DISCUSSION PAPER SERIES

No. 5433

SPECULATIVE CONTRACTS

Kfir Eliaz and Rani Spiegler

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP5433.asp

ISSN 0265-8003

SPECULATIVE CONTRACTS

Kfir Eliaz, New York University and CEPR
Rani Spiegler, Tel Aviv University

Discussion Paper No. 5433
January 2006

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 5433

January 2006

ABSTRACT

Speculative Contracts*

We propose to view action-contingent contracts as bets, motivated by different prior beliefs between the contracting parties (rather than, say, as an instrument for overcoming moral hazard problems). Such differences in prior beliefs may arise from inherent biases such as over-optimism. Menus of contingent contracts that arise in principal-agent relationships are thus interpreted as a consequence of the principal's attempt to screen the agent's prior belief. Thus, an employer may offer his worker to choose between fixed-wage and profit-sharing schemes, in order to screen the worker's degree of optimism. We present a model of bilateral contracting which captures these ideas, characterize the optimal menu and apply it to a number of economic settings.

JEL Classification: D42, D82 and L12
Keywords: menus, non-common priors and speculative trade

Kfir Eliaz
Economic Department
New York University
269 Mercer Street
New York, NY 10003
USA

Tel: (1 212) 998 8912
Fax: (1 212) 995 3932
Email: kfir.eliaz@nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=153104

Rani Spiegler
School of Economics
Tel Aviv University
Tel Aviv 69978
ISRAEL

Tel: (972 03) 6405827
Fax: (972 03) 6409908
Email: rani@post.tau.ac.il

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=148706

* Financial support from the US-Israel Binational Science Foundation, Grant No. 2002298 is gratefully acknowledged.

Submitted 06 December 2005

1 Introduction

When designing the terms of a bilateral contract, parties need to take into account differences in opinion regarding the likelihood of future events. Standard models assume that these differences are a result of informational asymmetries. However, they could also be purely differences in prior belief. An entrepreneur seeking to finance a new project may be more optimistic than an investor about the quality of the project; a sales person may be more optimistic about his salesmanship than a prospective employer; an advertising agency and a client may disagree over which type of campaign would be most successful; a project manager and a contractor often have different prior expectations about which variety of a product will be desired by consumers.[1]

In principal-agent relations, the principal may exploit such differences in prior beliefs by writing a contingent contract, which is essentially a bet on the future. The reason parties can agree on the bet is that each of them is willing to make a concession in the contingency that he deems less likely. The traditional view of contingent contracts in economics is either as a tool for screening agents according to their preference types, or as a tool for overcoming moral hazard. However, when the principal and the agent have different prior beliefs, contingent contracts can serve as bets. The question we pose in this paper is, how should such a contract be designed when the agent's prior belief is his private information.

We study a two-period principal-agent contracting model. The principal enables the agent to choose in period 2 from some set of actions, conditional on signing a contract in period 1. If the agent refuses to sign a contract, he chooses some outside option. A contract is a function that specifies a monetary transfer between the principal and the agent for every second-period action. Other than that, we place no restriction on the space of contracts. The agent has quasi-linear vNM utility over action-transfer pairs. However, his utility from actions may either be u or v, depending on a state of nature, which will be revealed in the second period only to the agent. The principal assigns probability p to state u, while the agent assigns probability θ to this state. The principal does not observe θ, but he believes that it is drawn from some distribution on $[0, 1]$. Hence, θ plays the role of the agent's "type". The principal's problem is to design a *menu of contracts* that maximizes his expected profit.

The following situation illustrates the model. Suppose that the principal is a manager and the agent is a sales person. The manager can devise a pay scheme that conditions on the amount of sales that the agent generates. The effort that the agent needs to incur in order to generate sales depends on the state of the world. In one state, there are zero costs. In another state, every dollar of sales is exactly offset by the cost of effort it requires. The manager assigns prior probability $\frac{1}{2}$ to the low-cost state.

If the sales person shared the manager's prior, the manager could offer him a fixed wage contract, that would pay him his reservation payoff 0 in each state, independently of the

[1] For other examples, see Bazerman and Gillespie (1999).

sales volume. Of course, the sales person would accept such a contract. Given the contract, it would be optimal for the sales person to generate the maximal sales volume $a = 1$ in the low-cost state, and zero sales in the high-cost state.[2] The induced expected revenue for the manager (calculated according to his own prior) is $\frac{1}{2}$. In fact, the fixed wage contract maximizes the expected revenue that the manager can attain, given the sales person's prior.

Now suppose that the manager knew that the sales person is more optimistic - specifically, that he assigns probability $\frac{3}{4}$ to the low-state cost. The manager could still offer the above fixed wage contract, which would be accepted by the sales person and generate an expected revenue of $\frac{1}{2}$. However, another possibility would be to sell the project to the sales person, for a price of $\frac{3}{4}$. To see why the sales person would accept this contract, observe that his expected payoff from the contract (according to his own prior) is $\frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 0 - \frac{3}{4} = 0$. From the manager's point of view, the expected revenue is $\frac{3}{4}$, which is better than the wage contract. In fact, it can be shown that it is an optimal contract given a sales person whose type is $\frac{3}{4}$.

Thus, when faced with a sales person who shares his level of optimism, the manager opts for a fixed-wage contract, while offering the ownership contract to a more optimistic sales person. However, so far we constructed the contracts as if the manager knew the sales person's degree of optimism. If he does not know whether the sales person's prior on the low-cost state is $\frac{1}{2}$ or $\frac{3}{4}$, he can ask him to choose between the above pair of contracts. The optimistic type is indifferent between them, whereas the type who shares the manager's belief strictly prefers the fixed-wage contract, because $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 0 - \frac{3}{4} < 0$.

This example highlights some features of our model. First, the optimal menu contains a high-risk scheme, side by side with a scheme that guarantees the agent his reservation value.[3] Second, the sales person ends up taking the same actions under both contracts: $a = 1$ in the low-cost state and $a = 0$ in the high-cost state. Thus, the multiplicity of contracts has nothing to do with designing second-period incentives - its objective is to screen the agent's degree of optimism regarding second-period cost of effort. Third, the menu of contracts consists of the first-best contracts that the manager would offer if he knew the agent's type. This is possible because type $\frac{1}{2}$ exerts no informational externality on type $\frac{3}{4}$.

In Section 4 we characterize the optimal menu of contracts for a large class of utility functions u and v, when the distribution of types is continuous over $[0, 1]$. The key features of this menu are the following.

The optimal menu is the union of two distinct menus. The principal divides the problem of designing the optimal menu into two separate problems: the optimal menu for the agents who assign a higher prior to u, and the optimal menu for the agents who assign a higher prior to v. The two menus can be designed in such a way that the only incentive constraints that the principal needs to worry about, are those within each category: the agents in one category

[2] Note that the sales person is indifferent among all actions in the second period.

[3] The example is special, however, in that each contract leaves one party totally insured. As we shall see, in other settings the menu contains contracts that impose risk on both parties.

have no incentive to pretend to be agents in a different category (e.g., an "optimist" has no incentive to pretend to be a "pessimist").

Speculative and non-speculative contracts. The contracts offered to agents whose beliefs are significantly different from the principal's have a speculative component: these agents earn less than their reservation payoff in the state which they find less likely compared to the principal. Put differently, had these agent shared the principal's prior, they would not have accepted the contracts offered to them. The contract offered to the agents who more or less agree with the principal, involves no speculation: it is the optimal contract when priors are common.

Exclusion of bets. There is no exclusion of types in the usual sense of not transacting with them. But there *is* exclusion in the sense that the principal does not bet with agents whose beliefs are close to his. Moreover, if the agent's payoff in state u is uniformly higher than in state v - in other words, when θ reflects the agent's degree of optimism - the principal chooses to speculate *only with the more optimistic types.*

Ex-post efficiency. The contract chosen by each agent type induces him to choose the efficient action in the state he finds less likely than the principal. In Section 4 we identify a sufficient condition on u and v that guarantees ex-post efficiency in all states. In these cases, all agent types end up choosing precisely the same action in each state, yet they make different payments. This result cannot be rationalized by a model with common priors.

Section 5 provides examples that demonstrate the relevance of our framework in real-life economic settings. First, we study in greater detail the manager-worker example described above, and demonstrate how our model can illuminate the coexistence of fixed-wage and profit-sharing schemes. Next, we show that in a buyer-seller environment, the model can give rise to menus that contain unlimited consumption contracts side by side with contracts that specify quantity-dependent payments. We use this example to interpret deals offered by mobile phone companies, DVD rental stories and restaurants.

Finally, in subsection 5.3 we examine another buyer-seller example, in which the two parties have different priors regarding the buyer's future ideal variety of the seller's product. We use this example to illustrate how speculative contracts might have an adverse effect on ex-post efficiency, such that the larger the difference between the parties' belief, the greater the efficiency loss. In such environments, a central planner who cares about ex-post efficiency would prefer that the principal did not know the agent's prior belief. In this sense, the asymmetric-information environment is socially desirable relative to the complete-information environment, because it mitigates distorting effects of speculative contracting.[4]

Related literature

Our paper is related to the contracting literature that studies non-linear pricing with imperfectly informed consumers (most notably, Baron and Besanko (1984), Armstrong (1996)

[4]However, these effects are not general, and rely on fine details of the buyer's risk attitudes.

and Courty and Li (2000)), a.k.a. "sequential screening". Both this literature and our work study the problem of a monopolist who offers consumers a menu of contingent pricing schedules. In both frameworks, the monopolist's objective is to screen agents according to their unobservable prior beliefs over their future tastes.

However, in contrast to our model, the sequential screening literature assumes common priors. More specifically, it is common knowledge that the agent is better informed about his future tastes than the principal. As we show in Section 3, we focus on an environment where this distinction has important implications. Proposition 1 establishes that non-common priors are necessary to generate price discrimination in our environment, i.e., the optimal menu under common priors is a singleton: there is no discrimination.

Our paper also follows a few recent papers on contract theory with non-common priors. Landier and Thesmar (2004) examined debt contracts that are signed between investors and entrepreneurs who differ in their degree of optimism. Assuming a competitive environment, in which investors earn zero profits, the authors construct a separating equilibrium in which entrepreneurs who are more optimistic than the investor choose short-term debt, while entrepreneurs who share the investor's belief choose long-term debt. In relation to our model, short-term debt may be interpreted as a bet in which the entrepreneur concedes cash flow rights in the low state in return for claims on the good state. The authors also provide empirical evidence suggesting that short-term debt is correlated with optimistic expectation errors of entrepreneurs.

Moscarini and Fang (2005) also study the implication of non-common priors on the design of wage contracts. Unlike us, they analyze contract design with an informed principal. More specifically, they ask how should a principal design a wage contract when he holds the correct prior about his workers' ability, whereas their priors are biased upwards. A key assumption in the paper is that an optimistic belief has a positive effect on a worker's productivity. The principal, therefore, faces the following trade-off. On the one hand, he would like to provide the appropriate monetary incentives for his workers to exert effort. On the other hand, he is concerned that workers may infer their true ability from the contract he offers (what the authors call "morale hazard"). Hence, in contrast to the present paper, Moscarini and Fang (2005) focus on the signaling aspect of contracts.

Sandroni and Squintani (2004) modify the Rothsechild-Stiglitz insurance market model, to allow for consumers who are over-optimistic regarding their probability of an accident. They show that although in equilibrium these consumers are under-insured, compulsory insurance need not be Pareto-improving.

