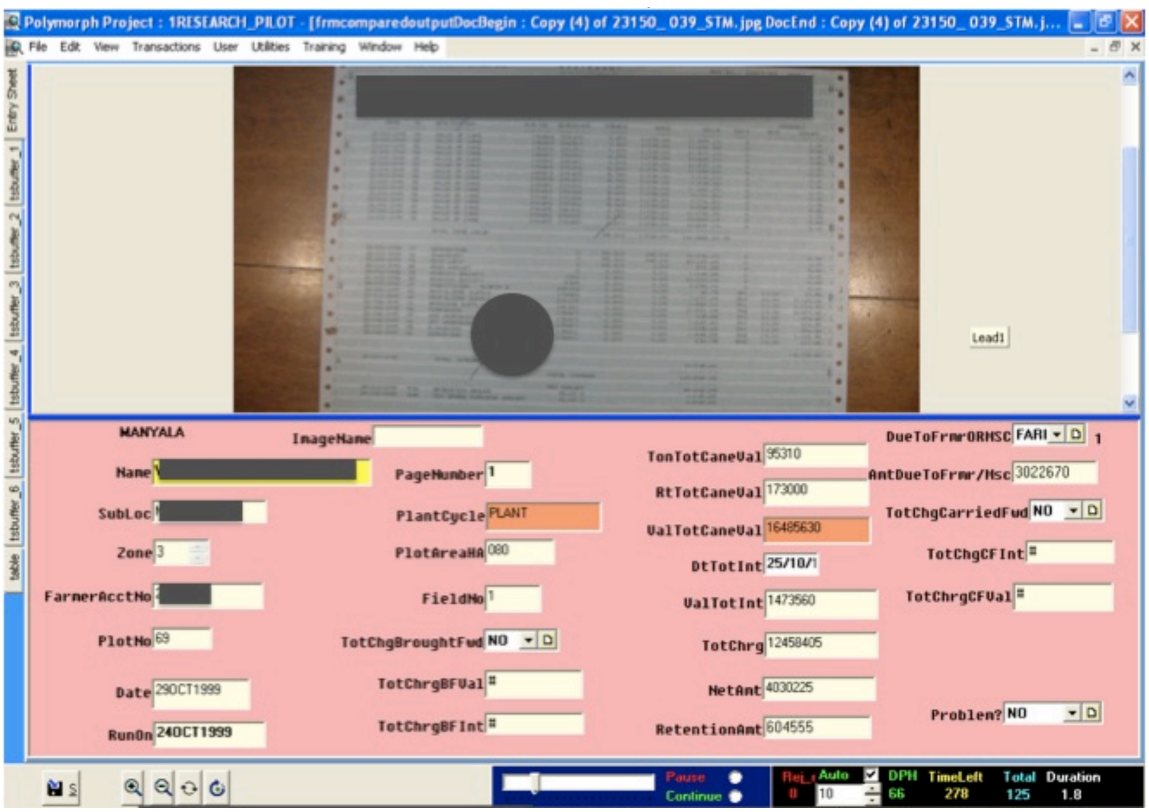

Figure D.2: Workers' Data Entry Screen

Notes: The figure displays a screen shot of a typical data entry screen. Workers viewed scanned images of records in the top half of their screen and entered information from these images into the appropriate fields at the bottom half of the screen. Identifying information from the records has been covered for confidentiality.

