The London School of Economics and Political Science

Essays on Contract Design in Behavioral and Development Economics

Jonathan de Quidt

February 2014

A thesis submitted for the degree of Doctor of Philosophy in Economics at the London School of Economics

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 50,000 words.

Statement of Conjoint Work

I certify that chapters 2 and 3 of this thesis were co-authored with Thiemo Fetzer and Maitreesh Ghatak. I contributed 50% of the work in chapter 2, and 33% of the work in chapter 3.

Abstract

This thesis consists of three chapters that fall under the broad banner of contract theory, applied to topics in behavioral and development economics.

Empirically, labor contracts that financially penalize failure induce higher effort provision than economically identical contracts presented as paying a bonus for success. This is puzzling, as penalties are infrequently used in practice. The most obvious explanation is selection: loss averse agents are unwilling to accept such contracts. In the first chapter, I formalize and experimentally test this intuition. Surprisingly, I find that workers were 25 percent *more* likely to accept penalty contracts, with no evidence of adverse selection. Penalty contracts also increased performance on the job by 0.2 standard deviations. Finally, I outline extensions to the basic theory that are consistent with the main results.

The second chapter analyzes the effect of market structure in microfinance on borrower welfare and the types of contracts used. We find that market power can have severe implications for borrower welfare, while despite information frictions, competition delivers similar borrower welfare to non-profit lending. We also find that for-profit lenders are less likely to use joint liability than non-profits, which is consistent with some empirical stylized facts suggesting a decline in use of joint liability. We simulate the model to evaluate quantitatively the importance of market structure for borrower welfare.

The third chapter contrasts individual liability lending with and without groups to joint liability lending, motivated by an apparent shift away from joint liability lending. We show under what conditions individual liability can deliver welfare improvements over joint liability, conditions that depend on the joint income distribution and social capital. We then show that lower transaction costs that mechanically favor group lending may also encourage the creation of social capital. Finally, we again simulate the model to quantify our welfare conclusions.

Acknowledgements

The work contained in this thesis has benefited from the help and support of many people. I am immensely grateful to Maitreesh Ghatak, Greg Fischer and Gharad Bryan for their guidance and encouragement throughout, as well as to Oriana Bandiera, Tim Besley, Robin Burgess, Matthew Levy, and Gerard Padró i Miquel and many others for helpful comments and discussions.

I gratefully acknowledge financial support from the Economic and Social Research Council, and funding from STICERD for the experiment described in the first Chapter.

Thanks go to all my friends and colleagues, especially Michael Best, Konrad Burchardi, Can Celiktemur, Jonathan Colmer, Kara Contreary, Laura Derksen, Erika Deserranno, Thiemo Fetzer, Jason Garred, Selim Gulesci, Sam Marden, Luis Martinez, Ana McDowall, Tara Mitchell, Fabio Pinna, Miriam Sinn, Munir Squires, Pedro Souza, and Dimitri Szerman.

This thesis would not exist without the support of my parents, Marion and Simon, and parents-in-law, Lois and Dave, especially in the last few months with childcare, moving home and country. And most of all I thank my beautiful wife and closest friend Nyree for her boundless patience, love and sense of fun.

Little Eliana gets the last word, although she can't read it just yet.

Contents

List of Tables

List of Figures

Preface

This thesis consists of three chapters, each of which addresses a separate issue in applied contract theory, and which reflect a common interest in how effective contract design can improve welfare.

In the first chapter I describe the results of a randomized experiment motivated by the question of why firms seem to prefer to present financial incentives to workers in terms of bonuses rather than penalties. This is somewhat puzzling given the empirical evidence that penalty contracts (contracts that offer a high base wage but deduct pay if a performance target is not met) are more effective at motivating workers than economically equivalent bonus contracts (with a low base wage and bonuses for meeting the target). This finding is usually attributed to loss aversion on the part of the worker. The low base pay in the bonus contract sets the worker's reference point low, so the bonus payments fall in the gain domain, while the high base pay of the penalty contract sets a high reference point, such that the penalty falls in the loss domain. Since losses loom larger than gains, workers exert more effort to avoid a penalty than to achieve an equivalent bonus.

The most obvious explanation for why firms prefer bonus contracts is selection: loss averse workers will be unwilling to accept penalty contracts. I formalize this intuition, then experimentally test it using a two-stage randomized experiment in which workers are recruited under economically equivalent contracts that are framed either in terms of penalties or bonuses. Surprisingly I find that workers were 25 percent *more* likely to accept the penalty framed contract. Despite this, I find no evidence of advantageous or adverse selection induced by the framing treatment, indeed the distributions of payoff-relevant characteristics, such as measures of ability, are essentially identical between those who accepted the bonus framed offer and those who accepted the penalty framed offer. Third, consistent with the existing literature, the penalty contract improved on the job performance by 0.2 standard deviations.

I outline two plausible extensions to the model that bring its predictions more in line with the data. One explanation is that the high reference point under the penalty contract also makes the job offer feel more attractive, or the outside option less attractive. Alternatively, it may be that workers face a self-control problem, and like penalty contracts because they motivate them to exert higher effort. Finally, I argue that more research is needed on the long-term effects of incentive framing if we want to understand why firms prefer not to use penalties.

The second and third chapters analyze a pair of related problems in the economics of microfinance, in particular related to group lending. Group lending refers to the common practice in microfinance of disbursing loans and collecting repayments at group meetings of several borrowers, rather than individually. It may go hand in hand with joint liability lending, in which case group members also act as guarantors for one another's loans.

In chapter two, my coauthors and I explore the consequences for borrower welfare of different market structures in microfinance. The majority of theoretical work in microfinance assumes a single, benevolent non-profit lender, and studies how the lender is able to use innovative contract structures, such as group lending, to overcome asymmetric information or weak enforcement problems that otherwise lead to credit rationing. Joint liability lending can harness social capital, encouraging mutual insurance and increasing repayment rates. However, in the hands of a lender with market power these instruments may actually be harmful to the borrowers. We show that a lender with market power can leverage the joint liability mechanism to charge higher interest rates to borrowers with more social capital, making them worse off.

Next we ask to what extent competition solves this problem, under the realistic assumption of weak information sharing between lenders. This information structure limits lenders' ability to recover loans, because borrowers can simply default and move to another lender. As a consequence, the competitive market is characterized by credit rationing. However, higher social capital relaxes the borrowers' repayment constraints, enabling entry and moving welfare toward the first best. Nevertheless, the welfare ranking of competition and monopoly lending is ambiguous, depending on the relative severity of credit rationing and social capital exploitation.

Finally, we find that for over reasonable parameter ranges both monopolistic and competitive for-profit lenders are less likely to offer joint liability loan contracts than non-profits. This is consistent with empirical stylized facts: a recent decline in joint liability lending which goes hand-in hand with growth in commercialization of the microfinance industry.

To assess the empirical relevance of the mechanisms studied, we then simulate the model for estimated values of the key parameters. We confirm that welfare losses under monopoly are large, but find that welfare losses due to information frictions in competition are relatively modest. We also confirm that for the parameters we estimate, commercial lenders are indeed less likely to use joint liability than a non-profit lender.

The third chapter is also motivated by the previously mentioned decline in joint liability lending. Some microfinance lenders are switching from joint liability to individual liability lending, but retaining the group lending structure. The model is motivated by recent findings that individual liability lending to groups can achieve the same high repayment rates as joint liability lending. We show that provided borrowers have sufficient social capital, they are able to sustain what we call "implicit joint liability", guaranteeing one another's loans without a stipulation from the lender that they do so. Furthermore, implicit joint liability has the potential to improve on traditional joint liability, since it avoids one perverse effect of joint liability lending, namely that a borrower who cannot afford to guarantee her unsuccessful partner's loan may default on her own loan as well; this does not happen under implicit joint liability. However, implicit joint liability is more demanding in terms of social capital, implying that abandoning explicit joint liability may be more successful in some contexts than others.

Next, motivated by evidence that social capital can be built through repayment group meetings, we analyze an extension in which the transaction cost structure of group lending leads borrowers to invest in social capital creation, increasing welfare. Finally, once again we simulate the model to evaluate the welfare effects of the contract types considered.

Chapter 1

Your Loss Is My Gain: A Recruitment Experiment With Framed Incentives[1]

Consider two otherwise identical job offers, of which the first pays a base wage of \$100, plus a bonus of \$100 if a performance target is reached, while the second pays a base wage of \$200, minus a penalty of \$100 if the target is not reached. Rational agents will behave identically under either of these two contracts. However, a large body of empirical evidence suggests that behavior does respond to framing manipulations such as this. In particular several lab and field studies find that workers exert higher effort under the penalty framed contract than the bonus framed one. The leading explanation for these findings is reference dependence and loss aversion (Kahneman and Tversky (1979)), where the frame influences the reference point. The low base pay of the bonus frame sets her reference point low, so bonuses are perceived as gains, while the high base pay of the penalty frame sets her reference point high, so penalties are perceived as losses. Since losses loom larger than gains, penalties are more motivational than bonuses.

This finding raises a puzzle. If penalties are more motivational than bonuses, why are they not more widely used?[2] The most obvious explanation is that while penalties are effective motivators of existing workers, they unlikely to be accepted at the recruitment stage. I formalize this intuition in a simple model based on Kőszegi and Rabin (2006, 2007), showing that a forward-looking loss averse agent who is subject to framing effects will be less willing to accept a penalty contract than an equivalent bonus contract, because penalties increase her reference point, reducing her utility in all states of the world. Under the bonus contract she feels elated when successful and not too disappointed when unsuccessful, while under

[1]I gratefully acknowledge financial support from STICERD for the experiment described in this chapter.

[2]See e.g. Baker et al. (1988), Lazear (1995). Although I am not aware of any datasets addressing this issue, a glance through any job vacancy listing reveals many jobs that specify potential bonuses and almost no mention of penalties.

the penalty contract she feels only contented when successful and very disappointed when unsuccessful, hence the penalty contract is less attractive.[3]

The main contribution of this paper is a real-effort randomized experiment with 1,450 participants, designed to test this intuition. I use a two-stage design that separates selection and incentive effects of penalty framing relative to bonus framing. It consists of a common first stage in which I measure workers' types, followed by a second stage a few days later in which workers are offered framed incentive contracts and choose whether or not to accept. In each stage workers perform a data entry task and are assessed on their accuracy.

In contrast with the theoretical predictions, workers offered a penalty framed contract were 25 percent *more* likely to accept than those offered an equivalent bonus contract. Both the direction and magnitude of this effect are striking. By way of comparison, doubling the size of the contingent component of pay had a small and statistically insignificant effect on acceptance rates.

Second, despite the large effect on recruitment, the penalty contract did not lead to adverse or advantageous selection, indeed the distributions of the main payoff-relevant observables are essentially identical between those who accepted the bonus and those who accepted the penalty. Reported reservation wages are slightly higher on average among participants who accepted the penalty contract, consistent with the penalty frame relaxing participation constraints.

Third, consistent with the existing literature, performance on the incentivized task was significantly higher under the penalty treatment, around 6 percent higher accuracy on the data entry task (0.2 standard deviations). The coefficient estimate is unchanged when including controls, consistent with the absence of selection on observables. The improvement in performance is observed across the whole distribution, so is not driven by a small number of strongly affected individuals. Participants exerted more effort, measured as time spent, under the penalty frame, although this measure is very noisy so the coefficient estimates are not always statistically significant. Once again, it is striking that doubling the contingent component of pay had a smaller and statistically insignificant effect on performance.

While I cannot of course rule out selection on unobservables, the broad range of controls gives confidence that the observed effect is an incentive effect and not driven by selection. Further confidence is gained from the fact that I do observe significant selection on ability into the incentivized second stage of the experiment (low ability participants are less likely to accept the job offer under both frames), and that increasing the non-contingent component of pay in the second stage led to adverse selection, suggesting that the experimental design is able to detect selection effects where they are present.

In addition to controlling for selection, the experiment is designed to rule out two key confounds, that workers might infer something about the task from the

[3]More generally, in almost any model where she chooses her effort provision optimally, a manipulation that leads her to exert more effort without changing the economic terms of the contract must make her worse off.

contract they are offered, and that inattention when reading the job offer might make the penalty contract seem relatively attractive. I also check for inference by testing whether workers perceived the task to be more or less difficult under the penalty frame, finding no difference.

Since the basic theory cannot explain the relative popularity of the penalty contract I outline two extensions that bring the model in line with the data. One possibility is that the high reference point under the penalty contract simply makes the job appear more lucrative, or makes the outside option appear less lucrative. Alternatively, it could be that workers like penalty contracts because they enable them to overcome a self-control problem: the worker would like to exert more effort and the penalty motivates her to do so. The first extension suggests that workers are making a mistake, and may ultimately come to regret their choice, while in the second the penalty contract takes the form of a commitment device that can increase welfare.[4]

There are two possible responses to the results. It may indeed be that firms can gain by increasing their use of penalties, including pre-announcing them at the recruitment stage. This is most likely to be the case in contexts similar to the experimental environment, short-term recruitment of workers to perform routine tasks with minimal screening. For example, hiring one-off data entry workers, field officers, tradesmen, event or campaign staff.

The second response is to go back to the original question. The penalty contract recruited more workers, who then exerted greater effort, why then are they not more widely used? The results suggest that selection is not the answer. Perhaps the explanation lies in the fact that while most employment relationships are long-term, the effects of framing manipulations may be short-lived. Over time, workers' reference points are likely to adjust, eroding the performance advantages and perhaps leading workers who were recruited under the penalty frame but would not have accepted the bonus frame to quit.[5] It is possible too that actually incurring a penalty (as opposed to not receiving a bonus) is particularly discouraging, leading to lower effort and higher attrition in future periods.[6] I show that the framing effect did not wear off over the course of the one incentivized stage of effort provision, but more theoretical and empirical work on the long-term effects of incentive framing is needed to understand whether these conjectures are correct.

Turning to related literature, existing work on incentive framing focuses on incentive effects, that is, its effect on effort provision among a sample of already-recruited workers or lab subjects.[7] Hossain and List (2012) in the field, and Ar-

[4]Recent evidence suggests that workers may prefer financially dominated incentive schemes that incorporate a similar form of commitment, see Kaur et al. (2013).

[5]Druckman (2001), Hossain and List (2012) and Jayaraman et al. (2014) discuss the issue of short-lived "behavioral" effects.

[6]This point relates to the literature on wage cuts, see e.g. Fehr et al. (2009) and also Mas (2006).

[7]The one exception of which I am aware is Luft (1994). In her study, lab participants indicated a preference between each of an increasing sequence of fixed payments, and a contingent contract which was either bonus or penalty framed. The mean valuation of the bonus contract was higher, in contrast to my results. However, the sample size is only 27 so the difference is not significant. Brooks et al. (2013) study only penalty framed incentives, varying the size of the target below which penalties are

mantier and Boly (2012), Hannan et al. (2005) and Church et al. (2008) in the lab consistently find higher effort provision under penalty incentives than equivalent bonus incentives.[8] However, Fehr and Gächter (2002) find in a buyer-seller experiment that penalty-framed performance incentives led to more shirking among sellers than equivalent bonus-framed offers.

The paper relates to the literature on behavioral contract design. For example, DellaVigna and Malmendier (2004) explore contract design for time-inconsistent agents, Von Thadden and Zhao (2012) for unaware agents and de Meza and Webb (2007) and Herweg et al. (2010) for loss-averse agents, the latter applying the model of Kőszegi and Rabin (2006, 2007). Hilken et al. (2013) theoretically analyze an incentive framing problem related to the one outlined in this paper. Empirical papers studying the effect of loss aversion on effort provision (without framing effects) include Camerer et al. (1997), Farber (2005, 2008), Crawford and Meng (2011), Pope and Schweitzer (2011), Abeler et al. (2011) and Gill and Prowse (2012).

Finally, it fits into the smaller empirical literature on selection effects of employment contracts. Lazear (2000), Eriksson and Villeval (2008) and Dohmen and Falk (2011) find that performance pay tends to select in high-ability types, a result that I also observe in my experiment, while Guiteras and Jack (2014) observe adverse selection. Propper and Van Reenen (2010) and Dal Bó et al. (2012) find that high (relative) wages attract higher quality employees, while I observe suggestive evidence that increasing the non-contingent component of pay in my experiment attracted lower ability workers. Deserranno (2014) studies the effect on worker selection of manipulating earning expectations by reporting different moments of the earning distribution.

The remainder of the paper is as follows. Section 1.1 sets up the basic theoretical framework and derives three testable predictions. Section 1.2 outlines the experiment design, the experimental platform (Amazon Mechanical Turk), and the data collected. Section 1.3 describes the main results on acceptance rates, selection and performance. Section 1.4 presents secondary results including two tests of possible mechanisms: inference and inattention. Section 1.5 discusses extensions to the model that bring it closer in line with the data, and areas where further theoretical and empirical work would be particularly valuable. Finally, Section 1.6 concludes. Two Appendices contain additional results and experimental details.

1.1 A simple model

Consider a standard moral hazard problem in which a principal (P) wants to hire an agent (A) to perform a task, the success of which depends on the agent's effort.

incurred. They show that setting the target extremely high reduces acceptance of the job offer and performance.

[8] Fryer et al. (2012) test a closely related but stronger manipulation than pure framing on school teachers: in the penalty treatment teachers were paid their bonuses upfront, to be clawed back if student performance fell below target. They find strong positive effects on teacher performance under the penalty relative to the bonus equivalent.

Effort is non-contractible so P must write a contract that incentivizes effort. However, the agent's utility is reference-dependent and loss-averse, and the principal can influence the agent's reference point by altering how the contract is framed. Lastly, there is limited liability, such that the payment to the agent in any state must exceed some lower bound $\underline{w}$, which I normalize to zero for convenience.

A chooses an effort level $e \in [0, 1]$ which equals the probability that the task is successful. If successful, P earns a payoff v, otherwise he earns 0. In the absence of any framing effect, this implies that the optimal contract (ruling out contracts that pay lotteries) consists of a pair, (w, b), such that w is a non-contingent payment, and b is a bonus paid if the agent is successful. In addition, I assume that P can choose a frame, F, that influences A's reference point under the contract. Thus, P offers A a triple, (w, b, F), A accepts or rejects the contract, then exerts effort if she accepted and is paid according to the outcome. P's payoff is simply:

$$\Pi = e(v - b) - w.$$

I follow Kőszegi and Rabin (2006, 2007) (henceforth, KR) in assuming that the agent's utility function is a sum of a standard component, expected consumption utility, and a gain-loss component that values payoffs against a stochastic reference point, less an effort cost that is not reference-dependent for simplicity. In KR, the reference point is the expected distribution of outcomes. For example, an agent holding a lottery ticket that pays \$50 with 50% probability has a stochastic reference point that equals \$50 with probability one half and \$0 with probability one half.

The most natural specification of the reference point in work on moral hazard under loss aversion applies Kőszegi and Rabin (2007)'s concept of choice-acclimating personal equilibrium (CPE). Under CPE, the agent's choices influence her reference point, and she takes account of this effect when making her decision, i.e. her reference point and choices are simultaneously determined. For instance, under CPE, an agent knows that buying a lottery ticket will increase her expected winnings, which will be incorporated into her reference point. This creates the possibility of experiencing a loss when she does not win, which might ultimately lead her not to purchase a ticket. In the moral hazard context, higher effort leads to a higher chance of receiving b, which increases the reference point, and the agent anticipates this when choosing her effort level.

To focus the analysis, I impose a number of standard simplifying assumptions on the KR utility function, described in Appendix 1.A. The resulting basic framework is very similar to that used by Herweg et al. (2010) and Gill and Prowse (2012).[9] A's utility over consumption c, which is distributed according to $G(c|e)$, and reference

[9] In particular, I assume that the agent places equal weight on consumption and gain-loss utility, that utility is linear except for a kink at the reference point, and that the agent does not weight probabilities.

point r, distributed according to $H(r|e,F)$, where F represents the frame, is:

$$\begin{aligned} U(G|H,e,F) = E(c|e) + \alpha e - \frac{e^2}{2\gamma} \\ + [1 + \mathbb{1}\{c \geq r\}(\lambda - 1)] \iint (c-r)dG(c|e)dH(r|e,F) \end{aligned} \tag{1.1}$$

The first line, $E(c|e) + \alpha e - \frac{e^2}{2\gamma}$ is standard expected consumption utility, plus a term representing intrinsic motivation, αe, less the cost of effort which I assume is quadratic and decreasing in an ability parameter $\gamma > 0$. I assume $\alpha \in [0, \bar{\alpha}]$. The second line is A's gain-loss utility, depending on A's coefficient of loss aversion, λ. If $\lambda = 1$ (no loss aversion), then gain-loss utility is simply the expected difference between consumption and the reference point. If $\lambda > 1$, then losses are weighted more heavily than gains, i.e. states in which $c < r$ are overweighted relative to states in which $c > r$.

In the KR model, H is the true distribution of outcomes, so framing has no effect. To allow for framing effects, I make a simple modification to the specification of the reference point. I assume that the reference point takes the form of a weighted sum of the KR reference point and the "base pay" supplied by P, which depends on the choice of frame. Formally, the marginal distribution of r becomes:

$$h(x|e,F) = Pr(r = x|e,F) = \begin{cases} 1-e & x = \phi w + (1-\phi)(w + Fb) \\ e & x = \phi(w+b) + (1-\phi)(w+Fb) \end{cases} \tag{1.2}$$
$$F \in [0,1], \phi \in [0,1).$$

$w + Fb$ is the base pay supplied as a reference point by P, and F corresponds to the fraction of the bonus b that is presented as part of the base pay. $F = 0$ is a pure bonus frame with base pay w; $F = 1$ is a pure penalty frame with base pay $w + b$.[10]

The parameter ϕ captures how susceptible the agent is to framing effects. $\phi = 1$ coincides with KR's model, in which framing has no effect on behavior. When $\phi = 0$ the reference point is simply a choice variable for P. For intermediate values of ϕ increases, A's reference point lies in between w and $w + b$ and she experiences mixed emotions whether she receives or does not receive the bonus.

Incorporating the above assumptions, A's gain-loss utility sums over four states of the world. With probability e^2, her consumption is $w + b$ and her reference point is $\phi(w + b) + (1-\phi)(w + Fb) \leq w + b$, putting her in the gain domain. With probability $e(1-e)$ her consumption is $w + b$ and her reference point is $\phi w + (1-\phi)(w + Fb) \leq w + b$, again putting her in the gain domain. With probability $e(1-e)$ her

[10]The assumption that $F \in [0,1]$ rules out unrealistic expectations, that she might earn less than w or more than $w + b$. $F \in [0,1]$ also allows for intermediate frames. These would be contracts that incorporate both a bonus and a penalty region, varying the point on the incentive scheme that is reported to A as "base pay". I allow for intermediate frames for analytical convenience; in the simple two-outcome setup they are not particularly intuitively appealing since the contract would be of the form "base pay of $w + Fb$, with a bonus of $(1-F)b$ for success, and penalty of Fb for failure." However, in settings with a richer distribution of potential outcomes intermediate frames become more relevant.

consumption is w and her reference point is $\phi(w+b)+(1-\phi)(w+Fb) \geq w$, putting her in the loss domain, and with probability $(1-e)^2$ her consumption is w and her reference point is $\phi w+(1-\phi)(w+Fb) \geq w$, again putting her in the loss domain. Plugging these into (1.1) and simplifying, I can write her utility as:

$$U(e,w,b,F) = e[\alpha + b((2-\lambda\phi)+(\lambda-1)(1-\phi)F)] - e^2\left[\frac{1}{2\gamma} - (\lambda-1)\phi b\right] + w - \lambda(1-\phi)Fb. \tag{1.3}$$

1.1.1 Effort choice

Suppose A accepts the contract. I impose two parameter restrictions that are sufficient for the optimal equilibrium effort choice to be interior (i.e. lie in the interval $(0,1)$) and given by the first-order condition.

Assumption 1 *No dominance of gain-loss utility:* $\lambda < \frac{2}{\phi} \equiv \hat{\lambda}$.

Assumption 2 *No top-coding in effort:* $\gamma < \bar{\gamma} \equiv \frac{1}{\bar{\alpha}+v[1+\lambda-\phi]}$.

Under Assumptions 1 and 2, A's optimal effort choice is equal to:[11]

$$e^*(b,F) = \frac{\gamma[\alpha + b((2-\lambda\phi)+(\lambda-1)(1-\phi)F)]}{1-2\gamma(\lambda-1)\phi b} \tag{1.4}$$

Assumption 1 ensures that A is not so strongly loss averse that she is discouraged from exerting any effort at all. This can happen because her reference point is increasing in her effort choice, as higher effort increases her expected earnings. Intuitively, if λ is large, the increase in her disappointment when unsuccessful outweighs the higher expected earnings, such that she prefers to not to exert any effort. Assumption 2 simply ensures $e^* < 1$.

Inspection of (1.4) immediately reveals the incentive effect of framing. The penalty frame increases F and thus A's reference point. This in turn increases the magnitude of the loss that she experiences when unsuccessful. In response, she increases her effort provision to reduce the probability of experiencing the loss. The first testable prediction of the model is therefore:

Prediction 1 *Suppose $\lambda > 1$ and $\phi < 1$, so A is loss averse and susceptible to framing effects. Then, her effort provision is higher under a penalty framed contract than an economically equivalent bonus framed contract.*

[11] To see that Assumptions 1 and 2 imply that e^* is the unique, interior maximizer of (1.3), first note that the second order condition can be written as $2\gamma(\lambda-1)\phi b < 1$. Assuming this holds, we have $e^* > 0$ for all F, α and $b > 0$, provided $2-\lambda\phi > 0$, which is Assumption 1. Second, $e^* < 1$ for all F, α provided $\gamma b[(2-\lambda\phi)+(\lambda-1)(1-\phi)+2(\lambda-1)\phi] < 1$, which can be easily seen to imply the second order condition when Assumption 1 holds, and therefore sufficiency of the first-order condition. Last, note that non-negative profits imply $b \leq v$, so substituting v for b and simplifying, the expression reduces to Assumption 2, which is therefore sufficient for $e^* < 1$ in equilibrium. Assumption 1 is the analog of the assumption of the same name in Herweg et al. (2010), whose model corresponds to the case $\phi = 1$.

The second prediction relates e^* to A's coefficient of loss aversion. Differentiating e^* with respect to λ yields:

$$\frac{de^*}{d\lambda} = \frac{\gamma b[(1-\phi)F - \phi(1-2\gamma(\alpha+b(2-\phi)))]}{(1-2\gamma(\lambda-1)\phi b)^2} \tag{1.5}$$

The sign of this expression is ambiguous. In particular, suppose $F = 0$. Then $\frac{de^*}{d\lambda}$ is negative if $\phi > 0$ and $1 > 2\gamma(\alpha + b(2-\phi))$. However it is clear that provided A is susceptible to framing effects ($\phi < 1$), $\frac{d^2e^*}{d\lambda dF} > 0$.

Two opposing forces drive the ambiguity of the relationship between effort and loss aversion. On the one hand, more loss averse agents will work harder to avoid any given loss. On the other hand higher effort increases the probability of success, and therefore the reference point. This in turn increases the magnitude of the loss experienced when unsuccessful, discouraging higher effort provision. The net effect is ambiguous. However the intuition behind the effect of F on this relationship is clear: since more loss averse agents work harder to avoid any given loss, increasing the reference point by increasing F elicits a greater effort response the higher is λ.

Prediction 2 *The relationship between effort and loss aversion (λ) can be positive or negative, but is more positive (less negative) under the penalty than the bonus frame.*

1.1.2 Participation

Now consider A's participation decision. Substituting for (1.4) yields A's utility when she accepts the contract (w, b, F):

$$U^*(w,b,F) = \frac{\gamma[\alpha + b((2-\lambda\phi) + (\lambda-1)(1-\phi)F)]^2}{2\,[1-2\gamma(\lambda-1)\phi b]} + w - \lambda(1-\phi)Fb. \tag{1.6}$$

She accepts the job offer if her participation constraint is satisfied:

$$U^*(w,b,F) - \bar{u} \geq 0 \tag{1.7}$$

$\bar{u}$ is a fixed amount given up when taking on the task, reflecting A's outside option which may depend upon her type $(\alpha, \phi, \lambda, \gamma)$. Note that in addition, γ may directly depend on her outside option if the cost of effort partly consists of foregone effort on other tasks.

Applying the envelope theorem to (1.3) it is clear that A's utility is decreasing in F if $\phi < 1$. This is because now, for any level of effort, her reference point is higher and therefore her elation when she receives b is reduced, and her disappointment when she does not receive it is increased. This delivers a third testable prediction of the model:

Prediction 3 *Suppose A is susceptible to framing effects ($\phi < 1$). Then she is less willing to accept a penalty contract than the equivalent bonus contract.*

The implication of Prediction 3 for optimal contracts is given in the following Proposition:

Proposition 1 *Suppose P wishes to recruit one agent of known type. P prefers bonus framing to penalty framing whenever A's participation constraint is binding.*

The proof of Proposition 1 simply uses the fact that when the participation constraint is binding P is the residual claimant of any surplus generated by the relationship. Since by Prediction 3 A's utility is decreasing in F, P will prefer bonus frames. A marginal decrease in F can be offset by a decrease in w. The complement of this proposition is that P will only use some form of penalty framing ($F > 0$) when the participation constraint is not binding, i.e. when the limited liability constraint ($w \geq 0$) is binding.

1.2 Experimental design

In order to identify the effect of incentive framing on selection and effort, the experiment needs to have three key features. First, we need to be able to observe the types of workers that select into each contract and compare these with the population of potential workers. Second, we need to observe their effort provision under the contract, controlling for selection. Third, we need as far as possible to shut down any channel through which the contract offered is interpreted as informative about the nature of the task or the type of the principal.

To separate selection and incentive effects I use a two-stage design similar to Dohmen and Falk (2011). In the first stage, workers are surveyed and perform a practice task under flat incentives, giving me an opportunity to measure their types, and giving them experience at the task. Flat incentives were used to avoid workers being exposed to more than one form of incentive pay during the experiment. In the second stage they are then offered an opportunity to perform the same task under randomly varied performance-related incentives, which they can accept or reject. Their outside option is simply determined by their value or leisure or alternative tasks available. Selection effects can be examined by comparing the types that accept different offers, while incentive effects are estimated by comparing behavior conditional on acceptance and type.[12]

[12] A prominent alternative approach would be to use Karlan and Zinman (2009)'s methodology. Offer workers a contract, but upon their acceptance, randomly change the offer for some in a way that does not make them worse off. By comparing workers who accepted the same contract but ultimately received different ones, the researcher can test for incentive effects, holding selection constant. By comparing workers who accepted different contracts but ultimately received the same one, selection effects can be identified holding incentives constant. However, this approach is not suited to the current context. First and foremost, a switch from a bonus to penalty frame (or vice versa) is likely to make transparent the equivalence of the two. Secondly, it is not obvious in which direction one should make the switch. Moving from a bonus to a penalty frame involves a simultaneous increase in "base pay" and a swap from a reward for success to a penalty for failure. Depending on how workers interpret these changes, even a switch that makes them better off in cash terms might not be perceived as better.

Turning to confounds, a concern in experiments on incentive pay is that agents may interpret the contract offered as informative about the task, or about characteristics of the principal that are relevant to their payoff. For example, Bénabou and Tirole (2003) analyze an asymmetric information context whereby if the principal offers a larger bonus for a task (in equilibrium), the agent will believe the task to be more difficult. Similarly, in my context it seems intuitive to think that agents might perceive a penalty as designed to punish failure at an easy task (or where shirking is easy to detect).[13] In some contexts, participants might believe that the choice of bonus or penalty reflects the principal's transaction costs, whereby a bonus (penalty) is chosen when bonuses will be paid (penalties deducted) infrequently, to minimize the number of transactions conducted. Trust might also play an important role. It might plausibly be argued that agents could perceive a principal offering a penalty-framed contract as more or less trustworthy than one offering a bonus-framed contract.

While it is probably not possible to eliminate these concerns completely, I believe that they are addressed at least in part by two features of the experiment design. Firstly, all participants have already interacted with me, the principal, through the first stage of the experiment. They have agreed to an informed consent form that states this is a research project from an internationally well-known university, gives my name and contact details. They were paid promptly after completing the first stage, and received a personalized email from me after completion. The invitations to the second stage were received 6 days later, again by personalized email. This should address the trust concerns.

To address concerns about inference, after the first stage of the experiment, workers were sent an email informing them of the percentage of strings that they typed correctly in the first stage, to ensure that they have a good sense of the task difficulty. In the second stage they are paid the bonus (avoid the penalty) if they entered a randomly selected item correctly, so a worker with a 60% accuracy rate knows that she has a 60% chance of receiving the bonus (avoiding the penalty) in stage 2 if she performs at the same level. As for transaction costs, workers knew that they would receive their full payment for the incentivized task in a single transaction. This should address concerns that bonuses and penalties are perceived as being enacted infrequently.[14]

To check whether workers beliefs about the task difficulty is affected by the frame, I ask them at the beginning of stage 2 to estimate the average accuracy rate

[13] Note that this argument is difficult to formalize in the context of choice of contract frame. The reason is that in Bénabou and Tirole (2003), the choice of incentive plays the role of a costly signal from principal to agent, whereas a frame in my context is pure cheap talk. Therefore although it may be possible to construct equilibria in which bonuses are taken to signal harder tasks, it is equally possible to construct the reverse equilibrium, so the model is not predictive.

[14] An alternative way to address these concerns would be to make it explicit that incentives are randomized and what the set of possible incentives is. However this reveals the existence of an equivalent bonus (or penalty) contract to the one they received, which is potentially problematic. Moreover, I wanted workers to consider their job offer with respect to their outside option, and not with respect to the offers that other workers received.

from stage 1. If workers who receive one contract perceive the task to be more difficult than those who receive a different one, they should estimate a lower mean performance from the first task.

1.2.1 Experimental Platform: Amazon Mechanical Turk (MTurk)

The experiment was run on the online platform Amazon Mechanical Turk (MTurk, for short). MTurk is an online labor market for "micro outsourcing". For example, a "requester" that needs data entered, audio recordings transcribed, images categorized, proofreading, or many other possible types of tasks can post a Human Intelligence Task (HIT) on MTurk, and recruit "workers" to carry it out. The HIT terminology reflects that the purpose is typically to outsource to human workers tasks that are not currently suited to computerization. Workers are typically paid a flat fee for a task, with the potential to earn bonuses for good work, or to have their work rejected unpaid if it is poor quality.

MTurk has many attractive features for research. For example, a short survey can be prepared, posted and completed by hundreds of workers in a matter of hours, typically for much smaller incentives than might be used in a laboratory experiment. Bordalo et al. (2012) test their theory of salience using MTurk surveys. Barankay (2011) uses MTurk to study the effect on willingness to undertake more work of telling workers about their rank in an initial task. Horton et al. (2011) and Amir et al. (2012) replicate some classic experimental results with MTurk subjects. A good introduction to how to go about using MTurk for research is given by Paolacci et al. (2010) and Mason and Suri (2012). The demographics of MTurk workers are analyzed in Ipeirotis (2010).

1.2.2 Effort task

In each stage of the experiment, subjects were asked to transcribe 50 text strings, gradually increasing in length from 10 characters to 55 characters. The strings were generated using random combinations of upper and lower case letters, numbers and punctuation and distorted to give the appearance of having been scanned or photocopied.

The task was chosen to be implementable online, to be reasonably similar to the types of tasks that participants are used to doing in the course of their work on MTurk, and to be sufficiently difficult to generate variation in performance (accuracy) without putting the workers under time pressure.[15] Time pressure was not used to maintain similarity with other MTurk tasks which typically allow workers to work in their own time. In each stage there were 10 possible sets of strings

[15]The task closely resembles the kind of garbled text that individuals must type to solve a CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) puzzle on the web. Such puzzles are used in web forms as an attempt to prevent bots and spammers from accessing sites; in fact this has led to some spammers recruiting MTurk workers to solve the CAPTCHAs that are blocking their access. See e.g. New York Times blog, March 13, 2008: `http://bits.blogs.nytimes.com/2008/03/13/breaking-google-captchas-for-3-a-day/`.

and participants were randomly assigned to one of these.[16] An example screen is reproduced in Figure 1.1.

1.2.3 Design specifics

A flowchart summarizing the design and timings is given in Figure 1.2. Two experimental sessions were conducted, each of which consisted of two stages. The first stage of the experiment recruited workers on MTurk for a "Typing task and survey" for a flat pay of $3. Recruitment was restricted to workers located in the US. Participants performed the typing task then filled out the survey, which is described below. Once all participants had been paid, they were sent an email informing them of their performance in the typing task, example text is given in Appendix 1.C.2.

Six days later, all participants from were sent a second email, inviting them to perform a new typing task, this time under experimentally varied incentives. Each contract has three components: a fixed pay component that does not depend on performance, a variable pay component that does depend on performance, and a frame that is either "bonus" or "penalty". Participants were told that the task would remain open for four days, and that they could only attempt the task once.

Participants were randomized into one of three possible financial incentives, and either bonus or penalty frame. The treatments are detailed in Table 1.1, and consist of either low fixed and variable pay, low fixed and high variable pay or high fixed and low variable pay. The reasoning behind the particular rates of pay used is discussed in Appendix 1.C.1.

Participants were told that after completion of the task I would select, using a random number generator, one of the 50 strings that they had been assigned to type, and that they would receive the bonus (avoid the penalty) conditional on that item being entered correctly. I avoided using emotive terms like "bonus" and "penalty". For example, a penalty framed offer in experimental session 1 was worded as follows: "The basic pay for the task is $3.50. We will then randomly select one of the 50 items for checking. If you entered it incorrectly, the pay will be reduced by $1.50." Examples of the full email text are given in appendix 1.C.3. Experimental sessions 1 and 2 differed in the exact phrasing of the email, in order to check whether the results from Session 1 were driven by inattention.

1.2.4 Data

This section describes the key variables collected in the survey and effort tasks. Summary statistics are given in Tables 1.2 and 1.3 and summary distributions of key variables plotted in Figure 1.3. I describe the main variables below.

The measure of loss aversion I use is similar to that of Abeler et al. (2011), but unincentivized. Participants are asked to consider a sequence of 12 lotteries of the form "50% chance of winning $10, 50% chance of losing $X," where X varies from

[16]This was done because other experimenters report occasionally participants posting answers to tasks on the web. I found no evidence of this occurring in this experiment.

\$0 to \$11. For each lottery, they are asked whether or not they would be willing to play this lottery if offered to them by someone "trustworthy". The main measure of loss aversion is the number of rejected lotteries.[17] I also record whether participants made inconsistent choices, accepting a lottery that is dominated by one they rejected. A screenshot of the lottery questions is given in Appendix 1.C.5.[18]

Two other key variables that I attempt to measure are participants' reservation wages and their perceptions of what constitutes a "fair" wage (see e.g. Fehr et al. (2009)). A measure of reservation wages is useful in considering how the framed incentives affect willingness to accept a job offer. All else equal (in particular, controlling for ability), if one contract is perceived as less attractive it should particularly discourage those with a higher reservation wage. To elicit reservation wages I ask participants what is the minimum hourly wage at which they are willing to work on MTurk. This is a meaningful concept in the MTurk context because at any point in time there are typically many HITs that a worker could perform, with varying rates of pay. Moreover, 93% of participants stated that their main reason for working on MTurk is to earn money, suggesting that reservation wages should be important.

Fehr and Gächter (2002) find in a buyer-seller experiment that penalty-framed performance incentives led to more shirking among sellers than equivalent bonus-framed offers, and argue that this is because the penalty contracts are perceived as less fair. I ask participants what they think is the minimum fair wage that requesters "should" pay on MTurk, and use this measure to proxy for fairness concerns. Reservation wages are typically lower than fair wages.

Participants were also asked to report the zipcode of their current location, which I map in Appendix 1.B.8. The distribution of participant locations closely resembles the population distribution across the US.

Now I turn to measuring effort and performance on the effort task. The main performance measure is "Accuracy Task X", the fraction of text strings that participants entered correctly in stage X. I also construct a second accuracy measure, "Scaled Distance Task X", which can be thought of as the error rate per character typed.[19] In the regressions I use the natural log of this measure since it is is heav-

[17]Note that by Rabin (2000) aversion to risk in small stakes lotteries is better explained by loss aversion than standard concave utility.

[18]The lottery choices were not incentivized because of concerns that this would interfere with studying selection effects and willingness to accept job offers for the effort task alone. Offering financial incentives large enough for participants to potentially lose \$10 is problematic because it would interfere with the selection effects I am trying to measure: if the incentives were advertised upfront they might attract high reservation wage participants who would not participate in the stage 2 effort task; if they were not pre-announced, subsequently revealing them might lead the participants to expect unannounced rewards in stage 2 and thus be more likely to accept in stage 2. Camerer and Hogarth (1999) argue that: "In the kinds of tasks economists are most interested in, like trading in markets, bargaining in games and choosing among gambles, the overwhelming finding is that increased incentives do not change average behavior substantively."

[19]For each text string I compute the Levenshtein distance between the participant's response and the correct answer, and divide by the length of the correct answer. The Levenshtein distance between two strings, A and B, is the minimum number of single character insertions, deletions, or swaps needed to convert string A into string B. This then translates into the probability of error per character for that string. I then take the average over all text strings for that participant to find their per-character average error rate

ily skewed by a small number of participants who performed poorly (per-string accuracy rates are sensitive to small differences in per-character error rates).

I also attempt to measure how much time participants spent on their responses. This measure is imperfect as I cannot observe how long participants were actually working on their responses, only how long the web pages were open for, and participants may leave the page open while performing other tasks. As previously discussed this is a common feature of how workers work on MTurk HITs and I did not want to introduce artificial time pressure. To deal with this issue, I take the median time spent on a page of the typing task (there are 10 pages of 5 items each), multiplied by 10. This measure is still heavily skewed, however.

Finally, participants were asked at the beginning of stage 2 to estimate the mean accuracy rate (across all participants) from stage 1, which I use to test whether the frame changed their perceptions of the task, a variable I label "Predicted Accuracy".

In total 1,465 participants were recruited, of which 693 returned for stage 2. 15 participants are dropped from all of the analysis, six because I have strong reasons to suspect that the same person is using two MTurk accounts and participated twice and nine because they scored zero percent accuracy in the stage 1 typing task, suggesting that they did not take the task seriously (of the six of these who returned for stage 2, five scored zero percent again).

1.2.5 Randomization

To test for selection effects it is important to have good balance on the variables for which selection might be a concern. In this context, the obvious candidates are the participants' performance in stage 1, measured loss aversion and reservation wage, so I stratify on these variables. I was slightly concerned that some participants might know one another (for example, a couple who both work on MTurk), so the treatments were randomized and standard errors clustered at the zipcode-session level.[20] In regressions that drop participants who share a zipcode this is equivalent to using robust standard errors, since then each cluster is of size one.

As a graphical check of balance, Appendix 1.B.2 plots the CDFs for task 1 accuracy, reservation wage, fair wage and rejected lotteries, separately for the bonus and penalty treatments, confirming near-perfect balance on these variables. Also, Appendix 1.B.3 plots the distributions of these variables by experimental session.

Table 1.4 gives the results of the statistical balance tests. I perform two exercises. The first tests the joint significance of the full set of treatment dummies in explaining each baseline characteristic. The second performs a t-test for comparison of means between pairs of treatments, where each pair considered differs only in terms of its bonus/penalty frame (groups (0,1), (2,3), (4,5) and (6,7) as labeled in Table 1.1). Both exercises confirm good mean balance on all characteristics with the exception of the minimum fair wage (F-statistic p-value 0.01), which is driven by differences between

[20] In the stage 1 data 187 individuals report being located in the same zipcode as another participant from the same session.

sessions 1 and 2, and the number of MTurk HITs completed (p-value 0.07), driven by a small number of participants with very large numbers of HITs completed. In addition, I run Kolmogorov-Smirnov or Mann-Whitney equality of distributions tests between bonus and penalty frames for stage 1 performance time spent on stage 1, loss aversion, reservation wage and fair wage, none of which reject the null of equal distributions.

1.3 Main Results

In this section I describe the main results of the experiment. The key results of interest are the effect of the penalty frame on participants' willingness to accept the contract, the effect of the frame on the types of participants who select into the contract, and the effect on performance on the job. I discuss the relation between the key observable characteristics, and between characteristics and performance, in Appendix 1.B.1.

1.3.1 Acceptance

First I consider participants' decision whether or not to accept the stage 2 job offer. Figure 1.4 graphs acceptance rates by treatment. The striking pattern in these data is that penalty framed contracts were much more likely to be accepted than equivalent bonus framed contracts. The relationship is pronounced for the four groups with fixed pay of \$0.50, and weaker for the two groups with fixed pay equal to \$2. In addition, acceptance is substantially higher under higher fixed pay, while the relation between variable pay and acceptance appears weak at best.

This result is particularly striking because it directly contradicts model Prediction 3. The simple loss aversion model implies that penalty contracts should be relatively unattractive compared with equivalent bonus contracts. I discuss this finding in relation to the theory in section 1.5.

Moving on to the regression results, the basic regression specification is a linear probability model with dependent variable $Accept_i \in \{0,1\}$, individuals indexed by i:

$$Accept_i = \beta_0 + \beta_1 * Penalty_i + \beta_2 * HighFixedPay_i + \beta_3 * HighVariablePay_i + \beta_4 * X_i + \epsilon_i \tag{1.8}$$

Penalty is a dummy equal to 1 if the contract is penalty framed and zero if bonus framed. *HighFixedPay* is a dummy indicating fixed pay equal to \$2 (alternative: \$0.50). *HighVariablePay* is a dummy indicating variable pay of \$3 (alternative: \$1.50). Since there are only two levels of fixed and variable pay, it is straightforward to compute the implied linear effects (per dollar), by dividing the coefficient on *HighFixedPay* by 1.5 (\$2 - \$0.50) and the coefficient on *HighVariablePay* by 1.5 (\$3 - \$1.50).

X_i is a vector of variables measured in stage 1. In particular, I include accuracy and time spent on the stage 1 effort task, to jointly proxy for ability and intrinsic motivation. All results are robust to additionally including the ratio of accuracy to time spent (not reported). X_i also includes dummies for the set of items assigned to be typed by that participant (10 possible sets). Note that the main specifications pool the effects of each component of the contract to increase power. I break down the treatment effects below.

Table 1.5 presents the main results. I find that a penalty framed contract increases acceptance rates by approximately 10 percentage points over the equivalent bonus frame. This implies a 25 percent higher acceptance rate under the penalty frame than (the mean acceptance rate under the bonus frame was 42 percent), a large effect for a simple framing manipulation. High fixed pay increases acceptance by around 15-16 percentage points, or around 10 percentage points per dollar. Surprisingly, the effect of high variable pay is positive but much smaller at around 3 percentage points greater take-up (around 2 percentage points per dollar), and not statistically significant. The results are robust to dropping participants who made inconsistent choices in stage 1, who spent a very long time on the first task, have very high reservation or fair wages, or are from zipcodes with more than one respondent. Near-identical average marginal effects are obtained using logistic instead of linear regression.

Participants who performed better on the unincentivized stage 1 were significantly more likely to accept the stage 2 job offer, as is clear from Table 1.2 and Figure 1.10. This is consistent with the common finding that performance pay differentially selects more able or motivated workers, and which I discuss further in Section 1.4.4. Participants with a higher reservation wage were significantly less likely to accept the offer, suggesting that this measure is informative. Interestingly, when controlling for this measure, the coefficient on "minimum fair wage" is not statistically significant, suggesting that fairness concerns (as measured by this variable) were not of primary importance for willingness to accept the contract.

The number of hypothetical lotteries rejected by participants is not predictive of acceptance, whether or not I drop participants who made inconsistent choices in the lottery questions. This is surprising as the stage 2 contract is risky, so one would expect more risk/loss averse participants to be less willing to accept it. Figure 1.10 shows that the distributions of rejected lotteries are essentially identical for participants who did and did not accept the stage 2 job offer. This could be because the measure is poorly capturing loss aversion, although similar unincentivized measures have been successful in other studies (see e.g. Camerer and Hogarth (1999)).

Next I unpack the effects of the contract terms on acceptance rates. Visual inspection of Figure 1.4 suggests that the effect of the penalty frame is larger when the variable pay component is larger and smaller when the size of the fixed pay component is larger. Table 1.6 reports the relevant interaction effect estimates. The point estimates do indeed suggest that the effect of the penalty frame is smaller for high fixed pay and larger for high variable pay, however neither estimate is statistically significant when estimated separately or simultaneously. In addition, the point estimate on "high variable pay" is essentially zero for participants under the bonus frame, implying that the potential for a $3 bonus as opposed to a $1.50 bonus did not make the job offer significantly more attractive.

1.3.2 Selection

Now I turn to the effect of the penalty frame on the types of workers that select into the contract. Figure 1.5 plots CDFs of stage 1 task performance, time spent on stage 1 task, rejected lotteries, reservation wage and fair wage, comparing those who accepted the bonus frame with those who accepted the penalty frame. Surprisingly, the distributions are essentially overlapping for all variables except for reservation wages, implying no differential selection on these variables. The lack of selection on fair wages suggests that fairness concerns do not drive the difference in acceptance rates between contracts.

I do observe suggestive evidence that the penalty contract attracted workers with higher reservation wages. This is consistent with either of the mechanisms proposed in section 1.5. Note that selecting in workers with higher reservation wages is not inconsistent with no selection on other characteristics, since as shown in Appendix 1.B.1, the correlation between reservation wages and other characteristics is small.

Table 1.7 tests for selection effects of penalty framing by interacting the penalty frame with the key observables in acceptance regressions. The interaction coefficients estimate the extent to which a given characteristic more or less strongly predicts acceptance under the penalty frame. In each case the interaction terms are not statistically significant, whether estimated separately or jointly. A joint test fails to reject the null that all interaction coefficients are equal to zero (p-value 0.90).

In Appendix 1.B.4 I check if the results are robust to dropping outliers for time spent on task 1, reservation wage and fair wage, or dropping participants who made inconsistent lottery choices. The only difference is that now the interaction between penalty frame and reservation wage is significant at the 10 percent level, consistent with the penalty screening in participants with higher reservation wages. Overall there is little evidence of selection on these key observables between contract types.

As for the other covariates, men and women are equally likely to accept the penalty contract, but men are eleven percentage points less likely to accept the bonus contract. Participants who reported that their main reason for working on MTurk is to earn money (93 percent of participants) are significantly relatively more likely to accept the penalty contract than those who gave another reason, similar applies to those who report mostly working on research HITs. Participants with more HITs completed are relatively less likely to accept the bonus contract. Regression tables showing these results are available on request. Despite the fact that being male and citing money as the main reason for working are positively associated with performance on stage 1, none of these results seems to be consequential for performance, as illustrated by the lack of selection on stage 1 performance measures and the evidence presented in the next section.

It is quite surprising that there seems to be no selection effect of penalty framing. One possibility is that selection is hard to detect in this context. If workers could choose between the bonus and penalty framed job, assuming this did not undo the framing effect, they should select into the job they preferred. However in this experiment, workers cannot choose jobs, only whether to accept the one that they are offered. As a result there will be many participants in each pool whose participation constraints would be satisfied under either contract: any selection effect would have to be driven by the fraction of participants whose participation constraints are satisfied under one but violated under the other. However, given the large difference in acceptance rates between contracts it seems unlikely that this is what is driving the lack of detectable selection. In addition, as documented in Section 1.4.4, I am able to observe the more standard result of differential selection by type into the incentivized task: workers with high performance in the first stage are more likely to accept the job in the second stage.

1.3.3 Performance

Now I turn to the incentive effects of contract framing on worker effort and performance. This section directly relates to the existing literature on framed contracts which considers incentive effects for an already recruited sample of workers or participants.

