

IZA DP No. 2296

Competition, Hidden Information and Efficiency: An Experiment

Antonio Cabrales
Gary Charness
Marie-Claire Villeval

September 2006

Competition, Hidden Information and Efficiency: An Experiment

Antonio Cabrales

Universidad Carlos III de Madrid

Gary Charness

University of California, Santa Barbara

Marie-Claire Villeval

CNRS-GATE and IZA Bonn

Discussion Paper No. 2296
September 2006

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

Email: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit company supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Competition, Hidden Information and Efficiency: An Experiment

We devise an experiment to explore the effect of different degrees of competition on optimal 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 these different degrees of competition affect outcomes and efficiency. Informational asymmetries generate inefficiency. In an environment where principals compete against each other to hire agents, these inefficiencies remain. 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. Competition leads to a substantially higher probability of trade, and that, overall, competition between agents generates the most efficient outcomes.

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

Keywords: experiment, hidden information, competition, efficiency

Corresponding author:

Marie-Claire Villeval
CNRS-GATE
93, Chemin des Mouilles
69130 Ecully
France
E-mail: villeval@gate.cnrs.fr

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 being able to observe the agent’s true type; in this case, it is typically necessary to devise ‘second-best’

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.

contracts that lead to separation of types, but which are somewhat distorted and less than fully efficient.

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. Prendergast (1999)'s and Chiappori and Salanié (2003)'s surveys show that the econometrics of contracts has recently become a burgeoning field of research. The latter point however that a number of empirical tests suffer from selection and endogeneity biases. In addition, many papers use similar data because of a lack of data on contracts. These difficulties explain that only few empirical tests of the adverse selection 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.

We perform experiments designed to test the influence of competition on the management of hidden-information problems. 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.

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

⁴ Although we use the standard static screening model, it is worth noting that Kanemoto and MacCleod (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.

Our approach is to systematically vary 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 include the induced contract menus for each treatment among the six feasible choices for the principal. We 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 degrees of competition in the market, affect outcomes and efficiency. In an environment where principals compete against each other to hire agents, inefficiencies remain. In contrast, 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 various competitive frameworks on the design of contracts in the presence of both heterogeneous agents and hidden information, and their subsequent efficiency. 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 either no competition or a competition among principals, our experiment finds that the institutional environment in which agents compete against each other the most efficient 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 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.⁵ 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. Our experimental design and implementation are presented in section 4, with the results given in section 5. We discuss welfare and efficiency considerations in section 6, and conclude in section 7.

⁵ As we discuss later, this parallels the Cooper and Kagel (2005) signaling result in an entry game.

2. RELATED LITERATURE

Previous experimental studies on asymmetric information have typically examined the behavioral issues present with individual contracting with hidden action (also known as moral hazard), where effort is not contractible; these studies include Berg, Daley, Dickhaut, and O'Brien (1992), Keser and Willinger (2000), Anderhub, Gächter, and Königstein (2002), and Königstein (2001).⁶ They observe that contracts are usually more generous than theoretically predicted and some suggest adding an equity constraint to the standard participation and incentive-compatibility constraints in the design of optimal contracts. Charness and Dufwenberg (forthcoming) consider the hidden-action problem and find that cheap-talk statements of intent (*promises*) help to achieve desirable outcomes (the Nash bargaining solution).

There is very little experimental work on hidden information and certainly no study considers the effect of competition. Lynch, Miller, Plott and Porter (1986) confirmed the existence of a market for “lemons” in experimental oral double auctions, and Holt and Sherman (1990) do so in posted–offer auctions. Experiments by Brandts and Holt (1992) and Banks, Camerer and Porter (1994), however, provide mixed evidence about the ability of subtle equilibrium refinements to predict players’ behavior in simple signaling games. DeJong, Forsythe, and Lundholm (1985) observe that the low quality/low price outcome is less frequent when the buyer cannot distinguish between a rip-off and bad luck. Cabrales and Charness (2004) study the static principal-agent problem with hidden information in the context of team

⁶ Other studies involving moral hazard include Bull, Schotter, and Weigelt (1987), who examine the incentive effects of piece rate and tournament payment schemes, and Nalbantian and Schotter (1997), who investigate group incentive contracts. Plott and Wilde (1982), and DeJong, Forsythe, Lundholm, and Uecker (1985) consider moral hazard problems with multiple buyers and sellers. Güth, Klose, Königstein, and Schwalbach (1998) consider a dynamic moral hazard problem where trust and reciprocity issues impede obtaining the first-best outcome.

production. They observe that when more equitable menus are proposed, rejection rates are less important and agents select actions according to their types. There are two studies of the dynamic contracting problem: Chaudhuri (1998) and Cooper, Kagel, Lo, and Gu (1999) study the problem of the *ratchet effect*, where the agent has an incentive to conceal her true type, as the principal may use this information to ratchet up the demands for performance in later periods. Nevertheless, principal-agent interactions in the field are frequently one-shot affairs; furthermore, if the principal could commit to an *ex ante* contract, it would be optimal to implement the one-shot problem in the dynamic setting.⁷

