

<p>Volume 77, Issue 2, February 2011</p> <p>Founding Editors: RICHARD H. DAY BENJAMIN G. WINTER</p> <p>Honorary Editors: GEORGE A. AKERLOF RICHARD A. EASTERLIN ALAN KRUMAN REINHARD SELTEN VERNON L. SMITH OLIVER E. WILLIAMSON</p> <p>Editor: William S. Nafziger</p> <p>Co-Editors: CATHERINE ECKEL GEORG KIRCHSTEIGER CLAUDE MENARD JÖRG SCHMIDT J. BARKLEY ROSSER, JR. NICOLAUS J. VRIENS</p> <p>Book Review Editor: ROGER KOEPL</p> <p>Associate Editors: JOHN A. CANTWELL JEFFREY P. CARPENTER GARY CHARNES SHAJEHUO CHEN ANDREW CLARK JOHN CONLISK RACHEL GROSSMAN SATYA P. DAS HERBERT DAVID DAVID DEQUECH ZENKUN DI M. DUPPENBERG D. ENGELMANN ENRIQUE FIAS DARREN FILSON DUNCAN K. FOLEY BRIAN GAECHTER LAURA GARIBAY HERBERT GINTIS CAROL SPANHAM WILLIAM C. GRANT DAVID GREYER GEOFFREY M. HOODSON GARY HORNES ED HOPRINS STEVEN HORWITZ DANIEL HOUSER STEPHEN HUCK PENKA IMAKULINAS GILJA IORI MICHAEL KOPEL TAMIR KURAN CHRISTOPH KUZMCS JOHN LERNER JOHN A. LIST THOMAS LUX INES MACHO-STADLER AKIO MATSUMOTO BERNARDO MUELLER WILLIAM NELSON CHARLES NOUSBAIR KARINE NYBORG AXEL OCKENFELS MARK PINKLE VICTOR POLTERNOVICH JASON POTTIS LUCA RIGOTTI ABDOLKARIM SADREH RAJY SETHI JASON F. SHOGREN PETER SICKT STEPHEN C. SMITH REINE STEINBUCKA MARIE THURSBY MARE CLAUDE VILLEVAL FRANK WESTERHOLZ WILLIAM C. WOOD</p>	<p>ISSN 0167-2681</p> <h1 style="color: red;">JOURNAL OF Economic Behavior & Organization</h1> <p>CONTENTS</p> <p>E.G. Haug and N.N. Taleb, Option traders use (very) sophisticated heuristics, never the Black-Scholes-Merton formula 97</p> <p>M. Madajewicz, Joint liability versus individual liability in credit contracts 107</p> <p>G. Bottazzi, G. Davettag and F. Pancotto, Does volatility matter? Expectations of price return and variability in an asset pricing experiment 124</p> <p>D.R. Stockman, Chace and capacity utilization under increasing returns to scale 147</p> <p>A. Cabrales and G. Charness, Optimal contracts with team production and hidden information: An experiment 163</p> <p>C.R. Proaño, Exchange rate determination, macroeconomic dynamics and stability under heterogeneous behavioral FX expectations 177</p> <p>K. Peren Arin, O. Lorz, O.F.M. Reich and N. Spagnolo, Exploring the dynamics between terrorism and anti-terror spending: Theory and UK-evidence 189</p> <p>D. Li, Commitment and compromise in bargaining 203</p> <p>P.J. Zak, Moral markets 212</p> <p>E.L. Khalil, How not to think about moral values 234</p> <p>P.J. Zak, Reply to letter to the Editor 239</p>
---	---

www.elsevier.com/locate/jebo

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jeboOptimal contracts with team production and hidden information: An experiment[☆]Antonio Cabrales^a, Gary Charness^{b,*}^a Department of Economics, Universidad Carlos III de Madrid, Spain^b Department of Economics, UC Santa Barbara, United States

ARTICLE INFO

Article history:

Received 23 March 2009

Received in revised form 9 June 2010

Accepted 18 October 2010

Available online 30 October 2010

JEL classification:

A13

B49

C91

C92

D21

J41

Keywords:

Experiment

Hidden information

Optimal contract

Production team

Wage rigidity

ABSTRACT

We devise an experiment to explore optimal contracts in a hidden-information context. A principal offers one of three possible contract menus to a team of two agents of unknown skill levels, with both agents' participation needed for production. We observe numerous rejections of the more lopsided menus, and principals respond by offering more favorable menus. Apart from rejections, we see almost complete separation in agent choices according to the agent types. Behavior converges towards a consensus in which one of the more equitable menus is proposed and agents accept a contract. The consensus menu differs across two treatments in which we vary the payoffs resulting from a rejection. We find strong evidence of social learning by low-skill agents (but only for low-skilled agents), in that a low-skill agent is more likely to reject a contract menu if her teammate rejected a contract menu in the previous period. In addition, low-skilled agents have a particularly adverse reaction to reduced wage offers.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

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

[☆] This research was undertaken while Charness was affiliated with Universitat Pompeu Fabra. We gratefully acknowledge the financial support of Spain's Ministry of Education under grants CONSOLIDER INGENIO 2010 (CSD2006-0016), ECO2009-10531 and SEJ2006-11665-C02-00. Charness also gratefully acknowledges support from the MacArthur Foundation. We thank George Akerlof, Robert Anderson, Rachel Croson, Brit Grosskopf, Ernan Haruvy, Botond Koszegi, Luca Rigotti, Reinhard Selten, Klaus Schmidt, Joel Sobel, seminar participants at the Micro Theory workshop at Universitat Pompeu Fabra, UC Berkeley, Bonn, and the Stanford Institute for Theoretical Economics for comments, and Ricard Gil for research assistance. All errors are our own.

* Corresponding author at: University of California Santa Barbara, Department of Economics, 2127 North Hall, Santa Barbara, CA 93106-9210, United States. Tel.: +1 805 893 2412; fax: +1 805 893 8830.

E-mail addresses: antonio.cabrales@uc3m.es (A. Cabrales), charness@econ.ucsb.edu (G. Charness).

¹ Mas-Colell et al. (1995) provide a good introduction to this topic, as well as its numerous applications.

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

of contracting. In the standard scenario, a firm hires a worker but knows less than the worker does about her innate work disutility.³

Private information leads to inefficiency because it is effectively a form of monopoly power (of information). Sometimes it is possible to introduce competition (such as auctions) as a method of reducing informational rents. If competition is not a possibility, mechanism design can still effectively minimize the rents of the privately informed, provided that there are more dimensions in preferences than in the informational problem. If a principal knows workers care both about wages and the number of hours worked, he can devise a contract menu of hours and wages that induces more truthful revelation and reduced inefficiency.⁴

However, while agency theory has traditionally sought incentive-compatible mechanisms on the assumption that people care only about their own material wealth, economic interactions frequently are associated with social approval or disapproval. In dozens of experiments (e.g., Güth et al., 1982; Charness and Rabin, 2002), many people appear to be motivated by some form of social preferences, such as altruism, difference aversion, or reciprocity. Recently, contract theorists such as Casadesu-Masanell (2004) and Rob and Zemsky (2002) have expressed the view that contract theory could be made more descriptive and effective by incorporating some form of behavioral considerations into the analysis. von Siemens (2004) considers optimal contracts in relation to hidden information and social preferences, finding that inequity aversion (Fehr and Schmidt, 1999) “increases the distortions originating from asymmetric information in a monopsony.”

We are only aware of two experimental studies of the static principal-agent problem with hidden information. Güth et al. (2001) conduct an experiment in which a principal faces two agents with unequal productivity functions. The principal finding is that making work contracts observable leads to a greater degree of pay compression. Effort choices differ systematically from the “rational” choices in relation to concerns of horizontal fairness. Charness and Dufwenberg (2009) consider the hidden-information problem in an experiment, and find that cheap-talk statements are very useful in achieving desirable outcomes when less-able agents can participate in a Pareto-improvement over the principal's outside option, but not when this Pareto-improvement is not feasible.

