

1. Hidden Information, Bargaining Power, And Efficiency: An Experiment

Antonio Cabrales, Gary Charness and Marie Claire Villeval

Forthcoming in *Experimental Economics*, 2011

Abstract: We devise an experiment to explore the effect of different degrees of bargaining power on the design and the selection of contracts in a hidden-information context. In our benchmark case, each principal is matched with one agent of unknown type. In our second treatment, a principal can select one of three agents, while in a third treatment an agent may choose between the contract menus offered by two principals. We first show theoretically how different ratios of principals and agents affect outcomes and efficiency. Informational asymmetries generate inefficiency. In an environment where principals compete against each other to hire agents, these inefficiencies may disappear, but they are insensitive to the number of principals. In contrast, when agents compete to be hired, efficiency improves dramatically, and it increases in the relative number of agents because competition reduces the agents' informational monopoly power. However, this environment also generates a high inequality level and is characterized by multiple equilibria. In general, there is a fairly high degree of correspondence between the theoretical predictions and the contract menus actually chosen in each treatment. There is, however, a tendency to choose more 'generous' (and more efficient) contract menus over time. We find that competition leads to a substantially higher probability of trade, and that, overall, competition between agents generates the most efficient outcomes.

Keywords: Experiment, Hidden Information, Bargaining Power, Competition, Efficiency

JEL Classification: A13, B49, C91, C92, D21, J41

Contact: Antonio Cabrales, Universidad Carlos III de Madrid, antonio.cabrales@uc3m.es
Gary Charness, University of California at Santa Barbara, charness@econ.ucsb.edu
Marie Claire Villeval, University of Lyon, CNRS, and IZA, Bonn, villeval@gate.cnrs.fr

Financial support from the *Agence Nationale de la Recherche* (ANR BLAN07-3_185547 "EMIR" project) is gratefully acknowledged.

1. INTRODUCTION

The theory of markets with asymmetric information has been a “vital and lively field of economic research” (2001 Nobel Prize committee) for decades. The classic ‘lemons’ paper (Akerlof 1970) illustrated the point that asymmetric information led to economic inefficiency, and could even destroy an efficient market. Since the seminal works of Vickrey (1961) and Mirrlees (1971), research on mechanism design has sought ways to minimize or eliminate this problem.¹ In an environment with *hidden information* (sometimes characterized as adverse selection), each agent knows more about her² ‘type’ than the principal does at the time of contracting. In the standard labor scenario, a firm hires a worker but knows less than the worker does about her innate work disutility. Other typical applications include a monopolist who is trying to price discriminate between buyers with different (privately known) willingness to pay, or a regulator who wants to obtain the highest efficient output from a utility company with private information about its cost.³

The fact that agents know their own ability levels while principals may not causes difficulties in contracting, as an agent may not choose the action that is in the best interest of the principal. If outcomes are related to actions, firms with complete information could design ‘first-best’ contracts that theoretically induce truthful revelation of types and generate economic efficiency by making the contract contingent on the outcome. However, in contracting under hidden information, the problem is how to induce the efficient action without observing the

¹ Applications include public and regulatory economics (Laffont and Tirole 1993), labor economics (Lazear 1999), financial economics (Freixas and Rochet 1997), business management (Milgrom and Roberts 1992), and development economics (Ray 1998).

² Throughout this paper we assume that the principals are male and the agents are female.

³ One-shot contracts are common in consumer transactions. In the public sector, government procurement is often conducted on a one-shot basis.

agent's true type; in this case, it is typically necessary to devise 'second-best' contracts that lead to separation of types, but which are somewhat distorted and less than fully efficient.

In this paper we report the results of experiments designed to test the influence of competition when there is hidden information. This can be seen as a question of organizational or institutional design – what effects do different rules and markets have on performance and efficiency?⁴ We examine how differing degrees of relative bargaining power between principals and agents affect outcomes and efficiency when there is a problem of hidden information.

Our approach is to consider three environments that differ according to the type of competition present in the environment. In our benchmark case, each principal is matched with one agent of unknown type. In our second treatment, a principal can select one of three agents, while in a third treatment an agent may choose between the contract menus offered by two principals. Principals can choose to offer one of six feasible contract menus, which are held constant across our treatments; in turn, agents can select high or low effort, or reject the contract menu entirely and receive reservation payoffs. We derive the equilibrium predictions for each environment and examine the outcomes in each treatment, ranking the institutions as a function of their relative efficiency, both in terms of effort and the probability of trade.

In this respect theory provides a first answer. To understand the theoretical efficiency ranking, it is important to realize that incomplete information in markets creates inefficiencies because the agents have a certain monopoly power. More precisely, they are the sole 'owners' of a valuable resource – information about their type. We first show from a theoretical point of view how different degrees of bargaining power between principals and agents, related to various

⁴ Although we use the standard static screening model, it is worth noting that Kanemoto and MacLeod (1992) examine the effect of competition in a dynamic environment and find that one obtains the first-best outcome if there is sufficient competition for workers, even with asymmetric information. Perhaps there is some empirical analog to this result in the static case.

degrees of competition in the market, affect outcomes and efficiency. In an environment where principals compete against each other to hire agents, these inefficiencies may disappear, but they are insensitive to the number of principals. In an environment where agents compete to be hired, efficiency improves dramatically and increases in the relative number of agents because competition reduces the agents' informational monopoly power. However, this environment also generates a high inequality level and is characterized by multiple equilibria, which may have important behavioral implications in the field if people have social preferences such as inequality aversion.

Our experiment constitutes the first test of the impact of varying the relative bargaining power between principals and agents on the selection of contracts in the presence of both heterogeneous agents and hidden information. Our results are mostly supportive of the theory and the major implication is that the bargaining power directly affects the choice of contract menus. In comparison with environments in which there is no competition, our experiment finds that the institutional environment in which agents compete against each other improves efficiency as far as we consider the contracting pairs.

Even though, in general, there is a fairly high degree of correspondence between the theoretical predictions and the contract menus actually chosen in each treatment, there is a tendency to choose more 'generous' (and more efficient) contract menus over time. We find that competition leads to a substantially higher probability of trade, and that, overall, competition between agents generates the most efficient outcomes. We observe a fairly high degree of separation of agents' types in the choices made in response to the various contract menus; interestingly, with agent competition we observe the more able agents strategically foregoing the option that would pay them more (if they are chosen), in order to signal their type by choosing

the option that less able agents should never choose. Our data also show considerable evidence of changes in behavior over time, as participants learn what is effective and what is not.

The remainder of the paper is structured as follows: We review the relevant literature in section 2, and we describe our theoretical model and derive its predictions in section 3. We present our experimental design and implementation in section 4, with the results given in section 5. We conclude in section 6.

2. RELATED LITERATURE

Perhaps due to the complexity of business relationships, it is difficult to find support from field data for principal-agent theory. While there has been considerable theoretical work on contracts in recent decades, empirical tests of the theory have long remained scarce, particularly as far as hidden information is concerned. The Prendergast (1999) and Chiappori and Salanié (2003) surveys show that the econometrics of contracts has recently become a burgeoning field of research. However, the latter study points out that a number of empirical tests suffer from selection and endogeneity biases. In addition, due to the lack of data on contracts many theoretical models have not yet been tested. These difficulties explain why few empirical tests of the hidden-information problem are available in the literature (see notably Cawley and Philipson, 1999; Chiappori and Salanié, 2000; Dahlby, 1983; Dione and Doherty, 1994; Finkelstein and Poterba, 2000; Genesove, 1993; Puelz and Snow, 1994; Young and Burke, 2001). Given the difficulties inherent with field data in this area, laboratory experiments offer a complementary approach that offers some promise, since it is possible to isolate and vary the factors of interest while keeping all others constant. In this paper we use experimental methodology to empirically test a model of the ratchet effect for which no survey data are available.

Previous experimental studies on asymmetric information have typically examined contracting with hidden action (moral hazard), where effort is not contractible (see notably Anderhub, Gächter, and Königstein, 2002; Berg, Daley, Dickhaut, and O'Brien, 1992; Charness and Dufwenberg, 2006; DeJong, Forsythe, Lundholm, and Uecker, 1985; Keser and Willinger, 2000; and Königstein, 2001). However, there is little experimental work on contracting with hidden information about an agent's type. A few studies confirm the existence of a market for "lemons" (see Lynch, Miller, Plott and Porter, 1986, in oral double auctions, and Holt and Sherman, 1990, in posted-offer auctions), while Brandts and Holt (1992) and Banks, Camerer and Porter (1994) provide mixed evidence in signaling games. Miller and Plott (1985) show that market processes allow the buyers to extract private information from the sellers, in an environment in which the proportion of buyers varies but is never greater than the number of sellers. Beyond market processes, it has also been shown that agents are more likely to select actions according to their types when more equitable menus are proposed (Cabrales and Charness, forthcoming, in the context of team production) and when communication is allowed (Charness and Dufwenberg, forthcoming). Finally, a few articles (Chaudhuri, 1998; Cooper, Kagel, Lo, and Gu, 1999; and Charness, Kuhn, and Villeval, forthcoming) study the dynamic contracting problem. However, we are unaware of any previous study that considers the effect of varying, in a static framework, the relative bargaining power in contracts with either excess principals or excess agents.