One might predict that different forms of competition should lead to different contract menus being selected, given the differences in the players' bargaining power. However, results from the handful of experimental papers on the effects of unbalanced competition on the outcomes between firms and workers (or principals and agents) is somewhat mixed. Brandts and Charness (2004) find little difference in the gift-exchange outcomes according to whether there are more workers than firms or *vice versa*. Fehr, Kirchler, Weichbold, and Gächter (1998) show that an unbalanced market eliminates fairness when contracts are complete but not when they are incomplete. On the other hand, Roth *et al.* (1991) find that principals capture nearly the entire surplus when 10 agents compete in a "demand game" similar to the ultimatum game. Davis and Holt (1994) show that the ability of a buyer to switch between two sellers provides a strong incentive to develop reputation in a repeated game. In an ultimatum game with responders' competition, Grosskopf (2003) observes that although the initial demands are similar, demands in the game with competition become higher over time than in the game with no competition. Finally, Fischbacher, Fong and Fehr (2003) demonstrate that the introduction of

⁷ In addition, even though a relationship may actually involve repeated play, a firm could choose to pre-commit to a

even a little competition provokes large behavioral changes. A model combining heterogeneous social preferences with decision errors enables to predict most of the experimental evidence with various degrees of competition. Thus, it is not clear *ex ante* what effects unbalanced competition will have on the hidden-information problem, particularly in terms of economic efficiency.

3. THE MODEL

In this section we describe the theoretical model that serves as the basis for the experimental design. As a preview, we note that the case with competition between principals (more principals than agents) yields a Rothschild-Stiglitz type of solution, which is invariant to the number of agents and generally inefficient. 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 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-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 type $j \in \{H,L\}$, which is her private information:

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

contract, and perhaps cultivate a reputation for integrity by doing so.

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

From the utility functions of the principal and the agents we have that the first-best effort levels are:

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

We call \hat{e}_j the efficient level of effort.⁸ If we denote by \underline{U} 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 type is denoted respectively by p_H or p_L , then the expected (optimal) profits for the principal are given by:

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

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 (one principal-one agent) treatment results from the solution of the maximization program:

⁸ 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.

$$\max_{w_H, w_L, e_H, e_L} p_H(e_H - w_H) + p_L(e_L - w_L)$$

subject to

$$w_H - \frac{k_H}{2}(e_H)^2 \geq \underline{U} \quad (\text{IR}_H)$$

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

$$w_H - \frac{k_H}{2}(e_H)^2 \geq w_L - \frac{k_H}{2}(e_L)^2 \quad (\text{IC}_H)$$

$$w_L - \frac{k_L}{2}(e_L)^2 \geq w_H - \frac{k_L}{2}(e_H)^2 \quad (\text{IC}_L)$$

where (IR_j) and (IC_{ji}) are respectively the individual rationality and incentive compatibility constraints of an agent of type $j \in \{H, L\}$. As usual in these problems, it turns out that the active constraints in the optimal solution are (IR_L) and (IC_H) , so that the solution is:

$$e_H^B = \frac{1}{k_H} = 1; \quad e_L^B = \frac{p_L}{k_L - k_H(1 - p_L)} = \frac{1}{k_L + \frac{1 - p_L}{p_L}(k_L - k_H)}; \quad (2)$$

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

The high 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 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 3 agents (this is the Excess Agent treatment). Then an equilibrium contract menu results from the solution of a slightly different maximization program. Given that high 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 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.⁹

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

subject to

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

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

The (IC_H) can also be written

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

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

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

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

The effort of the low type of agent in the EA treatment e_L^{EA} is closer to the efficient effort (\hat{e}_L) than that in the Benchmark treatment (e_L^B). To see this note first that both e_L^{EA} and e_L^B are smaller than

$\hat{e}_L = \frac{1}{k_L}$, but we will now show $e_L^{EA} > e_L^B$, so that the distortion is lower in EA than in B:

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

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

The reason for this enhanced efficiency is that the principal distorts the low agent in order to lower the rents to the high 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 types, the presence of other high types that makes this less necessary. In fact, it is easy to check that in the general model where the principal confronts n agents, the difference

⁹ We only write the binding constraints, in what follows.

between the equilibrium and the efficient level of effort for the low 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 types assume that other high 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 all agents, both high and low, 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 type (thus she will do it independently of what other individuals of her type are doing). 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 type to choose the high option, and that is exactly what the IC_H constraint in equation (4) does.

Finally, we also have a treatment where several principals compete for one agent. In that case, the equilibrium of the game is such that the principals make zero profits for each type of contract, the low agent gets an undistorted contract and the high agent is held to his Incentive Compatibility constraint (see e.g. Mas-Colell, Green and Whinston, 1995, ch.14D).

$$\begin{aligned} e_H^{EP} &= w_H^{EP}; e_L^{EP} = w_L^{EP}; \\ e_L^{EP} &= \frac{1}{k_L}; e_H^{EP} = \frac{2}{k_H} - \frac{1}{k_L} \end{aligned} \quad (5)$$

where e_H^{EP} comes from the Incentive constraint for the high type which solves

$$w_H^{EP} - \frac{1}{2}k_H \left(e_H^{EP}\right)^2 = w_L^{EP} - \frac{1}{2}k_H \left(e_L^{EP}\right)^2 \text{ which from (5) can be rewritten as:}$$

$$e_H^{EP} - \frac{1}{2}k_H \left(e_H^{EP}\right)^2 = \frac{1}{k_L} - \frac{1}{2}k_H \left(\frac{1}{k_L}\right)^2 \Leftrightarrow \left(e_H^{EP}\right)^2 - \frac{2}{k_H}e_H^{EP} + \frac{2}{k_L k_H} - \left(\frac{1}{k_L}\right)^2 = 0,$$

$$\text{which leads to } e_H^{EP} = \frac{\frac{2}{k_H} + \sqrt{\frac{4}{k_H^2} - \frac{8}{k_H k_L} + \frac{4}{k_L^2}}}{2} = \frac{1}{k_H} + \sqrt{\left(\frac{1}{k_H} - \frac{1}{k_L}\right)^2} = \frac{2}{k_H} - \frac{1}{k_L}.$$

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) induces menu 6. 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 type, and

respects the IC constraint for the high type. Menu 5 is fully efficient, for both the high and the low types.