Another consideration is that of social learning. Ellison and Fudenberg (1993, p. 612) use the term social learning to describe contexts where agents' decisions reflect the experience of others. We are particularly concerned with whether social norms are susceptible to learning, or whether one's social preferences are immune to information about the actions of others. Previous experimental work in this area has primarily involved social learning in intergenerational games with explicit advice,⁵ while some recent work (e.g., Charness et al., 2007; Eckel and Wilson, 2007; Zafar, 2009) has found evidence that people change their social behavior upon mere observation of the behavior of others, as in Ellison (2002).

In the light of these considerations, we conduct an experimental test of contracts with hidden information. In our design, there are two types of agents (high-skill and low-skill, with effort less costly for the high-skill agent) and it is common information that these types are equally prevalent. A principal selects one of three menus, each having two possible contracts, to a pair of agents of unknown types. Each individual agent, who knows her own type and the menus available to the principal, then independently selects one of the two contracts offered on the menu or rejects both. Incentive-compatibility based on standard preferences should separate the types' optimal choices for every menu and no rejections should ever be observed. The menus are ranked with respect to how much they favor the principal.

If both agents accept a contract, the contracts are implemented; if either agent rejects, both the agents and the principal receive symmetric payoffs that differ across treatments. By introducing contracts that must be accepted by both workers, we contemplate the common situation where contracts must be negotiated with a union and then approved by the workers.⁶ Besides this feature, our environment, with 3-person groups and interactive preferences, leads to a realistic structure (without otherwise altering the theoretical contractual benchmark) for the way in which subjects receive feedback; this is important since one of our purposes is to understand the way in which agents learn about social norms.

Since people frequently do not act as pure money-maximizers in experiments, there is the immediate conjecture that the usual theoretical predictions will be rejected; nevertheless, we observe that principals usually initially propose the theoretically predicted contract, with this being significantly more likely in the treatment with higher reservation payoffs (those in the case of rejection). When these early-period contracts are rejected sufficiently often (how often depends very much on the individuals and on the reservation payoffs), the principals who were offering them instead choose progressively less self-favorable alternatives, until rejections cease and an ‘equilibrium’ menu is reached; this menu differs across the two treatments. One might expect principals to recognize that rejections are more costly for agents when the rejection payoff is reduced, so that they offer the most selfish contract menu with some impunity. However, instead the contract menus initially offered are less selfish when the rejection payoff is lower and the contract menu is more stable over time.

We also find a nearly complete separation across types of agent when the agent does not reject the menu that is offered, with the high-type skill choosing low effort less than 6% of the time and the low-skill agent choosing high effort less than

³ 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.

⁴ This is with respect to a situation where the agent cares about only one variable (as in Akerlof, 1970).

⁵ See for example, Schotter and Sopher (2003, 2007), Ballinger et al. (2003), Chaudhuri et al. (2006) and Çelen et al. (2010).

⁶ In essentially all of Europe, collective bargaining involving trade unions covers more than 75% of all workers (Layard et al., 1994). Our design assumes that a contract structure that affects all workers needs to be approved by a supermajority rule.

1% of the time (and only in the very first period). Apart from the menu offered *per se*, agents are less likely to reject a menu when it is an improvement over the previous menu offered. Finally, we observe that whether a low-skill agent rejects either of the less favorable menus is strongly influenced by whether the other agent rejected the menu offered in the previous period; however, this form of social learning does not apply to the high-skill agent. This suggests that the social preferences of high-skill agents who reject are less malleable than those of low-skill agents who reject, which makes sense since a high-skill agent who rejects a menu stands to lose more than a low-skill agent who rejects a menu.

Our paper offers several contributions to the literature. First, since there has been little experimental work on hidden information in a principal-agent setting, the observed separation by types is new evidence on this important empirical issue. Second, we find evidence of social learning that varies according to the type of agent; we are unaware of any previous evidence along these lines. Third, we observe a definite trade-off between overall efficiency and the distribution of earnings in relation to the rejection payoffs. In the long run, agents receive a higher share of earnings when they can reject with less cost; however, if the time horizon is short, the efficiency consideration is more relevant. Thus, the socially preferred outside option (and the resulting contract menu) may well be a function of the time horizon.

The remainder of the paper is organized as follows: we describe the model in Section 2 and our experimental methodology in Section 3. Our experimental results are presented in Section 4 and we offer some discussion and our conclusion in Section 5.

2. The model

In this section we describe the theoretical model that serves as the basis for the experimental design. Imagine that a firm needs two workers in order to be able to operate. The profits for the firm when it is operating are:

$$\Pi = e^1 - w^1 + e^2 - w^2$$

where e^i , w_i are, respectively, the effort levels and wages of worker $i \in \{1,2\}$. Each worker i has a utility function which depends on her type $j \in \{H,L\}$, which is her private information:

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

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 efforts levels are:

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

We call \hat{e}_j the efficient level of effort.⁷ 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:

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

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:

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

The second-best optimal contracts, when the types are private information of the agents result from the solution of the maximization program:

$$\max_{w_H, w_L, e_H, e_L} 2(p_H(e_H - w_H) + p_L(e_L - w_L))$$

subject to:

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

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

⁷ 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.

Table 1
Parameter values.

Menu	k_L	p_L	e_H	e_L	w_H	w_L
1	2	1/2	1	0.33	0.69	0.24
2	2	1/2	1	0.50	0.88	0.39
3	2	1/2	1	0.33	0.94	0.36

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

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

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:

$$e_H^* = \frac{1}{k_H} = 1; \quad e_L^* = \frac{1 - p_H}{k_L - k_H p_H}; \quad w_L^* = \underline{U} + \frac{k_L}{2} \left(\frac{1 - p_H}{k_L - k_H p_H} \right)^2; \quad w_H^* = \frac{1}{2} + w_L^* - \frac{1}{2}(e_L^*)^2 \quad (2)$$

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:

$$w_H^* - \frac{1}{2} - \underline{U} = \frac{k_L - 1}{2} \left(\frac{1 - p_H}{k_L - k_H p_H} \right)^2$$

The effort of the low type of agent is ‘inefficiently’ low and she obtains no rents. This is a subgame-perfect equilibrium.⁸

We implemented the theoretical model in our experiment by choosing values for the parameters in the three permitted contract menus 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. While we thus limit the possibilities available to the principal, a continuous strategy space would make the data analysis problematic (even ignoring the increased complexity of the decisions of the experimental participants), without adding much insight.

We chose $k_L = 2$ for all menus, in order to give relatively large rents to the H type (under her preferred contracts). Menu 1 is the ‘theoretically predicted’ menu; it is not first-best efficient (since $e_L \neq \frac{1}{2}$) and has the most unequal payoffs. Here the values for e_i and w_i are obtained from Eq. (2).⁹ An H agent could obtain moderate rents (if she chose the ‘right’ contract and one of the contracts was accepted by the other agent) and an L agent could receive very small rents.¹⁰ In Menu 2 the effort choices were the efficient ones, computed from Eq. (1). The value for w_L is set so that the L agent could receive small rents, while the value for w_H provides the H agent with higher rents than in Menu 1. In Menu 3, both types of agents can receive substantial rents, and (as in Menu 1) the efforts of both types correspond to the optimal ones in the theoretical model.¹¹ The parameters, efforts, and wages for the different menus in the experiment are summarized below.

One of the criticisms of models of optimal contract design in adverse selection contexts is that the theoretically predicted 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. This would be quite complicated to design for the principal, and even the choice of the agent would not be simple. While we have selected a very simple structure (only two types), we feel that a ‘simple’ menu can serve as an approximation for the fully-optimal schedule. As Wilson (1993) points out (p. 146) in a representative example: “The firm’s profits from the 5-part and two-part tariffs are 98.8% and 88.9% of the profits from the nonlinear tariff.”