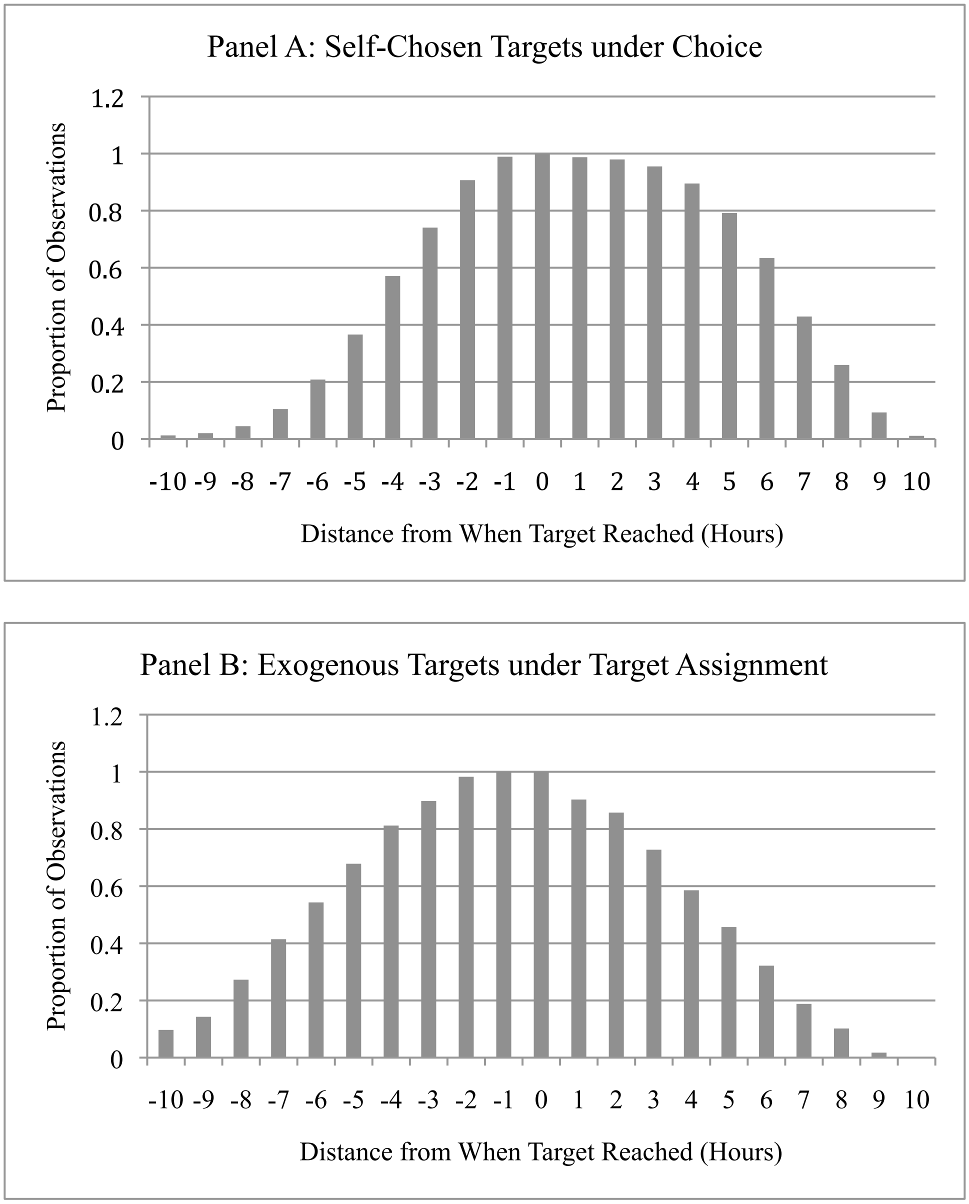

Figure D.3: Proportion of Observations around Target Thresholds

Notes: This figure shows in what proportion of worker-days we observe workers at a given distance from when they reach their targets. The x-axis measures distance in hours from when a worker reached her target for the day. The y-axis measures the proportion of observations for which that distance falls between 8 am – 6 pm (the hours of operation of the office). Panel A computes these statistics for observations in which workers were assigned to Choice, chose positive targets, and were present. Panel B computes these statistics for observations in which workers were assigned to a Target and were present.

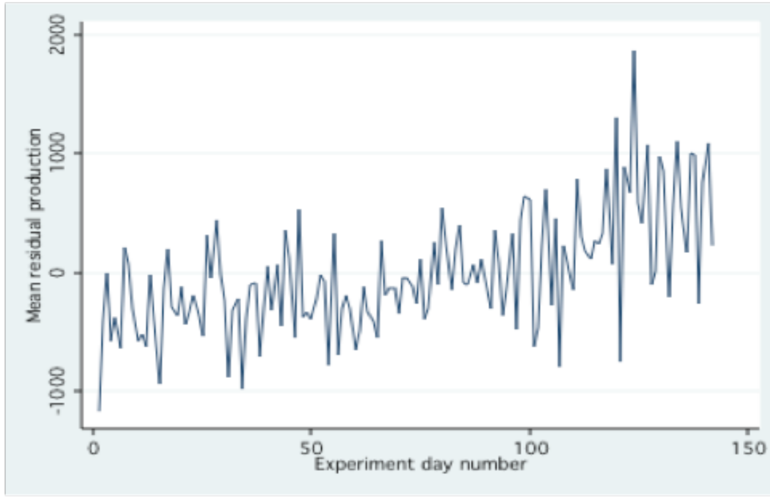

Figure D.4: Production Volatility

Notes: This figure shows how production varies across days in the Analysis Sample. Residual production is defined as the residual from a regression of production on a vector of worker dummies and lagged production controls. The x-axis measures the day number of the experiment and the y-axis measures the mean of the production residuals for that day.

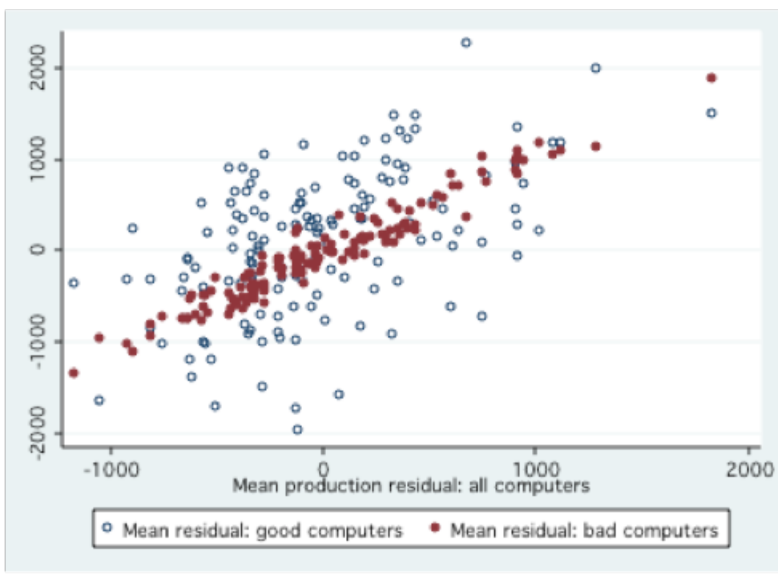

Figure D.5: Correlation of Computer Production Volatility with Office-level Shocks

Notes: This figure shows the relative production volatility of good and bad computers. Residual production is defined as the residual from a regression of production on a vector of worker dummies and lagged production controls. The x-axis measures the mean of the production residuals across all observations within a day. The open circles plot the mean production residual of workers assigned to good computers within a day. The closed circles plot the mean production residual of workers assigned to bad computers.

Table D.1: Randomization Example

	Worker 1	*Worker 2*	*Worker 3*	*Worker 4*	*Worker 5*
Day 1	Evening Choice	Evening Choice	Control	Target	Evening Choice
Day 2	Morning Choice	Control	Evening Choice	Morning Choice	Control
Day 3	Control	Evening Choice	Morning Choice	Control	Target
Day 4	Morning Choice	Morning Choice	Control	Morning Choice	Control
Day 5	Target	Target	Morning Choice	Target	Morning Choice
Day 6	Control	Control	Evening Choice	Control	Evening Choice
Day 7	Target	Control	Target	Target	Morning Choice
Day 8	Morning Choice	Evening Choice	Target	Control	Target
Day 9	Control	Target	Control	Morning Choice	Control
Day 10	Target	Target	Evening Choice	Evening Choice	Morning Choice
Day 11	Evening Choice	Morning Choice	Morning Choice	Evening Choice	Evening Choice
Day 12	Evening Choice	Morning Choice	Target	Evening Choice	Target
Day 13	Evening Choice	Control	Evening Choice	Evening Choice	Control
Day 14	Morning Choice	Target	Target	Morning Choice	Evening Choice
Day 15	Target	Evening Choice	Morning Choice	Morning Choice	Evening Choice
Day 16	Target	Control	Control	Target	Control
Day 17	Evening Choice	Morning Choice	Target	Evening Choice	Target
Day 18	Morning Choice	Morning Choice	Evening Choice	Evening Choice	Target
Day 19	Target	Control	Target	Target	Morning Choice
Day 20	Control	Target	Control	Control	Target
Day 21	Control	Morning Choice	Morning Choice	Morning Choice	Evening Choice
Day 22	Evening Choice	Evening Choice	Morning Choice	Target	Morning Choice
Day 23	Morning Choice	Evening Choice	Control	Control	Control
Day 24	Control	Target	Evening Choice	Control	Morning Choice