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

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. 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. As Wilson (1993) points out (p. 146) in a representative example: “The firm’s profits

from the five-part and two-part tariffs are 98.8% and 88.9% of the profits from the nonlinear tariff.”

4. EXPERIMENTAL DESIGN

We conducted three different treatments, which differed according to the numbers of principals and agents in the treatment. In our *Benchmark (B)* treatment, there were 10 principals and 10 agents in each session. In the *Excess Agents (EA)* treatment, there were four principals and 12 agents, while in the *Excess Principals (EP)* treatment, there were 12 principals and six agents.¹⁰ In all cases, there were equal numbers of high (H) agents and low (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.

Table 2 – Treatments and sessions

	Participants per session			
--	--------------------------	--	--	--

¹⁰ We chose 3 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 because the theoretical predictions are not sensitive to the size of the competing pool.

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

The sessions were conducted at the Groupe d'Analyse et de Théorie Economique (GATE), CNRS, France. The subjects were recruited from undergraduate courses in local Engineering and Business schools. Some of the subjects had participated in previous experiments, but all of the subjects were inexperienced in this particular type of experiment. No subject participated in more than one session of the study. On average, a session lasted 60 minutes, including initial instructions and payment of subjects. The experiment was computerized using the REGATE program developed at GATE (Zeiliger, 2000).

The participants were privately informed of their role; 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 low type and high type agents in the room. 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 the two principals 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 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.¹¹ Unmatched principals or agents received 125 points in the period. This process leads to Table 3 that was distributed to the subjects to help them to make their decisions (except that the term “menu” was replaced by that of “offer”).¹²

We used a conversion rate of 100 points for each Euro. At the end of each session, we selected (at random) four of the 40 periods for actual payment. In this way, we avoided possible income effects from having already accumulated a known amount of financial remuneration in the session. The average payoff was 14.9 Euros in the Benchmark treatment, and 13.5 Euros in both the Excess Agent and the Excess Principal treatments; on average the principals received 17 Euros, the high agents 13 Euros and the low agents 10 Euros, including a 4 Euro show-up fee.

¹¹ 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

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

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

principal competition enhances efficiency as it reduces the envy-driven rejections that hurt efficiency in the benchmark.

¹² See also the instructions in the Appendix A.

for principal behavior and agent behavior, supplemented with charts. We then consider the determinants of such behavior, providing statistical tests and regression analysis.

4.1 Descriptive statistics

Principal behavior

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

Table 4 – Menus offered, by treatment

Menu Offer	<i>B</i> treatment	<i>EA</i> treatment	<i>EP</i> treatment
1	89 (11.12)	181 (28.28)	56 (3.89)
2	17 (2.12)	34 (5.31)	28 (1.94)
3	242 (30.25)	348 (54.38)	57 (4.96)
4	361 (45.13)	58 (9.06)	136 (9.44)
5	75 (9.38)	16 (2.50)	490 (34.03)
6	16 (2.00)	3 (0.47)	673 (46.74)
Total	800 (100)	640 (100)	1440 (100)

Note: Equilibrium menus are in **bold**. Percentages are in parentheses.

The *EA* treatment has the ‘lowest’ contract menus (the ones most favorable to the principal), the *EP* treatment has the ‘highest’ menus (the ones most favorable to the agents), and the *B* treatment features intermediate menus. There is some support for the equilibrium predictions, as 30% of the menus offered in the *B* treatment, 34% of the menus offered in the *EA* treatment, and 47% 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, menu 4 is the most common choice in the *B* treatment and menu 3 is the most common choice in the *EA* treatment.

As menu 3 has a more egalitarian distribution than menu 1 (*EA* treatment) 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; a similar comment applies vis-à-vis menu 4 and menu 3 (*B* treatment). However, to the extent that these represent social preferences on the part of the principals, we should expect to see these choices from the beginning. Instead, the menus offered evolves over time, as can be seen in Figures 1-3.

Figure 1: Menu Offered over Time, Benchmark




Figure 2: Menu Offered over Time, EA





Figure 3: Menu Offered over Time, EP



In each of these treatments, there is some distribution across menus in the initial periods, but we typically see that two menus comprise the great majority of the menus offered thereafter in each treatment. In every case, it is the ‘higher’ menu that grows in frequency over time. Thus, it seems that some other force is inducing principals to choose more generous menus over time, even though principals’ social preferences may perhaps account for some of the generous menu choices. In every treatment, one menu predominates in the last group of periods, averaging about 60% of all menus offered (although there is a dip in favor of the equilibrium menu in *EA*).

While we postpone an in-depth analysis of the determinants of the contract menus offered by principals, Table 5 below shows the ex post profitability of these offers in our treatments;

It is clear that the profitability of a particular contract menu depends greatly on the competitive environment. Menus 1 and 2 yield highest profits for the principal in the *EA* treatment, but generate low profits in the *B* and *EP* treatments. Similarly, menu 6 is quite unattractive for the principal in the *EA* and *B* treatments, but provides nearly the best profits in the *EP* treatment (and the best in the final periods). Overall, we observe a good correspondence

between the most frequent offers made and their profitability; thus, to a large extent, it seems that principals are influenced by considerations of their own profits.