Finally, this paper extends our own previous work. Eliaz and Spiegler (2004) study optimal contract design with dynamically inconsistent agents. An agent type is his degree of naivete, modeled as his prior belief that his current preferences will change in the future. While the agent draws his belief from some distribution, the principal believes that the agent's preferences are sure to change. As in the present paper, the principal's objective in

Eliaz and Spiegler (2004) is to screen agents according to their prior beliefs. However, time-inconsistency has important implications for the design of the optimal menu. In particular, sophisticated types who believe that their tastes will change with high probability are assigned a contract that serves as a perfect commitment device: it induces them to choose the action that maximizes their current utility.

2 The model

A principal offers an agent the opportunity to choose an action from the set $[0,1]$. However, in order to have access to this set of actions, the agent must sign a contract with the principal one period beforehand. If the agent does not sign a contract with the principal, he chooses some outside option. We refer to the period in which a contract is signed as period 1, and to the period in which the action is chosen as period 2. A contract is a function $t:[0,1]\to\mathbb{R}$ that specifies for every second-period action, a (possibly negative) transfer from the agent to the principal. The principal is perfectly able to monitor the agent's second-period action.

The agent has quasi-linear preferences over action-transfer pairs. We assume that his net utility in period 1 from the outside option is zero. However, his preferences over second period actions depend on the state of nature. There are two possible states: in state u the agent's preferences are represented by the continuous function $u:[0,1]\to\mathbb{R}$, and in state v they are represented by the continuous function $v:[0,1]\to\mathbb{R}$. We assume that these functions have the property that one lies above another at a point in which the former attains a global maximum. Denote $u^*=\max_{a\in[0,1]}u(a)$ and $v^*=\max_{a\in[0,1]}v(a)$. We assume the u and v satisfy the following properties:

Condition 1 *There exist $a_u^*\in\arg\max_{a\in[0,1]}u(a)$ and $a_v^*\in\arg\max_{a\in[0,1]}v(a)$, such that $u(a_u^*)\geq v(a_u^*)$ and $v(a_v^*)\geq u(a_v^*)$ with at least one strict inequality.*

Condition 2 $u^*,v^*\geq 0$

We discuss the methodological contribution of Condition 1 in the next section. Condition 2 simply guarantees that there would be a reason for the two parties to write a contract, even when priors are common.

The agent believes that state u occurs with probability θ, while the principal believes that this probability is p. These are pure differences in prior opinion. We assume that it is common knowledge that neither party is better informed about the state of nature. When $u(a)>v(a)$ for every a, it makes sense to refer to an agent with a higher θ as a more "optimistic" agent. The value of θ is not known to the principal. However, he believes this value to be distributed on $[0,1]$ according to a continuous c.d.f. $F(\theta)$.

The principal's objective is to maximize expected revenue.[5] Because the principal does not observe θ, he offers the agents a menu of contracts, where a menu is set of transfer functions $t : [0,1] \to \mathbb{R}$. Given a menu, each agent type θ computes his indirect utility from each contract in the menu,

$$\theta \max_{a\in[0,1]} [u(a) - t(a)] + (1-\theta) \max_{a\in[0,1]} [v(a) - t(a)] \tag{1}$$

and then picks the contract with the highest indirect utility. The principal's problem is to design a menu of contracts that maximizes his expected revenue, taking into account that each agent type would select his most desired contract from this menu.

By the revelation principle, a solution to this problem may be obtained via a direct revelation mechanism in which agents are asked to report their type, and each reported type ϕ is assigned a contract $t_\phi : [0,1] \to \mathbb{R}$. The optimal menu of contracts $\{t_\theta(a)\}_{\theta\in[0,1]}$ is given by the solution to the following maximization problem:

$$\max_{\{t_\theta(a)\}_{\theta\in[0,1]}} \int_0^1 [pt_\theta(a_\theta^u) + (1-p)t_\theta(a_\theta^v)]dF(\theta)$$

subject to the constraints,

$$\theta [u(a_\theta^u) - t_\theta(a_\theta^u)] + (1-\theta)[v(a_\theta^v) - t_\theta(a_\theta^v)] \geq 0 \tag{IR_θ}$$

$$\theta [u(a_\theta^u) - t_\theta(a_\theta^u)] + (1-\theta)[v(a_\theta^v) - t_\theta(a_\theta^v)] \geq \theta \left[u\left(a_\phi^u\right) - t_\phi(a_\phi^u)\right] + (1-\theta)\left[v\left(a_\phi^v\right) - t_\phi(a_\phi^v)\right] \tag{$IC_{\theta,\phi}$}$$

for all $\phi \in [0,1]$, where

$$\begin{aligned} a_\theta^u &= \arg\max_{a\in A} [u(a) - t_\theta(a)] && (UR_\theta) \\ a_\theta^v &= \arg\max_{a\in A} [v(a) - t_\theta(a)] && (VR_\theta) \end{aligned}$$

The first and second constraints are the standard individual rationality and incentive compatibility constraints. Condition IR_θ says that an agent of type θ is at least as well off with his assigned contract than with the default option. Condition $IC_{\theta,\phi}$ says that an agent of type θ cannot be better off by pretending to be of type ϕ and signing the contract assigned to that type.

The novel conditions are UR_θ and VR_θ. These conditions represent the fact that an agent's indirect utility from a contract is determined by the actions he would choose in the two states. If the realized state in period 2 is u (an event to which the agent assigns a probability of θ), then he will choose the optimal action for him according to the utility function u. This is represented by UR_θ. If, on the other hand, the state in period 2 is v

[5] The model can easily be extended to accommodate costs for providing each action. Since this would complicate the analysis with no qualitative effect on our results, we present the analysis for zero costs.

(an event to which the agent assigns a probability of $1-\theta$), then he will choose the optimal action for him according to the utility function v. This is precisely the condition VR_θ.

It follows that any contract t can be identified with a pair of actions: a_θ^u and a_θ^v. The former action is consistent with u-maximization in the second period, while the latter action is consistent with v-maximization in the second period. Without loss of generality, we may assume that $t(a) = +\infty$ for every $a \notin \{a_\theta^v, a_\theta^u\}$.

The constraints IR_θ and $IC_{\theta,\phi}$ can be written more compactly by introducing the following notation. Let $U(\phi,\theta)$ denote the utility of a type θ agent who pretends to be of type ϕ, i.e.,

$$U(\phi,\theta) \equiv \theta\left[u\left(a_\phi^u\right) - t_\phi(a_\phi^u)\right] + (1-\theta)\left[v\left(a_\phi^v\right) - t_\phi(a_\phi^v)\right]$$

Then IR_θ and $IC_{\theta,\phi}$ can be rewritten as $U(\theta,\theta) \geq 0$ and $U(\theta,\theta) \geq U(\phi,\theta)$ for all θ and ϕ.

3 A benchmark: common priors

Before we proceed to analyze the solution to the principal's problem in our model, it is instructive to consider a "benchmark" in which the principal and the agent agree on the probability of each state. There are two candidates for such a benchmark depending on which side is believed to be better informed. Since we are interested in understanding how non-common priors affect the screening motives of a monopolist, we consider the benchmark to be a situation in which the principal believes that state u occurs with probability θ. As in our model, the principal does not know the value of θ, but he believes that this value is distributed on $[0,1]$ according to $F(\theta)$.

This benchmark is essentially the framework used in the sequential-screening literature mentioned in the Introduction. This literature focused on the case in which a higher type corresponds to either a higher expected willingness to pay (for each action), or to a higher variance in the willingness to pay (for each action). In other words, the agents' types can be ordered according to first- or second-order stochastic dominance. The optimal menu in such an environment typically contains more than one contract. More importantly, this menu is ex-post inefficient in the following sense: (i) there is a set of low types which are excluded, and (ii) all types (but the highest one) do not choose the action that maximizes their utility in each state. In contrast, the environment we focus on leads to the following result.

Proposition 1 *If u and v satisfy Conditions 1 and 2, it is optimal for the principal to offer a single, ex-post efficient contract.*

Thus, under common priors, the principal cannot do better than to offer a contract that charges u^* if $a = \arg\max u$ and v^* if $a = \arg\max v$ (and an arbitrarily large amount for any other action). In particular, if $u^* = v^*$, the principal can simply "sell the project" to the agent in return for an up-front payment of u^*. As we show in the next section, when the principal and the agent have different priors, both agents would strictly prefer an action-contingent contract, even when $u^* = v^*$.

We are now in a better position to discuss the importance of Condition 1. This condition allows us to isolate the effect of non-common priors on contract design. As we have shown, when this condition is satisfied in a world with common priors, it is optimal for the principal to offer only a single, non-contingent contract. As we shall show in the next section, if the principal and the agent hold different beliefs, then typically the optimal menu would consist of a number of contracts that discriminate between agents with diverse beliefs. In environments that fails to satisfy Condition 1, it would be difficult to disentangle the price discrimination that is due to speculation and that which arises for standard reasons.

4 The optimal menu with non-common priors

4.1 Qualitative features of the menu

In this section we characterize the optimal menu of contracts, when the principal's prior on u is p and the agent's prior on u is θ, and the principal believes that θ is drawn from a distribution F on $[0, 1]$. We begin by illustrating the qualitative features of the optimal menu with a simple stylized example of a "backup agreement" between a supplier and a retailer.

Consider a retailer who buys from a supplier one unit of a good that is made up of two components, labeled U and V. The retailer has to decide the proportion of each component in the unit that he orders. Let $a \in [0, 1]$ denote the proportion of component U in the good. The retailer's revenue as a function of a depends on the state of nature: in state u the revenue is given by $u(a) = a$, while in state v it is given by $v(a) = 1 - a$. The supplier believes that each state is equally likely, hence, he wishes to maximize the sum of transfers that he receives in each state. The retailer, on the other hand, holds one of three possible beliefs: $\theta \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$.

Suppose the supplier knows that he is facing a retailer of type $\theta = \frac{1}{4}$. The supplier then wishes to maximize the sum of transfers, $t_{\frac{1}{4}}(a^u) + t_{\frac{1}{4}}(a^v)$, subject to two constraints: (i) given these transfers, the retailer would choose an action that maximizes his net payoff in each state, and (ii) the retailer weakly prefers to sign the contract. The first constraint represents the UR and VR requirements mentioned in the previous section. In this example, they imply:

$$-(a^u - a^v) \leq t_{\frac{1}{4}}(a^u) - t_{\frac{1}{4}}(a^v) \leq a^u - a^v$$

Notice that to relax this constraint as much as possible, the supplier can offer an ex-post efficient contract in which $a^u = 1$ and $a^v = 0$. Given this, the second constraint implies that the expected transfer - according to the agent's beliefs - is equal to the highest willingness to pay of the agent in each of the states, i.e. to one:

$$\frac{1}{4}t_{\frac{1}{4}}(a^u) + \frac{3}{4}t_{\frac{1}{4}}(a^v) = 1$$

It follows that it is optimal for the supplier to offer the following two-part tariff:[6]

$$t_{\frac{1}{4}}(a) = \frac{3}{4} + a$$

The interpretation is that the retailer buys one unit that contains only the V component at a price of $\frac{3}{4}$, and any amount of V can later be exchanged to U at a rate of one.