⁸ There is one slightly non-standard feature of this model that should be mentioned. Since the agents’ decisions are simultaneous, and a rejection implies that both agents receive the outside option, there exist subgame-perfect equilibria of the game, whose outcomes are different from the one we have just described. If one agent expects the other to reject the contract menu, it is a best response to reject contracts that give her a higher utility than \underline{U} . This can be used to construct a variety of subgame-perfect equilibria. However, notice that any strategy that rejects a contract yielding a higher utility than \underline{U} is weakly dominated. While such equilibria are subgame-perfect, they are not trembling-hand perfect (Selten, 1965), and do not survive one round of deletion of weakly dominated strategies (Dekel and Fudenberg, 1990).

⁹ All payoffs were rounded to the nearest 25 units in our payoff table.

¹⁰ In the theoretical model the rents for the L player are exactly zero. We chose to make the rents positive (but very small) to make acceptance strictly dominant while remaining very close to the “theoretical prediction.”

¹¹ In Menu 2, the high-type agent is given a wage that respects incentive compatibility, and an extra .25 is added. This was done primarily to see if a low-type agent will refuse to reject an unfavorable menu for fear of hurting an innocent bystander who is getting a fair deal. Menu 3 is just like Menu 1, but each type of agent receives this gift of .25 to the wage. This is still incentive compatible, acceptable, and asymmetric.

Payoff table.

	2 HIGH responders					1 HIGH, 1 LOW responder					2 LOW responders				
	11	12	21	22	VETO	11	12	21	22	VETO	11	12	21	22	VETO
Proposer	3950	3075	3075	2175	500	3950	3075	3075	2175	500	3950	3075	3075	2175	500
Responder 1	775	775	725	725	500	775	775	725	725	500	-1275	-1275	525	525	500
Responder 2	775	725	775	725	500	-1275	525	-1275	525	500	-1275	525	-1275	525	500
Proposer	2500	2400	2400	2300	500	2500	2400	2400	2300	500	2500	2400	2400	2300	500
Responder 1	1450	1450	1050	1050	500	1450	1450	1050	1050	500	-550	-550	550	550	500
Responder 2	1450	1050	1450	1050	500	-550	550	-550	550	500	-550	550	-550	550	500
Proposer	2050	1625	1625	1175	500	2050	1625	1625	1175	500	2050	1625	1625	1175	500
Responder 1	1725	1725	1225	1225	500	1725	1725	1225	1225	500	-325	-325	1000	1000	500
Responder 2	1725	1225	1700	1225	500	-325	1000	-325	1000	500	-325	1000	-325	1000	500

3. Experimental procedures

Six sessions were conducted at Universitat Pompeu Fabra in Barcelona. There were 12 people (four principals, four high-type agents, and four low-type agents) in each session, so 72 participants in all. Groups of three (one principal and two agents) were matched randomly in each of the 15 periods, subject to the restriction that no group was ever repeated in consecutive periods. While there were few repeated 3-groups, each agent could expect to be matched with each principal several times during the experimental session. On average, each participant received about 13 euros, including a show-up fee of 4 euros. Sessions lasted less than 2 h.

At the beginning of a session, the instructions and a decision sheet were passed out to each subject. The decision sheet stated the subject number and the role (principal, high-type agent, or low-type agent). Instructions (presented in Appendix A) covered all rules used to determine the payoffs to each player in the group; these were read aloud to the entire room. We included a table showing the monetary payoffs for every possible combination of actions. We verbally reviewed every case in detail, asking random participants questions along the way to ensure that the process was understood. We proceeded only after the structure of the decision task and the payoffs seemed to be clear.

When the instructional phase was concluded, we proceeded with the session. In each period the principals first selected a menu on their decision sheets. Each matched agent could accept choice 1 or 2 from this menu, or reject both options. If both agents in the group accepted contracts, each obtained the corresponding payoff for an agent of her type. If either of the agents rejected both choices 1 and 2, then the payoffs for both the principal and the agents were the same (500 units or 250 units depending on the treatment). Payoff units converted to euros at the rate of 125 units to a euro.

The experimenter went around the room collecting this information, with care taken to preserve the anonymity (with respect to experimental role) of the principals by going to each cubicle and at least appearing to write information on each participant's record sheet. Once the principals' menu selections were recorded, the experimenter again went around the room, this time providing the information about the menu to the agents (again preserving anonymity). The agents then made their choices and the experimenter collected this information; finally, the experimenter privately informed each participant about the choices and types (but not the identities) of both agents in the group.

Participants knew that there would be 15 periods in all. At the end of the session, participants were paid privately, based on the payoffs achieved in a randomly selected round, as was indicated in the instructions.¹² As mentioned earlier, two types of sessions were conducted, and these differed only with respect to the reservation payoffs for a rejection. We conducted three sessions for each treatment.

4. Experimental results

Summarizing our main results, many principals choose Menu 1, the most self-favoring contract menu, in the early periods. However, both types of agents often reject Menu 1; there is evidence of significant social learning by low-skill agents in this regard, as we see that many agents do not reject unfavorable menus in the beginning, but begin to do so after observing another agent doing so. This process induces the principals to offer a more equitable menu, leading to a substantial degree of convergence on a 'community consensus' by the end of 15 periods. However, the consensus menu differs across treatments, with Menu 3 becoming the most frequent choice in Treatment 1 and Menu 2 becoming the most frequent choice in Treatment 2. We also find a clear separation in contract selection among the agents who do not reject the contract menu proposed.

4.1. Principal behavior

Figs. 1 and 2 visually illustrate the patterns of menu proposals over time. As can readily be seen, these proposals change over the course of the sessions, with the more dramatic changes taking place in Treatment 1.

¹² This was done in an effort to make payoffs more salient to the subjects, as this method makes the nominal payoffs 15 times as large as would be the case if payoffs were instead aggregated over 15 periods, and it also avoids possible wealth effects from accumulated earnings.




Fig. 1. Proposals over time (Treatment 1).

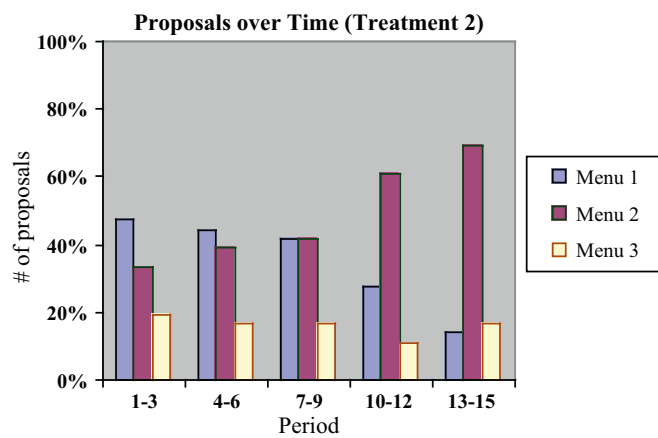


Fig. 2. Proposals over time (Treatment 2).

Table 2
Menu proposals by period.

	Period															Total
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
Treatment 1																
Menu 1	12	10	9	4	3	3	2	3	2	1	3	2	2	1	5	62
Menu 2	0	2	3	4	6	4	4	4	1	3	1	1	4	2	1	40
Menu 3	0	0	0	4	3	5	6	5	9	8	8	9	6	9	6	78
Treatment 2																
Menu 1	7	5	5	6	6	4	5	7	3	4	3	3	2	1	2	63
Menu 2	3	5	4	5	5	4	5	2	8	7	7	8	7	10	8	88
Menu 3	2	2	3	1	1	4	2	3	1	1	2	1	3	1	2	29

Table 2 provides period-by-period detail of the contract menus proposed in each treatment, along with the total number of times each menu is proposed in each treatment:

Simple inspection shows that the most common proposal in Treatment 1 is Menu 3, which is chosen at a rate of 43.3%; the most common proposal in Treatment 2 is Menu 2, which is chosen at a rate of 48.9%. The percentage of Menu 2 (Menu 3) contracts offered is lower (higher) in each and every Treatment 1 session than in each and every Treatment 2 session, and so the difference across treatments is significant at $p=0.05$ using the Wilcoxon–Mann–Whitney test, even with the conservative statistical approach of treating each of our sessions as being only one independent observation.

