Self Selection and Market Power in Risk Sharing Contracts*

Kislaya Prasad†
University of Maryland

Timothy C. Salmon‡
Florida State University

January 2007

Abstract

There is now a well established literature dealing with the theoretical aspects of contract formation and the principal-agent problem. Yet empirical testing of this theory has yielded results that are mixed at best. We use economic experiments to test for the empirical relevance of three possible confounding factors that could explain why the theory has not received stronger empirical support. First, parameters of interest may be biased if agents self-select into projects with differing risk profiles based on risk preferences. Second, differing levels of market power on either side of the market could play a confounding role. Finally, there is the prospect that preferences for fairness or reciprocity affect price formation, making fundamentals such as risk aversion and project risk profile secondary, if not irrelevant. In general, we find support for classical contract theory augmented to accommodate self-selection based on risk preferences. We also find that market power has a significant impact on wages. We find little to suggest that a preference for fairness or reciprocity motivates behavior. We identify a set of conditions—namely, an excess supply of agents—under which the predictions of the classical contracting model are very closely approximated.

JEL Codes: D86, C90
Key Words: Incentive contracts, principal-agent model, self-selection, market power, experiments.

1 Introduction

Risk sharing arrangements take many different forms and exist in a wide variety of situations. The optimal allocation of risk is of course central to the study of insurance and stock markets. Risk-sharing is also important in understanding the employment relationship, occupational choice, entrepreneurship and the optimal organization of production. In primitive societies non-market institutions allowed individuals to pool risks and smooth

*The authors would like to thank various attendees at the 2007 AEA Winter meetings, the 2006 ESA International meetings as well as Svetlana Pevnitskaya for many useful comments and suggests. We also thank the National Science Foundation for research support in funding these experiments.

†Robert H. Smith School of Business, University of Maryland, College Park, MD 20742, kprasad@rhsmith.umd.edu. Phone: 301-405-9637 Fax: 301-405-9655

‡Department of Economics, Florida State University, Tallahassee, FL, 32306-2180, tsalmon@fsu.edu. Phone: 850-644-7207 Fax: 850-644-4535.

food consumption over time. Nowadays, doctors and lawyers form medical and legal practices that share earnings risks among the partners. But as an individual's earnings become divorced from their productive investments and effort, the optimal sharing of risks creates a conflict with incentives. Such a conflict is inevitable when investments are unobservable, which is often the case when there is uncertainty in the environment. The study of this conflict is central to contract theory, which posits a fundamental trade-off between risk and incentives: the optimal incentive intensity must necessarily be a compromise between the twin goals of providing optimal incentives and sharing risks efficiently. Starting from this result, an extensive theoretical literature has developed. But, as Chiappori and Salanie (2000) assert, empirical validation of the theory has lagged behind. In this paper, we describe the results of an experimental study designed to test the theory's predictions along multiple dimensions. The principal focus, for reasons articulated below, is on two questions. First, when there are tasks of varying riskiness available, is there any evidence of endogenous matching of people to tasks on the basis of risk preferences? And second, are there tangible effects when the environment is changed in a way that affects the relative market power of contracting parties? Additionally, we test the theory's predictions about contract terms, and examine whether effort choices respond to incentives.

A straightforward case of a risk sharing relationship that has received much attention is the standard principal–agent game in which a principal wishes to hire an agent to perform some task whose outcome is subject to risk. The economic problem is how to design a contract that will allow the principal to induce the agent to make decisions that maximize the welfare of the principal while ensuring that the agent is still willing to participate. Classic analysis of the theoretical problem dates back to Mirrlees (1974), Mirrlees (1975) and Holmstrom (1979), and the now extensive theoretical literature is surveyed in Bolton and Dewatripont (2005). Empirical evidence to date provides support that is mixed at best. Most notably, a negative relationship between risk and incentives, whereby greater uncertainty in the environment leads to weaker incentive intensity, is hard to detect in the data. Prendergast (2002) summarizes the state of affairs: "Empirical work testing for a negative trade-off between risk and incentives has not had much success: the data suggest a positive relationship between measures of uncertainty and incentives rather than the posited negative trade-off." One explanation of the anomaly relies on the fact that there may well be heterogeneity of risk preferences in the population of workers. Chiappori and Salanie (2000) emphasize the need to control for such unobserved heterogeneity, failing which "the combination of unobserved heterogeneity and endogenous matching of agents to contracts is bound to create selection biases on the parameters of interest."

The problem can be illustrated by considering the example of sharecropping. Tenants bear all of the risk with fixed rent contracts, while risks are shared with sharecropping contracts. A standard prediction of incentive theory is that sharecropping contracts are more likely to be associated with more risky crops.[1] This reasoning leads researchers to run regressions of contract choice on some measure of the riskiness of the crop. Contrary

[1] In the context of the model of Holmstrom and Milgrom (1987), "the optimal piece rate is $1/[1+rC''(e)\sigma^2]$, where r is the degree of absolute risk aversion, $C(e)$ is the cost of effort of the agent, and σ^2 is the variance of the measurement error on performance" Prendergast (2002), p. 1076). Tests of the theory are based on the observation that the optimal incentive intensity is inversely related to the variance of the project. Note that the formula holds exactly only for the case of linear contracts in the normal–exponential model.

to the predictions of incentive theory, Allen and Lueck (1992) shows that the fraction of output that goes to sharecroppers is *positively* related to measures of riskiness. Studies of executive pay and franchising lead to similarly disturbing results. A possible explanation for the anomaly is the absence of adequate controls for risk-aversion. The regression of contract shares on crop riskiness is valid conditional on risk-aversion, but if agents select crops as well as contracts based on their risk preferences, then the endogeneity would need to be accounted for. In fact, as Chiappori and Salanie (2000) show, such self-selection could easily lead to a positive relationship, as the more risk-tolerant individuals choose both more risky crops and more risky (fixed rent, rather than sharecropping) contracts. A paper that attempts to correct for bias arising from self-selection is Ackerberg and Botticini (2002), who use a data set on agricultural contracts from early Renaissance Tuscany. Risk aversion coefficients being unavailable, they take wealth as a proxy. If the more wealthy were also less risk-averse, they would be more likely to choose fixed rent contracts. But wealth is an imperfect proxy, and they need to control for endogenous matching (which is done by instrumenting crop riskiness). The results obtained are more supportive of contract theory, suggesting that endogenous matching could be an empirically relevant phenomenon.

Our work is motivated in part by this empirical literature on contracts. Whether matching by risk–preferences occurs is, in and of itself, an important question. But a systematic look at the phenomenon would also be useful for the help it could provide in the empirical study of contracts. If we are to adequately control for matching or self-selection, we need to know more about the mechanics of the process, and how it depends on institutional details of the environment. One objective of the present experimental study is to detect whether endogenous matching of agents to contracts is likely to occur. To this end we consider a standard principal–agent game in which the principal controls two types of tasks, safe and risky, for which he is attempting to contract agents to perform. We measure the degree of risk–aversion of agents using their choices between lotteries using a technique developed in Holt and Laury (2002), and examine whether more risk averse agents end up on safer projects. Beyond this, we also consider a variety of environments by varying the market power of each side of the market to develop a better understanding of the circumstances that would make endogenous matching more or less likely. Players in our experiment negotiate contracts, and the data on contracts allows for a direct test of contract theory. We can test whether terms of contracts vary in predicted ways with the measured degree of risk aversion of principals and agents.[2] Finally, given the negotiated contracts, agents are allowed to choose investment or effort, and we examine the extent to which these are consistent with incentives.

In addition to the empirical literature on contract theory there is a growing literature on using economic experiments to analyze contract formation in greater detail (see Fehr and Falk (2002), Fehr and Gachter (2002), Guth, Klose, Konigstein, and Schwalbach (1998) and Fehr, Gachter, and Kirchsteiger (1997)). The thrust of many of these papers is that the traditional contract design literature overlooks key aspects of human behavior. Traditional theory assumes purely self-interested preferences, while these studies find that people have other motives such as inequity aversion, or a sense of fairness, or "the desire to reciprocate

[2] We should point out that we do not test for the negative relationship between risk and incentives. That is an indirect test, used because risk preferences are typically unavailable, whereas we measure the degree of risk aversion and can conduct more direct tests.

or the desire to avoid social disapproval", Fehr and Falk (2002). The practical implication for contracting relationships is that principals might offer contracts that are more generous than they need to be to satisfy an agent's participation constraint, and some agents may be willing to engage in high effort even if the offered contract does not satisfy the incentive compatibility constraint. This suggests an important complication for the field literature in attempting to examine contract formation. If in fact these issues of other regarding preferences are leading to contracts that do not satisfy the constraints in the way researchers expect, it could lead them to make incorrect inferences. This is particularly important in regards to the degree to which agents self-select into tasks or are matched with principals. If principals are offering contracts that are "too generous," they may attract even relatively risk averse individuals to relatively risky projects because the overall terms are generous enough to make up for the agent's dislike of risk. It is unclear whether the logic of such models generate self–selection and, if they do, what the nature of this selection is. It is also not clear how contract terms would vary, if at all, with fundamentals such as risk preferences.

There is, however, an additional and potentially important detail left unexamined in most of these previous experimental studies. Typically there is little consideration given to competition between agents for a task, or competition between principals for an agent. A notable exception is Brandts and Charness (2003), who examine the impact of market conditions in a gift exchange game between workers and managers and find it not to matter. In the conclusion we will return to addressing the difference between our study and Brandts and Charness (2003). Most of the other studies use a structure in which supply equals demand in the sense that there is typically one principal bargaining with a single agent over the terms of a single contract.[3] The findings are generally that the subgame perfect equilibrium does not hold. This is not surprising. A now well-established artifact of the experimental literature, summarized in Roth (1995), is that observed payoffs in ultimatum bargaining tend to be less extreme than would be predicted by theory. One would expect some version of this result to carry over to the case of negotiation over a contract as well. The market prediction in such cases is that the terms of the contract should be between the willingness to pay of the principal and the willingness to accept of the agent and thus the findings are consistent with this standard model. While subgame perfection may be a compelling argument to game theorists, evidence from countless experiments shows that it may not be as powerful to others.