The basic regression equation is:

$$Y_i = \delta_0 + \delta_1 * Penalty_i + \delta_2 * HighFixedPay_i + \delta_3 * HighVariablePay_i + \delta_4 * X_i + \epsilon_i \tag{1.9}$$

Where Y_i is a measure of effort or performance. The key measures are summarized in Table 1.2 and distributions plotted in Figure 1.3. As before, X_i is a vector of variables measured in stage 1, including accuracy and time spent to jointly proxy for ability and intrinsic motivation (results are robust to also including the ratio of these variables).

In general one would expect the estimates of δ_1, δ_2 and δ_3 to be biased by selection: if the workers that accept one type of contract are different from those that accept another, then performance differences may simply reflect different types rather than different effort responses to incentives. However as already documented, I do not observe differential selection between frames, which would bias the estimate of the key coefficient of interest, δ_1. Moreover, since I have stage 1 measures of performance and characteristics, I can control for selection on observables by including these. The estimates are very stable with respect to inclusion of controls, giving confidence that selection effects are not driving the results.

Recall that the incentivized performance measure is accuracy in task 2, which can be interpreted directly as the probability of receiving the bonus (avoiding the penalty) since participants were paid if one randomly selected item was entered correctly. Figure 1.6 presents the mean performance on the stage 2 task by treatment group. As with acceptance rates, a clear pattern is visible: performance is always higher in the penalty framed than in the bonus framed treatments. Figure 1.7 plots CDFs of accuracy, the log distance measure (recall that this is interpreted as the log of the per-character error rate) and time spent on stage 2, each of which demonstrate a small increase in performance and effort between frames consistent with performance or effort first-order stochastic dominating that under the bonus frame. These results are consistent with the existing experimental studies of framed incentives in both direction and magnitude.

The main regression results are given in Table 1.8. I find that accuracy under the penalty frame is around 3.6 percentage points (around 0.18 standard deviations or 6 percent of the mean accuracy of 0.59) higher than under the bonus frame, statistically significant at 5 percent without/1 percent with controls. The coefficient estimate is robust to dropping participants who made inconsistent lottery choices, participants from zipcodes with multiple respondents, and outliers on the reservation and fair wage questions (although a little smaller and only significant at 10 percent). Crucially, the estimated penalty effect is unaffected by the inclusion or exclusion of controls, consistent with the contract frame not inducing significant outcome-relevant selection, at least on observables. For selection to explain the results, there would have to be a substantial unobserved driver of performance that differentially selected under the penalty frame and orthogonal to the wide range of controls included in the regressions.

In addition, high fixed pay increases accuracy by around 2-4 percentage points, significant at 5 percent when controls are included. The point estimate doubles when controls are included, suggesting there may be adverse selection induced by the higher fixed pay. If anything, the fact that a selection effect *is* observed here gives comfort that the lack of observed selection between bonus and penalty reflects a true lack of selection in the data. Lastly, high variable pay increases accuracy by around 1.4-2.5 percentage points, although this is never significant at conventional levels.

As for the other key variables, performance in the first stage very strongly predicts performance in the second stage, while the coefficient on time spent in the previous task is negative, small in magnitude and not significant. As in stage 1, a higher reservation wage is associated with poorer performance. Controlling for the reservation wage, the reported minimum fair wage has no effect on performance, again consistent with fairness concerns not being of primary import.

Lastly, and interestingly, in this stage the number of rejected lotteries is negatively associated with performance, significant at ten percent when including participants who made inconsistent choices and five percent when they are dropped. A one standard deviation increase in the number of rejected lotteries is associated with around 1-2 percentage points worse performance. This result relates to Prediction 2 and is expanded upon when I consider heterogeneous effects below.

In table 1.9 I regress the distance measure of accuracy (log errors per character typed) and time spent on treatment dummies and covariates. The estimates imply that the penalty frame led to participants committing 20 percent fewer errors per character (from a mean of 0.066), and spending around two to three minutes longer on the task (mean 41 minutes), although the latter is not significant when controls are included. The point estimates on fixed and variable pay mirror their counterparts in the main regressions, and once again there is evidence of adverse selection induced by the higher fixed pay.

Next I break down the differential effects between bonus and penalty frames of varying fixed and variable pay, interacting the penalty dummy with the high fixed/high variable pay dummies. Regressions are reported in Table 1.10. Increasing the fixed pay seems to have the same effect under both bonus and penalty contracts (the interaction effect is an imprecisely estimated zero). Increasing the size of the variable pay is associated with higher effort under both frames, with a smaller effect under the penalty frame. However neither estimate is statistically significant.

Table 1.11 reports estimates of heterogeneous effects of the penalty treatment by the main variables. In each case the individually estimated interaction effect is not statistically significant: there is little evidence of strong heterogeneous effects. Including all of the interaction effects estimated together gives one unexpected and difficult to interpret result: a negative interaction effect for reservation wages and a positive one for minimum fair wages.[21]

Focusing on the coefficient on rejected lotteries, I note that although neither the main effect nor interaction coefficient are statistically significant. Nevertheless it is striking that the implied coefficient on rejected lotteries is close to zero under the bonus frame, and negative under the penalty frame (the combined effect under the penalty frame is statistically significant at the 5 percent level), while the model Prediction 2 implies that the coefficient should be more positive under the penalty frame.

Clearly I lack power to dig into this relation in depth. I do however perform one simple exercise. In Figure 1.9 I non-parametrically (LOWESS) plot accuracy against rejected lotteries separately under bonus and penalty frame, after partialling out the other variables, dropping participants with inconsistent choices and those who rejected or accepted all lotteries. The slopes are approximately equal over much of the range of values for rejected lotteries, but flattening out for high values under the bonus frame while becoming strongly negative under the penalty frame, which is what is driving the difference in the regression coefficients. For most participants the relationship between performance and loss aversion is similar between frames, but penalties seem to strongly *discourage* the most loss averse participants, an effect which is not in the model.[22]

[21] Dropping participants above the 99th percentile for reservation or fair wages, the p-value on the reservation wage interaction increases to 0.09 and that on the fair wage interaction increases to 0.14.

[22] In fact, the model does predict discouragement for highly loss-averse agents, which I ruled out by Assumption 1. However the discouragement effect is predicted to be stronger under the bonus than the penalty frame.

Summing up, I find strong and robust evidence that the penalty frame caused more participants to accept the stage 2 job offer, and performed better under penalty framed contracts. However I find no evidence of adverse or advantageous selection into the penalty contract, either directly from comparing distributions of characteristics, or indirectly by observing that the treatment effects are robust to inclusion of controls for participant type.

In addition, I find that higher fixed pay robustly leads to higher acceptance of the job offer, and to higher performance, with some evidence of adverse selection. I find a positive relationship between higher variable pay and acceptance and performance, but surprisingly this is never statistically significant.

The most surprising implication is that a simple switch from bonus to penalty framing is highly lucrative from the principal's perspective. My estimates suggest that recruitment and performance would be similar under a contract that pays \$0.50 fixed pay with \$1.50 variable pay framed as a penalty, as with \$1.50 fixed pay and \$1.50 variable pay, framed as a bonus.

1.4 Secondary Results

In this section I outline secondary results aimed at partly unpacking the mechanism behind the main results. First, I demonstrate that the penalty frame did not appear to change participants' perceptions of the task difficulty. Second, I argue that inattention is unlikely to explain the higher acceptance rate of the penalty contract. Third, I document that a standard selection result is present in my data: performance pay attracted more able workers, giving confidence that the no-selection result between frames is not driven by, for example, MTurk workers being atypical experimental subjects. Last, I briefly discuss persistence of the framing effect, exploring whether effort under the penalty remained high throughout the task, or whether participants quickly realize the equivalence of the two contracts.

1.4.1 Does framing affect perceived task difficulty?

Subjects were asked to predict the mean accuracy rating on the first stage at the beginning of the stage 2 task. If the frame influenced their beliefs about the task difficulty, it should also influence their belief about mean performance in the first stage. Figure 1.8 graphs the mean of these predictions by treatment group. There is no systematic relationship between the framing treatment and the predictions, but it seems that predictions were higher for the groups with high variable pay than with low variable pay.

In Appendix 1.B.5 I regress the participants' estimates on the treatment and the same controls as the main regressions. As expected, I find no evidence that the frame influenced participants beliefs about the task difficulty: the coefficient estimates are insignificant zeroes. Similarly I do not find any evidence that high fixed pay affected beliefs. I do find that participants assigned the high bonus predict that stage 1 accuracy was 3.5 percentage points higher. To check whether this affected stage 2 performance, I also regress stage 2 performance on the main regressors, now including the participant's prediction of stage 1 accuracy.

In the specification with only treatment indicators I find a strong and significant correlation between predicted and actual performance, perhaps because the predictions are positively (although not significantly) correlated with own performance on task 1, which is in turn a strong predictor of task 2 performance. Taken literally, the coefficient implies that a 10 percentage point higher belief about mean stage 1 performance is associated with 2 percentage points better performance in stage 2. When controls are included, this coefficient drops effectively to zero and is not statistically significant. In sum, it does not appear that the effect of the penalty contract is explained by different beliefs induced by the frame.

1.4.2 Inattention: Comparing Sessions 1 and 2

After running the first experimental session, one concern was that the way in which the job offer was phrased might differentially attract inattentive participants into the penalty frame, explaining the higher acceptance rate. The "pay" section of the offer email under the bonus (penalty) frame was written as follows:

> The basic pay for the task is $0.50 ($3.50). We will then randomly select one of the 50 items for checking. If you entered it correctly (incorrectly), the pay will be increased (reduced) by $3.00.

It is possible that an inattentive participant receiving the bonus email might glance at the first sentence, see a low amount and close the message, while an inattentive participant under the penalty frame might see the high amount and click through to the task.[23]

To alleviate this concern, the second session changed the email slightly. The pay was described as follows:

> The pay for this task depends on your typing accuracy. We will randomly select one item for checking, and if it was entered correctly (incorrectly), the pay will be increased above (reduced below) the base pay. The base pay is $0.50 ($3.50) which will be increased (reduced) by $3 if the checked item is correct (incorrect).

This phrasing pushes the pay information to the end of the paragraph and puts it all into one sentence. It also tells the participants immediately that the pay will depend on performance, and hence they should pay attention to the contingent component of pay. Of course, one challenge with such manipulations is that if no effect were observed, it is hard to tell whether the rephrasing eliminated inattention, or somehow changed the reference point.

[23]Note, however, that this is only part of the email content, that the link to the task came later in the, and that "base pay" should strongly signify that there is more payoff-relevant information to come.

As shown in Figure 1.4, the penalty framed offer was still strongly more attractive than the bonus framed offer in session 2, suggesting that inattention was not driving the main acceptance results. However, I also find that the performance difference between bonus and penalty was much smaller and not significant in session 2 (see Figure 1.6). I confirm in this by splitting the penalty treatment effect by session in regressions in Appendix 1.B.6. It is not possible to say whether this is a treatment effect (the rephrasing decreased the framing effect on performance) or sampling variation, particularly given the smaller sample size in session 2.

1.4.3 Does the framing effect wear off?

One interesting possibility discussed in Hossain and List (2012) is that the effect of an incentive frame might wear off over time as the agent realizes that the framing manipulation has no economic content.[24] Since my participants only perform the task once under framed incentives, I can only test for this within task, but nevertheless it is interesting to see whether performance under the bonus and penalty frame converges toward the end of the stage 2 task.

In Figure 1.11 I plot performance by typed item, or by page of five typed items in stage 1, and separately for each framing treatment in stage 2, including only participants who accepted the stage 2 offer. The lines slope down because the text items grow progressively longer between pages. The graph clearly illustrates the shift in performance in the penalty over the bonus frame is consistent throughout the task, there is no evidence of convergence. I confirm this in a performance regression in Appendix 1.B.7, where I find that the coefficient on item number interacted with the penalty dummy is a precisely estimated zero, while convergence would imply a negative coefficient.

Of course, this result does not imply that in the longer-run, for example after a couple of stages of incentive pay, participants' reference points would not shift to their true expected earnings, eliminating the framing effect. However over the short horizon of this experiment the penalty frame effect is persistent.

1.4.4 Standard Selection Effect

In this section I briefly show that I do find a standard selection effect: that higher ability workers are more likely to accept the stage 2 offers of incentive pay, as also found by Lazear (2000) and Dohmen and Falk (2011), for example.

Figure 1.3 shows the distribution of accuracy measures in stages 1 and 2, showing a large increase in performance in stage 2. This increase depends on three things: the effect of incentive pay on effort, the effect of incentive pay on selecting in motivated or able workers, and learning by doing. I cannot separate out learning by doing since I do not have a flat pay incentive treatment in stage 2, however I can illustrate the effect of selection.

[24] See also Jayaraman et al. (2014) who find a short lived "behavioral" response to a contract change.

Figure 1.10 plots CDFs of stage 1 variables, comparing acceptors with rejectors (pooling all treatments). Acceptors performed better in stage 1, spent less time, and have lower reservation and fair wages, each one suggesting a first-order stochastic dominance shift. As previously noted I see little difference in rejected lotteries, which is surprising since the incentive pay is inherently risky. The differences in ability are also demonstrated by comparing stage 1 performance measures between acceptors and rejectors in Table 1.2

Ideally, to demonstrate the selection effect on performance I would regress accuracy on a "round 2" dummy, with and without controlling for ability. This is not possible since my main ability measure is performance on the stage 1 task, which is also an outcome.[25] Instead, I perform the following simple exercise. I assume that the true performance model is linear and equal to that estimated in Table 1.8 column (4) (excluding the "set" dummy variables for the stage 2 task, since these are only observed for those who performed the task). Using this model I impute task 2 accuracy for rejectors.

The results are as follows. Mean accuracy across all participants in stage 1 is equal to 0.46, while in stage 2 it is equal to 0.59. Ignoring selection, a naive estimate of the combined effect of learning by doing and incentive pay would therefore be equal to a 13 percentage point improvement. However, the mean fitted stage 2 accuracy for all participants, including rejectors, is 0.56. Under the strong assumption that the fitted model is the true model, this implies that three percentage points of the combined effect can be attributed to advantageous selection of workers into incentive pay. An alternative way to control for selection is to compare mean performance of acceptors in stage 2 with mean performance of acceptors in stage 1, equal to 0.48. This gives me a similar effect of incentives and learning equal to 11 percentage points.

1.5 Revised model

It is clear that the model outlined in section 1.1 does not fit the data. Effort on the incentivized task is higher under the penalty frame, consistent with Prediction 1. Effort is decreasing in the number of rejected lotteries which is possible for $\phi > 0$. However, the data strongly contradict Prediction 3: the acceptance rate is substantially higher under the penalty frame. There is also suggestive evidence that Prediction 2 is violated: loss aversion does not appear to predict a stronger positive effect of framing on effort, if anything, the opposite.

In this section, I outline two extensions to the model that deliver theoretical predictions more closely in line with the experimental results. One allows for the framing treatment, which increases the agent's reference point, to also increase the perceived utility from accepting the job offer, or decreases the perceived value of the outside option. In this case, the increased willingness to accept the penalty

[25] Furthermore, I do not have a variable that will plausibly satisfy the exclusion restriction needed for a selection model.

contract can be thought of as a mistake; ultimately she will be more disappointed with outcomes under the penalty than under the bonus.

The second modification instead assumes that the agent has "Planner-Doer" preferences. As "Planner", she decides whether to accept the job offer. As "Doer", she chooses her effort level and performs the task. Planner and Doer's preferences conflict. Doer is loss-averse according to the original model, but planner only cares about material outcomes. Planner prefers the penalty contract, since the motivating effect of the penalty commits Doer to a higher effort level. The intuition is similar to that in Kaur et al. (2013), who argue that workers in their experiment select a particular form of dominated incentive scheme that helps them commit to exert higher effort. Of course the tension in my context is that effort is exerted shortly after acceptance, implying a rather short horizon self-control problem.

1.5.1 Salience

One possible explanation for the acceptance results is that the base pay both sets the worker's reference point and is somehow more salient than the incentive component of the contract, such that a high base pay makes the job offer appear more attractive. The simplest way to model this is to simply assume that the reference point set by the contract affects the perceived value of the agent's outside option, a higher reference point making the outside option appear relatively less attractive.

The expected reference point is $E(r) = \phi(w + eb) + (1 - \phi)(w + Fb)$, and the expected pay under the contract is $E(c) = w + eb$. Call the difference $S \equiv E(r) - E(c)$. The agent's modified participation constraint is $U^*(w, b, F) - \bar{u}(S) \geq 0$, or:

$$U^*(w, b, F) - \bar{u}(b(1 - \phi)(F - e^*)) \geq 0. \tag{1.10}$$

Applying the envelope theorem, a necessary condition for the agent to always prefer the penalty frame to the bonus frame is $\bar{u}'(S) + \lambda < 0$. Intuitively, the effect of the increased reference point on her perception of the value of her outside option must be sufficiently large to overcome the disutility of the penalty when she accepts the job, which depends on her loss aversion.[26]

1.5.2 Penalties as a commitment device

An alternative explanation for the popularity of the penalty framed contract is that participants see the motivational power of penalties as a commitment device that helps them overcome weak self-control to exert higher effort and increase their earnings. This requires a multiple-selves view of preferences where there is a preference conflict between the self who accepts the job, "Planner", and the self who performs

[26] It is also possible to derive a condition under which the penalty frame does not differentially attract more or less loss-averse workers, $\frac{d^2}{dFd\lambda}[U^*(w, b, F) - \bar{u}(b(1 - \phi)(F - e^*))] = 0$. This can be written as $\frac{de^*}{d\lambda}[(\lambda - 1) + b(1 - \phi)\bar{u}''(S)] - (1 - e^*) - \frac{\partial^2 \bar{u}(S)}{\partial S \partial \lambda} = 0$. However, it is difficult to conceive of an economic intuition for why this condition might or might not be satisfied.

the task, "Doer". Here I briefly outline how such a model might be structured.

Suppose that if she accepts the job, A's "Doer" preferences are described by (1.3), and thus her effort is equal to e^* as defined by (1.4). However, when considering whether to accept the job, her "Planner" self only considers material payoffs and the cost of effort. Planner accepts the job if the following participation constraint is satisfied:

$$w + e^*b - \beta\left(\frac{e^{*2}}{2\gamma} - \alpha e^*\right) - \bar{u} \geq 0 \tag{1.11}$$

where β is the weight applied by Planner to Doer's cost of effort. If $\phi = \lambda = \beta = 1$, Planner and Doer's preferences coincide. If $\beta < 1$ then in the absence of reference dependence and loss aversion ($\phi = \lambda = 1$), A suffers from weak self-control: Doer exerts less effort than Planner would like her to, and Planner will demand commitment devices that move Doer's effort closer to her first-best effort, which is equal to $e^{PFB} \equiv \gamma\left(\frac{b}{\beta} + \alpha\right)$.

Planner will prefer a marginal increase in base pay (through a marginal increase in F) whenever $e^{PFB} > e^*$. A sufficient condition for her to prefer the pure penalty contract to the pure bonus contract is $e^{PFB} > e^*(b, 1)$, which simplifies to $\beta < \frac{1-2\gamma(\lambda-1)\phi b}{2\gamma(\lambda-1)\phi\alpha+(1+\lambda+\phi-2\lambda\phi)}$. For example, when $\phi = 0$, this reduces to $\beta < \frac{1}{1+\lambda}$.

1.5.3 Discussion

The two modifications to the model outlined above enable it to produce the most striking empirical result, that the penalty contract is more popular than the bonus contract. However, four results remain outside the model.

First, participants exert higher effort when the fixed pay component is increased. The most obvious explanation here is a reciprocity or efficiency wage story. Participants may respond to perceived generosity from the principal (increasing the fixed pay) by reciprocating with higher effort. Alternatively, perhaps some participants believed that maintaining a good reputation with the principal would be rewarded in future, and higher fixed pay increased the perceived value of a good reputation. Participants were explicitly told in the first stage that the chance of being invited for future tasks would not depend on their performance, but this was not reiterated in the second stage. However, to explain the higher effort in all treatments under the penalty contract this reputation mechanism would have to be stronger for penalties than bonuses.

Second, it is surprising that increasing the size of the variable pay component seems to have at most a small and not statistically significant effect on acceptance and performance. This may be explained by diminishing sensitivity (Kahneman and Tversky (1979)), which is not in the model for simplicity. Under diminishing sensitivity the gain-loss component of utility is concave in the gain domain and convex in the loss domain. Then, a marginal change in the reference point may have a sharper effect on incentives than an equal marginal change in the actual financial incentive, holding the reference point fixed. Alternatively, it may be an example of “coherent arbitrariness” (Ariely et al. (2003)), whereby workers do not have a good sense of what an appropriate bonus may be so respond similarly to bonuses of varying sizes.

Third, as demonstrated previously, there is evidence of advantageous selection on ability in the second stage of the experiment, and adverse selection in the high fixed pay treatment. Therefore, given the size of the effect of the penalty frame on acceptance rates, it is very striking that there is apparently no selection effect of incentive framing. All three versions of the model outlined predict some form of selection on ability and loss aversion (although not necessarily monotone). It would be particularly interesting to see whether this result replicates in other contexts.

Fourth, the relation between incentivized performance and loss aversion, measured by rejected lotteries, remains something of a puzzle. The model allows for a negative relationship, as observed, but predicts a relatively more positive relationship under the penalty frame, which is not observed. Indeed, since loss aversion is the proposed driving mechanism by which penalty framing increases effort, it is hard to imagine any model in which framing does not have a stronger effect on more loss averse workers, nor any effect on selection. I presented evidence (albeit suggestive at best) that penalties may have a discouragement effect on the most loss-averse participants. This again would benefit from further research to unpack how loss aversion drives the effect of penalties.

1.6 Conclusion

This paper analyzes the effects of framed incentive pay on worker recruitment, selection and performance. I find that penalty framed incentives increased the number of workers who accepted the job by 25 percent relative to economically equivalent bonus framed incentives. In addition the penalty frame increased performance on the job by around 0.2 standard deviations. The effect sizes are large relative to those for standard manipulations of incentive size: increasing the non-contingent pay from \$0.50 to \$2 increased the acceptance rate by 36 percent and performance by 0.2 standard deviations, while the effect on both outcomes of increasing the contingent component of pay from \$1.50 to \$3 was small and statistically insignificant. Interestingly, I find no evidence of either adverse or advantageous selection into the penalty framed contract, while I do find advantageous selection into the incentivized task as a whole.

I present further evidence that the relative attractiveness of the penalty frame is not driven by changed perceptions of the difficulty of the task, nor is it driven by the wording of the job offer that might differentially attract inattentive workers. I also demonstrate that a standard selection result obtains in this context, namely that the job offer with incentive pay attracted relatively high ability participants. Lastly, I present evidence that the effect of the framed incentive on performance did not wear off during the course of the experiment, although this does not imply that it would persist in a repeated contracting environment.

To explain the main results, I propose an extension of the Kőszegi and Rabin (2006, 2007) model of loss aversion that allows for the agent's reference point to depend on the principal's choice of incentive frame. The most natural specification of the model is counterfactual: it predicts that bonus contracts would be more attractive than penalty contracts. Two possible modifications bring the model more in line with the data. One allows the agent's perception of the value of the job offer to depend on their reference point: the higher reference point under the penalty contract makes the job appear more attractive, or the outside option appear less attractive. An alternative model is one of self-control, in which the agent values the penalty contract as a commitment device that encourages her to exert higher effort, again via the loss aversion mechanism.

Overall, the results are surprising, suggesting that penalty contracts may strongly outperform bonus contracts in some settings. In environments similar to the experimental context, such as short-term one-off engagements, firms may be able to gain through greater use of penalties. Future research could work to unpack the candidate mechanisms described in this paper, which have different implications for welfare: penalties may be welfare enhancing if they help workers to overcome self-control problems, but harmful if they lead workers to mistakenly over-value the job offer. The relationship between performance and loss aversion could also be explored in more depth in a targeted lab study using careful, incentivized measures of loss aversion.

However to answer the motivating question of why firms seem reluctant to use penalties in general, it is clear that more research is needed. Probably the most promising place to start is to study longer-term engagements or repeated contracting. This allows for two possible further effects. First, the effect of the frame may wear off over time as the workers' reference points adjust. If the frame led some workers to over-value the job offer, they may then drop out, implying that penalty contracts are bad for staff retention. Second, actually incurring a penalty (as opposed to not receiving a bonus) may have discouragement effects that are outside the model in this paper, plausibly harming future effort provision and staff retention.

1.7 Figures

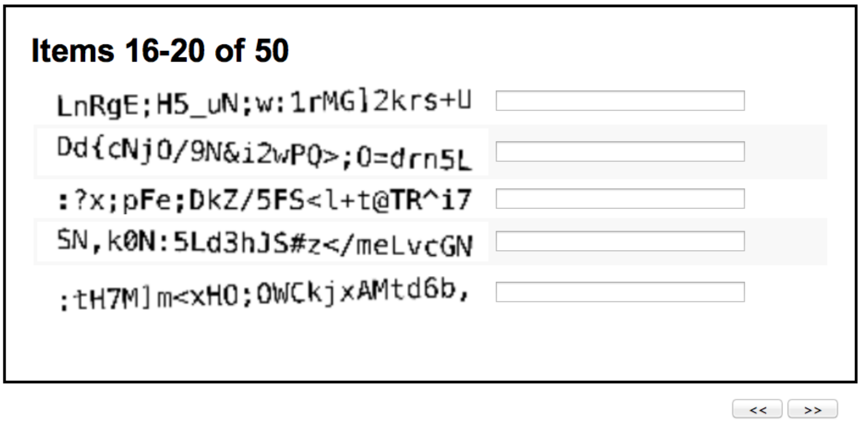

Figure 1.1: Example screen from the typing task.

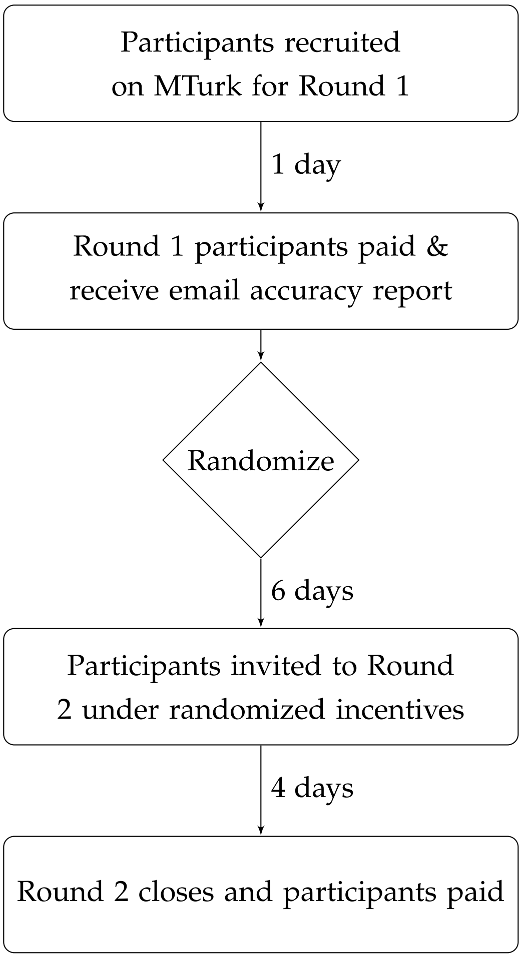

Figure 1.2: Experiment design flowchart.

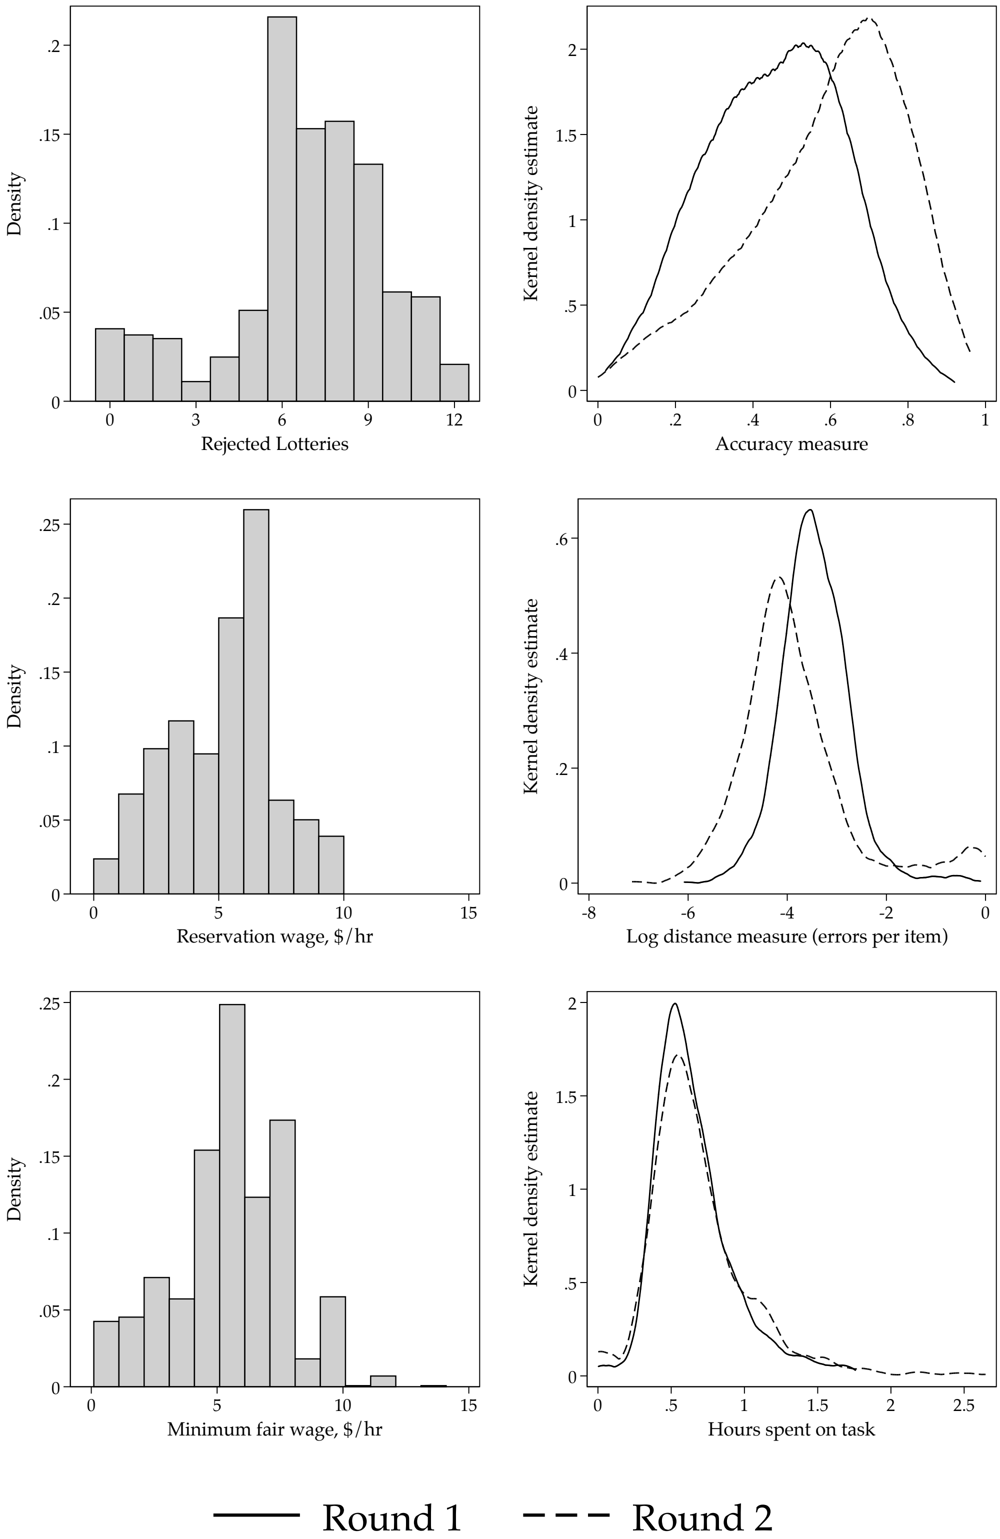

Figure 1.3: Distributions of survey responses and task performance. Reservation wage, fair wage and time spent trimmed at the 99th percentile (minimum acceptable wage = \$10, minimum fair wage = \$14, time = 1.75 hours (round 1), 2.65 hours (round 2)).

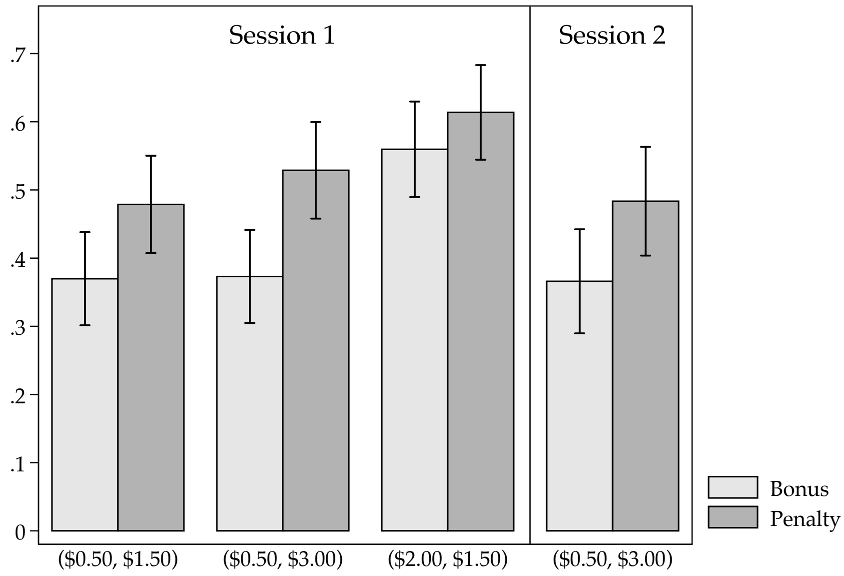

Notes: financial incentive levels given in parentheses as (fixed pay, variable pay). Error bars indicate 95% confidence intervals.

Figure 1.4: Acceptance Rates by treatment.

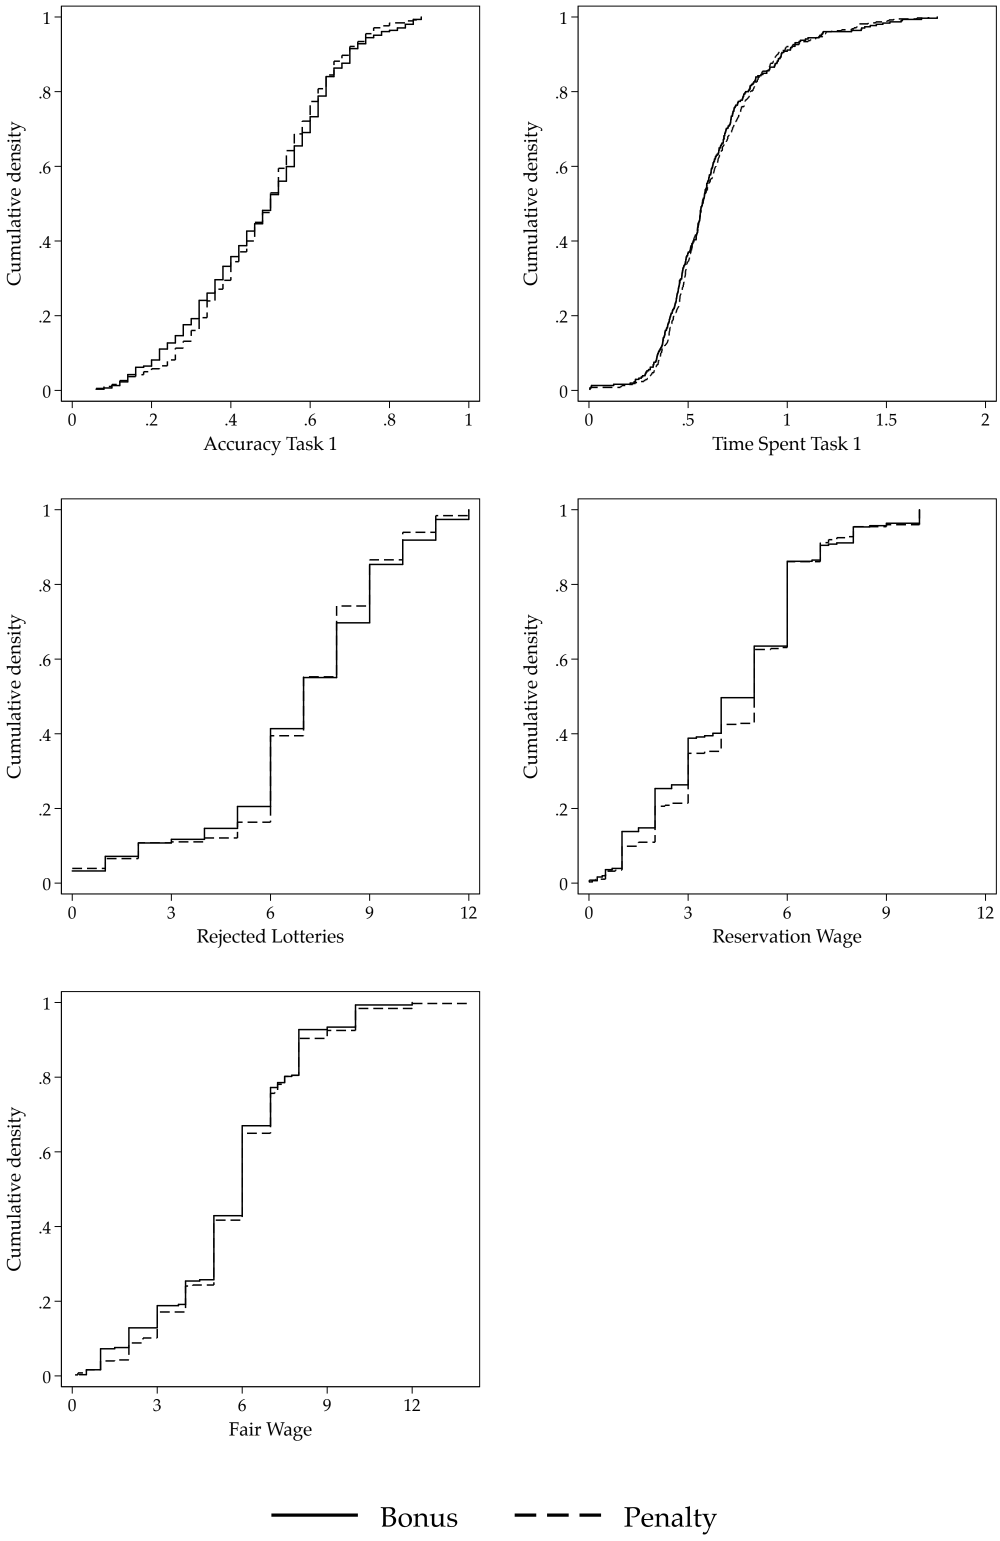

Figure 1.5: Comparing acceptors under Bonus and Penalty Frame. Reservation and fair wage trimmed at the 99th percentile.

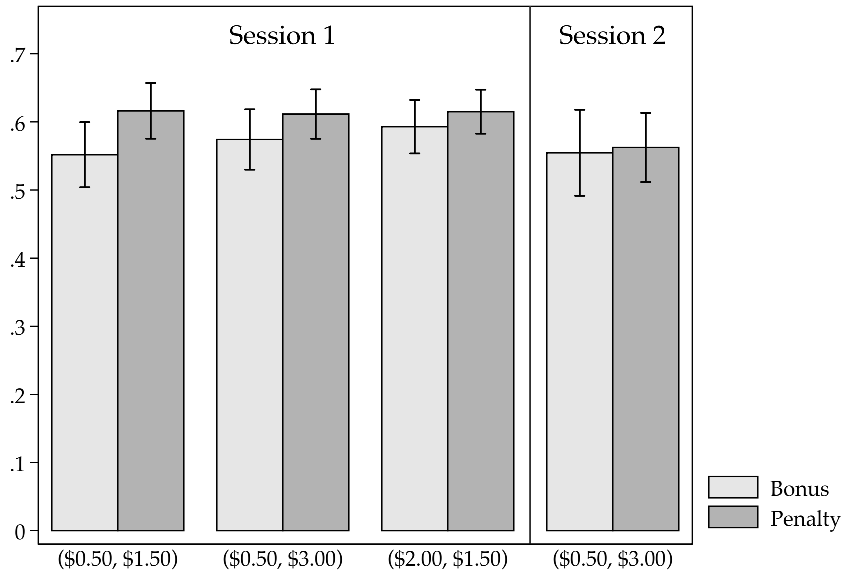

Notes: financial incentive levels given in parentheses as (fixed pay, variable pay). Error bars indicate 95% confidence intervals.

Figure 1.6: Accuracy by treatment.

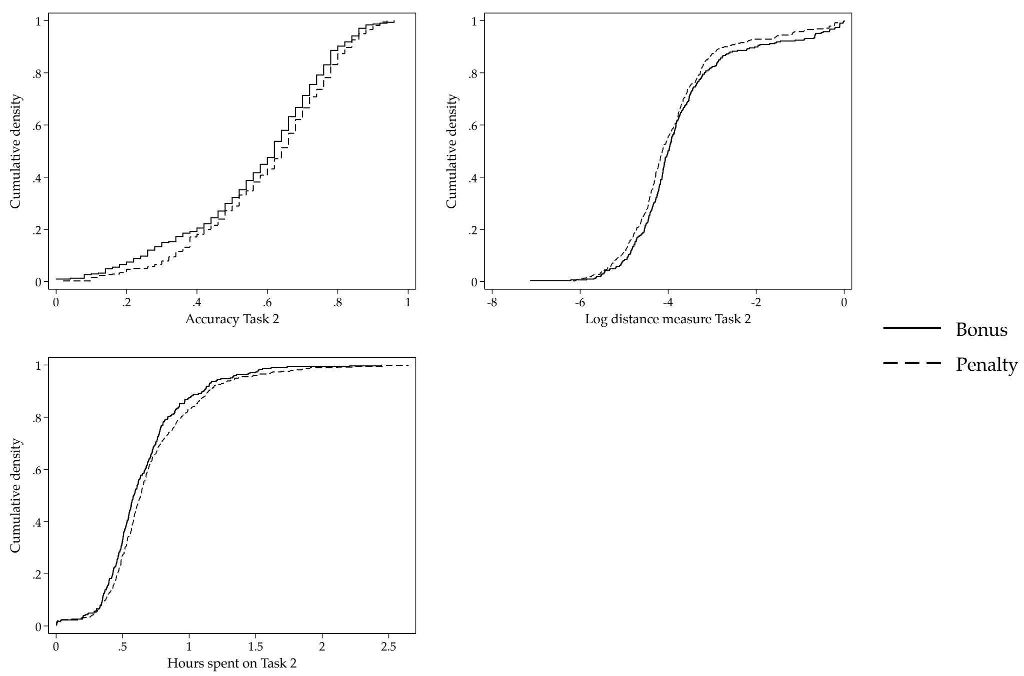

Figure 1.7: Performance and effort measures, comparing bonus vs penalty frame. Time spent is trimmed at the 99th percentile.

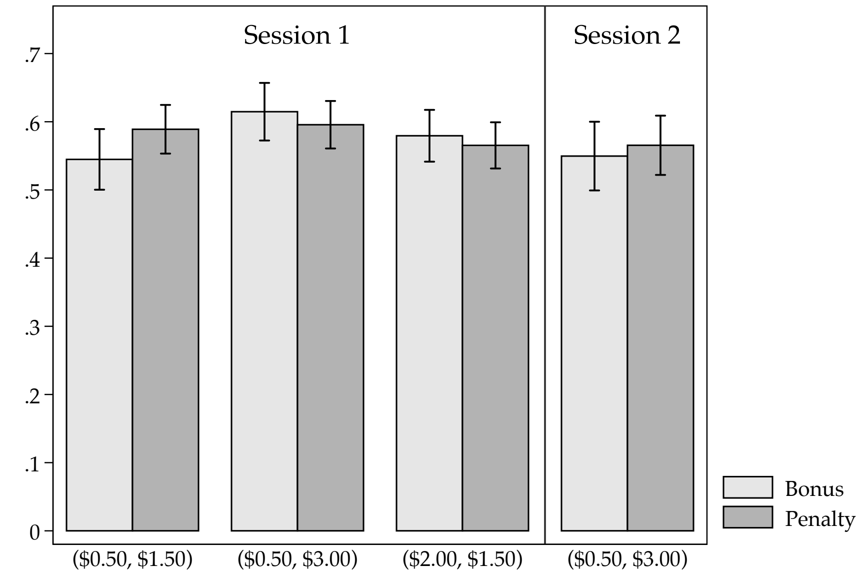

Notes: financial incentive levels given in parentheses as (fixed pay, variable pay). Error bars indicate 95% confidence intervals. The true mean task 1 performance was 0.46.

Figure 1.8: Participants' predictions of average task 1 accuracy by treatment.

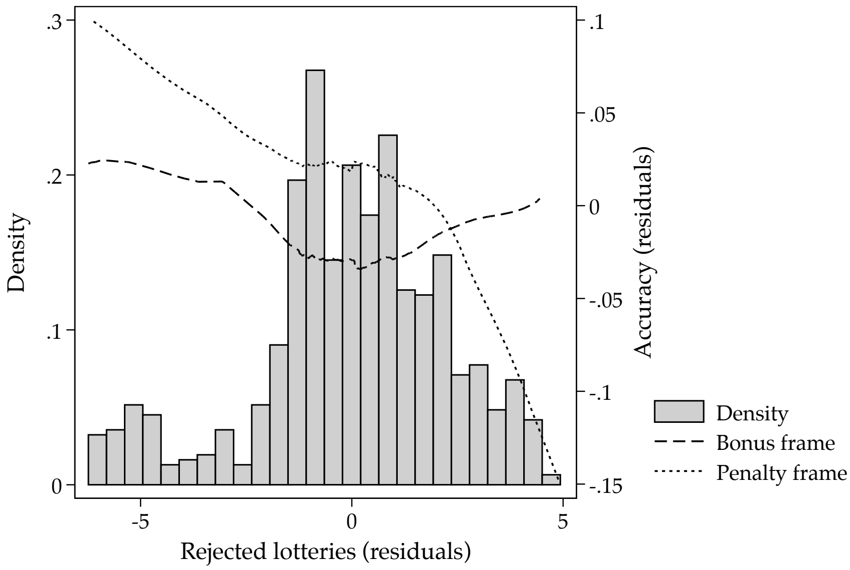

Notes: figure plots the residual of stage 2 task accuracy against residuals of rejected lotteries, after partialling out all other controls and treatment effects. Participants who rejected no lotteries or all lotteries, and participants who made inconsistent choices are dropped.

Figure 1.9: Relationship between Rejected Lotteries and performance.

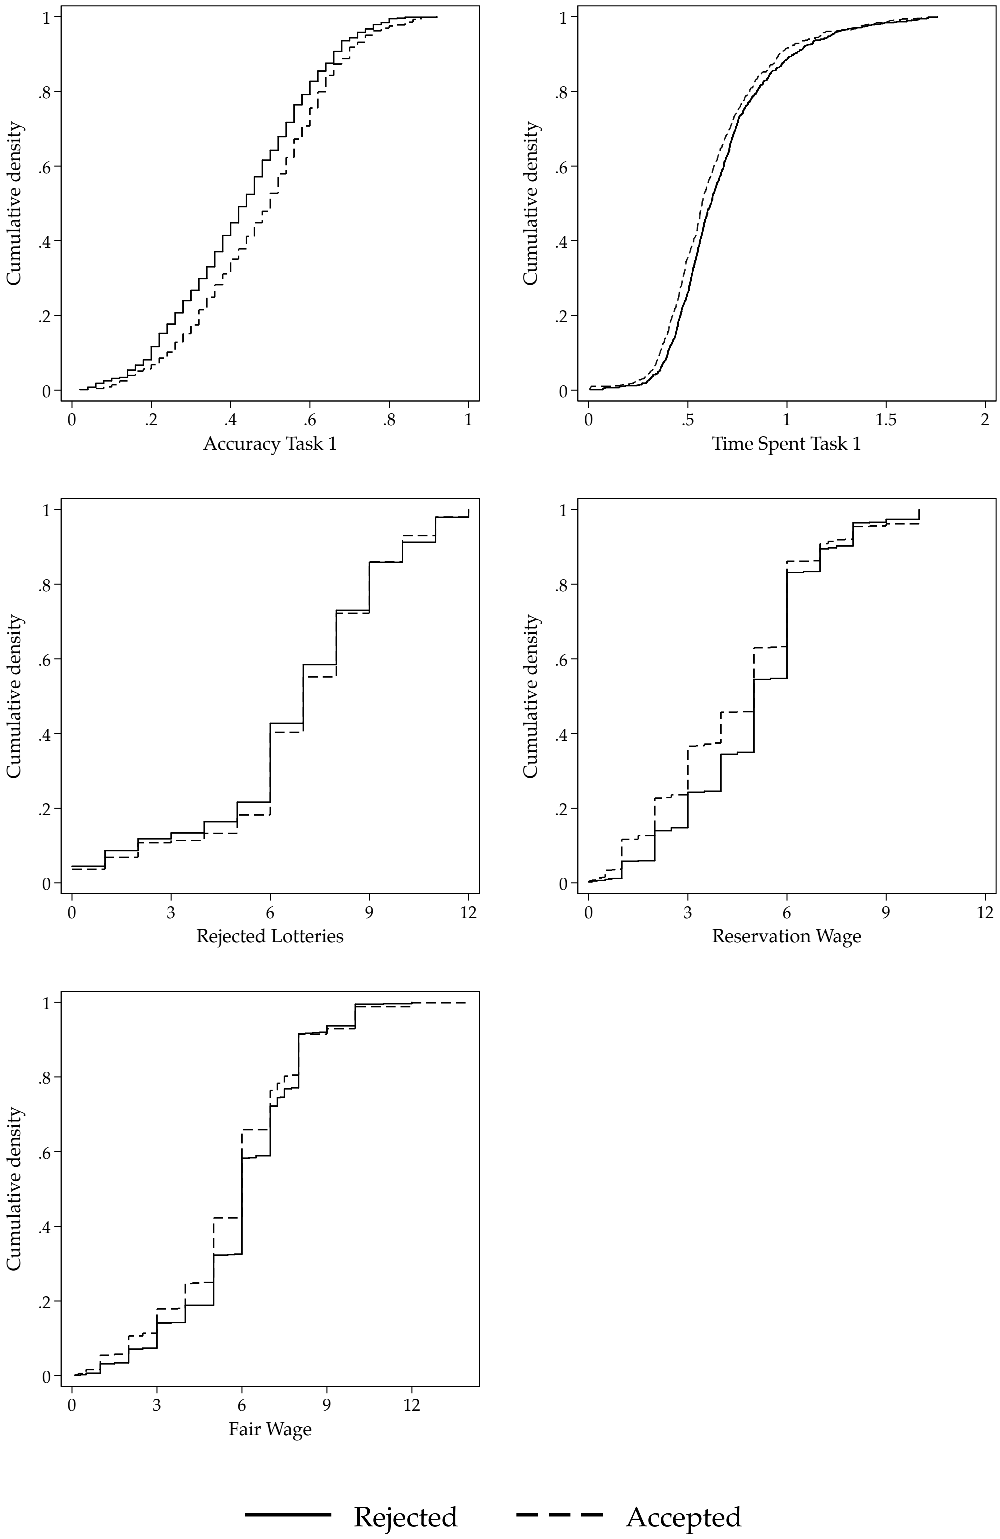

Figure 1.10: Comparing Acceptors vs Rejectors. Reservation and fair wage trimmed at the 99th percentile.