Nevertheless, these overall proportions do not reflect the considerable changes over time. For example, note that in the very first period Menu 1 was chosen by each of the 12 principals in Treatment 1 and seven of 12 principals in Treatment 2.¹³ For Treatment 1, Menu 1 was chosen at a rate of 86.1% in the first three periods, but at a rate of only 22.2% in the last three

¹³ Note that the test of proportions (Glasnapp and Poggio, 1985) finds a significant difference between these proportions ($Z=2.51, p=0.012$, two-tailed test).

Table 3
Menu changes by principals.

	Higher menu	Same menu	Lower menu
No rejection in prior period			
Treatment 1	10 (10%)	66 (65%)	25 (25%)
Treatment 2	16 (11%)	99 (70%)	26 (18%)
Rejection in prior period			
Treatment 1	37 (55%)	18 (27%)	12 (18%)
Treatment 2	12 (44%)	13 (48%)	2 (7%)

Table 4
Lagged rejections and menu choice.

Independent variables	Dependent variable					
	Menu up (all data) (1)	Menu up (T1 only) (2)	Menu up (T2 only) (3)	Menu down (all data) (4)	Menu down (T1 only) (5)	Menu down(T2 only)(6)
Lagged rejection	1.360*** (0.192)	1.539*** (0.270)	1.467*** (0.388)	-0.549** (0.228)	-0.504* (0.302)	-0.648 (0.466)
N	336	168	168	336	168	168
LL	-141.7	-71.7	-58.1	-151.4	-75.8	-64.4

Standard errors are in parentheses. Period dummies are included in the regressions, but are suppressed in the table, as none were significant.

- * Statistical significance at the 10% level, with two-tailed tests.
- ** Statistical significance at the 5% level, with two-tailed tests.
- *** Statistical significance at the 1% level, with two-tailed tests.

periods; by contrast, Menu 3 was never chosen in the first three periods, but was chosen at a rate of 58.3% in the last three periods. There is also considerable change over time in Treatment 2: Menu 1 was chosen at a rate of 47.2% in the first three periods, but at a rate of only 13.9% in the last three periods; by contrast, Menu 2 was chosen at a rate of 33.3% in the first three periods, but at a rate of 69.4% in the last three periods.

In Treatment 1 (Treatment 2), when Menu 1 is proposed, it is rejected by at least one of the two agents 69.4% (39.7%) of the time. Principals appear to respond to rejections, rather than changing their behavior in a vacuum.¹⁴ Table 3 presents the data concerning whether or not a principal changed the contract menu after observing either joint acceptance or a rejection by at least one agent (14 observations for each principal). Principals are substantially more likely to select a higher-numbered menu after a rejection than after no rejection, with the likelihood of a change to a more 'generous' menu being four or five times greater.

We use random-effects ordered-probit regressions to account for the 15 observations for each participant; we also include period dummies to account for possible time trends.¹⁵ We consider whether the principal was more likely to choose a contract menu less favorable to himself depending on whether there was a rejection in the previous period, and also whether the principal was more likely to choose a contract menu less favorable to himself after a contract menu was not rejected in the previous period.

We find that a principal is significantly more likely to choose a less aggressive menu after a rejection in the previous period, for both the pooled data and the separate treatments ($p < 0.001$ in all cases). It appears that rejections drive the changes in principal behavior over time (Table 4).¹⁶

Finally, principal choices vary considerably across individuals. Table A.3 in Appendix B shows the number of times that each principal chose each contract menu. Differences across individuals are immediately apparent. For example, one principal in Treatment 2 chose Menu 1 in all 15 periods, another principal chose Menu 2 in all 15 periods, and yet another principal chose Menu 3 in all 15 periods; the differences in Treatment 1 are a bit less stark, but we nevertheless see that one principal chose Menu 3 only once, while two other principals chose it in 11 of 15 periods.

4.2. Agent behavior

Although agents who are concerned only with maximizing their own material reward should never reject a contract menu, rejections are quite common. Table 5 provides a summary of rejections by session, contract menu, and responder type:

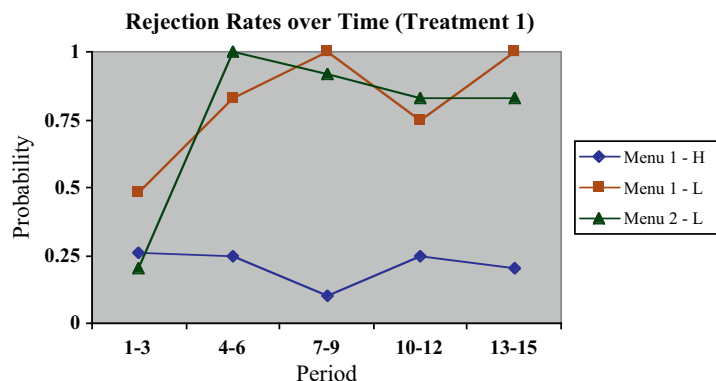
¹⁴ Tables A.1 and A.2 in Appendix B shows each principal's menu choice in every period of Treatments 1 and 2, respectively, along with the responses received.

¹⁵ We choose period 8 as the baseline, as this seemed a reasonable baseline (early exploration is largely finished and any potential unraveling should not yet be a factor). In fact, the results are robust to any intermediate period being chosen as the baseline, and no significant period effects were found.

¹⁶ We also consider whether a lagged rejection (or the lack thereof) plays a role in whether a principal chooses a lower-numbered menu (i.e., makes a more aggressive menu choice). We find, for the pooled data, that a principal is significantly more likely to make the aggressive change when there is no lagged rejection. However, this effect is not quite significant when treatments are considered separately, although it does achieve 5% significance with the Treatment 1 data, using an (appropriate) one-tailed test.

Table 5
Rejections.

Session	Menu 1		Menu 2		Menu 3	
	H	L	H	L	H	L
1	5/20	13/18	0/13	11/15	0/27	0/27
2	7/27	16/23	0/12	15/16	0/21	0/21
3	3/20	9/16	0/11	10/13	0/29	0/31
Treatment 1 total	15/67	38/57	0/36	36/44	0/77	0/79
4	0/21	6/17	0/21	2/29	0/18	0/14
5	10/26	0/24	0/28	0/30	0/6	0/6
6	7/17	5/21	0/36	1/32	0/7	0/7
Treatment 2 total	17/64	11/62	0/85	3/91	0/31	0/27



Rejection rates were always 0% for Menu 2 – H, Menu 3 – H, and Menu 3 – L.

Fig. 3. Rejection rates over time (Treatment 1). Rejection rates were always 0% for Menu 2–H, Menu 3–H, and Menu 3–L.

We first note that Menu 3 is never rejected (0 of 214 times) by any agent, despite the fact that principals on average earned considerably more than the agents even with this contract menu. Thus, it is not inequality *per se* that drives rejections.

In Treatment 1, low-skill agents have by far the highest rate of rejection, with a rate of 66.7% for Menu 1 and 81.8% for Menu 2, reflecting the fact that accepting Menu 1 (Menu 2) improves payoffs by only 5% (10%).¹⁷ It may seem surprising at first that the rejection rate is actually higher for Menu 2 than for Menu 1, although the acceptance payoff to low-skill agents is not much higher with Menu 2. The higher rate for Menu 2 perhaps reflects that low-skill agents learn to reject over the course of a session, and that Menu 2 is rare in the early periods. Fig. 3 shows the rejection rates over time in Treatment 1, aggregated over three periods for smoothing.


Note the large increase in rejection rates for low-skill agents from the first three periods to the next three periods. High-skill agents also reject Menu 1 in Treatment 1, even though it is considerably more costly for them to do so. The rejection rate is 22.4%, and is relatively stable over time. No high-skill agent ever rejected Menu 2.