One might predict that different relative bargaining power for principals and agents should lead to different contract menus being selected. However, the evidence from the handful of experimental papers on the effects of unbalanced competition on the outcomes between principals and agents is somewhat mixed. Fischbacher, Fong and Fehr (2009) demonstrate that

the introduction of even a small amount of competition leads to large behavioral changes. Roth *et al.* (1991) find that principals capture nearly the entire surplus when agents compete in a demand game, while Grosskopf (2003) shows, in an ultimatum bargaining game with responders' competition, that demands evolves more over time than in the game with no competition. Davis and Holt (1994) find that the ability of a buyer to switch between sellers provides a strong incentive to develop reputation in a repeated game. Charness, Kuhn and Villeval (forthcoming) show how competition can eliminate the ratchet effect. In contrast, other studies find that competition does not eliminate fairness when contracts are incomplete (see Brandts and Charness, 2004, and Fehr, Kirchler, Weichbold, and Gächter, 1998). Thus, it is not clear *ex ante* what effects competition will have on the static hidden-information problem. This is even more the case when competition is unbalanced, with either excess firms or excess workers.

3. THE MODEL

In this section we describe the theoretical model that serves as the basis for the experimental design. In this game, the principal offers one contract that is intended for low-ability-type agents and one contract intended for high-ability-type agents; these contracts are designed such that the agents have an incentive to self-select the appropriate contract. We vary the bargaining power by altering the relative proportion of principals and agents. As a preview, we note that the case with competition between principals (more principals than agents) yields a solution that is invariant to the number of agents and is sometimes efficient. On the other hand, the case of competition between agents is not invariant to the (relative) number of agents. The presence of more agents relaxes the binding incentive-compatibility constraint (for the high-ability type), yielding a level of effort that decreases towards the efficient level with the number

of agents. In the limit, the only relevant constraint for the high-ability type agent is the participation constraint. As a result, there are no inefficiencies.

Imagine that a firm needs one worker in order to be able to operate. The profits for the firm when it is operating are:

$$\Pi = e - w$$

where e, w are the effort levels and wages of the worker. Each worker has a utility function which depends on her ability type $j \in \{H, L\}$, which is her private information:

$$u_j(e, w) = w - \frac{k_j}{2} e^2$$

where $k_H = 1$ and $k_L = k > 1$. That is, the high-ability agent has a lower cost of effort than the low-ability type. Thus, only the individual agent knows j , but e is observable and contractible.

From the utility functions of the principal and the agents, the first-best effort levels are:

$$\hat{e}_j = \frac{1}{k_j}, j \in \{H, L\} \quad (1)$$

We call \hat{e}_j the efficient level of effort.⁵ Let \underline{U} be the outside option of the worker (which we assume for simplicity to be type-independent); we can induce optimal effort, with:

$$\hat{w}_j = \underline{U} + \frac{1}{2k_j}, j \in \{H, L\}$$

If the (independent) probability that an agent is a high- or low-ability type is denoted respectively by p_H or p_L , then the expected (optimal) profits for the principal are given by:

$$\Pi^E = \frac{p_L}{2k_L} + \frac{p_H}{2k_H} - \underline{U}$$

⁵ This is an appropriate terminology because in all the Pareto-efficient allocations of this problem (with complete information) the level of effort is always \hat{e}_j . This is so because of the quasi-linearity of the utility function of the agents, a common assumption in this field. Thus, the Pareto-efficient allocations only differ in the wages and profits of the principal and agent.

In order to make some comparisons across treatments we hold this first-best contract fixed in all the treatments. However, the second-best optimal equilibrium contracts, when the types are private information of the agents, depend on the structure of the market, which is our treatment variable. Then the equilibrium contract menu in the Benchmark (B) treatment, with one principal and one agent, results from the solution of the maximization program:

$$\begin{aligned} & \max_{w_H, w_L, e_H, e_L} p_H(e_H - w_H) + p_L(e_L - w_L) \\ \text{subject to} & \\ & w_H - \frac{k_H}{2}(e_H)^2 \geq \underline{U} \quad (\text{IR}_H) \\ & w_L - \frac{k_L}{2}(e_L)^2 \geq \underline{U} \quad (\text{IR}_L) \\ & w_H - \frac{k_H}{2}(e_H)^2 \geq w_L - \frac{k_H}{2}(e_L)^2 \quad (\text{IC}_H) \\ & w_L - \frac{k_L}{2}(e_L)^2 \geq w_H - \frac{k_L}{2}(e_H)^2 \quad (\text{IC}_L) \end{aligned}$$

where (IR_j) and (IC_j) are respectively the individual rationality and incentive compatibility constraints of an agent of ability type $j \in \{H, L\}$. As usual in these problems (see for example Mas-Colell, Whinston, and Green, 1995, ch.14C), it turns out that the active constraints in the optimal solution are (IR_L) and (IC_H) , so that the solution is:

$$\begin{aligned} e_H^B &= \frac{1}{k_H} = 1; \quad e_L^B = \frac{p_L}{k_L - k_H(1 - p_L)} = \frac{1}{k_L + \frac{1 - p_L}{p_L}(k_L - k_H)}; \\ w_L^B &= \underline{U} + \frac{k_L}{2}(e_L^B)^2; \quad w_H^B = \frac{1}{2} + w_L^B - \frac{1}{2}(e_L^B)^2 \end{aligned} \quad (2)$$

The high-ability type of agent provides the ‘efficient’ level of effort and obtains utility above \underline{U} . These informational rents (rents are defined here as the utility an agent gets above her reservation utility) are equal to:

$$w_H^B - \frac{1}{2} - \underline{U} = \frac{k_L - 1}{2}(e_L^B)^2 \quad (2e)$$

The effort of the low-ability type of agent is ‘inefficiently’ low and she obtains no rents, because she is held to the reservation value (the (IR_L) constraint is binding). This is the subgame-perfect equilibrium of this game.

Assume now that the each principal is matched with three agents (this is the Excess Agent, or EA , treatment). Then an equilibrium contract menu results from the solution of a slightly different maximization program. Given that high-ability types are ‘harder-working’ (they have a lower disutility of effort), they cost less per unit of output. Thus when any of the matched agents chooses the contract designed for the high-ability type, the principal always chooses her. If more than one agent chooses the high contract, the principal chooses randomly among those selecting the high contract.⁶⁷

$$\max_{w_H, w_L, e_H, e_L} (1 - p_L^3)(e_H - w_H) + p_L^3(e_L - w_L)$$

subject to

$$w_L - \frac{k_L}{2}(e_L)^2 \geq \underline{U} \quad (IR_L)$$

$$\left(\frac{(1 - p_L)^2}{3} + 2 \frac{(1 - p_L)p_L}{2} + p_L^2 \right) \left(w_H - \frac{k_H}{2}(e_H)^2 \right) + \left(1 - \frac{(1 - p_L)^2}{3} - 2 \frac{(1 - p_L)p_L}{2} - p_L^2 \right) \underline{U} \geq \frac{p_L^3}{3} \left(w_L - \frac{k_H}{2}(e_L)^2 \right) + \left(1 - \frac{p_L^3}{3} \right) \underline{U} \quad (IC_H)$$

The (IC_H) can also be written

$$w_H - \frac{k_H}{2}(e_H)^2 \geq q \left(w_L - \frac{k_H}{2}(e_L)^2 \right) + (1 - q)\underline{U} \quad (IC_H)$$

⁶ We only write the binding constraints, in what follows.

⁷ In our model the principal always prefers the high-ability agent. For this to be true we need to show:

$w_H - w_L \leq e_H - e_L$, but because of the binding IC_H , $w_H - w_L = \frac{k_H}{2}(e_H - e_L)(e_H + e_L)$, which means that $w_H - w_L \leq e_H - e_L$ is satisfied if $\frac{k_H}{2}(e_H + e_L) \leq 1$. This is true because $\frac{k_H}{2}(e_H + e_L) \leq \frac{k_H}{2}(2e_H) = \frac{k_H}{2} \left(2 \frac{1}{k_H} \right) = 1$.

where $q = \frac{p_L^3}{\frac{(1-p_L)^2}{3} + 2\frac{(1-p_L)p_L}{2} + p_L^2}$. The solution is now:

$$e_H^{EA} = \frac{1}{k_H} = 1; \quad e_L^{EA} = \frac{1}{k_L + \frac{(1-p_L)^3}{p_L^3} q(k_L - k_H)}; \quad (3)$$

$$w_L^{EA} = \underline{U} + \frac{k_L}{2} (e_L^{EA})^2; \quad w_H^{EA} = \frac{1}{2} + w_L^{EA} - \frac{1}{2} (e_L^{EA})^2$$