The unanswered question relates to how the forces determining negotiated contracts (be it classical rationality, or "other regarding" preferences) react to the presence of competition. The bargaining position of a player becomes stronger if there is competition on the other side of the transaction, and we would expect them to profit from this. As they do so, and drive others players down to reservation utility levels, some players may exit from the competition before others. To the extent that fundamentals influence the exit decision, variables such as risk preferences may begin to have explanatory power. A reasonable conjecture, which was the basis for the design of our study, is that market competition induced by a mismatch between the numbers of principals and agents will be more likely to generate outcomes of the kind predicted by traditional economic models of contracts than is a design with bargaining

[3] Some studies of principal agent games (e.g. Guth, Konigstein, Kovacs, and Zala (2001); Konigstein and Zala-Mezo (2003)) do allow for two agents, but the principal can feasibly contract with all of them.

power conferred on agents by ultimatum bargaining arrangements.[4] It is precisely in such environments that we should observe endogenous matching of players to contracts, as well as contract terms that vary with fundamentals. To test this hypothesis, we use experiments in which we vary the supply/demand conditions of the market, which allows us to determine the importance of market power.

Our stock situation has one principal bargaining with two agents in an attempt to contract with them to perform two tasks with varying levels of riskiness. We then change the market configuration by adding two treatments. In one we add an additional agent, leading to a situation in which there is more supply of labor than there is demand. The second treatment then adds another principal to the mix, leading to a total of four tasks but only three agents. This situation involves having more demand for labor than there is supply. As expected, these two treatments push the terms of the contract in ways that are less favorable to agents in the first case, and more favorable in the latter. The more general point to emerge is that the predictions of contract theory hold up well when there is competition. In this case, endogenous matching by risk preferences also appears as a robust phenomenon. There are two significant implications from our results for empirical research. First, our results reiterate the need to adequately control for endogenous matching. They also point to the importance of an issue that appears to be overlooked in past empirical studies regarding the institutional environment and the extent to which this is conducive to market competition. Our results clearly indicate that some arrangements should be more conducive to finding the results from standard contract theory than others and gives insight to why the standard predicted relationships may fail to be found in field data.

Our results also tie into another developing line of experimental literature which involves the stability of risk preferences across tasks and also whether such self-selection occurs. An early paper in this area, Isaac and James (2000), finds that risk preferences obtained through a BDM procedure do not match well with those obtained through a first price auction. On the other hand, Ivanova-Stenzel and Salmon (2004) and Ivanova-Stenzel and Salmon (2005) investigate auction choice behavior and find evidence in support of consistency in cross-task risk preferences, though with tasks perhaps more behaviorally similar. These latter two studies do not find any self-selection according to risk preferences into auction formats though there is no strong theoretical prediction that it should be observed in those environments. Palfrey and Pevnitskaya (2004) allows for subjects to choose between participating in a first price auction or accepting a determinate outside option and their results are consistent with self-selection occurring due to risk preferences. That study does not, however, contain a secondary measurement of risk preferences which would allow for a stronger confirmation of that result. There is also evidence in favor of the general temporal stability of risk preferences, Andersen, Harrison, Lau, and Rutstrom (2005).

The rest of the paper is organized as follows. In section 2 we describe the experimental design. In section 3 we discuss the theory behind our experiments and outline our hypotheses in some detail. In section 4 we discuss the results and Section 5 brings together our main conclusions.

[4] An uneven number of players could make a difference even when players are motivated by considerations such as fairness or inequity-aversion (Guth, Konigstein, Kovacs, and Zala (2001); Konigstein and Zala-Mezo (2003); Bolton and Ockenfels (2000)). Notions of fairness could be different, and may even reflect the market power of agents.

2 Experimental Design

2.1 Risk Preference Assessment Phase

Our experiments consisted of two phases. The first phase involved running our subjects through the instrument for characterizing risk preferences described in Holt and Laury (2002). The second phase involved the subjects interacting with each other in a principal agent framework in which each principal possessed two tasks he needed to contract with agents to perform. These tasks differed in their risk profiles which could lead to some sort of selection of agents into tasks according to their risk preferences. In the analysis later, we will take the risk preference indicated in the Holt-Laury module as the risk preferences of the subjects. While this is not an uncontroversial assumption, it was the best mechanism for the task. We did our best to make the choices among lottery options in the principal agent framework as similar to those as in the Holt-Laury treatment as feasible in an attempt to maximize the degree to which the risk preference identified through that technique should transfer to the principal-agent context.[5]

The Holt-Laury technique consists of having subjects make a choice for 10 different pairs of lotteries. At the end, one of those lotteries is chosen at random to be played for actual earnings. Each lottery pair involves choosing either lottery A or lottery B. Lottery A can result in one of two outcomes, \$2.00 or \$1.60, while lottery B can result in one of two different outcomes, \$3.85 or \$0.10. The first lottery choice gives the subjects a 10% chance at the good outcome and a 90% chance at the bad outcome for whichever lottery they pick. In that case, it is clearly a good idea to pick lottery A with expected value of \$1. 64 instead of lottery B with expected value of \$0.475. The other nine lotteries involve increasing the chance of the good outcome occurring by 10% each time (and of course decreasing the chance of the bad outcome) until at the end there is a 100% chance of the good outcome. At that point, it is clear that choosing lottery B is the best choice as \$3.85 for certain is better than \$2.00 for certain regardless of your risk preferences. The question is at what point along that path a subject will choose to switch from picking lottery A to lottery B. The earlier a subject switches, the more risk preferring are his preferences while the later the more risk averse.

Our implementation differed from the Holt-Laury implementation mainly because we computerized the choices using z-Tree (Fischbacher (1999)) which in this case did not allow subjects to go back and review decisions as they were able to do in the original experiments. This likely lead to more inconsistency in the answers as it was probably not as clear to the subjects that they were supposed to have a single switch-over point and they could not change a previous response to ensure they had only a single switch-point if they figured it out late in the set of choices. We recognize this is an issue, but do not see it as much of a problem. We should actually be able to use any consistency in choices in this part of the experiment as a proxy for the general level of sophistication of a subject that may allow us to get some traction on understanding some behavior in the other phase of the experiment.

[5]This is due in part to the noted results found in Isaac and James (2000) showing that risk aversion estimates obtained through different tasks may show little similarity.

2.2 Principal-Agent Phase

In the principal-agent phase of the experiment, each principal has two tasks that they are attempting to write a contract with an agent to have an agent perform. In the experiment, we use the labels "manager" for principal and "worker" for agent and generally used language specifically to evoke the context of a manager proposing a contract to a worker. The reason we chose to go with the contextualized framing is that the environment is fairly cognitively demanding for the subjects to interact in and it is our belief that the added context was useful in orienting them to the situation and helping them to process the information.

The two tasks owned by the principal were referred to as tasks A and B. Once an agent is contracted to a task he or she has to choose between two different ways of performing the task, described as X and Y to the subjects but corresponding to high and low effort. There are two outcome states that can result from an agent's choice labeled as "Success" or "Failure."[6] Table 1 outlines all of the costs and benefits to both principal and agent based on the combination of choice by agent and the outcome state. A table with exactly this form was provided to and explained to the subjects in the experiment. The probability of the two states occurring was a function of the agent's choice. If the worker chose X, success occurred with an 80% probably and Failure 20%. If the worker chose B, success occurred with 25% probability and failure 75%. It was more costly for the worker to choose X resulting in a direct cost of 35 cents for it while Y only cost 5 cents. The state probabilities only depended on the effort choice of the agent, not on the task. Task A was the relatively risky task because the principal would receive 240 cents if success occurs but only 40 cents if failure occurred. Task B was safer due to the fact that the principal received 205 cents upon a success and 180 cents in the result of a failure occurring. Both principals and agents possessed outside options. In any round that an agent did not end up with a contract, they earned 50 cents. For a task that a principal is unable to contract an agent to perform, the principal earns 25 cents per task. This gives both sides of the market equivalent outside options and therefore equivalent minimum possible earnings.

The contract between a principal and an agent consists of two parts: a fixed wage, w, to be paid regardless of the outcome and a bonus, b, paid only if the success occurred. The managers earnings were based on the money earned from the outcome state less the wages paid to the worker. The workers' earnings were determined by the terms of the contract. In the instructions for the experiment we gave the subjects a handout with a sample set of contracts and went carefully through an explanation of the consequences to each party from any choice by the agent. In our software demonstration, we also allowed all of the subjects, including those who would become principals, to spend time making practice decisions of X vs Y for a variety of randomly chosen task and contract arrangements so that they could form a better understanding of how the incentives fit together. A full set of our instruction scripts can be found in a supplemental appendix.

The bargaining framework of our experiments is substantially richer than what is found in standard ultimatum game based experiments. The reason for that is that given the

[6]There is a reasonable question about the advisability of using such contextually loaded terms in an experiment of this nature. As explained elsewhere, our general view is that the context was necessary to assist subjects in understanding a complex design. Still using the terms "success" and "failure" could well bias subjects toward the high effort choice to avoid a state described as "failure." As our later results will show, such an effect does not seem to have been created.

		Task A	Task B
If Worker Chooses X	**Probability Success**	80 %	
	Probability Failure	20 %	
	Cost to Worker	$0.35	
If Worker Chooses Y	**Probability Success**	25%	
	Probability Failure	75%	
	Cost to Worker	$0.05	
Earnings to Manager for Success		$2.40	$2.05
Earnings to Manager for Failure		$0.40	$1.80
Earnings to Manager if No Worker Contracted		$0.25	$0.25
Earnings to Worker if No Contract Agreed To		$0.50	

Table 1: Parameters for experiment environment.

complexities of this environment we believe it is important for subjects to bargain in an iterative manner as that should be expected to allow the subjects time and the opportunity to explore the possibilities better and to therefore form a better understanding of the underlying incentives. Consequently we allowed bargaining to proceed by the principals making initial contract offers with agents indicating their interest in those offers followed by potential revisions to those offers by the principals until essentially a market equilibrium is realized. In the beginning of the principal-agent phase, subjects would be assigned a role as either worker or manager and this would remain persistent throughout the rest of the experiment. Once assigned roles, subjects would then be assigned into groups according to the arrangements used for that treatment. These group assignments would also remain persistent through the rest of the experiment which means that each group was completely independent of the other groups in the session and can be treated as such in the statistical analysis.