Suppose next that the supplier knew he was facing a retailer of type $\theta = \frac{3}{4}$. By essentially the same argument given above for type $\theta = \frac{1}{4}$, it is optimal for the supplier to offer this retailer a contract,

$$t_{\frac{3}{4}}(a) = \frac{3}{4} + (1 - a)$$

interpreted as an agreement to purchase a unit having only the U component at a price of $\frac{3}{4}$, and charging a conversion rate of one on any subsequent substitution of U into V.

Finally, suppose the supplier knew he was facing a retailer who believes that each state is equally likely. By the $IR_{\frac{1}{2}}$ constraint, a retailer of this type is not willing to trade-off a loss in one state for a smaller gain in another. Hence, the supplier cannot extract from this retailer more than a surplus of one in each state. This can be achieved with a contract that charges $t(a) = 1$ for all $a \in [0, 1]$.

Now assume the supplier cannot identify the retailer's type. Suppose that he invites him to choose among the above three contracts: $\{t_{\frac{1}{4}}, t_{\frac{1}{2}}, t_{\frac{3}{4}}\}$. All contracts induce the retailer to choose $a = 0$ in state u and $a = 1$ in state v, hence each of the contracts is ex-post efficient. However, each retailer type weakly prefers his designated contract to the other contracts in the menu. For $\theta = \frac{1}{4}$, the expected payoff from both $t_{\frac{1}{4}}$ and $t_{\frac{1}{2}}$ is zero, while the expected payoff from $t_{\frac{3}{4}}$ is $-\frac{1}{2}$. Similarly, for $\theta = \frac{3}{4}$, the expected payoff from both $t_{\frac{3}{4}}$ and $t_{\frac{1}{2}}$ is zero, while the expected payoff from $t_{\frac{1}{4}}$ is $-\frac{1}{2}$. Finally, for $\theta = \frac{1}{2}$, the expected payoff from $t_{\frac{1}{2}}$ is zero, while the expected payoff from both $t_{\frac{1}{4}}$ and $t_{\frac{3}{4}}$ is $-\frac{1}{4}$. It follows that the above menu is optimal. The fact that all three types display the same behavior in period 2, while ending up making different payments, cannot be replicated by a standard model with a common prior.

This example illustrates the key features of the optimal menu summarized in the Introduction. First, the optimal menu is the union of two separate menus: the optimal menu for types who assign a weakly higher prior on u, and the optimal menu for types who assign a weakly higher prior on v. The agent types belonging to one category exert no informational externality on the types belonging to the other category. Specifically, type $\frac{1}{4}$ strictly prefers his contract to the contract selected by type $\frac{3}{4}$, and vice versa.

[6] To rule out the possibility that the retailer is indifferent between all actions in state u, we can simply set $t(a)$ to be arbitrarily high for interior values of a.

The fact that each type is assigned the same contract he would be offered under perfect information is only due to the fact that in the example there is a single type $\theta > p$ and a single type $\theta < p$. As we shall see in the next subsection, when there are multiple agent types in each category, incentive constraints within each category may prevent the principal from attaining the same revenue he would attain under perfect information.

Second, the contracts signed by the types $\theta \neq p$ are essentially bets. Each of these types accepts a loss in the state which he deems less likely than the principal, in return for a gain in the other state. The payoffs that these types expect to obtain from their contracts are "speculative", in the sense that they would not have signed these contracts had they shared the principal's beliefs. Similarly, the principal would not have offered these contracts had he held the agents' priors.

In contrast, the contract signed by the type $\theta = p$ involves no speculation: it guarantees a payoff of one, regardless of the agent's prior. While the principal transacts with all agent types, he does not "bet" with the type who shares his prior belief. Note that the non-speculative contract in our example would be an optimal contract for *all* agent types in the benchmark model, in which it is commonly known that the agent's type is the true probability of u.

Another important feature of the optimal menu concerns ex-post efficiency. The contract signed by the agent who assigns a higher prior to u (v) induces him to choose the efficient action in v (u). Our example has the more special feature that all contracts in the menu are ex-post efficient, such that all types behave identically in period 2 yet pay different amounts. The next subsection provides a sufficient condition for this feature of the optimal menu.

Finally, the example demonstrates the following property of the optimal menu, which will turn out to be useful in characterizing the menu. An agent with a *higher* prior on u than the principal, who signs a speculative contract, strictly prefers a^u to a^v in state u, yet he is indifferent between a^u and a^v in state v. Similarly, an agent with a *lower* prior on u than the principal, who signs a speculative contract, strictly prefers a^u to a^v in state v, yet he is indifferent between a^u and a^v in state u. Put differently, the VR constraint is binding for speculative agents who are more confident than the principal about state u, while the UR constraint is binding for speculative agents who are more confident than the principal about state v.

4.2 Characterization

Our goal in this subsection is to provide a complete characterization of the optimal menu of contracts. In the process of reaching this goal, we shall uncover some interesting properties of this menu.

In the example of the previous subsection, the agent types who held different beliefs than the principal were offered "speculative contracts", which they would never accept had they shared the principal's prior. Formally, we define such contracts as follows.

Definition 1 *A contract t is speculative if*

$$p\left[u\left(a^{u}\right)-t\left(a^{u}\right)\right]+(1-p)\left[v\left(a^{v}\right)-t\left(a^{v}\right)\right]<0 \tag{2}$$

The optimal menu in our example included both speculative and non-speculative contracts. The non-speculative contract was the optimal contract to offer to an agent who shares the principal's beliefs. Our first result establishes that this property also extends to the case with a continuum of agent types.

Lemma 1 *If the optimal menu includes a non-speculative contract, then that contract must be the first-best contract for an agent of type $\theta = p$.*

Lemma 1 implies that w.l.o.g. the optimal menu may include only one non-speculative contract. This contract is ex-post efficient and extracts from the agent $\max_a u(a)$ in state u and $\max_a v(a)$ in state v. Hence, an agent who signs this contract expects to obtain his reservation utility *regardless of his prior.*

The intuition for this result is as follows. The revenue that the principal expects to obtain from any contract t is $pt\left(a^{u}\right)+(1-p)t\left(a^{v}\right)$, where $a^u \in \arg\max_a[u(a) - t(a)]$ and a^v is similarly defined. By definition, if t is non-speculative, then this expected revenue is at most $pu\left(a^{u}\right)+(1-p)v\left(a^{v}\right)$. This expression is maximized at $a^u = \arg\max u$ and $a^v = \arg\max v$. Therefore, the following contract t^* is the optimal non-speculative contract:

$$t^{*}(a)=\begin{cases}\max(u) & if \quad a \in \arg\max_a(u) \\ \max(v) & if \quad a \in \arg\max_a(v) \\ \infty & if \quad other\end{cases}$$

By Condition 1, this contract satisfies the UR and VR constraints of all agent types. This is the optimal contract to offer to an agent of type $\theta = p$. If the principal's menu included non-speculative contracts that generate a lower expected revenue, the principal could omit them from the menu, without causing agent types to switch from speculative contracts to t^*. The reason is that t^* yields an expected payoff of zero to all types, hence the constraint that forbids an agent type to prefer t^* to his designated contract is indistinguishable from his IR constraint.

The optimal menu in our example contained both speculative and non-speculative contracts. A natural question that arises is whether this is a general feature of an optimal menu. By IR_p, type $\theta = p$ cannot be assigned a speculative contract. The question is, does an optimal menu necessarily include a speculative contract?

Lemma 2 *If $0 < p < 1$, the optimal menu contains at least one speculative contract. If $p = 0$, the optimal menu includes at least one speculative contract if and only if $u(a) > v(a)$*

for some a. If $p = 1$, the optimal menu includes at least one speculative contract if and only if $u(a) < v(a)$ for some a.

To understand the intuition for this result, consider how a speculative contract exploits the disagreement between the principal and the agent. For the sake of the argument, let $p = 0$. If $u(a) > v(a)$ for some a, then the principal can offer a contract that charges $u(a) - \varepsilon$ for choosing a, but charges $v^* + u(a) - v(a) - \varepsilon$ for choosing any action that maximizes v (and an infinite amount for any other action). Any agent who accepts this contract, gains a surplus of ε in state u, but loses a surplus of $u(a) - v(a) - \varepsilon$ in state v. High-θ types weight the gain more than the loss, and hence prefer this contract to a non-speculative contract that leaves them with zero expected payoff. For the principal, who believes that state v is certain to occur, this contract generates more surplus than the agent's highest willingness to pay in state v. Hence, he prefers this contract to the non-speculative contract. If, however, $v(a) \geq u(a)$ for all a, the principal cannot entice optimistic types by offering them an "imaginary win" in state u in return for a "real fine" in state v.

If the optimal menu contains both speculative and non-speculative contracts, which agent types are assigned the former and which are assigned the latter? Our next result provides a simple answer.

Lemma 3 *There exists a pair of types, $\underline{\theta} \in [0, p)$ and $\bar{\theta} \in (p, 1]$ such that: (i) t_θ is non-speculative for every $\theta \in (\underline{\theta}, \bar{\theta})$, and if $\underline{\theta} = 0$ then t_0 is also non-speculative, and similarly if $\bar{\theta} = 1$ then t_1 is non-speculative; (ii) t_θ is speculative for every $\theta < \underline{\theta}$; (iii) t_θ is speculative for every $\theta > \bar{\theta}$.*

This result stems from ordinary single-crossing arguments. The contract offered to $\theta = p$ must be non-speculative, by IR_p. Suppose that the optimal menu assigns a speculative contract to some type $\theta > p$, but assigns a non-speculative contract to a type $\phi > \theta$. By Definition 1,

$$
\begin{aligned}
p[u(a^u_\phi) - t(a^u_\phi)] + (1-p)[v(a^v_\phi) - t(a^v_\phi)] &\geq 0 \\
p[u(a^u_\theta) - t(a^u_\phi)] + (1-p)[v(a^v_\theta) - t(a^v_\theta)] &< 0
\end{aligned}
$$

Because $\theta > p$, these inequalities imply that $IC_{\theta,\phi}$ is violated. A similar argument shows that if type $\theta' < p$ is assigned a speculative contract, then every type $\phi' < \theta'$ must also be assigned such a speculative contract.

Lemma 3 implies that the "threshold" types, $\bar{\theta}$ and $\underline{\theta}$, should be indifferent between a speculative contract and a non-speculative contract. Since, by Lemma 1, a non-speculative contract gives a zero expected payoff to all types, the speculative contract offered to $\bar{\theta}$ and $\underline{\theta}$ should also give these types an expected payoff of zero. This implies the following result.

Corollary 1 *$IR_{\underline{\theta}}$ and $IR_{\bar{\theta}}$ must be binding.*

Recall that in the example of the previous subsection, the principal chose to speculate with types whose priors lie on both sides of his own prior. In some situations, however, in some cases the principal speculates only with agents whose priors lie on one side of some threshold type.

Proposition 2 *If $u(a) \geq v(a)$ for all a, then $\underline{\theta} = 0$. Similarly, if $v(a) \geq u(a)$ for all a, then $\bar{\theta} = 1$.*

The meaning of this result is that if payoffs in one state are always higher than in another state - such that θ may be viewed as the agent's degree of optimism - then the principal chooses to speculate only with the agents who are more optimistic than he is. Put differently, if it is possible to interpret θ as a degree of optimism, the principal will never choose to speculate with the pessimists.

Another distinctive feature of the example of the previous subsection is that the speculative contract designed for type $\theta > p$ is designed independently of the contract designed for $\theta < p$. Our next result establishes that this is a general property of the optimal menu.