Rejection rates are considerably lower in Treatment 2, and significantly so. At the session level, the percentage of rejections is (much) higher in every Treatment 1 session than in every Treatment 2 session, and so the difference across treatments is significant at $p = 0.05$ using the Wilcoxon–Mann–Whitney test. Since it costs a low-skill agent 11 times as much (275 versus 25) to reject Menu 1 in Treatment 2 compared to Treatment 1, the rejection rate by low-skill agents is only 17.7%, much less than in Treatment 1; once again, this rate is higher in every Treatment 1 session than in every Treatment 2 session, confirming statistical significance. The rejection rate (26.6%) for Menu 1 for high-skill agents is higher than for low-skill agents; this rate is also slightly higher than the high-skill-agent rejection rate of Menu 1 in Treatment 1. In quite a contrast to Treatment 1, low-skilled agents very rarely (3.3% of the time) reject Menu 2, while once again no high-skill agent ever rejects this contract menu.

Fig. 4 shows the rejection rates over time in Treatment 2, aggregated over three periods for smoothing.

Once again, we see an increase over time in the likelihood that a low-skill agent will reject Menu 1, although this is delayed (perhaps because there are fewer rejections to observe from the other agent). As with Treatment 1, there is no drop in any rejection rates over time. While one standard supergame explanation for rejections would imply that rejection rates

¹⁷ The rejection rate for Menu 2 was higher than for Menu 1 for seven of the low-skill agents in Treatment 1, compared to the reverse for two of the low-skill agents (with no difference for the other two low-skill agents).



Rejection rates were always 0% for Menu 2 – H, Menu 3 – H, and Menu 3 – L.

Fig. 4. Rejection rates over time (Treatment 2). Rejection rates were always 0% for Menu 2–H, Menu 3–H, and Menu 3–L.

Table 6
Determinants of rejection.

Variable	Coefficient	Z-statistic	p-value	Marg. effect
Treatment 1	1.136	3.95	0.000	0.321
Menu	−0.531	−2.69	0.007	−0.137
Lagged rejection	0.656	3.16	0.002	0.202
High_skill*Lagged rejection	−0.821	−2.25	0.025	−0.145
Menu_up	−0.382	−1.78	0.076	−0.088
Menu_down	0.782	3.22	0.001	0.218
High_skill*Menu_down	−1.096	−3.29	0.001	−0.201
Period	−0.018	−0.83	0.405	−0.005
Constant	−0.472	−1.22	0.224	−

N = 462; log pseudo-likelihood = −184.852. All p-values reflect two-tailed tests.

drop over time,¹⁸ we see no evidence in our Treatment 1 data of this occurring. Thus, strategic motivations alone do not provide an explanation for the observed behavior.¹⁹

Aside from rejections, the mechanism does successfully separate the types of agent in the types of contract accepted. Overall, of the 600 contract acceptances, 577 (96.2%) correctly mapped the agent to the predicted type. Low-type agents only chose high effort in three cases of 360 (0.8%), all in the first period; there were 20 cases of 360 (5.6%) where a high-type agent chose low effort.

As with principals, we find that there is considerable heterogeneity among the agents in the population; this can be seen in Tables A.4–A.6 in Appendices C and D. For example, one low-skill agent in Treatment 1 never rejected a contract menu, while another low-skilled agent rejected Menu 1 and Menu 2 all 10 times they were proposed. In Treatment 2, 13 agents never rejected a contract menu, while one high-skilled agent rejected Menu 1 six times and two low-skilled agents rejected contract menus five times.

Overall, 16 of 24 low-skilled agents and 11 of 24 high-skilled agents rejected at least one proposed menu. Among the 27 agents who rejected at least one proposal, only six did so in the first period. For the other 21 agents who began to reject contract menus later, in 10 cases the first rejection occurred the first time (not counting Menu 3 proposals, since these were never rejected) after the agent observed the other agent rejected a contract menu. We also find that the likelihood of a low-skilled agent rejecting a particular contract menu depends on the menu the agent was offered in the previous period; this pattern is much weaker for high-skilled agents.²⁰

To investigate the influences of changes in menu and observed rejections by other agents while accounting for other factors, we perform a random-effect probit regression (with robust standard errors), with rejection as the dependent variable. Since Menu 3 was never rejected, we exclude observations in which this menu was offered (Table 6).

In this regression, lagged rejection = 1 if the other agent rejected the contract menu in the previous period and was otherwise 0; Menu_up = 1 if the contract menu was more favorable to the agent than the one in the previous period and was otherwise 0; Menu_down = 1 if the contract menu was less favorable to the agent than the one in the previous period and was otherwise 0; High_skill = 1 if the agent had high skill and was otherwise 0.

¹⁸ However, this is not the case for the Binmore “game of life” game, where agents behave as if the strategic interactions had a positive (and sufficiently large) probability of continuing at every point in time.

¹⁹ One specific bit of evidence is that, in the very last round, seven principals tried Menu 1, perhaps thinking that rejections were only being made for strategic purposes; however, these were rejected by all L types (6/6) and 25% of H types (2/8).

²⁰ In Treatment 1, low-skilled agents rejected Menu 1 18 of 22 times (81.8%) in the period after being offered a better contract menu, compared to 20 of 35 times (57.1%) otherwise; the corresponding figures for Treatment 2 are 9 of 30 (30.0%) and 2 of 32 (6.2%). High-skilled agents types rejected Menu 1 14 of 53 times (26.4%) in the period after being offered a better contract menu versus 17 of 78 (21.8%) times otherwise.

The regression confirms the strong treatment effect (more rejections in Treatment 1), the effect of the menu chosen on the rate of rejection (the higher the menu, the lower the rejection rate). We also see highly significant evidence of social learning on the part of the low-skilled agents, with an observed lagged rejection leading to an increase of 20 percentage points in the likelihood the agent chooses rejection. However, we do not observe social learning for the high-skilled agent. If a given contract menu is more favorable than the previous menu offered, an agent is slightly less likely to reject, but this effect is only marginally significant. On the other hand, holding the menu constant, getting a worse offer than in the last period leads low-skilled agents to be more likely to reject; however, this effect is not present for high agents. Finally, we see a small and insignificant decline in the rejection rate over time.

5. Conclusion

We explore the problem of optimal contract menus with hidden information and team production in a laboratory experiment matching a principal with two agents of unknown level of skill. Agents of both types often reject the most unfavorable contract menu, and principals respond to rejections of their proposals by offering more agent-favorable menus. After the principals learn the (evolving and heterogeneous) standard for menu acceptability, the production team functions in a relatively efficient manner, with agents choosing contracts in accordance with their types and nearly complete separation. Rejection rates are much higher in Treatment 1, where the reservation payoffs are higher. This difference in reservation payoffs also leads to a different prevailing contract menu in our two treatments, and there is a substantial degree of heterogeneity in the behavior of both principals and agents.

We also find strong social learning by low-skill agents, who are more likely to reject a contract menu if they have observed a rejection from the other agent in the previous period, perhaps updating their views about the social norms and adjusting their values accordingly. The likelihood of rejection by high-skill agents is not sensitive to observed rejections by other agents, perhaps because only the more motivated high-skill agents will sacrifice the larger offered payoff. Agents are less likely to reject a menu when it is an improvement over the one assigned in the previous period. Low-skill agents are considerably more likely to reject a particular menu when it offers a lower wage than the menu offered in the previous period, although we do not see this effect for high agents. Thus, we see downward rigidity in wages for the disadvantaged agent.

The nearly complete separation by agent types suggests that incentive-compatible contract menus can be very effective, provided that they are not considered unfair. Since there is so much heterogeneity regarding what is unfair, it may be difficult to develop a contract menu that is optimal for all agents; however, it may be that social learning is helpful in eventually achieving a consensus. In any event, since more effective contracts are likely to lead to better economic outcomes, we echo the view (Güth et al., 2001, p. 85) that “there is a need for behavioral contract theory, based on empirical findings.” It is clear that further evidence on contracts and agent behavior is needed.