Table 5 – *Ex post* principal payoffs, by menu and treatment

Menu Offer	<i>B</i> treatment	<i>EA</i> treatment	<i>EP</i> treatment
1	270 (194)	516 (527)	146 (125)
2	210 (215)	528 (527)	125 (125)
3	334 (318)	520 (519)	147 (125)
4	392 (384)	478 (500)	180 (170)
5	346 (353)	372 (---)	219 (185)
6	250 (250)	250 (250)	214 (205)

Note: The most frequent menus are in **bold**. Payoffs for the last eight periods are in parentheses.

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. More favorable menus are more likely to be accepted in the *B* treatment, with 142 rejections of the 800 contract menus offered. The acceptance rate is 90% for menu 4, 75% for menu 3, but only 61% for menu 1. In addition, less favorable menus are not likely to be selected by agents in the *EP* treatment although rejections of offers from both principals occurred only twice. The punishment for offering an unfavorable contract menu is simply that no agent will select it. Menus 1 to 3 are only accepted about 5% of the time; in fact, these menus were never accepted in the final eight periods, and menu 4 was only accepted twice out of the 144 times it was offered in these periods. Finally, rejections in the *EA* treatment are not actually costly for the principal, since an offer has always been accepted by at least one agent.

We summarize agent behavior by listing the number of accepted menus (column 1), the rate of acceptance (column 2), and the proportion of X option chosen (column 3) in each treatment in Table 6 for the high agents and Table 7 for the low agents.

Table 6 – High agents' choices, by treatment

Menu Offer	<i>B</i> treatment			<i>EA</i> treatment			<i>EP</i> treatment		
	Nb	%accept	% X	Nb	%accept	% X	Nb	%accept	% X
1	36	80.00	0	271	94.76	63.10	1	5.00	0
2	6	60.00	0	46	93.88	71.74	0	0.00	-
3	97	94.17	57.73	513	98.28	94.54	2	8.70	100
4	191	95.50	79.06	82	100	95.12	10	16.95	80.00
5	37	100	86.49	16	100	93.75	102	39.23	98.04
6	5	100	60.00	5	100	80.00	245	70.61	68.16
Total	372	93.00	65.06	933	97.19	84.24	360	50.00	76.94

Table 7 – Low agents' choices, by treatment

Menu Offer	<i>B</i> treatment			<i>EA</i> treatment			<i>EP</i> treatment		
	Nb	%accept	% X	Nb	%accept	% X	Nb	%accept	% X
1	18	40.91	11.11	212	82.49	1.42	4	11.11	0
2	2	28.57	0	47	88.68	0	0	0.00	0
3	85	61.15	2.35	446	85.44	0	2	5.88	0
4	134	83.23	2.98	92	100	0	18	23.38	0
5	36	94.74	5.56	32	100	0	99	43.04	0
6	11	100	9.09	4	100	0	235	72.09	0
Total	286	71.50	3.85	833	86.77	0	358	49.72	0

Note: Equilibrium menus are in **bold**.


We do see that the rejection rate for low agents is always much higher than the rate for high agents, more than four times as high in both treatments *B* and *EA*. When they accept an offer, low agents rarely chose option X, which would generate negative earnings. The behavior

of high agents is more complex. Note that option X pays more than option Y for high 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 type with menu 6. In the *EP* treatment, high 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, recall that agents' choices are known to the principals prior to selecting an agent, so that agents must compete in their choices to be selected. High 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 his principal will not select him by if another agent has chosen option X. Since some agents appear to realize this, we observe that 64% of the high agents choose option X when accepting menus 1 or 2. Figure 4 shows that high agents appear to be learning to do so over time, as this likelihood increases during the first half of the sessions.¹³ We also chart menu 3 for comparison, to show a 'best-case scenario' for the X option, since a high agent trying to compete should always choose it; note the lack of a time trend.

¹³ An indirect consequence of this pattern is that principals in the *EA* treatment are increasingly successful in obtaining option X in response to menu 1 or menu 2. Overall, option X was obtained 67% of the time in response to these menus; this proportion was 63% during the first 20 periods and 72% during the last 20 periods.

**Figure 4: Likelihood of X Option over Time in EA,
High agents**



As a consequence of these decisions, the proportion of high 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 agents, in *EA* this reflects the process of selection related to the competitive environment. In *EP*, the proportion of high agents (50.14%) corresponds roughly to the initial distribution of the population since almost all the agents accept an offer. Regarding the evolution over time of the proportion of high agents, the selection process does not take a long time before being fully operative in *EA*, whereas there is almost no evolution in *B*. The percentage of high type among the agents recruited in *EA* rises from 71% in periods 1 to 8 up to 87% in periods 9 to 16 and stays above 81% in the following periods, whereas the proportion stays between 55 and 57% in *B* throughout the game.