Lemma 4 *The optimal menu is the union of two sets of menus: (i) the optimal menu for the distribution F conditional on the restricted support $[p, 1]$, and (ii) the optimal menu for the distribution F conditional on the restricted support $[0, p]$.*

The essence of Lemma 4 is that the speculative contracts offered to types $\theta < p$ do not exert an informational externality on types $\theta > p$ who speculate, and vice versa. Hence, the only incentive constraints that the principal needs to worry about are those that prevent agents whose prior is on one side of p to misrepresent the proximity of their belief to the principal's.

To see the intuition for this, consider two types $\phi < p < \theta$ who are both assigned speculative contracts. For any type x, let $\Delta_x^u = u(a_x^u) - t(a_x^u)$ and $\Delta_x^v = u(a_x^v) - t(a_x^v)$. By Definition 1,

$$\begin{aligned} p(\Delta_\phi^u) + (1-p)\Delta_\phi^v &< 0 \\ p(\Delta_\theta^u) + (1-p)\Delta_\theta^v &< 0 \end{aligned}$$

By IR_θ and IR_ϕ,

$$\begin{aligned} \phi(\Delta_\phi^u) + (1-\phi)\Delta_\phi^v &\geq 0 \\ \theta(\Delta_\theta^u) + (1-\theta)\Delta_\theta^v &\geq 0 \end{aligned}$$

Therefore,

$$\begin{aligned}\theta(\Delta_{\phi}^{u})+(1-\theta)\Delta_{\phi}^{v} &< 0\\ \phi(\Delta_{\theta}^{u})+(1-\phi)\Delta_{\theta}^{v} &< 0\end{aligned}$$

Hence, each type has no incentive to pretend to be the other type. This is essentially a single-crossing argument.

Lemma 4 simplifies the derivation of the optimal menu in that it breaks it down into three separate problems: (i) solving for the optimal menu for $\theta \geq \bar{\theta}$, (ii) solving for the optimal menu for $\theta \leq \underline{\theta}$, and (iii) solving for the non-speculative contract for $\underline{\theta} \leq \theta \leq \bar{\theta}$. By Lemma 1, the non-speculative contract is immediate. The question is, how does one design the optimal menus of speculative contracts? In standard models of price discrimination a menu consists of *pairs* of numbers: quantity-price, quality-price or a probability of obtaining a good and an expected payment. In contrast, a menu in our model consists (w.l.o.g.) of *two pairs* of numbers: $(a_{\theta}^{u}, t_{\theta}^{u})$ and $(a_{\theta}^{v}, t_{\theta}^{v})$. We, therefore, need to further simplify our problem in order to be able to apply standard tools of optimal price discrimination. The first step in this simplification is to show that w.l.o.g. we may fix a_{θ}^{v} for $\theta \geq p$ and fix a_{θ}^{u} for $\theta \leq p$.

Lemma 5 *W.l.o.g. we may restrict attention to an optimal menu in which* $a_{\theta}^{v} \in \arg\max_{a} v(a)$ *for* $\theta \geq p$*, and* $a_{\theta}^{u} \in \arg\max_{a} u(a)$ *for* $\theta \leq p$*.*

Lemma 5 implies that the agent chooses the ex-post efficient action in the state he deems *less* likely than the principal. The intuition for this is that a type $\theta > p$ (respectively, $\theta < p$) is willing to sacrifice a dollar in state v (respectively, u) for a dollar in state u (respectively, v), whereas the principal has the opposite preferences. To be able to increase the revenue in the state on which he puts more weight, without disrupting any of the constraints (IC, IR, UR and VR), the principal can increase the utility of the agent by the same amount. The highest revenue increase is thus obtained when the agent's utility reaches its maximum.

As their name suggests, speculative contracts may be interpreted as bets between the principal and the agent. Consider the speculative contract assigned to agent type $\theta > \bar{\theta}$. We may interpret this contract as the following bet. The agent gives the principal an up-front payment of $-\Delta_{\theta}^{v}$. If state u is realized, the principal pays the agent the amount Δ_{θ}^{u}. We may thus interpret the difference, $\Delta_{\theta}^{u} - \Delta_{\theta}^{v}$, as the "speculative gain" of the agent from the bet.

Let $q(\theta) \equiv \Delta_{\theta}^{u} - \Delta_{\theta}^{v}$. The principal's expected profit from the bet is $-pq(\theta) - \Delta_{\theta}^{v}$, while the agent's expected payoff is $\theta q(\theta) + \Delta_{\theta}^{v}$. The two parties would agree to this bet if, and only if, both earn non-negative payoff. For this to be true, Δ_{θ}^{v} must be non-positive, while $q(\theta)$ must be non-negative. If t_{θ} is a speculative contract, then by Definition 1, the principal must be earning a strictly positive profit from the bet. Thus, in order for the agent to accept the bet, $q(\theta) > 0$. Applying the same argument to types $\theta \leq \underline{\theta}$, we obtain the following result.

Lemma 6 *$q(\theta) > 0$ for all $\theta \geq \bar{\theta}$ and $q(\theta) < 0$ for all $\theta \leq \underline{\theta}$.*

A standard technique in the mechanism design literature is to transform the incentive-compatibility constraints into an integral representation of $U(\theta, \theta)$ (the expected utility of type θ from truthfully reporting his type). Lemma 6 is instrumental in adapting this technique to our framework. To do this, we rewrite $U(\theta, \theta)$ as follows,

$$U(\theta, \theta) = \begin{cases} \theta q(\theta) + \Delta_\theta^v & if \quad \theta \geq p \\ (1-\theta)[-q(\theta)] + \Delta_\theta^u & if \quad \theta < p \end{cases} \tag{3}$$

By Lemma 6, we may apply standard arguments to obtain the following result.

Lemma 7 *Assume the optimal menu assigns a speculative contract to type θ. If this contract satisfies $IC_{\theta,\phi}$ for all θ and ϕ, then*

$$U(\theta, \theta) = \begin{cases} \int_{\bar{\theta}}^{\theta} q(x)\, dx & if \quad \theta \geq \bar{\theta} \\ \int_{1-\underline{\theta}}^{1-\theta} [-q(x)] dx & if \quad \theta \leq \underline{\theta} \end{cases} \tag{4}$$

Recall that in our example of the previous sub-section, type $\theta > p$ was indifferent between a_θ^v and a_θ^u in state v, while type $\theta < p$ was indifferent between a_θ^v and a_θ^u in state u. That is, agent types who signed speculative contracts were indifferent between the two actions in the state they deemed less likely than the principal. Our next result establishes that this is a general property of the optimal menu.

Lemma 8 *The VR_θ constraint is binding for all types $\theta \geq \bar{\theta}$, while the UR_θ constraint is binding for all types $\theta \leq \underline{\theta}$.*

Lemmas 5 and 8 reduce the number of decision variables per type to only *two*: (a_θ^v, t_θ^v) for $\theta \leq \underline{\theta}$ and (a_θ^u, t_θ^u) for $\theta \geq \bar{\theta}$. Thus, we may apply standard mechanism-design tools to derive the optimal speculative menu - independently for the ranges $[0, \underline{\theta}]$ and $[\bar{\theta}, 1]$, by Lemma 4. For this we shall need the following notations. Let $z = 1 - x$ and define $G(z) = 1 - F(1 - x)$. Define

$$\psi(x) = x - \frac{1 - F(x)}{f(x)} \tag{5}$$

and define $\psi(z)$ accordingly.

Proposition 3 *If F satisfies the monotone hazard rate condition, then a speculative contract intended for a type $\theta > p$ induces the actions,*

$$\begin{aligned} a_\theta^v &\in \arg\max_a v(a) \\ a_\theta^u &\in \arg\max_a \left[(\psi(\theta) - p) \cdot (u(a) - v(a)) + p \cdot u(a) \right] \end{aligned} \tag{6}$$

If G satisfies the monotone hazard rate condition, then a speculative contract intended for a type $\theta < p$ induces the actions:

$$\begin{aligned} a_\theta^u &\in \arg\max_a u(a) \\ a_\theta^v &\in \arg\max_a \left[(\psi(1-\theta) - (1-p)) \cdot (v(a) - u(a)) + (1-p) \cdot v(a) \right] \end{aligned} \tag{7}$$

The speculative contracts in the optimal menu induce an *inefficient* action in the state, which is assigned a higher prior by the agent than by the principal. To interpret this action, consider the case of $p = 0$. The principal would like the agent to choose the efficient action in state v (the state he believes with certainty), but he could exploit his disagreement with the agent to charge more than $\max v$ for this action.

To do this, imagine that in period 1 the principal offered a menu of the following contracts. Each contract has three components: an action a^u, an up-front payment t_θ, and a changing fee of τ_θ. The contracts in the menu differ only in the last two components. An agent who signs a contract $(a^u, t_\theta, \tau_\theta)$ pays t_θ and commits to choosing a^u. However, if in period 2 the agent wants to switch to $\arg\max v$, he may do so only after paying the changing fee τ_θ.

The principal believes that the agent would pay the fee for sure, while the agent believes that this would happen only with probability θ. The principal's objective is to have $t_\theta + \tau_\theta > \max v$, and as high as possible. To meet this objective, the principal should choose a^u such that the agent would be willing to pay as much as possible in state v to switch to $\arg\max v$. Hence, $a^u \in \arg\max_a [u(a) - v(a)]$.

This intuition extends to the case of $p \in (0,1)$. This case is a little more complicated. First, we need to distinguish between optimistic and pessimistic types. The former pay the "changing fee" in state v, while the latter pay this fee in state u. Second, since the principal assigns positive probability to both states, the action a^u for the optimistic types maximizes a weighted average of $u - v$ and u. Similarly, a^v for the pessimistic types maximizes a weighted average of $v - u$ and v.

Notice that if the principal observed the agent's type, then in order to derive a_θ^u in the first-best contract for a type $\theta > p$, we would simply need to replace $\psi(\theta)$ with θ in (6). Likewise, we would replace $\psi(1-\theta)$ with $1 - \theta$ in (7) when deriving a_θ^u in the first-best contract for a type $\theta < p$.

To complete our characterization of the optimal menu, it remains to derive the cutoffs, $\bar{\theta}$ and $\underline{\theta}$. Recall that by Proposition 2, if $u(a) \geq v(a)$ (respectively, $v(a) \leq u(a)$) for all a, then the principal weakly prefers not to speculate with any of the types above p (below p), in which case $\bar{\theta} = 1$ ($\underline{\theta} = 0$). The following proposition characterizes the cutoffs for the case in which $u(a) > v(a)$ for some a and $v(a) > u(a)$ for some a.

Proposition 4 *(i) Suppose that F satisfies the monotone hazard rate condition, and $u(a) > v(a)$ at some a. Then, $\bar{\theta}$ is the unique solution to $\psi(\theta) = p$.*

(ii) Suppose G satisfies the monotone hazard rate condition, and $v(a) > u(a)$ at some a. Then, $\underline{\theta}$ is the unique solution to $\psi(1-\theta) = 1-p$.

Our next collection of results addresses the question of ex-post efficiency of the optimal menu. We say that the optimal menu is ex-post efficient if $a^u_\theta \in \arg\max u$ and $a^v_\theta \in \arg\max v$ for every type θ.