Each round would begin with principals making an initial set of contract offers on their tasks. Agents would see those offers and be able to click a button to indicate in which, if either, they were interested. Agents could only indicate interest in a single contract/task at any point in time. Principals would observe these interest states dynamically and were allowed to alter the terms of any offer if there were currently no agents interested in it or if there was more than one. Principals could not alter the terms of an offer if only a single agent was currently interested. Agents were allowed to change their interest state (including to a "no interest in any contract" state) at any time of their choosing. The idea of course, is that if no agents are interested the principal can make the terms more generous to attract an agent while if there is more than one agent interested they have the bargaining position to make the offer less generous. A bargaining round could end in two ways. First, there was a five minute hard close rule which ended the round in whatever configuration existed at the end of five minutes. Second, after the principals put in their first offers a 30 second countdown clock began. This clock would reset each time a principal changed the terms of an offer or an agent changed their interest state. If the clock reached 0, the round would end for that group before the five minute hard close. Regardless of how the round closed, any agents who were the only ones interested in a task at the close would be assigned that

Treatment	Group Composition	Num Sessions	Total Groups
Baseline	1 Principal - 2 Agents	4	15
Excess Supply	1 Principal - 3 Agents	5	15
Excess Demand	2 Principals - 3 Agents	6	16

Table 2: List of experimental treatments.

task while if more than one is currently interested in the same task the system randomly allocated the task to one while the other(s) received the outside option. Any tasks left without an agent went unfulfilled leaving the principal to only earn their outside option on the unfulfilled task. Any agents receiving a task would then be asked to make a choice regarding how to perform the task, choose X or Y. Then the agent was informed in detail about the outcome while the principal only observed whether the outcome was a success or failure. A principal never observed the X vs Y choice of the agent. This phase consisted of 8 rounds.

To examine market competition issues more carefully, as well as to get a better overall picture of contract formation, we conducted these experiments according to three different treatments as summarized in table 2. Our Baseline treatment is the one most similar to prior principal-agent experiments in that there are an equal number of agents and tasks, thus supply is exactly equal to demand. The main difference from previous experiments is the fact that each principal has two tasks and the agents will get to choose among those tasks and this could lead to some differences in behavior by itself. Our other two treatments are Excess Supply (ES) and Excess Demand (ED). In ES we increase the number of agents by one in each grouping so that there are more agents than tasks. This should give the principal maximal bargaining/market power and the ability to force subjects down to their willingness to pay if he desires. The ED treatment involves each group being composed of 2 principals and 3 agents which means there is one more task than there are agents to fulfill them. Now the agents should have maximal bargaining/market power and may be able to force the principals up to their willingness to pay.

The ED treatment should be the one with the least risk selection occurring as if the agents can force the principals to make generous offers, the incentives for risk selection should be small. If self selection according to risk preferences will occur in any treatment, it should occur in the ES treatment since this should be the treatment most likely for the standard optimal contracts to be realized. The Baseline treatment could go either way.

Subjects for these experiments were recruited from wide range of courses, mostly in economics. All subjects received a $10 show-up fee in addition to their earnings from the experiments. All monetary amounts mentioned in the experiment, and those described above, were denominated in US $. Sessions lasted about one and a half hours. Average total earnings per subject were around $20.50 with a maximum of around $32 and minimum of $12.

3 Theory and Hypotheses

The theoretically optimal contracts in this environment assuming risk neutral players are quite straightforward. We first solve for the standard principal-optimal contracts that are found by maximizing principal welfare subject to the standard incentive compatibility and participation constraints. For task B, it is not in the principal's interest to induce high effort, i.e. a choice of X, so the optimal contract involves setting $b_B^* = 0$. The only requirement is that w_B^* must satisfy the participation constraint assuming a choice of Y, so $w_B^* = \$0.55$. For task A, it is in the principal's interest to induce a risk neutral agent to choose high effort. The optimal contract for doing so implies setting $w_A^* = .41$ and $b_A^* = .55$. Given a contract with this structure, a risk neutral agent is just better off accepting the contract and choosing X rather than either rejecting the contract or accepting and choosing Y. If we introduce risk aversion on the part of the principal (but not the agents), the optimal contract for A has the principal keeping the fixed amount of \$1.15 and letting the agent have the rest. This corresponds to a fixed payment of -\$0.75 and a bonus of \$2.00. Similarly, the optimal contract on B involves the principal keeping \$1.31 and giving the agent the remainder. If we introduce risk aversion on the part of the agents (but not the principal), then the optimal contract for task B remains unchanged from the initial contract of $w_B^* = \$0.55$ and $b_B^* = 0$ because it is a fixed wage contract. The contract for task A must be made more generous in order to compensate the agent for bearing risk. The exact terms will depend upon functional form assumptions and the degree of risk aversion. If the agent's risk aversion is too great, then the principal may be better off inducing low effort.

If risk-aversion coefficients of agents are identical then it does not matter how agents are assigned to tasks. But if these are different, one might expect some selection between tasks based on the risk preferences of the agents. The principal would like to find the least risk averse agent possible to contract with over task A because it is that agent that will require the least amount of extra inducement to accept the task and work with high effort. That would leave task B for the relatively more risk averse agent. Efficiency dictates this—in a full information world with observable efforts, total surplus is maximized by exactly such an assignment. Even with unobservable efforts, such sorting occurs under very general conditions (the issue of self-selection from a menu of contracts in an environment such as ours was studied first by Rothschild and Stiglitz (1976)). Selection into tasks by risk preference will occur if the contracts offered by the principals satisfy the condition that the contract for the more risky project is attractive enough that it is desirable to the less risk averse agent, but not so attractive that it draws the more risk-averse agent. It is important to observe that principals do not have any prior knowledge of risk preferences of agents, and these are revealed only during the process of reaching an agreement. The iterative procedure that we have designed not only mimics how negotiations might actually proceed in the real world, but also facilitates discovery of risk-preferences.

The contracts computed above are optimal from the point of view of the principal. The agent ends up being indifferent between taking his outside option and working for the principal. When contracts are bargained over in a market setting, it is not certain that this will be the agreed upon structure. For one, the agent could have leverage in bargaining. There is a wide range of potential contracts that give the principal and agent at least as much as their outside option. A better bargaining position would give the agent a more generous contract from this set. We find the ultimatum game and other bargaining protocols unappealing as a way in which to create bargaining power for our purposes.[7] Instead we place the contracting relationship within a framework of market competition, and create bargaining power by manipulating the relative scarcity of subjects on each side of the transaction. We have two types of projects on which the principal is allowed to offer different contracts. In our baseline treatment there are as many agents as there are tasks. Market opportunities for the two types of players are quite different when there is a mismatch in the number of projects and agents. Suppose there are only two projects, but three agents, as in our Excess Supply treatment. Now at least one agent is guaranteed to be left without a project. With market motives, this should create more intense competition among agents, leading to less generous contracts. In our Excess Demand treatment there are two principals, each with two tasks, competing for three agents. In this case at least one task must go unfulfilled and the principals must compete to attract agents to their tasks. Competition for workers is most intense in this excess demand case, which should lead to contracts that are most generous to agents. We examine whether the generosity of contracts varies with the market power of players. Are payments to agents smaller in the Excess Supply case? Are they larger in the Excess Demand case? We would certainly expect so, leading us to the *market power* hypothesis:

Prediction 1 - *The contracts for the Baseline and the Excess Demand treatment will be more generous to the agent than for the Excess Supply treatment, with the terms being most generous in the Excess Demand treatment.*

Among the three treatments, the Excess Supply case gives the principal the most bargaining power. If agents are driven down to their reservation utility levels, this should yield negotiated contracts closest to those computed at the beginning of this section. The following prediction is based on the assumption of agent risk neutrality, but with typical assumptions about functional form the results should be close.

Prediction 2 - *In the Excess Supply treatment the contract for task A will have a fixed wage of 41 and bonus of 55. The contract for task B will have a fixed wage of 55 and a bonus of 0 or at least less than 43. Optimal efforts are still elicited if the bonus on task B is something less than 43.*

Creating market competition could affect the scope of the selection argument. More generous contracts could eliminate the need for agents to select into tasks by risk preferences. This would suggest that the best chance for endogenous matching of agents to contracts would occur in the Excess Supply case. Of course there still could be some selection pressure if contracts on the risky task A involved greater variation in payments across states (for whatever reason). While more risk averse agents may be willing to accept the task at those terms, they may still find the relatively less risky task B to be the better alternative. Since we measure agents' risk aversion, and observe their choice of projects, we can test the *endogenous matching* hypothesis:

[7] We have two objections to it, one practical and one philosophical. The practical issue is that it is not clear what is a tractable way in which one can study selection by risk-preferences within such an environment. The philosophical point is that market outcomes require market competition. The setting of market scarcity thus seems to be a much more compelling source of bargaining power, and would be more likely to generate outcomes consistent with the classical economic model.

Prediction 3 - *More risk tolerant agents should end up working on task A.*

In the principal-agent model, the role of fundamental variables such as risk aversion has been clearly articulated. For instance, the footnote in the introduction suggests that the bonus should be inversely related to the agent's risk aversion parameter. That formula is derived for a continuous effort model, but the basic logic can be applied in our framework. To provide incentives it is necessary to shift some risk on to agents. When only the agent is risk-averse, the optimal contract departs minimally from fixed payments. The risk imposes a cost on the agent for which he must be compensated. If this cost, which increases with the degree of risk aversion, is sufficiently high the principal may well settle for a lower level of effort, requiring smaller incentive intensity but better sharing risks. In our parametrization, the likelihood of decreasing efforts is made negligible. But the rest of the logic of efficient risk-sharing remains. In the typical interaction, both principals and agents are risk-averse. All other things being equal, an increase in the risk aversion coefficient of a subject should lead them to more highly value a reduction in variation in income. So a more risk averse agent would try to induce a higher fixed payment with perhaps a relatively lower bonus, while a more risk averse principal would attempt to negotiate for perhaps a relatively lower fixed payment but a higher bonus. Under a range of scenarios, from among the many different contracts that implement a desired level of effort, negotiations should lead to contracts that reflect this relative preference.[8] So higher agent risk aversion should imply a contract with higher fixed payments, and smaller bonus. Similarly, higher principal risk aversion should lead to a contract with smaller fixed payments and higher bonus. This logic of efficient risk-sharing also suggests why we should not expect only fixed payments on task B when the principal is risk-averse. There may be room for mutually beneficial trade where the agent takes on some of the risk via a positive bonus in return for a slightly higher expected payment. So long as the bonus does not provide perverse effort incentives, such a contract would be efficient. This leads to the *efficient risk sharing* hypothesis:

Prediction 4 - *The bonus should be negatively related to the degree of risk aversion of the agent and should be increasing in the risk aversion of the principal, while the fixed wage should be the opposite.*

Once contracts are settled upon, a remaining question concerns the choice of efforts. The standard theory tells us that agents will choose the action which yields the highest net benefits. Risk neutral agents will choose high effort if the gain in expected payments exceed the increased cost of high effort, while risk averse agents will choose high effort if the expected utility of high effort exceeds the expected utility of low effort. This suggests that once the agent has decided to participate effort choices will depend only upon the size of the bonus. We state this as the *optimal effort choice* hypothesis:

Prediction 5 - *Conditional upon accepting a contract, effort choice should depend only upon the bonus offered under the contract.*

All of the predictions enumerated above are consistent with the standard model of contracting, driven by individual rationality and incentives. People may also be guided by motives such as inequity aversion, fairness, reciprocity etc. For instance, high effort choices may be driven by a norm (for "gift-exchange"), in which the principal offers a high payment

[8] It is easy to generalize the formula on optimal incentive intensity from footnote 1 and allow for risk averse principals. The optimal intensity then increases with the principal's risk-aversion coefficient.

and the agent provides high effort in return. One test of the presence of such a norm would be if instances of high fixed payments (but insufficient bonuses) are related with high efforts. This would be an instance of a "gift exchange" norm, or a form of reciprocity where people are nice to others who are also nice. Reciprocity can also take the form of a preference for punishing people who are not nice. The literature on such non-classical preferences does not take a consistent position on how behavior responds to things like changes in market power or risk preferences. But there is some suggestion in models with norms that fundamentals should matter less, e.g. Young and Burke (2001). Anyhow there is little in these models to suggest that the above predictions should hold. For instance, consider market power in the Excess Supply treatment. If a principal finds that there are multiple agents interested in one of her projects, should she exploit this market position by changing terms in a manner detrimental to agents? If notions of fairness are made flexible, so that the notion of what is fair reflects every market advantage, it is unlikely to have much predictive content. So indirectly, the success of the above predictions would cast doubt on these alternative explanations for the institutional arrangements of our experiments. We have little doubt that there are situations in which these alternative motives would be important, just as we are persuaded by the results of Young and Burke (2001) on sharecropping norms. As we see it, the compelling task is to develop a mapping between institutions and the forms of behavior they engender. Findings on the above predictions will move us forward in this effort.

4 Results

4.1 Risk Preferences of the Subjects

The first issue we will address is the general nature of the risk preferences of our subjects. The manner in which the Holt-Laury mechanism is intended to measure risk preferences is that subjects are expected to start the sequence of choices by choosing the "safe" option, A, and then at some point they should be willing to switch to the risky option, B. The point at which they switch from choosing one option to the other indicates a subject's willingness to tolerate risk. The difficulty with applying this technique is that some subjects will switch back and forth between the two options more than once along the decision path. That makes it tricky to identify precisely the nature of the risk preferences of those subjects. We have chosen to identify the risk preferences of our subjects by using a measure equal to the question number at which they first begin a series of unrevoked choices of the risky option, B. For the subjects who switch only once, this is the ideal measure of their risk preferences. For subjects who switch back and forth between A and B multiple times, it is not clear what the ideal measure would be so we will use this one with full understanding that it may be a less precise measure for these subjects.[9] Further, we choose not to use some specific utility function to deliver an estimate of a specific risk aversion parameter for our subjects. Such a step is unnecessary for our purposes and would needlessly impose a parametric assumption. This data allow us to establish our first important result.

[9] In Holt and Laury (2002), their measure was the total number of safe choices. We have checked the robustness of our regressions to this alternate definition and find little in the way of differences.

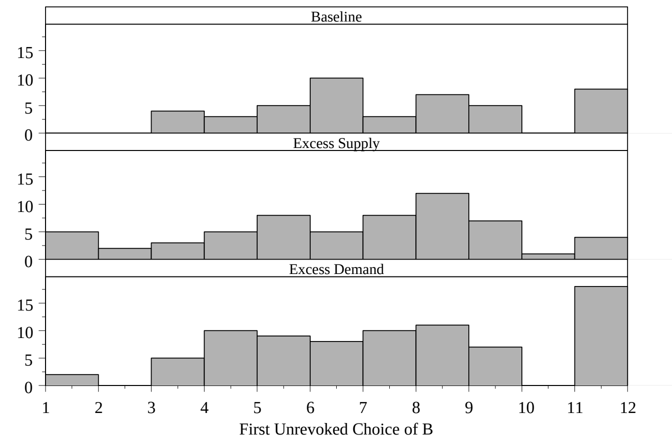

Figure 1: Histograms of risk preferences of subjects by treatment.

Result 1 - *There are no systematic differences in risk preferences between the subjects used for different treatments.*

Using this measure of risk preferences we find that the subjects used in different treatments do not appear to have structurally different risk preferences. In our baseline treatment we find an average of 7.09 for our Baseline treatment, for our Excess Supply treatment 6.28 and 7.11 for Excess Demand. A histogram of the population of values is found in figure 1. In each case, the distributions are approximately uniform over the range $[3, 10]$. We do see some subjects with a risk preference measure of 11. These correspond to subjects who never switched to the risky option irrevocably even on the last question. Wilcoxon ranked sum tests show that there is no statistically significant difference between these distributions.[10]

In regard to the issue of subjects switching more than "rational," 90 out of 185 subjects across all treatments exhibit the ideal pattern of a single switch while another 41 switch 3 times indicating perhaps a minor deviation in a range due to the expected utilities of the two options being too close to bother with precision. This means that around 54 people or 30% of our sample chose in ways that are not consistent with the model presumed by this measurement technique. This is higher than the approximately 10% of such subjects found to do so in the original Holt and Laury (2002) experiments. We suspect that this is mainly due to the difference in elicitation procedure. In our implementation of the procedure, we did not allow subjects to go back and make changes to ensure consistency while the procedures used in Holt and Laury (2002) had the subject making choices for all lottery pairs at once which encourages consistency. While our technique and the observed behavior does perhaps decrease the accuracy with which we can say we have estimated the risk preferences of our

[10]Baseline vs. Excess Supply: $p-$value of 0.233
Baseline vs. Excess Demand: $p-$value of 0.942
Excess Supply vs. Excess Demand: p-value of 0.157

subjects, we believe it also generates a useful measure for us as a proxy for the sophistication of the subject. A sophisticated subject is likely to recognize the structure of the incentives and only make a single switch. Less sophisticated subjects will likely switch more as an indication that they have less of an idea of what they are doing or are in general paying less attention to the details of the experiment (such as one subject who switched 9 times!). If a subject is not paying close attention to the incentives in this part of the experiment, then they are also not likely to be doing so in the main part of the session and so having observed this transparent lack of sophistication may help us in understanding any “erratic” behavior observed in the principal-agent phase of the experiment.

4.2 Impact of Market Structure on Contract Choice and Earnings

Our first prediction, from section 3, is the market power hypothesis. The generosity of contracts is expected to vary with market power of players. Table 3 contains some basic summary statistics regarding the contracts that we observe across treatments. The most striking thing here is that total compensation reflects the predicted effects of market power. If we focus on averages for the last four periods, total promised payments (i.e. fixed wages plus bonuses) are highest in the excess demand case, and lowest in the excess supply case. In the case of averages across all rounds, payments are still highest in the excess demand treatment, but there is not much difference between the baseline and excess supply case. This leads to our second result:

Result 2 - *The fixed wage offer on task B is on average higher than that offered on task A while the reverse is true for the bonus. The levels of the fixed and the bonus wages in the Baseline and Excess Supply treatments are approximately those of Prediction 2. Contracts in Excess Demand treatment are on average more generous to the agent.*

Our prediction for optimal contracts in the excess supply environment are a fixed wage of 41 and bonus of 55 for task A and then a fixed wage of 55 with bonus of 0 for task B (Prediction 2). As can be seen in table 3, the average contract in the baseline and the excess supply treatments are quite close to these optimal levels. Given the complexity of the environment and the interactions, it is quite surprising that the agents appear to converge to these levels. The only important difference is that in both treatments the bonuses on Task B are substantially greater than 0 (but lower than the bonuses on Task A). The indication here is that the principals are able to use their bargaining position in the Baseline treatment and their market power (as well as bargaining position) in the Excess Supply treatment to force the agents down to approximately the value of their outside option as the optimal contract would imply. While the bonuses on the B contract are greater than 0, the only real requirement, for optimal efforts, is they be less than $35/.8 = 43.75$ which would be the level required to induce high effort. On average these bonuses are less than that and the bonus decreases towards the end of the experiments pulling them even farther below this level as the principals learn that they do not want to induce high effort on that task. We conclude that in these two treatments, the standard optimal contract is approximately reached in our sessions.

In the Excess Demand treatment, however, the average bonus for task A is much greater than the minimum necessary to induce high effort and it is rising over the course of the experiment. Similarly the fixed wages for task B are also greater than the minimum required

		Baseline		Excess Supply		Excess Demand	
		Fixed	**Bonus**	**Fixed**	**Bonus**	**Fixed**	**Bonus**
All Rounds	**Task A**	48.25	55.45	40.73	58.00	44.28	70.81
	Task B	60.50	41.99	61.35	37.29	76.14	44.94
Last 4	**Task A**	47.14	55.02	35.21	50.21	38.59	89.21
	Task B	60.00	39.45	52.50	36.04	81.41	46.68

Table 3: Summary statistics on contracts.

	Baseline		Excess Supply		Excess Demand	
	Coeff.	**P-Value**	**Coeff.**	**P-Value**	**Coeff.**	**P-Value**
Constant	230.42	0.275	642.47	<0.001	913.22	<0.001
Principal	884.58	0.026	281.23	0.154	-377.64	0.051
Principal, Risk Pref	-12.86	0.660	17.53	0.416	7.82	0.434
Agent, Risk Pref	40.30	0.048	3.30	0.791	1.80	0.915
Principal, LS	29.67	0.872	339.45	0.014	119.92	0.055
Agent, LS	174.15	0.114	-69.04	0.416	44.71	0.618
Number of Obs	45		60		75	
Adj R^2	0.491		0.636		0.365	

Table 4: OLS regressions of total earnings on likely determinants.

to induce acceptance and are also rising over time. The bonuses on the B tasks are also higher than in the other treatments and in fact are at about the right level, on average, to just induce high effort. This demonstrates that the market power of the agents is working against the bargaining position of the principals leading to the agents getting substantially more generous contracts. It is not the case, however, that the principals are driven to their outside option which would be earning 25 per contract. For task B that would involve a fixed wage of around 160 and a bonus of 0 while for task A that would involve a combination of fixed and bonus wages that would pay the subject 175 on average. Workers are not receiving wages anywhere near this level of generosity even by the end of the session. This is an interesting demonstration of the interaction of bargaining position and market power.