4.2 Regression analysis

It appears that there are substantial and significant differences in both principal and agent behavior across treatments. We now turn to multiple-regression analysis of the determinants of the observed behavior, first considering the menus offered by principals.¹⁴ In order to

¹⁴ 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

disentangle the motivations underlying offers of the different categories of menus, we estimate multinomial logit models with robust standard errors, in which the reference category is the offer of the equilibrium menu. To better understand the choice of the most frequent menu that is not necessarily the equilibrium one, we also estimate ordered probit models with robust standard errors in which the switch to the most frequent menu offer is the dependent variable, equal to +1 if the menu increased, 0 if unchanged, and -1 if it decreased. In all the regressions, we include a time trend to identify a possible evolution over time. The first three columns of Table 8 presents the results of the multinomial logit model and the fourth column the results of the ordered probit model for the contract menus chosen in the *B* treatment.

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 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 agents is higher in *EA* than in *B*.

Table 8: Determinants of the choice of menus in the Benchmark Treatment

Ref.: offer of the equilibrium menu	Multinomial Logit model			Ordered Probit model
	Offer of menus 1 - 2	Offer of menu 4	Offer of menus 5 - 6	Switch to menu 4
Time trend	-0.051*** (0.015)	0.057*** (0.009)	0.008 (0.015)	-0.018** (0.007)
Lagged rejection rate	0.043*** (0.013)	-0.053*** (0.011)	0.014 (0.013)	0.004 (0.011)
X option chosen in (t-1)	-1.366*** (0.383)	0.176 (0.197)	-0.355 (0.301)	-0.321** (0.139)
Constant	-0.376 (0.340)	-0.053 (0.227)	-1.183*** (0.333)	
Nb observations		641		316
Log Likelihood		-701.339		-219.847
Wald χ^2		87.55		9.960
Prob> χ^2		0.000		0.028
Pseudo R ²		0.087		0.028

Note: Robust standard errors in parentheses; ***, and ** denote two-tailed statistical significance at the 1%, 5%, and 10% level, respectively.

Menus 1 and 2 are chosen less frequently over the course of the sessions, while menu 4 is increasingly preferred to menu 3 in the latter periods of the sessions.¹⁵ The offer of very generous menus (11% of offers) does not follow a clear trend.

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 the motivations driving 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. Moreover, an agent choosing the X option in the previous period (which never occurs in reaction to menus 1-2) strongly decreases the likelihood of menu 1-2 being chosen by the principal. This makes sense because offering these menus would induce both types to switch

to the low option. The ordered probit regression regarding switching from any menu to the most frequent menu 4 indicates that such a switch 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. It shows no significance for whether the principal has experienced a high proportion of rejections in the past. In contrast, a similar regression in which the lagged option is replaced by the rejection of the menu offer in the previous period (not reported here) indicates that a rejection strongly favors such a switch. Overall, these regressions suggest that offering a more generous menu than the equilibrium is driven by the experience of recent rejections and the selection of low-type agents, whereas offering very generous menus (5 or 6) does not depend on any of these variables, suggesting a preference for more egalitarian outcomes.

Table 9 presents the results for the menus chosen in the EA treatment. In the first two columns, we pool together the offers of menus 4 to 6 since they represent only 12% of the observations. The third column displays the results of the ordered probit model in which the switch to the most frequent menu 3 is explained. In contrast with the previous regression, we do not include the lagged rejection rate as an independent variable since all the contracts are accepted.

¹⁵ The same regression in which we replaced the time trend by period range dummies indicate that, regarding menus

Table 9: Determinants of the choice of menus in the Excess Agent Treatment

	Multinomial Logit model		Ordered Probit model
Ref.: offer of the equilibrium menu	Offer of menu 3	Offer of menus 4 - 6	Switch to menu 3
Time trend	0.019* (0.010)	-0.021* (0.013)	0.004 (0.007)
X option chosen in (t-1)	1.673*** (0.256)	0.345 (0.329)	-0.758*** (0.236)
X option * menu 1-2 in (t-1)	-3.679*** (0.305)	-2.228*** (0.403)	
Constant	-0.235 (0.281)	-0.121 (0.305)	
Nb observations	624		345
Log Likelihood	-466.574		-194.053
Wald χ^2	192.19		10.39
Prob> χ^2	0.000		0.006
Pseudo R ²	0.212		0.041

Note: Robust standard errors in parentheses; *** and * denote two-tailed statistical significance at the 1% and 10% level, respectively.

Menu 3 is chosen more frequently over time (we know from an alternative regression that this is particularly the case in periods 17-32, significant at 1%), whereas the frequency of menus 4-6 decreases (especially in the final eight periods with a coefficient significant at the 5% level). 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; this makes sense, since the X option leads to a higher payoff for the principal, and so there is less motivation to offer a more generous menu. On the other hand, if the X option was chosen in the previous period (in reaction to whichever offers), this increases the likelihood of menu 3 being chosen but this has no significant effect on the offer of menu 4-6. If we separate out the case when menu 3 is chosen (column 3), we see that switching to menu 3 is much *less* likely when the X option has been chosen in the previous period.

1, 2 and 4, all period ranges are significant at the 1% level except periods 9-16, significant at the 10% level.