Appendix A. Instructions

Thank you for participating in this experiment. The experiment will consist of a series of 15 decision periods. In each period you will be randomly and anonymously matched with two other persons; the action you choose and the action chosen by the persons with whom you are matched will jointly determine your payoffs in each period.

You have been assigned a subject number. Please retain this number, as we will need it to pay you at the end of the experiment.

A.1. Process

There are two classes of players: *proposers* and *responders*. The responders can be one of two types: HIGH or LOW. The class to which the player is assigned (proposer or responder) and the type of the players (in the case the player is a responder) are chosen randomly at the beginning of the game. Each responder has an equal initial probability to be of either type HIGH or type LOW. Half of all responders will be of each type. Each responder knows her type, but no other participant does. Your role (class and type) will not change during the experiment. Your subject number, class and type (if you are a responder) are printed on your decision sheet.

In each period you will be randomly matched in groups of three players, according to subject numbers. All groups will be composed of a proposer and two responders of any combination of types; *ex ante*, there is a 25% chance that both responders are HIGH, a 50% chance that one is HIGH and the other is LOW, and a 25% chance that both are LOW.²¹ The identity of the other players in the group is unknown to you and the composition of the groups will change randomly every period. While you will not know the matching process, we would be happy to show you (at the end of the experiment) how the matches were created.

Once the period begins each proposer must make a selection from one of 3 possible choices {1,2,3} and will do so by checking a box for that period on the decision sheet provided. We will come around the room and record each proposer

²¹ In fact, these were the actual probabilities, given our matching scheme (see Appendix 2). The actual *ex ante* probabilities are 3/14, 4/7, and 3/14.

selection. Next we will go around the room and mark the proposer selections on the decision sheets of the responders in the appropriate groups. At this point, the two responders in each group must each choose one of the three available options {1,2,VETO} by checking the corresponding box on the decision sheet. (For both proposers and responders, we ask that you do not fill in the spaces clearly marked as EXPERIMENTER.) We will then record these choices. Finally, we will once again go around the room and mark the responder decisions (and the type of responders) for each group on the decision sheets for all members of that group. At this point, you can calculate your payoff from the period from the table.

A.2. How choices depend on points

The payoffs will be a function of the proposer's choice and the responders' responses. The conversion rate from payoff units to euros is 125 units to one euro. Please refer to the table provided and we will offer some examples of how this process works. [This Table is at the end of Appendix A.]

First, you should understand that, unless one of the responders chooses to VETO the proposer's choice, the payoff for any responder depends **only** on the proposer's choice and the responder's choice. No person will ever receive a negative payoff *unless she chooses it herself*.

If **either** responder chooses to VETO the proposal, then the VETO payoffs (shown in the columns shaded in gray on the payoff table in your packet) would result.

If you are a Responder, you may be wondering how you can tell if you are Responder 1 or Responder 2. There is an algorithm you can use which will make your task easier: if you are a Responder of the HIGH type, simply consider yourself to be Responder 1; if you are a Responder of the LOW type, simply consider yourself to be Responder 2. In all cases, this will ensure that your payoffs correspond to your choices.

Suppose the proposer chooses option 1 and faces responders who are both type HIGH. Suppose further that both responders choose option 1. First, find the rows corresponding to Proposer Choice 1. Next, find the 5 columns corresponding to the case where both responders are HIGH. The column that is relevant in this case is headed by "11". As you can see, the Proposer would receive 3950 units, Responder 1 would receive 775 units and Responder 2 would receive 775 units. Suppose instead that Responder 1 chooses option 1 and Responder 2 chooses option 2. The column that is now relevant is headed by "12". In this case the Proposer would receive 3075 units, Responder 1 (who chose option 1) would receive 775 units, and Responder 2 (who chose option 2) would receive 725 units. If instead Responder 1 chooses option 2 and Responder 2 chooses option 1, the column that is now relevant is headed by "21". In this case the Proposer would receive 3075 units, Responder 1 (who chose option 2) would receive 725 units, and Responder 2 (who chose option 1) would receive 775 units. If instead Responder 1 chooses option 2 and Responder 2 chooses option 2, the column that is now relevant is headed by "22". In this case the Proposer would receive 2175 units, Responder 1 would receive 725 units, and Responder 2 would receive 725 units. Suppose instead that **either** Responder chooses to VETO the proposer's choice. In this case, the Proposer would receive 500 units and each Responder would receive 500 units.

Suppose the Proposer chooses option 2 and faces two LOW Responders. First, find the rows corresponding to Proposer Choice 2. Next, find the 5 columns corresponding to the case in which both responders are LOW. Suppose further that both responders choose option 1. The column that is relevant in this case is headed by "11". As you can see, the Proposer would receive 2500 units, Responder 1 would receive -550 units and Responder 2 would receive -550 units. Suppose instead that Responder 1 chooses option 1 and Responder 2 chooses option 2. The column that is now relevant is headed by "12". Then the Proposer would receive 2400 units, Responder 1 (who chose option 1) would receive -550 units, and Responder 2 (who chose option 2) would receive 550 units. If instead Responder 1 chooses option 2 and Responder 2 chooses option 1, the column that is now relevant is headed by "21". Then the Proposer would receive 2400 units, Responder 1 (who chose option 2) would receive 550 units, and Responder 2 (who chose option 1) would receive -550 units. If instead Responder 1 chooses option 2 and Responder 2 chooses option 2, the column that is now relevant is headed by "22". Then the Proposer would receive 2300 units, Responder 1 would receive 550 units, and Responder 2 would receive 550 units. Suppose instead that **either** Responder chooses to VETO the proposer's choice. In this case, the Proposer would receive 500 units and each Responder would receive 500 units.

Suppose the Proposer chooses option 3 and faces one HIGH responder and one LOW responder (by the way the table is written, the type HIGH is necessarily Responder 1 and the type LOW is necessarily Responder 2). First, find the rows corresponding to Proposer Choice 3. Next, find the 5 columns corresponding to the case where one responder is HIGH and the other is LOW. Suppose further that both responders choose option 1. The column that is relevant in this case is headed by "11". As you can see, the Proposer would receive 2050 units, Responder 1 would receive 1725 units and Responder 2 would receive -325 units. Suppose instead that Responder 1 chooses option 1 and Responder 2 chooses option 2. The column that is now relevant is headed by "12". Then the Proposer would receive 1625 units, Responder 1 (who chose option 1) would receive 1725 units, and Responder 2 (who chose option 2) would receive 1000 units. If instead Responder 1 chooses option 2 and Responder 2 chooses option 1, the column that is now relevant is headed by "21". Then the Proposer would receive 1625 units, Responder 1 (who chose option 2) would receive 1225 units, and Responder 2 (who chose option 1) would receive -325 units. If instead Responder 1 chooses option 2 and Responder 2 chooses option 2, the column that is now relevant is headed by "22". Then the Proposer would receive 1175 units, Responder 1 would receive 1225 units, and Responder 2 would receive 1000 units. Suppose instead that **either** Responder chooses to VETO the proposer's choice. In this case, the Proposer would receive 500 units and each Responder would receive 500 units.

A.3. Payment

Each person will be paid individually and privately. Only one of the 15 periods will be chosen at random for actual payment, using a die with multiple sides. In addition, you will receive 4 euros for participating in the experiment. If, in the period selected your payoff is negative, it will be deducted from the 4 euro show-up fee; however, no one will receive a net payoff less than 0.

If you have questions raise your hand and one of us will come and answer your question. Direct communication between participants is strictly forbidden. Please ask questions if you do not fully understand the instructions. Are there any questions?

Appendix B. Individual principal choices and responses

Tables A.1–A.3.

Table A.1

Individual principal choices over time in Treatment 1.