We can also investigate the outcomes of these contracts in terms of their effects on earnings. This analysis is summarized in our next result.

Result 3 - *Principals earn more money in the Baseline and Excess Supply treatments than agents while the opposite is true in the Excess Demand treatment. Risk aversion has little impact on earnings while more sophisticated principals earn more money in the Excess Supply and Demand treatments than non-sophisticated principals.*

Overall average profits per subject for this phase of the experiment did not vary much from treatment to treatment: \$7.86 for Baseline, \$7.76 for Excess Supply and \$7.82 for Excess Demand. On the whole, principals did better than agents. Principals earned \$8.66 and Agents \$7.20, but these varied from treatment to treatment. In the Baseline treatment the earnings were \$10.49 vs. \$6.54 for principals vs. agents and in the excess supply case the earnings were \$12.79 vs \$6.08. In the Excess Demand treatment, the ranking reversed with the Principals now earning only \$6.68 and the Agents \$9.54. In addition to these gross treatment level effects it might also be interesting to determine if subject characteristics impact their overall earnings. Table 4 presents the results of cross-section OLS regressions of total earnings (in cents) over the course of the experiment as a function of whether they were a Principal or Agent (1 if Principal), their risk preference separated by whether they were a principal or agent and then dummy variables equal to 1 if they switched 1 or 3 times in the Holt-Laury procedure ("Low Switchers" or LS), again separated for Principals and Agents. The results clearly show that Principals earn more than Agents in the first two treatments while the reverse is true in the last treatment. Risk preferences appear to only impact earnings for agents in the Baseline treatment. The LS or Low Switcher variable used as a proxy for sophistication shows that sophisticated principals earn more than non-sophisticated principals in the Excess Supply and Excess Demand treatments.

As a start to explaining those earnings differentials, we can examine the negotiated contracts in more careful terms to understand how they may also vary with fundamentals such as risk preferences. As explained in Prediction 4, we expect more risk-averse subjects to have a relative preference for smaller variation in incomes. So more risk averse agents would rather get a larger fixed payment and smaller bonus, whereas more risk averse principals would rather pay smaller fixed payments with larger bonuses.

To develop a better understanding of contract formation, we run regressions of contract terms on the characteristics of players. Our next result summarizes our tests regarding the efficient risk–sharing hypothesis (Prediction 4), on temporal effects, and on the "sophistication" of subjects as measured by the number of times switched (with more switches taken to mean lower sophistication).

Result 4 - *In the Excess Supply treatment, the fixed wage on negotiated contracts is increasing in the risk aversion of the agents and is decreasing in the risk aversion of the principal, while the bonus is decreasing in the risk aversion of the agent and increasing in the risk aversion of the principal. In the Excess Demand treatment, risk preferences of both principal and agent have no impact on the fixed wage while the bonus is decreasing in the risk aversion of the agent and increasing in the risk aversion of the principal. In the baseline treatment we do not find robust support for the efficient risk-sharing hypothesis.*

We examine every completed contract, *viz.* contracts for which a principal and an agent reached agreement. In the ensuing tables, we focus on the Low Switchers, i.e. players who switch either one or three times. The reason for doing these restricted regressions is that when looking for the effects of risk aversion on outcomes, we want to focus on the cases in which there is some reason to believe that our risk aversion estimates are meaningful. At the end of this section we describe the ways in which results change if we consider the unrestricted sample. Since fixed payments and bonuses are chosen together, the error terms are likely to be correlated. To account for this relationship we conduct our regressions using a seemingly unrelated regression (SUR) framework. The SUR structure in our environment involves defining a system of two contract formation equations, one for the fixed and one for the bonus wage. These are estimated simultaneously allowing for the error terms across equations to be correlated while the coefficients on each variable may be unrelated. The equations include multiple explanatory variables including the principal's characteristics, the agent's characteristics, and a dummy for whether the players are negotiating over a

	Fixed		Bonus	
	Coeff.	P-Value	Coeff.	P-Value
Constant	59.12	<0.001	22.21	0.478
Agent RA	4.19	0.003	3.69	0.190
Agent Times Switched	-0.09	0.979	-1.51	0.833
Principal RA	1.97	0.191	0.31	0.916
Principal Times Switched	-7.08	0.075	-2.50	0.750
Task A Dummy	-20.39	<0.001	39.09	<0.001
Time	-2.70	0.011	0.30	0.889
Number of Obs		86		
R^2	Fixed:	0.3190	Bonus:	0.1771

Table 5: SUR regression for baseline treatment.

contract for task A. In addition, we include a control for time.

The results for the Baseline treatment are presented in table 5. We find that more risk-averse agents end up with higher fixed payments, but the coefficient for bonus is not significantly different from zero. The former is consistent with an efficient distribution of risk. Even within the narrow range of switches included in the data, a principal's propensity to switch is negatively related with the size of the fixed payment offered. We find, that task A contracts are associated with higher bonuses, but lower fixed payments. In other words, some further risk is shifted on to the agent for task A. Finally, fixed payments decline with time leading to the contracts getting even closer to the theoretically optimal contracts over time.

Results for the Excess Supply case are presented in table 6. More risk-averse workers negotiate contracts that involve higher fixed payments and lower bonuses. More risk-averse principals negotiate contracts with lower fixed payments and higher bonuses. This accords exactly with Prediction 4. The principal's propensity to switch matters – those who switch more often negotiate higher fixed payments and lower bonuses. As before, task A contracts involve lower fixed payments and higher bonuses, and fixed payments decline with time.

We turn finally to the Excess Demand case. Now, consistent with prediction, more risk-averse agents negotiate lower bonuses, but we do not measure any effect of agent risk-aversion on fixed payments. Similarly, more risk-averse principals negotiate higher bonuses, with no significant fixed payment effects. The effect of the number of times a principal switches is exactly as in treatment 2. Finally, bonuses tend to increase over time.

Overall, our results on contract structure suggest that risk preferences matter. Indeed, the results for treatment 2 are striking in the support they provide for conventional wisdom. All other things being equal, more risk-averse people negotiate contracts that involve smaller variation in income. That treatment 2 should provide the sharpest support for standard theory seems reasonable. This is the case in which principals have the market power necessary to drive workers down to their reservation utility levels. Competition among workers, as well as competition for projects, is likely to be more intense. Fundamentals appear to matter in such an environment. The other two treatments present interesting

	Fixed		Bonus	
	Coeff.	P-Value	Coeff.	P-Value
Constant	86.42	<0.001	-6.92	0.657
Agent RA	2.32	0.018	-2.84	0.067
Agent Times Switched	-1.53	0.473	-4.88	0.146
Principal RA	-10.27	<0.001	15.86	<0.001
Principal Times Switched	18.90	<0.001	-16.84	0.001
Task A Dummy	-13.55	<0.001	27.74	<0.001
Time	-3.76	<0.001	1.97	0.147
Number of Obs	109			
R^2	Fixed:	0.4085	Bonus:	0.3481

Table 6: SUR regression for Excess Supply treatment.

	Fixed		Bonus	
	Coeff.	P-Value	Coeff.	P-Value
Constant	52.96	0.001	38.20	0.039
Agent RA	0.76	0.955	-2.89	0.075
Agent Times Switched	7.76	0.014	3.08	0.418
Principal RA	-0.21	0.915	4.30	0.068
Principal Times Switched	11.46	0.002	-8.01	0.073
Task A Dummy	-26.14	<0.001	66.34	<0.001
Time	-0.97	0.436	2.59	0.083
Number of Obs	104			
R^2	Fixed:	0.3429	Bonus:	0.5027

Table 7: SUR regression for Excess Demand treatment.

contrasts. Broadly, in treatment 1 it appears that fundamentals affect only the terms of fixed payments, but there is only marginal support for contract theory. In treatment 3 fundamentals matter only for bonuses. The results here are consistent with Prediction 4. It is also noteworthy that principals offer higher bonuses on task A, but make up for it by offering a lower fixed payment.

Robustness. As mentioned earlier, the analysis presented here restricts the number of switches of contracting parties to be one or three. We have also analyzed the unrestricted data. In all three treatments, the task A dummy continues to behave as before, with minor changes in magnitude. For treatment 1, time and principal risk-aversion continue to have the same signs, but are smaller in magnitude and significant only at the 10% level. In treatment 2, the principal's risk aversion measure continues to be as before (the magnitude of the bonus coefficient is reduced to 5.82, but this is significant at the 1% level). However, the agent's risk aversion measure is no longer significantly different from zero. That the principal's risk-aversion should matter more is not really surprising, since the principal makes the offers, and agents are only allowed to respond by communicating their interest. In treatment 3, we actually obtain a positive result in the unrestricted data. The coefficient of the agent's risk-aversion measure becomes positive and significant (Coef. 1.63, P-Value 0.020) in the fixed regression and the coefficient in the bonus regression remains negative and significant. But now the principal's risk-aversion measure is no longer significant. The effect of time is about the same.

A second issue of robustness relates to controlling for non-independence across observations, which we are unable to do in the SUR framework. However, we can run separate OLS regression for bonus and fixed payments. As expected, this increases standard errors (and in fact, the Breusch–Pagan test allows us to reject the hypothesis that the correlation of residuals in the two equations is zero). With OLS, we can control for non-independence of different observations of the same agent (in treatments 1 and 2 this is equivalent to clustering on principal–agent pairs).[11] This will also increase standard errors. So results from OLS are the most conservative, since we give up the efficiency gains of SUR, but get the increase in standard errors from assuming independence only across clusters. Our key results continue to hold, with one important exception. The agent risk-aversion measure is no longer significant (P-values of 0.146 and 0.190 in the fixed and bonus regressions respectively).

The final issue relates to the fact that we have left out of our regression any of the characteristics of players who are not party to the contracts. There is an argument to be made that this should matter for the final terms (although omitting the variables should not introduce a bias since they would be expected, *a priori*, to be uncorrelated with included regressors). Again, including other players does not affect our key relationships. We summarize other key results.