Table 10 presents the results for the contract menus chosen in the Excess Principal treatment. In the first two columns, we pool together menus 1 to 4 since they only represent 19% of the observations. In contrast with the previous treatments, we include among the independent variables both the lagged acceptance of the principal's offer and the lagged rejection rate 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. Column 3 displays the results of a random-effects probit model in which the switch to both the equilibrium and most frequent menu offer is the explained variable.¹⁶

Table 10: Determinants of the choice of menus in the Excess Principal Treatment

Ref.: offer of the equilibrium menu	Multinomial Logit model with robust S.E.		Random-effect Probit model
	Offer of menus 1-4	Offer of menu 5	Switch to menu 6
Time trend	-0.083*** (0.008)	-0.062*** (0.006)	-0.068*** (0.009)
Acceptance of the offer in (t-1)	-0.166 (0.167)	0.111 (0.136)	-2.266*** (0.204)
Lagged rejection rate	0.033*** (0.005)	0.012*** (0.004)	0.001 (0.009)
Constant	-0.929*** (0.351)	0.392 (0.243)	2.185*** (0.649)
Nb observations	1404		671
Log Likelihood	-1312.510		-221.422
Wald χ^2	215.390		147.600
Prob> χ^2	0.000		0.000
Pseudo R ²	0.090		

Note: *** denote two-tailed significance at the 1% level.

There is a strong time trend in favor of menu 6 being chosen, while the offer of less egalitarian menus decreases more strongly over time (in the detailed regression, not reported

¹⁶ Indeed, an ordered probit model would not be appropriate because data are censored at menu 6.

here, all the period range coefficients are statistically significant at the 1% level and decrease 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 probit regression regarding switching to menu 6 shows that the rejection of an offer in the previous period increases the likelihood of such a switch, whereas the overall lagged rejection rate exerts no specific impact. 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.

Another approach is to focus directly on whether a principal *changes* the contract menu from one period to the next.¹⁷ We might expect that a principal who is primarily interested in maximizing his own payoff would change the contract menu offered on the basis of expected profit, so that the payoff received from the previous menu offered would be critical. We consider this factor, along with a time trend in the following ordered probit regressions with robust standard errors in which data from all treatments are pooled. Menu change is the dependent variable, equal to +1 if the menu increased, 0 if unchanged, and -1 if it decreased. The results of these regressions are reported in Table 11:

¹⁷ A possible limitation is that the menu change is censored if the prior menu was 1 or 6 in these regressions.

Table 11: Contract menu changes and lagged payoffs

Independent variable	(1)	(2)	(3)
Lagged principal's payoff	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)
<i>B</i> * lagged principal's payoff	-	-	-0.001*** (0.000)
<i>EP</i> * lagged principal's payoff	-	-	-0.002*** (0.000)
Period	-	-0.005*** (0.001)	-0.006*** (0.001)
Nb observations	2808	2808	2808
Log pseudo-likelihood	-2671.8	-2668.8	-2636.9
Wald χ^2	26.270	32.600	60.800
Prob> χ^2	0.000	0.000	0.000
Pseudo R ²	0.006	0.007	0.019

Note: Robust standard errors in parentheses; *** denote two-tailed statistical significance at the 1% level.

In all cases, the lagged payoff is a very significant determinant of whether the principal left the contract menu unchanged or chose to increase or decrease it. Menus increased (decreased) when the previous payoff was lower, as with rejection or the Y option being chosen. When we consider lagged payoffs separately for each of our three treatments, we find that the coefficient is significantly negative in the *EA* treatment, and is even more negative (and significant) in the *B* and *EP* treatments. The effect of lagged payoffs is strongest in the *EP* treatment, as not being chosen generates a low payoff and so is to be avoided. We also observe a clear negative time trend in all specifications. An additional regression (not reported here) indicates little difference in this trend across treatments.

We next analyze the agents' decisions of whether to accept a contract menu and (for high agents) which option to choose if accepting a contract. Table 12 shows random-effects probit

regressions for the three treatments. Note that regarding the Excess Principal 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.

Table 12: Determinants of the decision to accept a contract offer (Random-effects 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 (0.004)	-0.012* (0.007)	0.007 (0.006)	-0.025*** (0.004)
Excess Agent Treatment	1.245*** (0.361)			
High-type agent	1.120*** (0.292)	1.399*** (0.549)	1.189*** (0.353)	-0.110 (0.072)
Offer	0.489*** (0.048)	0.759*** (0.080)	0.255*** (0.064)	0.734*** (0.049)
Lagged offer	-0.073* (0.045)	-0.130* (0.070)	-0.018 (0.063)	
Lagged % of no selection			0.008** (0.004)	
Constant	-0.263 (0.411)	-0.772 (0.543)	0.545 (0.382)	-3.255*** (0.242)
Nb observations	2652	780	1832	1440
Log Likelihood	-496.467	-204.915	-262.181	-822.616
Wald χ^2	118.650	96.830	29.110	219.910
Prob> χ^2	0.000	0.000	0.000	0.000

Note: ***, **, and * denote two-tailed significance at the 1%, 5%, and 10% level, respectively.

In all three treatments, the table confirms that higher offers (indexed from 1 to 6) are more likely to be accepted. We also see a large difference in acceptance rates, as high agents in both the *B* and *EA* treatments are significantly more likely to accept an offered contract than are low agents in these treatments. This cannot influence offers' acceptance in the *EP* treatment,

since each type of agent is able to accept the best offer. A couple of more ‘psychological’ factors come into play as well. In the *B* treatment, agents are influenced by the previous contract offered to them (although this effect is marginally significant); holding the current offer constant, the better the previous offer, the less likely the agent will accept a contract offer. This effect nearly vanishes in the *EA* treatment, where we instead observe that an agent who was hired less often in the prior periods is significantly more likely to accept the current contract offer. When we pool the *EA* and *B* data, we find that a contract is much more likely to be accepted in the *EA* treatment, illustrating the effect of competitive pressure.