Pr. #	Period														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	1*/	2	2	1*	2**	2	3	2**	1*	2*	1*	3	3	3	1*
2	1	2	1	1*/	1*	2*	1**	1*	3	2*	3	3	1	2	1**
3	1	1*	1*	2*	3	3	3	3	3	3	3	3	3	3	3
4	1*	1**	1*	2*	3	2*	2*	3	3	3	3	3	3	3	3
5	1*	1	1**	2*	1**	3	2	2**	1*	3	1*	3	2**	3	2
6	1	1*	1	1**	2*/	1*	3	1	3	3	3	3	3	3	3
7	1/	1**	1	1*	2*	1**	1	2**	2*	1*	2**	1*	2*	3	1*
8	1	1*	1*	3	2	2*	2*	3	3	3	3	3	3	3	1**
9	1	1/	1**	3	2*	3	3	3	3	3	3	3	3	3	3
10	1/	1*	1*/	2*	1*/	3	3	3	3	3	3	3	1/	3	3
11	1*	1*	2	3	2*	3	3	2**	3	3	1*	2	2	1*	3
12	1*	1*	2	3	3	1	2*	1	3	2*	3	1/	2*	2*	1*

*Means a rejection and ** means both agents rejected the proposal. /Means a low play by an H type and // means two low plays by H types. Principals 1–4 were in the first session in the treatment, principals 5–8 were in the second session, and principals 9–12 were in the third session.

Table A.2

Individual principal choices over time in Treatment 2.

Pr. #	Period														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	2	2	2	2	2*	2	2	2	2	2	2	2	2	2	2
2	1	2	3	2*	1	2	2	1	2	2	2	1**	2	2	2
3	1/	1	1	1	1	1*	1	1*	1	1	1*	1	1	1	1*
4	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3
5	1	1	1	1	1	1*	1*	1	1/	2	2	2	2	2	2
6	2	1	3	2	2	3	2	1*	2	2	3	2	3	2	3
7	3	2	2	2	2	1	1*/	2	2	2	2	2	2	2	2
8	1	1*	1	1*	1	1	1*	1*	1/	1*	1//	2	1*	2	2
9	1	3	2	1*	1	3	2	3	2	1*	2	1*	3	2*	1
10	2	2	2	1	1**	3	3	3	2	2	1	2	2	2	2
11	1	2	1*	2	2	2	2	1*	2	2	2	2	2	2	2
12	1*	1	1	1*	2	2	1	1*	2	1*	2	2	2	2	2

*Means a rejection and ** means both agents rejected the proposal. /Means a low play by an H type and // means two low plays by H types. Principals 1–4 were in the first session in the treatment, principals 5–8 were in the second session, and principals 9–12 were in the third session.

Table A.3

Aggregate contract menu proposals, by principal.

	1	2	3	4	5	6	7	8	9	10	11	12	Total
Principal #—Treatment 1													
5		8	3	3	6	6	9	4	3	5	4	6	62
6		4	1	3	5	1	5	3	1	1	5	5	40
4		3	11	9	4	8	1	8	11	9	6	4	78
Principal #—Treatment 2													
0		4	15	0	9	2	2	12	6	3	3	7	63
15		10	0	0	6	8	12	3	5	9	12	8	88
0		1	0	15	0	5	1	0	4	3	0	0	29

Appendix C. Individual agent choices by period

Tables A.4 and A.5.

Table A.4
Menus and agent choices in Treatment 1 (500).

	Period														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
L1	1/2	2/2	1/2	1/2	2/3	3/2	2/3	3/2	3/2	2/3	1/3	3/2	3/2	3/2	1/3
L2	1/3	2/2	1/2	1/3	2/3	2/3	1/3	2/3	3/2	2/3	3/2	3/2	3/2	3/2	3/2
L3	1/2	1/3	2/2	2/3	3/2	2/3	3/2	2/3	1/3	2/2	3/2	3/2	3/2	3/2	3/2
L4	1/3	1/3	1/3	2/3	1/3	3/2	1/3	3/2	3/2	3/2	3/2	3/2	3/2	3/2	1/3
L5	1/2	1/2	1/2	1/3	1/3	1/3	2/3	2/3	2/3	3/2	1/3	3/2	2/3	3/2	3/2
L6	1/3	1/3	1/3	2/3	1/3	2/3	3/2	2/3	3/2	3/2	2/3	3/2	2/3	3/2	1/3
L7	1/2	1/3	1/3	1/3	2/3	1/3	2/2	2/3	1/3	3/2	3/2	3/2	2/3	3/2	1/3
L8	1/2	1/2	1/2	3/2	2/3	1/3	3/2	2/3	3/2	3/2	2/3	3/2	3/2	3/2	3/2
L9	1/2	1/3	1/3	2/3	2/3	3/2	2/3/	2/3	3/2	3/2	3/2	3/2	3/2	1/3	3/2
L10	1/1	1/2	2/2	3/2	2/2/	1/2	3/2	3/2	3/2	3/2	1/2	3/2	2/2	3/2	3/2
L11	1/3	1/2	1/3	3/2	2/3	3/2	3/2	3/2	3/2	3/2	3/2	3/2	3/2	2/3	1/3
L12	1/2	1/3	2/3	3/2	1/3	3/2	3/2	2/3	3/2	2/3	1/3	3/2	2/3	3/2	3/2
H1	1/2	2/1	1/1	1/1	3/1	2/1	3/1	1/1	1/1	3/1	3/1	3/1	3/1	3/1	3/1
H2	1/1	2/1	1/1	2/2	1/1	2/1	3/1	3/1	3/1	2/1	3/1	3/1	1/1	2/1	1/1
H3	1/1	1/1	2/1	1/3	3/1	2/1	2/1	3/1	3/1	3/1	3/1	3/1	3/1	2/1	3/1
H4	1/2	1/3	1/3	2/1	3/1	2/1	3/1	1/3	3/1	3/1	1/1	3/1	1/1	3/1	1/3
H5	1/1	1/1	1/1	2/1	2/2	1/1	1/1/	1/1	1/1	3/1	3/1	1/1	2/2	3/1	1/1
H6	1/1	1/1	1/1	3/1	2/2	3/1	1/1	3/1	3/1	3/1	3/1	3/1	3/1	3/1	3/1
H7	1/2	1/3	1/3	1/3	2/1	3/1	2/1	3/1	3/1	1/3	3/1	1/3	3/1	3/1	1/3
H8	1/1	1/3	1/1	1/1	2/1	2/1	2/1	1/1	2/1	1/1	1/1	3/1	3/1	3/1	2/1
H9	1/1	1/1	2/1	3/1	2/1	3/1	3/1	3/1	3/1	2/1	3/1	2/1	2/1	1/1	3/1
H10	1/2	1/2	1/2	3/1	1/2	3/1	3/1	1/1	3/1	3/1	3/1	1/2	1/2	3/1	3/1
H11	1/3	1/3	1/3	2/1	3/1	3/1	2/1	1/1	3/1	3/1	3/1	2/1	2/1	3/1	1/1
H12	1/1	1/1	2/1	3/1	3/1	1/1	3/1	3/1	3/1	3/1	3/1	1/1	1/1	2/1	3/1

In this table, “x/y” indicates menu x and response y, where 1 means “high effort”, 2 means “low effort” and 3 means rejection.

Table A.5
Menus and agent choices in Treatment 2 (250).