Result 5 - *The fixed wage is decreasing with time in the baseline and excess supply treatments. The bonus is increasing with time in the excess demand treatment.*

Result 6 - *In the baseline treatment, the fixed wage is lower the less sophisticated is the*

[11] To take into account the interdependence of individual responses within the households, we employ regression with robust cluster variance estimator. It specifies that the observations are independent across clusters, but not necessarily within clusters, and corrects for the interrelated standard errors due to the observations from multiple individuals within a cluster and thus affects the estimated standard errors and variance-covariance matrix of the estimators (VCE), but not the estimated coefficients.

Treatment	Task A		Task B		Neither	
	All	**LS Only**	**All**	**LS Only**	**All**	**LS Only**
Baseline	6.74	6.37	7.35	6.47	7.12	6.13
Excess Supply	6.28	6.04	6.32	6.36	6.47	6.59
Excess Demand (P1)	7.34	6.36	7.17	6.44	7.32	6.45
Excess Demand (P2)	6.36	5.23	6.95	6.33		-

Table 8: Average RA Measure of agents in different tasks.

principal. In the excess supply case, the fixed wage is higher while the bonus is lower for less sophisticated principals. In the excess demand treatment, the fixed wage is larger when the principal is less sophisticated, while the bonus is smaller. Fixed payments are higher when the agent is less sophisticated.

4.3 Agent's Choices of Tasks

One of the central questions that motivated this study was the degree to which it would be possible to observe any selection of agents to contracts on the basis of their risk preferences. To get a handle on this issue we need to take a general look at what tasks agents were choosing to be interested in and then analyze those choices to determine what motivated them and in particular if the risk preferences identified using the Holt-Laury technique can help to explain any of the observed sorting. We begin with some summary statistics. In Table 8 we see, for each treatment, the average risk-aversion measure of agents who expressed interest in the various tasks. This includes the statistics for all data and then the same numbers, but only for the Low Switcher (LS) agents who switched one or three times. For the excess demand treatment we report results separated by the principal offering the task. With rare exception (Principal 1 in All Data case) the average risk aversion of agents engaged in Task A is smaller than that of agents engaged in Task B. In the Excess Supply and Excess Demand treatments, agents who pick neither task tend to be more risk-averse than agents who pick a task. A striking feature of the data is that in the Excess Demand treatment, there is clear evidence of agents self-selecting by risk-preferences into choosing contracts offered by principals. We find that the more risk tolerant agents systematically prefer the contracts offered by the second principal. Since the actual tasks of each principal are identical and principals are assigned to the two possible roles at random, it is surprising to find a structural difference emerging due to the theoretically meaningless distinction of labeling one as principal 1 and the other principal 2. Differences are accentuated when we focus only on people who switch a small number of times. So there is strong preliminary evidence that there is matching of agents to tasks. We investigate this further.

Our next result summarizes the task choice behavior of the agents.

Result 7 - *Across all treatments more agents are interested in B tasks than A tasks at the end of the bargaining phase and in all treatments there is a significant number of subjects showing interest in none of the proposed contracts at the conclusion of the bargaining period.*

Table 9 shows the average number of workers exhibiting interest in Task A, Task B or neither at the end of the bargaining phase during the experiments. What we generally

		Baseline	Excess Supply	Excess Demand
All Rounds	**Task A**	0.62	0.88	0.83
	Task B	1.11	1.68	1.57
	Neither	0.28	0.44	0.60
Last 4	**Task A**	0.75	0.88	0.82
	Task B	1.07	1.65	1.60
	Neither	0.18	0.47	0.58

Table 9: Number of subjects interested in tasks.

find is that more workers are interested in the B contracts than the A contracts. The low level of interest displayed by workers in Task A contracts during the Baseline sessions, may explain the similarity in contracts offered between that treatment and the Excess Supply treatment. In the Excess Supply treatment there were 3 agents but only 2 tasks which meant one worker necessarily ended up without a task at the end of the bargaining phase. This excess supply was expected to allow principals to bargain agents down to their outside option. In the Baseline treatment, there were 2 agents and 2 tasks which we expected could allow agents to end up with more generous contracts. However, since subjects were primarily interested in Task B contracts in the Baseline treatment, that essentially lead to excess supply for that task even in the Baseline treatment and principals exhibited the ability to exploit the resulting competition.

We also find that there are a number of workers displaying interest in none of the offered contracts at the end of the bargaining phase. This is particularly intriguing in the Excess Demand treatment because there were already 4 tasks and only 3 workers meaning that one task would necessarily go unfulfilled and all workers could end up assigned to tasks. Over the last half of the sessions, it was the case that about half of the time one subject was still holding out for a better offer from the principals and the principals were both unwilling to make a better offer to entice the worker. Since withholding interest in tasks is how workers can attempt to induce principals to offer more generous contracts, this is evidence that workers are bargaining to increase their wages by holding out from indicating interest interest. Since both worker and manager would benefit from agreeing to contract terms, this is also a display of a hold-up problem.

A more detailed examination of the contract choice behavior of the agents is shown in table 10. This provides support for our next three results on the endogenous matching hypothesis (Prediction 3).

Result 8 - *The terms of the contract have the expected effect on the likelihood of agents selecting them. More favorable terms for A increase chance of choosing it relative to B and more favorable terms for B decrease the likelihood of choosing A.*

Result 9 - *Risk preferences have no effect on task choice in Baseline and Excess Demand but we find that in the Excess Supply treatment an agent's likelihood of showing interest in task A decreases as their risk aversion increases.*

Result 10 - *In the excess demand treatment, less risk averse agents contract with Principal 2. Among these agents, the less risk averse pick task A.*

Table 10 contains multinomial logit regressions for the three treatments with the dependent variable being the contract choice (Neither, Task A or Task B) using Task B choices as the comparison group. Thus the coefficients on all variables represent the relative shift in behavior toward that category from a choice for Task B. The table contains the results for "low switchers" (i.e. workers who switch either one or three times) again to focus on those subjects for whom our measure of risk preferences may be accurate. The independent variables include the measure of the risk preferences of the subject (a higher value of this measure indicates more risk averse), the number of times a subject switched between lotteries in the risk preference assessment (used as a proxy for sophistication), details of the contracts, the period/time in the experiment, and lagged earnings. We control for the fact that different observations on the same worker are not independent. Both the number of observations as well as the number of clusters are indicated in table 10.

The findings show that the terms of the contract matter as expected. That is, more generous task A terms increase the likelihood of someone choosing it instead of task B, while more generous task B terms decrease the likelihood of the agent choosing A instead of B. In all the treatments, the fixed wage on task B is the main determinant of a subject choosing to be interested in B rather than withholding interest. In each case, a higher fixed payment on task B makes it less likely that a worker will not show an interest in either project.

A key result is that for the excess supply treatment, we observe a negative and large coefficient on the measure of risk preferences that is significant while the coefficient on this variable is insignificant in the other two treatments. This would indicate that subjects were self-selecting into contracts based on their risk preferences in the excess supply treatment. The negative value of the coefficient means that the more risk averse is the agent, the less likely they are to show interest in the more risky task A rather than the safer task B. That is precisely the relationship expected.

We do not get corresponding results in the other two treatments. In the case of the excess demand treatment, an examination of table 8 suggests that a more complex phenomenon is at play. Clearly, there is evidence of sorting by Principal. A number of alternative specifications allow us to capture this effect. First if we run a multinomial logit of choice of principal, with identical controls, we find that more risk-averse agents are less likely to contract with principal 2. Also, conditioning on principal 2, we find that more risk-averse agents are less likely to pick task A. Alternatively, running a logit model for the choice of task A versus all other choices, and with identical controls, we get a strong and significant selection effect of less risk-averse agents into task A of principal 2. Upon examining the contracts offered by principal 2, we find that principal 2 offers, on average, higher bonuses and lower fixed payments. This would indeed attract the less risk-averse workers. When we focus on the accepted contracts only, principal 2 contracts promise a much higher bonus on task A (to agents who switch one or three times). It appears that the relative position of principal 2 on the screen is significant. A plausible explanation is that Principal 2 attracts relatively risk-tolerant agents because intensive search runs the risk of time running out. The more risk-averse workers accept the contracts scanned first, whereas the more risk-tolerant workers search for more attractive options. At the same time, since it is harder to attract workers for them, people in the role of principal 2 offer more systematically different contracts. What is abundantly clear is that the extent of sorting is too strong to be random. While it is not certain what accounts for the sorting, the most plausible explanation is the

		Baseline		Excess Supply		Excess Demand	
		Coeff.	P-Value	Coeff.	P-Value	Coeff.	P-Value
Neither	**Constant**	6.05	0.023	2.45	0.087	14.27	0.006
	Risk Preference	0.50	0.082	-0.06	0.624	-0.01	0.970
	Times Switched	-1.24	.015	-0.07	0.788	0.10	0.811
	Fixed Task A	-0.01	0.539	0.01	0.371	0.04	0.032
	Bonus Task A	-0.01	0.512	-0.01	0.098	-0.06	<0.001
	Fixed Task B	-0.08	0.009	-0.07	0.001	-0.19	<0.001
	Bonus Task B	-0.01	0.707	<0.01	0.866	-0.05	0.090
	Fixed Task A (P2)					0.09	<0.001
	Bonus Task A (P2)					-0.07	0.012
	Fixed Task B (P2)					-0.10	0.001
	Bonus Task B (P2)					-0.09	<0.001
	Time	-0.17	0.583	0.14	0.926	2.06	<0.001
	Lagged Earning	>-0.01	0.354	<0.01	0.883	>-0.01	0.231
Task A	**Constant**	-4.03	0.082	-2.22	0.23	1.94	0.100
	Risk Preference	>-0.01	0.995	-0.20	0.046	-0.12	0.278
	Times Switched	0.11	0.591	0.20	0.312	-0.50	0.094
	Fixed Task A	0.09	<0.001	0.04	<0.001	0.05	0.002
	Bonus Task A	0.04	0.002	0.02	0.022	0.02	0.073
	Fixed Task B	-0.08	<0.001	-0.03	0.001	-0.04	<0.001
	Bonus Task B	-0.02	0.037	-0.01	0.389	-0.04	0.001
	Fixed Task A (P2)					0.01	0.269
	Bonus Task A (P2)					0.01	0.037
	Fixed Task B (P2)					-0.02	0.092
	Bonus Task B (P2)					<0.01	0.437
	Time	0.22	0.088	0.01	0.949	0.22	0.026
	Lagged Earning	<0.01	0.499	<0.01	0.393	>-0.01	0.128
	Number of Obs	184		288		240	
	Number of Clusters	23		36		30	
	Pseudo R^2	0.39		0.18		0.32	

Table 10: Multinomial logit regression of agent choice between Task A, Task B or neither. Task B is used as comparison group.

notion that agents' risk aversion correlates with their willingness to search the screen and that the second principal's observe that they have to offer contracts that appeal to the risk tolerant in order to attract workers. This demonstrates substantial sophistication on the part of the principals that would have been unlikely to have been observed had we run the experiments with simpler ultimatum offers.