We then explore the determinants of the option chosen by an agent who accepts a contract. Since low agents incur serious losses if they choose the X option, we focus only on high agents. Since the payoff for the high agent is the same for menu 6, regardless of the option chosen, we also omit this menu (5 observations in *B* and in *EA*) and the *EP* treatment, where it prevails, from the analysis.

Table 13: Determinants of the choice of option X by the high-type agents who accepted an offer (Random-effects Probit models)

Choice of option X	Benchmark Treatment	Excess Agent Treatment
Time trend	-0.020* (0.012)	0.017*** (0.006)
Accepted contract	2.342*** (0.232)	0.878*** (0.083)
Lagged selection		0.304** (0.138)
Constant	-7.241*** (0.919)	-1.045*** (0.298)
Nb observations	367	886
Log Likelihood	-58.564	-244.578
Wald χ^2	101.84	119.98
Prob> χ^2	0.000	0.000

Note: ***, **, and * denote two-tailed significance at the 1%, 5%, and 10% level, respectively.

We see that the higher the contract, the more likely the high agent is to choose option X. Option X is chosen less frequently over time in the *B* treatment, as it seems that high 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 contrast, option X is chosen more and more frequently over time in the *EA* treatment since the high 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. Finally, we see a significant effect in the *EA* treatment of not having been hired in the previous period, suggesting that the agents try to catch up by choosing the X option.

6. WELFARE AND EFFICIENCY COMPARISONS

Welfare comparisons across treatments are complicated for at least two reasons. One is that 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 importantly, the numbers of matches with the high agents is likely to be higher in the *EA* treatment, which necessarily gives a boost to the total payoffs in this treatment. In order to deal with this, we supply separate comparisons matches involving high and low agents. Finally, the kinds of theoretical distortions at the equilibrium are different in the *EP* treatment (where the high agents' contract forces them to work more than the efficient level) than in either the *EA* or the *B* treatments (where the low types work less than the efficient level).

Having said this, the equilibrium prediction (given by equation (3)) of the *EA* treatment seems unambiguously better when compared with the *B* treatment (given by equation (1)).¹⁸ There is a higher proportion of the matches with the more productive high agents in *EA* than in *B*. In addition, in the equilibrium of the *EA* treatment there is less distortion for the high agent than in the equilibrium of *B*; this can be seen in both equation (3') and Table 1.

The comparison between the *B* and the *EP* treatment is theoretically more complicated, since the distortion occurs for different individuals, the high agents in the *EP* treatment and the low agents in the *B* treatment. But since we expect an equal number of high and low agents to be matched in both cases, we can compare the relative value of the distortions of high and low types in both treatments. The loss in welfare with respect to first best (in equilibrium) for the low

agents' contract in *B* is $\hat{e}_L - \frac{k_L}{2}(\hat{e}_L)^2 - \left(e_L^B - \frac{k_L}{2}(e_L^B)^2 \right)$ and the loss for the high agents'

contract under *EP* is $\hat{e}_H - \frac{k_H}{2}(\hat{e}_H)^2 - \left(e_H^{EP} - \frac{k_H}{2}(e_H^{EP})^2 \right)$. With our parameterization, the

loss in *B* is 0.03 whereas the loss in *EP* is 0.12. Thus, in our experimental framework, the equilibrium in *EP* would lead to lower overall welfare than the equilibrium in *B*, which in turn would lead to lower overall welfare than the equilibrium in *EA*.

With the contract menus and payoff parameters we chose, we find that the greatest benefits accrue to society from a match when there are many agents competing to be hired. If we consider the theoretical predictions regarding efficiency, we should observe menu 1 or 2 in *EA*, menu 3 in *B*, and menu 6 in *EP*. While only half of the matches in *B* and *EP* would involve high agents, 87.50% (on average) of the matches in *EA* should be with high agents. So in *EA* (with

¹⁸ As we noted previously, the *EA* treatment has other equilibria, and we selected the contracts of one of those equilibria for one of the available menus in the experiment. The equilibrium given by (3) is the one, however, that

menu 2, the worst case for efficiency), the average total payoff should be 720, in *B* the average total payoff should be 620, and in *EP* the average total payoff should be 572.50. *EA* gives the most efficiency; the $EA > B > EP$ ordering also matches the pattern for accepted contracts.

The data show a qualitatively similar pattern. The overall average payoffs per treatment, which are a direct measure of efficiency since they take into account both the type and the effort of the recruited agents, are ordered in the same manner as that predicted theoretically. These welfare considerations/average payoffs are shown in Table 14:

Table 14: Welfare comparisons (average total payoffs) across treatments

Variable	<i>B</i> treatment	<i>EA</i> treatment	<i>EP</i> treatment
Average total payoff	538.46	698.97	590.42
Total payoff if contract accepted	600.71	698.97	591.36
Total payoff if H accepts contract	687.89	749.33	672.67
Total payoff if L accepts contract	487.31	480.75	509.61

We see that the *EA* treatment easily yields the highest average total payoff (line 1). 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 2 agents reject the offer of both principals, while the *B* treatment has a substantial rejection rate. Nevertheless, if we ‘level the playing field’ by considering only accepted contracts, the *EA* treatment still generates the highest degree of efficiency. 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 for one of the two theoretical reasons: the possibility of selection means that there are more matches with high

seems more relevant for the data.