	Period														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
L1	1/2	2/2	3/2	2/3	2/3	1/3	3/2	1/3	1/2	2/2	2/2	1/3	2/2	1/2	2/2
L2	2/2	2/2	3/2	2/2	2/2	3/2	2/2	2/2	2/2	2/2	1/3	2/2	1/2	2/2	1/3
L3	3/2	1/2	2/2	1/2	1/2	2/2	2/2	2/2	2/2	2/2	2/2	1/3	2/2	3/2	3/2
L4	3/2	3/2	3/2	3/2	1/2	1/2	2/2	1/2	3/2	3/2	1/2	2/2	3/2	2/2	2/2
L5	3/2	1/2	3/2	2/2	1/2	1/2	1/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	3/2
L6	1/1	1/2	1/2	1/2	1/2	1/2	2/2	1/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2
L7	1/2	2/2	1/2	2/2	2/2	3/2	1/2	1/2	1/2	2/2	3/2	2/2	2/2	3/2	2/2
L8	1/2	1/2	1/2	1/2	2/2	1/2	2/2	2/2	1/2	1/2	2/2	2/2	1/2	2/2	3/2
L9	1/2	2/2	2/2	1/2	1/2	2/2	1/2	1/2	2/2	1/2	2/2	2/2	3/2	2/2	2/2
L10	1/1	2/2	1/2	1/2	1/2	2/2	3/2	3/2	2/2	2/2	2/2	1/2	2/2	2/2	2/2
L11	1/3	2/2	2/2	2/2	2/2	3/2	2/2	3/2	2/2	2/2	1/2	2/2	3/2	2/2	2/2
L12	1/2	1/2	1/2	1/2	1/3	2/2	3/2	1/3	2/2	1/3	2/2	1/3	2/2	2/3	2/2
H1	1/1	2/1	1/1	2/1	1/1	2/1	1/1	1/1	2/1	3/1	3/1	1/1	1/1	1/1	1/1
H2	1/1	2/1	3/1	3/1	1/1	2/1	1/1	3/1	3/1	2/1	3/1	3/1	2/1	2/1	2/1
H3	1/2	1/1	2/1	2/1	3/1	2/1	3/1	3/1	2/1	1/1	2/1	1/1	3/1	2/1	3/1
H4	2/1	3/1	1/1	1/2	3/1	3/1	2/1	1/1	1/1	1/1	2/1	3/1	2/1	3/1	2/1
H5	2/1	1/1	2/1	1/1	2/1	3/1	1/2	1/3	1/2	1/3	1/2	2/1	2/1	2/1	2/1
H6	2/1	1/1	3/1	1/3	2/1	1/3	1/3	1/3	1/2	2/1	1/2	2/1	3/1	2/1	2/1
H7	3/1	2/1	1/1	2/1	1/1	1/1	1/3	1/1	2/1	2/1	3/1	2/1	1/3	2/1	2/1
H8	1/1	1/3	2/1	2/1	1/1	1/1	1/3	1/1	2/1	2/1	2/1	2/1	3/1	2/1	2/1
H9	2/1	3/1	1/3	1/3	2/1	3/1	2/1	3/1	2/1	1/3	2/1	2/1	2/1	2/1	2/1
H10	2/1	3/1	2/1	1/3	1/3	3/1	2/1	1/3	2/1	1/3	2/1	2/1	2/1	2/1	1/1
H11	1/1	2/1	2/1	1/1	2/1	3/1	1/1	1/1	2/1	2/1	1/1	2/1	2/1	2/1	2/1
H12	1/1	1/1	1/1	2/1	2/1	2/1	2/1	3/1	2/1	2/1	2/1	2/1	2/1	2/1	1/1

In this table, “x/y” indicates menu x and response y, where 1 means “high effort”, 2 means “low effort” and 3 means rejection.

Table A.6

Rejection rates by individual agent.

Session	L1		L2		L3		L4	
	M1	M2	M1	M2	M1	M2	M1	M2
L types								
1	2/5	3/4	3/4	4/5	2/3	3/5	6/6	1/1
2	4/7	4/4	5/5	5/5	6/7	3/4	1/4	3/3
3	3/4	4/4	0/4	0/3	3/4	2/2	3/4	4/4
4	3/6	2/7	2/3	0/10	1/4	0/8	0/4	0/4
5	0/4	0/8	0/7	0/8	0/5	0/8	0/8	0/6
6	0/6	0/8	0/5	0/8	1/2	0/10	4/8	1/6
Session	H1		H2		H3		H4	
	M1	M2	M1	M2	M1	M2	M1	M2
H types								
1	0/5	0/2	0/5	0/5	1/3	0/4	4/7	0/2
2	0/9	0/3	0/4	0/1	6/7	0/2	1/7	0/5
3	0/3	0/5	0/7	0/0	3/5	0/4	0/5	0/2
4	0/9	0/4	0/3	0/6	0/4	0/6	0/5	0/5
5	2/7	0/7	4/7	0/6	2/6	0/7	2/6	0/8
6	3/3	0/9	4/5	0/8	0/5	0/9	0/4	0/10

x/y in each cell refers to # of times the agent chose rejection/# of times menu was offered.

Appendix D. Individual agent behavior

Table A.6.

References

- Akerlof, G., 1970. The market for lemons: qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics* 89, 488–500.
- Ballinger, T., Palumbo, M., Wilcox, N., 2003. Precautionary saving and social learning across generations: an experiment. *Economic Journal* 113, 920–947.
- Casadesus-Masanell, R., 2004. Trust in agency. *Journal of Economics & Management Strategy* 13, 375–404.
- Çelen, B., Kariv, S., Schotter, A., 2010. An experimental test of advice and social learning. *Management Science* 56, 1678–1701.
- Charness, G., Corominas-Bosch, M., Frechette, G., 2007. Bargaining and network structure: an experiment. *Journal of Economic Theory* 136, 28–65.
- Charness, G., Dufwenberg, M., 2009. Hidden Information, Participation, and Communication. Mimeo.
- Charness, G., Rabin, M., 2002. Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117, 817–869.
- Chaudhuri, A., Graziano, S., Maitra, P., 2006. Social learning and norms in a public goods experiment with inter-generational advice. *Review of Economic Studies* 73, 257–280.
- Dekel, E., Fudenberg, D., 1990. Rational play under payoff uncertainty. *Journal of Economic Theory* 52, 243–267.
- Eckel, C., Wilson, R., 2007. Social learning in coordination games. *Experimental Economics* 10, 317–330.
- Ellison, G., 2002. Evolving standards for academic publishing: a $q-r$ theory. *Journal of Political Economy* 110, 994–1034.
- Ellison, G., Fudenberg, E., 1993. Rules of thumb for social learning. *Journal of Political Economy* 101, 612–643.
- Fehr, E., Schmidt, K., 1999. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114, 817–868.
- Glasnapp, D., Poggio, J., 1985. *Essentials of Statistical Analysis for the Behavioral Sciences*. Merrill, Columbus.
- Güth, W., Königstein, M., Kovács, J., Zala-Mező, E., 2001. Fairness within firms: the case of one principal and multiple agents. *Schmalenbach Business Review* 53, 82–101.
- Güth, W., Schmittberger, R., Schwarze, B., 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3, 367–388.
- Layard, R., Nickell, S., Jackman, R., 1994. *The Unemployment Crisis*. Oxford University Press, Oxford.
- Mas-Colell, A., Whinston, M.D., Green, J.R., 1995. *Microeconomic Theory*. Oxford University Press, Oxford.
- Mirrlees, J., 1971. An exploration in the theory of optimum income taxation. *Review of Economic Studies* 38, 175–208.
- Rob, R., Zemsky, P., 2002. Cooperation, corporate culture and incentive intensity. *RAND Journal* 33, 243–257.
- Schotter, A., Sopher, B., 2003. Social learning and convention creation in inter-generational games: an experimental study. *Journal of Political Economy* 111, 498–529.
- Schotter, A., Sopher, B., 2007. Advice and behavior in intergenerational ultimatum games: an experimental approach. *Games and Economic Behavior* 58, 365–394.
- Selten, R., 1965. A re-examination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4, 25–55.
- Vickrey, W., 1961. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* 16, 8–37.
- von Siemens, F., 2004. *Social Preferences, Sorting, and Employment Contracts*. Chapter 2 of Ph.D. Dissertation. University of Munich.
- Wilson, R., 1993. *Nonlinear Pricing*. Oxford University Press, Oxford.
- Zafar, B., 2009. An experimental investigation of why people conform. Federal Reserve Bank of New York, Staff Report No. 365.