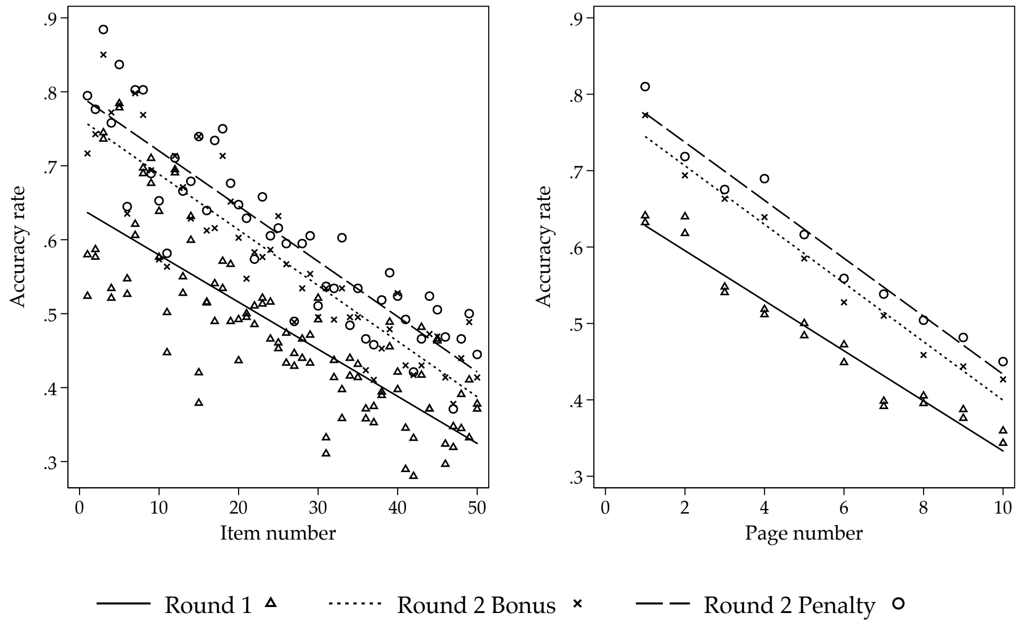

Notes: Plots the mean of performance for each typed item (page of 5 items) in stage 1, and each typed item (page) by framing treatment in stage 2.

Figure 1.11: Performance by item/page on the effort task.

1.8 Tables

Table 1.1: Treatments

		Contract[a]			Job offer[b]		
Group	N	Fixed pay	Variable pay	Frame	Base pay	Bonus	Penalty
Session 1							
0	192	$0.50	$1.50	Bonus	$0.50	$1.50	
1	188	$0.50	$1.50	Penalty	$2		$1.50
2	193	$0.50	$3	Bonus	$0.50	$3	
3	191	$0.50	$3	Penalty	$3.50		$3
4	193	$2	$1.50	Bonus	$2	$1.50	
5	189	$2	$1.50	Penalty	$3.50		$1.50
Session 2							
6	153	$0.5	$3.00	Bonus	$0.50	$3.00	
7	151	$0.5	$3.00	Penalty	$3.50		$3.00

[a] "Contract" details the three components of the contract offered: Fixed Pay (unconditional), Variable Pay (received if accuracy check is passed) and Frame.
[b] "Job offer" is the terms given in the email sent to subjects inviting them to the second stage.

Table 1.2: Performance on effort tasks

	N	Median	Mean	s.d.
Round 1				
Accuracy Task 1	1450	0.46	0.46	0.18
Errors per item Task 1	1450	0.031	0.044	0.063
Hours on Task 1	1448	0.60	0.67	0.31
Round 1, Rejectors only				
Accuracy Task 1	763	0.44	0.43	0.18
Errors per item Task 1	763	0.033	0.048	0.070
Hours on Task 1	761	0.62	0.69	0.32
Round 1, Acceptors only				
Accuracy Task 1	687	0.50	0.48	0.17
Errors per item Task 1	687	0.028	0.039	0.053
Hours on Task 1	687	0.57	0.64	0.28
Round 2				
Predicted Mean Round 1 Accuracy	686	0.60	0.58	0.19
Accuracy Task 2	687	0.62	0.59	0.20
Errors per item Task 2	687	0.017	0.066	0.17
Hours on Task 2	686	0.62	0.75	1.24

"Accuracy" is the fraction of effort task items a participant entered correctly. "Errors per item" is computed as the mean of the Levenshtein distance between entered and correct answers, scaled by string length. "Hours on Task X" is estimated by multiplying the median page time by 10 to account for outliers. "Round 1, Acceptors (Rejectors) only" gives the stage 1 scores of the participants who accepted (rejected) the job offer in stage 2. "Predicted Mean stage 1 Accuracy" is the participant's response to the question "Of the 50 items in the typing task you did before, how many do you think people entered correctly, on average?" Where timing data is missing (two observations in stage 1, one in stage 2) this is due to JavaScript errors. One participant did not report a predicted stage 1 accuracy.

Table 1.3: Summary Statistics from stage 1 Survey

	N	Mean	s.d.
Loss Aversion			
Rejected Lotteries	1450	6.82	2.78
Inconsistent Lottery Choices	1450	0.07	0.25
Reservation Wage			
Reservation wage, $/hr	1450	4.86	2.70
Minimum fair wage, $/hr	1450	6.00	2.56
MTurk Experience			
Hours working on MTurk per week	1450	16.5	15.3
Typical MTurk earnings, $100/week	1450	0.58	0.58
MTurk HITs completed	1450	7457	27548
Months of experience on MTurk	1450	11.7	13.6
Mainly participate in research HITs	1450	0.77	0.42
Work on MTurk mainly to earn money	1450	0.93	0.26
Demographics			
Age in 2013	1449	32.8	10.7
Male	1450	0.49	0.50
Household Income	1450	38362	27942
Zipcode cluster	1450	0.13	0.34
Employment Status			
Full time	1450	0.37	0.48
Part time	1450	0.14	0.35
Self employed	1450	0.10	0.30
Full time MTurk worker	1450	0.11	0.31
Unemployed	1450	0.12	0.33
Student	1450	0.11	0.31
Other	1450	0.05	0.21
Education			
Less than High School	1450	0.00	0.07
High School / GED	1450	0.11	0.31
Some College	1450	0.32	0.47
2-year College Degree	1450	0.12	0.33
4-year College Degree	1450	0.35	0.48
Masters Degree	1450	0.07	0.25
Doctoral/Professional Degree	1450	0.02	0.13

"Rejected lotteries" is the number (0-12) of 50-50 win-lose lotteries that participants report they would be unwilling to play. "Inconsistent Lottery Choices" indicates participants who indicated inconsistent preferences over lotteries (accepting a lottery dominated by one they rejected). "Minimum acceptable wage" is the minimum hourly wage at which participants are willing to work on MTurk. "Minimum fair wage" is the reported minimum fair hourly wage that requesters "should" pay on MTurk. "Mainly participate in research HITs" indicates participants who report mostly working on HITs posted by researchers. "Work on MTurk mainly to earn money" indicates participants main reason for working ("for fun", "something else" coded as zero). "Male" is a dummy, note that six transgender individuals are coded as zeros. "Household income" is calculated using midpoints of income bins ("0-$30,000", then $10,000 bins until "> $100,000"). Zipcode cluster is a dummy indicating at least one other participant in the same session reported the same zipcode. "Employment Status" and "Education" are sets of mutually exclusive dummy variables.

Table 1.4: Balance Check

	Joint		Groups 0 & 1		Groups 2 & 3		Groups 4 & 5		Groups 6 & 7	
	F-stat	p	Diff	p	Diff	p	Diff	p	Diff	p
Accuracy Task 1	0.90	0.51	-0.01	0.54	0.00	0.94	-0.00	0.84	0.01	0.59
Hours on Task 1	0.51	0.83	0.01	0.76	-0.00	1.00	0.01	0.68	-0.04	0.28
Rejected Lotteries	0.53	0.81	0.31	0.26	-0.08	0.79	-0.23	0.39	-0.17	0.60
Inconsistent Lottery Choices	0.62	0.74	0.01	0.60	0.02	0.42	-0.01	0.80	0.01	0.81
Reservation wage, $/hr	0.65	0.72	0.07	0.77	-0.02	0.93	0.10	0.77	0.06	0.84
Minimum fair wage, $/hr	2.87	0.01	0.08	0.73	0.05	0.88	-0.38	0.14	0.16	0.59
Hours working on MTurk per week	1.65	0.12	-0.28	0.83	-1.61	0.21	-1.47	0.29	0.80	0.75
Typical MTurk earnings, $100/week	0.35	0.93	-0.04	0.54	0.01	0.79	0.01	0.84	-0.00	1.00
MTurk HITs completed	1.85	0.07	-7499.66	0.09	3609.55	0.07	460.60	0.71	-1966.82	0.50
Months of experience on MTurk	0.67	0.70	1.04	0.49	0.50	0.68	-0.34	0.80	-1.55	0.37
Mainly participate in research HITs	0.44	0.87	-0.04	0.39	0.04	0.38	-0.00	0.99	0.02	0.73
Work on MTurk mainly to earn money	1.22	0.29	-0.03	0.31	-0.03	0.25	0.00	0.96	0.02	0.52
Age in 2013	0.73	0.65	1.38	0.18	-0.80	0.45	-0.90	0.42	0.53	0.69
Male	0.93	0.48	0.08	0.12	-0.00	0.93	0.01	0.92	-0.01	0.81

"Joint" reports the F-statistic and p-value from a joint test of the significance of the set of treatment dummies in explaining each relevant baseline variable. The remaining columns report the difference in means and p-value from the associated t-test between pairs of treatment groups, where pairs differ only in terms of the bonus/penalty frame.

Table 1.5: Acceptance decision

	OLS				Logit (marginal effects)	
	(1) Accepted	(2) Accepted	(3) Accepted	(4) Accepted	(5) Accepted	(6) Accepted
Penalty Frame	0.109***	0.110***	0.116***	0.104***	0.108***	0.103***
	(0.026)	(0.026)	(0.029)	(0.026)	(0.025)	(0.025)
High Fixed Pay	0.163***	0.160***	0.152***	0.154***	0.161***	0.153***
	(0.036)	(0.036)	(0.041)	(0.036)	(0.035)	(0.034)
High Variable Pay	0.027	0.025	0.008	0.020	0.015	0.022
	(0.036)	(0.036)	(0.040)	(0.036)	(0.032)	(0.035)
Accuracy Task 1		0.312***	0.136	0.337		0.282***
		(0.075)	(0.352)	(0.309)		(0.074)
Accuracy Task 1 ^2			0.106	-0.055		
			(0.380)	(0.337)		
Hours on Task 1		-0.139***	-0.197***	-0.164***		-0.170***
		(0.042)	(0.059)	(0.044)		(0.047)
Rejected Lotteries		0.005	-0.007	0.004		0.005
		(0.013)	(0.015)	(0.013)		(0.013)
Reservation wage		-0.021***	-0.013	-0.016**		-0.018**
		(0.005)	(0.009)	(0.005)		(0.006)
Fair wage		0.004	-0.006	0.005		0.007
		(0.006)	(0.009)	(0.006)		(0.006)
Session 2	-0.026	-0.011	-0.002	-0.012		-0.012
	(0.038)	(0.038)	(0.041)	(0.038)		(0.037)
Inconsistent Lottery Choices				-0.065		-0.068
				(0.050)		(0.051)
Set dummies	No	Yes	Yes	Yes	No	Yes
Controls	No	No	Yes	Yes	No	Yes
N	1450	1448	1146	1447	1450	1447
R-squared	0.031	0.066	0.119	0.114		
Pseudo R-squared					.022	.089
Mean dep. variable	0.474	0.474	0.480	0.475	0.474	0.475

Dependent variable is a dummy indicating whether the participant accepted the job offer in stage 2. Columns (1)-(4) report OLS linear probability regressions, columns (5) and (6) present average marginal effects logit equivalents to (1) and (4). Standard errors clustered at zipcode-session level in parentheses. + $p<0.10$, * $p<0.05$, ** $p<0.01$, *** $p<0.001$. Two participants are missing timing variable and one missing age variable. Column (3) drops participants who made inconsistent lottery choices or who are above the 99th percentile of reservation wage, fair wage or time spent on task 1, or who are from zipcodes with more than one respondent in that session. “Set dummies” indicate the set of strings participants typed in stage 1. “Controls” are Hours worked on MTurk per week, MTurk earnings per week, HITs completed, MTurk experience, “Mostly do research HITs”, “Mainly MTurk to earn money”, age, male, and income, occupation and education dummies. “Rejected lotteries” measured in standard deviations.

Table 1.6: Acceptance: Interactions

	(1) Accepted	(2) Accepted	(3) Accepted	(4) Accepted
Penalty Frame	0.104***	0.128***	0.072*	0.107*
	(0.026)	(0.030)	(0.036)	(0.050)
High Fixed Pay	0.154***	0.200***	0.155***	0.190***
	(0.036)	(0.045)	(0.036)	(0.050)
Penalty * High Fixed Pay		-0.092		-0.071
		(0.059)		(0.071)
High Variable Pay	0.020	0.020	-0.014	0.004
	(0.036)	(0.036)	(0.045)	(0.048)
Penalty * High Variable Pay			0.068	0.032
			(0.051)	(0.062)
Set dummies	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
N	1447	1447	1447	1447
R-squared	0.114	0.115	0.115	0.115
Mean dep. variable	0.475	0.475	0.475	0.475

Dependent variable is a dummy indicating whether the participant accepted the job offer in stage 2. Estimates from OLS linear probability model. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001. "Controls" are the full set of regressors from Table 1.5 specification (4). "Set dummies" indicate which set of strings participants typed in stage 1.

Table 1.7: Selection

	(1) Accepted	(2) Accepted	(3) Accepted	(4) Accepted	(5) Accepted	(6) Accepted	(7) Accepted
Penalty Frame	0.104***	0.104***	0.104***	0.104***	0.104***	0.104***	0.104***
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
Accuracy Task 1	0.288***	0.245*	0.289***	0.288***	0.288***	0.287***	0.239*
	(0.076)	(0.101)	(0.076)	(0.076)	(0.076)	(0.076)	(0.102)
Penalty * Acc. Task 1		0.086					0.098
		(0.145)					(0.148)
Hours on Task 1	-0.164***	-0.165***	-0.156**	-0.164***	-0.166***	-0.165***	-0.167**
	(0.044)	(0.044)	(0.058)	(0.044)	(0.044)	(0.044)	(0.058)
Penalty * Hours Task 1			-0.018				0.000
			(0.082)				(0.083)
Rejected Lotteries	0.004	0.004	0.004	0.008	0.004	0.004	0.008
	(0.013)	(0.013)	(0.013)	(0.018)	(0.013)	(0.013)	(0.018)
Penalty * Rej. Lotteries				-0.008			-0.007
				(0.025)			(0.025)
Reservation wage	-0.016**	-0.016**	-0.016**	-0.016**	-0.020**	-0.016**	-0.020**
	(0.005)	(0.005)	(0.005)	(0.005)	(0.007)	(0.005)	(0.007)
Penalty * Res. wage					0.010		0.011
					(0.010)		(0.011)
Fair wage	0.005	0.005	0.005	0.005	0.004	0.003	0.004
	(0.006)	(0.005)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
Penalty * Fair Wage						0.006	-0.001
						(0.010)	(0.012)
Set dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	1447	1447	1447	1447	1447	1447	1447
R-squared	0.114	0.114	0.114	0.114	0.114	0.114	0.115
Mean dep. variable	0.475	0.475	0.475	0.475	0.475	0.475	0.475

Dependent variable is a dummy indicating whether the participant accepted the job offer in stage 2. Estimates from OLS linear probability model. Standard errors clustered at zipcode-session level in parentheses. + $p<0.10$, * $p<0.05$, ** $p<0.01$, *** $p<0.001$. "Controls" are the full set of regressors from Table 1.5 specification (4) excluding Accuracy Task 1 squared. "Set dummies" indicate which set of strings participants typed in stage 1. Interaction variables are demeaned to stabilize interaction coefficients. "Rejected lotteries" is measured in standard deviations.

Table 1.8: Performance on stage 2

	(1) Accuracy	(2) Accuracy	(3) Accuracy	(4) Accuracy
Penalty Frame	0.033*	0.035**	0.025+	0.036**
	(0.015)	(0.012)	(0.014)	(0.012)
High Fixed Pay	0.018	0.032+	0.039*	0.038*
	(0.020)	(0.017)	(0.019)	(0.017)
High Variable Pay	0.007	0.023	0.018	0.025
	(0.021)	(0.018)	(0.022)	(0.018)
Accuracy Task 1		0.715***	0.986***	0.996***
		(0.037)	(0.188)	(0.172)
Accuracy Task 1 ^2			-0.296	-0.308+
			(0.190)	(0.176)
Hours on Task 1		-0.029	-0.043	-0.031
		(0.023)	(0.031)	(0.024)
Rejected Lotteries		-0.012+	-0.017*	-0.011+
		(0.006)	(0.007)	(0.006)
Reservation wage		-0.005+	-0.008*	-0.005+
		(0.003)	(0.004)	(0.003)
Fair wage		-0.001	0.000	-0.000
		(0.002)	(0.004)	(0.002)
Session 2	-0.036	-0.019	0.002	-0.016
	(0.025)	(0.019)	(0.021)	(0.019)
Inconsistent Lottery Choices				-0.073*
				(0.029)
Set dummies	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	687	687	550	687
R-squared	0.013	0.442	0.470	0.476
Mean dep. variable	0.590	0.590	0.596	0.590

Dependent variable is accuracy in the stage 2 typing task, measured as the fraction of items entered correctly. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001. Column (3) drops participants who made inconsistent lottery choices or who are above the 99th percentile of reservation wage, fair wage or time spent on task 1, or who are from zipcodes with more than one respondent in that session. “Set dummies” indicate which set of strings participants typed in stages 1 and 2. “Controls” are Hours worked on MTurk per week, MTurk earnings per week, HITs completed, MTurk experience, “Mostly do research HITs”, “Mainly MTurk to earn money”, age, male, and income, occupation and education dummies. “Rejected lotteries” is measured in standard deviations.

Table 1.9: Performance/effort on stage 2, alternative measures

	(1) Log distance	(2) Log distance	(3) Time spent	(4) Time spent
Penalty Frame	-0.192*	-0.209**	0.054*	0.028
	(0.091)	(0.080)	(0.027)	(0.023)
High Fixed Pay	-0.219+	-0.323**	0.005	0.023
	(0.118)	(0.113)	(0.035)	(0.031)
High Variable Pay	-0.015	-0.069	0.017	0.036
	(0.134)	(0.130)	(0.039)	(0.033)
Accuracy Task 1				0.152
				(0.384)
Accuracy Task 1 ^2				-0.084
				(0.382)
Log Errors per item Task 1		0.868***		
		(0.066)		
Hours on Task 1				0.735***
				(0.067)
Rejected Lotteries		0.036		-0.018
		(0.040)		(0.012)
Reservation wage		0.025		0.001
		(0.021)		(0.006)
Fair wage		-0.010		-0.001
		(0.020)		(0.005)
Session 2	0.158	0.031	0.008	-0.035
	(0.150)	(0.127)	(0.045)	(0.036)
Inconsistent Lottery Choices		0.305		-0.036
		(0.189)		(0.042)
Set dummies	No	Yes	No	Yes
Controls	No	Yes	No	Yes
N	687	687	680	680
R-squared	0.018	0.338	0.007	0.400
Mean dep. variable	-3.825	-3.825	0.687	0.687

Dependent variable in columns (1) and (2) is the log of the mean scaled Levenshtein distance (which can be interpreted as the mean number of errors per character in the typing task). Dependent variable in columns (3) and (4) is time spent on the stage 2 typing task, estimated as 10 times the median time spent per page of typed items, in addition dropping participants above the 99th percentile for time spent. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001. "Set dummies" indicate which set of strings participants typed in stages 1 and 2. "Controls" are Hours worked on MTurk per week, MTurk earnings per week, HITs completed, MTurk experience, "Mostly do research HITs", "Mainly MTurk to earn money", age, male, and income, occupation and education dummies. "Rejected lotteries" is measured in standard deviations.

Table 1.10: Round 2 Performance: Interactions

	(1) Accuracy	(2) Accuracy	(3) Accuracy	(4) Accuracy
Penalty Frame	0.036**	0.033*	0.044**	0.044+
	(0.012)	(0.015)	(0.016)	(0.027)
High Fixed Pay	0.038*	0.032	0.038*	0.039
	(0.017)	(0.022)	(0.017)	(0.026)
Penalty * High Fixed Pay		0.010		-0.001
		(0.026)		(0.034)
High Variable Pay	0.025	0.025	0.035	0.035
	(0.018)	(0.018)	(0.025)	(0.028)
Penalty * High Variable Pay			-0.017	-0.018
			(0.025)	(0.034)
Set dummies	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
N	687	687	687	687
R-squared	0.476	0.476	0.477	0.477
Mean dep. variable	0.590	0.590	0.590	0.590

Dependent variable is accuracy in the stage 2 typing task measured as the fraction of items entered correctly. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001. "Controls" are the full set of regressors from Table 1.8 specification (4). "Set dummies" indicate which set of strings participants typed in stages 1 and 2.

Table 1.11: Round 2 Performance: Heterogeneous effects

	(1) Accuracy	(2) Accuracy	(3) Accuracy	(4) Accuracy	(5) Accuracy	(6) Accuracy	(7) Accuracy
Penalty Frame	0.037**	0.037**	0.037**	0.037**	0.037**	0.037**	0.037**
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Accuracy Task 1	0.708***	0.745***	0.709***	0.708***	0.708***	0.705***	0.741***
	(0.037)	(0.053)	(0.037)	(0.037)	(0.037)	(0.038)	(0.054)
Penalty * Acc. Task 1		-0.071					-0.082
		(0.075)					(0.076)
Hours on Task 1	-0.031	-0.031	-0.019	-0.030	-0.031	-0.033	-0.010
	(0.024)	(0.024)	(0.036)	(0.024)	(0.024)	(0.024)	(0.037)
Penalty * Hours Task 1			-0.024				-0.040
			(0.044)				(0.046)
Rejected Lotteries	-0.011+	-0.011+	-0.011+	-0.005	-0.011+	-0.011+	-0.004
	(0.006)	(0.006)	(0.006)	(0.010)	(0.006)	(0.006)	(0.010)
Penalty * Rej. Lotteries				-0.012			-0.012
				(0.013)			(0.013)
Reservation wage	-0.005+	-0.005+	-0.005	-0.005	-0.004	-0.005+	-0.001
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.004)
Penalty * Res. wage					-0.002		-0.009+
					(0.005)		(0.005)
Fair wage	-0.000	-0.000	-0.000	-0.001	-0.000	-0.003	-0.004+
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.002)
Penalty * Fair Wage						0.006	0.010*
						(0.004)	(0.004)
Set dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	687	687	687	687	687	687	687
R-squared	0.473	0.474	0.474	0.474	0.474	0.475	0.479
Mean dep. variable	0.590	0.590	0.590	0.590	0.590	0.590	0.590

Dependent variable is accuracy in the stage 2 typing task measured as the fraction of items entered correctly. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001. "Controls" are the full set of regressors from Table 1.8 specification (4) excluding Accuracy Task 1 squared. "Set dummies" indicate which set of strings participants typed in stage 1. Interaction variables are demeaned to stabilize interaction coefficients. "Rejected lotteries" is measured in standard deviations.

1.A Model Setup

In this appendix I briefly outline the standard simplifying assumptions that I apply to the general KR model.

Formally, for stochastic consumption c distributed according to $G(c|e)$, reference point r distributed according to $H(r|e,F)$, weight η on gain-loss utility and effort cost $c(e)$, A's utility function is:

$$U(G|H,e,F) = \int m(c)dG(c|e) + \eta \iint \mu(m(c) - m(r))dG(c|e)dH(r|e,F) - c(e). \tag{1.12}$$

To focus the analysis I make a number of standard assumptions on U. Firstly, following Herweg et al. (2010) and Gill and Prowse (2012), I ignore risk aversion in the conventional sense, as well as Kahneman and Tversky (1979)'s diminishing sensitivity. Formally, risk aversion is ruled out by assuming that $m(x) = x$.[27] Diminishing sensitivity, under which the gain-loss function $\mu(x,r)$ is concave for $x \geq r$ and convex for $x < r$ is ruled out by adopting the standard (KR) functional form:

$$\mu(x) = \begin{cases} x & x \geq 0 \\ \lambda x & x < 0 \end{cases} \qquad \lambda \geq 1. \tag{1.13}$$

where λ is A's loss aversion parameter, reflecting the extent to which losses (outcomes below the reference point, r) are overweighted relative to gains.

Furthermore, following KR, Herweg et al. (2010) and Gill and Prowse (2012), I ignore Kahneman and Tversky (1979) and Tversky and Kahneman (1992)'s probability weighting, under which agents evaluate prospects according to decision weights rather than true probabilities (in particular overweighting unlikely events and underweighting likely ones). Thus the marginal distribution of c, g, is

$$g(c|e) = Pr(c = x|e) = \begin{cases} 1 - e & x = w \\ e & x = w + b \end{cases}. \tag{1.14}$$

Lastly, for simplicity I set $\eta = 1$.[28]

1.B Additional Results

1.B.1 Relation between stage 1 variables.

Before analyzing stage 2 behavior, it is instructive to analyze the relationships between the key variables measured in stage 1. The main variables of interest are the

[27] Loss aversion still implies *first order* risk aversion: an agent with reference point r is averse to fair lotteries with possible payoffs either side of the reference point.

[28] In many applications, η and λ are not separately identified; Herweg et al. (2010) also set $\eta = 1$. The presence of framing effects in my framework implies that η and λ play slightly different roles, however it is intuitively easier and analytically simpler to consider the role of loss aversion by focusing on comparative statics involving λ.

accuracy score, time spent on the effort task, the number of rejected lotteries, the reservation wage and the minimum fair wage reported by each participant. Table 1.12 reports correlations between these variables with and without dropping participants who made inconsistent lottery choices. Higher performance on the stage 1 task is significantly negatively correlated with time spent on the task. Reservation wage and minimum fair wage are strongly positively correlated with one another. Reservation wage is negatively correlated with time spent on task 1, but not with performance. Lastly, there is no economically or statistically significant correlation between reservation wages, the number of rejected lotteries, and performance on the first stage task (rejected lotteries are significantly positively correlated with performance, but only when including participants who made inconsistent choices). These results suggest that selection on, say, loss aversion, will not necessarily imply selection on ability.

Table 1.13 regresses stage 1 performance on various key variables. Time spent on stage 1 is robustly negatively correlated with performance on the effort task, with each additional hour spent associated with around 10 percentage points lower accuracy. Unlike the simple correlations, the minimum acceptable wage is now robustly negatively related with performance, perhaps because participants with higher reservation wages may be more likely to rush the task and therefore make more errors. A 1s.d. increase in the minimum acceptable wage ($2.70) is associated with 1.5-2.5 percentage points lower accuracy. In one specification, the minimum fair wage is positively associated with accuracy, controlling for reservation wages. Lastly, the number of rejected lotteries is positively associated with performance, although this is not statistically significant when participants who made inconsistent choices are dropped. The coefficient estimates imply a 0.6 to 1.3 percentage point improvement for a 1s.d. increase in the number of rejected lotteries.

Of the other variables, higher weekly Mturk earnings are negatively associated with performance, (although higher hours on MTurk positively associated, perhaps because those higher earnings are attained by expending less effort on each individual task). Participants who report that they mainly work on MTurk to earn money (93% of participants, other options "for fun", "other") perform around 5-6 percentage points better, men perform around four percentage points better, and participants who made inconsistent lottery choices around 8 percentage points worse.

Table 1.12: Correlation between key stage 1 variables

Panel A

	Accuracy T1	Hours on T1	Rej. lotteries	Res. wage	Fair wage
Accuracy T1	1				
Hours on T1	-0.174***	1			
Rej. lotteries	0.0678**	-0.0275	1		
Res. wage	-0.0403	-0.107***	-0.0864***	1	
Fair wage	-0.00588	-0.0426	-0.0446	0.513***	1
Observations	1450				

Panel B

	Accuracy T1	Hours on T1	Rej. lotteries	Res. wage	Fair wage
Accuracy T1	1				
Hours on T1	-0.173***	1			
Rej. lotteries	0.0429	-0.0160	1		
Res. wage	-0.0348	-0.114***	-0.0786**	1	
Fair wage	0.00456	-0.0414	-0.0386	0.507***	1
Observations	1351				

Correlation matrices. * $p<0.05$, ** $p<0.01$, *** $p<0.001$. Accuracy T1 is accuracy on the stage 1 typing task. Hours on T1 is the median time spent on a page of the stage 1 task, multiplied by 10. Panel A includes all participants, Panel B drops participants who made inconsistent lottery choices.

Table 1.13: Performance on stage 1

	(1) Accuracy	(2) Accuracy	(3) Accuracy	(4) Accuracy
Hours on Task 1	-0.107***	-0.131***	-0.096***	
	(0.015)	(0.022)	(0.018)	
Rejected Lotteries	0.009*	0.006	0.011*	0.013**
	(0.005)	(0.005)	(0.005)	(0.005)
Reservation wage	-0.005*	-0.009**	-0.005**	-0.005*
	(0.002)	(0.003)	(0.002)	(0.002)
Fair wage	0.002	0.006*	0.002	0.002
	(0.002)	(0.003)	(0.003)	(0.003)
Session 2	-0.020+	-0.017	-0.022+	-0.022+
	(0.011)	(0.013)	(0.012)	(0.012)
Hours working on MTurk per week			0.001	0.000
			(0.000)	(0.000)
Typical MTurk earnings, $100/week			-0.017+	-0.017+
			(0.009)	(0.009)
MTurk HITs completed			-0.000	-0.000
			(0.000)	(0.000)
Months of experience on MTurk			0.000	0.000
			(0.000)	(0.000)
Mainly participate in research HITs			0.002	-0.002
			(0.011)	(0.011)
Work on MTurk mainly to earn money			0.051**	0.057**
			(0.019)	(0.020)
Age			-0.000	-0.001
			(0.001)	(0.001)
Male			0.037***	0.038***
			(0.010)	(0.010)
Inconsistent Lottery Choices			-0.078***	-0.084***
			(0.017)	(0.017)
Constant	0.549***	0.562***	0.391***	0.337***
	(0.021)	(0.025)	(0.057)	(0.057)
Set dummies	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes
N	1448	1147	1447	1449
R-squared	0.053	0.056	0.107	0.082
Mean of dependent variable	0.456	0.463	0.456	0.455

Dependent variable is accuracy rate on stage 1 typing task. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001. All columns include dummies for the set of items typed. Column (2) drops participants who made inconsistent lottery choices, and those above the 99th percentile for time spent, reservation wage and fair wage. "Controls" are dummy variables for income, education and employment status. Rejected lotteries measured in standard deviations.

1.B.2 Balance checks

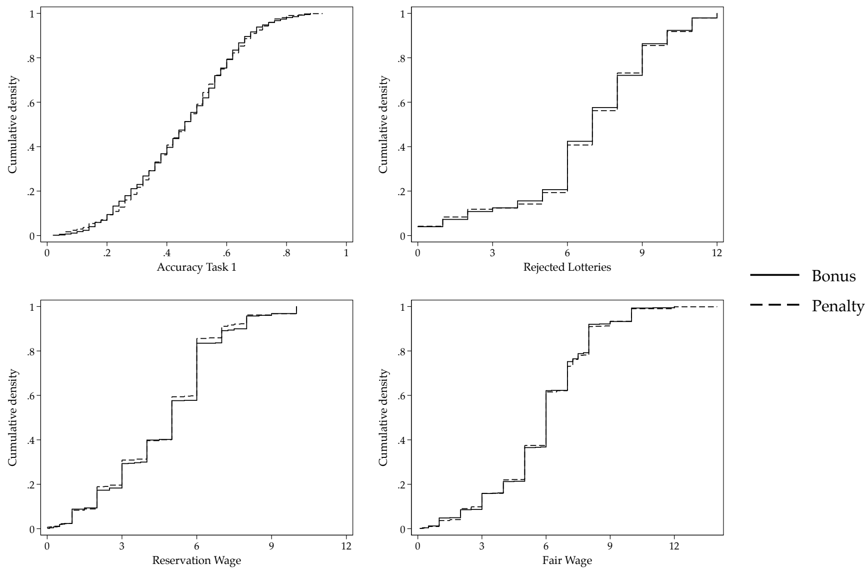

Figure 1.12: Balance between bonus and penalty treatments. Reservation wage and fair wage trimmed at the 99th percentile.

1.B.3 Comparing Session 1 and 2 Participants

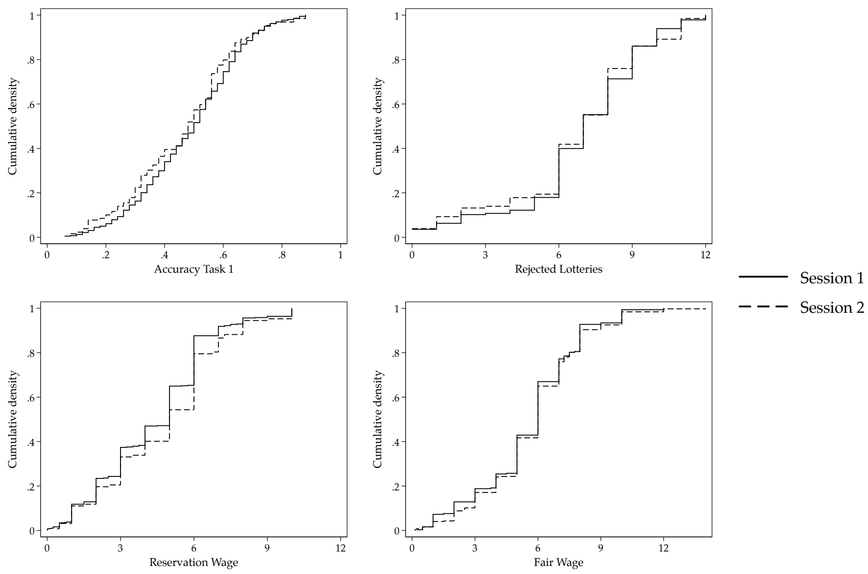

Figure 1.13: Comparing session 1 and 2 participants. Reservation wage and fair wage trimmed at the 99th percentile.

1.B.4 Selection: Robustness

Table 1.14: Selection: Robustness

	(1) Accepted	(2) Accepted	(3) Accepted	(4) Accepted
Penalty Frame	0.108***	0.102***	0.102***	0.105***
	(0.026)	(0.027)	(0.026)	(0.026)
Accuracy Task 1	0.286***	0.281***	0.278***	0.289***
	(0.076)	(0.078)	(0.076)	(0.076)
Hours on Task 1	-0.182**	-0.180***	-0.173***	-0.174***
	(0.069)	(0.046)	(0.043)	(0.044)
Penalty * Hours Task 1	0.018			
	(0.097)			
Rejected Lotteries	0.004	0.013	0.002	0.005
	(0.013)	(0.019)	(0.013)	(0.013)
Penalty * Rej. Lotteries		-0.021		
		(0.026)		
Reservation wage	-0.016**	-0.013**	-0.033***	-0.015**
	(0.005)	(0.005)	(0.009)	(0.005)
Penalty * Res. wage			0.021+	
			(0.012)	
Fair wage	0.005	0.003	0.004	-0.007
	(0.006)	(0.006)	(0.006)	(0.009)
Penalty * Fair Wage				0.014
				(0.012)
Set dummies	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
N	1433	1348	1433	1433
R-squared	0.113	0.114	0.117	0.115
Mean dep. variable	0.477	0.481	0.473	0.472

Dependent variable is a dummy indicating whether the participant accepted the job offer in stage 2. Estimates from OLS linear probability model. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001. Column (1) drops participants above the 99th percentile of time spent on task 1. Column (2) drops those who made inconsistent lottery choices. Column (3) drops participants above the 99th percentile of reservation wage. Column (4) drops participants above the 99th percentile of fair wage. “Controls” are the full set of regressors from Table 1.5 specification (4) excluding Accuracy Task 1 squared. “Set dummies” indicate which set of strings participants typed in stage 1. Interaction variables are demeaned in each specification to stabilize interaction coefficients. “Rejected lotteries” is measured in standard deviations.

1.B.5 Participants' predictions of Task 1 accuracy

Table 1.15: Participants' Predictions of stage 1 Accuracy

	(1) Predicted Acc.	(2) Predicted Acc.	(3) Accuracy	(4) Accuracy
Penalty Frame	0.004	0.001	0.033*	0.036**
	(0.015)	(0.014)	(0.015)	(0.012)
High Fixed Pay	0.003	0.010	0.016	0.037*
	(0.020)	(0.019)	(0.020)	(0.017)
High Variable Pay	0.034+	0.036+	-0.000	0.024
	(0.020)	(0.020)	(0.021)	(0.019)
Session 2	-0.045*	-0.035	-0.028	-0.017
	(0.022)	(0.022)	(0.025)	(0.019)
Predicted Acc.			0.188***	-0.006
			(0.041)	(0.032)
Set dummies	No	Yes	No	Yes
Controls	No	Yes	No	Yes
N	686	686	686	686
R-squared	0.008	0.150	0.043	0.476
Mean of dependent variable	0.577	0.577	0.590	0.590

Dependent variable in (1) and (2) is participant's prediction of mean stage 1 accuracy, and in (3) and (4) is participant's actual stage 2 accuracy. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001.

"Set dummies" indicate which set of strings participants typed in stage 1 (column (2)) or stages 1 and 2 (column (4)). "Controls" are the full set of regressors from the main specifications.

1.B.6 Comparing sessions 1 and 2

Table 1.16: Comparing Sessions 1 and 2

	(1) Accepted	(2) Accepted	(3) Accuracy	(4) Accuracy	(5) Log distance	(6) Log distance	(7) Time spent	(8) Time spent
Penalty	0.156**	0.137**	0.037	0.039+	-0.317+	-0.238	0.051	0.040
	(0.052)	(0.052)	(0.029)	(0.023)	(0.186)	(0.168)	(0.055)	(0.049)
Session 2	-0.007	-0.015	-0.020	-0.000	-0.032	-0.062	-0.043	-0.061
	(0.053)	(0.053)	(0.039)	(0.030)	(0.242)	(0.198)	(0.057)	(0.052)
Penalty * Session 2	-0.038	0.006	-0.029	-0.028	0.330	0.149	0.089	0.045
	(0.076)	(0.075)	(0.051)	(0.038)	(0.308)	(0.256)	(0.087)	(0.071)
Set dummies	No	Yes	No	Yes	No	Yes	No	Yes
Controls	No	Yes	No	Yes	No	Yes	No	Yes
N	688	1447	302	687	302	687	300	680
R-squared	0.021	0.115	0.012	0.477	0.013	0.343	0.017	0.402
Mean dep. variable	0.439	0.475	0.580	0.590	-3.729	-3.825	0.698	0.687

This table compares outcomes between treatment groups 2 and 3 (session 1) with groups 6 and 7 (session 2). Each received the same economic incentives (fixed pay of \$0.50 and variable pay of \$3.00), but the job offer was rephrased in session 2 to explore whether inattention might explain the session 1 results. Dependent variable in columns (1) and (2) is a dummy indicating whether the participant accepted the job offer in stage 2; in columns (3) and (4) is the participant's accuracy score in stage 2; in columns (5) and (6) is the "distance" accuracy measure described in table 1.9 and in columns (7) and (8) is the time spent by the participant on the task in stage 2. Columns (7) and (8) drop participants above the 99th percentile of time spent. Standard errors clustered at zipcode-session level in parentheses. + $p<0.10$, * $p<0.05$, ** $p<0.01$, *** $p<0.001$. Odd-numbered columns do not include controls and use data from groups 2, 3, 6 and 7. Even-numbered columns include the full sets of controls from the equivalent main specifications, and use data from all participants to increase precision of estimated coefficients on control variables, differencing out the main effects for treatment groups 0, 1, 4 and 5 using dummy variables.

1.B.7 Does the framing effect wear off?

Table 1.17: Performance by item in effort task

	(1) Item correct
Penalty Frame	0.038**
	(0.014)
High Fixed Pay	0.028+
	(0.017)
High Variable Pay	0.025
	(0.018)
Item	0.006***
	(0.002)
Item x Penalty	0.000
	(0.000)
Set dummies	Yes
Page dummies	Yes
Controls	Yes
N	34350
R-squared	0.126
Mean of dependent variable	0.590

Dependent variable is a dummy indicating whether an item was entered correctly. Standard errors clustered at zipcode-session level in parentheses. + p<0.10, * p<0.05, ** p<0.01, *** p<0.001. "Set dummies" indicate which set of strings participants typed in stage 1 and 2. "Page dummies" indicate which page the current item is found on. "Controls" are the full set of regressors from the main specifications.

1.B.8 Participant locations

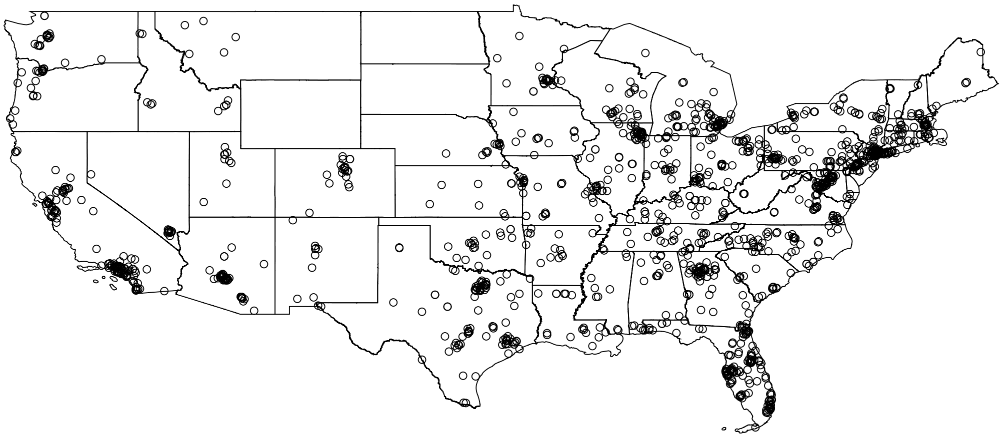

Figure 1.14: Participant locations by geocoded zipcodes

1.C Experimental details

1.C.1 Rates of pay

One potential concern about the study is the low size of the incentives used: the maximum a participant could earn in stage 1 is $3 and the maximum in stage 2 is $3.50. If a participant spent 30 minutes on stage 1 this implies a $6 hourly wage, significantly lower than the typical rates of pay in experimental labs, for example. The size of the incentive was selected, based on trial runs of the task, to match Amazon's suggested rate of pay of $6; the concern was that paying more would be perceived as unusual or not comparable to other jobs on MTurk. More critically, to study selection effects it is critical that rates of pay are low enough that some participants do not want to accept the job. The 47% acceptance rate in stage 2 suggests that rates were pitched about right. Taking into account the level of the incentives, the relative sizes of the fixed and variable pay in stage 2 were selected to be reasonably high-powered to emphasize the role of the penalty component of the contract.

1.C.2 Accuracy report after first stage

Shortly after the first stage was completed, subjects were sent an email informing them of their performance. The purpose of this was to ensure that they understood the difficulty level of the task and had at least a sense of their ability. An example message is given below:

```
Thanks for doing the typing task + survey HIT.
```

```
We have now processed the data and approved your work.  We estimate
that out of the 50 items, you entered 31 (62%) without errors.

Best wishes

Jon
```

1.C.3 Invitation to second stage

Session 1

One week after the first stage, all participants from the first stage were sent a second email inviting them to the second stage task under their randomly assigned incentive.

```
Thanks for participating in our recent typing task and survey.  You
are invited you to do another typing task (typing 50 text strings)
exactly like the one you did before.  There is no survey this time.

Pay:
The basic pay for the task is $3.50.  We will then randomly select
one of the 50 items for checking.  If you entered it incorrectly,
the pay will be reduced by $3.00.

If you would like to perform this task, please use the following
personalized link which will take you straight to the task.

https://lse.qualtrics.com/SE/?SID=SV_a5HXEhTVyucdg1f&MID=
XXXXXXXXXXXXX

Your MTurk ID (XXXXXXXXXXXXX) will be recorded automatically.  If
you don't want to do the task, you can just ignore this message.

The task will remain open for 4 days from the time of this message.
Payments will be made through the MTurk "bonus system" within 48
hours of the task closing.

Best wishes

Jon de Quidt
```

```
PS: We'll select the line to be checked using a random number
generator.  If you attempt the task more than once, only the first
attempt will be counted.
```

Following the link in the email took them to the experimental task, in which the first page contained the same text.

Session 2

In session 2, the invitation was altered slightly to explore whether inattention might be driving the results. The concern was that inattentive participants might quickly read the first sentence of the "Pay" paragraph, which gives the high (low) base pay under the penalty (bonus) frame, then ignore the rest and click through to the task. Then the high acceptance rate under the penalty frame might be explained by inattentive participants reading only the high base pay. This would be a concern for external validity since for longer term job offers where the stakes are higher, workers presumably read their contracts carefully. The challenge in addressing this concern is to find a manipulation that encourages participants to read the invitation carefully without also changing their reference point, since then it would not be clear which was driving any change in outcome.

The invitation and introduction were kept identical to session 1 except for the "Pay" paragraph which was rephrased as follows:

```
Pay:
The pay for this task depends on your typing accuracy.  We will
randomly select one item for checking, and if it was entered
incorrectly, the pay will be reduced below the base pay.  The base
pay is $3.50 which will be reduced by $3 if the checked item is
incorrect.
```

This phrasing first sets out the structure of the pay scheme, then gives all of the financial details in a single sentence, keeping the base pay and bonus/penalty amount as close to one another as possible.

1.C.4 Informed consent

After clicking the link on MTurk, but before beginning stage 1 of the experiment, participants were required to read and agree to an informed consent statement, which is reproduced below.

The task you are going to complete forms part of a study of the behavior of workers on MTurk. Data on your answers in the following task will be collected and analyzed by researchers at the London School of Economics.

Your participation is anonymous and no sensitive data will be collected. In addition, worker IDs will be deleted from any published data. Participation is voluntary and you can choose to stop at any time. There are no risks expected from your participation.

We would like you to complete a typing task and a short survey. For an average typing speed this should take around 30 minutes to complete. At the end you will be given a completion code. Please copy and paste the code into the HIT on MTurk to be paid. The payment for completion of the HIT is $3.

We may also contact you through MTurk to invite you to complete other HITs. This will not be affected by what you do in this HIT.

If you have any questions or concerns at any time, please feel free to contact the researcher, Jon de Quidt, at <MTurk contact address>.

If you are happy to proceed, please type "ACCEPT" into the box below and click through to the next page.

1.C.5 Lottery Questions

Survey part 2 of 3: Lottery Questions

In this section we are going to describe a series of choices to you. Each choice is a decision to play or not play an imaginary lottery in which you win with a 50% chance and lose with a 50% chance (for example, based on a coin toss).

For example, you might be asked if you would play the lottery "50% chance of winning $10 and 50% chance of losing $5".

Although the lotteries are imaginary, please carefully consider what you would choose if someone offered you the chance to play for real money.

>>

If someone trustworthy offered you the following lottery, would you accept?

50% chance of winning $10 | 50% chance of losing $0

YES I would play the lottery ○ NO I would not play the lottery ○

If someone trustworthy offered you the following lottery, would you accept?

50% chance of winning $10 | 50% chance of losing $1

YES I would play the lottery ○ NO I would not play the lottery ○

If someone trustworthy offered you the following lottery, would you accept?

50% chance of winning $10 | 50% chance of losing $2

YES I would play the lottery ○ NO I would not play the lottery ○

If someone trustworthy offered you the following lottery, would you accept?

50% chance of winning $10 | 50% chance of losing $3

YES I would play the lottery ○ NO I would not play the lottery ○

Figure 1.15: Introduction and examples of lottery questions.

Chapter 2

Market Structure and Borrower Welfare in Microfinance

> *Commercialization has been a terrible wrong turn for microfinance, and it indicates a worrying "mission drift" in the motivation of those lending to the poor. Poverty should be eradicated, not seen as a money-making opportunity.*
>
> Muhammad Yunus, New York Times, January 14th 2011[1]

Lately, microfinance has often been in the news for the wrong reasons. The success of microfinance institutions (henceforth, MFIs) across the world has been tremendous over the last two decades, culminating in the Nobel Peace Prize for the Grameen Bank and its founder Dr. Muhammad Yunus. However, in the last few years there has been some controversy about the activities of some MFIs that has stirred a broader debate about commercialization and mission drift in the sector.[2] There are concerns that some MFIs are profiteering at the expense of poor borrowers, attracted by the high repayment rates, and charging very high interest rates which seemingly contradicts the original purpose of the MFI movement, namely making capital accessible to the poor to lift them out of poverty.[3] While the discussion has been mostly about "commercialization", there is an implicit assumption that these lenders enjoy some market power, for example, in Yunus's statement that microcredit has "[given] rise to its own breed of loan sharks".[4] This critique is acknowledged within the MFI sector and has led to calls for tougher regulations, for example, in the form of a new bill, entitled the "Micro Finance Institutions (Development and Regulation) Bill" being tabled in the Indian Parliament.

[1] Accessible at `http://www.nytimes.com/2011/01/15/opinion/15yunus.html`

[2] For instance, SKS in Andhra Pradesh, India, Banco Compartamos of Mexico, LAPO of Nigeria. See, for example, MacFarquahr (New York Times, April 13, 2010), and Sinclair (2012).

[3] In addition, the results from several randomized experiments in India, Mongolia, Morocco, and the Philippines suggest that while microfinance has a positive effect in starting small businesses, but it did not have a statistically significant effect reducing poverty. See Banerjee et al. (2010), Attanasio et al. (2011), Crépon et al. (2011), and Karlan and Zinman (2009). By design these studies look at a single MFI and its borrowers rather than addressing industry or market level issues. Nevertheless the results suggest the need to look at factors that might be limiting the impact of microfinance on its stated goal of poverty alleviation.

[4] The New York Times, 14th January 2011.

The academic literature on microfinance, both theoretical and empirical, has not kept pace with these developments.[5] It has typically assumed lenders to be non-profits or to operate in a perfectly competitive market, and which more generally ignores the issue of market structure in considering the welfare effects of microfinance (see for example, the recent review by Banerjee (2013)). Our paper aims to fill this gap both theoretically and empirically. It is the first paper, to the best of our knowledge, that looks at the issues of commercialization and market power in the context of microfinance in a framework that allows us to understand and interpret many of the important issues that the recent controversies have thrown up.

Most of the existing work has looked at the remarkable repayment rates achieved by MFIs. In a world where lenders are not necessarily acting in the best interests of borrowers, we need to look beyond repayment rates. More broadly, the existing literature, both theoretical and empirical, has typically adopted a partial equilibrium framework focusing on one MFI and a given set of borrowers, whereas a lot of the issues these debates have thrown up require looking at the broader market and institutional environment within which a MFI operates. This allows us to evaluate borrower welfare beyond repayment rates - looking at the types of loans offered, interest rates, and the extent of credit rationing. In addition, we can show that some of the changes in the lending patterns - for example, some suggestive evidence that there has been some decline in joint liability loans relative to individual liability loans (see Giné and Karlan (2013) and discussion below), may indeed be related to changes in market structure, e.g., increasing commercialization.

Our paper analyzes the consequences of for-profit or commercial lending in microfinance, with and without market power, compared to a benevolent non-profit maximizing borrower welfare subject to a break-even constraint on these outcome variables. Much of the microfinance literature has shown how joint liability lending can be used by MFIs to leverage borrowers' social capital and local information in order to lend to otherwise unbankable customers and increase their welfare. Using a simple, tractable model we show that when the lender is a for-profit with market power he can instead leverage these to extract higher rents at the borrowers' expense. In particular, borrowers with more social capital may be worse off than those with less. However, given that borrowers are credit constrained and have very few outside options, they are better off borrowing than not borrowing, and they are better off borrowing under joint liability (when the lender chooses to use it) than under individual liability.

We then show that competition between for-profit lenders can close down this channel, but has an ambiguous effect on borrower welfare as competition undermines borrowers' incentives to repay their loans and thus leads to credit rationing. One of the interesting trade-offs that emerges therefore is that of rent extraction under monopoly with the enforcement externality under competition.