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 agents in *EA*, it would still be the treatment with the highest total payoffs, either with or without rejections (615.04). Note, however, that our theory over-predicts payoffs in the presence of competition among agents and under-predicts them in the presence of competition between principals. Indeed, the total payoff in *EA* is slightly under the predicted 720 ($p = 0.040$, according to a t-test), whereas the total payoff in *B* is not different from the predicted 620 ($p = 0.420$) and it is even higher in the *EP* treatment than the predicted 572.50 ($p = 0.004$).

If we only consider welfare with low agents, the best treatment is the *EP* treatment. Perhaps this is because high and low agents each net the principal the same profits, regardless of the option chosen, and so are equivalent in some sense. However, option X is so relatively inefficient in menu 6 (total payoffs of 635, compared to 760 with all other menus) that the *EP* treatment cannot be the best overall.

One immediately wonders just how robust the *EA*-efficiency result is to alternative specifications. Excess agents are necessarily better than the benchmark, since when the number of excess agents goes to infinity, there is no distortion for even low agents, so we obtain first-best efficiency (proof upon request). The comparison with *EP* is a matter of parameterization, for fixed numbers (3) of Excess agents. With excess principals the outcome is the same no matter how many principals you have, so there would always be some inefficiency. Thus, there is a sense in which *EA* is ‘best’, if the relative number of agents can grow large.

However, this conclusion about the higher efficiency of the *EA* institutional environment should be qualified if we account for those agents who get unmatched in this environment.

Indeed, if the *EA* environment provides the biggest benefit from a match to the society, it may generate a higher social cost than the benchmark if the unmatched agents remain unemployed. Including all participants, the average individual payoff amounts to 269 in *B*, 237 in *EA* and 238 in *EP*. Therefore, there may be a dilemma: a higher competitive pressure among agents helps in getting closer to the first-best efficiency with the matched agents but could generate a higher social cost related to the unmatched agents if these ones, especially the low type, do not get an alternative occupation. We acknowledge that our measure of welfare, based on the total surplus, is purely utilitarian, and a Rawlsian approach would somewhat qualify our conclusions.

7. CONCLUSION

We have conducted an experiment based on a model of contracting under asymmetric information. We have shown theoretically that, in this context, various degrees of competition on the market 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 inequality level 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 theory. Our experiment finds that the institutional environment in which agents compete against each other is indeed the most efficient one. Our experiment also finds 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 show that principals

offer more generous menus than predicted, although less frequently in the context of competition.

These results question the interaction between the various degrees of bargaining power and social preferences. In many experimental papers, social preferences are shown to interfere with the predictions of standard contract theory. In our treatment without competition, we also observe that principals offer menus of contracts that are more generous than 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-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-type agents.

Finally, the superiority of the institutional environment with a 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.
- Ausubel, L.M. (1999). Adverse Selection in the Credit Card Market. Mimeo, University of Maryland.
- 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.
- Bull, C., Schotter, A., and Weigelt, K. (1987). Tournaments and Piece-Rates: An Experimental Study. *Journal of Political Economy*, 95(1), 1-33.
- Cabrales, A., and Charness, G. (2000). Optimal contracts with team production and hidden information: An experiment. Mimeo, University of California at Santa Barbara.
- 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 Rabin, M. (2002). Understanding Social Preferences With Simple Tests, *The Quarterly Journal of Economics*, 117(3), 817-869,
- Charness, G., and Dufwenberg, M. (forthcoming). Promises and Partnership. *Econometrica*.
- 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., and Kagel, J. (2005). Are Two Heads Better than One? Team versus Individual Play in Signaling Games. *American Economic Review*, 95(3), 477 – 509.
- Cooper, D., Kagel, J., Lo, W., and Gu, Q.L. (1999). Gaing 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., and Lundholm, R. (1985). Ripoffs, Lemons and Reputation Formation in Agency Relationships: A Laboratory Market Study, *Journal of Finance*, 40, 809-823.
- 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. (2003). Fairness, Errors, and the Power of Competition. Institute for Empirical Research in Economics, University of Zürich Working Paper n°133.
- 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.
- Grosskopf, B. (2003). Reinforcement and Directional Learning in the Ultimatum Game with Responder Competition. *Experimental Economics*, 6, 141-158.
- Güth, W., Klose, W., Koenigstein, M., and Schwalbach, J. (1998). An Experimental Study of a Dynamic Principal-Agent Relationship, *Managerial and Decision Economics*, 19:327-341.
- 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.
- Mirrlees, J.A. (1971). An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies*, 38(114), 175-208.
- Nalbantian, H.R., and Schotter, A. (1997). Productivity under Group Incentives: An Experimental Study. *American Economic Review*, 87, 314-41.
- Plott, C., and Wilde, L.L. (1982), Professional Diagnosis vs. Self-Diagnosis: An Experimental Investigation of Some Special Features of Markets with Uncertainty; *Research in Experimental Economics*, Vol.2; JAI Press, 63-112.
- 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.
- Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16, 8-37.
- Wilson, R. (1993). *Nonlinear Pricing*. Oxford: Oxford University Press.
- 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 treatment¹⁹

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?

¹⁹ The instructions for the other treatments are available upon request.

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 consists of only one independent observation, since there is considerable 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 $p = 0.03$ for the comparison 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 $p = 0.06$.²⁰ In addition, high agents in the *B* treatment are less likely to choose option X in response to menus 1 and 2 than are high 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 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

²⁰ We choose menus 1-3 as there are few rejections of menus 4-6 and these all occurred in the *B* treatment.

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 ($Z = 3.99, 5.71, \text{ and } 5.69$ 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 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 ($Z = 1.56, p = 0.11$).