Notes: This table provides an example of the daily contract randomizations. The four contract treatments were: Assignment to the Control contract; Assignment to a Target (at either the low, medium, or high target level), Evening Choice, and Morning Choice. The table shows the contract treatment assignments for five workers over a 24-day period of the study. Workers were assigned to each of the four treatments exactly 3 times over each 12-day period. The order of the assignments was random and changed every 12 days. The vectors of treatment assignments were independent across workers.

Table D.2: Randomization Balance

	Payday Treatments			Contract Treatments				Total
	Tuesday Payday	Thursday Payday	Saturday Payday	Control Contract	Target Assignmt	Evening Choice	Morning Choice	
Proportion of observations	0.33	0.33	0.34	0.25	0.25	0.25	0.25	1.00
Number of observations	2,788	2,809	2,826	2,117	2,113	2,088	2,105	8,423

Notes: This table shows how many observations from the Analysis Sample are in each treatment cells. The Payday Treatments randomly assigned workers into one of three payday groups—Tuesday, Thursday, and Saturday—which determined on which day of the week they were paid their weekly earnings. Workers were randomized into four contract treatments: linear control contract, dominated contract with an exogenously set Target, evening choice (in which workers chose their preferred contract the evening before the workday), and morning choice (in which workers chose their preferred contract the morning of the workday). Workers were randomly assigned to each of the 4 contract treatments 3 times over every 12 workdays.

Table D.3: Decomposition of Treatment Effects

Dependent variable	*Workday Length*	*Workday Length*	*Production*	*Production*	*Workday Length | Attend=1*	*Workday Length | Attend=1*	*Production | Attend=1*	*Production | Attend=1*
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
	Panel A: Payday Treatments							
Payday	20	40	-60	26	2	10	-69	-11
	(4)***	(7)***	(38)	(55)	(3)	(4)**	(38)*	(53)
1 day before payday		39		119		17		102
		(7)***		(54)**		(4)***		(53)*
2 days before payday		30		135		16		104
		(8)***		(68)**		(5)***		(67)
3 days before payday		20		169		9		137
		(8)**		(65)**		(5)*		(65)**
4 days before payday		22		87		-0		35
		(9)***		(71)		(5)		(69)
5 days before payday		14		30		4		-2
		(9)		(69)		(5)		(67)
Minutes logged into software			17	17			12	12
			(0.39)***	(0.39)***			(1)***	(1)***
(Minutes logged into software)^2			-0.01	-0.01			-0.003	-0.003
			(0.00)***	(0.00)***			(0.001)**	(0.001)**
Observations	8423	8423	8423	8423	7376	7376	7376	7376
R2	0.27	0.28	0.86	0.86	0.30	0.31	0.84	0.84
Dependent variable mean	404	404	5337	5337	461	461	6094	6094

Table continued on next page...

Table D.3 (Continued)

Dependent variable	*Workday Length*	*Workday Length*	*Production*	*Production*	*Workday Length | Attend=1*	*Workday Length | Attend=1*	*Production | Attend=1*	*Production | Attend=1*
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Panel B: Contract Treatments								
Assignment to choice	9	-1	48	-13	6.17	5.77	35	-8
	(4)**	(5)	(35)	(46)	(2.69)**	(3.33)*	(34)	(43)
Assignment to a target	4	-6	84	4	6.37	2.47	84	121
	(5)	(6)	(41)**	(53)	(3.14)**	(4.02)	(40)**	(70)*
Assignment to a choice *		27		158		1.37		4
High payday difference		(9)***		(71)**		(5.68)		(51)
Assignment to a target *		27		205		11.47		217
High payday difference		(11)**		(84)**		(6.43)*		(83)***
Minutes logged into software			17	17			12	12
			(0.39)***	(0.39)***			(1)***	(1)***
(Minutes logged into software)^2			-0.01	-0.01			-0.003	-0.003
			(0.00)***	(0.00)***			(0.001)**	(0.001)**
Observations	8240	8240	8240	8240	7213	7213	7213	7213
R2	0.27	0.28	0.86	0.86	0.30	0.30	0.84	0.84
Dependent variable mean	404	404	5355	5355	461	461	6118	6118

Notes: This table decomposes treatment effects on production into minutes worked and production conditional on minutes worked. The dependent variable in columns (1)-(2) is workday length—the number of minutes between a worker's first login and last logout in the data entry software within a day. The dependent variable in columns (3)-(4) is production—the number of accurate fields entered in a day. Both variables equal zero on days workers are absent. The dependent variables in columns (5)-(6) and (7)-(8) are workday length conditional on being present and production conditional on being present, respectively. Panel A investigates impacts of paydays. Panel B investigates behavior under the contract treatments. High payday difference is a binary indicator that equals 1 if a worker's mean payday production difference—the difference in production on paydays and non-paydays under control, divided by overall mean production under control—is higher than the sample average. All regressions include controls for lag production and worker, date, and seat fixed effects. Robust standard errors are reported in parentheses.