Lastly, we show that for-profit lenders - both with and without market power -

[5] Exceptions are Cull et al. (2007), Cull et al. (2009) and Baquero et al. (2012).

inefficiently *under-use* joint liability relative to the altruistic non-profit benchmark. The latter use joint liability whenever it is socially efficient, but the former use it only if it is profit-maximizing relative to using individual liability. Since joint liability is associated with tighter repayment incentive constraints (because larger amounts are due, when a group member is unable to pay her loan), it is relatively less attractive to for-profit lenders. This suggests that the apparent decline in the use of joint liability loans relative to individual liability loans may indeed be related to changes in market structure, e.g., increasing commercialization. This result is also consistent with the evidence presented in Cull et al. (2009) and in this paper that non-profits tend to use group-based lending methods, whereas for-profit lenders tend to use individual-based lending methods.

We then simulate the model using parameters estimated from the MIX Market (henceforth, MIX) dataset and existing research. The attempt to bridge theory and policy debates via quantitative analysis is a novel aspect of the paper. We initially expected that the monopolist's ability to leverage borrowers' social capital would have large welfare effects. We find that forcing the monopolist to use JL when he would prefer IL increases borrower welfare by a minimum of 12% and a maximum of 20%. Meanwhile, switching to a non-profit lender delivers a much larger welfare gain of between 54% and 73%. The qualitative sizes of these effects are robust to alternative parameter values. Secondly, we find that despite its effect on undermining repayment incentives, competition delivers similar borrower welfare to the non-profit benchmark. Taking these results together suggests that regulators should be attentive to lenders with market power, but that fostering competition rather than heavy-handed regulation can be an effective antidote. Thirdly, we find that for our parameter values, the non-profit lender would offer JL to all borrowers, irrespective of their level of social capital. The for-profit lenders, with and without market power, only switch to JL lending when borrowers have social capital worth around 15% of the loan size.

In support of the context and predictions of the model, we document three stylized facts. Firstly, there has been a steady increase in the market share of for-profit lenders. In our sample, this rises from 32 percent of institutions (38 percent of loans) in 1996, to 39 percent of institutions (46 percent of loans) in 2009. We graph these trends in Figure 2.1.[6]

The second stylized fact is that that non-profit lenders are more likely to use group-based lending methods than for-profit lenders. This observation is also documented by Cull et al. (2009). In our full sample of 712 MFIs with both legal status and lending methodology data for 2009, the mean share of "solidarity group" loans for non-profits is 37 percent, while for for-profits it is 34 percent. In the balanced panel used to construct Figure 2.2 (described below), the corresponding figures are

[6]Figure constructed from the cross-section of 1,106 MFIs that reported to the MIX Market dataset in 2009. "Unweighted" counts the number of institutions in existence at a given date, "weighted" weights institutions by size, measured as number of loans outstanding in 2009. See Appendix 2.C for further details including a discussion of weaknesses.

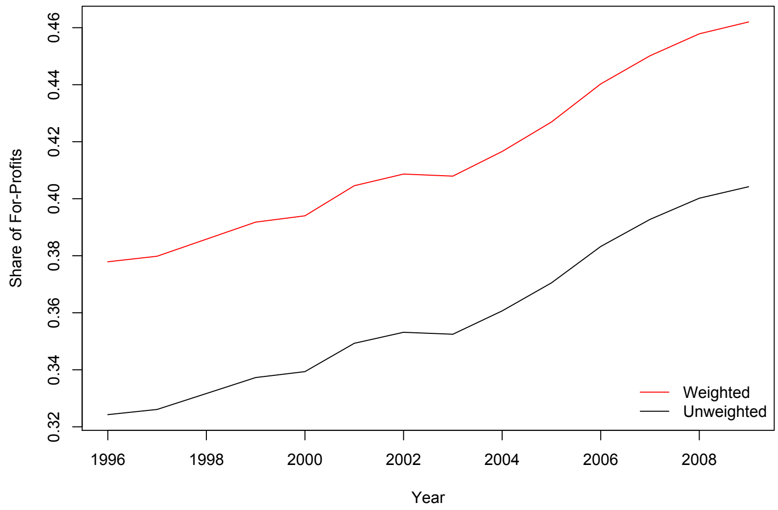

Figure 2.1: Market share of for-profit lenders by number of institutions ("unweighted") and number of loans ("weighted").

33 percent and 25 percent.[7]

Third, anecdotal evidence following the move by Grameen Bank and BancoSol, among others, to switch to IL lending, points to a decline in the use of joint liability. In the best data we have available, there does appear to have been a modest decline in the use of joint liability in recent years. Taking the balanced panel of 333 MFIs in the MIX Market dataset that report lending methodology information in 2008, 2009 and 2010, we show in Figure 2.2 that 31 percent of the average MFI's loans were made to solidarity groups in 2008, falling to 28.5 percent in 2010.[8] We see a much larger drop in the fraction of all loans made to solidarity groups, because two very large lenders (BRAC in Bangladesh and Bandhan in India) switch from solidarity group to individual in 2009. Dropping these two, we again find a modest decline in JL. We do note however, that groups are still widely used.[9]

Turning to related literature, our model is along the lines of Besley and Coate (1995) who show how JL can induce repayment guarantees within borrowing groups, with lucky borrowers helping their unlucky partners with repayment when needed.

[7]Since data on formal liability structure are not available, we treat loans recorded as "solidarity group" as joint liability, and "individual" as individual liability. See Appendix 2.C for discussion of how these numbers were calculated, and potential weaknesses.

[8]To construct the figure, we again use the measure "Solidarity Group Share" which is the ratio of each lender's number of solidarity group loans to solidarity group and individual loans. We then compute the mean of this measure across all MFIs to find the average solidarity group share ("unweighted"). We compute the weighted mean (by number of loans outstanding) to find the share of JL in all loans, with/without two outliers.

[9]Our numbers on the decline of JL relate to Giné and Karlan (2013). They use the MIX Market dataset as well and document a large fall in the fraction of MFIs using only JL lending, and rise in the number of MFIs using only IL lending, between 2007 and 2009. However the data they use has only 31 observations in 2007. For this reason we prefer to use 2008-2010 data. We also compute our solidarity group shares slightly differently. We thank Xavier Giné for sharing the data used to construct their figures.

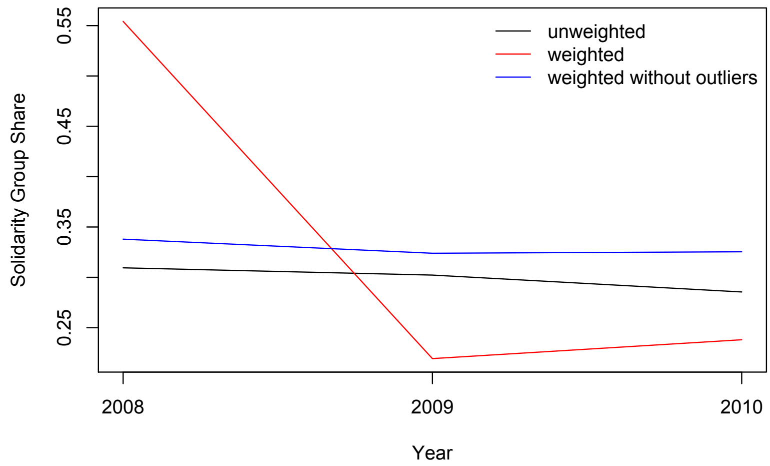

Figure 2.2: Fraction of loans disbursed via solidarity groups in the average lender's portfolio ("unweighted"), in all loans ("weighted"), and in all loans, excluding two large outliers ("weighted without outliers").

They show a trade-off between improved repayment through guarantees, and a perverse effect of JL, that sometimes a group may default *en masse* even though one member would have repaid had they received an IL loan. Introducing social sanctions, they show how these can help alleviate this perverse effect by making full repayment incentive compatible in more states of the world, generating welfare improvements that can be passed back to borrowers. Rai and Sjöström (2004) and Bhole and Ogden (2010) are recent contributions to this literature, both using a mechanism design approach to solve for efficient contracts (neither include the social capital channel). Within the microfinance literature there are various approaches to modeling social capital. For instance Besley and Coate (1995) model an exogenously given social penalty function, representing the disutility an agent can impose on her partner as a punishment. We model social capital in a similar reduced-form way as a sanction worth S that a borrower can impose on a partner in response to a violation of an informal contract, thus social capital in our model is a measure of the strength of informal contracting.[10]

There are a number of empirical studies of the role of social capital in group borrowing.[11] Feigenberg et al. (2013) study the effect of altering loan repayment frequency on social interaction and repayment, claiming that more frequent meetings can foster the production of social capital and lead to more informal insurance within the group. It is this insurance or repayment guarantee channel on which our

[10]Alternative approaches include Greif (1993), where deviations in one relationship can be credibly punished by total social ostracism. Bloch et al. (2008) and Karlan et al. (2009) Jackson et al. (2011) present models where insurance, favor exchange or informal lending are embedded in social networks such that an agent's social ties are used as social collateral to enforce informal contracts.

[11]See, for example Cassar et al. (2007), Wydick (1999), Karlan (2007), Giné et al. (2010).

model focusses. They also highlight that peer effects are important for loan repayment, even without explicit JL, through informal insurance, and that these effects are decreasing in social distance. There is also some emerging evidence on the relative roles of IL and JL. Giné and Karlan (2013) and Attanasio et al. (2011) find no significant difference between group and individual repayment probabilities, although repayment rates are very high under both control and treatment groups. They are not strictly comparable as in the first study groups were retained under IL while in the second groups are not used either under IL or JL.

The plan of the paper is as follows. In section 2.1 we lay out the basic model and analyze the choice of contracts by a non-profit lender who maximizes borrower welfare and a for-profit monopolist. In section 2.2 we analyze the effects of introducing competition to the market. We then simulate the model in section 2.3, allowing an empirical interpretation of the key mechanisms analyzed. Section 2.4 concludes.

2.1 The Model

We assume that there is a set of risk neutral agents or "borrowers", each of whom has access to a technology costing one unit of output each period that produces R units of output with probability $p \in (0, 1)$ and zero otherwise. Project returns are assumed to be independent. In each period the *state* is the vector of output realizations for the set of borrowers under consideration, so when we consider an individual borrower the relevant state is $Y \in \{0, R\}$, while for a pair of borrowers it is $Y \in \{(0, 0), (0, R), (R, 0), (R, R)\}$. The outside option of the borrower is assumed to be zero. Borrowers do not save and have no assets, so they must borrow 1 unit of output at the start of the period to finance production, and consume all output net of loan repayments at the end of the period. Since they have no assets their liability in any given period is limited to their income in that period. Borrowers have infinite horizons and discount the future with factor $\delta \in (0, 1)$. Throughout the paper, we will use "hat" notation ($\hat{x}$) to denote interest rates, utilities, etc arising under the non-profit, "tilde" ($\tilde{x}$) for the monopolist, and "double tilde" ($\tilde{\tilde{x}}$) for competition.

Each period, the state is common knowledge for the borrowers but not verifiable by any third party, so the lender cannot write state-contingent contracts. Borrowers can write contingent contracts with each other but these can only be enforced by social sanctions.

There is a single lender who may be a non-profit who is assumed to choose a contract that maximizes borrower welfare subject to a zero-profit condition, or alternatively a for-profit who maximizes profits.[12] The lender's opportunity cost of funds is $\rho \geq 1$ per unit. We assume (purely for simplicity) that the for-profit lender does not discount, i.e. he chooses the contract to maximize current-period profits only. We also assume that the lender has sufficient capacity to serve all borrowers that want credit.

[12]We do not explore the organizational design issues that might cause non-profits to behave differently than postulated above, as for example in Glaeser and Shleifer (2001).

Since output is non-contractible, lenders use dynamic repayment incentives as in, for example, Bolton and Scharfstein (1990). Following much of the microfinance literature we focus attention on IL or JL contracts. The IL contract is a standard debt contract that specifies a gross repayment r, if this is not made, the borrower is considered to be in default and her lending relationship is terminated. Under JL, pairs of borrowers receive loans together and unless both loans are repaid in full, both lending relationships are terminated. The lender can choose the interest rate and whether to offer IL or JL. Borrowers are homogeneous in the basic model so the lender offers a single contract in equilibrium.

We assume the lender commits to a contract in the first period by making a take-it-or-leave it offer specifying r and either IL or JL. Borrowers may then agree on an intra-group contract or *repayment rule*, specifying the payments each borrower will make in each possible state of the world.

Throughout the paper we assume the following timing of play. In the initial period:

1. The lender enters the community and commits to an interest rate and either IL or JL for all borrowers.
2. Borrowers may agree a repayment rule.

Then, in this and all subsequent periods until contracts are terminated:

1. Loans are disbursed, the borrowers observe the state and simultaneously make repayments (the repayment game).
2. Conditional on repayments, contracts renewed or terminated.

2.1.1 Intra-group contracting

Under JL, borrowers form groups of two individuals $i \in \{1, 2\}$, which are dissolved unless both loans are repaid. Once the loan contract has been written the borrowers agree amongst themselves and commit to a repayment rule or repayment guarantee agreement that specifies how much each will repay in each state in every future period.[13] In order to prevent the group from being cut off from future finance, a borrower may willingly repay the loan of her partner whose project was unsuccessful. We assume that deviation from the repayment rule is punished by a social sanction of size S, introduced in 2.1.2 below.[14] Some examples of possible rules are "both borrowers only repay their own loans," or "both repay their own loans when they can, and their partner's when she is unsuccessful."

[13]It is plausible that such agreements could expand to include others outside the group. For simplicity we assume that this is not possible, perhaps because borrowers' output realizations or borrowing and repayment behavior are only observable to other borrowers within their group.

[14]In a related paper, de Quidt et al. (2013), we show how social sanctions can enable borrowers to guarantee repayments even without an explicit JL clause. In the framework of this paper such behavior will not arise in equilibrium.

The agreed repayment rule can be seen as a device that fixes the payoffs of a two-player "repayment game" for each state of the world. Since the state is common knowledge to the borrowers, each period they know which repayment game they are playing. Either a borrower pays the stipulated amount, or she suffers a social sanction and may also fail to ensure her contract is renewed. The repayments stipulated in the rule must constitute a Nash equilibrium (i.e. be feasible and individually incentive-compatible). We assume that the pre-agreed rule enables the borrowers to coordinate on a particular equilibrium by fixing beliefs about their partner's strategy. This in turn implies that social sanctions never need to be enacted on the equilibrium path since there will be no deviations from the rule and both borrowers know the state.

For simplicity, we restrict attention to repayment rules that are symmetric (i.e., do not condition on the identities of the players), and stationary (depend only on the current state and social capital). Thus we can focus on a representative borrower with a time-independent value function. Symmetry prevents one borrower from taking advantage of the other using the threat of social sanction as leverage. Furthermore we focus on repayment rules that induce a joint welfare maximizing equilibrium. This implies that the total repayment in any state will be either zero or $2r$, and that social sanctions will not be used on the equilibrium path but only to punish off-equilibrium deviations.

2.1.2 Social Sanctions

A central theme in the microfinance literature is how the innovative lending mechanisms used by MFIs can harness social capital and local information among borrowers to overcome standard asymmetric information problems that make profitable lending to the poor difficult. With altruistic or competitive lenders the typical result is that the greater the lender's ability to access these, the better.

In this paper we show how under market power this result is reversed. Specifically, joint liability borrowers with a lot of social capital can be worse off than those with a little. This mirrors recent work on property rights (Besley et al. (2012)) that shows that in an insufficiently competitive market, an improvement in borrowers' ability to collateralize their assets can make them *worse off*, in contrast to the standard view of the "de Soto effect." This, however only represents a partial point against joint liability lending: it turns out that banning the use of joint liability would make borrowers still worse off.

There are many possible ways to model social capital. We adopt a very simple reduced form approach. We model social capital as borrowers' ability to enforce informal contracts amongst themselves. Such contracts specify actions that a borrower must take in certain states of the world, and if she deviated from the agreement she is punished by a sanction worth S in utility terms.[15] Borrowers in our model form

[15]This is closely related to the approach of Besley and Coate (1995) and the informal sanctions in Ahlin and Townsend (2007).

loan guarantee or informal insurance arrangements to assist one another in times of difficulty, backed by social sanctions. With an altruistic or competitive lender, the standard intuition follows: larger social sanctions can support more efficient lending contracts, increasing borrower welfare. When the lender has market power, more social capital still increases efficiency, but the lender can exploit the borrowers' sanctioning ability to extract more rents, potentially making them worse off.

In the core model we assume that S is observable and homogeneous across borrowers, and explore the comparative statics of varying the borrowers' informal enforcement ability. In an extension in section 2.1.4, we discuss the consequences of heterogeneity.

2.1.3 Loan contracts

With a single lender, contract termination means no credit ever again (unlike under competition in section 2.2, when a borrower cut off by one lender can later obtain a loan from another). Since borrowers must be given a rent for dynamic incentives to be effective, any incentive-compatible contract will satisfy their participation constraint.

If a borrower's contract is renewed with probability π, it must be that her expected per-period repayment is πr. Thus the value of access to credit for a representative borrower is $V = pR - \pi r + \delta \pi V$, which simplifies to:

$$V = \frac{pR - \pi r}{1 - \delta \pi}. \tag{2.1}$$

We can use (2.1) to derive the first incentive constraint on the lender. No borrower or group of borrowers will repay a loan if $r > \delta V$, i.e. if the benefit of access to future credit is worth less than the interest payment. This reduces to the constraint $r \leq \delta pR$, which we term Incentive Constraint 1 (IC1). We define r_{IC1} as the interest rate at which IC1 binds:

$$r_{IC1} \equiv \delta pR.$$

When IC1 binds, $V = pR$. This caps the lender's rent extraction: borrowers cannot be made worse off than if they took one loan and defaulted immediately.

Joint liability and social capital exploitation

First we consider joint liability lending. Recall that the lender observes the borrowers social capital, S, then offers a contract.

Suppose the lender offers a contract with interest rate r, satisfying IC1. The borrowers now have to agree a repayment rule to maximize joint welfare. Since IC1 is satisfied, joint welfare is higher when both loans are repaid than when both default, so the optimal rule will repay both loans in all states except $(0, 0)$ (when repayment is not possible). A minimal symmetric rule that achieves this is "both repay own loans in state (R, R), and the successful borrower bails out her partner in states $(R, 0)$ and $(0, R)$." Under this rule, each loan is repaid with probability $1-(1-p)^2$, which simplifies to:

$$q \equiv p(2-p).$$

Repayment of own loans in state (R, R) is incentive compatible by IC1 (borrower stands to lose at least δV if she does not repay r when her partner is also repaying r). Now suppose borrower j is called upon to assist i. If she does not, she loses future credit access, worth V, and is socially sanctioned next period, costing S. Thus the following incentive constraint (IC2) must hold: $R - 2r + \delta V \geq R - \delta S$. This reduces to an upper bound on the interest rate, which we call r_{IC2}:

$$r_{IC2}(S) \equiv \frac{\delta[pR + (1-\delta q)S]}{2-\delta q}.$$

In addition, a limited liability constraint must hold: $R \geq 2r$. We can ignore this without qualitatively affecting the results by the following parameter assumption:

Assumption 3 $\delta p \leq \frac{1}{2}$.

Now consider a lender offering a JL contract. If $r_{IC2}(S) \geq r_{IC1}$, the borrowers always guarantee one another when $r \leq r_{IC1}$, repaying with probability q, and always default if $r \geq r_{IC1}$. Suppose then that $r_{IC2}(S) < r_{IC1}$. If he offers $r \leq r_{IC2}(S)$, the borrowers will guarantee one another's loans and repay with probability q. If he sets $r \in (r_{IC2}(S), r_{IC1}]$, the borrowers will not be able to help one another with repayment, so will only repay in state (R, R), which occurs with probability p^2. Lastly, if he sets $r > r_{IC1}$, the borrowers always default. Clearly the latter cannot be an equilibrium. In addition, as we show when we discuss contract choice, a contract with $r > r_{IC2}(S)$ will always be dominated by an individual liability contract, so we ignore this possibility and focus on JL contracts under which borrowers repay with probability q.

Consider first a non-profit, altruistic lender offering joint liability loans. Since the repayment probability is q, the zero profit interest rate is $\hat{r} = \frac{\rho}{q}$. Plugging into (2.1), the equation for borrower welfare under the nonprofit is:

$$\hat{V}^{JL} = \frac{pR - \rho}{1 - \delta q}.$$

Note that $\hat{V}^{JL}$ does not depend on S.

Now consider a for-profit monopolist. The profit-maximizing interest rate binds the tighter of IC1 and IC2. We define the following threshold value of S:

$$\bar{S} \equiv pR.$$

For $S < \bar{S}$, IC2 is tighter than IC1, while for $S \geq \bar{S}$, IC1 is the tightest. Thus we obtain the monopolist's interest rate, $\tilde{r}^{JL}$, and borrower welfare, $\tilde{V}^{JL}(S)$ as follows.

$$\tilde{r}^{JL}(S) = \min\{r_{IC1}, r_{IC2}(S)\}$$
$$\tilde{V}^{JL}(S) = \frac{pR - q\min\{r_{IC1}, r_{IC2}(S)\}}{1-\delta q} \geq pR.$$

Note that for $S < \bar{S}$, $\tilde{r}$ is increasing in S, and therefore borrower welfare is decreasing in S, which we state as a proposition.

Proposition 2 *Under joint liability lending a monopolist for-profit lender exploits the borrowers' social capital by charging higher interest rates to borrowers with high social capital. Thus borrower welfare decreases in social capital.*

Another way of viewing this result is that the lender's motivation matters more as the amount of borrower social capital increases, as the difference between borrower welfare under the nonprofit and for-profit monopolist increases. We will return to this issue later on when we consider equilibrium contract choice.[16]

As discussed above, much of the microfinance literature has shown how different aspects of MFIs' lending methodologies can be thought of as leveraging social capital and local information among borrowers to address various asymmetric information or weak enforcement issues. Proposition 2 shows that this not need be a force for good from the perspective of borrowers: a monopolist may be able to use their social capital against them to extract more rents.

We have assumed that S is homogeneous and observable, so that the lender can choose the interest rate accordingly. Why can't the borrowers resist the lender's exploitation by refusing to use their ability to socially sanction one another? The problem is that threatening to do so is not credible. Conditional on the contract offered, the borrowers are better off using their ability to socially sanction to agree the most efficient repayment rule. Refusing to do so makes them less likely to be able to repay their loans and therefore worse off. The lender is a natural Stackelberg leader in this context - he simply commits to a single contract in period zero and the borrowers adjust accordingly. We consider the issue of heterogeneity of S as an extension below.

Individual Liability

Under individual liability the only incentive constraint is the one that ensures a borrower will repay her own loan, IC1. Provided IC1 holds ($r \leq \delta pR$), individual liability borrowers will repay whenever successful.[17] Then borrowers repay with probability p, so the nonprofit charges $\hat{r}^{IL} = \frac{\rho}{p}$, with borrower welfare $\hat{V}^{IL} = \frac{pR-\rho}{1-\delta p}$. The for-profit monopolist chooses r to bind IC1, giving the following interest rate

[16]The result also follows if we assume a lender that puts weight α on profits and $1-\alpha$ on borrower welfare (subject to a zero profit condition). By linearity in r of V and lender profits, there is an α threshold above which the lender behaves as a for-profit, and below which as a non-profit.

[17]The limited liability constraint, $R \geq r$ is implied by IC1.

and borrower welfare:

$$\tilde{r}^{IL} = r_{IC1}$$
$$\tilde{V}^{IL} = \frac{pR - pr_{IC1}}{1 - \delta p} = pR$$

It is clear that under the non-profit, borrower welfare under JL exceeds that under IL, due to the higher repayment probability. However, we also obtain a somewhat surprising result:

Proposition 3 *Despite the monopolist's exploitation of their social capital under joint liability, borrowers are still better off than under individual liability.*

Joint liability lending has received some negative press of late, in part due to perceptions of excessive peer pressure among borrowers. Our model captures this in one particular way: a lender with market power can exploit borrowers' ability to socially sanction one another to charge higher interest rates. It is thus surprising that the same lender would make borrowers worse off under individual liability.

The reason is straightforward. Under both contracts, the lender is constrained by IC1: it must be individually rational to repay a loan, at least when the partner is repaying. This constraint puts a lower bound on borrower welfare of pR. Under joint liability, for low levels of social capital the lender faces an additional constraint, IC2, that forces him to cut interest rates below individual liability levels in order to induce borrowers to guarantee one another's repayments. Furthermore, borrowers benefit directly from the higher repayment probability under JL.

One implication of this result is perhaps missed in the policy debates. Our model speaks to any lender with market power, using dynamic incentives to enforce repayment. Regulators should be alert to abuses by standard, IL-using lenders, who may or may not be formally registered as MFIs or even consider themselves to be MFIs.

Equilibrium contracts

So far we have analyzed IL and JL in isolation. Now we turn to the choice of contract in equilibrium. IL lending can earn non-negative profits as long as expected repayment at r_{IC1}, equal to pr_{IC1}, exceeds the opportunity cost of funds, ρ. To use IL lending as a benchmark, we retain this throughout as an assumption.

Assumption 4 $\delta p^2 R > \rho$.

JL can be used profitably provided that expected revenue when the tightest of IC1 and IC2 binds exceeds the opportunity cost of capital, i.e. $q \min\{r_{IC1}, r_{IC2}(S)\} \geq \rho$. This yields a threshold level of social capital, $\hat{S}$, above which JL lending can break even. Since borrowers are better off under JL, this is the switching point for the non-profit lender. We obtain:

$$\hat{S} \equiv \max\left\{0, \frac{(2-\delta q)\rho - (2-p)\delta p^2 R}{\delta q(1-\delta q)}\right\} < \bar{S}.$$

A simple sufficient condition that we shall make use of throughout for $\hat{S} = 0$ (i.e. JL is always profitable) is $p \leq \delta q$, or

$$1 + \delta p - 2\delta \leq 0. \tag{2.2}$$

Since the for-profit monopolist lender maximizes per-period profits, $\Pi = \pi r - \rho$, he chooses the contract offering the highest per-period revenue πr. Therefore he offers JL provided $q\tilde{r}^{JL}(S) \geq p\tilde{r}^{IL}$. This gives us a second threshold, above which JL is offered by the monopolist:

$$\tilde{S} \equiv \max\left\{0, \frac{p^2 R(p-\delta q)}{q(1-\delta q)}\right\}.$$

Condition (2.2), which was sufficient for JL to break even for all S, is *necessary and sufficient* for the monopolist to offer JL for all S. This is because of the following proposition.

Proposition 4 *$\tilde{S} \geq \hat{S}$, with the relation holding strictly if $p > \delta q$. Therefore, the for-profit monopolist lender offers JL over a (weakly) smaller range of S than the non-profit lender.*

This result is consistent with the current debate over the decline of joint liability lending in microfinance, which goes hand in hand with increasing commercialization of microfinance lending. The for-profit monopolist is less willing to offer joint liability loans than the non-profit, because when social capital is low the need to give borrowers incentives to help one another (IC2) constrains his rent extraction.

Intuitively, the non-profit is willing to offer JL whenever the borrowers have sufficient social capital for JL to break even, which requires $qr_{IC2}(S) \geq \rho$. The for-profit monopolist only offers JL when doing so is more profitable than IL, i.e. when $qr_{IC2}(S) \geq pr_{IC1}$, which is clearly a more restrictive condition. We will find that an analogous result carries over to competitive equilibrium in section 2.2.

This result is the classic rent-extraction/efficiency tradeoff with market power. We define an efficient contract as one that maximizes $V(S) + \frac{\pi r - \rho}{1-\delta\pi}$, i.e. discounting profits at the borrowers' discount rate. The following observation is then straightforward:

Observation 1 *Monopoly for-profit lending is inefficient when $S \in [\hat{S}, \tilde{S})$.*[18]

The use of group lending to leverage borrowers' social capital has been criticized for putting stress on borrowers and suggested as an important motivation for the

[18]The reader might be concerned that this result and Proposition 4 are artifacts of the assumption that the lender is myopic, only maximizing per-period profit. In fact this is not the case as we formally show in the appendix.

tendency of some lenders to move toward individual loans.[19] In our model, a monopolist using JL is bad for borrowers, but he is even worse with IL. The problem is market power, not the form of lending, and restricting contract choice without paying attention to this may be bad for both efficiency and equity.

In the simulation section we analyze the welfare implications of market power in detail. However the model allows us to easily make one policy-relevant remark on the effect of interest rate caps (a key component of some of the regulatory efforts, e.g., the Indian Microfinance Bill). The first-order effect is that the lender will be forced to cut his rates, essentially a transfer to the borrowers, increasing borrower welfare. There is a second-order effect on contract choice as well. If the lender is offering JL he will continue to do so. However, if he is offering IL but the cap lies below $\tilde{r}^{JL}(S)$, he will switch to JL, further improving borrower welfare. The reason is that the lender must now charge the same rate under IL and JL, but the JL repayment rate is higher. Thus in our framework, correctly calibrated interest rate caps can be an effective tool for borrower protection.

We have assumed that individual liability borrowers cannot side-contract among themselves to guarantee one another's repayments. However, this may be an overly strong assumption as they have an incentive to do so if this enables them to repay more frequently. We have a related paper on the effects of such side-contracting (de Quidt et al. (2013)) which we term "implicit joint liability" or IJ.

It turns out that in our simple framework, IJ plays no role. For IJ to enable borrowers to repay more frequently than JL, there need to be states of the world where one borrower would repay were she under individual liability, but defaults due to the joint liability burden (because she cannot afford to bail out her partner). This does not arise in equilibrium due to the simple production function we use; under JL either both borrowers succeed and repay, one succeeds and repays both loans, or both fail.[20] In de Quidt et al. (2013) we analyze this contractual form in detail in an environment where IJ can play a role.

2.1.4 Heterogeneity

The analysis so far assumes that social capital S is homogeneous and observable across borrowers. Suppose that this is not the case. To keep things simple, suppose there are two possible values of S. A fraction $\theta \in (0, 1)$ of borrowers have $S = 0$, and $1 - \theta$ have $S = S_h > 0$. The lender cannot observe social capital so must screen borrowers by offering an appropriate menu of contracts. We will first characterize the candidate pooling and separating equilibria, then solve for the equilibrium contract offer as a function of S_h. Also, to keep things brief, we only consider the monopolist for-profit lender.

[19]See, for example, Grameen II at `http://www.grameen.com/`.

[20]In addition, the incentive constraint for an IJ borrower to help her partner is tighter than the one for an JL borrower to do so (because there is no pressure from the lender encouraging her to do so). As a result, the lender always prefers to offer a contract conforming to the JL or IL contracts detailed above, rather than one that might induce the borrowers to engage in IJ.

In a pooling equilibrium, the monopolist offers a single interest rate r, and either IL or JL. There are three possible pooling equilibria. Equilibrium A uses IL and the interest rate will be $\tilde{r}^{IL} = r_{IC1}$. Equilibrium B uses JL with interest rate $\tilde{r}^{JL}(0)$, in which case all groups are able to guarantee one another's loans. Equilibrium C uses JL and interest rate $\tilde{r}^{JL}(S_h)$, in which case only the high S groups can do so, (in this case, the low S groups will only repay when both are successful, with probability p^2). We show that these are the only possible pooling equilibria in the appendix.

Now we turn to the separating equilibrium. We use the following notational convention. Where the interest rate corresponds to that from the basic model, we retain the same $\tilde{r}$ notation. Where the interest rate function differs, it is denoted by subscript "sep", as in $\tilde{r}_{sep}$. In a separating equilibrium the lender offers the following menu of contracts: one JL contract at interest rate $\tilde{r}^{JL}(S_h)$, and one IL contract at interest rate $\tilde{r}^{IL}_{sep}(S_h)$. Note that the IL interest rate depends upon the social capital of the high S types. High S borrowers take the JL contract and low S borrowers take the IL contract. When $S_h \geq \bar{S}$, the lender charges the same interest rate under both contracts, namely $\tilde{r}^{JL}(S_h) = \tilde{r}^{IL}_{sep}(S_h) = \tilde{r}^{IL}$, i.e. all borrowers are charged the "maximum" interest rate $\tilde{r}^{IL} = r_{IC1}$. When $S_h < \bar{S}$, we find that $\tilde{r}^{JL}(S_h) < \tilde{r}^{IL}_{sep}(S_h) < \tilde{r}^{IL}$. The lender cannot charge the maximum interest rate to the JL borrowers any more as they do not have sufficient social capital to guarantee one another. In addition, the truth-telling constraint that induces low-S borrowers to choose IL rather than JL constrains the lender from charging the maximum interest rate to IL borrowers either.

In the appendix we derive the $r^{IL}_{sep}(S_h)$ function, obtaining the following: $\tilde{r}^{IL}_{sep}(S_h) \equiv \phi\tilde{r}^{IL} + (1-\phi)\tilde{r}^{JL}(S_h)$, where $\phi \equiv \frac{1-p}{1-\delta p^2} < 1$. Note that $r^{IL}_{sep}(S_h)$ is increasing in S_h. The higher is S_h, the higher the interest rate the lender can charge under JL and thus the higher he can charge under IL as well. This gives us an observation analogous to the "exploitation" results earlier:

Observation 2 *In a separating equilibrium with heterogeneous social capital, the interest rate faced by the individual liability borrowers (who have low social capital) is increasing in the social capital of the joint liability borrowers (who have high social capital).*

The main addition to the benchmark model is that now more social capital among one type of borrowers has spillover effects on the other type, enabling the lender to exploit them more as well.

Now we derive the equilibrium contract. As before, the lender maximizes per-period profits, which is equivalent to choosing the contract or menu that yields the highest per-period revenue. When S_h is high, the lender will have a strong incentive to separate borrowers by type, so the separating equilibrium prevails. When S_h is low, we need to check which of the pooling equilibria (A, B or C) will be chosen.

We can immediately rule out pooling equilibrium C (JL with interest rate $\tilde{r}^{JL}(S_h)$). This yields revenue of $(\theta p^2 + (1-\theta)q)\tilde{r}^{JL}(S_h)$, while the separating menu yields strictly larger revenue of $\theta p \tilde{r}^{IL}_{sep}(S_h) + (1-\theta)q\tilde{r}^{JL}(S_h)$, using the fact that $\tilde{r}^{IL}_{sep}(S_h) >$

$\tilde{r}^{JL}(S_h)$. Intuitively, pooling equilibrium C is unattractive as it leads the low S borrowers to default very frequently.

Now note that revenue does not depend on S_h in either of the remaining two pooling equilibria (A or B), so these can be ranked based on model parameters only. Determining which the lender prefers (IL at $\tilde{r}^{IL}$, or JL at $\tilde{r}^{JL}(0)$) is equivalent to determining whether $\tilde{S} \gtrless 0$, which reduces to our usual condition (2.2) or $p \gtrless \delta q$. Essentially, when $p \geq \delta q$, JL is always attractive so the lender will offer JL in the pooling equilibrium (B), otherwise he offers IL (A). For brevity, we analyze the $p \geq \delta q$ case here, the other is similar and is discussed in the appendix.[21]

If $p > \delta q$, then the IL pooling equilibrium (A) is more profitable than the JL one (B). It is easy to check that for any θ, there exists a threshold $S_h^{p>\delta q}$ for S_h above which the lender offers the separating contracts and below which he offers the pooling contract.[22] We have the following result:

Proposition 5 *When $p > \delta q$, for $S_h < S_h^{p>\delta q}$, the lender offers IL at interest rate $\tilde{r}^{IL}$. For $S_h \geq S_h^{p>\delta q}$ he offers IL at $\tilde{r}^{IL}_{sep}(S_h)$ and JL at $\tilde{r}^{JL}(S_h)$, low S borrowers take the IL contract and high S borrowers take JL. Welfare of both types of borrowers increases discontinuously at $S_h^{p>\delta q}$, and decreases in S thereafter.*

Low S types initially have IL contracts and utility $V = pR$ as usual. When the lender switches to the separating contract, they continue to receive IL loans but their interest rate decreases from $\tilde{r}^{IL}$ to $\tilde{r}^{IL}_{sep}(S_h)$, so they are discontinuously better off. Then, as noted in observation 2, this interest rate increases in S_h thereafter, reducing their welfare. High S types exactly mirror the borrowers in the homogeneous model: they receive IL up to $S^{p>\delta q}$ (as opposed to $\tilde{S}$), then switch to JL with a lower interest rate, making them better off, but this interest rate subsequently increases in S_h.

Proposition 5 is closely analogous to our earlier results. Enough social capital to induce the lender to offer JL is beneficial: the high S borrowers receive a more efficient JL contract at a lower interest rate, while the low S borrowers continue to receive IL but also benefit from a lower interest rate. However, above $S_h^{p>\delta q}$, Observation 2 kicks in and more social capital makes borrowers worse off.

We see the results in this section as broadly supporting the main conclusion that understanding market structure is critical for how we think about the role of social capital in influencing borrower welfare. With heterogeneity, there are also spillovers:

[21]The key qualitative difference between the two cases is that when $p \geq \delta q$, the lender offers IL when S_h is low, and welfare discontinuously increases when he switches to the separating equilibrium. When $p < \delta q$ he offers the most favorable JL contract when S_h is low, and therefore welfare for both types discontinuously *decreases* when he switches to the separating contract, and then further decreases in S_h thereafter, by Observation 2. Furthermore, when $p < \delta q$ the lender may always prefer the pooling equilibrium.

[22]For $S_h = 0$ and $p > \delta q$ the lender earns earns strictly lower per-borrower revenue from each type in the separating equilibrium than under the IL pooling contract. For $S_h \geq \bar{S}$, $\tilde{r}^{IL}_{sep}(S_h) = \tilde{r}^{JL}(S_h) = \tilde{r}^{IL} = \delta pR$, so the interest rate is the same under both pooling and separating equilibrium, but the repayment probability is higher under the separating equilibrium, (formally revenue under the separating contract is $[\theta p + (1-\theta)q]\delta pR$ which is superior to pooling IL). The existence of the threshold $S_h^{p>\delta q}$ then follows by continuity.

the more social capital held by the high types, the higher the interest rate faced by the low types.

2.2 Competition

The previous section showed how relaxing the assumption of altruistic non-profit lending affects borrower welfare. A for-profit lender with market power charges higher interest rates, inefficiently under-uses joint liability and exploits the social capital of joint liability borrowers by charging higher interest rate to those with more social capital (although some social capital may be beneficial if it leads the lender to use JL). In this section we explore to what extent competition can mitigate these problems.[23]

It turns out that competition is not guaranteed to deliver an improvement in borrower welfare over monopoly lending. Moreover, competitive lenders also inefficiently underuse joint liability, although to a lesser extent than a monopoly for-profit lender. However, competition does eliminate the exploitation of social capital: in a competitive market more social capital unambiguously improves borrower welfare.

Firstly, suppose that competitive lenders share information on defaulting borrowers, for example through a credit bureau, and agree not to lend to any borrower with a bad history. In that case, competition is identical to our nonprofit lender: free entry ensures that lenders break even and all borrowers can access credit. The problem becomes more interesting when we allow for imperfect information sharing among lenders. We do this in the simplest possible way, in a setup analogous to Shapiro and Stiglitz (1984). Entry by competitors imposes an enforcement externality on existing lenders by tightening the borrowers' repayment constraints. This happens because the probability of a defaulting borrower being able to obtain a loan from another lender increases the more lenders operate in the market. The higher is this probability, the weaker the incentive effects of the existing lender's termination threat.

Our framework gives a simple and tractable model of competitive equilibrium in microfinance that allows us to quantify the potential welfare gains from competition. Although competitive lenders earn zero profits, there is always credit rationing in equilibrium to preserve the repayment incentives of existing borrowers. The trade-off between higher interest rates and credit rationing means it is theoretically ambiguous whether borrower welfare will be higher under competition than under the for-profit monopolist. However, in a competitive market social capital is unambiguously beneficial to borrowers. Higher social capital slackens the lenders' enforcement constraints, leading to less credit rationing.

[23] Recent work on competition in microfinance has studied issues of adverse selection and multiple borrowing. For example, in McIntosh and Wydick (2005) competition can can be harmful by preventing lenders from cross subsidizing their bad borrowers with profits on good borrowers. In contrast, our framework directs us to focus on enforcement problems created by competition, as in Hoff and Stiglitz (1997). Entry by competitors makes it more difficult for a lender to incentivize his borrowers to repay their loans, since they can more easily obtain credit elsewhere.

To keep the model as simple as possible, we assume a very large number of lending "branches" that may belong to the same or different lenders, with no information sharing between branches. Therefore defaulters can go on to borrow at another branch, the source of the enforcement problem.[24] The assumption of atomistic branches means that we do not need to track the credit histories of individual borrower-lender pairs. Each branch is capable of serving two IL borrowers or one JL pair. The population mass of branches is l, while we normalize the population of borrower *pairs* to 1.[25] If $l < 1$ there will be rationing in the credit market: not all borrowers can obtain a loan in a given period. If $l > 1$ then some branches will have excess capacity.

Every borrower has a large number of potential partners, so even after being socially sanctioned a borrower is assumed to be able to form a new group with social capital S.[26] At the start of a period, borrowers will be either "matched", in an existing relationship with a lender, or "unmatched", waiting to find a lender. Since branches are atomistic the probability of a borrower rematching to a branch at which she previously defaulted is zero, so an unmatched borrower's matching probability does not depend on her history. Unmatched branches post a contract offer and are randomly matched to borrowers until all borrowers are matched or there are no more unmatched lenders. Each period, loans are made according to the contracts agreed, the repayment game is played, and any defaulters have their contracts terminated, rejoining the pool of unmatched borrowers. We note the following:

Observation 3 *There is credit rationing in equilibrium, i.e.* $l < 1$.

If this were not the case, there would be no dynamic repayment incentives, so all borrowers would default. Although formally trivial this result has an interesting implication. A common response to concern about commercialization is that access to funds from profit-motivated investors will enable much greater outreach.[27] By assuming all lenders face the same, constant opportunity cost of capital, we shut down that channel. An opposite tension emerges: the enforcement externalities in competitive credit provision lead to *lower* outreach than either under a single non-profit or for-profit lender.

Since there is rationing, every branch will be able to attract borrowers every period. Therefore each branch can act as a local monopolist, offering the more profitable of IL and JL at the highest r that satisfies the (modified) IC1 and IC2.[28] In

[24]Formally, this is identical to assuming a single lender who forgets the credit history of borrowers in the pool of potential customers. We do not consider other, more nuanced approaches to information sharing, which are analyzed in the growing literature on credit bureaus in microfinance, see for instance De Janvry et al. (2010) and references therein.

[25]It is conceptually slightly more convenient to think in terms of pairs.

[26]This assumption means that we do not need to track the social capital level of individual borrowers.

[27]See for example the letter from Michael Schlein, chief executive of Accion to The New York Times, January 23rd, 2011, available at `http://www.nytimes.com/2011/01/24/opinion/lweb24micro.html`.

[28]Instant costless replacement of defaulters means that even patient lenders would simply maximize per-period profits.

equilibrium, entry occurs until lenders earn zero profits, at the intersection of the zero-profit interest rate and the tightest repayment constraint. We assume that if both IL and JL break even, lenders offer the borrowers' preferred contract, JL, which rules out equilibria in which both IL and JL are offered.[29]

Suppose that proportion η branches offer IL loans, and $1-\eta$ offer JL. Therefore there are ηl IL branches[30] Each period, fraction $(1-p)$ of the IL borrowers default, creating vacancies in their respective branches. This is equivalent to there being $(1-p)\eta l$ vacant IL branches at the beginning of the next period (although note that in general there will be zero, one or two vacancies at a given branch). Similarly, there are $(1-\eta)l$ JL branches.[31] Of these, fraction $(1-q)$ of the borrower pairs will jointly default each period, leaving $(1-q)(1-\eta)l$ vacant JL branches at the beginning of the next period. The total proportion of unmatched borrower *pairs* at the beginning of a period is therefore $P \equiv (1-p)\eta l + (1-q)(1-\eta)l + (1-l)$, so an unmatched borrower matches with an IL branch with probability $\frac{(1-p)\eta l}{P}$, and a JL branch with probability $\frac{(1-q)(1-\eta)l}{P}$. In competitive equilibrium, $\tilde{\tilde{r}}^{IL} = \frac{\rho}{p}$ and $\tilde{\tilde{r}}^{JL} = \frac{\rho}{q}$. We denote the utility of an unmatched borrower by U. We obtain:

$$
\begin{aligned}
\tilde{\tilde{V}}^{IL} &= \frac{pR-\rho}{1-\delta p} + \frac{\delta(1-p)U}{1-\delta p} \\
\tilde{\tilde{V}}^{JL} &= \frac{pR-\rho}{1-\delta q} + \frac{\delta(1-q)U}{1-\delta q} \\
U &= \frac{(1-p)\eta l}{P}\tilde{\tilde{V}}^{IL} + \frac{(1-q)(1-\eta)l}{P}\tilde{\tilde{V}}^{JL} + \frac{\delta(1-l)}{P}U \\
&= \chi(l,\eta)\frac{pR-\rho}{1-\delta}.
\end{aligned}
$$

The function χ is defined as:[32]

$$
\chi(l,\eta) \equiv \frac{(1-p)(1-\delta q)\eta l + (1-q)(1-\delta p)(1-\eta)l}{(1-\delta p)(1-\delta q)(1-l) + (1-p)(1-\delta q)\eta l + (1-q)(1-\delta p)(1-\eta)l}
$$

$$
\chi(l,\eta) \in [0,1],\ \chi_l \geq 0,\ \chi_\eta \geq 0.
$$

Total welfare from microfinance is the combined welfare of matched and unmatched borrowers, equal to:

$$
\begin{aligned}
Z &\equiv \eta l \tilde{\tilde{V}}^{IL} + (1-\eta)l\tilde{\tilde{V}}^{JL} + (1-l)U \\
&= \left[\frac{\chi(l,\eta)}{1-\delta} + l(1-\chi(l,\eta))\left(\frac{\eta}{1-\delta p} + \frac{1-\eta}{1-\delta q}\right)\right](pR-\rho)
\end{aligned}
$$

[29] There is a single value of S, termed $\tilde{\tilde{S}}$, at which mixed equilibria could occur so this assumption is innocuous.

[30] i.e. $2\eta l$ IL borrowers, since we measure in terms of pairs.

[31] i.e. $2(1-\eta)l$ JL borrowers.

[32] $\chi_l \geq 0$ and $\chi_\eta \geq 0$ follow from the fact that greater scale or a higher proportion of (more frequently defaulting) IL borrowers increase the matching probability and thus welfare of an unmatched borrower. It is straightforward to check that borrower welfare is (weakly) higher under JL for all χ. Also note that as $\chi \to 1$, $\tilde{\tilde{V}}$ and U approach $\frac{pR-\rho}{1-\delta}$, which is the first-best welfare.

The modified framework implies that each lender will face a new IC1 (and IC2 under JL). The constraints now reflect the fact that the borrowers' outside option upon default is improved (they become unmatched and may re-borrow in future), and so are tighter than before. As χ, and thus U increases, the tightest of these two constraints becomes tighter. This is the competition effect that constrains existing lenders' interest rates. We derive the constraints in Appendix 2.A.5.

2.2.1 Equilibrium

In equilibrium, it must not be profitable to open a new branch offering either IL or JL. Two key thresholds in the following analysis are $\tilde{\tilde{S}} \equiv \frac{p-\delta q}{\delta q(1-\delta q)}\rho$, and $\bar{\bar{S}} \equiv \frac{\rho}{\delta q}$. The former is the analog of $\tilde{S}$, representing the level of social capital at which the competitive market switches from IL to JL lending, and the latter is the analog of $\bar{S}$, the level of social capital at which IC1 binds under JL (as opposed to IC2). Note also that $\bar{\bar{S}} > \tilde{\tilde{S}}$.

Proposition 6 *If $\tilde{\tilde{S}} \leq 0$, the competitive equilibrium is JL-only lending, with market scale strictly increasing in S for $S < \bar{\bar{S}}$, and equal to a constant, $\bar{\bar{l}}$ for $S \geq \bar{\bar{S}}$. If $\tilde{\tilde{S}} > 0$, the equilibrium for $S < \tilde{\tilde{S}}$ is IL-only lending at fixed scale $\underline{l}$. At $\tilde{\tilde{S}}$, all lending switches to JL at scale $\tilde{\tilde{l}} > \underline{l}$, then increases continuously in S to $\bar{\bar{l}}$, at $\bar{\bar{S}}$. Welfare, Z, is strictly increasing in scale, l, and therefore weakly increasing in S.*

The proof can be found in Appendix 2.A.5. The intuition of the proof is simple. For a given contract type, lender entry occurs until the tightest repayment constraint (either IC1 or IC2) binds. For low levels of S, IC2 under JL is tight so lenders prefer IL. As S increases, the JL IC2 is relaxed to the point that all lending switches to JL. Thereafter, JL is offered and scale increases in S until IC1 binds. Aggregate welfare from microfinance, Z, is improved as S increases because this enables a relaxation of credit rationing.

Comparing the key S thresholds, we have the following result:

Proposition 7 *$\hat{S} \leq \tilde{\tilde{S}} \leq \tilde{S}$, with both inequalities strict when $p > \delta q$. Therefore the competitive market is weakly less likely to offer JL than the non-profit, but more likely than the for-profit monopolist.*

This result is interesting because it is consistent with the (perceived) trend away from JL. Our model predicts that a commercialized microfinance market will exhibit this tendency, whether competition is weak or strong. A common story that is told to explain this trend is that lenders are responding to borrowers' preference for more flexible IL loans. Our model rules this channel out, but yields another: for-profit lenders (competitive or otherwise) benefit from the slacker repayment constraints under IL.

We have already outlined the intuition for why $\hat{S} \leq \tilde{S}$ (see Proposition 4). To see why $\tilde{\tilde{S}} \geq \hat{S}$ note that JL relies on two forms of enforcement. The first is the lender's threat of terminating lending, and the second is the partner's threat to use social sanctions. We can write a generalized form of IC2 as $\delta((V-U)+S) \geq 2r$, noting that with a single lender (monopolist for-profit or non-profit), $U=0$, while in competitive equilibrium $U \geq 0$ is an increasing function of l, the market scale. The threat of termination (loss of $\delta(V-U)$) is a substitute for social capital. Since the termination threat is weaker under competition than with a single lender, more social capital will be required for lenders to be able to break even in competitive equilibrium than for the non-profit to break even, and hence the non-profit switches to JL for a lower level of S than the competitive market.

Turning to the comparison between competition and monopoly, we note that in both cases, each lender chooses which contract to offer by comparing per-period revenues between the two options (in equilibrium the competitive lender will break even). The monopolist lender's decision problem is nested in the competitive framework by setting $U=0$, and an increase in the amount of competition through entry by competitors corresponds to an increase in U. $\tilde{\tilde{S}} \leq \tilde{S}$ implies that for a given level of S, JL is relatively more attractive to a competitive lender than it is to the monopolist lender. An increase in the amount of competition weakens the termination threat under both IL and JL, thus leading to lower interest rates under both. This effect is stronger under IL than JL, because the incentive constraint for IL concerns the repayment of only one loan, so a small decrease in $\delta(V^{IL}-U)$ maps one-to one into a decrease in r^{IL}. Meanwhile, under JL, IC2 concerns the repayment of two loans, so a small decrease in $\delta(V^{IL}-U)$ leads to a one-to-one decrease in $2r^{JL}$, i.e. r^{JL} only decreases by half that amount. As a result, for a given level of S, JL becomes relatively more attractive as competition increases, which implies that JL will be used for lower levels of S in competitive equilibrium than under monopoly.[33]

Note that Proposition 7 is not quite enough to argue that the lower use of JL by competitive lenders is inefficient. The structure of the market implies that dynamic incentives are somewhat different under competition than in the core model - a defaulting borrower in competition can expect to borrow again in future. In appendix 2.A.7 we extend the basic model to allow the nonprofit to offer a "stochastic renewal" contract that mimics the competitive market, and show that the competitive market does indeed under-use joint liability. Note also that in the neighborhood of $\tilde{\tilde{S}}$, aggregate welfare in a JL-only equilibrium is strictly higher than in an IL-only equilibrium.