The effort of the low-ability type of agent in the *EA* treatment, e_L^{EA} , is closer to the efficient effort (\hat{e}_L) than that in the *B* treatment (e_L^B). To see this note that both e_L^{EA} and e_L^B are smaller than $\hat{e}_L = \frac{1}{k_L}$; we now show $e_L^{EA} > e_L^B$, so that the distortion is lower in *EA* than in *B*:

$$e_L^{EA} = \frac{1}{k_L + \frac{(1-p_L)^3}{p_L^3} q(k_L - k_H)} > e_L^B = \frac{1}{k_L + \frac{1-p_L}{p_L} (k_L - k_H)} \Leftrightarrow \frac{(1-p_L)^2}{p_L^2} q < 1 \quad (3')$$

$$\frac{(1-p_L)^2}{p_L^2} q = \frac{(1-p_L)^2}{p_L^2} \frac{\frac{p_L^3}{3}}{\frac{(1-p_L)^2}{3} + 2\frac{(1-p_L)p_L}{2} + p_L^2} = \frac{(1-p_L)^2 \frac{p_L}{3}}{\frac{(1-p_L)^2}{3} + 2\frac{(1-p_L)p_L}{2} + p_L^2} < \frac{(1-p_L)^2 \frac{p_L}{3}}{\frac{(1-p_L)^2}{3}} \leq 1$$

The reason for this enhanced efficiency is that the principal distorts the low-ability agent in order to lower the rents to the high-ability agent. To see this, note that the informational rents in the Benchmark treatment (equation 2e) are increasing in e_L^B , so the principal prefers to lower e_L^B (thus reducing efficiency) in order to get higher profits. But in the *EA* treatment there is a competitive pressure on the high-ability types. In fact, it is easy to check that in the general model where the principal confronts n agents, the difference between the equilibrium and the efficient level of effort for the low-ability type goes to zero as n goes to infinity.

Nevertheless, there is an additional problem with this treatment. We have found the equilibrium by assuming that the high-ability types assume that other high-ability types choose the high contract. But that is not the unique equilibrium here. In the second stage, where a menu

is offered, it is also possible that both types of agents select the low option for the menu. If all agents are choosing the low option, it is indeed a best response to choose low for all of them. But in this case, it need not be optimal to propose the menu of contracts specified in (3). In the design of the experiment we provide another menu, which is the equilibrium under the assumption that whenever there is multiplicity of equilibria in the second stage, the worst equilibrium for the principal is selected. The equilibrium menu in that case would solve:

$$\max_{w_H, w_L, e_H, e_L} (1 - p_L^3)(e_H - w_H) + p_L^3(e_L - w_L),$$

subject to

$$w_L - \frac{k_L}{2}(e_L)^2 \geq \underline{U} \quad (\text{IR}_L) \quad (4)$$

$$w_H - \frac{k_H}{2}(e_H)^2 \geq \frac{1}{3}\left(w_L - \frac{k_H}{2}(e_L)^2\right) + \frac{2}{3}\underline{U} \quad (\text{IC}_H)$$

where the incentive constraint now ensures that it is dominant to choose the high option for a high-ability type (thus she will do it independently of what other individuals of her type are doing).⁸ Choosing the high contract when the low contract gives a higher payoff makes sense to reduce competition from other workers. In fact the ‘attractiveness’ of the high contract increases with the probability that a competing worker also chooses the high option. So the worst-case scenario for the principal is when no competitor chooses the high contract. If even in that case a high type should choose the high over the low contract, then it is dominant for a high-ability type to choose the high option, and that is exactly what the IC_H constraint in equation (4) does.

⁸ If the contract offered by the principal did not satisfy the (IC_H) constraint in program (4), there would be a pooling equilibrium in the contract acceptance subgame, where both the H and L types would accept the L contract. To see this, notice that the utility for the H type of accepting the H contract would be: $w_H - \frac{k_H}{2}(e_H)^2$, since by being the only agent departing from the pooling strategy, he would guarantee being chosen. On the other hand, if he chose to accept the L contract he would get a utility of: $\frac{1}{3}\left(w_L - \frac{k_H}{2}(e_L)^2\right) + \frac{2}{3}\underline{U}$, since he would only be chosen one-third of the time. Thus, if the (IC_H) constraint in program (4) is violated it is indeed optimal for the H type to pool with the L type. If the pooling equilibrium is always selected when available, the optimal way to screen types is given by the solution to (4). As usual, one still has to check that screening types is optimal for the principal.

Finally, we also have a treatment (Excess Principals, or *EP*) where two principals compete for one agent. In that case, competition among principals forces the equilibrium of the game to give agents the maximum possible rents subject to incentive compatibility. A candidate for equilibrium would thus be a menu of contracts where agents make the efficient level of effort and principals make zero profits for each type of contract.

$$\begin{aligned} e_H^{EP} &= w_H^{EP}; & e_L^{EP} &= w_L^{EP}; \\ e_L^{EP} &= \frac{1}{k_L}; & e_H^{EP} &= \frac{1}{k_H} \end{aligned} \quad (5)$$

We just need to check at this point that this menu of contracts is incentive compatible to show that it is an equilibrium.⁹ This is true whenever:

$$\frac{1}{k_H} - \frac{1}{2}k_H \left(\frac{1}{k_H} \right)^2 \geq \frac{1}{k_L} - \frac{1}{2}k_H \left(\frac{1}{k_L} \right)^2 \quad (6)$$

$$\frac{1}{k_L} - \frac{1}{2}k_L \left(\frac{1}{k_L} \right)^2 \geq \frac{1}{k_H} - \frac{1}{2}k_L \left(\frac{1}{k_H} \right)^2 \quad (7)$$

But notice that (6) and (7) are equivalent to

$$1 \geq \frac{1}{2}k_H \left(\frac{1}{k_H} + \frac{1}{k_L} \right) \quad (8)$$

$$\frac{1}{2}k_L \left(\frac{1}{k_H} + \frac{1}{k_L} \right) \geq 1 \quad (9)$$

One can easily check that (8) and (9) are indeed satisfied for the parameters of our experiment, $k_L=2$, $k_H=1$.

We implemented the theoretical model in our experiment by choosing a single set of six menus allowable in all contracts. For the parameter values $k_L=2$, $k_H=1$, $p_L=1/2$, equation (3) leads to menu 1, equation (4) induces menu 2, equation (2) leads to menu 3, and equation (5)

⁹ This equilibrium can be shown to be unique. See, e.g. Von Siemens (2010).

induces menu 6.¹⁰ We provide the details of these mappings into experimental payoffs in section 3. In addition we chose two non-equilibrium menus, in order to provide a richer contractual environment. Menu 4 is similar to menu 3, but has a little more effort for the low-ability type, and respects the IC constraint for the high-ability type. Menu 5 is fully efficient, for both types.

Each menu consisted of a choice of two (enforceable) effort levels and payments that depend on the type of agent involved; if neither choice seemed attractive to the agent, she could veto the contract menu. We chose $k_L = 2$ for all menus, in order to give relatively large rents to the high-ability type (under her preferred contracts). The parameters, efforts, and wages for the six different menus in the experiment are summarized in Table 1:

Table 1 – Parameter Values

Menu	k_L	p_L	e_H	e_L	w_H	w_L
1	2	1/2	1	0.36	0.64	0.25
2	2	1/2	1	0.23	0.64	0.18
3	2	1/2	1	0.33	0.70	0.25
4	2	1/2	1	0.4	0.75	0.33
5	2	1/2	1	0.50	0.85	0.44
6	2	1/2	1.50	0.50	1.50	0.50

One of the criticisms of models of contract design with hidden information is that the contract menus are more ‘complex’ than one observes in reality. In an environment like ours, these often employ a nonlinear structure and a very large number of possible choices of pairs of wages and efforts. Using a continuous strategy space would be quite complicated to design for the principal, and even the choice of the agent would not be simple without adding much insight; this would also make the data analysis problematic. We have selected a relatively small number

¹⁰ Due to a misunderstanding, the correct theoretical parameter values were not used for Menus 5 and 6. Instead of these menus, there should be a menu with parameters $e_H = 1$, $e_L = 0.5$, $w_H = 1$ and $w_L = 0.5$. This would be the equilibrium menu in the excess-principals treatment. Nevertheless, given the available menus in Table 3, Menu 6 is

of menus of contracts; since they include all forms of equilibrium in some versions of the game, this number is sufficiently large for exploring how the choice of contracts is affected by relative bargaining power. While we have selected a very simple structure (only two types), we feel that a ‘simple’ menu can serve as an approximation for a full schedule.

4. EXPERIMENTAL DESIGN

We conducted three different treatments, which differed according to the numbers of principals and agents in the treatment to study how offered contracts depend on the structure of bargaining power. In our *B* treatment, there were 10 principals and 10 agents in each session. In the *EA* treatment, there were four principals and 12 agents, while in the *EP* treatment, there were 12 principals and six agents.¹¹ In all cases, there were equal numbers of high-ability (H) agents and low-ability (L) agents and this was made common knowledge among the participants. In order to observe roughly similar numbers of observations (matches) in each treatment, we conducted four sessions of the *EA* treatment, three sessions of the *EP* treatment, and two sessions of the *B* treatment. Each session consisted of 40 periods of play to allow for possible learning dynamics, with random and anonymous re-matching after every period. The re-matching procedure was common information to the participants.

The organization of our sessions is summarized in Table 2. The participants were privately informed of their roles; agents were also informed of their type. One’s role and/or type were kept constant throughout the session. The participants also knew that there were the same number of L and H agents in the room.

the equilibrium one in the *EP* treatment, as it is the menu that leads to the highest utility for the agents and it is incentive compatible. The only real difference is that the rents for the high-type agent are somewhat smaller.