Table D.4: Production Conditional on Attendance over the Pay Cycle

Dependent variable:	Production | Attend=1	Production | Attend=1
Sample	*Analysis Sample*	*Payday Sample*
	(1)	(2)
Payday	35	44
	(69)	(65)
1 day before payday	268	242
	(68)***	(64)***
2 days before payday	218	169
	(85)**	(81)**
3 days before payday	205	193
	(81)**	(78)**
4 days before payday	25	38
	(90)	(80)
5 days before payday	37	63
	(85)	(80)
Lag production controls	Yes	Yes
Worker fixed effects	Yes	Yes
Date fixed effects	Yes	Yes
Seat fixed effects	Yes	Yes
Observations	7376	10341
R2	0.74	0.73
Dependent variable mean	6093	6432

Notes: This table reports the effects of paydays and distance from paydays on worker production conditional on attendance, defined as the number of accurate fields completed by the worker in a day (conditional on the worker being present). Each column shows results from an OLS regression of the dependent variable on indicators of distance from payday, lag production controls, and worker, date, and seat fixed effects. Column (1) reports results from observations in the Analysis sample in which workers were present. Column (2) reports results from observations in the Payday sample in which workers were present. Robust standard errors are reported in parentheses.

Table D.5: Take-up of Dominated Contracts: Summary Statistics

Statistic	*Worker-day means*	*Worker means*
	(1)	(2)
Proportion choosing a positive target (conditional on attendance)	0.35	0.36
	(0.48)	(0.31)
Proportion choosing a positive target (target=0 when absent)	0.28	0.28
	(0.45)	(0.26)
Number of observations	4193 worker-days	101 workers

Notes: This table reports the proportion of times workers selected positive targets when assigned to a Choice treatment. The first row of the table summarizes choice behavior when the worker was present both the day before and the day of Choice assignment. The second row of the table includes absentee observations, and defines target choice to be 0 if a worker was absent the day before or day of Choice assignment. Column (1) presents means for the Analysis sample of 4,193 choice observations as a whole. Column (2) summarizes the worker means for each statistic, computed over the 101 workers that were assigned to Choice at least once during their employment. Standard deviations for each statistic are reported in parentheses.

Table D.6: Magnitude of Chosen Targets: Summary Statistics

Statistic	*All workers*	*High payday impact workers*	*Low payday impact workers*
	(1)	(2)	(3)
Probability of missing chosen target if assigned to control contract	0.091	0.118	0.073
	(0.120)	(0.146)	(0.096)
Proportion of times chosen target was actually missed	0.026	0.052	0.008
	(0.122)	(0.187)	(0.036)
Number of observations: worker-days	1132	514	618
Number of observations: workers	78	31	47

Notes: This table reports means and standard deviations of statistics that describe the aggressiveness of targets chosen by workers. Row 1 reports the probability that workers would have missed their chosen targets if they had been assigned to the control contract that day. This is computed as follows. For observations where workers were in attendance, we estimate a regression of production on worker, date, and computer fixed effects; lag production controls; payday distance dummies; contract assignment dummies; and log experience. For each observation in which a worker was assigned to Choice, selected a positive target, and was present, we predict the worker's production under the control contract on that day using the estimates from the above regression. To this predicted value, we add the worker's vector of residuals from the above regression to arrive at a vector of potential production values, which we fit to a lognormal distribution. Evaluating the CDF of this distribution at the chosen target level gives an estimate of the probability that the worker would have missed her chosen target under the control contract. Row 2 reports the actual proportion of times workers' production fell below their chosen targets. Both rows summarize worker means of these statistics. Column (1) presents these statistics for the workers that chose a positive target at least once and for whom the payday difference can be computed; columns (2) and (3) report these statistics separately for high and low payday difference workers. High payday impact workers are those whose payday impact measure—defined as the difference between mean production on paydays and non-paydays under Control divided by mean production under assignment to Control—is above the sample average.

Table D.7: Hourly Production around Target Thresholds

Observations	Self-Chosen Targets (Assignment to Choice)			Exogenous Targets (Assignment to a Target)		
Dependent variable	*Hourwise prodn*	*Hourwise prodn*	*Hourwise prodn*	*Hourwise prodn*	*Hourwise prodn*	*Hourwise prodn*
	(1)	(2)	(3)	(4)	(5)	(6)
3 hours before target reached		135	-13		189	48
		(19)***	(15)		(14)***	(12)***
2 hours before target reached		236	14		224	80
		(23)***	(87)		(16)***	(14)***
1 hour before target reached		432	88		398	176
		(22)***	(20)***		(16)***	(16)***
Hour in which target was reached	597	934	498	543	722	510
	(51)***	(33)***	(30)***	(25)***	(26)***	(25)***
1 hour after target reached		664	223		439	247
		(29)***	(28)***		(19)***	(19)***
2 hours after target reached		584	170		336	206
		(26)***	(24)***		(43)***	(22)***
3 hours after target reached		588	198		313	226
		(34)***	(31)***		(27)***	(25)***
4+ hours after target reached		270	133		54	164
		(21)***	(18)***		(20)***	(20)***
Worker's mean production under control in current hour			0.90			0.85
			(0.02)***			(0.02)***
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,948	13,948	13,948	20,174	20,174	20,174
R2	0.26	0.29	0.42	0.24	0.27	0.41
Dependent variable mean	847	847	847	730	730	730
F-test p-value: (Hour in which target reached) = (1 hour after target reached)		0.00	0.00		0.00	0.00

Notes: This table describes production levels around the hours in which workers achieved their targets.

The dependent variable in each regression is hourwise production, defined as production by worker i on day t in calendar hour h. For each date in the sample, there are 11 calendar hours, which correspond to the times the office was open and it was possible for employees to work: the first hour is 8-9 am and last is 6-7 pm. Production is coded as 0 if a worker did not work in a certain hour.

The table displays 8 binary variables that capture distance from when the target was reached. For example, the "Hour in which target was reached" is an indicator variable that equals 1 in the calendar hour in which a worker achieved her target for the day, and equals 0 otherwise. In cases where workers failed to achieve their target, the hour in which they reached their target is coded as the hour after the office closed (7-8 pm). "Worker's mean production under control in current hour" is a covariate that equals the sample mean of worker i's production in hour h on days in the Analysis Sample when worker i was assigned to the Control contract and was present. The last row of the table displays p-values from an F-test of whether the coefficient on Hour in which target was reached equals the coefficient on 1 after target reached in the regression shown in that column. Standard errors are corrected to allow for clustering by worker-day.

Columns (1)-(3) report estimates from days when workers selected positive targets under Assignment to Choice and were present. Columns (4)-(6) report estimates from days when workers were assigned to exogenous targets under Target Assignment and were present.