2.2.2 Comparing market structures

Lastly, we turn to the question of whether competition is necessarily beneficial for borrower welfare in the presence of weak information sharing as in the framework

[33] A more formal way to see this is to consider the point at which the lender is indifferent between IL and JL. Using the expressions for IC1 and IC2, this occurs when $p\delta(V^{IL}-U) = \frac{1}{2}q\delta((V^{JL}-U)+S)$. The left hand side decreases in U at rate δp, while the right hand side decreases at rate $\frac{\delta q}{2}$. Since $\frac{q}{2} = \frac{p(2-p)}{2} < p$, revenues under JL are less responsive to changes in U than those under IL, and therefore the level of S required to maintain indifference is decreasing in U.

outlined here:

The following proposition shows that the ranking (by borrower welfare) of the market structures considered in this paper is ambiguous. It is straightforward to see the following result:

Proposition 8 *The ranking of total borrower welfare under competition and monopoly for-profit lending is ambiguous.*

Under the monopolist, all borrowers receive loans in the first period so total welfare is equal to $\tilde{V} \geq pR$. Under the competitive equilibrium, total welfare is W, which depends on the degree of credit rationing in equilibrium. When credit rationing is low (l is close to 1, for example because ρ is small and S is large), Z approaches the first-best welfare $\frac{pR-\rho}{1-\delta}$, and so dominates the monopolist. Meanwhile when credit rationing is high (l is close to zero, for example because ρ is large and S is small), Z approaches zero and is dominated by monopoly lending.

Proposition 8 follows from the observation that when market scale under competition is small, the cost of credit rationing outweighs the benefits of lower interest rates and the potential to borrow again after defaulting. By eliminating the enforcement externality generated by competition, the monopolist solves the credit rationing problem. When scale is large, borrowers are essentially able to borrow every period, so there is no longer the inefficiency generated by dynamic incentives.

Proposition 8 reflects the genuine concern about externalities in uncoordinated competition.[34] A key purpose of the simulations performed in the next section is to understand the scope of this ambiguity - we will rank borrower welfare under our key market structures when we under a reasonable parameterization of the model.

2.2.3 Credit Bureaus

It is worth briefly considering how our results map into the growing literature on credit bureaus in microfinance. Much of the existing work primarily concerns adverse selection issues - credit bureaus help banks to screen out bad types. They also play a role in enforcement frameworks such as ours: credit bureaus enable lenders to damage a borrower's reputation following default, reducing her access to credit from other lenders and thus increasing repayment incentives.

[34]It is also possible in our simple framework that competition might dominate the non-profit, arises because of the assumption that the non-profit must use strict dynamic incentives, while competition mimics a contract with probabilistic termination on default, akin to Bhole and Ogden (2010). We discuss relaxing the assumption of strict dynamic incentives in Appendix 2.A.7. We assume strict dynamic incentives in the main analysis because this is what lenders seem to use in practice and because the analysis is much simpler. However, if the non-profit chose to use stochastic renewal, he could achieve at least the same welfare as competition. For example by choosing the appropriate renewal probability he can mimic the contract faced by the matched borrowers under competition. However, he can do better by offering this contract to all borrowers. Moreover, sometimes the competitive market offers IL when JL would be better for the borrowers. In Appendix 2.A.7 we analyze a relaxed dynamic incentive, namely, renewing the group's contracts with certainty following repayment and with probability $\lambda \in [0, 1]$ following default. We find that the monopolist and competitive market always set $\lambda = 0$, while the nonprofit does use stochastic renewal, achieving higher borrower welfare than the competitive market.

Our single-lender and competitive frameworks represent two extremes of possible forms of information sharing. The competitive environment is one without any form of credit bureau - borrowers' reputations are not harmed by default so the only cost of default is that the borrower might be credit rationed for a few periods until a new lender has spare capacity.

The single-lender environment is analogous to one with a strong credit bureau: defaulters cannot borrow again. When the lender is a monopolist, this enables them to extract rents using the strong dynamic incentives invoked. The non-profit equilibrium can be thought of as analogous to a competitive equilibrium with full information sharing, where all lenders commit not to lend to defaulters.

If we take the latter case seriously, our simulations could allow us to calculate the welfare effect of introducing a credit bureau to a competitive market. However, it is worth noting that this is not really an equilibrium. Equilibrium requires lenders to prefer not to lend to former defaulters, i.e. for the commitment to be self-enforcing. This is not the case - former defaulters are no different from non-defaulters (in equilibrium, they are just unlucky) and are no less attractive to lenders. Formally modeling information sharing in this environment requires a number of modifications to the framework along the lines of Greif (1993), which are beyond the scope of this paper.

To summarize, the single non-profit lender is equivalent to perfect competition with full information sharing whereby defaulters can never borrow again. The for-profit monopolist is equivalent to a single (myopic) profit-maximizing lender or cartel with full information sharing. Lastly the competition model represents competition with no information sharing. Although reality will of course be a more complex mixture of these cases, interestingly they are not strictly ordered in a welfare sense. We now simulate the model for real-world parameters to further explore the welfare effects of changing market structure.

2.3 Simulation

In this section we carry out a simple simulation exercise to get a sense of the order of magnitude of the effects analyzed in the theoretical analysis. We draw on plausible values for the key parameters of the model, mostly estimated using 2009 data from MIXMarket.org, an NGO that collects, validates and publishes financial performance data of MFIs around the world.

Throughout the analysis the numeraire is the loan size, so borrower welfare and social capital are measured in multiples of this. Loan sizes of course vary widely but in South Asia a typical microfinance loan is of the order of $100-200. The full sample results give a good picture of the basic empirical predictions of the model. The non-profit always offers JL, at a net interest rate of 15.9%, while the for-profit monopolist's interest rate is 38.2% when he offers IL, which occurs for social capital worth less than 0.15 in present discounted value (or 15% of the loan size). When he switches to JL, the interest rate falls to 34.5%, but this difference is eroded as social capital increases, until eventually IC1 binds at social capital worth 0.40 and IL and JL interest rates equalize. Borrower utility from access to microfinance, V, is 2.76 with a non-profit lender, while the maximum value with a monopolist (at the point of switching from IL to JL) is only 1.80, reducing to 1.60 under IL or when S is large.

Under competition, IL is offered for social capital worth less than 0.13, and JL thereafter. Market scale varies from 67% of borrowers served under IL, to 78% under JL when S is sufficiently large (note that these predictions should be thought of as local rather than national or regional market penetration). IL is offered for social capital worth less than 0.13 and aggregate welfare from microfinance, Z (which includes matched and unmatched borrowers) is 2.49. This is higher than welfare under a monopolist, so the welfare effect of credit rationing is clearly not too severe. For social capital worth more than 0.13, JL is offered at increasing market scale, with welfare increasing to a maximum of 2.90 for $S \geq 0.33$, higher even than welfare under the non-profit.

The discussion proceeds as follows. First we modify the model to allow for larger group sizes, and also discuss the possibility that the LLC may bind (which we ruled out for the theoretical discussion for simplicity). Second, we describe the estimation of the model parameters. Third, we discuss the results using the full global sample of MFIs. Fourth, we perform some sensitivity checks and finally we discuss the results when parameters are estimated at the regional level.

2.3.1 Group size and limited liability condition

We make one modification to the framework, modeling larger groups of size five instead of two.[35] Theoretically, small groups disadvantage JL, since they require very large "guarantee payments" and hence a very tight LLC. For simplicity, we retain the notion of S from the benchmark model - a deviating member loses social capital with the other members worth a total of S. In addition to this, we need to allow for the possibility that the LLC might be tighter than IC1. This is straightforward to implement in the simulations.

With a group of size n, borrowers will agree to guarantee repayment provided at least some number, m, of members are successful, defining a guarantee payment of $\frac{nr}{m}$ per successful member, so for example if $n = 5$ and $m = 4$, each successful member would repay 1.25r when one member fails. It is easy to see that the group size does not affect IC1, that is, $\delta pR \geq r$ is still necessary. There will be a different IC2 for each value of m, corresponding to the payment that must be made when only m members are successful. In equilibrium, borrowers will repay for every $m \geq m^*$, where m^* is the smallest m such that repayment is incentive compatible. By reducing the interest rate the lender can increase the number of states of the world in which repayment takes place, generating a (binomial) repayment probability of $\pi(n, m, p)$.

[35] Five was the group size first used by Grameen Bank and by other prominent MFIs. An unexplored extension would be to allow the lender to optimally choose the group size.

Table 2.1: Summary Statistics and Parameter Estimates

	MFIs	Loans (m)	% Full Sample	IL share (num)[a]	IL share (value)	Interest rate	p	R	ρ	δ
Full Sample	715	65.217	100.0%	46.0%	81.9%	1.206	0.921	1.737	1.098	0.864
Central America	60	1.671	2.6%	93.8%	98.8%	1.190	0.881	1.816	1.112	0.860
South America	133	6.884	10.6%	97.7%	99.3%	1.237	0.928	1.724	1.102	0.874
Eastern Africa	20	2.439	3.7%	38.7%	70.4%	1.152	0.831	1.925	1.115	0.848
Northern Africa	20	1.735	2.7%	37.5%	59.2%	1.227	0.984	1.626	1.115	0.871
Western Africa	48	1.184	1.8%	60.5%	89.2%	1.306	0.882	1.814	1.173	0.896
South Asia	133	44.067	67.6%	34.8%	33.3%	1.180	0.926	1.728	1.083	0.856
South East Asia	85	4.296	6.6%	45.7%	68.3%	1.389	0.988	1.619	1.164	0.922
South West Asia	61	0.865	1.3%	75.0%	93.8%	1.272	0.967	1.655	1.106	0.885

Notes: IL shares are the fraction of the total number or total value of loans reported as IL loans. The interest rate column reports the weighted mean risk-adjusted portfolio yield, i.e. $r_i = \frac{RealPortfolioYield}{1-PAR}$. [a] For 10 observations we use the share by value to compute the overall figures in this column.

We discuss the derivation of the constraints in detail in Appendix 2.A.8.

2.3.2 Data and Parameter values

The model's key parameters are R, p, ρ and δ. The numeraire throughout is the loan size, assumed to be identical between IL and JL, and the loan term is assumed to be 12 months. Since social capital and market structure are our key independent variables, we perform the various exercises for the non-profit, for-profit monopolist and competition cases while varying the level of S, computing welfare, interest rates and market scale. Changes in contract choice at the various thresholds of S lead to discontinuous jumps in the value functions, interest rate and market scale. Throughout we use weighted means or regression techniques, weighting by the number of loans outstanding (each observation is a single MFI). We use these weights since our unit of analysis is the borrower, thus, assuming one loan per borrower. We work with 2009 data from 715 institutions from the MIX to estimate the parameters and perform extensive sensitivity checks. Details of the construction of the dataset can be found in Appendix 2.C.

Table 2.1 summarizes parameters in the full sample and across the regions. In addition we report the number of MFIs, number of loans outstanding (million) and the weighted mean interest estimate that was used to calibrate δ. We later compare these interest rate estimates with the non-profit rates predicted by the model. One immediate observation is the extent to which South Asia dominates the sample, comprising 68% of the full sample by number of loans (India comprises 41% of the full sample, and Bangladesh 22%). This observation partly motivates the decision to repeat the exercise by region.

Estimating p We estimate p using cross-sectional data from the MIX on Portfolio At Risk (PAR), the proportion of an MFI's portfolio more than 30 days overdue, which we use as a proxy for the unobserved default probability. This is not an ideal measure for two reasons. Firstly, PAR probably exaggerates final loan losses, as some overdue loans will be recovered. However, MFIs' portfolios are typically growing rapidly (see the discussion of the estimation of ρ below). If loans become delinquent late in the cycle, they will be drowned out by new lending, understating the fraction of a cohort that will subsequently default.

We also need to be mindful of the lending methodology, since the model predicts that JL borrowers will repay more frequently than IL borrowers. Our data allow us to separate the portfolio by lending methodology. Let θ denote the IL fraction of the lender's portfolio. Then we have $1 - PAR = \theta p + (1-\theta)\pi(n, m, p)$. We estimate this equation by Nonlinear Least Squares (NLS), obtaining full sample estimates of $p = 0.921$ and $m = 3$.[36] Since we do not observe detailed contractual information we treat all "solidarity group" lending as JL, and all individual lending as IL, see the Appendix for more details.

Estimating ρ We estimate ρ using data from the MIX on administrative (x_a) and financial expenses (x_f). To obtain the cost per dollar lent, we need to divide expenses by the total disbursals of that MFI during the year. Since MIX does not report data on disbursals, we hand-collected disbursal data from annual reports of the largest MFIs listed on MIX, for which the (weighted) mean ratio of disbursals to year-end portfolio was 1.91.[37] Therefore, for MFI i we estimate $\rho_i = 1 + \frac{x_{a,i}+x_{f,i}}{GrossLoanPortfolio*1.91}$. Our full sample estimate is $\rho = 1.098$. Although we calibrate pR and δ using data in real terms, we do not deflate our estimate of ρ since we do not know the timing of expenses throughout the loan term or year.

Estimating δ Since the lender's only instrument to enforce repayment is the use of dynamic incentives, the borrowers' time preferences play an important role in the analysis. Unfortunately, it is not obvious what value for δ to use. Empirical estimates in both developed and developing countries vary widely, and there is little consensus on how best to estimate this parameter (see for example Frederick et al. (2002)). Due to this uncertainty, we calibrate δ as the mid-point of two bounds. We take the upper bound for all regions to be $\delta^U = 0.975$, since in a long-run equilibrium with functioning capital markets $\delta = \frac{1}{1+r^{rf}}$, where r^{rf} is the risk-free real rate of return which we take to be 2.5%, the mean real return on US 10-year sovereign bonds in 1962-2012. For the lower bound we use the model's prediction that $r \leq \delta pR$ by IC1. We estimate the real interest rate charged by MFIs in the MIX data as $r_i = \frac{RealPortfolioYield}{1-PAR}$. To avoid sensitivity to outliers, we then calibrate $\delta^L = \frac{\bar{r}}{pR}$, where $\bar{r}$ is the weighted mean interest rate. Using our calibrated value for pR of 1.6 (see below), we obtain $\delta^L = 0.753$ in the full sample. The midpoint of δ^U and δ^L gives us $\delta = 0.864$.

[36]However, if p varies at the MFI level, contract choice (reflected in θ) may also be a function of p, so the restriction that the underlying p for IL and JL is the same will be violated. An OLS regression of $1 - PAR$ on θ and a constant estimates a separate probability for IL and JL, but ignores the nonlinearity in the the model and sometimes yields estimates of π that exceed one. In practice, both approaches give very similar estimates, so we focus on the NLS results.

[37]We looked up annual reports, ratings or MFI websites for the 50 largest MFIs by number of outstanding loans. For 26 we were able to obtain data, comprising 60% of the loans in our sample.

Estimating R There are few empirical studies that exploit exogenous variation in microentrepreneurs' capital stocks to estimate the returns to capital. We draw our full sample value for the returns to capital from De Mel et al. (2008). They randomly allocate capital shocks to Sri Lankan micro enterprises, and their study suggests annual expected real returns to capital of around 60%.[38] Since expected returns in our model are pR, we use $pR = 1.6$, dividing by our estimate of p to obtain $R = 1.737$.

2.3.3 Results

Figure 2.3 graphically presents the results for the baseline simulation, which were discussed in detail in the introduction to this section. The values for the full sample and all regions of the S thresholds, corresponding interest rates, market scale and welfare are reported in Tables 2.2 and 2.3. Table 2.2 also reports the contracts used by each type of lender, showing that the non-profit exclusively offers JL in the majority of cases, while the monopolist and competitive market typically offer IL for low S and JL for high S, although sometimes only IL is offered, corresponding to cases when the JL LLC is tight.

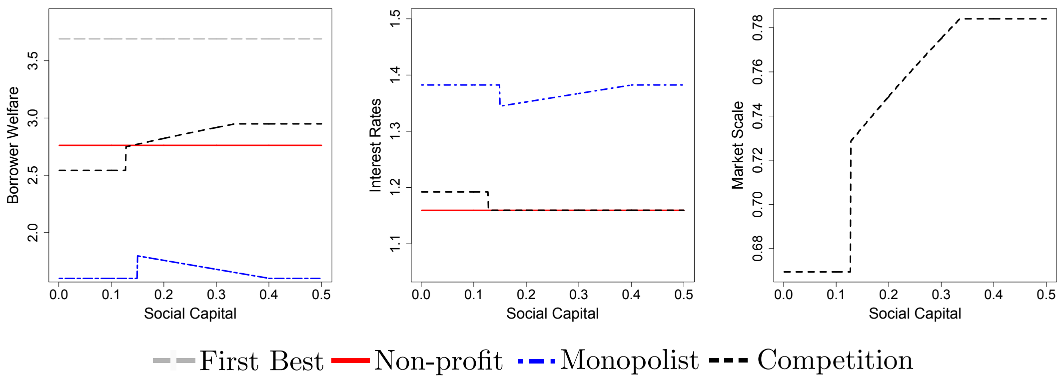

Figure 2.3: Full Sample Welfare, Interest Rates and Market Scale.

The first graph depicts borrower welfare, $\hat{V}$, $\tilde{V}$ and Z, and we also indicate the first-best borrower welfare level, $\frac{pR-\rho}{1-\delta}$. At jumps in the graph the contract switches from IL to JL. The welfare differences between the different market forms are substantial, with the interesting result that competition and non-profit lending are not strictly ordered. As discussed in section 2.2, this follows from the assumption that the non-profit uses strict dynamic incentives; in our view the key lesson is that non-profit and competition achieve similar performance despite the externality under competition. See also further discussion below.

The second panel depicts the interest rates offered by the monopolist and non-profit (competitive interest rates are not reported, but correspond to the zero-profit interest rate for the relevant contract and value of m). We observe that monopolist

[38] In a similar study in Ghana, they find comparable figures. Udry and Anagol (2006) find returns around 60% in one exercise, and substantially higher in others.

Table 2.2: Lending methods and S thresholds across regions

	Lending Methods			S thresholds			
	M	NP	C	M		C	
				$\tilde{S}$	$\bar{S}$	$\tilde{\tilde{S}}$	$\bar{\bar{S}}$
Full Sample	IL-JL	JL	IL-JL	0.148	0.400	0.126	0.334
Central America	IL-JL	JL	IL-JL	0.333	0.400	0.307	0.363
South America	IL-JL	JL	IL-JL	0.112	0.263	0.097	0.330
Eastern Africa	IL	IL	IL				
Northern Africa	IL	JL	IL-JL			0.188	0.319
Western Africa	IL-JL	JL	IL-JL	0.317	0.400	0.296	0.367
South Asia	IL-JL	JL	IL-JL	0.143	0.400	0.123	0.331
South East Asia	IL	JL	IL-JL			0.146	0.315
South West Asia	IL	JL	IL-JL			0.077	0.315

Notes: M, NP and C denote Monopoly, Non-Profit and Competition, respectively. Lending methods denotes which contract forms are used in equilibrium for some S. For example, "IL-JL" means that the lender uses IL for low S and JL for high S. "IL" means that JL is never used, and vice versa. S thresholds denote switch points from IL to JL, where a switch occurs. The Non-Profit never switches lending method for our parameter values.

Table 2.3: Interest Rates, Market Scale and Borrower Welfare

	Interest Rates			Market Scale		Borrower Welfare			
	M		NP	C		M	NP	C	
	$\tilde{r}^{IL}$	$\tilde{r}^{JL}(\tilde{S})$	$\hat{r}$ [a]	$\underline{l}$	$\bar{l}$	$\tilde{V}^{JL}(\tilde{S})$ [b]	$\hat{V}$ [a]	$Z(0)$	$Z(\bar{\bar{S}})$
Full Sample	1.382	1.345	1.159	0.669	0.784	1.796	2.761	2.543	2.949
Central America	1.376	1.363	1.251	0.432	0.476	1.648	2.074	1.581	1.733
South America	1.398	1.359	1.154	0.712	0.826	1.830	3.016	2.883	3.313
Eastern Africa	1.357		1.342[a]	0.063			1.642[a]	0.217	
Northern Africa	1.394		1.118	0.935	0.990		3.698	3.535	3.725
Western Africa	1.434	1.419	1.317	0.399	0.449	1.662	2.116	1.706	1.914
South Asia	1.370	1.331	1.137	0.698	0.813	1.800	2.803	2.573	2.967
South East Asia	1.475		1.166	0.955	0.995		5.498	5.351	5.562
South West Asia	1.416		1.117	0.879	0.963		3.983	3.811	4.150

Notes: [a] This is the JL interest rate or borrower welfare with a non-profit except where annotated with [a], in which case the values corresponds to the IL case as there is only IL lending in equilibrium. [b] $\tilde{V}^{IL}$ is equal to $pR = 1.6$ in every case, so not reported.

rates are substantially higher. Furthermore, leverage of social capital affects the interest rate and borrower welfare for $S \in [0.15, 0.40]$ in the full sample. Moving to the third panel, we see how market scale under competition varies with S. Market scale follows the same pattern as the Z function and ranges between 67% and 78%. As previously discussed, this should be interpreted as a measure of local, and not national penetration.

We can now analyze the welfare implications of market power and the lender's choice of contractual form. When the monopolist voluntarily switches from IL to JL at $\tilde{S}$, borrower welfare increases by approximately 12%. If we go further, forcing the monopolist to always use JL the gain is 20% at $S = 0$ (and declining in S). Switching to a non-profit lender delivers a minimum gain of 54% (at $\tilde{S}$) and a maximum of 73% for $S < \tilde{S}$ or $S \geq \bar{S}$. Thus our results underline the importance of constraining market power where it exists.

Similarly, we can consider the effect of mandating JL under competitive lending, since for $S < \tilde{\tilde{S}}$ the market equilibrium is IL only. We find that welfare would increase by 2% at $S = 0$, with this gain increasing as S increases, up to 16% at $\tilde{\tilde{S}}$. This illustrates one aspect of the inefficiency of the competitive equilibrium discussed in section 2.3.6. We graph the welfare effects of mandating JL or IL under monopoly and competition in Figure 2.5 in the Appendix.

2.3.4 Sensitivity analysis

We check the sensitivity of the results by varying each parameter over a reasonable range, while holding the others constant. For simplicity we focus on the results for $S = 0$. The results of these exercises are presented in Figure 2.4. We only plot the parameter regions in which the model predicts any lending, hence at $S = 0$, there is no lending for $\delta < 0.773$, $p < 0.887$, $\rho > 1.273$ and $R < 1.515$.

It becomes clear that welfare under a monopolist lender is not sensitive to any of the parameters, varying little in comparison with the larger effects under competition or non-profit lending (in particular, the monopolist's contract offer does not depend on ρ). For example, as R increases with a non-profit lender, all of the welfare gains are enjoyed by the borrowers. The monopolist, on the other hand, simply increases his interest rate, extracting almost all of the gains. Borrower welfare under competition typically tracks that under non-profit lending quite closely, so our conclusion that non-profit and competition have similar performance seems robust. The large welfare difference between non-profit and monopolist varies in each parameter, but is reasonably robust in the neighborhood of our estimates. It is also interesting to note that for low R, low p and δ welfare may be lower under competition than with a for-profit monopolist, as was theoretically predicted in Proposition 8. The patterns for interest rates are of course similar, and provide a useful check on the results.

In the third row we plot the three key S thresholds, $\hat{S}$, $\tilde{S}$, and $\tilde{\tilde{S}}$, at which non-profit, monopolist and competition switch from IL to JL lending. As predicted by

Proposition 7, the non-profit is the most likely to use JL, with $\hat{S} = 0$ over most of the parameter ranges. The monopolist is the least likely, and when δ is large abandons JL altogether as the LLC is too tight relative to IC1. Moreover, our predictions for the thresholds vary very little over most of the parameter ranges. The small non-monotonicities under competition arise due to switching between different values of m in the neighborhood of $\tilde{\tilde{S}}$.

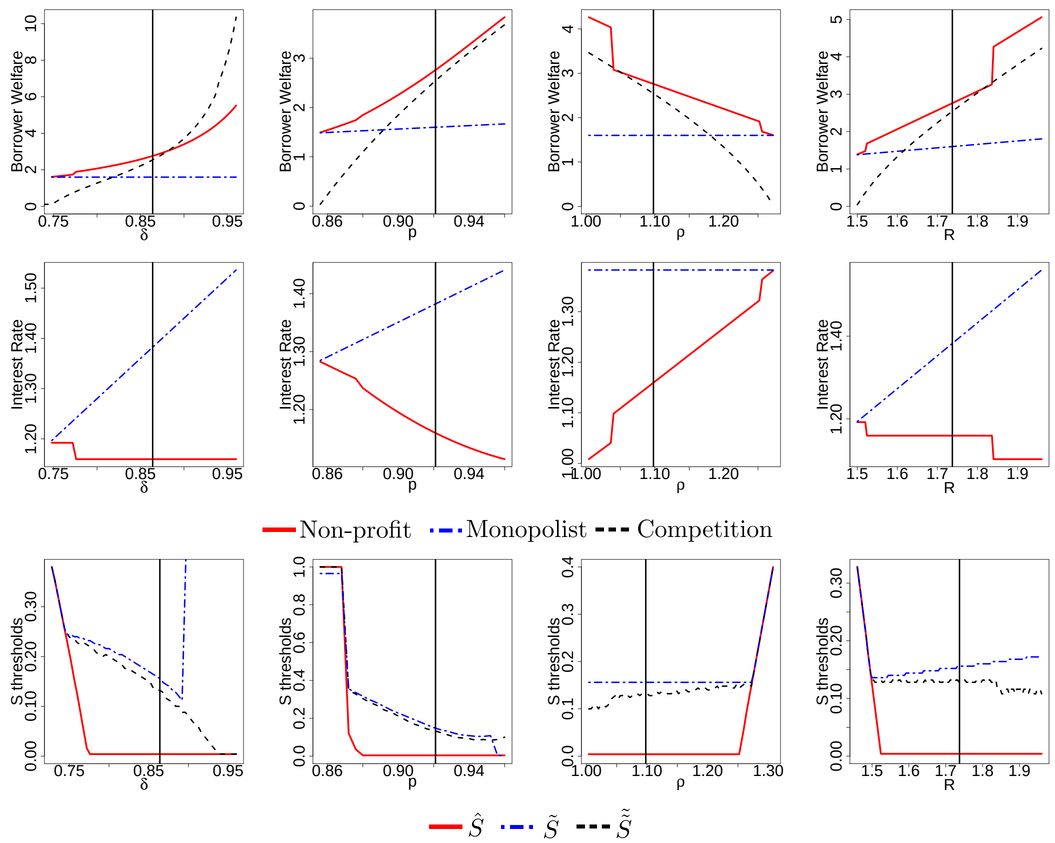

Figure 2.4: Sensitivity Analysis. Vertical lines indicate full sample parameter estimates.

Overall, the model gives some fairly robust predictions about contracts offered and the ordering of borrower welfare, and in particular the results highlight the conclusion that market power matters more than contract choice for borrower welfare. The interest rate predictions are more sensitive, but remain reasonable in all cases. With this in mind, we explore the patterns across regions. This allows us to comment on the respective lending types that we would expect to prevail in certain regions and it allows us to comment on the variation in borrower welfare across regions due to different market structures.

2.3.5 Regional analysis

We now turn to the results at the regional level, presented in Tables 2.2 and 2.3. We graph the predicted borrower welfare functions in Figure 2.7 in Appendix 2.A.9. We focus on seven regions with at least 1% of the total number of outstanding loans, comprising 94.2% of the total. We first observe that our parameter estimates always satisfy Assumption 4, so the model predicts at least IL lending in every region. However, the pattern of contracts offered depends on the market structure. In Eastern Africa, the model predicts only IL lending under all three market structures; the JL LLC is too tight for JL to even break even in these regions, primarily since the low success probability requires high interest rates. In Northern Africa, South East Asia and South West Asia, the non-profit would always offer JL, while the monopolist always offers IL. The relatively high success probabilities mean that the guarantee effect of JL is small relative to the cost to the lender of lower interest rates. In these cases, uncoordinated competition delivers IL for low S and JL for high S.

In all regions except Central America and Eastern Africa we observe that in welfare terms the non-profit and competition achieve similar outcomes. This observation is highlighted in the sensitivity analysis. In Central America the for-profit monopolist outperforms competition for S sufficiently small. In Eastern Africa, which has a very low success probability rendering repayment guarantees very costly for borrowers, competition performs very poorly, while non-profit and for-profit monopolist are almost identical in welfare terms. In line with Proposition 7, we see that the monopolist is less likely to offer JL than the competitive market.

Finally we compare the interest rates computed from the data (reported in Table 2.1) to the model predictions.[39] The broad pattern we observe is that mean interest rates seem reasonably close to the predicted non-profit rates, suggesting that most MFIs are operating close to their zero-profit constraints. For example, in South Asia the mean observed (net) rate is 18%, while our prediction is 14%. In South America, the difference is larger; the mean rate is 24%, versus a prediction of 15%, and the difference is 11 percentage points in Northern Africa. The results suggest that all three of these have the potential for significant rent-extraction by lenders. South East Asia and South West Asia are the most striking cases, with observed rates 22 and 16 percentage points higher than predicted, respectively. Particularly in South Asia, where Grameen is based, our model predicts that abuse of market power by for-profit lenders would have severe consequences. However, as a whole, interest rates in South Asia are very close to our predictions for a non-profit, suggesting that either competition or pro-borrower motivation of the lenders is constraining this abuse.

2.3.6 Comparing competition and non-profit lending

A result that emerges from the simulations is that frequently the competitive market dominates the non-profit in welfare terms, despite the enforcement externality that leads to credit rationing. The reason for this is the relatively high repayment probabilities ensure that the population of unmatched borrowers is small, while all

[39] Note that we cannot compare these values with the monopolist interest rates. To see this, recall that the IL interest rate is δpR. Since δ^L was calibrated from $\frac{\bar{r}}{pR}$, by construction the monopolist IL rate will exceed the mean rate in the data, $\bar{r}$.

borrowers benefit from the ability to reborrow in future. Under the non-profit this is not available due to the assumption that strict dynamic incentives are used.

As we have argued, we do not consider this an unreasonable assumption. However, as mentioned in section 2.2.2, a benevolent non-profit can deliver at least the same welfare as the competitive market, by renewing borrowers' contracts with probability λ upon default. A simple way to achieve this would be to choose the renewal probability upon default to mimic the competitive outcome. In this case, the value function of the non-profit would be the envelope of the matched utility. We have computed this example and illustrate it in Figure 2.6 in Appendix 2.A.7. The welfare effect is not dramatic; borrower welfare under the non-profit increases from 2.761 with strict dynamic incentives to a maximum of 3.239 with the new contract, a 17% gain. Chosing λ optimally, the non-profit can perform even slightly better. We have simulated this as well and observe that the only difference arises because the non-profit would switch to JL for a lower value of S.

2.3.7 Discussion

Collecting the simulation results, a picture emerges supporting the discussion in the theoretical analysis. The monopolist for-profit lender does exploit the borrowers' social capital and this has economically meaningful effects on interest rates and welfare. However, these are substantially smaller than the change in interest rates and welfare when switching to a large non-profit lender. The severe "mission drift" implied by a switch to for-profit lending with market power has large consequences for borrower welfare, consistent with the concerns of Muhammad Yunus raised earlier.

The competition results are more positive. The theoretical welfare effects of competition are ambiguous, as shown in Proposition 8, due to the trade-offs between credit rationing, lower interest rates and the ability of borrowers to reborrow after an involuntary default. However, for the parameters estimated from our full sample and most regions considered, welfare under competition is approximately the same as under non-profit lending. Despite the negative press and industry concerns about competition in microfinance, here modeled as a enforcement externality, our results suggest a more positive view in which competition is able to mitigate the problems of market power. However, some discussion, particularly surrounding the recent crisis in Andhra Pradesh, India, has centered on multiple borrowing and over-indebtedness. These are important issues but not ones that we can address in our framework, so we leave them to future research.

Lastly, the findings corroborate the theoretical prediction that for-profit lenders are less likely to offer JL than the non-profit. In low social capital areas (S smaller than around 0.13-0.15) our competitive or monopolistic lenders would offer IL, while the non-profit continues to offer JL.

2.4 Conclusion

Motivated by recent debates about commercialization and the trade-off between the objectives of making profits and alleviating poverty this paper studies the consequences of market power in the context of microfinance. We focus on the consequences for borrower welfare going beyond the usual focus on repayment rates and interest rates. The existing literature on microfinance starts with the premise that MFIs are competitive or motivated by borrower welfare and in this paper we showed that there are interesting implications for relaxing this assumption. A lender with market power can extract rents from repayment guarantee agreements between his borrowers, but is ultimately constrained from making those borrowers worse off in the process. We compare borrower welfare under a for-profit with market power, a benevolent non-profit, and a competitive credit market. One of the interesting trade-offs that emerges is that of rent extraction under monopoly with the enforcement externality under competition. We simulated the model using empirical parameter estimates, and found that the consequences of market power for borrower welfare are significant, while the choice of lending method itself is somewhat less important. Competitive for-profits typically do not perform much worse than our non-profit benchmark, especially when the level of social capital is high. Furthermore, commercial lenders with or without market power are less likely to use JL than the non-profit lender. Our findings suggest that Yunus appears to be correct to be concerned about abuses by for-profit lenders - particularly in South Asia, where Grameen is based, our model predicts considerable scope for abuse of market power by for-profit lenders. However, as a whole, interest rates in South Asia are very close to our predictions for a non-profit lender, weakening the case for concluding there is systemic abuse in practice in this region.

There are several directions for future work that we believe might be promising. For example, Muhammed Yunus argues that the shift from non-profit to for profit, with some institutions going public, led to aggressive marketing and loan collection practices in the quest for profits to serve the shareholders equity. Our paper does not model coercive loan collection methods by lenders, and allowing this might create an additional channel for for-profit and non-profits to behave differently, in a manner similar to the cost-quality trade-off as in the non-profits literature (see, for example, Glaeser and Shleifer (2001)).

2.A Proofs, Derivations and Simulation Results Omitted in the Paper

2.A.1 Patient monopolist inefficiently under-uses JL

Proposition 4 and Observation 1 point out that the monopolist lender inefficiently under-uses joint liability relative to the non-profit lender. One concern might be that this is due to the fact that the lender is assumed (for simplicity) to be myopic, choosing the contract based only on per-period revenue. JL's relatively high repayment rate makes it relatively more attractive to a patient lender. However this does not overcome the basic inefficiency result as we show here.

Suppose the lender discounts profits from a given borrower with factor $\beta \in [0,1]$. Now the lender's discounted profits per borrower are

$$\Pi = \frac{\pi r - \rho}{1 - \beta\pi}.$$

The only ingredient of the benchmark model that will change is the monopolist lender's contract choice. The constraints and thus interest rates for a given contract as a function of S remain the same. The monopolist now prefers JL whenever $\frac{q\hat{r}^{JL}(S)-\rho}{1-\beta q} \geq \frac{p\hat{r}^{IL}-\rho}{1-\beta p}$. We can solve this condition for a new $\tilde{S}(\beta)$ which is the value of S at which the lender switches from IL to JL. This is:

$$\tilde{S}(\beta) \equiv \max\left\{0, \frac{p^2R(p-\delta q)}{q(1-\delta q)} - \frac{\beta(1-p)(2-\delta q)(\delta p^2 R - \rho)}{\delta(1-\beta p)(2-p)(1-\delta q)}\right\}.$$

$\tilde{S}(\beta)$ is strictly decreasing in β. Therefore, as intuitively argued above, the monopolist becomes more willing to offer JL as β increases. However, this does not reverse the inefficiency result.

When $p < \delta q$, we know that $\hat{S} = \tilde{S} = 0$ and there is no inefficiency, so we focus on the case where $p > \delta q$. We want to show that the monopolist is less willing to offer JL than the nonprofit. Recall that $\hat{S} \equiv \max\left\{0, \frac{(2-\delta q)\rho-(2-p)\delta p^2 R}{\delta q(1-\delta q)}\right\}$. Subtracting the non-zero term in the max in $\hat{S}$ from that in $\tilde{S}(\beta)$ we obtain $\frac{(\delta p^2 R-\rho)(2-\delta q)(1-\beta q)}{\delta q(1-\delta q)(1-\beta p)}$, which is positive for all β, so we know that $\tilde{S}(\beta) \geq \hat{S}$ for all β, and that this inequality is strict for $\tilde{S}(\beta) > 0$.[40] Therefore the monopolist is less willing to offer JL than the nonprofit, and thus potentially inefficient, even when fully patient ($\beta = 1$).

2.A.2 Derivation of pooling contract under heterogeneity

Under IL it is obvious that charging $\tilde{r}^{IL} = r_{IC1}$ is optimal from the lender's perspective.

Under JL, if he charges less than $\tilde{r}^{JL}(0)$, repayment will not increase but revenue will be lower than charging $\tilde{r}^{JL}(0)$. If he charges between $\tilde{r}^{JL}(0)$ and $\tilde{r}^{JL}(S_h)$, the low S borrowers will still only repay with probability p^2 and the high S with probability

[40] I.e. for $\beta < \frac{\delta p^2 R(p-\delta q)(1-\delta q)}{\delta p^2 Rq(1-\delta p(1-\delta q))-\rho p(1-p)(2-\delta q)}$.

q, and revenue will be lower than charging $\tilde{r}^{JL}(S_h)$. Lastly, charging more than $\tilde{r}^{JL}(S_h)$ results in a repayment probability of p^2 from all borrowers and revenue lower than that attainable under IL.

2.A.3 Derivation of separating contract under heterogeneity

In a separating equilibrium, it must be that the lender offers one IL and one JL contract.[41] Define the interest rates as $\tilde{r}^{IL}_{sep}(S_h)$ and $\tilde{r}^{JL}_{sep}(S_h)$ where S_h is the social capital of the high S group.

By IC1, the lender will never charge more than δpR under either contract. If he did so under both, all borrowers would default, and if he did so under only one, then all borrowers would take the other contract.

Since a borrower's utility from a *given* JL contract (i.e. when holding *r constant*) is increasing in S, the high S types must choose JL and the low S types choose IL in equilibrium. Furthermore, since utility from a *given* IL contract does not depend on S, both types value IL equally. Denoting the utility in separating equilibrium of a borrower with social capital S under contract j by $V^j_{sep}(S)$, we have:

$$V^{IL}_{sep}(S_h) = V^{IL}_{sep}(0) \equiv V^{IL}_{sep}$$
$$V^{JL}_{sep}(S_h) \geq V^{JL}_{sep}(0)$$

the truth-telling constraints are:

$$V^{JL}_{sep}(0) \leq V^{IL}_{sep} \tag{2.3}$$
$$V^{JL}_{sep}(S_h) \geq V^{IL}_{sep}. \tag{2.4}$$

with one strict. These reduce to $V^{JL}_{sep}(S_h) > V^{JL}_{sep}(0)$ and either $V^{JL}_{sep}(S_h) = V^{IL}_{sep}$ or $V^{JL}_{sep}(0) = V^{IL}_{sep}$.

$V^{JL}_{sep}(S_h) > V^{JL}_{sep}(0)$ requires that $\tilde{r}^{JL}_{sep}(S_h) \leq \tilde{r}^{JL}(S_h)$. Otherwise, the high types would not be able to guarantee one another under JL, so all types would repay with probability p^2 in which case $V^{JL}_{sep}(S_h) = V^{JL}_{sep}(0)$. Similarly, $\tilde{r}^{JL}_{sep}(S_h) > \tilde{r}^{JL}(0)$ since otherwise low types would be able to guarantee under JL which would again mean that $V^{JL}_{sep}(S_h) = V^{JL}_{sep}(0)$. Clearly, then, the lender wants to charge the highest possible interest rate under IL and JL subject to these constraints. This is easy to find. Under JL he charges $\tilde{r}^{JL}_{sep}(S_h) \leq \tilde{r}^{JL}(S_h)$. This minimizes $V^{JL}_{sep}(S_h)$ and $V^{JL}_{sep}(0)$. Then he charges the highest possible interest rate under IL, such that $V^{JL}_{sep}(0) = V^{IL}_{sep}$.

Solving (2.3), we obtain the following expression for $r^{IL}_{sep}(S_h)$:

$$\tilde{r}^{IL}_{sep}(S_h) = \frac{\delta pR(1-p)}{1-\delta p^2} + \frac{p(1-\delta p)}{1-\delta p^2}\tilde{r}^{JL}_{sep}(S_h)$$

[41] Note that there can be no "screen-out" equilibrium, since the borrowers' participation constraints are slack: a borrower can always obtain utility pR by taking a loan and defaulting. Thus the separating equilibrium must involve a contract offer for both types.

substituting for $\tilde{r}^{JL}_{sep}(S_h) = \tilde{r}^{JL}(S_h)$, $\delta pR = \tilde{r}^{IL}$ and $\phi \equiv \frac{1-p}{1-\delta p^2} < 1$, we obtain

$$\tilde{r}^{IL}_{sep}(S_h) \equiv \phi \tilde{r}^{IL} + (1-\phi)\tilde{r}^{JL}(S_h).$$

This concludes the derivation.

2.A.4 Equilibrium with heterogeneous social capital and $p \leq \delta q$

If $p \leq \delta q$, the lender prefers the JL pooling contract to the IL one. Moreover, it is not guaranteed that he will ever choose the separating contract. To see this, note that if $S_h > \bar{S}$, revenue in the separating equilibrium is equal to $(\theta p + (1-\theta)q)\delta pR$. This is only greater than revenue in the pooling equilibrium, $\frac{\delta pqR}{2-\delta q}$, if $\theta < \frac{(2-p)(1-\delta q)}{(2-\delta q)(1-p)}$, a threshold strictly smaller than one when $p < \delta q$. Hence when the fraction of low S types is large, the lender may never offer the separating contract. Intuitively, in the pooling contract these borrowers receive a JL contract at interest rate $\tilde{r}^{JL}(0)$, while in the separating contract they receive IL, and we know that the former earns higher revenue than the latter when $p \leq \delta q$.

For simplicity, suppose $\theta < \frac{(2-p)(1-\delta q)}{(2-\delta q)(1-p)}$, so that by an analogous argument to that given above (footnote 22) there exists a threshold, $S_h^{p \leq \delta q}$, such that for S_h above the threshold the lender offers the separating contract. We have the following:

Proposition 9 *When $p \leq \delta q$, for $S_h < S_h^{p \leq \delta q}$, the lender offers JL at interest rate $\tilde{r}^{JL}(0)$. For $S_h \geq S_h^{p \leq \delta q}$ he offers IL at $\tilde{r}^{IL}_{sep}(S_h)$ and JL at $\tilde{r}^{JL}(S_h)$, low S borrowers take the IL contract and high S borrowers take JL. Welfare of both types decreases discontinuously at $S_h^{p > \delta q}$. Welfare of both types of borrowers further decreases in S thereafter.*

To see why welfare now *decreases* at the switching threshold, simply note that the pooling equilibrium in this case is the most favorable contract the monopolist ever offers the borrowers - it achieves the highest possible repayment probability, q, at the lowest interest rate the lender will ever charge. At the switching point, the low types switch to IL (lower repayment probability) at a higher interest rate, and the high types keep JL but at a higher interest rate. Thereafter, Observation 2 applies as before. Now, higher social capital among high S borrowers is doubly bad for borrower welfare.

2.A.5 Competition

Consider a repayment probability π. IC1 requires that the value of future access to credit from the current lender, less the repayment amount, exceeds the borrower's outside option which is to return to the pool of unmatched borrowers. At the zero profit interest rate the condition is:

$$\delta V - \frac{\rho}{\pi} \geq \delta U.$$

Simplifying, we obtain

$$\delta pR\frac{1-\chi(l,\eta)}{1-\delta\pi\chi(l,\eta)} \geq \frac{\rho}{\pi}.$$

We denote the left hand side by $r^{IL}_{IC1}(\chi)$ under IL (when $\pi = p$) and $r^{JL}_{IC1}(\chi)$ under JL ($\pi = q$).

Unlike the single-lender case, we have a different IC1 for IL and JL. Note that $r^{JL}_{IC1}(\chi) > r^{IL}_{IC1}(\chi)$, so it is not possible for both IC1s to bind simultaneously. Also note that for all $\chi > 0$, $\frac{dr_{IC1}(\chi)}{d\chi} < 0$ for IL and JL. This is the competition effect through improvements in the borrowers' outside option. Also note that, as before, provided IC1 holds JL borrowers will always be willing to repay their own loans provided their partner is also repaying.

The IC2 under JL requires that repayment of both loans is preferred to losing access to the current lender (rejoining the unmatched pool) and losing the social capital shared with the current partner. The condition is

$$\delta(V+S) - 2\frac{\rho}{\pi} \geq \delta U,$$

which simplifies to

$$\frac{\delta[(1-\chi(l,\eta))pR + (1-\delta q)S]}{2-\delta q-\delta q\chi(l,\eta)} \geq \frac{\rho}{q}.$$

We denote the left hand side by $r_{IC2}(S,\chi)$. $r_{IC2} \leq r^{JL}_{IC1}$ for $S \leq \frac{1-\chi(l,\eta)}{1-\delta q\chi(l,\eta)}pR$ (we compute the equilibrium value of this threshold below). $\frac{dr_{IC2}(S,\chi)}{d\chi} < 0$ whenever $\chi > 0$ and IC2 is tighter than IC1.

Therefore, in conclusion, the tighter of $r_{IC2}(S,\chi)$ and r^{JL}_{IC1} is downward sloping in χ. This is the effect of competition on the repayment incentive constraints.

Recall now the two key thresholds, stated in the text and derived below: $\tilde{\tilde{S}} \equiv \frac{p-\delta q}{\delta q(1-\delta q)}\rho$, the analog of $\tilde{S}$, representing the level of social capital at which the competitive market switches from IL to JL lending; and $\bar{\bar{S}} \equiv \frac{\rho}{\delta q}$, the analog of $\bar{S}$, the level of social capital at which IC1 binds under JL (as opposed to IC2). Note also that $\bar{\bar{S}} > \tilde{\tilde{S}}$.

Proposition 6 (restated) If $\tilde{\tilde{S}} \leq 0$, the competitive equilibrium is JL-only lending, with market scale strictly increasing in S for $S < \bar{\bar{S}}$, and equal to a constant, $\bar{\bar{l}}$ for $S \geq \bar{\bar{S}}$. If $\tilde{\tilde{S}} > 0$, the equilibrium for $S < \tilde{\tilde{S}}$ is IL-only lending at fixed scale $\underline{l}$. At $\tilde{\tilde{S}}$, all lending switches to JL at scale $\tilde{\tilde{l}} > \underline{l}$, then increases continuously in S to $\bar{\bar{l}}$, at $\bar{\bar{S}}$. Welfare, Z, is strictly increasing in scale, l, and therefore weakly increasing in S.

Proof. In equilibrium, at most one of IL and JL breaks even (except for at the switching threshold as discussed below). So for a lender, who takes χ as given, the following condition will hold:

$$\rho = \max\{pr^{IL}_{IC1}(\chi), \min\{qr^{JL}_{IC1}(\chi), qr^{JL}_{IC2}(S,\chi)\}\}.$$

Next note that $pr^{IL}_{IC1}(\chi) < qr^{JL}_{IC1}(\chi)$, so to determine the equilibrium contract, we only need to compare $pr^{IL}_{IC1}(\chi)$ with $qr^{JL}_{IC2}(\chi, S)$. If $\rho = pr^{IL}_{IC1}(\chi) > qr^{JL}_{IC2}(S, \chi)$, only IL will be used in equilibrium, and if $\rho = qr^{JL}_{IC2}(S, \chi) \geq pr^{IL}_{IC1}(\chi)$ only JL will be used (we assume that JL will be offered when both IL and JL break even).

Solving $\rho = pr^{IL}_{IC1}(\chi)$, for χ we obtain the equilibrium value of χ under IL:

$$\tilde{\tilde{\chi}} = \frac{\delta p^2 R - \rho}{\delta p^2 R - \delta p \rho}.$$

Next we solve $pr^{IL}_{IC1}(\tilde{\tilde{\chi}}) = qr^{JL}_{IC2}(\tilde{\tilde{\chi}}, S)$ to find $\tilde{\tilde{S}}$, the switching threshold value of S at which both IL and JL break even:

$$\tilde{\tilde{S}} = \frac{p - \delta q}{\delta q(1 - \delta q)} \rho \gtreqless 0.$$

Lastly, we solve $\rho = qr_{IC2}(S, \chi)$ to find the equilibrium value of χ under JL when IC2 is binding. This is $\chi = \psi(S)$ where we define $\psi(S)$ as:

$$\psi(S) \equiv \tilde{\tilde{\chi}} + \frac{1 - \delta q}{pR - \rho}(S - \tilde{\tilde{S}}).$$

There is no equilibrium with JL lending for $S < \tilde{\tilde{S}}$, since then $\chi(S) < \tilde{\tilde{\chi}}$, in which case the IL IC1 would be slack and new lenders would enter offering IL. By a symmetric argument there is no equilibrium with IL lending for $S > \tilde{\tilde{S}}$. At $\tilde{\tilde{S}}$, lending switches from IL to JL, so η changes discontinuously from 1 to 0. This enables us to solve for market scale, using $\chi(\underline{l}, 1) = \chi(\tilde{\tilde{l}}, 0) = \tilde{\tilde{\chi}}$. We obtain the market scale under IL, equal to

$$\underline{l} = \frac{\delta p^2 R - \rho}{\delta p^2 R - p\rho},$$

and the market scale after the switch to JL, equal to

$$\tilde{\tilde{l}} = \frac{(\delta p^2 R - \rho)(1 - \delta q)}{(\delta p^2 R - p\rho)(1 - \delta q) + p(1 - p)(1 - \delta)\rho}.$$

Given that by the definition of $\underline{l}$ and $\tilde{\tilde{l}}$, the following condition holds (as shown above): $\chi(\underline{l}, 1) = \chi(\tilde{\tilde{l}}, 0) = \tilde{\tilde{\chi}}$, and that $\chi_l > 0$ and $\chi_\eta > 0$ for all $l > 0$, it follows that $\tilde{\tilde{l}} > \underline{l}$.[42] Intuitively, in a JL equilibrium, borrowers default less frequently, so for a given market scale it is less likely that a "slot" will come available at a lender for a currently unmatched borrower. As a result, for a given value of χ, a higher level of l can be sustained under JL than under IL.

Lastly note that if $\tilde{\tilde{S}} < 0$, there is never IL lending in equilibrium. Even for $S = 0$, the JL IC2 is more slack than the IL IC1 and therefore $\chi = \psi(0) > \tilde{\tilde{\chi}}$. Thus, market scale at $S = 0$ exceeds $\tilde{\tilde{l}}$.

[42] The interested reader may note that there are many mixed equilibria at $\tilde{\tilde{S}}$, defined by a one-to-one function $l(\eta)$, $\eta \in [0, 1]$, of which $l = \tilde{\tilde{l}}$, $\eta = 0$ is the welfare-maximizing case.

Now consider $S > \tilde{\tilde{S}}$. Lending is JL-only (i.e. $\eta = 0$). Since IC2 is relaxed as S increases, entry will occur to compensate, so l and hence χ are strictly increasing in S as long as IC1 is slack. IC2 must then intersect IC1 at some $\bar{\bar{S}} > \tilde{\tilde{S}}$, where χ reaches a maximum $\bar{\bar{\chi}}$. For $S \geq \bar{\bar{S}}$, IC1 is tighter than IC2, and therefore market scale has reached its maximum.