¹¹ We chose three agents per principal in the *EA* treatment because the theoretical model shows that the distortion between the efficient level of effort and the equilibrium effort reduces in the number of competing agents; in

addition, it increases the probability to be matched with at least one high agent. In the *EP* treatment, we chose to match two principals with one agent, as the theoretical predictions are unaffected by adding more principals.

Table 2 – Treatments and sessions

Treatment	Participants per session			Sessions	Periods	Observations
	Principals	H-agents	L-agents			
Benchmark	10	5	5	2	40	800
Excess Agent	4	6	6	4	40	640
Excess Principal	12	3	3	3	40	720
Total	72	43	43	9	-	2160

In our *B* treatment, the “proposer” (principal) first makes a selection from among the six “offers” (feasible contract menus). The “responder” (agent) is informed of this choice, and then selects “option X” (high contract), “option Y” (low contract), or rejects the contract menu. Each person then learns his or her payoff and play then continues on to the next period. The sequence in the *EA* treatment is similar, except that the principal is informed of the options chosen by each of the three agents and then selects one of these agents. No agent is informed about the choices of the two other agents. The *EP* treatment has the same sequence as the *B* treatment, with the proviso that an agent can accept at most one offer from the two principals with whom she is paired. When both principals make the same offer, the agent chooses at random between them if she is willing to accept the offer. The principal is not informed of the offer of the other principal.

We used the parameter values in Table 1 to generate experimental payoffs for the feasible contract menus. We first derived the payoffs from these parameters to three decimals and then multiplied these by one thousand. We next rounded these payoffs to the nearest multiple of 5. In the case of the principals, we added 250 to each of the non-rejection payoffs; this reflects the notion that setting up the firm requires some capital, and the minimum level of revenues that are needed to recoup the cost of capital is 250. In the case of the agents, we added 10 to each non-rejection payoff, in order to provide some minimal separation (avoiding the possibility of equilibrium failure due to indifference) between the payoff for a low agent who accepts the least

favorable offer and her payoff from rejecting the contract menu in its entirety.¹² Unmatched principals or agents received 125 points in the period.¹³ This process leads to Table 3:

Table 3 – Payoff Table

		Option X	Option Y	Reject
Menu 1	P	610	355	125
	H	150	200	125
	L	-350	135	125
Menu 2	P	605	305	125
	H	155	160	125
	L	-345	135	125
Menu 3	P	550	335	125
	H	210	200	125
	L	-310	145	125
Menu 4	P	500	350	125
	H	260	230	125
	L	-240	160	125
Menu 5	P	400	310	125
	H	360	325	125
	L	-140	200	125
Menu 6	P	250	250	125
	H	385	385	125
	L	-740	260	125

A form of Table 3 was distributed to the subjects to help them to make their decisions.¹⁴

The sessions were conducted at the Groupe d'Analyse et de Théorie Economique (GATE), Lyon, France. Participants were recruited from undergraduate courses in local

¹² As it happens, we inadvertently added 20 points to the L payoffs from option Y with menu 4. Perhaps this turns out to be useful for testing what is needed to obtain efficiency. The reason is that even with this extra kick, the *B* treatment is least efficient once rejections are considered. Thus, there is an argument that competition between agents is good for efficiency because it reduces informational rents, both in theory and in practice. And that principal competition enhances efficiency as it reduces the envy-driven rejections that hurt efficiency in the benchmark.

¹³ We note that adding the same constant to the payoffs for all agents or to the payoffs for all principals cannot change the equilibrium, since these are Von Neumann-Morgenstern utilities.

Engineering and Business schools using ORSEE (Greiner, 2004). All the participants were inexperienced in this type of game and no one participated in more than one session. On average, a session lasted 60 minutes, including initial instructions. The experiment was computerized using the REGATE program developed at GATE (Zeiliger, 2000). We used a conversion rate of 100 points for each Euro. At the end of each session, we selected randomly four of the 40 periods for actual payment. In this way, we avoided possible income effects from having already accumulated a known amount of money in the session. The average payoff was 14.9 Euros in the *B* treatment, and 13.5 Euros in both the *EA* and the *EP* treatments; on average the principals received 17 Euros, the high-ability agents 13 Euros and the low-ability agents 10 Euros, including a four Euro show-up fee.

5. EXPERIMENTAL RESULTS

An overview of our experimental results is that we find substantial treatment effects in our sessions, with large differences in the contract menus offered and accepted, substantially in line with the equilibrium predictions. The menus that are offered (and accepted) evolve over time. In general, rejections and competition drive behavior. We first give descriptive statistics for principal behavior and agent behavior. We then consider the determinants of such behavior, providing statistical tests and regression analysis.

5.1 Descriptive statistics

Principal behavior

While there is certainly heterogeneity present among the principals, we do observe some clear patterns and differences for the menus chosen in each treatment. Table 4 lists the menus

¹⁴ See the instructions in Appendix A.

offered in each treatment; the left panel considers the full 40 periods of the game, while the right panel considers only the last 20 periods to show the evolution of contract offers over time.

Table 4 – Menus offered by treatment

Menu	Menus offered in all periods			Menus offered in the last 20 periods		
	<i>B</i> treatment	<i>EA</i> treatment	<i>EP</i> treatment	<i>B</i> treatment	<i>EA</i> treatment	<i>EP</i> treatment
1	89 (11.12)	181 (28.28)	56 (3.89)	29 (7.25)	75 (23.44)	14 (1.94)
2	17 (2.12)	34 (5.31)	28 (1.94)	7 (1.75)	14 (4.38)	9 (1.25)
3	242 (30.25)	348 (54.38)	57 (3.96)	103 (25.75)	199 (62.19)	23 (3.19)
4	361 (45.13)	58 (9.06)	136 (9.44)	218 (54.50)	24 (7.50)	42 (5.83)
5	75 (9.38)	16 (2.50)	490 (34.03)	36 (9.00)	7 (2.19)	191 (26.53)
6	16 (2.00)	3 (0.47)	673 (46.74)	7 (1.75)	1 (0.31)	441 (61.25)
Total	800 (100)	640 (100)	1440 (100)	400 (100)	320 (100)	720 (100)

Note: Equilibrium menus are in **bold** and percentages are in parentheses.

Table 4 provides some support for the equilibrium predictions, as 30.25% of the menus offered in the *B* treatment, 33.59% of the menus offered in the *EA* treatment, and 46.74% of the menus offered in treatment *EP* are equilibrium menus (menus 1 and 2 with *EA*, menu 3 with *B*, and menu 6 with *EP*). However, the mode of the distribution corresponds to the equilibrium menu only in the *EP* treatment. Indeed, menu 4 is the most common choice in the *B* treatment (45.13% of the offers) and menu 3 is the most common choice in the *EA* treatment (54.38% of the offers). As these menus offer a more egalitarian distribution than the equilibrium menus in both the *B* and *EA* treatments and also provides greater efficiency (higher total payoffs), one might suspect that social preferences such as those expressed in the Charness and Rabin (2002) model play a role here. However, if these deviations were only driven by the principals' intrinsic social preferences, we should expect to observe them from the beginning of the game. Instead, the menus offered evolve over time, as indicated by the right panel of Table 4.

Indeed, in both the *B* and the *EA* treatments, the contracts offered deviate from the equilibrium even further in the last 20 periods. In this set of periods, menu 4 represents 54.50% of the offers and the equilibrium menu 3 only 25.75% of the offers in the *B* treatment; menu 3 represents 62.19% of the offers while the equilibrium menus 1 or 2 represents only 27.82% of the offers on the *EA* treatment. Since it is the *ex ante* more generous menus that grow in frequency over time, it seems that some other force is inducing principals to choose them more often over time, even though principals' social preferences possibly also account for these choices. This is not contradicted by the observations from the *EP* treatment, since if the share of the equilibrium menu 6 increases over time (representing 46.74% across all periods but 61.25% in the last 20 periods), this menu also corresponds to the most generous menu.

While we postpone an in-depth analysis of the determinants of the contract menu choices, it appears that the profitability of a particular contract menu depends greatly on the environment. We find that menus 1 and 2 yield among the highest average profits to the principals in the *EA* treatment (516 and 528 points, respectively), but generate among the lowest average profits in the *B* (270 and 210 points, respectively) and *EP* treatments (146 and 125 points, respectively). Similarly, menu 6 is quite unattractive for the principal in the *EA* and *B* treatments (yielding 250 points on average in both), but provides nearly the best profits in *EP* (214 points compared to 219 provided by menu 5). Overall, we observe a good correspondence between the most frequent offers made and their profitability. For example, menu 4 *ex post* yields higher profits for the principals than the equilibrium menu 3 in the *B* treatment (392 vs. 334 points). Thus, to a large extent, it seems that principals are influenced by considerations of their own profits.