One possible reason for the apparent absence of endogenous matching in the baseline treatment is that, with only two agents, the range of risk-aversion coefficients within a group becomes quite small on average. This makes it harder to detect significance in differences. Further, the contract for the risk tolerant agent needs to be just different enough from the other contract to only attract the risk tolerant agent which means that the terms are based on the risk preferences of the more risk averse individual. Thus finding much of an effect based on the risk preferences of the more risk tolerant seems unlikely.

Finally, we should note that the effects of time are significant in the baseline and excess demand treatments. As time increases, task A is more likely to be chosen relative to task B.

4.4 Choice of Effort

To complete our analysis we must look at how agents made effort decisions. Standard decision theory tells us that risk neutral decision makers will only engage in high effort, that is choose option X, if the expected return from high effort is greater than the return from low effort or $.8 \times b - 35 \geq .25 \times b - 5$. For risk averse agents, the condition is similar. Other previously cited literature on experiments with principal-agent games suggests that agents might be willing to engage in high effort as a reward to a principal for offering a high fixed wage, w. Our results on the effort choices of agents are summarized by our next two results.

Result 11 - *In all treatments and tasks except the B tasks in the Excess Demand treatment, agents predominately choose their effort level as a best response to the incentives of the contract.*

Result 12 - *The primary variable for determining effort level is the bonus wage. There is no observed relationship between the fixed wage and effort choice. In the Excess Demand treatment there is also an inertia effect that leads to some subjects choosing high effort repeatedly even in cases for which it is not a best response.*

Table 11 contains a basic summary of the key statistics relating to effort choice. We do find that subjects engage in high effort much more often on task A than task B in all treatments though in the Excess Demand treatment subjects choose high effort even on B tasks quite often. Looking at this with respect to the contract structure from table 3 is likely to explain much of this because task B contracts in the Excess Demand treatment had much higher bonus wages than in the other treatments. Table 11 confirms this showing that high effort is the optimal choice in 88 and 74% of the accepted contracts for A and B tasks respectively in the Excess Demand treatment (over the last 4 periods) while the numbers for the other two treatments are 62/40 and 55/13. We further see that except for agents accepting task B contracts under the Excess Demand treatment, agents best respond to their incentives around 80% of the time.

		Baseline		Excess Supply		Excess Demand	
		Task A	**Task B**	**Task A**	**Task B**	**Task A**	**Task B**
All Rounds	**High Eff Choices**	44	25	48	12	65	44
	High Eff is BR	54	39	53	15	78	72
	Agent Choose BR	80	74	79	88	75	39
Last 4	**High Eff Choices**	49	21	52	12	71	43
	High Eff is BR	62	40	55	13	88	74
	Agent Choose BR	81	81	83	73	78	38

Table 11: Summary statistics on effort choices as well as when high effort was a best response and how often subjects chose whatever option was the best reponse.

This is a strong indication that subjects are simply best responding to contract structures but the behavior on task B in the Excess Demand suggests the possibility of gift exchange or some other reciprocal behavior on the part of the agents. Figure 2 gives some indication in regard to whether this is the case. It shows a scatter plot of the fixed wage versus the bonus for all accepted contracts by treatment and indicates whether the effort choice by the agent was high or low. It includes two key lines along the Bonus axis with one being the lower bound on the bonus for which high effort has a positive expected value while the second line indicates the lower bound on the bonus for high effort to have a higher expected value than low effort. In the Baseline and Excess Supply treatments, there are very few choices of high effort below both of these lines. There are in fact more cases in which high effort is a best response and unchosen than when it is not a best response but chosen anyway. There is no indication that high fixed wages induce high effort.

Something is quite different for the Excess Demand case. Here there are quite a number of cases in which agents engage in high effort even though it is not in their self-interest to do so. There is, however, still no indication that the agents were more likely to engage in high effort based on seeing a high fixed wage which would be necessary to support an explanation involving some sort of reciprocity or gift exchange explanation.

While figure 2 gives a nice graphical representation of how effort choice varies with contract structure, there is a need for a more precise statistical examination of the determinants of effort choices. To provide this we have conducted a series of random effects logit panel regressions. The first one, contained in table 12 is intended to look at the overall effects and any main differences between treatments. The regressors include the fixed wage and bonus faced by the agent, a dummy variable equal to 1 if the task were of type A, the previous effort choice of the agent, treatment dummies and then our measures of risk aversion and sophistication. The sophistication measure, High Switcher is a dummy variable equal to 1 if the subject switched anything other than 1 or 3 times during the Holt-Laury procedure. The overall conclusion is that the fixed wage has no significant effect on effort choices while the bonus does. There is also an indication that some subjects may have some inertia in their effort decisions as picked up by the lagged effort variable and then agents in the Excess Demand treatment choose high effort more often than in the other two treatments.

Table 13 shows the same random effects logit regression performed separately for each treatment. In general we observe the same effects across treatments, except for the Excess

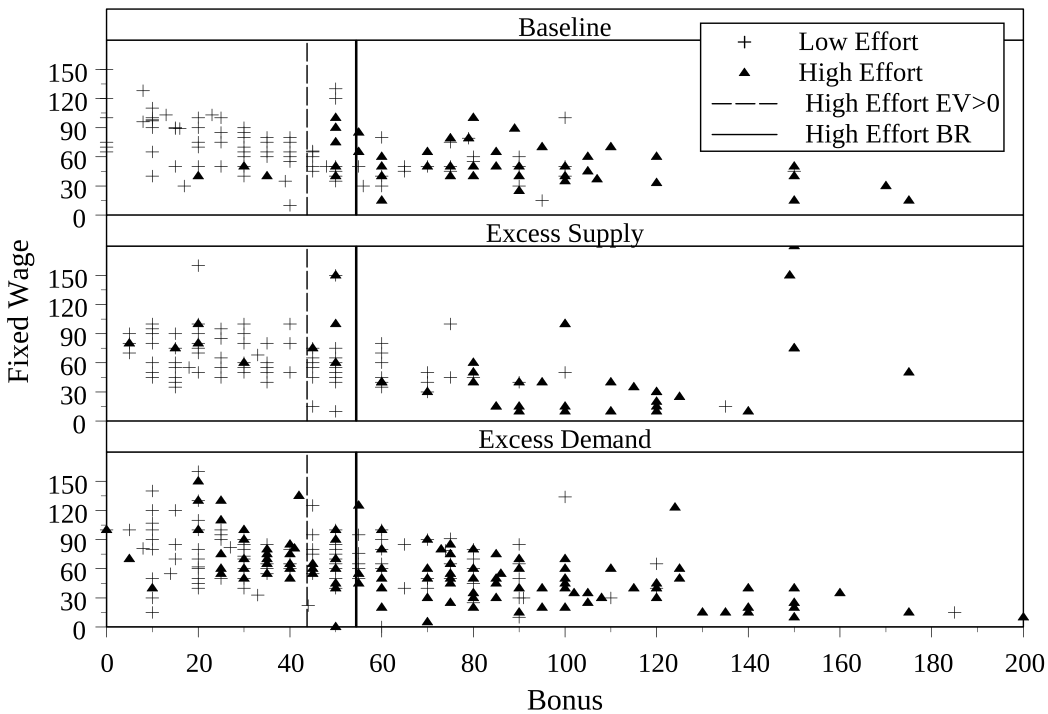

Figure 2: Scatterplot of accepted contracts by treatment with an indication of the effort choice by the agent.

	Coeff.	P-Value
Constant	-3.447	<0.001
Fixed	-0.003	0.477
Bonus	0.035	<0.001
Task A Dummy	0.144	0.607
Lagged Effort	1.054	<0.001
Excess Supply Dummy	-0.112	0.758
Excess Demand Dummy	0.847	0.007
Risk Preference	0.051	0.302
High Switcher	-0.289	0.336
Number of Obs (Groups)	486 (120)	
Log Likelihood	-223.57	

Table 12: Panel logit regression of choice of effort (high=1, low=0).

	Baseline		Excess Supply		Excess Demand	
	Coeff.	**P-Value**	**Coeff.**	**P-Value**	**Coeff.**	**P-Value**
Constant	-3.932	0.002	-5.72	<0.001	-1.745	0.051
Fixed	0.001	0.969	0.003	0.786	-0.009	0.234
Bonus	0.047	<0.001	0.056	<0.001	0.029	<0.001
Task A Dummy	0.377	0.493	1.862	0.004	-0.605	0.201
Lagged Effort	0.405	0.449	0.389	0.547	1.297	0.002
Risk Preference	-0.015	0.921	0.010	0.935	0.054	0.433
High Switcher	0.171	0.850	0.620	0.434	-0.240	0.577
Number of Obs (Groups)	127 (30)		136 (43)		223 (47)	
Log Likelihood	-52.35		-40.72		-118.76	

Table 13: Panel logit regression of choice of effort (high=1, low=0) by treatment.

Demand treatment. For the first two treatments, the coefficient on the bonus is about the only significant effect (though the task A dummy is significant for Excess Supply). For the Excess Demand treatment, we again see that the bonus is highly significant but we also find that the lagged effort choice is positive and highly significant. As indicated by figure 2, the fixed wage is not significant for any treatment.

The overall result is that effort choices are well explained by classical decision theory in that the primary determinant of whether the subject chooses high effort or not is whether it has the higher expected value. We see no evidence of agents choosing high effort in response to receiving contracts with high fixed wage components from principals.

5 Conclusion

Although there is an extensive theoretical literature on the subject, empirical validation of contract theory provides support that is mixed at best. In response a number of alternative foundational assumptions, such as inequity aversion, fairness, reciprocity, norms, and so forth are being explored in the current literature as a means of providing alternative explanations for observed data. Against this background, we designed our study with two related aims in mind. The first was to conduct very detailed experimental tests of several predictions that arise out of standard contract theory. The second was to explore some of the reasons behind the empirical failures of the theory that could inform on the design of future field studies.