To find $\bar{\bar{S}}$ and $\bar{\bar{\chi}}$ we simply need to solve for the values at which IC1 intersects IC2 in competitive equilibrium. In other words, we solve the following condition

$$\rho = qr^{JL}_{IC1}(\bar{\bar{\chi}}) = qr^{JL}_{IC2}(\bar{\bar{S}}, \bar{\bar{\chi}})$$

obtaining:

$$\begin{aligned}\bar{\bar{S}} &= \frac{\rho}{\delta q} \\ \bar{\bar{\chi}} &= \frac{\delta pqR - \rho}{\delta pqR - \delta q\rho}.\end{aligned}$$

Lastly, to obtain the maximum scale, $\bar{\bar{l}}$, we solve $\chi(\bar{l}, 0) = \bar{\bar{\chi}}$ yielding $\bar{l} = \frac{\delta pqR - \rho}{\delta pqR - q\rho}$.

For $S \in (\tilde{\tilde{S}}, \bar{\bar{S}})$, we find the market scale by setting $\chi(l, 0) = \psi(S)$, obtaining equilibrium l equal to $\frac{(1-\delta q)\psi(S)}{(1-q)+q(1-\delta)\psi(S)}$ which is strictly increasing in S. Collecting results, we can write the equilibrium market scale as the following function of S:

$$l(S) \equiv \max\left\{ \underline{l}, \min\left\{ \frac{(1-\delta q)\psi(S)}{(1-q) + q(1-\delta)\psi(S)}, \bar{\bar{l}} \right\} \right\}.$$

Where there is IL-only lending for $S < \tilde{\tilde{S}}$ and JL-only for $S \geq \tilde{\tilde{S}}$.

■

2.A.6 Mandating JL or IL

In section 2.3.3 we discussed the welfare effects of mandating JL or IL under monopoly or competitive lending. These are illustrated in Figure 2.5.

2.A.7 Stochastic Renewal

Suppose the lender offers either JL or IL, but renews the group's contracts with certainty following repayment and with probability $\lambda \in [0, 1]$ following default. One complication immediately arises. Suppose the state is $(R, 0)$ and the interest rate is r. If borrower 1 defaults, her social capital is lost but the group might survive, so her IC2 is $\delta(V(S, r) + S) - 2r \geq \delta\lambda V(0, r)$. For a given interest rate r, $V(S, r) \geq V(0, r)$, since without social capital repayment guarantees may not be possible. This may be a key reason why such flexible penalties are not widely used - the borrowing group dynamic may be too badly damaged following a default. To retain the basic structure of our benchmark model, we make the simplifying assumption that if the borrowers' contracts are renewed following a default, the group is dissolved and members matched up with new partners with whom they share the same value of

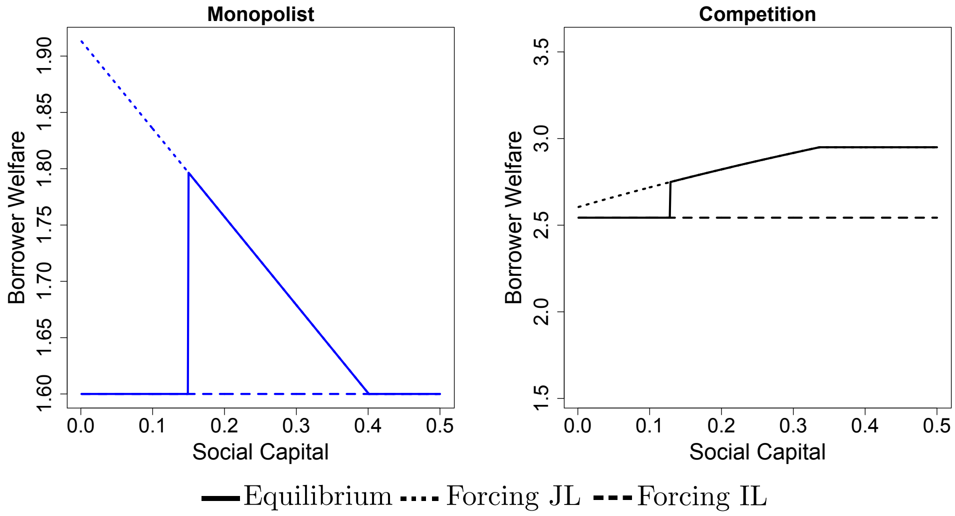

Figure 2.5: Mandating contractual form. Social capital ranges on horizontal axes, borrower welfare on vertical axes.

social capital. Default is still costly, since it destroys the social capital of the existing group, but does not adversely affect the dynamic of the group if it survives. This assumption is the analogue of the group reformation assumption in the competition framework.

It is easy to see that the stochastic renewal setup closely mirrors the competition framework. Specifically, for a given S, a single lender could offer a the same contract (IL or JL and the same interest rate) as offered under competition, that renews with probability $\lambda = \frac{U(S)}{\tilde{\tilde{V}}(S)}$ following default. The tightest of IC2 and IC1 would bind, and all borrowers would receive utility $\tilde{\tilde{V}}(S)$.[43] However, the contracts that emerge in equilibrium are quite different, as shown in the following proposition.

Proposition 10 *Consider the following modification to the contracting setup: the lender renews the borrowers' contracts with certainty after repayment, and probability λ following default. Equilibrium contracts are as follows:*

1. *Neither the monopolist nor competitive lenders use stochastic renewal: $\lambda = 0$.*

2. (a) *If $\frac{\delta p^2 r}{\rho} \leq \frac{1-\delta p}{1-p}$, the nonprofit offers JL for all $S \geq \hat{S}$ as before, and $\lambda > 0$ for all S (unless the JL IC2 binds at $S = \hat{S}$, in which case $\lambda = 0$ at $\hat{S}$).*

   (b) *If $\frac{\delta p^2 r}{\rho} > \frac{1-\delta p}{1-p}$, there is an $\hat{\hat{S}} \in (\hat{S}, \frac{\rho}{\delta q})$ such that the nonprofit offers IL for all $S < \hat{\hat{S}}$, JL otherwise, and $\lambda > 0$ for all S.*

   (c) *When JL is used, λ and thus borrower welfare V is strictly increasing in S for all $S < \frac{\rho}{\delta q}$.*

3. *Borrower welfare is always higher with the nonprofit lender than under competition.*

[43] To see this, note that for repayment probability π and $r = \frac{\rho}{\pi}$, $\tilde{\tilde{V}} = pR - \rho + \delta(\pi V + (1-\pi)U)$, while the stochastic renewal contract yields $V = pR - \rho + \delta(\pi + (1-\pi)\lambda)V$.

Proof. The key relationship to check is the effect of λ on IC1 and IC2. For a given V, higher λ implies weaker penalty for default. However, higher λ increases V by improving the borrower or group's renewal probability. It turns out that the former effect dominates; the constraints are strictly tighter as λ increases.

First consider the single (non-profit or for-profit) lender case. Borrower utility with stochastic renewal and repayment probability π is

$$V = \frac{pR - \pi r}{1 - \delta(\pi + (1-\pi)\lambda)}.$$

The LLC is unchanged. The IC1 is $\delta(1-\lambda)V \geq r$ or

$$\frac{1-\lambda}{1-\delta\lambda}\delta pR \geq r.$$

The IC2 under JL is $\delta[(1-\lambda)V + S] \geq 2r$ or

$$\frac{\delta[(1-\lambda)pR + (1-\delta(q+(1-q)\lambda))S]}{2-\delta(q+\lambda(2-q))} \geq r.$$

Both are strictly tighter as λ increases. To see this for IC2, suppose IC2 binds. Rearranging, we obtain $\frac{dr}{d\lambda} = \frac{pR-\pi r}{2+\pi a(\lambda)}a'(\lambda)$ where $a(\lambda) = \frac{\delta(1-\lambda)}{1-\delta(q+(1-q)\lambda)} > 0$, $a'(\lambda) < 0$. Thus the monopolist always sets $\lambda = 0$, since increasing λ forces him to decrease the interest rate.

With competition, the corresponding constraints are

$$\delta(1-\lambda)(V-U) \geq r \tag{IC1}$$

$$\delta[(1-\lambda)(V-U) + S] \geq 2r \tag{IC2}$$

U is exogenous from the lender's perspective, and $V - U > 0$ in equilibrium. Using $V = pR - \pi r + \delta[(\pi + (1-\pi)\lambda)V + (1-\pi)(1-\lambda)U]$, we obtain $\delta(1-\lambda)(V-U) = a(\lambda)(pR - \pi r - (1-\delta)U)$, from which it is straightforward to check that both IC1 and IC2 are strictly tighter as λ increases. Thus stochastic renewal is never used in competition. To see this, consider an equilibrium with $U = U^*$ where some lender offers IL with $\lambda^* > 0$ and breaks even. This implies that, for his borrowers, $\delta(1-\lambda^*)(V(\lambda^*) - U^*) = \frac{\rho}{p}$. But then an entrant could offer IL with $\lambda' < \lambda^*$ and earn positive profits since $\delta(1-\lambda')(V(\lambda') - U^*) > \frac{\rho}{p}$. An analogous argument rules out equilibria with stochastic renewal and JL, and rules out entry by lenders using stochastic renewal in an equilibrium with no stochastic renewal.

The non-profit lender will use stochastic renewal whenever the tightest repayment constraint is slack at the zero-profit interest rate, since increasing λ improves borrower welfare without violating the constraint. We first analyze contract choice under IL and JL, then the choice of contract type.

Under IL, the lender chooses λ to bind IC1. The solution to $\frac{1-\lambda}{1-\delta\lambda}\delta pR = \frac{\rho}{p}$ is $\hat{\lambda}^{IL} \equiv \frac{\delta p^2 R - \rho}{\delta p^2 R - \delta\rho}$, which is strictly positive by Assumption 4.

Under JL, the lender chooses λ to bind the tighter of IC1 and IC2. Just as in the competition setup, IC1 and IC2 intersect at $S = \frac{\rho}{\delta q}$. If IC1 is binding, $\hat{\lambda}^{JL}(S) = \frac{\delta pqR-\rho}{\delta pqR-\delta\rho}$. If IC2 is binding, $\hat{\lambda}^{JL}(S) = \frac{\delta q[pR+(1-\delta q)S]-(2-\delta q)\rho}{\delta q[pR+\delta(1-q)S]-\delta(2-q)\rho}$. λ is strictly increasing in S until $S = \frac{\rho}{\delta q}$. However, note that if $S < \hat{S}$, JL is not usable even with $\lambda = 0$, and for $S > \hat{S}$, $\hat{\lambda}^{JL}(S) > 0$. Therefore, we have:

$$\hat{\lambda}^{JL}(S) = \begin{cases} 0 & S < \hat{S} \\ \frac{\delta q[pR+(1-\delta q)S]-(2-\delta q)\rho}{\delta q[pR+\delta(1-q)S]-\delta(2-q)\rho} & S \in [\hat{S}, \frac{\rho}{\delta q}) \\ \frac{\delta pqR-\rho}{\delta pqR-\delta\rho} & S \geq \frac{\rho}{\delta q} \end{cases}$$

The nonprofit chooses JL whenever

$$\hat{V}^{JL}(S, \hat{\lambda}^{JL}(S)) \geq \hat{V}^{IL}(\hat{\lambda}^{IL}).$$

Since the numerator is $pR - \rho$ in both cases, JL is used if and only if

$$1 - \delta(q + (1-q)\hat{\lambda}^{JL}(S)) \leq 1 - \delta(p + (1-p)\hat{\lambda}^{IL})$$

or

$$\hat{\lambda}^{JL}(S) \geq \frac{\hat{\lambda}^{IL} - p}{1-p}.$$

At $\hat{S}$ (i.e. $\hat{\lambda}^{JL}(S) = 0$) this reduces to $\frac{\delta p^2 R}{\rho} \leq \frac{1-\delta p}{1-p}$. If this condition holds, the lender offers JL for all $S \geq \hat{S}$, just as before. Otherwise, he offers JL for $S \geq \hat{\hat{S}}$, with $\hat{S} < \hat{\hat{S}} < \frac{\rho}{\delta q}$, defined implicitly by $\hat{\lambda}^{JL}(\hat{\hat{S}}) = \frac{\hat{\lambda}^{IL}-p}{1-p}$.

To see the last part of the proposition, we have already noted that by mimicking the competitive market the nonprofit can give utility $\tilde{\tilde{V}}$ to each borrower. However, as he is unconstrained by the market equilibrium conditions, he may be able to offer an alternative contract that yields higher borrower welfare. Secondly, since he uses stochastic renewal instead of credit rationing as a motivating device, this contract can be offered to all borrowers, instead of just the matched borrowers as under competition. ■

Stochastic renewal is more efficient than strict dynamic incentives. Nevertheless we find that the for-profit monopolist and competitive lenders will never use it. As a result, the nonprofit organizational form achieves the highest borrower welfare.[44] Figure 2.6 shows borrower welfare and λ under the simulated stochastic renewal contract.

[44]If the monopolist also valued future profits from a given borrower (non-myopic), he *would* use stochastic renewal, since there is now a tradeoff between higher interest rates and increasing the renewal probability. The result for the competitive market only relies on free entry and zero-profit equilibrium and therefore does not depend on the lenders' time horizon.

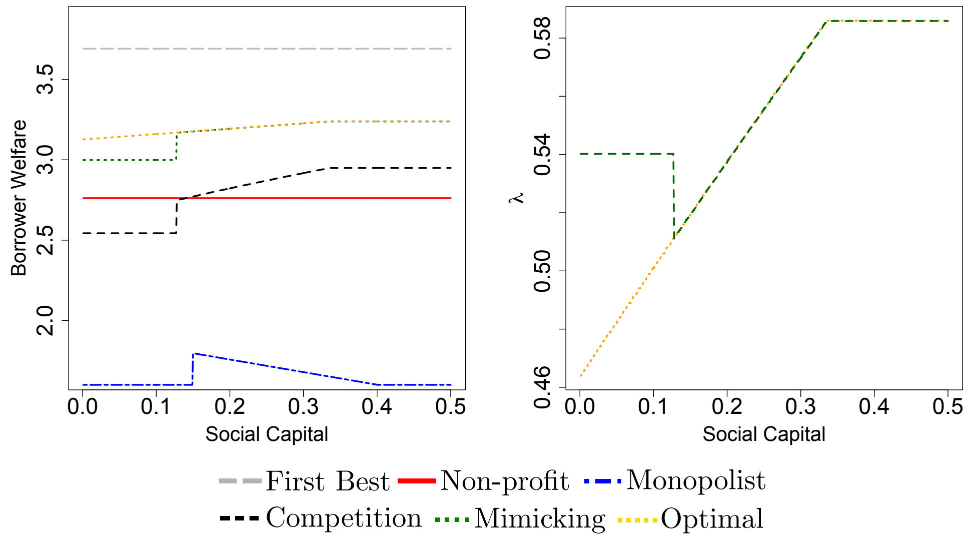

Figure 2.6: Simulating the stochastic renewal contract. Social capital ranges on horizontal axes, borrower welfare on vertical axes.

2.A.8 Group size and binding limited liability condition

Consider a group of size n, and suppose the group's loans are repaid whenever at least m members are successful. Then the repayment probability is

$$\pi(n,m) = \sum_{i=m}^{n} \binom{n}{i} p^i (1-p)^{n-i},$$

so

$$V = \frac{pR - \pi(n,m)r}{1 - \delta\pi(n,m)}.$$

IC1 is unchanged: $r_{IC1} = \delta pR$. For the successful borrowers to be willing to repay when exactly m are successful, each repaying $\frac{nr}{m}$, we must have $r \leq r_{IC2}(S,n,m)$, which we can derive as:

$$r_{IC2}(S,n,m) \equiv \frac{\delta m[pR + (1-\delta\pi(n,m))S]}{n - (n-m)\delta\pi(n,m)}.$$

The LLC requires that the m successful borrowers can afford to repay all 5 loans, i.e. $nr \leq mR$ yielding

$$r_{LLC}(n,m) \equiv \frac{mR}{n}.$$

For a given $r \leq r_{IC1}$, borrowers will choose the lowest m such that to IC2 and LLC are satisfied, so equilibrium m^* is determined by

$$\min\{r_{LLC}(n,m^*), r_{IC2}(S,n,m^*)\} \geq r > \min\{r_{LLC}(n,m^*-1), r_{IC2}(S,n,m^*-1)\}.$$

This m^* then defines the repayment probability function $\pi^*(S,n,r)$.

The non-profit lender chooses the lowest r such that $\pi^*(S,n,r)r = \rho$. The for-profit chooses r to maximize $\pi^*(S,n,r)r$.

Despite this modification, it may be that LLC at m^* is tighter than IC1, in which case the highest interest rate the lender can charge under JL will now be dictated by the LLC and smaller than r_{IC1}. If this is the case and the lender is a for-profit monopolist, borrowers will be *strictly* better off under JL than IL. However, if the LLC is very tight, JL may never be offered. This has three implications for the simulations. Firstly, the value of $\bar{S}$, obtained from the point at which the lender can no longer leverage social capital, depends on whether IC1 or LLC are tightest. Formally, with the group size modification,

$$\bar{S} = \min\left\{\frac{(n-m)pR}{m}, \frac{[n(1-\delta p)-(n-m)\delta\pi^*]R}{\delta n(1-\delta\pi^*)}\right\}.$$

Secondly, the interest rate and borrower welfare at $\bar{S}$ are be lower and higher respectively than the corresponding values under IL, when $r_{LLC} < r_{IC1}$. Thirdly, if r_{LLC} is very tight for every m there may be no value $\tilde{S}$ at which the lender is willing to offer JL.

2.A.9 Regional welfare predictions

Figure 2.7 plots the predicted borrower welfare in each of the regions considered in the simulations, as was discussed in section 2.3.3.

2.B Simulation Methodology

This Appendix outlines the algorithm used to simulate the core model. The simulation was implemented in Scilab, an open-source alternative to Matlab. Rather than solving the model explicitly, which becomes increasingly complicated with larger groups, we chose to simulate the optimization problem numerically. As the objective functions are all linear, this is a computationally tractable and simple task.

The simulation consists of two parts. The first part computes the optimal contracts of a non-profit and a monopolist lender, while varying the level of social capital S. The second part computes the competition section.

The section proceeds by presenting annotated pseudo-codes, that illustrates how the code proceeds to arrive at the optimal contracts.

Non-Profit and Monopolist

Here the optimization is very simple, as we do not have to study an entry condition, but just have to evaluate a set of constraints. The optimization procedure is carried out for each level of social capital, which then gives us the value functions we use for the main plots in the paper. Since $n = 5$ throughout we drop the n notation.

For each value of S:

Non-Profit

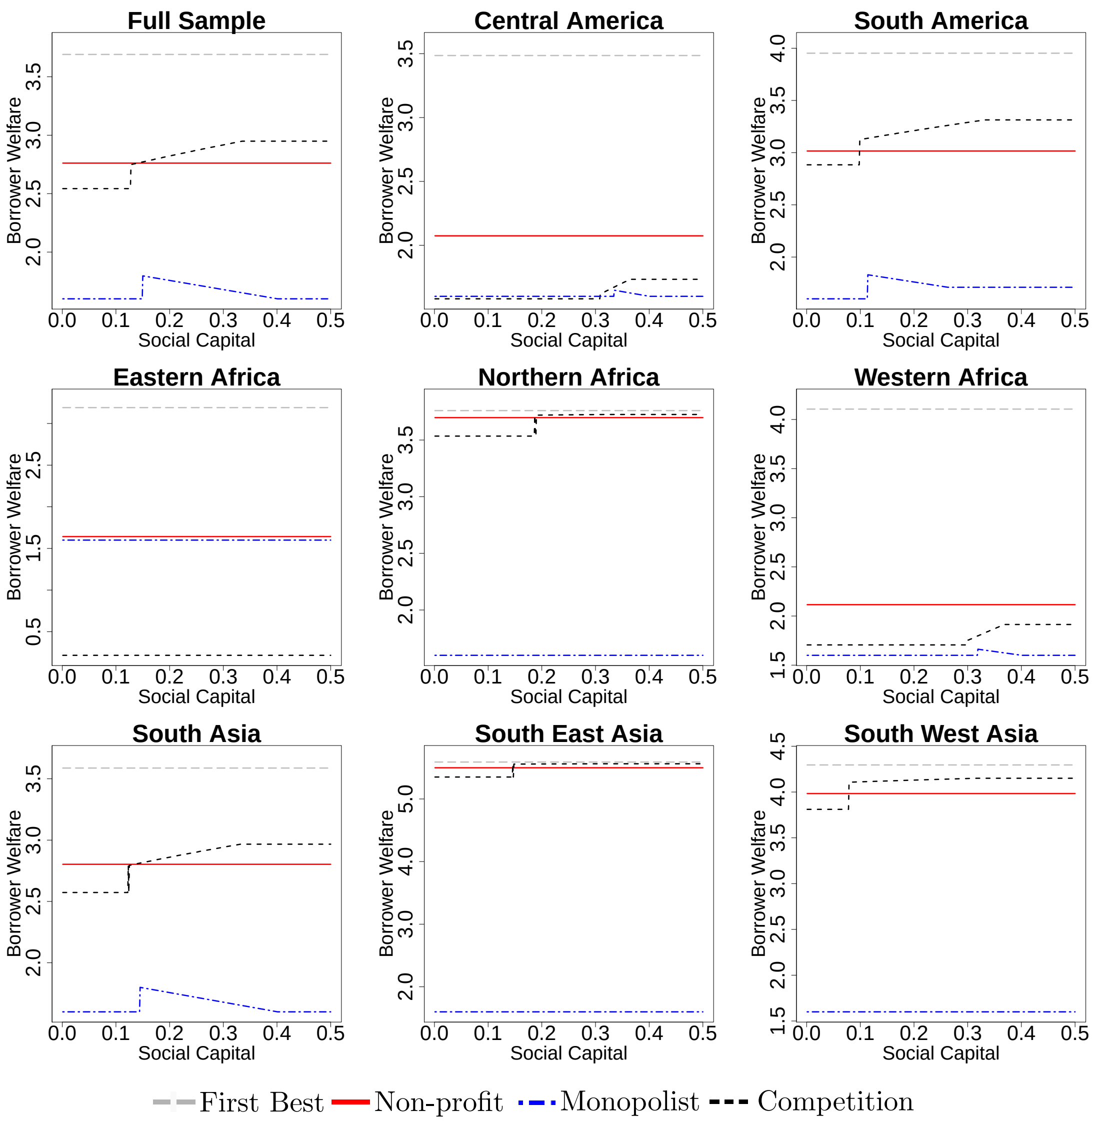

Figure 2.7: Borrower Welfare: Regional Differences. Social capital ranges on horizontal axes, borrower welfare on vertical axes.

1. JL: find the set M_{JL}^{ZP} of values for m that satisfy $r_{LLC}(m) \geq \rho/\pi(m)$ and the associated functions $\hat{V}^{JL}(m)$.

2. IL: Find, if it exists, the IL zero-profit equilibrium and the associated $\hat{V}^{IL}$.

3. Choose the contract (IL/JL), value of m and corresponding interest rate that gives borrowers maximal utility.

Monopolist

1. JL: For each $m \in M_{JL}^{ZP}$ find the maximal interest rate $\tilde{r}(m)$ such that $\tilde{r}_{JL}(m) = \min\{r_{IC2}(m), r_{LLC}(m), r_{IC1}\}$ and compute the associated profits $\tilde{\Pi}(m)^{JL} = \pi(m)\tilde{r}^{JL}(m) - \rho$.

2. IL: Compute the maximal interest rate $\min\{r^{IC1}, r^{LLC}\}$ and compute the associated profits $\tilde{\Pi}^{IL} = p\tilde{r}^{IL} - \rho$.

3. Choose the contract that maximizes profits.

Competition

For the competition model, we simulate the entry condition for lenders. For each value of S and U we check whether an entrant could earn positive profits with some contract (recall that in equilibrium there is always excess demand for credit). This will happen as long as the relevant constraints (see below) are slack at the relevant zero-profit interest rate. Hence, for each S we proceed by iteratively increasing U until the most profitable contract breaks even. The details are provided in the following pseudo-code:

For each value of S:

1. Initialize $U = 0$.

2. JL: for all $m = 1, ..., n$, check that all three constraints (LLC, IC2, IC1) are satisfied at the zero-profit interest rate.

3. IL: check that IC1 is satisfied at the zero-profit interest rate.

4. If there exists at least one contract such that all relevant constraints are satisfied, increase U by one unit and repeat from step 2. Otherwise, we have found the equilibrium value of U. The equilibrium contract (either IL or JL and the appropriate value of m) is the one for which all three constraints were satisfied in the previous round of iteration. If two or more contracts are feasible, pick the one that delivers the highest borrower welfare.

5. Given the equilibrium contract, solve U for the equilibrium market scale, and thus find Z.

Optimal Contract with Stochastic Renewal

The algorithm to determine the optimal level of λ is very similar to the one that determines the level of U in the competition simulation. The idea is, that a non-profit adjusts λ as long as the relevant constraints are slack. The key difference is that the non-profit finds the binding level of λ for for all different levels of m and then choses the level of m that provides borrowers with maximal utility. Free-entry competition may not yield the welfare-maximizing level of m. The reason is that entry continues until the slackest constraints eventually binds, which gives a single value for U. Under the optimal stochastic renewal contract, we find the optimal λ for each level of m respectively and then let the non-profit chose the welfare-maximizing contract. The details are provided in the following pseudo-code:

For each value of S:

1. Initialize $\lambda = 0$.
2. JL: for all $m = 1, ..., n$, check that all three constraints (LLC, IC2, IC1) are satisfied at the zero-profit interest rate.
   (a) if for any m, a constraint is violated, we record the current λ as the optimal one for that particular m.
3. IL: check that IC1 is satisfied at the zero-profit interest rate.
   (a) if the constraint is violated, we record the current λ as the optimal one for IL.
4. As long as there exists at least one contract such that all relevant constraints are satisfied (either IL or all JL), increase λ by one unit and repeat from step 2.
5. Evauate the value functions at the respective optimal λ and chose the contract that maximizes utility.

2.C Data Appendix

The dataset we work with comes from MIXMarket.org, an organization that collects, validates and publishes financial performance data of MFIs around the world. The MIX provides a set of reports and financial statements for each MFI reporting to it. The financial statements and reports were downloaded in March 2011, the relevant data was then extracted into a database using an automated script. The variables we use in this paper come from the MFIs' Overall Financial Indicators, the Income Statement, the Balance Sheet and the Products and Clients report. The Balance sheet and the Income statements are regular financial statements, while the Financial Indicators report variables such as Portfolio at Risk and the Products and Clients report include the number of loans by methodology.

The variables we use from the Balance Sheet are Value of IL Loans, Value of Solidarity Group Loans and overall Gross Loan Portfolio. From the Income statement we use the Operating Expense and the Financial Expense to compute the expense per dollar lent as described in the main text. From the Financial Indicators report, we use the Portfolio at Risk numbers, along with the Real Portfolio Yield to compute the risk adjusted real yields. From the Products and Clients report, we extract the Number of IL Loans and Number of Solidarity Group Loans, which we refer to in the main table and the text.

We work with a sample of 715 institutions for the year 2009. We chose the year 2009 as that is the year for which we have the largest number of institutions reporting lending methodology.[45]

The MIX data does not give us information whether JL is used, but they state that "loans are considered to be of the Solidarity Group methodology when some aspect of loan consideration depends on the group, including credit analysis, liability, guarantee, collateral, and loan size and conditions." We will refer to the share of loans falling into this category as JL share loans.

Sometimes the data on lending methodology by number of loans or by volume does not correspond exactly to the reported total portfolio or number of loans outstanding because of data entry errors, missing data or number of borrowers rather than number of loans reported. In these cases we assume that the errors are not biased toward either IL or JL, so we compute the share from the data we have. For example, if a lender reports $1m of loans, but $450k IL and $450k solidarity group lending, we compute an IL share of 50% and apply this to the whole portfolio. Of the 715 institutions in the sample, 143 have such incompleteness in the value data, 16.7% of the total Gross Portfolio is unaccounted for. As for the number of loans (which are not used in the estimation), 10 have no data so we use the value shares as a proxy, and 222 institutions have incomplete data; a total of 11.4% of the number of loans are unaccounted for. In total 304 institutions have some incompleteness in these data.

The relationship between the two is illustrated in Figure 2.8. Points lying on the 45 degree line correspond to lenders where the IL share by value is the same as the IL share by number. Each point corresponds to an MFI, with those in red, the "portfolio data incomplete", corresponding to the observations where the methodology breakdown does not exactly match the portfolio figures as discussed in the previous paragraph. From this graph we learn three things. Firstly, the pattern of the data is very similar when we compare "complete" and "incomplete" observations, which suggests we need not be concerned about the incomplete cases. Secondly, most points lie to the north west of the 45 degree line, indicating that IL loans tend to be larger than JL loans (an issue we do not explore in this paper). This has been

[45] In 2009, 911 (out of a total of 1106) provide some data on lending methodology by volume coming from the Balance Sheets. Of these, we exclude 154 "village banks" for which lending methodology is unclear. Furthermore, we lose 41 observations due to missing data on the key variables used for the simulation: Portfolio at Risk, Operating Expense, Financial Expense and Real Portfolio Yield. Lastly, we drop one MFI that reports PAR greater than 100%.

previously observed in Cull et al. (2007). Thirdly, although we do observe some lenders offering both IL and JL, the majority of lenders use predominantly one or the other. 72% of lenders (accounting for 68% of loans by number and 84% by value) have 95% of their portfolio in either IL or solidarity lending.

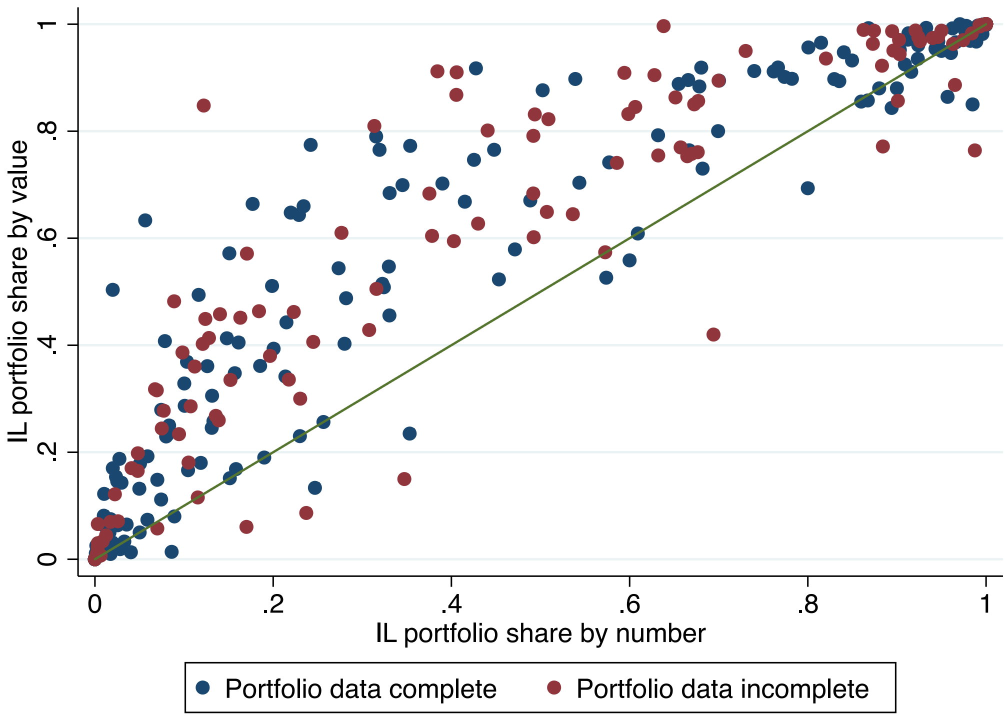

Figure 2.8: IL Share by Value and by Number

2.C.1 Stylized Facts

Construction of Figure 2.1

The figure is based on the cross-section of 1,106 MFIs that reported to the MIX Market dataset in 2009. Time variation comes from the recorded years of establishment. Legal status (for-profit/non-profit) is the status in 2009. "Unweighted" counts the number of institutions in existence at a given date, "weighted" weights institutions by their number of loans outstanding in 2009. Note that there are three key sources of error in this figure. First, institutions choose whether or not to report to MIX, so the sample may not be representative of the universe of MFIs in existence. Secondly there may be a survivor bias as we only observe institutions still in existence in 2009. Thirdly, some institutions may have changed legal status over the period. In each case, to overturn the case for growth in for-profit lending, each source of error would have to be *growing* differentially for for-profits and non-profit lenders.

Calculation of group lending shares by legal status

There are two key weaknesses with these data. First, loans are classified by disbursal method: "individual", "solidarity group" and "village bank". As discussed elsewhere, we only observe whether loans are "individual" or "solidarity group", which might not always correspond to liability structure. To construct the statistics, we compute a measure "Solidarity Group Share" which is the ratio of each lender's number of solidarity group loans to the total of solidarity group and individual loans (we drop "village bank" loans as it is even less clear what the lending method is here). Second, the sample of MFIs that choose to report methodology data (and in particular to report it in every year) is probably non-random; indeed we observe that we have methodology information for more for-profit loans than non-profit loans, even though the for-profit share of all loans is less than 50 percent in the larger sample used to construct Figure 2.1.

The numbers reported in the text are unweighted. Weighting the mean by number of loans we actually find 34 percent of non-profit loans are solidarity group loans, versus 70 percent for for-profits, driven by a few large for-profit group lenders. For example in the balanced panel sample, dropping two outliers, BRAC in Bangladesh and Bandhan in India, we find that 43 percent of non-profit loans are solidarity group loans, and 24 percent of for-profit loans. It is not clear whether the weighted or unweighted means are more informative about the respective behavior of non-profits and for-profits: the theory predicts behavior at the institution level. We also note that our theory predicts that for-profits are more likely to use individual liability, all else equal - these simple averages cannot control for local social capital, the key unobservable in our model.

Chapter 3

Group Lending Without Joint Liability

While joint liability lending by microfinance institutions (MFIs) continues to attract attention as a key vehicle of lending to the poor, recently some MFIs have moved away from explicit joint liability towards individual lending. The most prominent such institutions are Grameen Bank of Bangladesh and BancoSol of Bolivia.[1] However, interestingly, Grameen and others have chosen to retain the regular group meetings that traditionally went hand-in-hand with joint liability lending.

Now it should be pointed out that in the absence of good panel data on lending methods it cannot be conclusively said that there has been a significant overall decline in joint liability among MFIs worldwide just on the basis of various anecdotes about a handful of high-profile MFIs. Indeed, existing evidence suggests that joint liability continues to be widely used. For example, de Quidt et al. (2012) use a sample of 715 MFIs from the MIX Market (Microfinance Information Exchange) database for 2009, and estimate that 54% of loans are made under "solidarity group" lending as opposed to "individual" lending.[2]

Nevertheless, these phenomena raise the question of the costs and benefits of using joint liability, and the choice between group loans with and without (explicit) joint liability. Besley and Coate (1995) is one of the first papers to point out both benefits and costs of joint liability: joint liability can increase repayment rates by inducing borrowers to repay on behalf of their unsuccessful partners but there are also states of the world where an individual borrower may default because of this burden, even if she was willing to pay back her own loan. Using a limited en-

[1] For a discussion of the reasons for the shift in Grameen Bank's lending strategy, see Muhammad Yunus's article "Grameen Bank II: Lessons Learnt Over Quarter of A Century," at `http://www.grameen.com/index.php?option=com_content&task=view&id=30&Itemid=0`, accessed 18 December 2012.

[2] An earlier study Cull et al. (2009) puts this number at 51% using 2002/04 data involving 315 institutions. The year 2009 is one for which the largest cross-section of lending methodologies is available. Solidarity group loans defined by MIX as those for which "some aspect of loan consideration depends on the group, including credit analysis, liability, guarantee, collateral, and loan size and conditions." Individual loans are "made to individuals, and any guarantee or collateral required comes from that individual." We excluded 154 "village banks" for which lending methodology is unclear. See `http://www.mixmarket.org/about/faqs/glossary`.

forcement or "ex-post moral hazard" framework introduced by Besley and Coate (1995) in the group lending context, in this paper we study two issues raised by this apparent shift.

First, we analyze how by leveraging the borrowers social capital, individual liability lending (henceforth, IL) can mimic or even improve on the repayment performance and borrower welfare of explicit joint liability (EJ). When this occurs, we term it "implicit joint liability" (IJ). For this argument to work, there is no need for group lending per se - borrowers can, in theory, sustain this without any explicit effort on the part of the lender. Second, to understand better the logic of group lending, we introduce a purely operational argument for its use under IL, namely, it simply reduce the lender's transactions costs, shifting the burden to the borrowers. This is valuable because lower interest rates relax the borrowers' repayment incentive constraints, increasing repayment and welfare. We then show how this related to first issue: group lending may contribute to the creation of social capital, and therefore, may induce IJ.[3]

Next we carry out some simple simulation exercises using empirically estimated parameters. The goal is to complement the theoretical analysis and to get a quantitative sense of the welfare effects as well as the relevant parameter thresholds that determine which lending method is preferred. Our key findings are as follows. First, in low social capital environments, EJ does quite well compared to IJ. For example, when the standard deviation of project returns of 0.5, for social capital worth 10% of the loan size, the welfare attainable under IJ is 32.4% lower compared to the welfare under EJ. However, with social capital worth 50% of the loan size, the welfare attainable under EJ is 5% lower to the one attainable under IJ. Second, we find that the interest rate, repayment rate and borrower welfare are all rather insensitive to social capital under EJ, whereas in the case of IJ, they are all highly sensitive. This is what we would expect, since the only sanction available under IJ is coming through social capital. Third, when project returns are high variance, the welfare gains from higher social capital are quite large under IJ, which is not the case under EJ. To illustrate consider the case where project returns have a standard deviation of 0.5. If borrowers share social capital worth 10% of the loan size, borrower welfare under IJ is 35.9% lower than that of borrowers who share social capital worth 50% of the loan size.

Our analysis is motivated by two influential recent empirical studies. Giné and Karlan (2011) found that removing the joint liability clause, but retaining the group meetings, of a random subset of borrowing groups of Green Bank in the Philippines had no meaningful effect on repayment rates. In our model, this outcome arises when the newly individually liable groups have sufficient social capital to continue to assist one another with repayments, as under EJ. Secondly, Feigenberg

[3]It could even be that without the group, borrowers would be less able to interact. Indeed, in some conservative societies, social norms may prevent women from attending social gatherings (for instance under the Purdah customs in some parts of India and the Middle East). Then externally mandated borrowing groups can be a valuable vehicle for social interaction. See, for instance Sanyal (2009), Anderson et al. (2002), Kabeer (2005).

et al. (2011) randomly varied the meeting frequency of individually liable borrowing groups of the Village Welfare Society in India. They found that groups who met more frequently had subsequently higher repayment rates. In particular, they present evidence suggesting that this is due to improved informal insurance among these groups due to higher social capital. Both Giné and Karlan (2011) and Feigenberg et al. (2011) find evidence for intra-group transfers to help a borrower repay her loan even without explicit joint liability.[4] We argue that more frequent group meetings give borrowers a stronger incentive to build social capital, and that this is then leveraged to generate IJ. Grameen Bank states that Grameen II is designed to "lean on solidarity groups: small informal groups consisting of co-opted members coming from the same background and trusting each other."[5] The emphasis on trust suggests that the group continues to play an important role in Grameen's lending methodology beyond simply moderating the lender's transaction costs.

The main conclusions of our analysis is that it is premature to write off EJ as a valuable contractual tool and group lending without (explicit) joint liability may still harness some of the benefits of joint liability via implicit joint liability. Thus far we have one high quality randomized study of contractual form (Giné and Karlan (2011)) in which EJ seems not to play an important role. However in our theoretical analysis there are always parameter regions over which EJ is the most efficient of the simple contracts we analyze. A recent randomized control trial by Attanasio et al. (2011) finds stronger consumption and business creation impacts under EJ (albeit no significant difference in repayment rates - note that in their context mandatory group meetings are not used under either IL or EJ). Carpena et al. (2010) analyze an episode in which a lender switched from IL to EJ and found a significant improvement in repayment performance. For the same reasons, Banerjee (2012) stresses the need for more empirical work in the vein of Giné and Karlan (2011) before concluding that EJ is no longer relevant.

It is instructive to briefly look at the types of contracts currently used by MFIs. As mentioned, from the MIX dataset, 54% of borrowers were borrowing under what are classified as solidarity group loans. Although the solidarity group loans might not correspond exactly to pure EJ, this is the best measure we have. Our concept of IJ is most relevant to the "individual" category; the MIX Market notes that "loans based on consideration of the sole borrower, but disbursed through and recollected from group mechanisms, are still considered individual loans." A notable example is the Indian MFI Bandhan, which is one of the top MFIs in India, and is listed as having 3.6m outstanding loans in 2011, all classified as "individual". Bandhan

[4] In table IX of Giné and Karlan (2011) we see that conversion to individual liability caused a decrease, significant at 10%, in side-loans between borrowers, although no significant effect on borrowers "voluntarily [helping others] repay loans". Note that one challenge of interpreting these results in the light of our analysis is that group composition changed in Giné and Karlan (2011)'s experiment, while our model analyzes contract choice for a given level of social capital. Converted centers tended to take in members that were less well-known by existing members, presumably because individual liability made doing so less risky.

[5] See `http://www.grameen-info.org/index.php?option=com_content&task=view&id=33&Itemid=107`, accessed 28th September, 2012.

does not use joint liability but disburses the majority of its loans through borrowing groups. Unfortunately, we do not have data on the method of disbursement of the full sample of loans classified as individual, but it seems likely that many institutions are indeed using groups to disburse individual loans. This paper highlights how this may improve welfare through two channels: first of all, borrowers with sufficient social capital can mutually insure one another and secondly, attending costly group meetings may give borrowers incentives to invest in social capital.

Much of the existing theoretical work has sought to show how explicit joint liability improves repayment rates (see Ghatak and Guinnane (1999) for a review). In the model of Besley and Coate (1995), joint liability gives borrowers an incentive to repay on behalf of their partner when the partner is unable to repay her own loan. If borrowers can threaten social sanctions against one another, this effect is strengthened further. However, there are two problems with EJ. Firstly, since repaying on behalf of a partner will be costly, incentive compatibility requires the lender to use large sanctions and/or charge lower interest rates, relative to individual liability.[6] Secondly, when a borrower is unsuccessful, sometimes EJ induces the successful partner to bail them out, but sometimes it has a perverse effect, inducing them to default completely, while under IL they would have repaid. Rai and Sjöström (2004) and Bhole and Ogden (2010) approach these issues from a mechanism design perspective - designing cross-reporting mechanisms or stochastic dynamic incentives that minimize the sanctions used by the lender. Baland et al. (2010) provide an alternative explanation of the apparent trend away from what we call EJ towards IL, based on loan size. They find that the largest loan offered under IL cannot be supported under joint liability and that the benefits of the latter are increasing in borrower wealth. We do not focus on this angle but briefly touch on the issue of loan size in section 3.1. Allen (2012) shows how partial EJ, whereby borrowers are liable only for a fraction of their partner's repayment, can improve repayment performance by optimally trading off risk-sharing with the perverse effect on strategic default. In contrast, we focus on how simple group lending with no joint liability can achieve some of these effects, as side-contracting by the borrowers can substitute for the lender's enforcement mechanism.

Our model is also related to Rai and Sjöström (2010). In that paper, borrowers are assumed to have sufficient social capital to support incentive-compatible loan guarantees through a side-contract between borrowers, provided they have sufficient information to enforce such side contracts. The role of groups is to provide publicly observable repayment so as to enable efficient side-contracting. In contrast, in our setting, repayment behavior is common knowledge among the borrowers, and it is the amount of social capital that is key. Groups play a role that depends on meeting costs introduced in the next two sections. Secondly, in our model, borrowers are better off when they guarantee one another as their probability of contract

[6] This issue is the focus of the analysis in Rai and Sjöström (2010). Because of this, de Quidt et al. (2012) show that with a for-profit monopolist lender borrowers are better off under EJ than IL lending, because the lender must typically charge lower interest rates under EJ.

renewal is higher. In Rai and Sjöström (2010) this is not the case as the lender is simply assumed to use a punishment that simply imposes a utility cost on the borrowers in case of default. In fact, the optimal contract delivers the same borrower welfare whether they guarantee one another or not.

Other than the above mentioned papers, our paper is also broadly related to the theoretical literature in microfinance that have emerged in the light of the Grameen Bank of Bangladesh abandoning explicit joint liability and switching to the Grameen II model, focusing on aspects other than joint liability, such as sequential lending (e.g., Chowdhury (2005)), frequent repayment (Jain and Mansuri (2003), Fischer and Ghatak (2010)), exploring more general mechanisms than joint liability (e.g., Laffont and Rey (2003)), and exploring market and general equilibrium (Ahlin and Jiang (2008); McIntosh and Wydick (2005) and de Quidt et al. (2012)).

The paper is structured as follows: in section 3.1 we present the basic model where in principle lending may take place with or without group meetings. We introduce our concept of implicit joint liability and show when it will occur and be welfare improving. Section 3.2 formalizes a key transaction cost in group and individual lending - the time spent attending repayment meetings. Section 3.3 then shows how meeting costs can give borrowers incentives to invest in social capital, and shows when this is welfare improving. Section 3.4 presents results of a simulation of the core model, while section 3.5 summarises the results and concludes.

3.1 Model

We model a lending environment characterized by costly state verification and limited liability. Borrowers are risk neutral, have zero outside option, no capital and limited liability. They have access to a stochastic production technology that requires 1 unit of capital per period with expected output $\bar{R}$, and therefore must borrow 1 per period to invest (we assume no savings for simplicity). There are three possible output realizations, $R \in \{R_h, R_m, 0\}$, $R_h \geq R_m > 0$ which occur with positive probabilities p_h, p_m and $1 - p_h - p_m$ respectively. We define the following:

$$\begin{aligned} p &\equiv p_h + p_m \\ \triangle &\equiv p_h - p_m \\ \bar{R} &\equiv p_h R_h + p_m R_m. \end{aligned}$$

We will refer to p as the probability of "success", and $\bar{R}$ as expected output.

We assume that output is not observable to the lender and hence the only relevant state variable from his perspective is whether or not a loan is repaid. Since output is non-contractible, the lender uses dynamic repayment incentives, as in Bolton and Scharfstein (1990). We assume that if a borrower's loan contract is terminated following a default, she can never borrow again. Under individual liability (IL), a borrower's contract is renewed if she repays and terminated otherwise. Under explicit joint liability (EJ), both contracts are renewed if and only if both loans are repaid.

Now we introduce the notion of social capital used in the paper. We assume that pairs of individuals in the village share some pair-specific social capital worth S in discounted lifetime utility, that either can credibly threaten to destroy. In other words, a friendship yields lifetime utility S to each person. If the social capital is destroyed it is lost forever. We assume that each individual has a very large number of friends, each worth S. Thus each friendship that breaks up represents a loss of size S.[7]

We assume a single lender with opportunity cost of funds equal to $\rho > 1$. In the first period, the lender enters the community, observes S and commits to a contract to all potential borrowers. The contract specifies a gross interest rate, r and EJ or IL. We assume the lender to be a non-profit who offers the borrower welfare maximizing contract, subject to a zero-profit constraint.[8]

In this section we ignore the role of groups altogether - being in a group or not has no effect on the information or cost structure faced by borrowers and lenders. Although borrower output is unobservable to the lender, we assume it is observable to other borrowers. As a result, they are able to write informal side contracts to guarantee one another's repayments, conditional on the output realizations. For simplicity, in the theoretical analysis we assume such arrangements are formed between pairs of borrowers.[9]

EJ borrowers will naturally side contract with their partner, with whom they are already bound by the EJ clause. Specifically, we assume that once the loan contract has been fixed, pairs of borrowers can agree a "repayment rule" which specifies each member's repayment in each possible state $Y \in \{R_h, R_m, 0\} \times \{R_h, R_m, 0\}$. Then in each period, they observe the state and make their repayments in a simultaneous-move "repayment game". Deviations from the agreed repayment rule are punished by a social sanction: destruction of S. The repayment rule, social sanction and liability structure of the borrowing contract thus determine the payoffs of the repayment game and beliefs about the other borrower's strategy. To summarize, once the lender has entered and committed to the contract, the timings each period are:

1. Borrowers form pairs, and agree on a repayment rule.

[7]One way to conceptualize S is as the net present value of lifetime payoffs in a repeated "social game" played alongside the borrowing relationship, similar to the multi-market contact literature, such as Spagnolo (1999), who models agents interacting simultaneously in a social and business context, using one to support cooperation in the other. As an illustration, suppose the borrowers play the following "coordination" stage-game each period: if both play A, both receive s. If one plays A and the other, B, both receive -1. If both play B, both receive 0. Clearly, both (A, A) and (B, B) are Nash equilibria in the stage-game. If players expect to play (A, A) forever, their expected payoff is $S \equiv \frac{s}{1-\delta}$. However, switching to (B, B) forever as a social sanction is always a credible threat, and can be used to support the repayment rule.

[8]We abstract from other organizational issues related to non-profits, see e.g. Glaeser and Shleifer (2001).

[9]This could be for example because there are two types of investment project available and returns within a project type are perfectly correlated, such that side-contracting with another borrower who has the same project type yields no benefit. In the simulations we extend the analysis to larger groups.

2. Loans are disbursed, borrowers observe the state and simultaneously make repayments (the repayment game).

3. Conditional on repayments, contracts are renewed or terminated and social sanctions carried out.

4. If an IL borrower's partner was terminated but she repaid, she rematches with a new partner.

We restrict attention to repayment rules that are stationary (depending only on the state) and symmetric (do not depend on the identity of the borrower). This enables us to focus on the stationary value function of a representative borrower. Stationarity also rules out repayment rules that depend on repayment histories, such as reciprocal arrangements. In addition, we assume that the borrowers choose the repayment rule to maximize joint welfare. Welfare maximization implies that social sanctions are never used on the equilibrium path, since joint surplus would be increased by an alternative repayment rule that did not punish this specific deviation.

Given repayment probability π, the lender's profits are:

$$\Pi = \pi r - \rho$$

and therefore the zero-profit interest rate is:

$$\hat{r} \equiv \frac{\rho}{\pi}. \tag{3.1}$$

By symmetry, each borrower i pays $\pi r = \rho$ per period in expectation.

There are two interesting cases that arise endogenously and determine the feasibility of borrowers guaranteeing one another's loans. In Case A $R_m \geq 2r$ and hence a successful borrower can always afford to repay both loans. In Case B we have $R_h \geq 2r > R_m \geq r$, thus it is not feasible for a borrower with output R_m to repay both loans. Case B will turn out to be the more interesting case for our analysis, since in this case there is a cost to using joint liability lending. Specifically there are states of the world (when one borrower has zero output and the other has R_m) in which under joint liability both borrowers will default, since it is not feasible to repay both loans and they will therefore be punished whether or not the successful partner repays her loan. Meanwhile under individual liability, the successful partner is able to repay her loan and will not be punished if she does so.

Consider Case A. If borrowers agree to guarantee one another's loans, they will repay in every state except $(0,0)$, so the repayment probability is $\pi = 1 - (1-p)^2 = p(2-p)$, in which case $\hat{r} = \frac{\rho}{p(2-p)}$. Therefore Case A applies if $R_m \geq \frac{2\rho}{p(2-p)}$, i.e. when the successful partner can afford to repay both loans even if her income is only R_m. If this condition does not hold, then it will not be feasible for the successful borrower to help her partner in this state of the world, and therefore Case B applies.

Definition 1 *Case A applies when* $R_m \geq \frac{2\rho}{p(2-p)}$. *Case B applies when* $R_m < \frac{2\rho}{p(2-p)}$.

Suppose that borrowers only repay when both are successful, i.e. when both have at least R_m, which occurs with probability p^2. If this is the equilibrium repayment rate, then $\hat{r} = \frac{\rho}{p^2}$. We make a simple parameter assumption that ensures that this will be the highest possible equilibrium interest rate (lowest possible repayment rate), by ensuring that even with income R_m, borrowers can afford to repay $\frac{\rho}{p^2}$.