Agent behavior

The menus accepted by the agents naturally mirror the menus that were offered; however, there are some substantial differences, due primarily to rejections in the *B* treatment and selection

pressures in the *EP* treatment. Table 5 displays the number of accepted contracts and the acceptance rate for each menu, by treatment and by agent's type. In total, there are 800 observations in the *B* treatment, 1920 in the *EA* treatment (since there is one offer for three agents), and 1440 observations in the *EP* treatment (since there are two principals for one agent).

Table 5 – Acceptance rates, by treatment and type

Menu Offer	<i>B</i> treatment		<i>EA</i> treatment		<i>EP</i> treatment	
	<i>High type</i>	<i>Low type</i>	<i>High type</i>	<i>Low type</i>	<i>High type</i>	<i>Low type</i>
1	36/45 (80)	18/44 (41)	271/286 (95)	212/257 (82)	1/20 (5)	4/36 (11)
2	6/10 (60)	2/7 (29)	46/49 (94)	47/53 (89)	0/11 (0)	0/17 (0)
3	97/103 (94)	85/139 (61)	513/522 (98)	446/522 (85)	2/23 (9)	2/34 (6)
4	191/200 (96)	134/161 (83)	82/82 (100)	92/92 (100)	10/59 (17)	18/77 (23)
5	37/37 (100)	36/38 (95)	16/16 (100)	32/32 (100)	102/260 (39)	99/230 (43)
6	5/5 (100)	11/11 (100)	5/5 (100)	4/4 (100)	245/347 (71)	235/326 (72)
Total	372/400 (93)	286/400 (72)	933/960 (97)	833/960 (87)	360/720 (50)	358/720 (50)

Note: These numbers represent accepted contracts over offered contracts. Equilibrium menus are in **bold** and percentages are in parentheses.

When they accept an offer, low-ability agents rarely (14 of 1477 times for all treatments pooled together, or 0.95%) chose option X (high effort), which would generate negative earnings. The behavior of high-ability agents is more complex. Note that option X pays more than option Y for high-ability types with menus 3-5, but option Y pays more with menus 1 and 2. In addition, both options give the same payoff to the high-ability type with menu 6. In the *EP* treatment, high-ability agents nearly always maximize own payoffs (356 of 360 non-rejections). While they also do so with menus 1, 2, 5, and 6 in treatment *B*, we observe a substantial proportion of option-Y choices with menus 3 and 4 (42% and 21% of the non-rejection choices, respectively). In the *EA* treatment, the principals know the agents' choices prior to selecting one,

so that agents must compete in their choices to be selected. High-ability agents do maximize own profits with menus 3-5, since this maximization also coincides with maximizing the profits for the principal. However, with menus 1 and 2, an agent who myopically chooses option Y runs the risk that she will not be selected if another agent has chosen option X. Since some agents appear to realize this, we see that 64% of the high-ability agents choose option X when accepting menus 1 or 2. There is some evidence of learning over time as this percentage increases from 58% in the first 20 periods to 73% in the last 20 periods. In contrast, in response to the offer of menu 3, the corresponding percentage (95%) for choosing option X is stable over time.

As a consequence of these decisions, the proportion of high-ability agents actually recruited is larger in *EA* (81%) than in *B* (57%). While in *B* this distortion of the initial distribution of the population reflects the higher frequency of rejections of offers by low-ability agents, in *EA* this reflects the process of selection related to the competitive environment. In *EP*, the proportion of high-ability agents (50%) corresponds roughly to the initial distribution of the population since almost all the agents accept an offer.

5.2 Regression analysis

We now turn to multiple-regression analysis of the determinants of the observed behavior in each treatment, first considering the menus offered by principals.^{15,16}

Principal behavior

¹⁵ All of these estimations used robust standard errors and clustering at the individual level to account for the fact that a same individual makes several decisions over time.

¹⁶ We also performed nonparametric Mann-Whitney rank-sum tests, with both session-level and individual-level data (see Appendix B). These tests find that the average menu offered is lowest in the *EA* sessions and highest in the *EP* sessions. They also indicate that rejection rates of menus below 4 are significantly higher in the *B* than in the *EA* treatments and that high-ability agents in *B* are more likely to choose option X in response to menus 1 and 2 than in *EA*, the reverse being true for menu 3. Overall, they show that the proportion of high-ability agents is higher in *EA* than in *B*.

In order to disentangle the motivations underlying offers of the different categories of menus in the *B* treatment, we estimate a multinomial Logit model in which the reference category of the dependent variable is the offer of the equilibrium menu 3 (see the first three columns of Table 6). The two other categories correspond either to the offer of menu 1 or 2, or to the offer of menu 5 or 6. This model does not make any assumption about the order between the menus offered. Next, we estimate multinomial Logit models to explain the offer of the most frequent contract, i.e. menu 4 in the *B* treatment (see the last two columns in Table 6). The reference category is the offer of menu 4 both in period t and in period $t-1$. The other categories correspond to the switch to menu 4 in period t from menus 1, 2 or 3 offered in period $t-1$ and to the switch to menu 4 from menus 5 or 6 offered in period $t-1$.

Table 6: Determinants of the choice of menus in the Benchmark Treatment

Baseline treatment	Ref.: Offer of the equilibrium menu 3			Ref.: Offer of the most frequent menu 4 in both $t-1$ and t	
	Offer of menus 1 - 2	Offer of menu 4	Offer of menus 5 - 6	Switch from menus 1,2 or 3 in $t-1$	Switch from menus 5 or 6 in $t-1$
Time trend	-0.051** (0.023)	0.057*** (0.014)	0.008 (0.017)	-0.080*** (0.022)	-0.041*** (0.015)
Current rejection rate	0.043** (0.022)	-0.053** (0.022)	0.014 (0.020)	0.044# (0.027)	0.045 (0.029)
X option chosen in ($t-1$)	-1.366*** (0.432)	0.176 (0.171)	-0.355 (0.334)	-1.557*** (0.457)	-0.765*** (0.257)
Constant	-0.376 (0.440)	-0.053 (0.292)	-1.183*** (0.462)	0.216 (0.526)	-1.467** (0.587)
Nb obs.	641			316	
Log-likelihood	-701.339			-200.879	
Wald χ^2	38.06			65.26	
Prob> χ^2	0.000			0.000	
Pseudo R ²	0.087			0.112	

Note: These estimations have been conducted with robust standard errors (in parentheses) and clustering at the individual level; *** and ** denote two-tailed statistical significance at the 1%, and 5% level, respectively. # indicates borderline significance at 10.2% level.

In all the regressions, the independent variables include a time trend (i.e. the period index). They also control for the agent's choice of high effort (option X) in the previous period and the current rejection rate (defined as the current proportion of periods in which an offer from the principal has been rejected). These three variables are intended to capture how principals modify their behavior depending on their cumulative experience over time.

Relative to the equilibrium menu 3, menus 1 and 2 are chosen less frequently over time. A principal who experiences a high proportion of rejections of his previous offers is more likely to offer menus 1-2 and less likely to offer menu 4. This may suggest some inertia in behavior. Indeed, a principal who is driven by social preferences makes more generous offers, experiences fewer rejections, and is encouraged to continue to offer generous menus. An agent choosing the X option in the previous period decreases the likelihood of menu 1 or 2 being chosen by the principal as offering these menus would induce both types to switch to the low option.

The last two columns show that switching from any menu to the most frequent menu 4 is more frequent at the beginning of the game. It is less likely if the agent has chosen the X option in the previous period. The experience of a high proportion of rejections in the past exerts a borderline significant effect on switching from menus 1, 2 or 3 to the more generous menu 4. In contrast, offering very generous menus 5 or 6 is not conditioned on these variables and could be motivated by an unconditional preference for more egalitarian outcomes (or a fear of rejection without having previously experienced this rejection).

Table 7 presents the estimates of two multinomial Logit models studying the determinants of the menus chosen in the *Excess Agent* treatment. In the first model, the reference category is the offer of the equilibrium menus 1 or 2 pooled together. The first two columns correspond respectively to the offer of menu of 3 and the offer of menus 4 to 6 (pooled, since they represent

only 12% of the observations). The second model studies the determinants of the switch to the most frequent menu offered. The reference category is the offer of menu 3 in both periods $t-1$ and t . The other categories correspond to the switch to menu 3 in period t from menus 1 or 2 offered in period $t-1$ and to the switch to menu 3 from menus 4, 5 or 6 offered in period $t-1$. The independent variables include a time trend and the choice of option X in the previous period, which is also interacted with the equilibrium menu. We do not consider the lagged rejection rate since all the contracts are accepted.

Table 7: Determinants of the choice of menus in the Excess Agent Treatment

Excess-Agent treatment (Multinomial Logit)	Ref.: Offer of the equilibrium menu 1-2		Ref.: Offer of the most frequent menu 3 in both t-1 and t	
	Offer of menu 3	Offer of menu 4, 5 or 6	Switch from menus 1,2 in t-1	Switch from menus 4, 5 or 6 in t-1
Time trend	0.019** (0.009)	-0.021 (0.013)	-0.022 (0.020)	-0.040* (0.022)
X option chosen in (t-1)	1.673*** (0.464)	0.345 (0.369)	-1.885*** (0.342)	-0.183 (0.744)
X option * menu 1-2 in (t-1)	-3.679*** (0.738)	-2.228*** (0.695)	-	-
Constant	-0.235 (0.473)	-0.121 (0.456)	-0.253 (0.440)	-1.432* (0.801)
Nb obs.	624		345	
Log-likelihood	-466.574		-187.251	
Wald χ^2	34.39		39.62	
Prob> χ^2	0.000		0.000	
Pseudo R ²	0.212		0.074	

Note: These estimations have been conducted with robust standard errors (in parentheses) and clustering at the individual level; ***, **, and * denote two-tailed statistical significance at the 1%, 5%, and 10% level, respectively.