Proposition 5 *If the optimal menu is ex-post efficient, then there exists such a menu with at most one speculative contract for types higher than p, and at most one such contract for types lower than p.*

If the optimal menu is ex-post efficient, then w.l.o.g. all types can be made to choose precisely the same action in each state ($a^u \in \arg\max u$ and $a^v \in \arg\max v$). But since each optimistic type is indifferent between a^u and a^v in state v, while each pessimistic type is indifferent between these two actions in state u (recall Lemma 8), the absolute value of the "speculative gain", $|q(\theta)|$ is constant for *all* types. Hence, the principal can only discriminate between types above p and types below p, but he cannot discriminate among the types within each of these groups.

The following result provides a sufficient condition for ex-post efficiency.

Proposition 6 *If $\arg\max u \cap \arg\max(u-v) \neq \varnothing$ and $\arg\max v \cap \arg\max(v-u) \neq \varnothing$, the optimal menu is ex-post efficient .*

The intuition for this result is simple. Consider the speculative contracts that the principal designs for types $\theta > p$. By Lemma 5, a^v_θ is ex-post efficient. In our discussion below Proposition 3, we saw that when designing a^u_θ, the principal attempts to compromise between maximizing u and maximizing $u-v$. When the same action maximizes both u and $u-v$, the principal can set a^u to be ex-post efficient without having to worry about this trade-off. Note that the sufficient condition identified in the proposition is satisfied in our example of the previous sub-section.

5 Examples

In this section we illustrate some features of the optimal menu using a collection of simple applications of our framework. Given the simplicity of our model, the examples are highly stylized and are not meant to serve as descriptive models of the concrete economic environments referred to. However, we believe they may illuminate some contractual arrangements that we observe in reality. Throughout this subsection, we assume $F(\theta) = \theta$ and $p = \frac{1}{2}$ for simplicity.

5.1 Ownership vs. fixed wages

We begin by revisiting the example given in the Introduction. A manager contemplates what contracts to offer to potential sales agents. Assume that the volume of sales can take any value in the interval $[0, 1]$. The cost that the agent incurs in generating a volume of sales a depends on the state of nature. In state u, it is not costly to generate higher sales, whereas in state v, the cost of generating sales exactly offsets the revenue. Thus, $u(a) = a$ and $v(a) = 0$ for every $a \in [0, 1]$. The agent's reservation value is interpreted as the highest wage he can earn, in case he declines all of the manager's contracts. We normalize this wage to zero.

The optimal menu can be easily derived by making the following observations. By Proposition 2, the manager will speculate only with agents who are more optimistic than he is. By Proposition 6, the optimal contract is ex-post efficient. Therefore, by Proposition 5, the menu will consist of one non-speculative contract (denoted t_{NS}) and one speculative contract (denoted t_S). All types lower than some $\bar{\theta}$ sign t_{NS}, while all higher types sign t_S.

By Lemma 1, the non-speculative contract extracts the entire surplus of the agent in each of the states. One such contract is $t_{NS}(a) = a$. This contract means that agent receives no share in the sales he generates, and is compensated by a fixed wage, equal to his best outside option.

Let us turn to characterizing t_S. By Proposition 4, $\bar{\theta} = \frac{3}{4}$. By Proposition 5, $a^u = 1$ and $a^v = 0$. The pair of transfers, $t_S(0)$ and $t_S(1)$, can now be derived using the following observations. By Lemma 8, the VR constraint is binding for all agents who sign t_S:

$$-t_S(0) = -t_S(1) \tag{8}$$

By Corollary 1, type $\frac{3}{4}$ is indifferent between signing a speculative contract and signing a non-speculative contract:

$$\frac{3}{4}[1 - t_S(1)] + \frac{1}{4}[-t_S(0)] = 0 \tag{9}$$

Taken together, (8) and (15) imply that $t_S(0) = t_S(1) = \frac{3}{4}$.

These transfers may be implemented by a contract, $t_S(a) = \frac{3}{4}$ for every a. This contract essentially transfers full ownership of the project to the sales person. Thus, the optimal contract consists two extreme risk-sharing schemes. In one scheme, the manager bears all the risk and the sales person receives a fixed wage. In another scheme, the manager sells the project to the sales person, who bears all the risk.

Both contracts provide precisely the same incentives in the second period. Under both contracts, the sales person will generate a volume $a^u = 1$ in state u and $a^v = 0$ in state v. Thus, the multiplicity of contracts in the menu has nothing to do with second-period incentives. Similarly, it has nothing to do with the agent's risk attitudes. Instead, the agent's selection of a payment scheme reveals his degree of optimism.

This example is special, in the sense that t_S leaves the manager totally insured. In this respect, the term "speculative contract" is perhaps a misnomer, because the manager does not bear any risk under t_S. This feature is due to the assumption that v is flat (rather than, say, a strictly decreasing function). In the examples that follow, the optimal menu will contain speculative contracts that impose risk on both parties.

5.2 Unlimited consumption vs. variable rates

A commonly-observed menu of contracts offers either unlimited consumption at a fixed fee, side by side with a variable-rate scheme that charges according to consumption. Menus of this type are offered by telecommunication companies, where firms often offer a choice between "unlimited calling plans" and plans that condition the rate per minute on the amount of minutes used. Similarly, DVD rental stores offer "unlimited plans" as well as "limited plans". Finally, restaurants sometimes offer diners a choice between an "all you can eat" buffet and a selection of dishes "a la carte".

We propose to interpret such a menu as a tool for screening consumers according to their prior beliefs regarding their future tastes. This interpretation fits situations in which the agent has a fixed budget for the product in question, yet he is unsure of his future satiation level: would he desire a large amount of airtime to communicate with his partners on the mobile phone, or will short conversations suffice? how many idle hours will he spend? how much would he need to eat in order to satisfy his appetite?

In these situations, a variable-rate contract may be viewed as a bet between the principal and the agent, where the agent "wins" if he manages to consume only a small amount and the principal "wins" if the agent ends up consuming a large amount. In contrast, an unlimited consumption contract has no speculative component, because the payment the agent makes is independent of his level of consumption. Consumers who believe that their satiation level is likely to be low would prefer a speculative, variable-rate contract. On the other hand, consumers who believe that their satiation level is likely to be high would opt for the it is very likely that they would feel satiated only after consuming a large amount, opt for the non-speculative, unlimited-consumption contract.

We illustrate this idea with a simple example. A consumer is about to purchase a calling plan from a monopolistic mobile phone company. Being a newcomer to the mobile phone market, the consumer does not know if his desired amount of airtime will be high or low.

Formally, let a denote the proportion of monthly minutes that the consumer uses (where $a = 1$ means that a consumer uses every available minute). Let $u(a)$ and $v(a)$ represent the consumer's willingness to pay for a minutes of airtime in the low and high state respectively. Specifically, let $v(a) = a$ and

$$u(a) = \begin{cases} 2a & for \quad 0 \leq a < \frac{1}{2} \\ 1 & for \quad \frac{1}{2} \leq a < 1 \end{cases}$$

Note that from the mobile phone company's point of view, a consumer of type $\theta > \frac{1}{2}$ ($\theta < \frac{1}{2}$) underestimates (overestimates) his desired amount of airtime. Because $u(a) \geq v(a)$ for every a, Proposition 2 implies that the mobile phone company will offer speculative contracts to types above $\frac{1}{2}$ only. All types below $\bar{\theta}$ will be assigned a non-speculative contract. This contract can be implemented with a calling plan that offers unlimited calls for a monthly fee of 1.

To derive the set of speculative contracts, we first note that by Proposition 3, for all $\theta \geq \bar{\theta}$, $a^v_\theta = 1$ and

$$a^u_\theta \in \arg\max_{a \in [0,1]} [(2\theta - 1)u(a) + (\frac{3}{2} - 2\theta)a]$$

Solving this optimization problem yields $a^u_\theta = \frac{1}{2}$. Since $a^u_\theta \in \arg\max_a u(a)$ and $a^v_\theta \in \arg\max_a v(a)$, it follows that the optimal menu is ex-post efficient. By Proposition 5, the optimal menu contains a single speculative contract t^*.

Recall that VR_θ is binding for every $\theta \geq \bar{\theta}$ and that IR_θ is binding for $\bar{\theta}$. By Proposition 4, $\bar{\theta} = \frac{3}{4}$. Therefore:

$$\begin{aligned} \frac{3}{4}[u(\frac{1}{2}) - v(1) + t^*(1) - t^*(\frac{1}{2})] + v(1) - t^*(1) &= 0 \\ t^*(1) - t^*(\frac{1}{2}) &= v(1) - v(\frac{1}{2}) \end{aligned}$$

These equations yield $t^*(1) = \frac{11}{8}$ and $t^*(\frac{1}{2}) = \frac{7}{8}$. This speculative contract can be implemented with the following three-part tariff. For a flat fee of $\frac{7}{8}$, the consumer receives $a = \frac{1}{2}$ at a marginal price of zero. For any additional units, the consumer has to pay a marginal price of 1 (i.e., $t(a) = \frac{3}{8} + a$ for $a > \frac{1}{2}$).

Thus, we have interpreted the coexistence of variable-rate and unlimited consumption contracts as a consequence of the principal's attempt to screen the agent's probabilistic assessment of his satiation point. The unlimited consumption contract is a non-speculative contract, whereas the variable-rate contract is a speculative contract directed at consumers who significantly underestimate (in a probabilistic sense) their satiation level of consumption.

Applied to the DVD rental example, our characterization of the optimal menu implies that unlimited rental plans are essentially non-speculative contracts, while limited plans are speculative contracts who underestimate their future amount of idle time. In the context of the restaurant example, our interpretation may be surprising. One could expect a priori that the all-you-can-eat deal would be a speculative contract, targeted at diners who overestimate their satiation level. However, in our example, the "all-you-can-eat" deal is a non-speculative contract, while the "a la carte" deal is a speculative contract, aimed at diners who enter the restaurant believing they will have only an entree and a main course, but find themselves at the end of the meal browsing the dessert menu.

5.3 Speculative contracts and inefficiency

In this sub-section we analyze an example that demonstrates the subtle effect of speculative contracts on ex-post efficiency. A seller can provide (at zero cost) a product at any variety $a \in [0,1]$ demanded by a buyer. The ideal variety for the buyer depends on the state of nature. Assume that u and v take the following functional forms:

$$\begin{aligned} u(a) &= 2-\left(a-\frac{1}{2}\right)^2 \\ v(a) &= 2-\left(a-\frac{3}{4}\right)^2 \end{aligned}$$

That is, the buyer's ideal variety is $a=\frac{1}{2}$ in state u and $a=\frac{3}{4}$ in state v. For expositional simplicity, we shall focus on the contracts that the seller designs for types in the set $[\frac{1}{2},1]$.

The non-speculative contract for types $\frac{1}{2} \leq \theta < \bar{\theta}$ may take the simple form of a flat payment of 2, which is independent of a. Let us turn to the speculative contracts for types $\theta \geq \bar{\theta}$. We first note that for each of these types, a_θ^v can be set to $\frac{3}{4}$. By Proposition 3, to compute a_θ^u for $\theta \geq \bar{\theta}$ we need to solve

$$\max_{a\in[0,1]}\left\{\left[(2\theta-1)-\frac{1}{2}\right]\left[\left(a-\frac{3}{4}\right)^2-\left(a-\frac{1}{2}\right)^2\right]+\frac{1}{2}\left[2-\left(a-\frac{1}{2}\right)^2\right]\right\}$$

yielding $a_\theta^u=\frac{5}{4}-\theta$. Finally, by Proposition 4, $\bar{\theta}=\frac{3}{4}$.