Table D.8: Test for Inter-temporal Substitution in Effort across Days

Dependent variable: Production

Observations: Today's assignment	*All contracts*	*Control & Choice*	*Control only*	*Control only*	*Control only*
	(1)	(2)	(3)	(4)	(5)
Assigned to choice yesterday	1	7	193		
	(67)	(78)	(163)		
Assigned to a target yesterday	19	12	243		
	(77)	(87)	(179)		
Assigned to a high target yesterday				372	
				(267)	
Assigned to choice the past 2 days in a row					54
					(144)
Assigned to a target the past 2 days in a row					-74
					(274)
Worker fixed effects	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	8423	6310	2117	2117	2117
R2	0.50	0.51	0.55	0.55	0.55
Dependent variable mean	5337	5311	5283	5283	5283

Notes: This table tests for evidence on whether there is inter-temporal substitution in worker effort across days. The dependent variable in each column is production, which is defined as the number of accurate fields completed by worker i on date t, and equals 0 if the worker was absent.

The observations in column (1) are all the observations in the Analysis sample. Column (2) restricts analysis to those observations in which a worker was assigned to the Control contract or to Choice on day t. Columns (3)-(5) further restrict analysis to only those observations in which a worker was assigned to the Control contract on date t.

Columns (1)-(3) show results from an OLS regression of production on an indicator for whether worker i was assigned to Choice on date t-1, and an indicator for whether worker i was assigned to a Target on date t-1. Column (4) shows an OLS regression of production on an indicator for whether worker i was assigned to the High Target on date t-1. Column (5) shows an OLS regression of production on an indicator for whether worker i was assigned to Choice on dates t-1 and t-2, and an indicator for whether worker i was assigned to a Target on date t-1 and t-2. All regressions include controls for worker, date, and seat fixed effects. Robust standard errors are reported in parentheses.

Table D.9: Predictive Power of Contract Assignment Probability

Dependent variable	*Production*	*Production*	*Production*
	(1)	(2)	(3)
Assignment to evening choice	150		161
	(69)**		(78)**
Assignment to morning choice	73		131
	(69)		(76)*
Assignment to low target	3		20
	(90)		(100)
Assignment to medium target	213		207
	(91)**		(102)**
Assignment to high target	334		405
	(150)**		(159)**
Pr(evening choice)		106	-65
		(161)	(182)
Pr(morning choice)		-155	-296
		(158)	(176)*
Pr(low target)		-63	-93
		(203)	(225)
Pr(medium target)		260	43
		(207)	(234)
Pr(high target)		-88	-487
		(373)	(396)
Worker fixed effects	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes
Lag production controls	Yes	Yes	Yes
Observations	8423	8423	8423
R2	0.59	0.59	0.59
Dependent variable mean	5337	5337	5337
F-test of joint significance of probability controls (p-value):		0.39	0.45

Notes: This table tests whether probability of assignment to a contract treatment predicts output. The dependent variable in each regression is production, which is defined as the number of accurate fields completed by worker i on date t, and equals 0 when a worker is absent. Column (1) reports results from an OLS regression of the dependent variable on dummies for each contract assignment treatment. Column (2) shows a regression of production on the probabilities of worker i receiving each contract treatment on date t. The probabilities are computed using the worker's previous assignments within the randomization block. Column (3) shows a regression of production on all the treatment assignment dummies and probability controls. All regressions include controls for lagged production and worker, date, and seat fixed effects. Robust standard errors are reported in parentheses. The bottom row of the table reports the p-value of an F-test of joint significance of the 5 probability controls in columns (2) and (3).

Table D.10: Contract Effects - High Payday Impact Workers

Dependent variable	*Production*	*Attendance*	*Production | Attend=1*
	(1)	(2)	(3)
Assignment to choice	395	0.044	144
	(116)***	(0.016)***	(73)**
Assignment to a target	452	0.023	376
	(121)***	(0.019)	(83)***
Worker fixed effects	Yes	Yes	Yes
Computer fixed effects	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes
Lag production controls	Yes	No	Yes
Observations	3216	3216	2706
R2	0.53	0.14	0.72
Dependent variable mean	4696	0.84	5581

Notes: This table estimates the treatments effects of the contracts on production for high payday impact workers. The table uses observations from only those workers in the Analysis sample whose payday production impact measure was above the sample average. The payday production impact measure is computed as the difference between the worker's mean production on paydays and non-paydays under Control divided by mean production under assignment to Control. The dependent variable in column (1) is production. Production is defined as the number of accurate fields completed by the worker in a day, and equals zero on days workers are absent. The dependent variable in columns (2) is attendance—a binary indicator that takes the value of 1 when workers are present and 0 otherwise. The dependent variable in column (3) is production conditional on attendance—it equals production when a worker is present and is missing otherwise. Each column shows the estimates from an OLS regression of the dependent variable on indicators for Choice and Target assignment. All regressions include fixed effects for each date in the sample, each worker in the sample, and each seating assignment. In addition, columns (1) and (3) also include controls for lagged production. Robust standard errors are reported in parentheses.

Table D.11: Heterogeneity in Payday Treatment Effects - Correlation with Contract Effects

Dependent variable	*Production*	*Attendance*	*Production*	*Attendance*	*Production*	*Attendance*
	(1)	(2)	(3)	(4)	(5)	(6)
Payday	156	0.044	-207	-0.007	-247	-0.013
	(166)	(0.025)*	(173)	(0.023)	(182)	(0.023)
High dominated contract take-up * Payday	-2	-0.001				
	(249)	(0.032)				
High choice treatment effects * Payday			802	0.107		
			(251)***	(0.035)***		
High target treatment effects * Payday					884	0.118
					(248)***	(0.033)***
Lag production controls	Yes	No	Yes	No	Yes	No
Observations	2117	2117	2117	2117	2115	2115
R2	0.65	0.24	0.65	0.24	0.65	0.24
Dependent variable mean	5283	0.87	5283	0.87	5186	0.87

Notes: This table tests whether behavior under the contract treatments predicts payday effects. All regressions report estimates from observations from the Analysis Sample in which the worker was assigned to the Control contract. Columns (1), (3), and (5) report estimates from an OLS model where the dependent variable is production. Production is defined as the number of accurate fields entered in a day and equals 0 when the worker is absent. Columns (2), (4), and (6) report estimates from an OLS Linear Probability Model in which the dependent variable is a binary indicator for attendance.