Menu 3 is chosen more frequently over time whereas the frequency of menus 4-6 decreases but not significantly. If the X option was chosen in the previous period in response to the offer of the equilibrium menu, this decreases the likelihood of non-equilibrium menus being

offered; since the X option leads to a higher payoff for the principal, there is less motivation to offer a more generous menu. On the other hand, if the X option was chosen in $t-1$ (in reaction to whichever offers), this increases the likelihood of menu 3 being chosen. If we separate out the case when menu 3 is chosen (columns 3 and 4), we see that switching from the equilibrium menu 1-2 to menu 3 is much *less* likely when the X option has been chosen in the previous period.

Table 8 presents first the estimates of a multinomial Logit model studying the determinants of the menus chosen in the Excess Principal treatment. The reference category is the offer of the equilibrium menu 6. The first two columns correspond respectively to the offer of menus 1 to 4 (pooled together since they only represent 19% of the observations) and to the offer of menu 5. The third column displays the results of a Logit model in which the switch from menus 1 to 5 to the most frequent (and equilibrium) menu offer is the explained variable. We include among the independent variables both the acceptance of the principal's offer in $t-1$ and the current rejection rate (as defined above) instead of the lagged option chosen by the agent, in order not to eliminate half of the observations and because both the principal and the high-type agent are indifferent between the two options with menu 6.

The offer of less egalitarian menus than menu 6 decreases strongly over time. Although the acceptance of the offer in the previous period has little effect on the choice of any kind of menu, there is a strong and positive correlation between a high rejection rate in the past and the offer of less generous menus. In contrast, the Logit regression regarding switching to menu 6 shows that an offer rejected in the previous period makes such a switch more likely, while the overall current rejection rate exerts is no longer significant. The switch is also more likely in the early periods of the game. The competitive pressure to choose a favorable menu is naturally a major factor in the principal's choice of contract to offer.

Table 8: Determinants of the choice of menus in the Excess Principal Treatment

Excess-Principal treatment	Ref.: Offer of the equilibrium menu 6 (Multinomial Logit)		Probability to switch from menus 1-5 in t-1 to menu 6 (Logit)
	Offer of menus 1-4	Offer of menu 5	
Time trend	-0.083*** (0.011)	-0.062*** (0.010)	-0.069*** (0.013)
Acceptance of the offer in (t-1)	-0.166 (0.223)	0.111 (0.175)	-3.202*** (0.329)
Current rejection rate	0.033*** (0.009)	0.012 (0.010)	0.022 (0.014)
Constant	-0.929* (0.517)	0.392 (0.516)	0.971 (0.879)
Nb observations	1404		671
Log Likelihood	-1312.510		-267.631
Wald χ^2	125.25		103.66
Prob> χ^2	0.000		0.000
Pseudo R ²	0.090		0.352

Note: These estimations have been conducted with robust standard errors (in parentheses) and clustering at the individual level; *** and * denote two-tailed statistical significance at the 1% and 10% level, respectively.

Agent behavior

We next analyze the agents' decisions of whether to accept a contract menu and which option to choose if accepting a contract. Table 9 reports the estimates of the probability of accepting an offer, using Probit regressions with robust standard errors and clustering at the individual level. Model (1) pools the data from the *B* and *EA* treatments and the next three models consider each treatment separately. Note that regarding the *EP* treatment, we explain the likelihood of an acceptance from the principal's and not from the agent's point of view since, with the exception of two rejections of both offers, the agent always accepts one offer and selects the best one among the two, or chooses at random in case of a tie.

The independent variables include a time trend to control for the evolution of behavior over time and the agent's type. They also include the menus offered in the current and in the

previous periods (except in *EP* since we consider the principal's point of view).¹⁷ Indeed, although not predicted by the model, it is possible that the likelihood of accepting an offer is affected by the quality of offers experienced in the past. Finally, in the regression for the *EA* treatment we also control for the relative frequency of no selection in the previous periods, as it may affect the current decision of accepting an offer whatever the menu offered.

Table 9: Determinants of accepting a contract offer (Probit models)

Decision to accept the offer	<i>B</i> and <i>EA</i> Treatments	Benchmark Treatment	Excess Agent Treatment	Excess Principal Treatment
Time trend	-0.001 (0.003)	-0.010** (0.004)	0.001 (0.005)	-0.025*** (0.003)
Excess Agent Treatment	0.793*** (0.280)	-	-	-
High-type agent	0.912*** (0.242)	1.010*** (0.395)	1.322*** (0.400)	-0.109 (0.078)
Menu offered (from 1 to 6)	0.314*** (0.057)	0.482*** (0.084)	0.186** (0.079)	0.729*** (0.066)
Menu offered in <i>t</i> -1 (from 1 to 6)	-0.029 (0.035)	-0.102** (0.049)	-0.021 (0.035)	-
Relative frequency of no selection	-	-	0.023*** (0.005)	-
Constant	-0.362 (0.350)	-0.485 (0.474)	-0.742 (0.566)	-3.235*** (0.342)
Nb observations	2652	780	1832	1440
Log Likelihood	-773.506	-291.556	-363.406	-822.684
Wald χ^2	49.43	38.19	66.09	127.62
Prob> χ^2	0.000	0.000	0.000	0.000
Pseudo R ²	0.157	0.209	0.199	0.176

Note: These estimations have been conducted with robust standard errors (in parentheses) and clustering at the individual level; ***, and ** denote two-tailed significance at the 1% and 5% level, respectively.

To analyze the determinants of the choice of option X once a contract has been accepted, we have estimated a Probit model with robust standard errors and clustering at the individual

¹⁷ We include one single menu variable instead of each menu value since the agents' payoffs increase with each

level on the data of the *B* and the *EA* treatments separately (not reported here but available upon request).¹⁸ The independent variables include both a time trend (the number of the period) and the number of the menu that has been accepted by the agent. Not surprisingly, the higher the contract, the more likely the high agent is to choose option X (*B* treatment: coeff. = 0.840, $p < 0.001$, marginal effect = 0.273; *EA* treatment: coeff. = 0.517, $p < 0.001$, marginal effect = 0.104). The marginal effect of the menu offered is lower in *EA* than in the *B* treatment, probably due to the competitive pressure. Option X is chosen less frequently over time in the *B* treatment (coeff. = -0.013, $p = 0.005$, marginal effect = -0.005), as it seems that high-ability agents develop a taste for more favorable contracts and make modest sacrifices by choosing Y instead of X (rather than costly rejections) to punish the principal for a lack of generosity in choosing menus that give them lower payoffs. In contrast, option X is chosen more frequently over time in the *EA* treatment (coeff. = 0.014, $p = 0.056$, marginal effect = 0.003) since the high-ability agents also learn to make modest sacrifices, by choosing X instead of Y when they accept contracts 1 or 2, to increase the likelihood of their hiring.

5.3 Welfare and efficiency comparisons

Welfare comparisons across treatments are complicated because in *EA* and *EP* there is an unbalanced structure of principals and agents, so that some parties remain unmatched. The right assumption in this case, we think, is to consider only the individuals who are actually matched; equivalently, one might think of this as how much benefit society derives from a match. More

higher menu value.

¹⁸ Since low-type agents incur serious losses if they choose the X option, we only consider here the high-type agents. We also omit the *EP* treatment since payoffs are identical for both options X and Y when the equilibrium menu is offered. In the regression for the *B* treatment, $N = 367$, log-likelihood = -183.352, Wald $\chi^2 = 24.52$, $\text{prob} > \chi^2 < 0.001$, and pseudo $R^2 = 0.227$. In the regression for the *EA* treatment, $N = 933$, log-likelihood = -341.993, Wald $\chi^2 = 37.74$, $\text{prob} > \chi^2 < 0.001$, and pseudo $R^2 = 0.159$.

importantly, the numbers of matches with the high-ability agents is likely to be higher in the *EA* treatment, which necessarily gives a boost to the total payoffs in this treatment.

With the contract menus and payoff parameters we chose, we find that the greatest benefits accrue to society when there are many agents competing. The average total payoffs are 698.97, 538.46, and 590.42 in the *EA*, *B*, and *EP* treatments, respectively; thus, the *EA* treatment easily yields the highest average total payoff. One reason for this is that (*ex post*) every contract is accepted by at least one of the paired agents in the *EA* treatment. The *EP* treatment slightly dominates *B* in terms of average total payoff since only two agents reject the offer of both principals, while the *B* treatment has a substantial rejection rate. ◦

Nevertheless, even if we consider only accepted contracts, the *EA* treatment still generates the highest degree of efficiency, with average total payoffs of 698.97, 600.71, and 591.36 in the *EA*, *B*, and *EP* treatments, respectively. Mann-Whitney tests find that the total payoffs are higher in *EA* than in both *B* ($Z = 1.85$ and $p = 0.06$) and *EP* ($Z = 2.14$ and $p = 0.03$). One reason why *EA* empirically yields higher payoffs than *B* or *EP* is that the possibility of selection means that there are more matches with high-ability agents in *EA* (81%, against 57% in *B* and 50% in *EP*) and these yield higher total payoffs per match. Nevertheless, even if there were an equal number of matches with high- and low-ability agents in *EA*, it would still be the treatment with the highest total payoffs, either with or without rejections (615.04).