Assumption 5 $R_m \geq \frac{\rho}{p^2}$.

We also assume that R_h is sufficiently large that a borrower with R_h could afford to repay both loans even at interest rate $\hat{r} = \frac{\rho}{p^2}$. Since this is the highest possible equilibrium interest rate, this implies that R_h is always sufficiently large for a borrower to repay both loans.

Assumption 6 $R_h \geq 2\frac{\rho}{p^2}$.

To summarize, together these assumptions guarantee that $R_m \geq r$ and $R_h \geq 2r$ on the equilibrium path.

We can now write down the value function V for the representative borrower, which represents the utility from access to credit. Suppose that borrower i's loan is repaid with some probability π. Since the repayment rule is assumed to maximize joint welfare, it follows that borrowers's loans are only repaid when repayment leads to the loan contracts being renewed, and therefore the representative borrower's contract is also renewed with probability π. Since the lender charges zero profit interest rate $\hat{r} = \frac{\rho}{\pi}$, the borrower repays $\pi\hat{r} = \rho$ in expectation. Hence, her welfare is:

$$\begin{aligned} V &= \bar{R} - \rho + \delta\pi V \\ &= \frac{\bar{R} - \rho}{1 - \delta\pi}. \end{aligned} \tag{3.2}$$

For any borrower to be willing to repay her loan, it must be that the value of access to future loans exceeds the interest rate, or $\delta V \geq r$. If this condition does not hold, all borrowers will default immediately. We refer to this condition as Incentive Condition 1 (IC1), and it must hold under any equilibrium contract.

Provided IC1 is satisfied, borrower welfare is maximized by achieving the highest repayment rate possible. To see this, suppose the lender charges some interest rate r. Then $V = \frac{\bar{R} - \pi r}{1 - \delta\pi}$. It can be verified that this is increasing in π if and only if IC1 holds. Therefore, in the subsequent discussion the ranking of welfare will be equivalent to the ranking in terms of the repayment probability.

Using (3.2) and $\hat{r} = \frac{\rho}{\pi}$ we can derive the equilibrium IC1 explicitly:

$$\rho \leq \delta\pi\bar{R}. \tag{IC1}$$

By Assumption 5, the lowest possible equilibrium repayment probability π is equal to p^2. For the theoretical analysis we make the following parameter assumption that ensures IC1 is satisfied in equilibrium:

Assumption 7 $\delta p^2 \bar{R} > \rho$.

Now that the model is set up we analyze the choice of contract type.

3.1.1 Individual Liability

Suppose first of all that the borrower does not reach a repayment guarantee arrangement with a partner. Since IC1 is satisfied, the borrower will repay her own loan whenever she is successful, so her repayment probability is p. Her utility V is then equal to $\frac{\bar{R}-\rho}{1-\delta p}$.

Now we consider when pairs of IL borrowers will agree a repayment guarantee arrangement. If this occurs, we term it implicit joint liability (IJ).

Since IC1 holds, the borrowers want to agree a repayment rule that maximizes their repayment probability. There are many possible such rules that can achieve the same repayment rate, so for simplicity we focus on the most intuitive one, whereby borrowers agree to repay their own loan whenever they are successful, and also repay their unsuccessful partner's loan if possible.[10]

We already know that repayment of the borrower's own loan is incentive compatible by IC1. For it to be incentive compatible for her to repay on behalf of her partner as well, it must be that social sanction outweighs the cost of the extra repayment, i.e. $r \leq \delta S$. This gives us a constraint which we term IJ Incentive Constraint 2, or IJ IC2. For equilibrium interest rate $\hat{r} = \frac{\rho}{\pi^{IJ}}$ IJ IC2 reduces to:

$$\rho \leq \delta \pi^{IJ} S. \tag{IJ IC2}$$

There is a threshold value of S, $\hat{S}^{IJ}$, such that IJ IC2 holds for $S \geq \hat{S}^{IJ}$:

$$\hat{S}^{IJ}_k \equiv \frac{\rho}{\delta \pi^{IJ}_k}, k \in \{A, B\},$$

where k denotes the relevant case. When $S \geq \hat{S}^{IJ}$, it is feasible and incentive compatible for borrowers to guarantee one another's loans, and therefore they will do so as this increases the repayment probability and thus joint welfare. Therefore IJ applies for $S \geq \hat{S}^{IJ}$.

Next we work out the equilibrium repayment probabilities and interest rates in cases A and B respectively. Assume $S \geq \hat{S}^{IJ}$. In Case A, a successful borrower can always afford to repay both loans, so both loans are repaid with probability $\pi^{IJ}_A \equiv 1 - (1-p)^2 = p(2-p)$. In Case B, both loans are repaid whenever both

[10] An example of an alternative, less intuitive rule that can sometimes achieve the same repayment rate but cannot do better is where borrowers agree to repay their partner's loan, and then repay their own as well if they can afford to do so.

are successful, and in states $(R_h, 0), (0, R_h)$. In state $(R_m, 0)$, borrower 1 cannot afford to repay borrower 2's loan, so she repays her own loan, while borrower 2 defaults and is replaced in the next period with a new partner. Therefore $\pi_B^{IJ} \equiv p^2 + 2p_h(1-p) + p_m(1-p) = p + p_h(1-p)$. Notice that both π_A^{IJ} and π_B^{IJ} are greater than p.

The lender observes whether Case A or Case B applies, and the value of S in the community, and offers an individual liability contract at the appropriate zero profit interest rate. Equilibrium borrower welfare under individual liability is equal to:

$$V_k^{IL}(S) = \begin{cases} \frac{\bar{R}-\rho}{1-\delta p} & S < \hat{S}_k^{IJ} \\ \frac{\bar{R}-\rho}{1-\delta \pi_k^{IJ}} & S \geq \hat{S}_k^{IJ} \end{cases}, k \in \{A, B\}.$$

It is straightforward to see that as S switches from less than $\hat{S}_k^{IJ}$ to greater than or equal to it, $V_k^{IL}(S)$ goes up as $\pi_k^{IJ} > p$.

3.1.2 Explicit Joint Liability

Now we analyze EJ contracts. Recall that under EJ, a pair of borrowers are offered a contract such that unless both loans are repaid, both partners lose access to credit in the future. The advantage of this contractual form is that it gives additional incentives to the borrowers to guarantee one another's loans. However, the disadvantage is that when borrower i is successful and j is unsuccessful, there may be states in which borrower i would repay were she under individual liability, but she will default under joint liability because she is either unwilling or unable to repay both loans.

The borrowers will agree a repayment rule, just as under IJ. Since this will be chosen to maximize joint welfare, it will only ever involve either both loans being repaid or both defaulting, due to the joint liability clause that gives no incentive to repay only one loan. Subject to this, because IC1 holds, joint welfare is maximized by ensuring both loans are repaid as frequently as possible.

IC1 implies that when both borrowers are successful, they will both be willing to repay their own loans. We therefore need to consider i's incentive to repay both loans when j is unsuccessful. Borrower i will be willing to make this loan guarantee payment provided the threat of termination of her contract, plus the social sanction for failing to do so, exceeds the cost of repaying two loans. Formally, this requires $2r \leq \delta(V^{EJ} + S)$. We refer to this condition as EJ IC2. Rearranging, and substituting for $\hat{r} = \frac{\rho}{\pi^{EJ}}$, we obtain:

$$\rho \leq \frac{\delta \pi^{EJ}[\bar{R} + (1 - \delta \pi^{EJ})S]}{2 - \delta \pi^{EJ}}. \quad \text{(EJ IC2)}$$

We can derive a threshold, $\hat{S}^{EJ}$, such that EJ IC2 is satisfied for $S \geq \hat{S}^{EJ}$:

$$\hat{S}_k^{EJ} \equiv \max \left\{ 0, \frac{\rho}{\delta \pi_k^{EJ}} - \frac{\delta \pi_k^{EJ} \bar{R} - \rho}{\delta \pi_k^{EJ}(1 - \delta \pi_k^{EJ})} \right\}, k \in \{A, B\}$$

where as before, k denotes the relevant Case.

Note that $\hat{S}^{EJ}$ can be equal to zero. This corresponds to the basic case in Besley and Coate (1995) where borrowers can be induced to guarantee one another even without any social capital. This relies on the lender's use of joint liability to give borrowers incentives to help one another, and is not possible under individual liability.

Provided $S \geq \hat{S}^{EJ}$, borrowers are willing to guarantee one another's repayments. The repayment rule will then specify that i repays on j's behalf whenever i can afford to and j is unsuccessful. If $S < \hat{S}^{EJ}$, borrowers will not guarantee one another. They will therefore only repay when both are successful.

We now derive the equilibrium repayment probability under each Case. Firstly, if $S < \hat{S}^{EJ}$, borrowers repay only when both are successful, so $\pi^{EJ} = p^2$ in either Case.

Now suppose $S \geq \hat{S}^{EJ}$. In Case A, both loans can be repaid whenever at least one borrower earns at least R_m. Thus the repayment probability is $\pi_A^{EJ} = p(2-p)$. In Case B, R_m is not sufficient to repay both loans. Therefore both loans are repaid in all states except $(0,0), (R_m, 0), (0, R_m)$. In these three states both borrowers default. The repayment probability is therefore $\pi_B^{EJ} = p^2 + 2p_h(1-p) = p + \triangle(1-p)$.

Borrower welfare is:

$$V_k^{EJ}(S) = \begin{cases} \frac{\bar{R}-\rho}{1-\delta p^2} & S < \hat{S}_k^{EJ} \\ \frac{\bar{R}-\rho}{1-\delta \pi_k^{EJ}} & S \geq \hat{S}_k^{EJ} \end{cases}, k \in \{A, B\}.$$

Note that $\hat{S}_A^{EJ} \leq \hat{S}_B^{EJ}$. This is because the interest rate is lower in Case A, and V is higher (due to the higher renewal probability), so the threat of termination is more potent.

Now that we have derived the equilibrium contracts assuming either IL or EJ, we turn to analyzing the lender's choice of contractual form in equilibrium, which will depend crucially on the borrowers' ability to guarantee one another's loans.

Let us define $V(S) \equiv \max\{V^{EJ}(S), V^{IL}(S)\}$ as the maximum borrower welfare from access to credit. Observe that the repayment probability and borrower welfare from access to credit, $V(S)$, are stepwise increasing in S.

3.1.3 Comparing contracts

In this section we compare borrower welfare under each contractual form. We have seen that EJ has the advantage that it may be able to induce borrowers to guarantee one another even when they have no social capital. However, in Case B it has a perverse effect: in some states of the world borrowers will default when they would have repaid under IL.

This is most acute when $p_m > p_h$. Then $\pi_B^{EJ} = p + \triangle(p_h - p_m) < p$. Therefore in Case B, EJ actually performs worse than IL for all levels of social capital - the perverse effect dominates. Thus for Case B, EJ would never be offered.

We have already derived thresholds for S, $\hat{S}^{IJ}$ and $\hat{S}^{EJ}$, above which borrowers will guarantee one another's loans under individual and joint liability respectively. The lender's choice of contract will depend on the borrowers ability to do so, so first we derive a lemma that orders these thresholds in Case A and Case B.

Lemma 1

1. $\hat{S}^{IJ}_A > \hat{S}^{EJ}_A$.
2. *Suppose* $p_h \geq p_m$. *Then* $\hat{S}^{IJ}_B > \hat{S}^{EJ}_B$.

Proof. See appendix. ■

Lemma 1 shows that supporting a loan guarantee arrangement requires more social capital under IL than under EJ. The reason for this is that the lender's sanction under EJ is a substitute for social capital in providing incentives to borrowers to guarantee one another.[11]

The lender is a non-profit who offers the borrower welfare-maximizing contract. Therefore he offers IL if $V^{EJ}(S) \leq V^{IL}(S)$ and EJ otherwise. This will depend on the Case (A or B), the sign of $\triangle$, and S. We summarize the key result of this section as:

Proposition 11 *The contracts offered in equilibrium are as follows:*

Case A: *IL is offered at* $\hat{r} = \frac{\rho}{p}$ *for* $S < \hat{S}^{EJ}_A$, *otherwise EJ is offered at* $r = \frac{\rho}{\pi^{EJ}_A}$.

Case B, $\triangle > 0$***:*** *IL is offered at* $\hat{r} = \frac{\rho}{p}$ *for* $S < \hat{S}^{EJ}_A$, *EJ is offered at* $\hat{r} = \frac{\rho}{\pi^{EJ}_B}$ *for* $S \in [\hat{S}^{EJ}_B, \hat{S}^{IJ}_B)$, *IL is offered at* $\hat{r} = \frac{\rho}{\pi^{IJ}_B}$ *for* $S \geq \hat{S}^{IJ}_B$.

Case B, $\triangle \leq 0$***:*** *IL is offered at* $\hat{r} = \frac{\rho}{p}$ *for* $S < \hat{S}^{IJ}_B$, *IL is offered at* $\hat{r} = \frac{\rho}{\pi^{IJ}_B}$ *otherwise.*

Whenever EJ is offered borrowers guarantee one another's repayments. Whenever IL is offered and $S \geq \hat{S}^{IJ}$ *borrowers guarantee one another's repayments.*

Proof. See appendix. ■

The result is summarized in Table 3.1, which gives the equilibrium contract and repayment probability π in alternate rows. Borrower welfare is not shown, but is easily computed as $V = \frac{pR-\rho}{1-\delta\pi}$, is strictly increasing in π.

This table shows that there are clear trade-offs in the contractual choice. In Case A, IJ has no advantage over EJ because in both cases borrowers repay both loans whenever successful. Therefore IL is offered for low S, and EJ for high S. In Case B when $\triangle \leq 0$, we have already remarked that EJ is always dominated by IL.

[11] A slight complication arises in the proof because in Case B the repayment probability is higher and therefore the interest payment is lower under IJ. As a result, the size of the guarantee payment that must be incentive compatible is actually smaller under IJ, but the net effect is still that borrowers are more willing to guarantee one another under EJ.

<table>
<tr><td></td><td>Case A</td><td>Case B, $\triangle > 0$</td><td>Case B, $\triangle \leq 0$</td></tr>
<tr><td>$S < \hat{S}^{EJ}$</td><td>IL (no IJ)
p</td><td>IL (no IJ)
p</td><td rowspan="2">IL (no IJ)
p</td></tr>
<tr><td>$S \in [\hat{S}^{EJ}, \hat{S}^{IJ})$</td><td>EJ
$p(2-p)$</td><td>EJ
$p + \triangle(1-p)$</td></tr>
<tr><td>$S \geq \hat{S}^{IJ}$</td><td>EJ
$p(2-p)$</td><td>IL (with IJ)
$p + p_h(1-p)$</td><td>IL (with IJ)
$p + p_h(1-p)$</td></tr>
</table>

Table 3.1: Equilibrium contracts and repayment probabilities

Therefore basic IL is offered for low S, and when S is high enough, borrowers will begin to guarantee one another, leading to an increase in the repayment rate and a fall in the equilibrium interest rate.

The most interesting case is Case B for $\triangle > 0$. Here there is a clear progression as S increases. For low S, borrowers cannot guarantee one another under either contract, so basic IL is offered. For intermediate S, EJ can sustain a loan guarantee arrangement but IL cannot, so EJ is offered. Finally for high S, borrowers are able to guarantee one another under IL as well. Since this avoids the perverse effect of EJ, the lender switches back to IL lending.

3.1.4 A remark on loan size

For simplicity, our core model assumes loans of a fixed size. However we can allow for variable loan size as a simple extension. To keep things simple, we assume that borrowers require a loan of size L. The relation between loan size and output is linear, that is, with a loan of size L, output is LR_h with probability p_h, LR_m with probability p_m, and 0 otherwise. Therefore we can simply scale $\bar{R}$ and r by L, so borrower welfare is now equal to LV. However, borrowers' social capital is derived from relationships external to the production function and therefore is assumed not to depend on L. Thus for a given amount of social capital S, borrowers are less willing to guarantee one another's loans as the loan size increases.[12] Thus we have the following observation:

Observation 4 *$\hat{S}^{EJ}(L)$ and $\hat{S}^{IJ}(L)$ are increasing in loan size, L. For a given S borrowers are less likely to guarantee one another's repayments as loan sizes increase. The repayment probability is thus decreasing in L.*

Note that the region $L(\hat{S}^{IJ} - \hat{S}^{EJ})$ is increasing in L. In particular, as L increases, the region $[0, \hat{S}^{EJ})$ expands. Over this region, borrowers are receiving "basic" IL, and not guaranteeing one another. Thus this result suggests a simple intuition for the stylized fact that IL loans tend to be larger. When loan sizes are small, the borrowers' social capital can be tapped to smooth out occasional small imbalances in income. As loan sizes and incomes increase, this becomes less feasible. As borrowers become

[12] Formally, the IJ IC2 is $Lr \leq \delta S$ and the EJ IC2 is $Lr \leq \delta(LV + S)$. Both are tighter as L increases. Replacing $\bar{R}$ with $L\bar{R}$, we observe that $\hat{S}^{EJ}(L) = L\hat{S}^{EJ}$ and $\hat{S}^{IJ}(L) = L\hat{S}^{IJ}$.

unwilling to guarantee one another's loans, EJ becomes unattractive as it induces the borrowers to default unless both are successful.[13]

3.1.5 Discussion

Borrowers form partnerships that optimally leverage their social capital to maximize their joint repayment probability. Thus when social capital is sufficiently high to generate implicit joint liability, IL lending can dominate EJ: borrower i no longer defaults in state $(R_m, 0)$. This does not however mean there is no role for EJ. In particular, for intermediate levels of social capital, EJ can dominate IL - social capital is high enough for repayment guarantees under EJ but not under IL. We analyze borrower welfare under EJ and IL/IJ quantitatively in the simulations.

The results of Giné and Karlan (2011) are consistent with our Case A. Here, IL and EJ lending can achieve the same repayment probability, provided S is sufficiently high. This does not imply that those same borrowers would repay as frequently if they were not able to side-contract. Giné and Karlan (2011) additionally find that borrowers with weak social ties are more likely to default after switching to IL lending - this is consistent with these borrowers having $\hat{S}^{EJ} \leq S < \hat{S}^{IJ}$, so they are unable to support implicit joint liability.

So far, we have ignored the use of groups for disbursal and repayment of loans. However, it is frequently argued (see e.g. Armendáriz de Aghion and Morduch (2010)) that group meetings generate costs that differ from those under individual repayment. In the next section we show that this may induce the lender to prefer one or the other. We then proceed to show that by interacting with the benefits from social capital, group meetings may induce the creation of social capital. This is consistent with the results of a field experiment by Feigenberg et al. (2011).

3.2 Meeting Costs

In this section we lay out a simple model of loan repayment meeting costs. This immediately suggests a motivation for the use of groups. Holding group repayment meetings shifts the burden of meeting costs from the lender to the borrowers. This enables the lender to reduce the interest rate, which in turn makes it easier for borrowers to guarantee one another. Then in the next section we explore how the use of groups might create social capital, and thus generate implicit joint liability.

Since we want to focus on the interplay between meeting costs and social capital under individual liability, we assume that Case B applies and $\triangle \leq 0$. Therefore we can ignore EJ and drop the A, B notation.

A common justification for the use of group meetings by lenders is that it minimizes transaction costs. Meeting with several borrowers simultaneously is less time-

[13] Baland et al. (2010) obtain a result that gives the same negative correlation between the use of IL and loan size. Our above result is different in a nuanced way. In their model the poorest borrowers need the largest loan. Hence, their model generates a positive correlation between loan size and poverty.

consuming than meeting with each individually. However, group meetings might be costly for the borrowers, as they take longer and are less convenient than individual meetings. We term IL lending to groups ILG and IL lending to individuals ILI.

We assume that loan repayment meetings have two components, each of which takes a fixed amount of time. For simplicity, we assume that the value of time is the same for borrowers and loan officers[14] Also, for simplicity, we assume that the cost of borrower time is non-monetary so that borrowers are able to attend the meeting even if they have no income. However, more time spent in meetings by the loan officer increases monetary lending costs, for example because more staff must be hired.

Each meeting incurs a fixed and variable cost. The fixed cost includes travel to the meeting location (which we assume to be the same for borrower and loan officer for simplicity), setting up the meeting, any discussions or advice sessions that take place at the meeting, reminding borrowers of the MFI's policies, and so on. This costs each borrower and the loan officer an amount of time worth γ_f irrespective of the number of borrowers in the group. Secondly there is a variable cost that depends on the number of borrowers at the meeting. This time cost is worth γ_v per borrower in the meeting. This covers tasks that must be carried out once for each borrower: collecting and recording repayments and attendance, reporting back on productive activities, rounding up missing borrowers, and so on. As with the fixed cost, each borrower and the loan officer incurs the variable cost. We assume that for group loans, each borrower also has to incur the cost having to sit through the one-to-one discussion between the loan officer and the other borrower, i.e., in a two group setting, the total variable cost per borrower is $2\gamma_v$ whereas under individual lending, it is γ_v.

Therefore, in a meeting with one borrower, the total cost incurred by the loan officer is $\gamma_f + \gamma_v$, and the total cost incurred by the borrower is the same, bringing the aggregate total time cost of the meeting to $2\gamma_f + 2\gamma_v$. In a meeting with two borrowers the loan officer incurs a cost of $\gamma_f + 2\gamma_v$, and similarly for the borrowers. Thus the aggregate cost in this case is $3\gamma_f + 6\gamma_v$. The lender's cost of lending per loan under ILI is $\rho + \gamma_f + \gamma_v$. Under ILG it is $\rho + \frac{\gamma_f}{2} + \gamma_v$. Therefore the corresponding zero-profit interest rates are $\hat{r}^{ILI} \equiv \frac{\rho+\gamma_f+\gamma_v}{\pi}$ and $\hat{r}^{ILG} \equiv \frac{\rho+\frac{\gamma_f}{2}+\gamma_v}{\pi}$.

Accounting for these costs, per-period expected utility for borrowers under ILI is $\bar{R} - \rho - 2(\gamma_f + \gamma_v)$. Under ILG, the per-period utility is $\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)$.[15]

Of course, the first thing to check is whether one lending method is less costly than the other in the absence of any loan guarantee arrangement between borrowers. This is covered by the following observation:

[14]This may not be too unrealistic. For example, the large Indian MFI, Bandhan, deliberately hires loan officers from the communities that they lend to.

[15]We need to adapt Assumptions 5, 6 and 7 to reflect the additional costs. We assume $R_m \geq \frac{\rho+\frac{1}{2}(\gamma_f+2\gamma_v)}{p^2}$, $R_h \geq 2\frac{\rho+\frac{1}{2}(\gamma_f+2\gamma_v)}{p^2}$, $\delta p^2 \bar{R} - \max\left\{(1+\delta p^2)(\gamma_f+\gamma_v), \left(\frac{1}{2}+\delta p^2\right)(\gamma_f+2\gamma_v)\right\} \geq \rho$.

Observation 5 *Suppose $S = 0$. The lender uses ILG if and only if $\gamma_v < \frac{\gamma_f}{2}$.*[16]

The intuition is straightforward. When $\frac{\gamma_f}{\gamma_v}$ is large, i.e., fixed costs are important relative variable costs (e.g., when a large part of repayment meetings is repetitious) it is economical to hold group meetings. However, the more time is spent on individual concerns, the more costly it is to the borrowers to have to attend repayment meetings in groups because they have to sit through all the bilateral exchanges between another borrower and the loan officer. Microfinance loans are typically highly standardized and so $\frac{\gamma_f}{\gamma_v}$ will be relatively large, which is consistent with the common usage of group lending methods in microfinance.

Now consider borrowers' incentives to guarantee one another's loans. First we observe that for a given γ_v, γ_f, half of the aggregate meeting cost per borrower is borne by the lender under ILI, while only a third is borne by the lender under ILG. The lender passes on all costs through the interest rate, so inspecting the value functions suggests that it is innocuous upon whom the cost of meetings falls. In fact this is not the case. Consider once again IJ IC2: $r \leq \delta S$. The only benefit a borrower receives from bailing out her partner is the avoidance of a social sanction, while the cost depends on the interest payment she must make. Therefore a lending arrangement in which the lender bears a greater share of the costs, and thus must charge a higher interest rate, tightens IJ IC2. This gives us the next proposition, which is straightforward:

Proposition 12 *Borrowers are more likely to engage in IJ under group lending than individual lending: $\hat{S}^{IJG} < \hat{S}^{IJI}$.*[17]

The implication of this result is that there is a trade-off between minimizing total meeting costs, and minimizing those costs borne by the lender. It may actually not be optimal to minimize total costs as shown by the following corollary, the proof of which is straightforward and given in the appendix. This arises from the fact that in an environment where the borrowers' participation constraints are not binding, the lender does not put weight on the disutility costs of meetings (individual or group) to the borrowers.

Corollary 1 *Suppose $S \in [\hat{S}^{IJG}, \hat{S}^{IJI})$. Borrower welfare under ILG may be higher than under ILI, even if $\gamma_v > \frac{\gamma_f}{2}$.*

We have now set the stage to analyze the interaction between meeting costs and social capital.

[16]Proof: $S = 0$ implies IJ is not possible so $\pi = p$ under ILI and ILG. The result then follows from comparison of per-period borrower welfare.

[17]Proof: Borrowers are willing to guarantee their partner's repayments provided $r \leq \delta S$. Plugging in for the interest rates under ILG and ILI, we obtain $\hat{S}^{IJG} = \frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{\delta \pi^{IJ}} < \frac{\rho + \gamma_f + \gamma_v}{\delta \pi^{IJ}} = \hat{S}^{IJI}$.

3.3 Social capital creation

In this section we show how group lending can actually generate social capital that is then used to sustain IJ. This analysis is motivated by the findings of Feigenberg et al. (2011). In their experiment, borrowers who were randomly assigned to higher frequency repayment meetings went on to achieve higher repayment rates. The authors attribute this to social capital being created by frequent meetings, social capital which can then support mutual insurance.

We show two main results. Firstly, group lending may create social capital where individual lending does not. The reason is simply that forcing the borrowers to spend time together in group meetings gives them an added incentive to invest in getting to know one another, as this makes the time spent in group meetings less costly. The knock-on effect is then that individual liability in groups may outperform individual liability with individual meetings because the groups are creating social capital that is then being used to support IJ.

Secondly, we turn to a comparative static more closely related to the Feigenberg et al. (2011) finding. Our simple framework does not easily allow us to model varying meeting frequency, so instead we study the effect on social capital creation of increasing the meeting costs (γ_f or γ_v). We find that an increase in the amount of time spent in group meetings can induce borrowers to switch to creating social capital, and can in fact be welfare-increasing.

Suppose that initially borrowers do not have any social capital, because creating social capital is too costly. For example, borrowers must invest time and effort in getting to know and understand one another, extend trust that might not be reciprocated, and so forth. Assume that social capital can take two values only, 0 and $S > 0$ and for a pair to generate social capital worth S in utility terms, each must make a discrete non-monetary investment that costs them η. To make the analysis interesting, we assume that in the absence of microfinance, they prefer not to do so, namely, $\eta > S$.

Once we introduce group lending, social capital generates an indirect benefit, by enabling the formation of a guarantee arrangement. This may or may not be sufficient to induce them to make the investment - that would depend on how $\eta - S$ compares with the insurance gains from

.[18]

Suppose the lender offers ILI and S is sufficiently large to sustain IJ. If the borrowers prefer to invest in social capital, each time their partner defaults they must invest in social capital with their new partner. We obtain the following result:

Lemma 2 *Borrowers will not invest in social capital under ILI if:*

$$\eta - S > G_1. \tag{3.3}$$

[18]Note that each time a borrower's partner defaults and is replaced, she must invest in social capital with the new partner in order to continue with IJ.

where

$$G_1 \equiv \frac{p_h(1-p)\left[\delta\left(\bar{R}-\frac{\rho}{\pi^{IJ}}\right)-\frac{1+\delta\pi^{IJ}}{\pi^{IJ}}(\gamma_f+\gamma_v)\right]}{(1-\delta p)(1-\delta(p+\triangle(1-p)))}.$$

The proof is given in the appendix. The greater the welfare gain from insurance, the higher is G_1 so the more likely the borrowers will invest in social capital. If (3.3) holds, the only equilibrium under ILI is one in which the borrowers do not invest in social capital, and therefore are not able to guarantee one another's loans.

Now assume that under ILG, the per-meeting cost to borrowers is decreasing in S. Attending group meetings is a chore unless the other group members are friends, in which case it can be a social occasion. By forcing the borrowers to meet together, the lender might give them an incentive to create social capital, benefiting them.

For simplicity, we assume that the cost to the borrowers of the time spent in group meetings is $(1-\lambda(S))(\gamma_f+2\gamma_v)$. In particular, $\lambda(0)=0$ and $\lambda(S)=\lambda>0$. The larger is λ, the smaller the disutility of group meetings, and when $\lambda>1$, borrowers actually derive positive utility from group meetings that is increasing in the length of the meeting. We can now check when social capital will be created in groups.

Lemma 3 *Borrowers invest in social capital under ILG if:*

$$\eta - S \leq G_2. \tag{3.4}$$

where

$$G_2 \equiv \frac{p_h(1-p)\left[\delta\left(\bar{R}-\frac{\rho}{\pi^{IJ}}\right)-\frac{1+2\delta\pi^{IJ}}{2\pi^{IJ}}(\gamma_f+2\gamma_v)\right]+\lambda(1-\delta p)(\gamma_f+2\gamma_v)}{(1-\delta p)(1-\delta(p+\triangle(1-p)))}.$$

The proof is given in the appendix. The greater the welfare gain from insurance, the higher is G_2, but in addition, G_2 is increasing in λ, which represents the reduction in the cost of attending group meetings when the borrowers have social capital. The larger is G_2, the more likely borrowers are to invest in social capital.

Lemmas 2 and 3 suggest that there may exist an interval, $(G_1, G_2]$ for $\eta - S$ over which groups create social capital but individual borrowers do not. The condition for this to be the case is derived in the next proposition, which follows from straightforward comparison of (3.3) and (3.4):

Proposition 13 *If the following condition holds:*

$$\lambda > \frac{p_h(1-p)(\delta\pi^{IJ}\gamma_v-\frac{\gamma_f}{2})}{4\pi^{IJ}(1-\delta p)(\gamma_f+2\gamma_v)} \tag{3.5}$$

then there exists a non-empty interval for $\eta - S$ over which both (3.3) *and* (3.4) *are satisfied. If $\eta - S$ lies in this interval, groups create social capital, and individual lending does not.*

This is a key result, as it shows that when creating social capital sufficiently offsets the cost to borrowers of attending group meetings, borrowing groups may create social capital and guarantee one another's loans, while individual borrowers may not do so. We can see that the threshold for λ in (3.5) is negative if $\frac{\gamma_f}{2} > \gamma_v > \delta\pi^{IJ}\gamma_v$ and so the condition (3.5) is always satisfied if group lending has a cost advantage to the lender. What can be checked is, even if this is not the case, and $\delta\pi^{IJ}\gamma_v - \frac{\gamma_f}{2} > 0$ the critical threshold for λ is always strictly less than 1 and therefore, there always exists a λ high enough (but strictly less than 1) such that the condition (3.5) would hold. However it does not yet establish that the use of groups is necessarily welfare-improving. In other words, observing that groups are bonding and creating social capital does not tell the observer that group lending is the welfare-maximizing lending methodology. All it tells us is that investment is preferred to no investment under ILG, and no investment is preferred to investment under ILI. The welfare ranking of these two will depend on the meeting costs, η and S. The following proposition addresses the welfare question.

Proposition 14 *Suppose condition* (3.5) *is satisfied and* $\eta - S \in (G_1, G_2]$*. Borrower welfare under ILG is higher than that under ILI if:*

$$\eta - S \leq G_3 \tag{3.6}$$

where

$$G_3 \equiv \frac{\delta p_h(1-p)\,(\bar{R}-\rho) + 2(1-\delta\pi^{IJ})(\gamma_f+\gamma_v) - \frac{1}{2}(1-\delta p)(\gamma_f+2\gamma_v)(3-2\lambda)}{(1-\delta p)(1-\delta(p+\triangle(1-p)))}.$$

The proof is given in the appendix. G_3 is higher the larger is the meeting cost under ILI relative to under ILG. It is also increasing in λ, representing the reduction in the cost of attending group meetings when the borrowers have social capital. Note that (3.6) is always satisfied for sufficiently large λ.

The expressions G_1, G_2 and G_3 are somewhat unwieldy. The following proposition establishes a sufficient condition under which $G_1 < G_2 < G_3$, i.e. there is guaranteed to exist an interval for $\eta - S$ over which groups invest in social capital and individuals do not, and over which borrower welfare is higher under group than individual lending:

Proposition 15 *Suppose total meeting costs per borrower are weakly lower under ILG than ILI, i.e.* $\gamma_v \leq \frac{\gamma_f}{2}$*. Then* $G_1 < G_2 < G_3$*, i.e.:*

1. *There always exists an interval for* $\eta - S$ *over which groups create social capital and individuals do not.*

2. *Borrower welfare is weakly higher under ILG than ILI for all values of* $\eta - S$*.*

The proof is given in the appendix. The condition $\gamma_v \leq \frac{\gamma_f}{2}$ implies that ILG has a (weak) cost advantage over ILI, as was discussed in Observation 5. In addition, when $G_1 < \eta - S \leq G_2$, groups invest in social capital while individuals do not, and this gives ILG a further advantage.

3.3.1 Meeting frequency and social capital creation

Now we take this basic framework and carry out one particular comparative-static exercise, motivated by the findings of Feigenberg et al. (2011). They find that groups that were randomly assigned to meet more frequently have better long-run repayment performance, which they attribute to higher social capital and informal insurance within the group. It is not possible to model repayment frequency in our simple setup, but nevertheless our model is able to capture some of this intuition.

We model an increase in meeting frequency as an increase in meeting costs, represented by an increase in either γ_f or γ_v. The more time spent in group meetings, the greater the benefit from social interaction within those meetings, captured by λ. Intuitively, it may not be too costly to attend meetings once a month with a stranger, but the more frequent those meetings are, the greater the incentive the borrowers have to build social capital.

However, more frequent meetings require more of the loan officer's time as well, leading to higher lending costs and a higher interest rate. This reduces the borrowers' incentive to invest in S, since the higher meeting costs reduce the value of maintaining access to credit.

The net effect on borrowers willingness to invest in S is positive if λ is sufficiently large, as shown by the following proposition.

Proposition 16 *Increases in γ_f or γ_v make borrowers under group lending more willing to invest in social capital if and only if the following condition holds:*

$$\lambda > \frac{p_h(1-p)(1+2\delta\pi^{IJ})}{2\pi^{IJ}(1-\delta p)}. \tag{3.7}$$

The proof is immediate from inspection of (3.4). This proposition implies an interesting corollary: an increase in meeting costs can actually be welfare-improving, by inducing borrowers to invest in social capital and thus engage in implicit joint liability.

Corollary 2 *Suppose* (3.7) *holds. Then there exists a threshold at which increases in the costs γ_f or γ_v cause group borrowers to switch to creating social capital, and this is welfare-improving.*

The proof is given in the appendix. The reason for this result is that in the neighborhood of (3.4) binding, the no-investment equilibrium is inefficient. A marginal increase in the meeting cost can be enough to give the borrowers sufficient incentive to switch to the investment equilibrium, generating a strict welfare increase.

It is worth explaining here why it is that there may not be an investment equilibrium even when utility is strictly higher under the investment than the no-investment equilibrium. In fact the reasoning is straightforward: the welfare cost of switching from investment to no-investment may be high. This is because of two things: the repayment rate is lower in the no-investment equilibrium, and the interest rate is higher. However, a borrower considering whether to deviate under the investment equilibrium does not consider the effect on the interest rate, since this only changes in equilibrium. Hence the cost of deviating from a hypothetical investment equilibrium is lower than the cost of switching from investment to no-investment.

Proposition 13 derives a condition on λ under which groups are better able to create social capital than individual borrowers. Proposition 16 simply focuses on group lending and asks when higher meeting costs actually lead to more social capital creation. As meeting costs increase, two things occur. Firstly, the lender must charge a higher interest rate, which reduces borrower welfare and tightens IJ IC2. Secondly, the cost to borrowers of being in a group with a stranger increase: by creating social capital the cost to borrowers of time spent in meetings decreases by $\lambda(\gamma_f + 2\gamma_v)$. If λ is sufficiently large, the second effect dominates and higher meeting costs increase the borrowers' incentive to invest in S.

Feigenberg et al. (2011) show that the improvement in repayment performance associated with higher meeting frequency approximately offset the increase in the lender's cost. This implies that among contracts with group meetings the total surplus was increasing in meeting frequency in their experiment. In our model, all surplus accrues to the borrower, so condition (3.7) is necessary for there to exist a region over which total surplus is increasing in the meeting frequency.

If the lender holds the interest rate fixed, as in Feigenberg et al. (2011), borrowers will be more willing to create social capital for a given increase in the meeting frequency (the extra cost is not passed on through a higher interest rate). However, a parallel condition must then hold for the increase in repayment frequency to offset the lender's costs.

3.4 Simulation

In this section, we simulate a simple extension of the model calibrated to empirically estimated parameters. This enables us to illustrate the costs and benefits of explicit joint liability and explore under which environments it will be dominated by individual liability lending that induces implicit joint liability.

We find that in low social capital environments, EJ does quite well compared to IJ. For example, when the standard deviation of project returns of 0.5, for social capital worth 10% of the loan size, the welfare attainable under IJ is 32.4% lower compared to the welfare under EJ. However, with social capital worth 50% of the loan size, the welfare attainable under EJ is 5% lower to the one attainable under IJ. We find that for social capital worth around 25% of the loan size, EJ and IJ perform approximately equally well in terms of borrower welfare. For lower values of S, EJ dominates, and for higher values of S, IJ dominates. This analysis thus gives us insights into the extent of the perverse effect of JL. With high S under IJ, the borrowers now have enough social capital to help one another when they can afford to do so, but are not penalized in states of the world where only some of the group can repay. We also find that the interest rate, repayment rate and borrower welfare are highly insensitive to social capital under EJ, whereas IJ is highly sensitive to social capital, since the only sanction available is coming through the social capital. For example, when the standard deviation of project returns is 0.5, the EJ net interest rate is 11.3%, while the IJ net interest rate ranges between 10.4% and 21.4% for levels of S valued at 10% to 50% of the loan size respectively. The difference in the interest rate translates correspondingly into borrower welfare. If borrowers share social capital worth 10% of the loan size, the attainable IJ welfare is $V^{IJ} = 2.29$, which is 35.9% lower compared to the IJ welfare of $V^{IJ} = 3.57$ attained by borrowers who share social capital worth 50% of the loan size. We also find that these welfare and interest rate differentials between low and high levels of social capital S are increasing in the variance of project returns.

From theory we know the basic trade off between EJ, Il and IJ and how that changes with social capital. What this analysis adds is to give a quantitative magnitude to the relevant thresholds and also suggests some policy implications. In low social capital environments, despite its well known costs (Besley and Coate (1995)) EJ is an effective device to induce repayment incentives and moreover, if the extent of social capital is not known *ex ante* it is a robust instrument. It also suggests a high payoff from encouraging investing in social capital given the welfare implications of higher S on borrower welfare in IJ.

3.4.1 Approach

We approach the simulations in a very similar way to de Quidt et al. (2012). Firstly, while it is theoretically convenient to model groups of size two, these require empirically implausibly high returns to investment for the borrowers to be able to repay on one another's behalf, so instead we extend the model to groups of size 5, the group size originally used by Grameen Bank and others. For simplicity, we carry over our concept of social capital unaltered to the larger groups. Previously a borrower who did not help her partner when the repayment rule stipulated she should was sanctioned by her partner. Now we simply assume she is sanctioned by the whole group, losing social capital worth S.

We express all units in multiples of the loan size and a loan term of 12 months. For example, if $S = 0.15$ this means the borrowers have social capital worth 15% of the loan size. We obtain our parameter values from the estimates in de Quidt et al. (2012). $\bar{R}$, the expected return to borrowers' investments, is set to 1.6, i.e. a 60% annual return, based on De Mel et al. (2008)'s preferred estimates of the rate of return to capital among microenterprises in Sri Lanka. The lender's cost of capital, ρ, is set to 1.098, which was estimated using lender cost data from the MixMarket database of financial information from MFIs around the world. Lastly, we set δ

equal to 0.864. This is the midpoint between the value implied by the return on US treasury bills and a lower bound implied by the model in de Quidt et al. (2012).

The two key ingredients that drive the trade-off between explicit and implicit joint liability are the level of social capital and the shape of the borrowers' return distribution function. We do not have data on social capital, so instead we estimate the equilibrium interest rate, repayment rate and welfare for a range of values for S. This enables us to say, for example, how much social capital is required for implicit and joint liability to perform as well or better than explicit joint liability.

It is more difficult to explore how the shape of the returns distribution affects the trade-off between EJ and IJ. In the theoretical analysis it was convenient to illustrate the key intuition using a simple categorical distribution with three output values and associated discrete probabilities. With larger groups, this distribution function is less useful. It no longer gives a simple and intuitive set of states of the world in which EJ does and does not perform well (with a group of size n, there are 3^n possible states of the world). More problematic is that the distribution has four parameters (p_m, R_m, p_h, R_h), only one of which can be tied down by our calibrated value of $\bar{R}$. As a result, it is very difficult to perform meaningful comparative statics - there are too many degrees of freedom.[19]

Therefore, for the main simulations we use the most obvious two-parameter distribution function, the Normal distribution.[20] Fixing the mean at $\bar{R}$, we can vary the shape of the distribution by changing the standard deviation. The range for σ was chosen to obtain the highest and lowest possible repayment rates at which the lender is able to break even. For the benchmark simulations, we assume the borrowers' returns are uncorrelated, but we also allow for positive and negative correlations in an extension.

To simulate the model, for each contract we work out a welfare-maximizing repayment rule for the borrowing group, i.e. one that maximizes the repayment rate, subject to the borrowers' incentive constraints. Solving analytically for the equilibrium repayment probability (which then gives us the interest rate and borrower welfare) is complex, so instead we simulate a large number of hypothetical borrowing groups and use these to compute the equilibrium repayment probability. We describe the simulation approach in detail in appendix 3.B.

3.4.2 Results

The main results for uncorrelated borrower incomes are presented in Figure 3.1. The standard deviation σ of individual borrower returns is varied on the horizontal axis of each figure.

For the distribution and parameter values used, it turns out that individual lia-

[19] We perform one exercise in the appendix, where we vary $p_h - p_m$ while holding p, R_h, R_m constant. The confound here is that the mean return also varies as we vary p_h and p_m.

[20] One complication arises, namely the possibility of negative income realizations. For simplicity, we allow these to occur, but we assume that only borrowers with positive incomes can assist others with repayment.

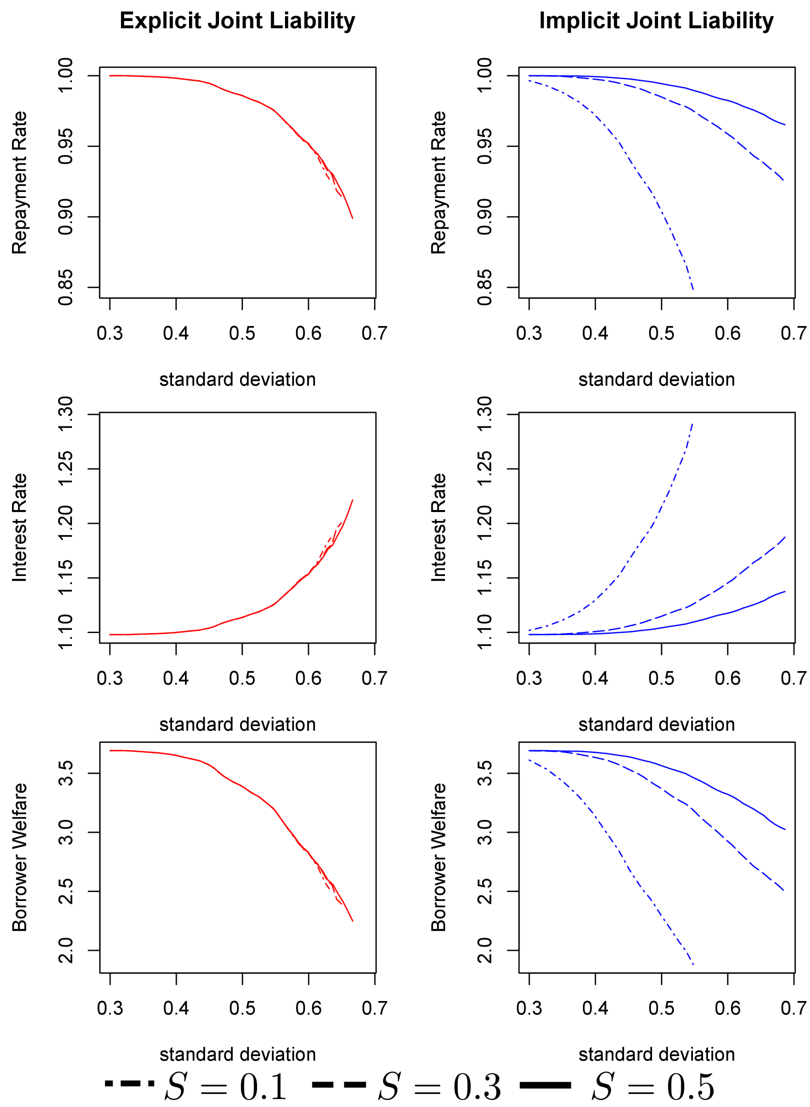

Figure 3.1: Simulation results for uncorrelated borrower returns. Explicit joint liability results are in the left column and implicit joint liability in the right column. Each figure plots the relevant object (repayment rate, interest rate and borrower welfare) for three levels of social capital, $S = 0.1, 0.3, 0.5$. The standard deviation of the individual borrower's income is varied on the horizontal axis of each figure.

bility is in fact marginally loss-making for all σ, so we just present results for implicit joint liability and explicit joint liability for values $S \in \{0.1, 0.3, 0.5\}$.

The figures show that increasing the variance of returns is bad for repayment and thus welfare under both contracts. This is unsurprising: higher variance income processes are more difficult to insure (the required transfers between members tend to be larger), so states in which members cannot or will not help one another out become more common. Increasing S partially mitigates this effect since it increases the size of incentive-compatible transfers between borrowers.

Our simulated repayment rates vary between around 85% to close to 100% as the variance of borrower income decreases. These high repayment rates follow from the fact that the calibrated mean return $\bar{R}$ is higher than the lender's cost of funds, ρ, so perfect repayment is attainable for sufficiently low variance. However, these values are fairly typical for microfinance repayment rates. For example, in de Quidt et al. (2012) we conservatively estimate a repayment rate in the MIX Market dataset of around 0.92. Using the simulated repayment rate, we can obtain the zero-profit interest rate and borrower welfare. The net interest rate varies between 10% and 30% per year (again, these are not unreasonable values for the microfinance context), while borrower welfare varies between around 1.8 and 3.7 multiples of the loan size.

One of the most striking lessons we learn from the graphs is that the interest rate, repayment rate and borrower welfare are highly insensitive to social capital under explicit joint liability. The reason is that social capital is only shifting the borrowers from default to repayment in states of the world where they can afford to help one another and where the joint liability penalty is not already sufficient. The probability that such a state occurs is lower, the bigger the sample of borrowers. Meanwhile, implicit joint liability is highly sensitive to social capital, since the only sanction available is coming through the social capital. For example, at $\sigma = 0.5$, the IJ repayment rate is 91% for $S = 0.1$, 98% for $S = 0.25$, and close to 100% for $S = 0.5$, while the EJ repayment rate is fixed at 98% throughout.[21]

In order to more easily compare EJ and IJ, in Figure 3.2 we overlay the welfare curves for EJ and IJ. The simulation exercise emphasizes much of the core intuition from the model. When S is low, explicit joint liability tends to dominate since the joint liability clause gives the borrowers an additional incentive to help one another. When S is high, implicit joint liability dominates, due to the perverse effect of JL - the borrowers now have enough social capital to help one another when they can afford to do so, but are not penalized in states of the world where only some of the group can repay.

To give a numerical example of the magnitudes of the welfare gains from EJ and IJ as a function of S, consider the case of a standard deviation of project returns of 0.5. Here for social capital worth 10% of the loan size for example, the welfare

[21]Note that in de Quidt et al. (2012) we find that the interest rate and borrower welfare are sensitive to social capital when the lender is a monopolist, since higher social capital relaxes IC2, and therefore enables the lender to increase the interest rate. The non-profit lender, as modeled in this paper, does not do this.

attainable under IJ, $V^{IJ} = 2.29$ is 32.4% lower compared to the welfare under EJ $V^{EJ} = 3.39$. This highlights the clear welfare gains that are possible under EJ in environments with low S. These gains disappear however for higher levels of S. With social capital worth 50% of the loan size, the welfare attainable under EJ $V^{EJ} = 3.39$ is in fact 5 % lower to the one attainable under EJ $V^{IJ} = 3.56$. The higher levels of social capital make it incentive compatible to help each other out, when they are able to, while not being punished when not the whole group is able to repay.

The graph also highlights that the EJ and IJ contracts are almost completely overlapping for intermediate values of $S = 0.3$ of the loan size, suggesting that in environments with intermediate levels of social capital both contracts can perform equally well.

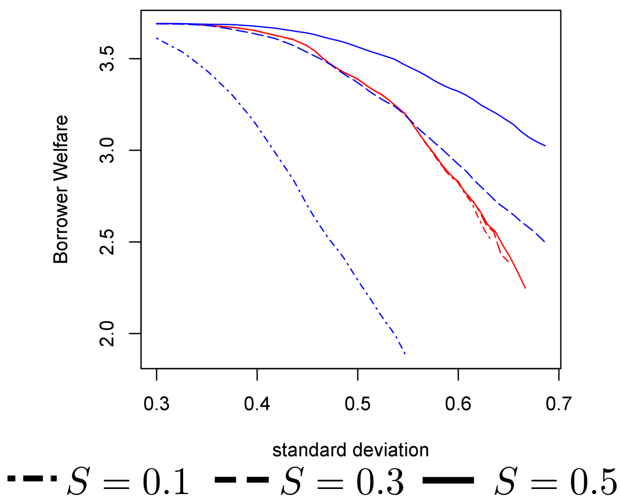

Figure 3.2: Simulation results for uncorrelated borrower returns. Explicit joint liability results are in red and implicit joint liability in blue. The figure plots borrower welfare for three levels of social capital, $S = 0.1, 0.3, 0.5$. The standard deviation of the individual borrower's income is varied on the horizontal axis.

While these results illustrate the problems with strict EJ,[22] we also interpret them as showing why EJ should not be prematurely dismissed as an important contractual tool (as also recently argued by Banerjee (2012)). Many of the candidates for alternative mechanisms discussed in the literature are complex and potentially difficult to implement, so we have focused on two extremely simple mechanisms that we feel are empirically relevant. What we find is that implicit joint liability can perform very well, provided borrowers have enough social capital: borrowers have to be willing to impose sanctions on one another worth at least 25% of their loan size. Meanwhile EJ functions well in our simulations even for low levels of social capital. This illustrates how important the lending environment, and in particular borrowers' social ties are for determining the preferred contract in our framework.

[22] Problems that have also received attention in Besley and Coate (1995), Rai and Sjöström (2004), Bhole and Ogden (2010), Rai and Sjöström (2010) and Allen (2012).