High dominated contract take-up rate is a dummy that equals 1 if the worker's take-up rate of the dominated contract—the proportion of times the worker chose a positive target under assignment to choice on nonpaydays—is above the mean take-up rate in the sample. Columns (1) and (2) show regressions of the dependent variable on an indicator for whether the day was the worker's assigned payday and an interaction of the payday dummy with the high take-up dummy.

High choice treatment effects is a dummy that equals 1 if the worker's choice differential—defined as the difference between mean production under assignment to Choice versus Control on nonpaydays—is above the sample average. Columns (3) and (4) show regressions of the dependent variable on the payday dummy and an interaction of the payday dummy with the high choice treatment effects dummy.

High target treatment effects is a dummy that equals 1 if the worker's choice differential—defined as the difference between mean production under assignment to a Target versus Control on nonpaydays—is above the sample average. Columns (5) and (6) show regressions of the dependent variable on the payday dummy and an interaction of the payday dummy with the high target treatment effects dummy.

All regressions include controls for worker, date, and seat fixed effects. Note that the number of observations changes between columns because contract treatments were assigned randomly and workers were in the sample for varying lengths of time. Each measure requires that we have observations for a worker on non-paydays under a certain set of contract assignments—assignment to Choice, assignment to Choice and Control, and assignment to a Target and Control, respectively. For some workers, these measures cannot be computed and they are therefore not included in the analysis. Robust standard errors are reported in parentheses.

Table D.12: Sensitivity of Bad Computers to Production Shocks

Dependent variable	*Production residual*	*Production residual*
	(1)	(2)
Bad computer indicator	-139	-130
	(72)*	(79)
Mean of day's residual production	0.454	
	(0.097)***	
Bad computer indicator * Mean of day's residual production	0.261	
	(0.116)**	
Mean of day's residual production for bad computers		0.373
		(0.098)***
Bad computer indicator * Mean of day's residual production for bad computers		0.263
		(0.117)**
Mean of day's residual production for good computers		0.088
		(0.087)
Bad computer indicator * Mean of day's residual production for good computers		-0.020
		(0.099)
Observations	8423	8423
R2	0.024	0.024
Dependent variable mean	0.000	0.000

Notes: This table reports estimates on whether workers assigned to bad computers are more sensitive to productivity shocks. The table shows results from OLS regressions in which the dependent variable is the residual from a regression of production on a vector of worker dummies and lagged production controls. Column (1) shows a regression of the production residual on: an indicator for whether the worker was assigned to a computer identified by workers and management as being likely to become slow during network fluctuations; the mean residual production for that day (by averaging the dependent variable across the observations for that day, excluding the worker's own observation); and an interaction of the 2 covariates. Column (2) shows a regression of the production residual on the bad computer indicator; the mean of residual production that day, computed separately for the good and bad computers (excluding the worker's own observation); and interactions of these latter two covariates with the bad computer indicator. Standard errors are clustered by computer (seat) assignment.

Table D.13: Correlation between Contract Quiz Score and Demand for Dominated Contracts

Dependent Variable:	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*
	(1)	(2)	(3)	(4)	(5)	(6)
Scored 100% on quiz	175	0.001				
	(182)	(0.058)				
Quiz score (percentage)			280	-0.11		
			(682)	(0.21)		
Education (years)					233	0.059
					(49)***	(0.015)***
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3582	3582	3582	3582	3582	3582
R2	0.10	0.11	0.10	0.11	0.15	0.14
Dependent variable mean	804	0.29	804	0.29	804	0.29

Notes: This table tests whether a worker's score on the contracts quiz, which tested comprehension of the contracts treatments and was administered before workers began randomizations, is correlated with demand for the dominated contract. Columns (1), (3), and (5) show estimates of OLS models in which the dependent variable is the target level selected by a worker when assigned to Choice. Columns (2), (4), and (6) show estimates of OLS Linear Probability Models in which the dependent variable is a binary indicator for whether the worker selected a target above zero when assigned to a Choice treatment. Both dependent variables are defined as 0 if a worker was absent the day before or day of Choice assignment. Columns (1)-(2) show regressions of the dependent variable on an indicator that equals 1 if the worker scored 100% on the contracts quiz, and equals 0 if the worker scored below 100%. Columns (3)-(4) show regressions of the dependent variable on the continuous measure of the worker's percentage score in the quiz. Columns (5)-(6) show regressions of the dependent variable on years of education. All regressions include date and seat fixed effects. Standard errors are clustered by worker. In 3 instances, workers who initially scored below 80% on the quiz were re-trained and retook the quiz before the start of contract randomizations; their quiz score in the data equals their score on their second quiz. Due to a clerical error, the quiz was not administered to all workers. Observations are for those workers in the Analysis sample that took the contracts quiz.

Table D.14:
Test for Income Targeting: Production Increases Before Paydays

Dependent variable	*Prodn*	*Prodn*	*Prodn*
	(1)	(2)	(3)
Payday	200	110	105
	(90)**	(145)	(129)
1 workday before payday	331	181	268
	(99)***	(147)	(132)**
Assigned to Target yesterday		-99	-76
		(103)	(92)
Payday *		259	179
Assigned to Target yesterday		(191)	(169)
1 workday before payday *		341	122
Assigned to Target yesterday		(205)*	(185)
Lag production controls	Yes	No	Yes
Observations	4344	4344	4344
R2	0.61	0.52	0.61
Dependent variable mean	5352	5352	5352

Notes: This table tests for income targeting by workers within their payweeks. All regressions report estimates from observations from the Analysis Sample in which the worker was assigned to the Control contract or to a Target on the previous day. All columns report estimates from an OLS model where the dependent variable is production. Production is defined as the number of accurate fields entered in a day and equals 0 when the worker is absent. Payday is an indicator variable for whether the current day is the worker's assigned payday. 1 workday before payday is an indicator for whether the current day is 1 workday before the worker's assigned payday. Assigned to Target yesterday is an indicator that equals 1 if the worker was assigned to a target yesterday, and equals 0 if the worker was assigned to the control contract yesterday. All regressions include controls for worker, date, and seat fixed effects. Columns (1) and (3) also include controls for lag production. Robust standard errors are reported in parentheses.