However, the conclusion about the higher efficiency of the *EA* institutional environment should be qualified if we account for those agents who get unmatched. If the *EA* environment provides the biggest benefit to society from a match, it may generate a higher social cost than the benchmark if the unmatched agents remain unemployed.

6. CONCLUSION

We conducted an experiment based on a model of contracting under asymmetric information. We show theoretically that, in this context, various degrees of relative bargaining power affect outcomes and efficiency. In this environment, efficiency improves in the relative number of agents because competition reduces the agents' informational monopoly power. However, this environment also generates high inequality levels and is characterized by multiple equilibria, which may have important behavioral implications in the field and suggests that empirical testing could produce valuable insights.

Our results provide qualitative support for the theory. We find that the institutional environment in which agents compete against each other improves efficiency with respect to a situation without competition. We also show that behavior evolves over time. People make errors, learn and adjust their decisions accordingly in order to increase their payoffs. In particular, the payoff obtained in the previous period, especially related to the ability to separate between types, is a driving force of the evolution of principals' menu offers. Our results also indicate that principals offer more generous menus than predicted, although less frequently in the context of competition.

These results offer some insight into the interaction between the various degrees of bargaining power and social preferences. In many experimental papers, the outcomes do not correspond to the predictions of standard contract theory; this is often considered to stem from some form of social preferences. In our treatment without competition, we observe that principals offer more generous contracts menus than in the equilibrium. When agents compete, principals tend also to offer more generous contract menus than the equilibrium, but this is less the case when they are able to separate the agents by type with the equilibrium menus; this calls

into question the true generosity of these offers. In addition, the existence of social preferences can hardly change the outcome for the low-ability type agents: due to the heterogeneity among agents, offering a more generous menu increases the selected agent's expected payoff but also increases the likelihood of the repeated exclusion of the low-ability agents.

Finally, the superiority of the institutional environment with competition among agents is shown in terms of total surplus of the matched pairs. The higher total surplus is achieved by making the payoff of the principal higher and lowering those of the agents. Thus, there is a genuine tradeoff between equity and efficiency in this environment, both theoretically and empirically.

REFERENCES

- Akerlof, G. (1970). The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics*, 84, 488-500.
- Anderhub, V., Gächter, S., and Königstein, M. (2002), Efficient Contracting and Fair Play in a Simple Principal-Agent Experiment. *Experimental Economics* 5(1), 5-27.
- Banks, J., Camerer, C.F., and Porter, D. (1994). An experimental analysis of Nash refinements in signaling games. *Games and Economic Behavior*, 6, 1-31.
- Berg, J.E., Daley, L.A., Dickhaut, J.W. and O'Brien, J. (1992) Moral Hazard and Risk Sharing: Experimental Evidence. *Research in Experimental Economics*, 5, 1-34.
- Brandts, J., and Holt, C. (1992). An experimental test of equilibrium dominance in signaling games. *American Economic Review*, 82, 1350-1365.
- Brandts, J., and Charness, G. (2004). Do Labour Market Conditions Affect Gift Exchange? Some Experimental Evidence. *Economic Journal*, 114, 684-708.
- Cabrales, A., and Charness, G. (forthcoming). Optimal Contracts with Team Production and Hidden Information: An Experiment. *Journal of Economic Behavior and Organization*.
- Cawley, J., and Philipson, T. (1999). An Empirical Examination of Information Barriers to Trade in Insurance. *American Economic Review*, 89, 827-846.
- Charness, G., and Dufwenberg, M. (2006). Promises and Partnership. *Econometrica*, 74, 1579-1601.
- Charness, G., and Dufwenberg, M. (forthcoming). Contracts and Communication. *American Economic Review*.

- Charness, G., Kuhn, P. and Villeval, M.C. (forthcoming), Competition and the Ratchet Effect. *Journal of Labor Economics*.
- Charness, G., and Rabin, M. (2002). Understanding Social Preferences With Simple Tests, *The Quarterly Journal of Economics*, 117(3), 817-869,
- Chaudhuri, A. (1998). The ratchet principle in a principal agent game with unknown costs: An experimental analysis. *Journal of Economic Behavior and Organization*, 37, 291-304.
- Chiappori, P.A., and Salanié, B. (2000). Testing for Asymmetric Information in Insurance Markets. *Journal of Political Economy*, 108, 56-78.
- Chiappori, P.A., and Salanié, B. (2003). Testing Contract Theory: A Survey of Some Recent Work in Dewatripont, M., Hansen, L., and Turnovsky, S. (Eds), *Advances in Economics and Econometrics*, Vol 1, Cambridge: Cambridge University Press.
- Cooper, D., Kagel, J., Lo, W., and Gu, Q.L. (1999). Gaming against managers in incentive systems: Experimental results with Chinese students and Chinese managers. *American Economic Review*, 89, 781-804.
- Dahlby, B. (1983). Adverse Selection and Statistical Discrimination: An Analysis of Canadian Automobile Insurance. *Journal of Public Economics*, 20, 121-130.
- Davis, D., and Holt, C.A. (1994). Equilibrium Cooperation in Three-Person, Choice of Partner Games, *Games and Economic Behavior*, 7, 39-53.
- DeJong, D., Forsythe, R., Lundholm, R., and Uecker, W.C. (1985). A Laboratory Investigation of the Moral Hazard Problem in an Agency Relationship, *Journal of Accounting Research*, 23, 81-120.
- Dione, G., Doherty, N. (1994). Adverse Selection, Commitment and Renegotiation: Extension to and Evidence from Insurance Markets. *Journal of Political Economy*, 102 (2), 210-235.
- Fehr, E., Kirchler, E., Weichbold, A., and Gächter, S. (1998). When social norms overpower competition: Gift exchange in experimental labor markets. *Journal of Labor Economics*, 16, 324-351.
- Finkelstein, A., and Poterba, J. (2004). Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market. *Journal of Political Economy*, 112 (1), 183-208.
- Fischbacher, U., Fong, C.M., and Fehr, E. (2009). Fairness, Errors, and the Power of Competition. *Journal of Economic Behavior and Organization*, 72, 527-545.
- Freixas, X., and Rochet, J.C. (1997). *Microeconomics of Banking*. Cambridge: MIT Press,
- Genesove, D. (1993). Adverse Selection in the Wholesale Used Car Market. *Journal of Political Economy*, 104, 644-665.
- Greiner, B. (2004). An Online Recruitment System for Economic Experiments. In: K. Kremer, V. Macho (Eds.): *Forschung und wissenschaftliches Rechnen 2003*. GWDG Bericht 63, Göttingen: Ges. für Wiss. Datenverarbeitung, 79-93.
- Grosskopf, B. (2003). Reinforcement and Directional Learning in the Ultimatum Game with Responder Competition. *Experimental Economics*, 6, 141-158.

- Holt, C.A., and Sherman, R. (1990). Advertising and product quality in posted-offer experiments. *Economic Inquiry*, 28, 39-56.
- Kanemoto, Y., and MacLeod, B.W. (1992). The Ratchet Effect and the Market for Secondhand Workers. *Journal of Labor Economics*, 10(1), 85-98.
- Keser, C., and Willinger, M. (2000), Principals' principles when agents' actions are hidden. *International Journal of Industrial Organization*, 18, 163-185.
- Königstein, M. (2001). Optimal Contracting with Boundedly Rational Agents, *Homo oeconomicus*, XVIII (2): 211-228.
- Laffont, J.J., and Tirole, J. (1993). A Theory of Incentives in Procurement and Regulation. Cambridge: MIT Press.
- Lazear, E.P. (1999). Personnel Economics: Past Lessons and Future Directions. *Journal of Labor Economics*, 17(2), 199- 236.
- Lynch, M., Miller, R.M., Plott, C.R., and Porter, R. (1986). Product Quality, Consumer Information, and 'Lemons' in Experimental Markets, in Ippolito, M., and Scheffman, D.T. (Eds.), *Empirical Approaches to Consumer Protection Economics*. Washington, D.C. : Federal Trade Commission, Bureau of Economics, 251-306.
- Mas-Colell, A., Whinston, M.D., and Green, J. R. (1995). Microeconomic theory. Oxford: Oxford University Press.
- Milgrom, P., and Roberts, J. (1992). Economics, Organization and Management. Englewood Cliffs: Prentice Hall.
- Miller, R.M., Plott, C.R. (1985). Product Quality Signaling in Experimental Markets. *Econometrica*, 53(4), 837-872.
- Mirrlees, J.A. (1971). An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies*, 38(114), 175-208.
- Prendergast, C. (1999). The Provision of Incentives in Firms. *Journal of Economic Literature*, 37, 7-63.
- Puelz, R., and Snow, A. (1994). Evidence on Adverse Selection: Equilibrium Signalling and Cross-Subsidization in the Insurance Market. *Journal of Political Economy*, 102, 236-257.
- Ray, D. (1998). Development Economics. Princeton: Princeton University Press.
- Roth, A., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991). Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study. *American Economic Review*, 81(5), 1068-1095.
- von Siemens, F. (2010). Social Preferences, Sorting, and Competition. *Oxford Economic Papers*, forthcoming.
- Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16, 8-37.
- Young, P.H., and Burke, M.A. (2001). Competition and Custom in Economic Contracts: A Case Study of Illinois Agriculture. *American Economic Review*, 91, 559-573.