3.4.3 Correlated returns

As an extension, we now present simulation results when borrowers' returns are correlated. A number of recent papers have analyzed how correlated returns affect repayment behavior under joint liability lending.[23]. As a simple extension, we consider how our EJ and IJ borrowers are affected by introducing positively or negatively correlated returns into the model. We simulate the borrowing group's per-period income vector $[Y_1, ..., Y_n]$ as a multivariate Normal distribution. We fix the standard deviation at 0.5, the midpoint of the range considered in the previous section, and vary the pairwise correlation between group members from -0.25 to 0.45.[24] We graph the results in Figure 3.4.[25]

The main conclusion from this analysis is that for a given level of social capital, EJ is sufficiently more sensitive to the strength of correlation between borrower incomes. EJ requires all loans to be repaid. When borrower incomes are only weakly correlated, there will typically only be a small number of failures in a group, which are relatively easy for the other members to assist with. With a strongly positive correlation this is no longer the case, it becomes more common to have large numbers of failures. In this environment IJ is an advantage because the borrowers are not penalized when their partners default. This becomes evident when comparing the gradient of the IJ curves relative to the EJ curves as the correlation increases.

3.5 Conclusion

Anecdotal evidence suggests that there has been a move away from explicit joint liability towards individual liability by some prominent institutions. Most of these institutions have retained the use of groups to facilitate credit disbursal. The key question now is whether groups do more than just facilitate the lender's operations. The interest in this question has been strengthened by two recent field experiments by Giné and Karlan (2011) and Feigenberg et al. (2011).

The first of these, Giné and Karlan (2011), found that removing the joint liability clause, but retaining the group meetings, of a random subset of borrowing groups of Green Bank in the Philippines had no meaningful effect on repayment rates, although borrowers with weak social ties to other borrowers were more likely to drop out.

In this paper we have shown that this outcome may result when the newly individually liable groups have sufficient social capital to continue to guarantee one another's repayments, as under EJ, which we call implicit joint liability (IJ). We show

[23]For example, Laffont (2003), Ahlin and Townsend (2007) and Allen (2012)

[24]For correlation smaller than -0.25 we essentially have 100% repayment everywhere, and for greater than 0.45 there is typically no lending equilibrium.

[25]Note that the graphs are less smooth than those in Figure 3.1. This is because for the benchmark simulations we are able to reuse the same underlying random draws for each set of output realizations, simply by rescaling as the standard deviation changes. This is not possible when considering variously correlated returns, so we need to generate a new sample of borrower output realizations for each value of the correlation coefficient, and this naturally introduces some extra noise.

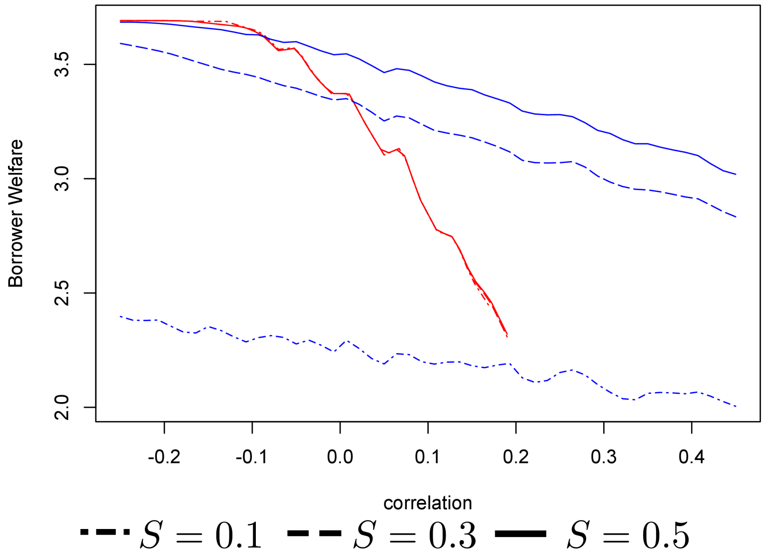

Figure 3.3: Simulation results for correlated borrower returns. Explicit joint liability results are in red and implicit joint liability in blue. The figure plots borrower welfare for three levels of social capital, $S = 0.1, 0.3, 0.5$. The correlation between pairs of borrower's returns is varied on the horizontal axis.

that this may even lead to higher repayment rates and borrower welfare. However this first result does not depend upon the use of groups, provided borrowers are able to side contract on loan repayments outside of repayment meetings.

We next show that when individual and group repayment meetings are costly, mutual insurance or IJ are easier to sustain under group lending, because IJ depends crucially on the interest rate, which in turn depends on the share of total meeting costs borne by the lender. Group meeting reduces the lender's share of meeting costs, enhancing the advantages of IJ.

The second experimental paper highlighting the role of groups is Feigenberg et al. (2011). They find that varying meeting frequency for a subset of individually liable borrowing groups seemed to have persistent positive effects on repayment rates. They suggest that this is due to improved informal insurance among these groups due to higher social capital.

We analyze situations under which microcredit might induce borrowers to create social capital, which in turn enables them to sustain IJ. We derive conditions under which group lending is more likely than individual lending to create social capital, and show when this is indeed welfare increasing. Finally, relating to one of the key findings of Feigenberg et al. (2011), we derive conditions under which more frequent meetings, modeled here as an increase in the amount of time borrowers and loan officers must spend in loan repayment meetings, increases borrowers' incentive to invest in social capital. This provides a theoretical foundation for Feigenberg et al. (2011)'s observation. We also carry out a simulation exercises to assess the quantitative magnitudes of the effects of alternative forms of lending, as well as some of the relevant thresholds of social capital.

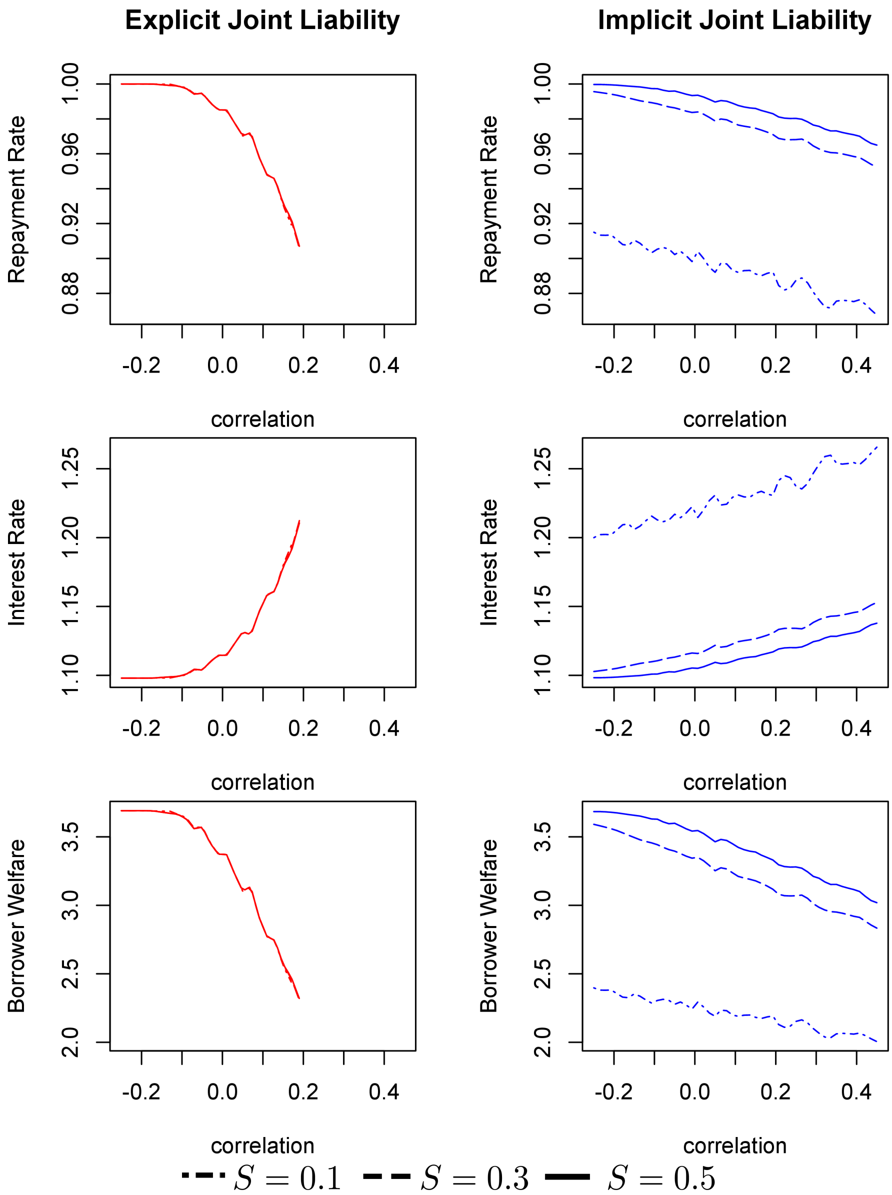

Figure 3.4: Simulation results for uncorrelated borrower returns. Explicit joint liability results are in the left column and implicit joint liability in the right column. Each figure plots the relevant object (repayment rate, interest rate and borrower welfare) for three levels of social capital, $S = 0.1, 0.3, 0.5$. The correlation between pairs of borrower's returns is varied on the horizontal axis of each figure.

3.A Mathematical appendix

Proof of Lemma 1

Proof. Comparing the expressions for $\hat{S}_A^{EJ}$ and $\hat{S}_A^{IJ}$, it is immediate that $\hat{S}_A^{EJ} < \hat{S}_A^{IJ}$ since $\pi_A^{EJ} = \pi_A^{IJ}$ and $\delta\pi_A^{EJ}\bar{R} - \rho > 0$ by Assumption 7.

Now consider Case B. It is obvious that if $\hat{S}^{EJ} = 0$, $\hat{S}^{IJ} > \hat{S}^{EJ}$, since $\hat{S}^{IJ} > 0$. Suppose therefore that $\hat{S}^{EJ} > 0$. It is straightforward to check that Assumptions 5, 6 and 7 imply that $\delta p \geq \frac{1}{2}$. Given this, and $p_h \geq p_m$, it follows that $\pi^{IJ} \geq p \geq \frac{1}{2}$ and $\pi^{EJ} \geq p \geq \frac{1}{2}$. Also using the fact that π_B^{EJ} can be written as $p^2 + 2p_h(1-p)$. We have:

$$
\begin{aligned}
\hat{S}^{IJ} - \hat{S}^{EJ} &= \frac{\delta\pi_B^{EJ}\bar{R} - \rho}{\delta\pi_B^{EJ}(1-\delta\pi_B^{EJ})} + \frac{\rho}{\delta\pi_B^{IJ}} - \frac{\rho}{\delta\pi_B^{EJ}} \\
&= \frac{\pi_B^{IJ}(\delta\pi_B^{EJ}\bar{R} - \rho) - p_m(1-p)(1-\delta\pi_B^{EJ})\rho}{\delta\pi_B^{IJ}\pi_B^{EJ}(1-\delta\pi_B^{EJ})} \\
&\geq \frac{(\delta\pi_B^{EJ}\bar{R} - \rho) - p_m(1-p)\rho}{2\delta\pi_B^{IJ}\pi_B^{EJ}(1-\delta\pi_B^{EJ})} \\
&= \frac{\delta p^2\bar{R} - \rho + p_h(1-p)(2\delta\bar{R} - \rho) + (p_h - p_m)(1-p)\rho}{2\delta\pi_B^{IJ}\pi_B^{EJ}(1-\delta\pi_B^{EJ})} \\
&> 0
\end{aligned}
$$

which follows from $2\delta\bar{R} - \rho > 0$ by Assumption 7. ■

Proof of Proposition 11

To compare IL and EJ, we consider first Case A, then Case B with $p_h > p_m$, and lastly Case B with $p_h \leq p_m$.

In Case A, borrower repayment guarantees under IL offer no advantage over EJ, so provided $S \geq \hat{S}_A^{EJ}$, EJ is the borrower welfare-maximizing contract (with indifference for $S \geq \hat{S}^{IJ}$). For $S < \hat{S}_A^{EJ}$, borrower will not mutually guarantee under EJ and also default unless their partner is successful, so IL is preferred to EJ:

$$
V_A^{EJ}(S) - V_A^{IL}(S) = \begin{cases} -\frac{\delta p(1-p)(\bar{R}-\rho)}{(1-\delta p)(1-\delta p^2)} & S < \hat{S}_A^{EJ} \\ \frac{\delta p(1-p)(\bar{R}-\rho)}{(1-\delta p)(1-\delta p(2-p))} & S \in [\hat{S}_A^{EJ}, \hat{S}_A^{IJ}) \\ 0 & S \geq \hat{S}_A^{IJ} \end{cases}
$$

In Case B, with $p_h > p_m$, EJ dominates IL when borrowers guarantee one another under EJ but not under IL, for $S \in [\hat{S}_B^{EJ}, \hat{S}_B^{IJ})$, so EJ is preferred in this region. However, once IJ is possible, for $S \geq \hat{S}_B^{IJ}$, it dominates EJ. This is because borrower

1 repays her own loan in state $(R_m, 0)$, while she would default under EJ. We have:

$$V_B^{EJ}(S) - V_B^{IL}(S) = \begin{cases} -\frac{\delta p(1-p)(\bar{R}-\rho)}{(1-\delta p)(1-\delta p^2)} & S < \hat{S}_B^{EJ} \\ \frac{\delta\triangle(1-p)(\bar{R}-\rho)}{(1-\delta p)(1-\delta(p+\triangle(1-p)))} & S \in [\hat{S}_B^{EJ}, \hat{S}_B^{IJ}) \\ -\frac{\delta p_m(1-p)(\bar{R}-\rho)}{(1-\delta(p+p_h(1-p)))(1-\delta(p+\triangle(1-p)))} & S \geq \hat{S}_B^{IJ} \end{cases}$$

Lastly, in Case B with $p_h \leq p_m$, EJ is always dominated by IL. This is because under EJ the highest possible repayment probability is $p + \triangle(1-p)$, which is weakly smaller than p, the lowest possible repayment probability under IL. Therefore we do not need to know the ordering of $\hat{S}_B^{EJ}$ and $\hat{S}_B^{IJ}$ for this case - EJ will never be used.

Proof of Corollary 1

Suppose total meeting costs are higher under ILG: $\frac{3}{2}(\gamma_f + 2\gamma_v) > 2(\gamma_f + \gamma_v)$ or $2\gamma_v > \gamma_f$. Suppose also that $S \in [\hat{S}^{IJG}, \hat{S}^{IJI})$. Then group lending sustains IJ but individual lending does not. Welfare is higher under group lending if:

$$\frac{\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)}{1 - \delta(p + p_h(1-p))} > \frac{\bar{R} - \rho - 2(\gamma_f + \gamma_v)}{1 - \delta p}$$

Taking the limit as $\gamma_f \to 2\gamma_v$, it is clear that this condition holds strictly, while $\hat{S}^{IJG} > \hat{S}^{IJI}$ continues to hold, thus the corollary follows for a non-trivial interval of costs by a standard open set argument.

Proof of Lemma 2

First, note that $\frac{\partial^2 V}{\partial r \partial \pi} < 0$. Therefore, the benefit of increasing π is higher when interest rates are low.

We want to find conditions under which ILI borrowers will not invest in social capital in equilibrium. To show this, we hypothesize a (low interest rate) equilibrium in which ILI borrowers do invest, and show that there exists a profitable deviation. Then, we know that in a (high interest rate) equilibrium in which borrowers do not invest, they will not wish to deviate to investing; this follows from $\frac{\partial^2 V}{\partial r \partial \pi} < 0$ as noted above.

Consider then a hypothetical equilibrium in which the borrowers do invest in social capital and repay with probability $\pi^{IJ} \equiv p + p_h(1-p)$. They are charged $\hat{r} = \frac{\rho + \gamma_f + \gamma_v}{\pi^{IJ}}$.

At the beginning of the first period, the borrower and her partner pay cost η and create social capital. Then, each period with probability $p + \triangle(1-p)$, both loans are repaid and both contracts renewed. With probability $p_m(1-p)$, only borrower i's loan is repaid. As a result, at the beginning of the next period, she must again pay cost η to create social capital with her new partner.[26]

[26]Since no social capital is destroyed on the equilibrium path, the S created with the original partner

Consider an ILI borrower in the first period, or one whose partner has just defaulted. We know that IC1 is satisfied, since by repaying her loan she can guarantee herself at least $\delta(\bar{R} - (\gamma_f + \gamma_v)) - \frac{\rho+\gamma_f+\gamma_v}{\pi^{IJ}}$ if she agrees with the new partner to simply take a loan and default immediately. This expression is positive by the modified Assumption 3 in footnote 15. Then we note that if it is an equilibrium for the borrower to invest in social capital, it must be that she does even better than this, and therefore IC1 must hold.

As we are considering an equilibrium in which she invests in social capital, we use an "IJI" superscript to denote the fact that IJ is taking place. If she invests in social capital with the new partner, she earns utility U, defined as follows:

$$U_1^{IJI} = S - \eta + W_1^{IJI}$$

where

$$W_1^{IJI} = \left(\bar{R} - \rho - 2(\gamma_f + \gamma_v)\right) + \delta(p + \triangle(1-p))W_1^{IJI} + \delta p_m(1-p)U_1^{IJI}.$$

The first term in W is the per-period utility under ILI. The second term represents the continuation payoff when both borrowers repay and have their contracts renewed. This occurs with probability $p + \triangle(1-p)$. In this case she earns W_1^{IJI} next period. The third term represents the continuation payoff if she repays but her partner defaults, which occurs with probability $p_m(1-p)$. In this case she matches with a new partner and therefore earns U_1^{IJI} next period.

Substituting for W, we can write U as:

$$\begin{aligned}
U_1^{IJI} &= S - \eta + \frac{\left(\bar{R} - \rho - 2(\gamma_f + \gamma_v)\right) + \delta p_m(1-p)(S-\eta)}{1 - \delta\pi^{IJ}} \\
&= \frac{\left(\bar{R} - \rho - 2(\gamma_f + \gamma_v)\right) + (1 - \delta(p + \triangle(1-p)))(S-\eta)}{1 - \delta\pi^{IJ}}.
\end{aligned}$$

Now we check for a one-shot deviation. In this context, a deviation is to defer investing in social capital by one period, i.e. to undergo one period without social capital (and therefore with repayment probability p), then invest in social capital next period. She prefers to deviate if:

$$U_1^{IJI} < \left(\bar{R} - p\frac{\rho + \gamma_f + \gamma_v}{\pi^{IJ}} - (\gamma_f + \gamma_v)\right) + \delta p U_1^{IJI}. \tag{3.8}$$

The first term on the right hand side represents the per-period utility of a borrower under ILI without social capital, paying an interest rate of $\hat{r} = \frac{\rho+\gamma_f+\gamma_v}{\pi^{IJ}}$ (intuitively, since the lender does not know she has deviated, the interest rate is not adjusted). With probability p her loan is repaid, and in the next period she invests in S, thus receiving continuation value U_1^{IJI}. Substituting for U_1^{IJI} and rearranging yields condition (3.3).

is not lost but cannot be leveraged in the credit contract.

Proof of Lemma 3

Hypothesize an equilibrium in which borrowers invest in social capital. We know that IC1 is satisfied, since by repaying her loan she can guarantee herself at least $\delta(\bar{R} - (\gamma_f + 2\gamma_v)) - \frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{\pi^{IJ}}$ if she agrees with the new partner to simply take a loan and default immediately. This expression is positive by the modified Assumption 3 in footnote 15.

We need to check that no borrower prefers to deviate by deferring their investment by one period, exactly as in Lemma 2. We define the value functions analogously to those in the proof of Lemma 2:

$$U_1^{IJG} = S - \eta + W_1^{IJG}$$
$$W_1^{IJG} = \left(\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda)\right) + \delta(p + \triangle(1-p))W_1^{IJG} + \delta p_m(1-p)U_1^{IJG}.$$

Where the possession of social capital reduces the borrowers' cost of group meetings by $\lambda(\gamma_f + 2\gamma_v)$. The appropriate substitutions yield:

$$U_1^{IJG} = \frac{\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda) + (1 - \delta(p + \triangle(1-p)))(S - \eta)}{1 - \delta\pi^{IJ}}.$$

There will be no deviation if $U_1^{IJG} \geq \left(\bar{R} - p\frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{\pi^{IJ}} - (\gamma_f + 2\gamma_v)\right) + \delta p U_1^{IJG}$. Simplifying yields condition (3.4).

Proof of Proposition 14

Total borrower welfare under ILI (where borrowers do not invest in social capital) is:

$$\begin{aligned} V^{ILI} &= \bar{R} - \rho - 2(\gamma_f + \gamma_v) + \delta p V^{ILI} \\ &= \frac{\bar{R} - \rho - 2(\gamma_f + \gamma_v)}{1 - \delta p} \end{aligned}$$

and when groups are used (and the borrowers do invest in social capital) it is:

$$U_1^{IJG} = \frac{\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda) + (1 - \delta(p + \triangle(1-p)))(S - \eta)}{1 - \delta\pi^{IJ}}.$$

as was derived in the proof of Lemma 3. The result then follows from comparison of these value functions.

Proof of Proposition 15

First, observe that if $\gamma_v \leq \frac{\gamma_f}{2}$, condition (3.5) is satisfied for all $\lambda \geq 0$, hence $G_1 < G_2$. From the proof of Lemma 3, $\eta - S \leq G_2$ if and only if $U_1^{IJG} \geq \frac{\bar{R} - p\frac{\rho + \frac{1}{2}(\gamma_f + 2\gamma_v)}{\pi^{IJ}} - (\gamma_f + 2\gamma_v)}{1 - \delta p}$. Call the RHS of this condition B. From the proof of Proposition 14, $\eta - S < G_3$ if and only if $U_1^{IJG} > V^{ILI}$. Finally, note that $B - V^{ILI} = \frac{p_h(1-p)(\rho+\gamma_f)+p\left(\frac{\gamma_f}{2}-\gamma_v\right)}{\pi(1-\delta p)}$, which is strictly positive if $\gamma_v < \frac{\gamma_f}{2}$. Thus, $\eta - S \leq G_2$ implies $\eta - S < G_3$, or $G_2 < G_3$.

Claim 1 follows immediately from $G_1 < G_2 < G_3$. Claim 2, that borrower welfare is *always* higher under ILG, can be broken into three parts. Firstly, if $\eta - S \leq G_1$, both groups and individuals invest in social capital. Then, the cost advantage of ILG ($\gamma_v \leq \frac{\gamma_f}{2}$) implies that welfare is higher under ILG. Secondly, if $\eta - S > G_2$, neither groups nor individuals invest in S, and again the cost advantage leads to ILG dominating. Lastly, if $G_1 < \eta - S \leq G_2$, groups invest and individuals do not, and thus ILG dominates by Proposition 14.

Proof of Corollary 2

Suppose condition (3.4) binds, such that a small decrease in γ_f causes borrowers to stop investing in social capital. We want to show that this leads to a discontinuous decrease in welfare.

Before the change, welfare is:

$$U_1^{IJG} = \frac{\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda) + (1 - \delta(p + \triangle(1-p)))(S - \eta)}{1 - \delta\pi^{IJ}}.$$

after the change (in the limit as the increase in γ_f approaches zero), it is:

$$V^{ILG} = \frac{\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)}{1 - \delta p}$$

since the borrowers can no longer sustain IJ, so the new equilibrium is one in which they repay with probability p and the interest rate is $\frac{\rho+\frac{1}{2}(\gamma_f+2\gamma_v)}{p}$. From condition (3.4) binding we know that:

$$\eta - S = \frac{p_h(1-p)\left[\delta\left(\bar{R} - \frac{\rho}{\pi^{IJ}}\right) - \frac{1+2\delta\pi^{IJ}}{2\pi^{IJ}}(\gamma_f + 2\gamma_v)\right] + \lambda(1-\delta p)(\gamma_f + 2\gamma_v)}{(1-\delta p)(1 - \delta(p + \triangle(1-p)))}. \quad (3.9)$$

For U_1^{IJG} to be strictly larger than V^{ILG} we require:

$$\frac{\bar{R} - \rho - \frac{1}{2}(\gamma_f + 2\gamma_v)(3 - 2\lambda) + (1 - \delta(p + \triangle(1-p)))(S - \eta)}{1 - \delta\pi^{IJ}} > \frac{\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)}{1 - \delta p}$$

which reduces to

$$\frac{\delta p_h(1-p)\left(\bar{R} - \rho - \frac{3}{2}(\gamma_f + 2\gamma_v)\right) + \lambda(1-\delta p)(\gamma_f + 2\gamma_v)}{(1-\delta p)(1 - \delta(p + \triangle(1-p)))} > \eta - S.$$

Substituting for $\eta - S$ from (3.9) and simplifying, we obtain:

$$2\delta\rho(1 - \pi^{IJ}) + (1 - \delta\pi^{IJ})(\gamma_f + 2\gamma_v) > 0$$

which is satisfied.

More generally, this demonstrates that the no-investment equilibrium is inefficient in the neighborhood of $\eta - S = G_2$. A marginal increase in the meeting cost that gives the borrowers greater incentive to invest in social capital can lead to a strict increase in borrower welfare.

3.B Simulation approach

This Appendix outlines the algorithm used to simulate the core model. The simulation was implemented in *R*. The intuition of the simulation procedure is very straightforward. We use a random sample of N groups with n members each. A group merely constitutes a vector of income realizations. These incomes are drawn from some distribution function F. We assume that F is a Normal distribution with $\mu = \bar{R} = 1.6$, however we allow the standard deviation σ to vary.

Given these income realizations, we compute the repayment rate that would arise under each contract for a given interest rate r. This process gives us a repayment probability function $\pi(r)$ under either contract.

Given this repayment probability function, we can then compute the break-even repayment rate and thus the break-even interest rate under each contract, along with borrower welfare. This then allows us to make comparisons between the two contractual forms.

We now describe in detail how the group-level repayment rate is computed, as this is different under each contract type due to the different incentive constraints.

We denote an income realization of a group i with n borrowers is represented by an n-vector, $Y_i = (y_1, ..., y_n)$, where y_j is group member j's income draw.

We want to find a repayment rule analogous to the one outlined in the theory that allows for larger groups and the continuous output distribution. The most obvious way to do this is to construct for each Y_i a "group bailout fund" that can be used for transfers between group members to assist with repayments. Since the incentive constraints differ between EJ and IJ, the construction of the group fund also differs and is described below.

Group Lending without Joint Liability

The relevant incentive constraint under group lending without joint liability implies that the maximum amount a group member j is willing to contribute to the group fund is $c_{ij} = max(y_{ij}, \delta S)$. All the transfers are put into a common pool C_j. This pool is then used to ensure the maximum possible number of repayments. The borrowers are sorted in ascending order of the amount of transfer they require to repay their own loan. and transfers made from the fund until it is exhausted.[27]. If m

[27]This in fact implies that in some cases the worse off borrowers will be bailing out the better off borrowers. In particular, it may be that an unlucky borrower gives her whole income to a partner to repay their loan, but defaults on her own loan. This is because the worse off borrowers require a larger transfer, which is thus less likely to be incentive compatible. This mechanism achieves the maximum possible repayment rate and therefore maximizes ex-ante expected utility.

group members repay, then we obtain a group level repayment rate $\pi_i = \frac{m}{n}$. As this procedure is repeated for a sample of N groups, we can then estimate the overall repayment probability as the simple average.

The procedure in pseudo-code:

Group Lending without JL

1. Generate a $N \times n$ matrix of income realizations from F.

2. For each possible value of the interest rate r:

   (a) For each Y_i: compute the maximum level of contributions that each group member is willing to make to the common pool as $c_{ij} = max(y_{ij}, \delta S)$. This pot amounts to $C_{ij} = \sum_n c_{ij}$

   (b) Compute the redistributions required by members to ensure repayment as $t_{ij} = max(0, r - y_{ij} - c_{ij})$.

   (c) Order the required transfer in ascending order and redistribute the pot C_{ij} until it is exhausted.

   (d) Compute the group level repayment rate $\pi_i(r)$.

3. Given all the π_i, compute $\pi(r) = \frac{\sum_i \pi_i}{N}$.

Group Lending with Joint Liability

The simulation of this contract is more involved, since the relevant incentive constraint is $c_{ij} \leq \delta(V + S)$. This implies that in order to construct the repayment rate π, a number for the continuation value V is needed. V however, is itself a function of π.

The method proceeds as follows, for each possible value of r. First, we construct a set of possible candidates for $\pi(r)$, denoted $\tilde{\pi}$ [28], we calculate the associated $V(\tilde{\pi})$. Given these candidate $\tilde{V}$'s, the group fund C_{ij} is computed as follows. Each member is willing to contribute at most $c_{ij} = max(y_{ij}, \delta(\tilde{V} + S))$ toward repayment of the group's loan obligations. Explicit joint liability implies that the group will only repay when $C_j = \sum_n c_{ij} \geq nr$. Thus a group's repayment rate is $\pi_i = \mathcal{I}[C_j \geq nr] \in \{0, 1\}$. Taking the average we obtain the simulated repayment rate given $\hat{\pi}(V(\tilde{\pi}))$. In other words, taking as given a value for $V(\tilde{\pi})$, the implied repayment rate $\hat{pi}$ is computed. Then, the true π (and thus the true V) is found by solving for the fixed point $\pi = \hat{\pi}(V(\hat{\pi}))$. By iterating over r, we obtain the schedule $\pi(r)$ and the associated $V(\pi(r))$.

The procedure in pseudo code:

This does not imply that a borrower with $y_j > r$ would ever default (i.e. be forced to choose between losing δV and δS. The reason is that all borrowers "above" her in the bail out chain also have $y > r$, so are making net positive contributions to the fund, which therefore has a positive "balance" when her turn comes

[28] These candidate π's exploit the monotonicity of the $\pi(r)$ schedule. The upper bound is given by the previous iteration for a higher r, while the lower bound is globally defined as $\frac{\rho}{(\delta R)}$

Group Lending with JL

1. Generate a $N \times n$ matrix of income realizations from F.
2. For each interest rate r:
   (a) Construct a set of candidates for $\tilde{\pi}(r)$.
   (b) For each $\tilde{\pi}(r)$:
      - For each Y_i: compute the maximum level of contributions that each group member is willing to make to the common pool as $c_{ij} = max(y_{ij}, \delta(S + V(\tilde{\pi})))$. This pot amounts to $C_{ij} = \sum_n c_{ij}$
      - The group defaults if $C_j = \sum_n c_{ij} < nr$
      - Compute the group level repayment rate $\hat{\pi}_i(\tilde{\pi})$.
3. Given all the $\hat{\pi}_i(V(\tilde{p}))$, compute $\hat{\pi}(V(\tilde{\pi}))$ as the average and find the fixed point π such that $\pi = \hat{\pi}(V(\pi))$.

3.C Simulation Results for Piecewise Returns

As discussed in the main text, there is no straightforward approach to simulate the model with the piecewise returns distribution. The problem is one of too many degrees of freedom. A sensible approach would be to vary the difference between the parameters p_h and p_m, as we saw in the main draft that for $p_h < p_m$, group lending with joint liability performs particularly bad. We can vary this difference, but still hold the sum $p_h + p_m = \bar{p}$ fixed, where $\bar{p} = 0.921$, as in de Quidt et al. (2012).[29]

We still have three parameters to tie down. Namely R_m, R_h and the mean return. There is no straightforward approach to tie down either of these parameters when varying the difference between p_h and p_m. This appendix will show the results from one pragmatic way. First, we tie down $R_m = \rho/p^2$. This condition is motivated by assumption 5 for the two player model. It implies that the medium return is high enough to repay a individual liability loan. Given this and the value of $\bar{R} = 1.6$, we compute R_h imposing the constraint that $p_h = p_m$. This thus gives us the value for R_h, when the difference between p_h and p_m is zero. Given these fixed values, we then simply vary the difference between p_h and p_m, holding everything else constant. This exercise thus maps somewhat into the table of the two-player model, where the model suggest that there is only an IL equilibrium for low S and only IJ equilibria for sufficiently high S. There is no EJ equilibrium in this case however. For $\Delta > 0$, the simple model would predict EJ lending for some range of parameter values. In the two-player model thus, the Δ is key. For groups with larger size, we would not expect this simple result to go through as now there are a lot more

[29]Please refer to this paper for details on how this value was estimated using cross-sectional data from the MIX Market database

states of the world. However, when plotting the simulation results as a function of the difference between p_h and p_m in figure 3.5, we do see that EJ performs better the larger $p_h - p_m$. However, this may simply be due to the fact that for higher p_h relative to p_m, the mean return in this case is changing as well.

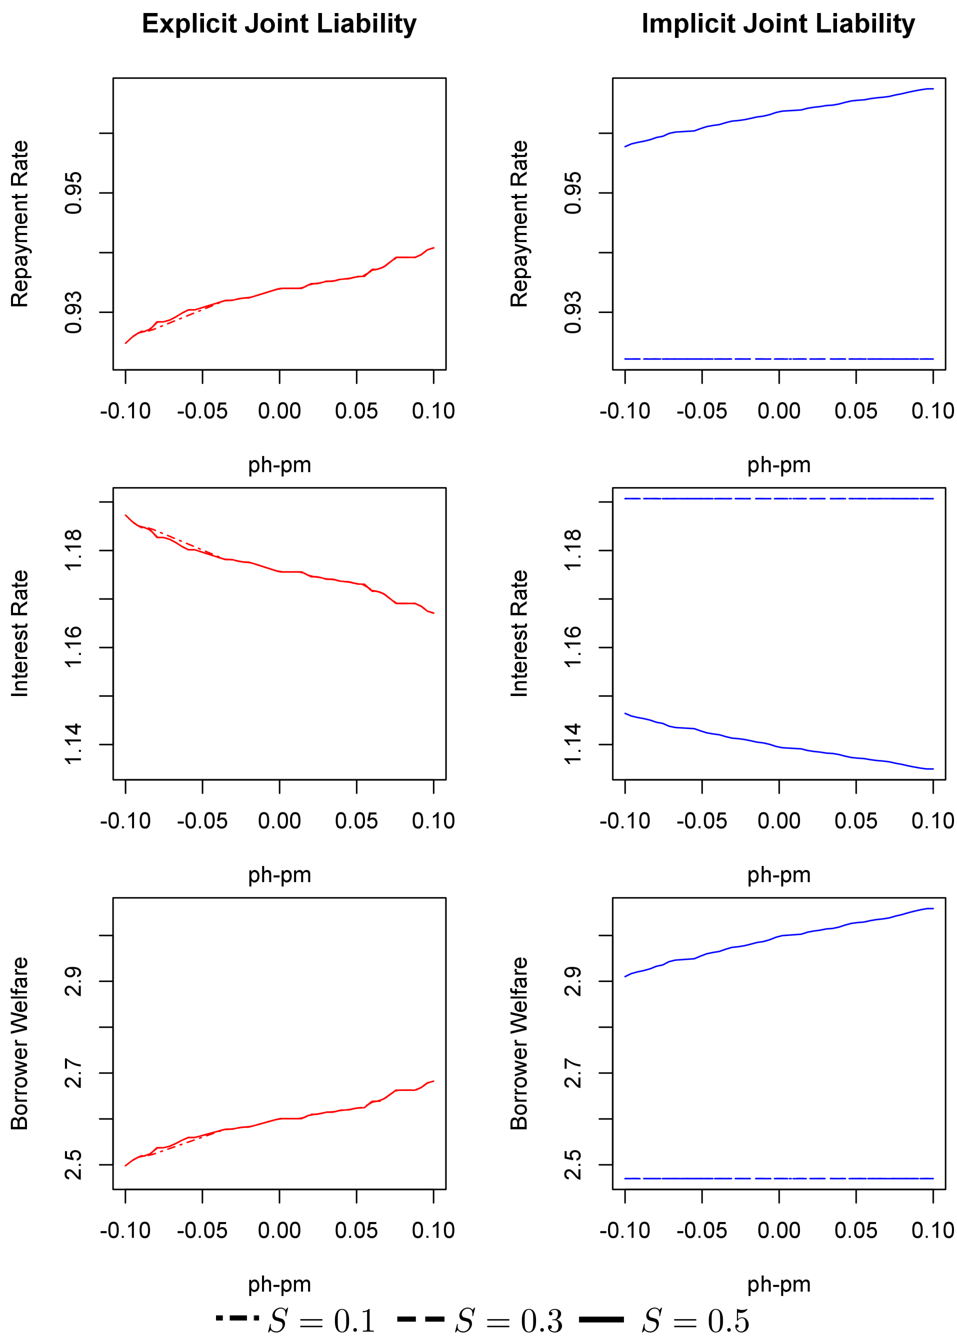

Figure 3.5: Simulation results for piecewise borrower returns distribution. Curves for explicit joint liability are drawn in red, and implicit joint liability in blue. Each figure plots the relevant object (repayment rate, interest rate and borrower welfare) for three levels of social capital, $S = 0.1, 0.3, 0.5$. The difference between p_h and p_m of individual borrower returns is varied on the horizontal axis of each figure.

Bibliography

Abeler, J., A. Falk, L. Goette, and D. Huffman (2011). Reference Points and Effort Provision. *American Economic Review 101*(2), 470–492.

Ahlin, C. and N. Jiang (2008, January). Can micro-credit bring development? *Journal of Development Economics 86*(1), 1–21.

Ahlin, C. and R. M. Townsend (2007, January). Using Repayment Data to Test Across Models of Joint Liability Lending. *The Economic Journal 117*(517), F11–F51.

Allen, T. (2012). Optimal (Partial) Group Liability in Microfinance Lending. *mimeo*.

Amir, O., D. G. Rand, and Y. K. Gal (2012, January). Economic games on the internet: the effect of $1 stakes. *PloS one 7*(2), e31461.

Anderson, C., L. Locker, and R. Nugent (2002, January). Microcredit, Social Capital, and Common Pool Resources. *World Development 30*(1), 95–105.

Ariely, D., G. Loewenstein, and D. Prelec (2003, February). "Coherent Arbitrariness": Stable Demand Curves Without Stable Preferences. *The Quarterly Journal of Economics 118*(1), 73–106.

Armantier, O. and A. Boly (2012). Framing of Incentives and Effort Provision. *mimeo*.

Armendáriz de Aghion, B. and J. Morduch (2010, May). *The Economics of Microfinance, Second Edition*. MIT Press Books. Cambridge, MA: The MIT Press.

Attanasio, O., B. Augsburg, R. De Haas, E. Fitzsimons, and H. Harmgart (2011). Group Lending or Individual Lending? Evidence from a Randomized Field Experiment in Mongolia. *EBRD Working Paper No. 136*.

Baker, G., M. Jensen, and K. Murphy (1988). Compensation and Incentives: Practice vs. Theory. *The Journal of Finance 43*(3), 593–616.

Baland, J.-M., R. Somanathan, and Z. Wahhaj (2010, January). Repayment incentives and the distribution of gains from group lending. *CDEDSE Working Paper* (192).

Banerjee, A. V. (2012, September). Microcredit Under the Microscope: What Have We Learned in the Past Two Decades, and What Do We Need to Know? *Annual Review of Economics 5*(1).

Banerjee, A. V. (2013, January). Microcredit Under the Microscope: What Have We Learned in the Past Two Decades, and What Do We Need to Know? *Annual Review of Economics 5*(1), 487–519.

Banerjee, A. V., E. Duflo, R. Glennerster, and C. Kinnan (2010). The miracle of microfinance? Evidence from a randomized evaluation. *Department of Economics, Massachusetts Institute of Technology (MIT) Working Paper, May*.

Baquero, G., M. Hamadi, and A. Heinen (2012). Competition, loan rates and information dispersion in microcredit markets. *ESMT Working Paper*.

Barankay, I. (2011). Rankings and Social Tournaments: Evidence from a Crowd-Sourcing Experiment. *mimeo*, 1–30.

Bénabou, R. and J. Tirole (2003, July). Intrinsic and Extrinsic Motivation. *The Review of Economic Studies 70*(3), 489–520.

Besley, T., K. Burchardi, and M. Ghatak (2012). Incentives and the de Soto Effect. *The Quarterly Journal of Economics 127*(1), 237–282.

Besley, T. and S. Coate (1995). Group lending, repayment incentives and social collateral. *Journal of Development Economics 46*(1), 1–18.

Bhole, B. and S. Ogden (2010). Group lending and individual lending with strategic default. *Journal of Development Economics 91*(2), 348–363.

Bloch, F., G. Genicot, and D. Ray (2008). Informal insurance in social networks. *Journal of Economic Theory 143*(1), 36–58.

Bolton, P. and D. Scharfstein (1990). A Theory of Predation Based on Agency Problems in Financial Contracting. *The American Economic Review 80*(1), 93–106.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012, April). Salience Theory of Choice Under Risk. *The Quarterly Journal of Economics 127*(3), 1243–1285.

Brooks, R. R. W., A. Stremitzer, and S. Tontrup (2013). Stretch It but Don't Break It: The Hidden Risk of Contract Framing. *Working paper*.

Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler (1997, May). Labor Supply of New York City Cabdrivers: One Day at a Time. *The Quarterly Journal of Economics 112*(2), 407–441.

Camerer, C. and R. Hogarth (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty 19*(1-3), 7–42.

Carpena, F., S. Cole, J. Shapiro, and B. Zia (2010). Liability Structure in Small-Scale Finance: Evidence from a Natural Experiment. *World Bank Policy Research Working Paper 5427*.

Cassar, A., L. Crowley, and B. Wydick (2007, February). The effect of social capital on group loan repayment: evidence from field experiments. *The Economic Journal 117*(517), F85–F106.

Chowdhury, P. R. (2005). Group-lending: Sequential financing, lender monitoring and joint liability. *Journal of Development Economics 77*(2), 415–439.

Church, B. K., T. Libby, and P. Zhang (2008, December). Contracting Frame and Individual Behavior: Experimental Evidence. *Journal of Management Accounting Research 20*(1), 153–168.

Crawford, V. P. and J. Meng (2011, August). New York City Cab Drivers' Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income. *American Economic Review 101*(5), 1912–1932.

Crépon, B., F. Devoto, E. Duflo, and W. Parienté (2011). Impact of microcredit in rural areas of Morocco: Evidence from a Randomized Evaluation. *mimeo*.

Cull, R., A. Demirgüç-kunt, and J. Morduch (2007, February). Financial performance and outreach: a global analysis of leading microbanks. *The Economic Journal 117*(517), F107–F133.

Cull, R., A. Demirgüç-kunt, and J. Morduch (2009). Microfinance meets the market. *Journal of Economic Perspectives 23*(1), 167–192.

Dal Bó, E., F. Finan, and M. Rossi (2012). Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service. *NBER Working Paper No. 18156*.

De Janvry, A., C. McIntosh, and E. Sadoulet (2010, November). The supply- and demand-side impacts of credit market information. *Journal of Development Economics 93*(2), 173–188.

De Mel, S., D. McKenzie, and C. Woodruff (2008). Returns to Capital in Microenterprises: Evidence from a Field Experiment.

de Meza, D. and D. Webb (2007). Incentive design under loss aversion. *Journal of the European Economic Association 5*(March), 66–92.

de Quidt, J., T. Fetzer, and M. Ghatak (2012). Market Structure and Borrower Welfare in Microfinance. *BREAD Working Paper No. 360*.

de Quidt, J., T. Fetzer, and M. Ghatak (2013). Group Lending Without Joint Liability. *STICERD - Economic Organisation and Public Policy Discussion Papers Series*.

DellaVigna, S. and U. Malmendier (2004, May). Contract Design and Self-Control: Theory and Evidence. *The Quarterly Journal of Economics 119*(2), 353–402.

Deserranno, E. (2014). Earning Expectations and Selection of Social Workers: Experimental Evidence from Uganda. *mimeo*.

Dohmen, T. and A. Falk (2011, April). Performance Pay and Multidimensional Sorting: Productivity, Preferences, and Gender. *American Economic Review 101*(2), 556–590.

Druckman, J. N. (2001, April). Using Credible Advice to Overcome Framing Effects. *Journal of Law, Economics, and Organization 17*(1), 62–82.

Eriksson, T. and M. C. Villeval (2008). Performance-pay, sorting and social motivation. *Journal of Economic Behavior & Organization 68*(2), 412 – 421.

Farber, H. S. (2005, February). Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers. *Journal of Political Economy 113*(1), 46–82.

Farber, H. S. (2008, May). Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers. *American Economic Review 98*(3), 1069–1082.

Fehr, E. and S. Gächter (2002). Do Incentive Contracts Undermine Voluntary Cooperation? *IEW Working Paper No. 34*.

Fehr, E., L. Goette, and C. Zehnder (2009, September). A Behavioral Account of the Labor Market: The Role of Fairness Concerns. *Annual Review of Economics 1*(1), 355–384.

Feigenberg, B., E. Field, and R. Pande (2011). The Economic Returns to Social Interaction: Experimental Evidence from Microfinance. *mimeo*.

Feigenberg, B., E. Field, and R. Pande (2013, April). The Economic Returns to Social Interaction: Experimental Evidence from Microfinance. *The Review of Economic Studies 80*(4), 1459–1483.

Fischer, G. and M. Ghatak (2010). Repayment Frequency in Microfinance Contracts with Present-Biased Borrowers. *mimeo*.

Frederick, S., G. Loewenstein, and T. O'donoghue (2002, June). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature 40*(2), 351–401.

Fryer, R. G., S. D. Levitt, J. List, and S. Sadoff (2012). Enhancing the Efficacy of Teacher Incentives Through Loss Aversion: A Field Experiment. *NBER working paper 18237*.

Ghatak, M. and T. W. Guinnane (1999). The economics of lending with joint liability: theory and practice. *Journal of Development Economics 60*(1), 195–228.

Gill, D. and V. Prowse (2012, February). A Structural Analysis of Disappointment Aversion in a Real Effort Competition. *American Economic Review 102*(1), 469–503.

Giné, X., P. Jakiela, D. S. Karlan, and J. Morduch (2010). Microfinance games. *American Economic Journal: Applied Economics 2*(3), 60–95.

Giné, X. and D. S. Karlan (2011). Group versus Individual Liability : Short and Long Term Evidence from Philippine Microcredit Lending Groups. *mimeo*.

Giné, X. and D. S. Karlan (2013). Group versus Individual Liability: Short and Long Term Evidence from Philippine Microcredit Lending Groups. *Working paper*.

Glaeser, E. L. and A. Shleifer (2001). Not-for-profit entrepreneurs. *Journal of Public Economics 81*(1), 99–115.

Greif, A. (1993). Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition. *The American Economic Review 83*(3), 525–548.

Guiteras, R. P. and B. K. Jack (2014). Incentives, Selection and Productivity in Labor Markets: Evidence from Rural Malawi. *NBER Working Paper 19825*.

Hannan, R. L., V. B. Hoffman, and D. V. Moser (2005). Bonus Versus Penalty: Does Contract Frame Affect Employee Effort? In A. Rapoport and R. Zwick (Eds.), *Experimental Business Research, Vol II*, Volume II, Chapter 8, pp. 151–169. Netherlands: Springer.

Herweg, F., D. Müller, and P. Weinschenk (2010, December). Binary Payment Schemes: Moral Hazard and Loss Aversion. *American Economic Review 100*(5), 2451–2477.

Hilken, K., K. D. Jaegher, and M. Jegers (2013). Strategic Framing in Contracts. *Tjalling C. Koopmans Research Institute Discussion Paper Series 13-04*.

Hoff, K. and J. Stiglitz (1997, April). Moneylenders and bankers: price-increasing subsidies in a monopolistically competitive market. *Journal of Development Economics 52*(2), 429–462.

Horton, J. J., D. G. Rand, and R. J. Zeckhauser (2011, February). The online laboratory: conducting experiments in a real labor market. *Experimental Economics 14*(3), 399–425.

Hossain, T. and J. A. List (2012, July). The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations. *Management Science 1909*, 1–17.

Ipeirotis, P. (2010). Demographics of Mechanical Turk. *Working Paper CeDER-10-01*.

Jackson, M. O., T. Rodriguez-Barraquer, and X. Tan (2011). Social Capital and Social Quilts: Network Patterns of Favor Exchange. *American Economic Review forthcomin*.

Jain, S. and G. Mansuri (2003, January). A little at a time: the use of regularly scheduled repayments in microfinance programs*. *Journal of Development Economics*.

Jayaraman, R., D. Ray, and F. de Vericourt (2014). Productivity Response to a Contract Change. *NBER working paper 19849*.

Kabeer, N. (2005). Is Microfinance a 'Magic Bullet' for Women's Empowerment? Analysis of Findings from South Asia. *Economic and Political Weekly 40*(44/45), 4709–4718.

Kahneman, D. and A. Tversky (1979, January). Prospect Theory: An Analysis of Decision under Risk. *Econometrica 47*(2), 263–292.

Karlan, D., M. Mobius, T. Rosenblat, and A. Szeidl (2009). Trust and Social Collateral. *Quarterly Journal of Economics 124*(3), 1307–1361.

Karlan, D. S. (2007). Social connections and group banking. *Economic Journal 117*(517), F52–F84.

Karlan, D. S. and J. Zinman (2009). Observing Unobservables: Identifying Information Asymmetries With a Consumer Credit Field Experiment. *Econometrica 77*(6), 1993–2008.

Kaur, S., M. Kremer, and S. Mullainathan (2013). Self-Control at Work. *mimeo*.

Kőszegi, B. and M. Rabin (2006, November). A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics 121*(4), 1133–1165.

Kőszegi, B. and M. Rabin (2007, September). Reference-Dependent Risk Attitudes. *American Economic Review 97*(4), 1047–1073.

Laffont, J. J. (2003, January). Collusion and group lending with adverse selection. *Journal of Development Economics*.

Laffont, J.-J. and P. Rey (2003). Moral Hazard, Collusion and Group Lending. *mimeo*.

Lazear, E. (2000, February). Performance pay and productivity. *The American Economic Review 90*(5), 1346–1361.

Lazear, E. P. (1995). *Personnel Economics*. MIT Press.

Luft, J. (1994, September). Bonus and penalty incentives contract choice by employees. *Journal of Accounting and Economics 18*(2), 181–206.

Mas, A. (2006). Pay, Reference Points, and Police Performance. *NBER Working Paper No. 12202*.

Mason, W. and S. Suri (2012, March). Conducting behavioral research on Amazon's Mechanical Turk. *Behavior research methods 44*(1), 1–23.

McIntosh, C. and B. Wydick (2005). Competition and microfinance. *Journal of Development Economics 78*(2), 271–298.

Paolacci, G., J. Chandler, and P. Ipeirotis (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making 5*(5), 411–419.

Pope, D. and M. Schweitzer (2011). Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes. *The American Economic Review 101*(February), 129–157.

Propper, C. and J. Van Reenen (2010). Can pay regulation kill? panel data evidence on the effect of labor markets on hospital performance. *Journal of Political Economy 118*(2), 222–273.

Rabin, M. (2000). Risk Aversion and Expected-Utility Theory: A Calibration Theorem. *Econometrica 68*(5), 1281–1292.

Rai, A. and T. Sjöström (2010). Redesigning Microfinance. *Working paper*.

Rai, A. S. and T. Sjöström (2004). Is Grameen Lending Efficient? Repayment Incentives and Insurance in Village Economies. *Review of Economic Studies 71*(1), 217–234.

Sanyal, P. (2009, January). From Credit to Collective Action: The Role of Microfinance in Promoting Women's Social Capital and Normative Influence. *American Sociological Review 74*(4), 529–550.

Shapiro, C. and J. E. Stiglitz (1984). Equilibrium unemployment as a worker discipline device. *American Economic Review 74*(3), 433–444.

Sinclair, H. (2012). *Confessions of a Microfinance Heretic: How Microlending Lost Its Way and Betrayed the Poor*. San Francisco, CA: Berrett-Koehler Publishers.

Spagnolo, G. (1999, January). Social relations and cooperation in organizations. *Journal of Economic Behavior & Organization 38*(1), 1–25.

Tversky, A. and D. Kahneman (1992, October). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty 5*(4), 297–323.

Udry, C. and S. Anagol (2006). The Return to Capital in Ghana. *Economic Growth Center Discussion Paper No. 932*.

Von Thadden, E.-L. and X. Zhao (2012, January). Incentives for Unaware Agents. *The Review of Economic Studies 79*(3), 1151–1174.

Wydick, B. (1999, July). Can Social Cohesion be Harnessed to Repair Market Failures? Evidence from Group Lending in Guatemala. *The Economic Journal 109*(457), 463–475.