We wish to make two observations regarding this result. First, the non-speculative contract specifies a flat payment, whereas the speculative contract specifies a contingent pricing schedule. Buyers who roughly share the seller's belief choose a contract that gives them the freedom to choose the exact variety of the product only after they learn their ideal point. In contrast, buyers who sharply disagree with the seller are willing to commit to the variety they will choose, and to pay a fine in case they change their mind in period 2.

Second, note that $a_\theta^u \leq \arg\max u$ for all $\theta \geq \bar{\theta}$. Since a_θ^u decreases with θ, the distance between a_θ^u and $\arg\max u$ increases with θ (in the range $\theta > \frac{3}{4}$). Thus, as the parties' beliefs become more polarized, the contract they sign becomes more inefficient ex-post (in state u). In other words, the more speculative the contract, the more inefficient the action that it induces in state u.

The latter aspect of our example has a further implication. If the seller could observe the buyer's type θ, he would assign to any buyer type $\theta > p$ a speculative contract that induces $a_\theta^v = \frac{3}{4}$ and $a_\theta^u = \arg\max\left[(\theta - p)(u - v) + pu\right] = \frac{3}{4} - \frac{\theta}{2}$. Compare this with our result that when the seller dose not observe θ, $a_\theta^u = \frac{1}{2}$ for $\theta \in [\frac{1}{2}, \frac{3}{4})$ and $a_\theta^u = \frac{5}{4} - \theta$ for $\theta > \frac{3}{4}$. It is easy to see that the outcome is "less inefficient" ex-post when the seller does not observe the buyer's type. Thus, if a social planner, who wishes to maximize social surplus (according to his own prior beliefs), had to choose between an environment in which the seller observes the buyer's prior and an environment in which the buyer's prior is his private information, he would prefer the latter environment.

These welfare implications are not general, but a consequence of certain features of the payoff structure: *(i)* $v(a) \equiv u(a - c)$, where c is the distance between the ideal points in the two states; *(ii)* u is concave - i.e., as the distance from the ideal point becomes larger, the marginal disutility from steering away from it increases; *(iii)* the Arrow-Pratt coefficient $-u''/u'$ increases with a (in the relevant range, in which a falls below the ideal point). The proof is elementary, involving first- and second-order conditions of the objective function $(\theta - p)(u - v) + pu$ when the agent's prior is known, and$(\psi(\theta) - p)(u - v) + pu$ when the prior is private information.

Note that in contrast to previous examples in the paper, the optimal menu in this subsection discriminates among optimistic types. Specifically, there is a continuum of speculative contracts. To characterize these contracts for $\theta > \frac{3}{4}$, recall that by Lemma 8, VR_θ is binding for these types. Therefore,

$$t(a_\theta^v) - t(a_\theta^u) = (\frac{1}{2} - \theta)^2 \tag{10}$$

This implies that for all $\theta \geq \frac{3}{4}$,

$$q(\theta) = \frac{\theta}{2} - \frac{5}{16}$$

Hence, by (17),

$$t(a_\theta^v) = \theta(\frac{\theta}{2} - \frac{5}{16}) + 2 - \int_{\frac{3}{4}}^{\theta} (\frac{x}{2} - \frac{5}{16})dx = \frac{\theta^2}{4} + \frac{61}{32}$$

Substituting this into (10) we obtain $t(a_\theta^u) = t(a_\theta^v) - (\frac{1}{2} - \theta)^2$. This means that the higher the buyer's prior on u, the lower the payment he makes in this state, and the higher the payment he makes in v.

6 Concluding remarks

We have argued that menus of contingent contracts can be explained as a consequence of a principal's attempt to screen the agent's prior belief, in environments with non-common priors. Whereas standard accounts highlight the role of contingent contracts in providing incentives for the agent, we interpreted them as bets. Indeed, in many of our examples the agent's actions are independent of the contract he selects from the menu.

Note that we do not take a stand as to which of the two parties, if any, holds the correct prior. Their differences in beliefs may or may not be a result of inherent biases such as over-optimism. For example, in the manager-worker example of the Introduction, each party may be over-optimistic or over-pessimistic regarding the agent's future cost of effort, and our analysis does not require us to make any judgment in this regard. Even if it is natural to assume that the agent will tend to be over-optimistic relative to the principal, it does not follow that we have to assume that the principal is correct.

In some applications, however - e.g., when the principal is a firm and the agent is a consumer - one might one want to interpret the principal's prior p as correct, and the agent's prior $\theta \neq p$ as resulting from a psychological bias such as over-optimism. In this case, in order to be consistent with the model, it must be assumed that the consumer does not believe that his belief is biased. Otherwise, he would regard the menu as a signal of the principal's prior, and he would update his beliefs. Also, such an interpretation would have to be confronted with the empirical question of whether consumers make systematic errors in anticipating their future tastes (see Miravete (2003), for example).

References

[1] Armstrong, M. (1996), "Nonlinear Pricing with Imperfectly Informed Consumers," mimeo, Nuffield College.

[2] Baron, D.P. and D. Besanko (1984), "Regulation and Information in a Continuing Relationship," *Information Economics and Policy,* 1(3), 267-302.

[3] Bazerman, M. H., and J. J. Gillespie (1999), "Betting on the Future: The Virtues of Contingent Contracts," *Harvard Business Review*, 155-160.

[4] Courty, P. and H. Li (2000), "Sequential Screening," *Review of Economic Studies,* 67(4), 697-718.

[5] Eliaz, K. and R. Spiegler (2004), "Contracting with Diversely Naive Agents," mimeo, New York University and Tel-Aviv University.

[6] ——————— (2005), "A Mechanism Design Approach to Speculative Trade," mimeo, New York University and Tel-Aviv University.

[7] Fang, H. and G. Moscarini (2005), "Morale Hazard," *Journal of Monetary Economics,* 52, 749-777.

[8] Krishna, V. (2002), *Auction Theory*, Academic Press, San Diego.

[9] Landier, A. and D. Thesmar (2004), "Financial Contracting with Optimistic Entrepreneurs," mimeo, University of Chicago GSB.

[10] Miravete, E.J. (2003), "Choosing the Wrong Calling Plan: Ignorance and Learning," *American Economic Review,* 93(1), 297-310.

[11] Sandroni A. and F. Squintani (2004), "The Over-Confidence Problem in Insurance Markets," mimeo, UCL.

Appendix

Proof of Proposition 1. Assume that the principal knows the value of θ. Then, an optimal contract is given by the solution to the problem,

$$\max_{\left(a_\theta^u, t_\theta^u, a_\theta^v, t_\theta^v\right)} \theta t_\theta^u + (1-\theta)\, t_\theta^v$$

subject to the IR_θ, UR_θ and VR_θ constraints. Because the principal does not need to worry about incentive compatibility constraints, IR_θ is binding. Therefore, an optimal contract can be obtained by solving

$$\max_{\left(a_\theta^u, a_\theta^v\right)} \theta u\left(a_\theta^u\right) + (1-\theta)\, v\left(a_\theta^v\right)$$

Noting that for all $a_\theta^u, a_\theta^v \in [0,1]$,

$$\theta u\left(a_\theta^u\right) + (1-\theta)\, u\left(a_\theta^v\right) \leq \theta \max_a \left[u\left(a\right)\right] + (1-\theta) \max_a \left[v\left(a\right)\right]$$

it follows that the following contract is optimal:

$$t_\theta\left(a\right) = \begin{cases} \max_a \left[u\left(a\right)\right] & if \quad a = a_u^* \\ \max_a \left[v\left(a\right)\right] & if \quad a = a_v^* \\ \infty & if \quad a \notin \{a_u^*, a_v^*\} \end{cases}$$

The above contract is independent of θ. Therefore, it is also optimal when the principal does not know the value of θ and needs to take into consideration the incentive compatibility constraints. ■

Proof of Lemma 1. Assume the optimal menu includes a set of non-speculative contracts T (which may be a singleton) with the property that every contract $t \in T$ induces a pair of actions $\left(a_t^u, a_t^v\right)$ such that $u\left(a_t^u\right) < \max_a u(a)$ and $v\left(a_t^v\right) < \max_a v(a)$. Then

$$pt\left(a_t^u\right) + (1-p)t\left(a_t^v\right) < pu\left(a_t^u\right) + (1-p)v\left(a_t^v\right)$$

By definition:

$$pu\left(a_t^u\right) + (1-p)v\left(a_t^v\right) \leq pu^* + (1-p)v^* \tag{11}$$

Hence, the principal's expected revenue from every type who is assigned a non-speculative contract is strictly less than the highest expected willingness to pay of type $\theta = p$.

Consider amending the original menu by replacing the set T with the following contract:

$$t^*(a) = \begin{cases} u^* & if \quad u(a) = u^* \\ v^* & if \quad v(a) = v^* \\ \infty & if \qquad other \end{cases}$$

Note that this is the first-best contract for type $\theta = p$, and it yields zero expected payoff to all types. By Condition 1, it satisfies the UR and VR constraints of any type who chooses it. Because the original menu satisfied the individual rationality constraints, types who chose contracts outside T have no incentive to choose t^*. In addition, every contract outside T must yield the principal an expected revenue of at least $pu^* + (1-p)v^*$.

To see why replacing T with t^* increases the principal's expected revenue, consider first a type who originally chose a contract in T but now chooses the contract t. By inequality (11), the expected payment of such a type has strictly increased. Now consider a type who originally chose a contract in T but who now deviates to a contract other than t (i.e., a contract that was outside of T in the original menu). The expected payment of this type will be at least $pu^* + (1-p)v^*$. ■

Proof of Lemma 2. Assume $0 < p < 1$ and yet the optimal menu does *not* include any speculative contract. Consider the case of $u(a_u^*) > v(a_u^*)$ for some $a_u^* \in \arg\max_a u(a)$. Choose ε such that $u(a_u^*) - \varepsilon > v(a_u^*)$ and

$$0 < \frac{\varepsilon}{u(a_u^*) - v(a_u^*) - \varepsilon} < \frac{1-p}{p}$$

Then

$$p(u^* - \varepsilon) + (1-p)[v^* + u^* - v(a_u^*) - \varepsilon] > pu^* + (1-p)v^* \tag{12}$$

where $u^* \equiv \max_a u(a)$ and $v^* \equiv \max_a v(a)$. In addition, there exists a type $\bar{\theta} < 1$ sufficiently close to one that satisfies

$$\bar{\theta}\varepsilon - (1-\bar{\theta})[u^* - v(a_u^*) - \varepsilon] = 0 \tag{13}$$

Suppose the principal added to the original menu the following speculative contract:

$$t(a) = \begin{cases} u^* - \varepsilon & if \qquad a = a_u^* \\ v^* + u^* - v(a_u^*) - \varepsilon & if \qquad a = a_v^* \\ \infty & if \quad a \notin \{a_u^*, a_v^*\} \end{cases}$$

By (13), all types $\theta > \bar{\theta}$ strictly prefer $t(a)$ to any non-speculative contract, while the opposite is true of types below $\bar{\theta}$. By (12) and the continuity of F, the expected revenue from the new menu is strictly higher than with the original menu, a contradiction.