Zeiliger, Romain, (2000). *A Presentation of Regate, Internet Based Software for Experimental Economics*. <http://www.gate.cnrs.fr/~zeiliger/regate/RegateIntro.ppt>, GATE.

APPENDIX A - Instructions for the Excess Agent treatment (*the instructions for the other treatments are available upon request*)

You are about to participate in an experiment on decision-making carried out by researchers from the Universitat Pompeu Fabra, the University of California at Santa Barbara and GATE. During this session, you can earn money. The amount of your earnings depends on your decisions and on the decisions of the other participants in this session. During the session, your earnings will be calculated in points,

with 100 points = 1 Euro

During the session, losses are possible. However, they can be avoided with certainty by your decisions.

The session consists of 40 independent periods. Only 4 periods will be chosen at random for actual payment, at the end of the session. The earnings you have made during these 4 periods will be added up and converted into Euros. In addition, you will receive € 4 for participating in the experiment. Your earnings will be paid to you in cash in private to preserve confidentiality. Your decisions are anonymous and confidential.

During this session, there are two categories of participants: 4 participants are proposers and 12 participants are responders. The responders can be of two types: A or B.

The category to which the participant is assigned (proposer or responder) and the type of participant in the case the participant is a responder are chosen randomly at the beginning of the session. Each responder has an equal initial probability to be of either type A or type B. Half of all responders will be of each type.

You will be informed of your category and of your type if you are a responder at the beginning of the session and you will keep the same category and the same type throughout the session. If you are a responder, no one knows your type.

Description of each period

At the beginning of each period, each proposer is randomly matched with 3 responders. The responders may be either type, but the proposer does not know their types when making a proposal. The identity of your co-participants is unknown to you. The composition of the group changes randomly every period.

Each period consists of four stages.

- ❑ In the first stage, the proposer makes a selection from one of 6 possible “offers” {1,2,3,4,5 or 6} by checking a box on his screen.
- ❑ In the second stage, the three responders are informed of this offer. Each can then choose one either option X or option Y or “rejection” by checking the corresponding box on his screen.
- ❑ In the third stage, the proposer is informed of the choices of the three responders. If more than one responder has accepted the proposer’s offer, the proposer will select one of the responders among those who accepted his offer. He can accept at most one responder. The responders are not informed about the choices made by the other responders. The responders who have not been selected receive a payoff of 125 points.
- ❑ In the fourth stage, each person is informed of his own payoff in that period.

How are payoffs calculated?

The payoffs depend on the offer made by the proposer, on the responders’ decisions and on the choice made by the proposer among the responders. When a proposer chooses a responder, his payoff depends only on his offer and on the option chosen by this responder; the responders who have not been selected do not provide him with any additional payoff.

Please refer to the Table provided. This Table displays the 6 possible offers and their associated payoffs.

Corresponding to each offer, you can see 3 rows:

- The first row, in blue, indicates the payoffs of the proposer.
- The second row, in yellow, indicates the responder's payoffs if his type is A.
- The third row, in pink, indicates the responder's payoffs if his type is B.

The 3 columns represent the decisions made by the responder:

- The column (1) corresponds to the choice of option X if the responder accepts the offer
- The column (2) corresponds to the choice of option Y if the responder accepts the offer
- The column (3) corresponds to the case of the responder rejects the offer.

At the intersection of a row and a column, you can read the payoffs associated with an offer and a choice as a function of the role of proposer or responder.

Here are some examples.

Example 1. The proposer has chosen the offer 1. One responder of type B has accepted this offer and chosen option Y. The two other responders have rejected this offer. In this case, the proposer will receive 355 points; the responder who has accepted the offer will receive 135 points; the responders who have rejected the offer will receive 125 points.

Example 2. The proposer has chosen the offer 3. One responder of type A and one responder of type B have accepted this offer and chosen option Y; the other responder of type A has also accepted the offer and chosen option X. The proposer chooses the responder who chose option X. The proposer will receive 550 points; the responder who has been chosen will receive 210 points; the responders who have not been chosen will receive 125 points.

Example 3. If the proposer has chosen the offer 6 and if no responder has accepted his offer, both the proposer and the responders receive 125 points.

To sum up, in each period, if you are a proposer, you choose an offer from among the six feasible options and you choose between the responders who have accepted your offer; you cannot accept more than one responder. If you are a responder, you choose either option X or option Y or you reject the offer. Your payoffs for the current period are then computed.

At the end of a period, a new period starts automatically. Each period is independent.

If you have any question regarding these instructions, please raise your hand. Your questions will be immediately answered in private. Throughout the entire session, direct communication between participants is strictly forbidden.

APPENDIX B - Nonparametric tests

Our nonparametric tests are Wilcoxon-Mann-Whitney rank-sum tests, conducted with both session-level and individual-level data. In a strict sense each session is only one independent observation, since there is interaction between parties over the course of each session. Table A presents a summary of principal and agent choices in each of our sessions.

Table A: Session-level data

Variables	<i>B</i> treatment		<i>EA</i> treatment				<i>EP</i> treatment		
	S1	S2	S3	S4	S5	S6	S7	S8	S9
Average offer	3.41	3.50	2.53	2.51	2.07	3.04	5.08	5.11	5.05
Rejection rates (M 1, 2 & 3)	0.23	0.38	0.10	0.08	0.13	0.04	-	-	-
High agent-option X (M 1& 2)	0.00	0.00	0.52	0.50	0.82	0.56	-	-	-
High agent-option X (M 3)	0.75	0.26	0.87	0.96	1.00	0.96	-	-	-
% high type (actual contracts)	0.56	0.57	0.74	0.84	0.84	0.83	0.50	0.50	0.50

The average contract menu offered is lowest in the four *EA* sessions and highest in the three *EP* sessions. Rank-sum tests find $Z = 2.12$ and $p = 0.03$, comparing between *EA* and *EP*, $Z = 1.85$ and $p = 0.06$ for the comparison between *EA* and *B*, and $Z = 1.73$ and $p = 0.08$ for the comparison between *B* and *EP*. The likelihood that $EA < B < EP$ (for average menu offered) is only $p = 0.004$. Principals offer significantly different contract menus in each treatment.

We also see that rejection rates of the less generous menus (1-3) are substantially higher in both *B* sessions than in any of the four *EA* sessions, yielding $Z = 1.85$ and $p = 0.06$.¹⁹ In addition, high-ability agents in the *B* treatment are less likely to choose option X in response to menus 1 and 2 than are high-ability agents in the *EA* treatment (insufficient observations in the *EP* treatment); recall that the myopic profit-maximizing choice in the *EA* treatment is Y. As the rate is lower in both *B* sessions than in any of the four *EA* sessions; this gives $Z = 1.88$ and $p = 0.06$. The rate of option X being chosen by the high agent is lower in both *B* sessions than in any of the four *EA* sessions, with $Z = 1.85$ and $p = 0.06$. This is consistent with the fact that the offer of menu 3 is already more ‘generous’ than the equilibrium in the *EA* treatment. Finally, the proportion of high-ability agents in the actual contracts is higher in the *EA* sessions than in the *B* sessions ($Z = 1.85$ and $p = 0.06$) and it is smaller than a random draw of 0.875 ($p = 0.04$, t-test). This proportion is also higher in the *B* than in the *EP* sessions ($p = 0.08$) and it is larger than a random draw of 0.50 ($p = 0.01$, t-test).

Since we have only a few sessions in each treatment, we supplement these tests by collapsing the 40 choices of each participant to one number; while this approach ignores the interaction between parties, we feel it is nevertheless informative. These results confirm the patterns above, but with a higher degree of statistical significance. The Wilcoxon-Mann-Whitney tests find that there are significant differences in the average menu offered between each pair of treatments ($Z = 3.99$, 5.71, and 5.69 for *EA* vs. *B*, *EA* vs. *EP*, and *B* vs. *EP*, respectively; all of these test statistics give $p < 0.001$). The test also indicates that the rejection rates of menus 1-3 are significantly higher in the *B* treatment than in the *EA* treatment ($Z = 2.60$, $p = 0.01$). Finally, the test confirms that the proportion of high-ability agents in the actual contracts is larger in the *EA* than in the *B* treatment ($Z = 4.36$, $p < 0.001$) but this test fails when comparing the *B* and *EP* treatments ($Z = 1.56$, $p = 0.11$).

¹⁹ We choose menus 1-3 as there are few rejections of menus 4-6 and these all occurred in the *B* treatment.