One of the keys to our study is our ability to obtain a separate measure of our subjects' risk preferences which we can use in understanding their contracting behavior. Field studies have been hampered by the unavailability of quality measures of risk preferences which has meant that tests of the underlying theory have been potentially confounded by issues with this key variable. By itself, this may have been a surmountable empirical problem, but the difficulty gets compounded if the environment allows for the endogenous matching of subjects to tasks by risk preference. In an experimental setting these problems can be avoided since we can measure the degree of risk aversion of subjects, who are then put in the position of negotiating over contracts, either in the role of a principal or of an agent. Other potential confounds that our design allows us to investigate involve understanding the effects when either side of the market possesses market power as well as being able to identify when contracting behavior may be affected by fairness or reciprocity norms. Our general findings largely confirm the validity of a standard model of contracting behavior with an allowance for the effects of market power and self-selection into tasks according to risk preferences.

It should be reiterated that one key difference between our study and prior experimental examinations of principal-agent theory is that in our design, principals bargain with agents over wages through a continuous bargaining protocol. This is substantially different than most prior studies which use a standard ultimatum game, with the principal making take-it-or-leave-it offers. This design difference could be at least part of the reason why our results show little support for the fairness and reciprocity norms found in other studies. We used this more extensive bargaining framework due to a belief that the overall complexity of the environment necessitated an iterative procedure to allow subjects on both sides of the market to develop a clearer understanding of the problem.

Although we will not re-enumerate all of our findings, it is worth emphasizing some of the key results. We found first that market power matters, as subjects take advantage of competition on the other side of the transaction. We also found that effort, by and large, responds to incentives just as standard theory predicts, although in one treatment we did find a result of some degree of inertia in effort choices which remains unexplained. In the excess supply treatment we find striking support for each of our key predictions – the market power hypothesis, the efficient risk-sharing hypothesis, the endogenous matching hypothesis and the optimal effort choice hypothesis. In the excess demand treatment market position is clearly exploited by agents, who end up with very generous contracts on average. The predictions of the efficient risk-sharing hypothesis are confirmed to a considerable degree, as bonuses accord with predictions. The results on endogenous matching was, to us, surprising, as we found matching of risk-tolerant agents to subjects in the role of principal 2. These principals offered contracts more likely to attract risk-tolerant agents. Among the agents who were attracted to principal 2 the less risk averse tended to choose the more risky task. In the baseline case results were somewhat ambiguous—although we find support for the market power hypothesis and the efficient effort choice hypothesis, there were few detectable effects relating to risk preferences.

Overall, the results provide support for the classical theory of contracts. What we find particularly compelling is how various pieces of evidence fit together. Consider the excess supply case, for instance. The principal offers higher bonuses, but lower fixed payments, on task A. This is just what is required to attract only the less risk averse agents, who do in fact select task A. The negotiated contracts also vary with risk preferences in ways consistent with the predictions of contract theory. At the same time, principals are also able to exploit competition among agents and bargain them down to very near their reservation utility. The bonuses are about what is needed for optimal effort levels, which tend to be provided by agents. It is indeed surprising that the choices along so many different dimensions all happen to be consistent with the classical theory based on individual rationality and incentives.

Endogenous matching by risk preferences is a robust phenomenon in our experiments. It occurs in expected ways in the excess supply treatment, and in unexpected ways in the excess demand treatment. Interesting as these results are by themselves, there are important implications for empirical work. Researchers have previously expressed concern about potential for bias in parameters of interest when unobserved heterogeneity in risk preferences is combined with endogenous matching. Our study confirms that the concerns are not misplaced. Endogenous matching occurs, and needs to be adequately controlled for in econometric studies of contracts. The excess demand result suggests, moreover, that the nature of sorting by risk preferences is likely to depend on institutional details such as which side of the market possesses market power. Without correcting a field study for the confounds of self-selection and market power it seems quite likely that the results could end up showing lack of support for the underlying theory even though the decisions made by all parties in that market could be conforming quite well to a properly specified version of the theory.

To return to the question of alternative preference specifications, on balance, the evidence from our experiments seems to favor classical preferences. Compensation reflects market opportunities, although it is true that people at a disadvantage are not always driven down to precisely their reservation levels. In a somewhat different environment than ours Brandts and Charness (2003) find that market conditions do not matter and treat this is as evidence against the classical model. The experiment in the Brandts and Charness (2003) centers on the "gift exchange game" in which a principal is essentially restricted to only being allowed to offer a fixed wage. In such an environment offering the highest fixed wage possible is the only way a principal might induce an agent to engage in high effort and this is true regardless of the state of the principal's market power.[12] Consequently it is perhaps not surprising that the actual wages offered did not depend on market power in that study while wages do depend on market power in ours. That paper and others do, however, suggest that in the presence of alternative institutions and alternative preference structures that fundamentals of the classical model, such as risk preference and market power, should matter less. Our finding that they often matter a lot, and in ways consistent with classical theory, might be interpreted as providing evidence against models based on alternative foundational assumptions. The better interpretation of our study relative to the others is that they all provide evidence that modes of behavior are contingent upon institutional details. Further work on outlining the conditions in which institutions will trigger classical preferences as in our results or preferences involving some degree of "other-regardingness" as suggested by others is certainly necessary to allow a finer exploration of these issues.

References

Ackerberg, D. and M. Botticini (2002). Endogenous Matching and the Empirical Determinants of Contract Form. *Journal of Political Economy 110*, x–x.

[12]The intuition for this argument can be found in the trust game results from Berg, Dickhaut, and McCabe (1995). In the trust game first movers were allowed to send some portion of $10 to a second mover with the amount of money sent being tripled. The second mover could then send some amount back. The typical result is that the second mover is much more likely to reciprocate by sending money back if they received full trust, i.e. $10, from the first mover rather than partial trust, i.e. $8-9. A worker in the gift-exchange game should be similarly much more likely to engage in high effort as a result of a "full" gift rather than only a partial one.

Allen, D. and D. Lueck (1992). Contract Choice in Modern Agriculture: Cash Rent versus Crop Share. *Journal of Law and Economics 35*.

Andersen, S., G. Harrison, M. Lau, and E. Rutstrom (2005). Lost in State Space: Are Preferences Stable? Working Paper UCF.

Berg, J., J. Dickhaut, and K. McCabe (1995). Trust, Reciprocity, and Social History. *Games and Economic Behavior 10*, 122–142.

Bolton, G. and A. Ockenfels (2000). ERC: A Theory of Equity, Reciprocity and Competition. *American Economic Review 90*, 166–193.

Bolton, P. and M. Dewatripont (2005). *Contract Theory*. The MIT Press.

Brandts, J. and G. Charness (2003). Do Labour Market Conditions Affect Gift Exchange? Some Experimental Evidence. Working Paper.

Chiappori, P. and B. Salanie (2000). Testing Contract Theory: A Survey of Some Recent Work. World Congress of the Econometric Society, Seattle.

Fehr, E. and A. Falk (2002). Psychological Foundations of Incentives. *European Economic Review 46*(4-5), 687–724.

Fehr, E. and S. Gachter (2002). Do Incentive Contracts Undermine Voluntary Cooperation? Working Paper No. 34, IERE University of Zurich.

Fehr, E., S. Gachter, and G. Kirchsteiger (1997). Reciprocity as a Contract Enforcement Device.

Fischbacher, U. (1999). z-Tree: Zurich Toolbox For Readymade Economic Experiments. Zurich University.

Guth, W., W. Klose, M. Konigstein, and J. Schwalbach (1998). An Experimental Study of a Dynamic Principal-Agent Relationship. *Managerial and Decision Economics 19*(4-5), 327–341.

Guth, W., M. Konigstein, J. Kovacs, and E. Zala (2001). Fairness Within Firms: The Case of One Principal and Multiple Agents. *Schmalenbach's Business Review*.

Holmstrom, B. (1979). Moral Hazard and Observability. *Bell Journal of Economics 10*, 74–91.

Holmstrom, B. and P. Milgrom (1987). Aggregation and Linearity in the Provision of Intertemporal Incentives. *Econometrica 55*, 303–28.

Holt, C. and S. Laury (2002). Risk Aversion and Incentive Effects. *American Economic Review 92*(5), 1644–1655.

Isaac, R. M. and D. James (2000). Just Who Are You Calling Risk Averse. *Journal of Risk and Uncertainty 20*(2), 177–187.

Ivanova-Stenzel, R. and T. C. Salmon (2004). Bidder Preferences Among Auction Institutions. *Economic Inquiry 42*, 223–236.

Ivanova-Stenzel, R. and T. C. Salmon (2005). Revenue Equivalence Revisited. Working Paper Florida State University.

Konigstein, Manfred, J. K. and E. Zala-Mezo (2003). Fairness in a one-principal-two-agents game a post-experimental questionnaire analysis. *Journal of Economic Psychology 24*, 491–503.

Mirrlees, J. A. (1974). Notes on Welfare Economics, Information and Uncertainty. In M. Balch, D. McFadden, and S. Wu (Eds.), *Essays in Equilibrium Behavior under Uncertainty*. Amsterdam: North-Holland.

Mirrlees, J. A. (1975). The Theory of Moral hazard and Unobservable Behaviour, Part 1. Mimeo, Oxford (Published in the Review of Economic Studies in 1999).

Palfrey, T. R. and S. Pevnitskaya (2004). Endogenous Entry and Self-selection in Private Value Auctions: An Experimental Study. Working Paper.

Prendergast, C. (2002). The Tenuous Trade-off between Risk and Incentives. *Journal of Political Economy 110*, 1071–1102.

Roth, A. E. (1995). Bargaining Experiments. In J. H. Kagel and A. E. Roth (Eds.), *The Handbook of Experimental Economics*, pp. 253–348. Princeton, New Jersey: Princeton University Press.

Rothschild, M. D. and J. E. Stiglitz (1976). Equilibrium in Competitive Insurance Markets: The Economics of Markets with Imperfect Information. *Quarterly Journal of Economics 90*(4), 629–649.

Young, H. P. and M. A. Burke (2001). Competition and custom in economic contracts: A case study of Illinois agriculture. *American Economic Review 91*, 559–573.