By essentially the same argument, we can show that if $v(a_v^*) > u(a_v^*)$ for some $a_v^* \in \arg\max_a v(a)$, then we can design a speculative contract with the following properties: (i)

there exists $\underline{\theta} > 0$ such that all types $\theta < \underline{\theta}$ strictly prefer this contract to any non-speculative contract, and (ii) the principal obtains an expected revenue strictly above $pu^* + (1-p)v^*$. Taken together, these properties imply that the principal can raise his expected revenue by including this speculative contract in his original menu, a contradiction.

Assume next that $p = 0$. Suppose $u(a') > v(a')$ for some a', yet the optimal menu has *no* speculative contracts. By Condition 1, there exists $a_v^* \in \arg\max_a v(a)$ such that $a' \neq a_v^*$. The principal can add to this menu the following contract:

$$t(a) = \begin{cases} u(a') - \varepsilon & if \quad a = a' \\ v^* + u(a') - v(a') - \varepsilon & if \quad a = a_v^* \\ \infty & if \quad a \notin \{a', a_v^*\} \end{cases}$$

For $0 < \varepsilon < u(a') - v(a')$ this contract yields the principal a higher revenue than any non-speculative contract. In addition, there exists a type $\bar{\theta} < 1$ such that $t(a)$ yields a strictly positive expected payoff to all types $\theta > \bar{\theta}$ and a strictly negative expected payoff to all types below $\bar{\theta}$. By the continuity of F, the new menu generates a strictly higher revenue than the original menu, a contradiction.

This has two implications. First, if $p = 0$ and $u(a') > v(a')$ for some a', then the optimal menu must contain at least one speculative contract. But this also implies the converse: if $p = 0$ and the optimal menu has no speculative contract, then $u(a) \leq v(a)$ for all a.

An argument along the same lines establishes that when $p = 1$, the optimal menu includes at least one speculative contract if, and only if, $u(a) < v(a)$ for some a. ■

Proof of Lemma 3. Note first that by IR_p, type $\theta = p$ cannot be assigned a speculative contract. We now show that if a type $\theta > p$ is assigned a speculative contract, then a higher type must also be assigned such a contract. Letting $x = \theta - p > 0$, it follows that

$$(\theta - x)\left[u\left(a_\theta^u\right) - t\left(a_\theta^u\right)\right] + (1 - \theta + x)\left[v\left(a_\theta^v\right) - t\left(a_\theta^v\right)\right] < 0$$

By IR_θ,

$$\theta\left[u\left(a_\theta^u\right) - t\left(a_\theta^u\right)\right] + (1 - \theta)\left[v\left(a_\theta^v\right) - t\left(a_\theta^v\right)\right] \geq 0$$

Hence,

$$u\left(a_\theta^u\right) - t\left(a_\theta^u\right) > v\left(a_\theta^v\right) - t\left(a_\theta^v\right)$$

and

$$u\left(a_\theta^u\right) - t\left(a_\theta^u\right) \geq 0$$

Consider a type $\phi > \theta$. Assume this type is assigned a non-speculative contract. Then by Lemma 1, $U(\phi, \phi) = 0$. By $IC_{\phi,\theta}$, an agent of type $\phi > \theta$ satisfies $U(\phi, \phi) \geq U(\theta, \phi)$. Because $U(\cdot, x)$ is affine in x, $U(\theta, \phi) > U(\theta, \theta)$. Since $U(\theta, \theta) \geq 0$, we reached a contradiction. A similar argument applies for types lower than p. ■

Proof of Corollary 1. Assume that $U(\bar{\theta},\bar{\theta}) > 0$. Since $U(\cdot,x)$ is affine in x, $U(\bar{\theta},\phi) > 0$ for all $\phi \geq \bar{\theta}$. Suppose the principal deviates by modifying all the speculative contracts as follows: for every $\phi \geq \bar{\theta}$, $t(a^u_\phi)$ and $t(a^v_\phi)$ are both raised by some arbitrarily small ε. This modification leaves all the IR, IC, UR and VR constraints intact, and generates a higher revenue, a contradiction. ■

Proof of Proposition 2. Assume $u(a) \geq v(a)$ for all a. By Lemma 2, if $p = 1$, there are no speculative contracts in the optimal menu. Consider then some $p < 1$. Assume, by contradiction, that type $\theta = 0$ is assigned a speculative contract t_0. Then there exists a pair of actions (a, a') such that $v(a) \geq t_0(a)$ and $u(a') < t_0(a')$. But from our assumption that $u(a) \geq v(a)$ it follows that $u(a') - t_0(a') < u(a) - t_0(a)$, which violates UR_0, a contradiction. By Lemma 3, if $\theta = 0$ is not assigned a speculative contract, then no type $0 < \theta < p$ can be assigned such a contract. A similar proof establishes the second part of the result. ■

Proof of Lemma 4. Assume the principal designs two separate menus such that one is optimal for the distribution of types F conditional on $\theta \in [p,1]$, and another is optimal for the distribution F conditional on $\theta \in [0,p]$. Denote the first set of contracts by T^+ and the second by T^-. We claim that these menus have the property that each type in $[p,1]$ (resp. $[0,p]$) weakly prefers his assigned contract to every contract in T^- (resp. T^+).

By Lemmas 1 and 3, there exist $\underline{\theta} \in [p,1]$ and $\bar{\theta} \in [0,p]$ such that $U(\theta,\theta) = 0$ for all $\theta \in \left(\underline{\theta},\bar{\theta}\right)$. By Definition 1, for all $\theta \leq \underline{\theta}$,

$$v\left(a^v_\theta\right) - t\left(a^v_\theta\right) > u\left(a^u_\theta\right) - t\left(a^u_\theta\right)$$

while for all $\theta \geq \bar{\theta}$,

$$u\left(a^u_\theta\right) - t\left(a^u_\theta\right) > v\left(a^v_\theta\right) - t\left(a^v_\theta\right)$$

In addition, every $\theta \notin \left(\underline{\theta},\bar{\theta}\right)$ satisfies inequality (2). This means that for every pair of types θ, ϕ such that $\theta \leq p$ and $\phi \geq \bar{\theta}$ we have that $U(\phi,\theta) < 0$. Similarly, $U\left(\phi',\theta'\right) < 0$ for every pair of types θ', ϕ' with $\theta' \geq p$ and $\phi \leq \underline{\theta}$.

By assumption, T^+ (resp. T^-) maximizes the principal's expected revenue, conditional on $\theta \in [p,1]$ (resp. $\theta \in [0,1]$). Hence, the union of T^+ and T^- maximizes the *unconditional* expected revenue of the principal. In addition, T^+ satisfies the IR, IC, UR and VR constraints of the types in $[p,1]$, while T^- satisfies the corresponding constraints for types in $[0,p]$. From the argument made in the previous paragraph, the set of contracts $T^+ \cup T^-$ also satisfies these constraints for *all* types in $[0,1]$. ■

Proof of Lemma 5. Consider the types in $[p,1]$. By Lemma 3, there exists $\bar{\theta} \in [p,1]$ such that all types $\theta < \bar{\theta}$ choose a non-speculative contract. By Lemma 1, this contract is ex-post efficient, hence, $a^v_\theta \in \arg\max_a u(a)$.

By Lemma 3, if $\bar{\theta} < 1$, then all types $\theta \geq \bar{\theta}$ are assigned speculative contracts. Assume there exists $\theta^* > \bar{\theta}$ who is assigned a contract with $a^v_\theta \notin \arg\max_a v(a)$. Let $b \in \arg\max_a v(a)$.

Consider the modified contract t'_θ, which differs from t_θ only in two actions, a^v_θ and b, such that $t'(a^v_\theta) = \infty$ and $t'_\theta(b)$ is designed such that $v(b) - t'_\theta(b) = v(a^v_\theta) - t_\theta(a^v_\theta)$. Note that $t'_\theta(b) > t_\theta(a^v_\theta)$. It follows that b is an optimal action for the agent in state v, under t'_θ. Because the agent's net payoff in state v from the modified contract is the same as under the original contact, all the IR, IC constraints, as well as VR_θ, continue to hold. To see that UR_θ is also satisfied, note that because t_θ is speculative, $[u(a^u_\theta) - t(a^u_\theta)] - [v(a^v_\theta) - t(a^v_\theta)] > 0$. Combining this inequality with the definition of $t_\theta(b)$, we obtain that UR_θ is satisfied.

A similar argument shows that w.l.o.g. we may set $a^u_\theta \in \arg\max_a u(a)$ for $\theta \leq p$. ■

Proof of Lemma 7. We adopt Krishna's (2002, pp. 63-66) derivation of incentive compatibility for direct mechanisms. Define $m(\theta) \equiv t_\theta(a^v_\theta) - v(a^v_\theta)$. The optimal menu is incentive compatible if for all types θ and ϕ,

$$V(\theta) \equiv \theta q(\theta) - m(\theta) \geq \phi q(\theta) - m(\phi)$$

By Observation 1, $q(\theta) \geq 0$ for all $\theta \geq \underline{\theta}$ (by the definition of $\underline{\theta}$, $q(\theta) = 0$ for all $\theta < \underline{\theta}$). Hence, the L.H.S of the above inequality is an affine function of the true value θ. Incentive compatibility implies that for all $\theta \geq \bar{\theta}$,

$$V(\theta) = \max_{\phi \in [0,1]} \{\theta q(\phi) - m(\phi)\}$$

I.e., $V(\theta)$ is a maximum of a family of affine functions, and hence it is convex on $[\bar{\theta}, 1]$.[7]

Incentive compatibility is equivalent to the requirement that for all $\theta, \phi \in [\bar{\theta}, 1]$,

$$V(\phi) \geq V(\theta) + q(\theta)(\phi - \theta)$$

This implies that for all $\theta > \bar{\theta}$, $q(\theta)$ is the slope of a line that supports the function $V(\theta)$ at the point θ. Because $V(\theta)$ is convex it is absolutely continuous, and thus differentiable almost everywhere in the interior of its domain. Hence, at every point that $V(\theta)$ is differentiable, $V'(\theta) = q(\theta)$. Since $V(\theta)$ is absolutely continuous, we obtain that for all $\theta > \bar{\theta}$,

$$V(\theta) = V(\underline{\theta}) + \int_{\bar{\theta}}^{\theta} q(x)\, dx$$

By Corollary 1, $V(\underline{\theta}) = 0$, and so we obtain that $U(\theta, \theta) = \int_{\underline{\theta}}^{\theta} q(x)\, dx$ for $\theta \geq \bar{\theta}$.

By a similar argument, if the speculative contract assigned to type $\theta \leq \underline{\theta}$ is speculative, then $U(\theta, \theta) = \int_{1-\underline{\theta}}^{1-\theta} [-q(x)] dx$. ■

Proof of Lemma 8. Consider first the types $\theta \geq \bar{\theta}$. Let δ_θ denote the slack in the VR_θ

[7] Because all types in $[p, \bar{\theta})$ are assigned a non-speculative contract, $V(\theta) = 0$ for all $\theta \in [p, \bar{\theta})$.

constraint of type θ. Then

$$t_\theta\left(a_\theta^v\right) - t_\theta\left(a_\theta^u\right) = v\left(a_\theta^v\right) - v\left(a_\theta^u\right) - \delta_\theta \tag{14}$$

By Lemma 6, if type $\theta \geq \bar{\theta}$ is assigned a speculative contract, which is incentive compatible, then

$$U(\theta,\theta) = \int_{\bar{\theta}}^{\theta} \left[u\left(a_x^u\right) - v\left(a_x^u\right) - \delta_x\right] dx \tag{15}$$

Assume the optimal menu satisfies that $\delta_\theta > 0$ for some positive measure of types in $[\bar{\theta}, 1]$. Consider amending the menu by changing only $t_\theta\left(a_\theta^v\right)$ for all types $\theta \geq \underline{\theta}$ such that the new transfer is equal to $t_\theta\left(a_\theta^v\right) + \delta_\theta$, making the VR_θ constraint binding for all these types. Clearly, this change does not violate the UR_θ constraint of these types. From (15), it follows that the incentive compatibility constraints are not violated, and that this change only raises $U(\theta,\theta)$. Hence, the IR_θ constraint is also not violated.