References

References for Introduction and Chapter 1

Akerlof, George. 2002. "Nobel Lecture: Behavioral Macroeconomics and Macroeconomic Behavior," *American Economic Review*.

Akerlof, George, William Dickens, George Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity*. 1: 1-76.

Akerlof, George, and Janet Yellen. 1990. "The Fair Wage-Effort Hypothesis and Unemployment," *Quarterly Journal of Economics*. 255-83.

Bardhan, Pranab. 1973. "Size, Productivity, and Returns to Scale: An Analysis of Farm-level Data in Indian Agriculture." *Journal of Political Economy*. 81(6): 1370-1386.

Behrman, Jere. 1999. "Labor Markets in Developing Countries," in ed. O. Ashenfelter and D. Card. (eds), *Handbook of Labor Economics*, Vol 3: 2859-2939.

Benjamin, Dwayne. 1992. "Household Composition, Labor Markets, and Labor Demand: Testing for Separation in Agricultural Household Models." *Econometrica*. 60(2): 287-322.

Bewley, Truman F. 1999. *Why Wages Don't Fall During a Recession*. Harvard University Press.

Blanchard, Olivier. 1990. "Why Does Money Affect Output? A Survey", in Ben Friedman and Frank Hahn, editors, *Handbook of Monetary Economics*. Vol 2.

Blinder, Alan and Don Choi. 1990. "A Shred of Evidence on Theories of Wage Stickiness." *Quarterly Journal of Economics*, pp. 1003-1015.

Card, David. 1990. "Unexpected Inflation, Real Wages, and Employment Determination in Union Contracts." *American Economic Review*.

Card, David and Dean Hyslop. 1997. "Does Inflation 'Grease the Wheels of the Labor Market'?", in Christina D. Romer and David H. Romer, editors, *Reducing Inflation: Motivation and Strategy*. University of Chicago Press.

Clarida, Richard, Jordi Galì, and Mark Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature*. 37: 1661-1707.

Dickens, William et al. 2006. "How Wages Change: Micro Evidence from the International Wage Flexibility Project." National Bank of Belgium Working Paper No. 96.

Fehr, Ernst, Georg Kirchsteiger and Arno Riedl. 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *Quarterly Journal of Economics*. 108(2): 437-459.

Fehr, Ernst and Armin Falk. 1999. "Wage Rigidity in a Competitive Incomplete Contract Market." *Journal of Political Economy*. 107(1): 106-134.

Fehr, Ernst, Lorenz Goette, and Christian Zehnder. 2009. "A Behavioral Account of the Labor Market: The Role of Fairness Concerns." *Annual Review of Economics*. 1: 355-384.

Galì, Jordi. 2002. "New Perspectives on Monetary Policy, Inflation, and the Business Cycle." NBER Working Paper No. 8767.

Gneezy, Uri and John List. 2006. "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments." *Econometrica*. 74(5): 1365-1384.

Greenwald, Bruce and Joseph Stiglitz. 1987. "Keynesian, New Keynesian and New Classical Economics." *Oxford Economic Papers*. 39(1): 119-133.

Hall, Robert and Paul Milgrom. 2008. "The Limited Influence of Unemployment on the Wage Bargain." *American Economic Review*. 98(4): 1653-1674.

Holzer, Harry and Edward Montgomery. 1993. "Asymmetries and Rigidities in Wage Adjustments by Firms." *Review of Economics and Statistics*. 75(3): 397-408.

Jayachandran, Seema. 2006. "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries." *Journal of Political Economy*. 114(3).

Kahn, Shulamit. 1997. "Evidence of Nomninal Wage Stickiness from Microdata." *American Economic Review*. 87(5).

Kahneman, Daniel, Jack Knetsch, and Richard Thaler. 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *American Economic Review*, 728-741.

Kochar, Anjini (1999). "Smoothing Consumption by Smoothing Income: Hours-of-Work Responses to Idiosyncratic Agricultural Shocks in Rural India." Review of Economics and Statistics, 81(1): 50-61.

McLaughlin, Kenneth. 1994. "Rigid Wages?" *Journal of Monetary Economics*. 34: 1983-414.

Paxson, Christina. 1992. "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand." *American Economic Review*. 82(1).

Rosenzweig, Mark. 1980. "Neoclassical Theory and the Optimizing Peasant: An Econometric Analysis of Market Family Labour Supply in Developing Countries." *Quarterly Journal of Economics.* (94): 31–56.

Rosenzweig, Mark. 1988. "Labour Markets in Low-income Countries", in Chenery H. and Srinivasan T. N. (eds), *Handbook of Development Economics*, Vol. I, North-Holland, Amsterdam.

Rosenzweig, Mark and Kenneth Wolpin. 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India." *Journal of Political Economy.* 101(2).

Shafir, Eldar, Peter Diamond, and Amos Tversky. 1997. "Money Illusion." *Quarterly Journal of Economics.* 112(2): 341-374.

Singh, I., L. Squire and J. Strauss, eds. 1986. *Agricultural Household Models: Extensions, Applications, and Policy.* The World Bank, Washington, DC.

Tobin, James. 1972. "Inflation and Unemployment." *American Economic Review.* 62: 1-18.

Townsend, Robert. 1994. "Risk and Insurance in Village India." *Econometrica.* 62(3): 539-591.

Udry, Chris. 1996. "Efficiency and Market Structure: Testing for Profit Maximization in African Agriculture." Mimeo. Northwestern University, Evanston, IL.

References for Chapter 2

Alchian and Demsetz. 1972. Production, information costs, and economic organization. *American Economic Review*, 62 (December): 777-795.

Ariely, Dan and Klaus Wertenbroch. 2002. "Procrastination, Deadlines, and Performance: Self-control by Precommitment." *Psychological Science*, 13(3): 219-224.

Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines." *Quarterly Journal of Economics*, 121(2): 635-672.