We claim that the above change in the menu, increases the principal's expected revenue, in contradiction to our assumption that the original menu was optimal. To see this, note that by Lemma 3, the optimal menu maximizes the expression

$$G\left(\bar{\theta}\right)\left[pu^* + (1-p)v^*\right] + \int_{\bar{\theta}}^{1} \left\{(1-p)\left[t_\theta\left(a_\theta^v\right) - t_\theta\left(a_\theta^u\right)\right] + t_\theta\left(a_\theta^u\right)\right\} dG(\theta)$$

Substituting (14) into the above expression, we obtain that reducing δ_θ for a positive measure of types in $[\bar{\theta}, 1]$ increases the principal's expected revenue.

A similar argument establishes that the UR_θ constraint is binding for all types $\theta \leq \underline{\theta}$. ■

Proof of Proposition 3. Assume F satisfies the monotone hazard rate condition. By Lemma 5, we can restrict attention w.l.o.g. to optimal menus that induce $a_\theta^v \in \arg\max v$ for $\theta \geq p$, and $a_\theta^u \in \arg\max u$ for $\theta \leq p$. Consider first the problem of designing an optimal menu for types distributed on $[p, 1]$. Denote $u^* \equiv \max_a u(a)$ and $v^* \equiv \max_a v(a)$. The above lemmas allow us to reduce this problem to the following optimization problem,

$$\max_{\bar{\theta}, \left\{a_\theta^u, t_\theta^u\right\}_{\theta\in[\bar{\theta},1]}} G\left(\underline{\theta}\right)\left[pu^* + (1-p)v^*\right] + \int_{\bar{\theta}}^{1} \left\{(1-p)\left[v^* - v\left(a_\theta^u\right)\right] + t_\theta\left(a_\theta^u\right)\right\} dG(\theta)$$

subject to the IR_θ, $IC_{\theta,\phi}$, UR_θ and VR_θ constraints for all types θ and ϕ. We adopt the standard practice in mechanism-design (see Krishna (2002)) of solving this problem under the assumption that the $IC_{\theta,\phi}$ and IR_θ constraints are satisfied, and then checking that this constraint is indeed satisfied by the solution we obtain.

Because VR_θ is binding (Corollary 2) for all $\theta \geq \bar{\theta}$, we can write

$$t_\theta\left(a_\theta^u\right) = t_\theta\left(a_\theta^v\right) + v\left(a_\theta^u\right) - v^* \tag{16}$$

Equating (3) and (4), we can obtain the following expression for $t_\theta\left(a^v_\theta\right)$,

$$t_\theta\left(a^v_\theta\right) = \theta q\left(\theta\right) + v^* - \int_{\bar{\theta}}^{\theta} q\left(x\right) dx \tag{17}$$

Substituting this expression into (16), we obtain

$$t_\theta\left(a^u_\theta\right) = \theta q\left(\theta\right) - \int_{\bar{\theta}}^{\theta} q\left(x\right) dx + v\left(a^u_\theta\right) \tag{18}$$

We can thus simplify the optimization problem by expressing the objective as a function only of the cutoff $\bar{\theta}$ and the actions $\{a^u_\theta\}_{\theta\in[\bar{\theta},1]}$,

$$\max_{\bar{\theta},\left\{a^u_\theta\right\}_{\theta\in[\bar{\theta},1]}} G\left(\bar{\theta}\right) pu^* + \int_{\bar{\theta}}^{1} \left[\theta q\left(\theta\right) - \int_{\bar{\theta}}^{\theta} q\left(x\right) dx + pv\left(a^u_\theta\right)\right] dG\left(\theta\right)$$

(note that we have taken the constant $(1-p)v^*$ out of the objective function). By interchanging the order of integration in $\int_{\bar{\theta}}^{1}\int_{\bar{\theta}}^{\theta} q\left(x\right) dx$, we can rewrite this expression as $\int_{\bar{\theta}}^{1}\left[1 - G\left(\theta\right)\right] d\theta$. We can thus replace $\int_{\bar{\theta}}^{1}\left[\theta q\left(\theta\right) - \int_{\bar{\theta}}^{\theta} q\left(x\right) dx\right] dG\left(\theta\right)$ with $\int_{\bar{\theta}}^{1} \psi\left(\theta\right) q\left(\theta\right) dG\left(\theta\right)$, where $\psi\left(\theta\right)$ is defined in (5). Because VR_θ is binding we can substitute $u\left(a^u_\theta\right) - v\left(a^u_\theta\right)$ in place of $q\left(\theta\right)$ to obtain the following objective function

$$\max_{\bar{\theta},\left\{a^u_\theta\right\}_{\theta\in[\bar{\theta},1]}} G\left(\bar{\theta}\right) pu^* + \int_{\bar{\theta}}^{1} \left[\psi\left(\theta\right)\left(u\left(a^u_\theta\right) - v\left(a^u_\theta\right)\right) + pv\left(a^u_\theta\right)\right] dG\left(\theta\right) \tag{19}$$

Because F satisfies the monotone hazard rate property, it follows that the optimal a^u_θ for each $\theta \in \left[\bar{\theta}, 1\right]$ is given by (6). A similar argument establishes that the optimal a^v_θ for each $\theta \in [0, \underline{\theta}]$ is given by (7).

It remains to verify that the menu we constructed is indeed individually-rational and incentive compatible. Note first that a non-speculative contract gives zero indirect utility to types in $\left(\underline{\theta}, \bar{\theta}\right)$, while speculative contracts, by definition, give negative indirect utility to these types. We now show that the speculative contracts assigned to types above p satisfy the IR and IC constraints of these types. Analogous arguments apply to speculative types below p.

By Corollary 1, the IR constraint of the lowest speculative type is binding. Hence, the speculative contracts assigned to all higher types is individually rational if, and only if they are incentive compatible. Incentive compatibility is equivalent to the requirement that $q\left(\theta\right)$ is non-decreasing (see Krishna (2002), p.68). Since by Corollary 2, VR_θ is binding for all $\theta \geq \bar{\theta}$, $q\left(\theta\right)$ is non-decreasing for these types if, and only if $u\left(a^u_\theta\right) - v\left(a^u_\theta\right)$ is non-decreasing for all $\theta \geq \bar{\theta}$. Consider a pair of types ϕ, θ such that $\phi > \theta$. By construction,

$$\psi\left(\phi\right)\left(u\left(a^u_\phi\right) - v\left(a^u_\phi\right)\right) + pv\left(a^u_\phi\right) \geq \psi\left(\phi\right)\left(u\left(a^u_\theta\right) - v\left(a^u_\theta\right)\right) + pv\left(a^u_\theta\right)$$

and

$$\psi(\theta)\left(u\left(a_\theta^u\right)-v\left(a_\theta^u\right)\right)+pv\left(a_\theta^u\right)\geq\psi(\theta)\left(u\left(a_\phi^u\right)-v\left(a_\phi^u\right)\right)+pv\left(a_\phi^u\right)$$

Adding these two inequalities and cancelling common terms yields

$$\left[\psi(\phi)-\psi(\theta)\right]\left[\left(u\left(a_\phi^u\right)-v\left(a_\phi^u\right)\right)-\left(u\left(a_\theta^u\right)-v\left(a_\theta^u\right)\right)\right]\geq 0$$

Because F satisfies the monotone hazard rate property, $\psi(\phi)>\psi(\theta)$. Therefore, $u\left(a_\phi^u\right)-v\left(a_\phi^u\right)\geq u\left(a_\theta^u\right)-v\left(a_\theta^u\right)$. ■

Proof of Proposition 4. We begin by proving part (i) of the proposition. Define

$$h(\theta)\equiv\max_{a\in[0,1]}\left[(\psi(\theta)-p)\cdot(u(a)-v(a))+p\cdot u(a)\right]-p\cdot u^* \tag{20}$$

and let $a_\theta\in\arg\max_{a\in[0,1]}\left[(\psi(\theta)-p)\cdot(u(a)-v(a))+p\cdot u(a)\right]$. It follows from the proof of Proposition 3 that the principal strictly prefers to speculate with agent type θ if and only if $h(\theta)>0$. By the monotone hazard rate property, $\psi(\theta)$ is a continuously increasing function. Note that $\psi(p)<p$ and $\psi(1)=1$. Therefore, there exists a unique $\theta^*>p$ that solves the equation $\psi(\theta)=p$. By Lemma 6, $q(\theta)>0$ for $\theta\geq\bar{\theta}$. By Lemma 8, VR is binding for $\theta\geq\bar{\theta}$. By the definition of $q(\theta)$, this means that $u(a_\theta)-v(a_\theta)>0$ for $\theta\geq\bar{\theta}$. Consider $\theta\geq\bar{\theta}$ for which $\psi(\theta)<p$. Then, the only way for $h(\theta)>0$ is if $u(a)-v(a)<0$, a contradiction. Now let $\theta=\theta^*$. Then, clearly $h(\theta^*)=0$. But this means that for every $\theta>\theta^*$ because $\psi(\theta)$ is increasing with θ, it must be the case that $h(\theta)>0$ for $\theta>\theta^*$. Therefore, $\bar{\theta}=\theta^*$. A similar argument establishes part (ii) of the proposition. ■

Proof of Proposition 5. By ex-post efficiency, a_θ^u can be chosen to be the same action for all types. Similarly, by Lemma 4, a_θ^v can be chosen to be the same action for all types. By Lemma 8, this means that $q(\theta)=u^*-v(\arg\max u)$ for all types $\theta\geq\bar{\theta}$. Hence, from equations (16) and (17) it follows that the transfers $t_\theta\left(a_\theta^u\right)$ and $t_\theta\left(a_\theta^v\right)$ are also constant for all $\theta\geq\bar{\theta}$. By an analogous argument it can be shown that $t_\theta\left(a_\theta^u\right)$ and $t_\theta\left(a_\theta^v\right)$ are constant for all $\theta\leq\underline{\theta}$. ■

Proof of Proposition 6. Suppose that a_θ^u does not maximize u. Then, a_θ^u also fails to maximize $u-v$. Denote $b=\arg\max u$. Consider a modified contract t'_θ, which differs from t_θ in two actions only, a_θ^u and b. Specifically, $t'_\theta(a_\theta^u)=\infty$, and $t'_\theta(b)$ is designed such that $u(b)-t'_\theta(b)=u(a_\theta^u)-t(a_\theta^u)$. Note that $t'_\theta(b)>t_\theta(a_\theta^u)$. By construction, the agent's net payoff in state u is the same under t_θ and t'_θ. Therefore, all the IR and IC constraints, as well as UR_θ, continue to hold. By the assumption that $b\in\arg\max(u-v)$, it can easily be checked that VR_θ also continues to hold. A similar argument shows that a_θ^v must maximize v.