Bandiera, Oriana, Iwan Barankay and Imran Rasul. 2007. "Incentives for Managers and Inequality Among Workers: Evidence From a Firm Level Experiment." *Quarterly Journal of Economics*, 122: 729-75.

Banerjee, Abhijit and Esther Duflo. 2006. "The Economic Lives of the Poor." MIT Working Paper No. 06-29.

Banerjee, Abhijit, and Sendhil Mullainathan. 2009. "The Shape of Temptation: Implications for the Economic Lives of the Poor." Mimeo, Harvard University.

Benhabib, Jess, Alberto Bisin, and Andrew Schotter. 2007. "Present Bias, Quasi-Hyperbolic Discounting, and Fixed Costs."

Bernheim, B. Douglas, Debraj Ray and Sevin Yeltekin, 2011, "Poverty and Self-Control", Working paper.

Burger, Nicholas, Gary Charness, and John Lynham. 2008. "Three Field Experiments on Procrastination and Willpower." Available: www.bec.ucla.edu/Charness_Procrastination.pdf

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *Quarterly Journal of Economics,* 111: 408-441.

Cheung, S.N.S., 1969. The Theory of Share Tenancy. University of Chicago Press, Chicago, pp: 188.

Cheung, Steven N. S.1983, "The Contractual Nature of the Firm," *Journal of Law and Economics*, Vol. 26, No. 1. April, pp. 1-21.

Clark, Gregory. 1994. "Factory Discipline." *Journal of Economic History*, 54: 128-163.

Chwe, Michael S.1990. Why Were Workers Whipped? Pain in a Principal-Agent Model, *Economic Journal,* 100, pps. 1109-1121.

DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47: 315-372.

DellaVigna, Stefano and Ulrike Malmendier. 2004. "Contract Design and Self-Control: Theory and Evidence." *Quarterly Journal of Economics*, 119: 353-402.

Fehr, Ernst and Lorenz Goette. 2008. "Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment." *American Economic Review*, 97(1): 298-317.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40(2): 351- 401.

Fudenberg, Drew and David Levine. 2006. "A Dual Self Model of Impulse Control." *American Economic Review*, 96: 1449-1476.

Giné, Xavier & Goldberg, Jessica & Yang, Dean, 2010. "Identification Strategy: A field Experiment on Dynamic Incentives in Rural Credit Markets," Policy Research Working Paper Series 5438, The World Bank.

Gneezy, Uri, and John List. 2006. "Putting Behavioral Economics To Work: Testing For Gift Exchange In Labor Markets Using Field Experiments", *Econometrica*, 74(5): 1365—1384.

Grossman, Sanford J., and Oliver D Hart 1983. "An Analysis of the Principal-Agent Problem", *Econometrica*, Vol. 51, No. 1 (Jan., 1983), pp. 7-45.

Gul, Faruk and Wolfgang Pesendorfer. 2001. "Temptation and Self-Control," *Econometrica*, 69(6): 1403-1435.

Holmstron, Bengt, 1979. “Moral hazard and Observability.” *The Bell Journal of Economics*, Vol. 10, No. 1, (Spring, 1979), pp. 74-91

Hossain, Tanjim and John List. 2009. “The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations.” NBER Working Paper No. 15623.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainthan. 2010. “Self-Control and the Development of Work Arrangements.” *American Economic Review Papers and Proceedings*, 100(2): 624-628.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainthan. 2012. “Works Well with Others: Peer Effects and Social Motivation.” Mimeo, Harvard University.

Laibson, David I. 1996. "Hyperbolic Discount Functions, Undersaving, and Savings Policy," NBER Working Paper 5635, National Bureau of Economic Research, Inc

Laibson, David. 1997. “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics*, 112(2): 443-477.

Loewenstein, George & Prelec, Drazen, 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation," *The Quarterly Journal of Economics*, MIT Press, vol. 107(2), pages 573-97, May

O’Donoghue, Ted and Matthew Rabin. 2006. “Incentives and Self-control.” *Econometric Society Monographs*, 42: 215-245.

O’Donoghue, Ted and Matthew Rabin. 2001. “Choice and Procrastination,” *Quarterly Journal of Economics*, 116(1): 121-160.

O’Donoghue, Ted and Matthew Rabin. 1999. “Doing It Now or Later,” *American Economic Review*, 89(1): 103-124.

O’Donoghue, Ted and Matthew Rabin. 1999. "Incentives for Procrastinators", *Quarterly Journal of Economics*, 114(3), 769-816.

Shearer, Bruce. 2004. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment," *Review of Economic Studies*, 71(2): 513-34.

References for Chapter 3

Banerjee, Abhijit and Esther Duflo. 2006. “The Economic Lives of the Poor”, *Journal of Economic Perspectives.* 21(1): 141-167.

Bardhan, Pranab and Chris Udry. 1999. *Development Microeconmics.* New York: Oxford University Press.

Behrman, Jere. 1999. "Labor Markets in Developing Countries," in ed. O. Ashenfelter and D. Card. (eds), *Handbook of Labor Economics*, Vol 3: 2859-2939.

Benjamin, Dwayne. 1992. “Household Composition, Labor Markets, and Labor Demand: Testing for Separation in Agricultural Household Models.” *Econometrica.* 60(2): 287-322.

Duflo, Esther, Michael Kremer and Jon Robinson. 2010. “Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya.” *American Economic Review*, forthcoming.

Jayachandran, Seema. 2006. “Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries.” *Journal of Political Economy.* 114(3).

Kaur, Supreet. 2012. “Nominal Wage Rigidity in Village Labor Markets.” Mimeo, Harvard University.

Rosenzweig, Mark. 1988. “Labour Markets in Low-income Countries”, in Chenery H. and Srinivasan T. N. (eds), *Handbook of Development Economics*, Vol. I, North-Holland, Amsterdam.

Shaban, Radwan. 1987. “Testing between Competing Models of Sharecropping,” *Journal of Political Economy.* 95(5): 893-920.

Singh, I., L. Squire and J. Strauss, eds. 1986. *Agricultural Household Models: Extensions, Applications, and Policy.* Washington, DC: The World Bank.

Udry, Chris. 1996. “Efficiency and Market Structure: Testing for Profit Maximization in African Agriculture.” Mimeo, Northwestern University.