Consumer optimism and price discrimination

Kfir Eliaz

Department of Economics, Brown University

Ran Spiegler

Department of Economics, University College London

We study monopolistic design of a menu of non-linear tariffs when consumers have biased prior beliefs regarding their future preferences. In our model, consumers are "optimistic" if their prior belief assigns too much weight to states of nature characterized by large gains from trade. A consumer's degree of optimism is his private information, and the monopolist employs the menu of non-linear tariffs to screen it. We characterize the optimal menu and show that the existence of non-common priors has significant qualitative implications for price discrimination and ex-post inefficiency. Finally, the characterization enables us to interpret aspects of real-life menus of non-linear tariffs.

Keywords. Contracts, speculative trade, screening, non-common priors, mechanism design, optimism, three-part tariffs.

JEL classification. D42, D84, D86.

1. Introduction

In many market situations, consumers need to create forecasts of their future preferences at the time they choose a supplier. Making travel and accommodation arrangements for a future vacation, signing up for a mobile phone or cable TV, DVD rental, and health care services, are instances of such situations. Given the time gap between the acceptance and realization of the deal, it makes sense to assume that consumers lack perfect knowledge of their future tastes and needs, especially when they are newcomers to the markets.

Given the prevalence of these situations, surprisingly little has been written on the question of how a monopolist would discriminate between consumers with diverse beliefs regarding their future tastes (notable exceptions are Baron and Besanko 1984, Armstrong 1996, Courty and Li 2000, and Miravete 2002, 2003). This literature invariably assumes that although consumers are imperfectly informed about their future tastes, their

---

Kfir Eliaz: `Kfir_Eliaz@Brown.Edu`
Ran Spiegler: `r.spiegler@ucl.ac.uk`

beliefs are unbiased. However, the diversity in consumers' beliefs may arise also from inherent biases in prior beliefs. For example, consumers may systematically underestimate the consumption quantity that will be required to satiate their need; inexperienced consumers may be unaware of future contingencies that will affect their willingness to pay for high quality products; finally, consumers may exaggerate the future benefits a service will bring. Different consumers will exhibit different belief biases. The firm may try to design its menu of pricing schemes to screen the degree of the consumer's bias.

In this paper we study a simple two-period monopolistic contracting model that addresses this non-standard price discrimination problem. In our model, the monopolist enables the consumer to choose in period 2 the quantity (or quality) of some good or service, conditional on signing a contract in period 1. If the consumer refuses to sign a contract, he is left with an outside option of known value. A contract is a non-linear pricing scheme—that is, a function that assigns a (possibly negative) monetary transfer from the consumer to the monopolist for every possible consumption quantity. Other than that, we place no restriction on the space of contracts.

The consumer has quasi-linear vNM utility over second-period outcomes. His utility from consumption is either u or v, depending on a state of nature, which is revealed to the consumer alone in period 2. The monopolist has a cost function c. The consumer assigns probability θ to state u. We assume that $\max(u - c) > \max(v - c)$—that is, maximal gains from trade are higher in state u than in state v. A higher θ thus represents a more "optimistic" belief. The monopolist's prior on the state u is p. The difference in the parties' beliefs is a pure difference in opinions; the consumer does not believe that the monopolist is better informed than him, and the monopolist need not update his beliefs as a result of the consumer's first-period decision. We assume that the value of θ is the consumer's private information. The monopolist believes that θ is drawn from some continuous cdf on $[0, 1]$. Thus, θ plays the role of a consumer "type." The monopolist's problem is to design a *menu of contracts* that maximizes his expected profit.

The following simple example illustrates the model. Imagine an individual without any driving experience who wishes to obtain a driver's license, yet is unsure about the number of driving lessons he will need to pass the driving test. In one state (the "low ability" state), the individual passes only if he takes two lessons: a basic lesson and a review lesson right before the test. In another state (the "high ability" state), the basic lesson is necessary and sufficient for passing. He receives a payoff of 1 (0) if he passes (fails). There is available one driving instructor, who incurs a cost of $c < \frac{1}{2}$ from each lesson. The instructor assigns probability p to the high-ability state. Note that the high-ability state is characterized by larger gains from trade. Hence, a prior belief that assigns a high probability to this state can be viewed as being "optimistic."

If the individual were known to share the instructor's prior, the latter could offer him a flat-rate deal, which allows the individual to take as many lessons as he wants for a fixed price of 1. This contract extracts the entire consumer surplus in each state, and yields an expected profit of $1 - [p \cdot c + (1 - p) \cdot 2c]$. Let us verify that this is indeed an optimal contract. Denote the prices for the first and second lessons by $t(1)$ and $t(2)$, respectively. In order for the individual to accept the contract, his expected payment

cannot exceed his willingness to pay for passing the test. Since the individual must take at least one lesson to pass, it must be the case that $t(1)+(1-p)t(2)\leq1$. In addition, the second lesson cannot be so expensive that the individual will refuse to take it if he learns that he is in the low-ability state. This means that $1-t(1)-t(2)\geq-t(1)$. Similarly, in the high-ability state the individual should not have an incentive to take a second lesson, hence $1-t(1)\geq1-t(1)-t(2)$. The above flat-rate deal—i.e., $t(1)=1$ and $t(2)=0$—maximizes the instructor's expected profit given the three constraints.

Suppose next that the instructor knows that the individual *overestimates* his ability—specifically, that he assigns probability $\theta^H>p$ to the high-ability state. The instructor could still offer the flat-rate scheme; the individual would accept it and it would generate the same expected profit as before. However, another possibility is to offer the individual a variable-rate scheme, in which the prices of a basic lesson and a review lesson are θ^H and 1, respectively. If the individual accepts this offer, then in the low-ability state he will take both a basic lesson and a review, while in the high-ability state he will take only the basic lesson. The individual's expected payoff from this offer, calculated according to *his own* prior belief, is zero, and so he would accept it. As for the instructor, this contract is superior to the flat-rate scheme because it generates an expected profit, calculated according to the *instructor's* prior belief, of $p\cdot(\theta^H-c)+(1-p)\cdot(\theta^H+1-2c)=(\theta^H-p)+1-[p\cdot c+(1-p)\cdot2c]$. Note that this conclusion relies on the fact that the state is not contractible. If it were, the instructor and the individual could sign what is essentially an infinite bet.

Alternatively, suppose that the instructor knows that the customer *underestimates* his ability—specifically, that he assigns probability $\theta^L<p$ to the high-ability state. The instructor can offer a variable-rate scheme such as the one offered to the optimistic type, except that the basic lesson's price is θ^L rather than θ^H. However, since $\theta^L<p$, this contract generates a lower expected profit than the original flat-rate deal, which thus constitutes the optimal contract for the pessimistic type.

The variable-rate scheme offered to the optimistic customer in this example is a "speculative contract"—namely, a bet on whether the individual will opt for a review lesson. An individual who shares the instructor's beliefs would never accept this contract. Note that if the instructor does *not* know which of the three types p, θ^H, or θ^L he is facing, he can simply ask the individual to choose between the flat-rate and variable-rate schemes. The optimistic customer type will weakly prefer the latter, while the other types will strictly prefer the former.

This example has several noteworthy features. First, the optimal menu of non-linear pricing schemes contains a risky, speculative scheme as well as a safe scheme that guarantees the consumer his reservation value in each state. Second, the consumer's behavior is independent of his choice of contract: he takes both lessons in the low-ability state, while skipping the review lesson in the high-ability state. Thus, the multiplicity of contracts has nothing to do with designing second-period incentives. Rather, its objective is to screen the consumer's prior belief. Third, the three types do not exert any informational externality on one another. This allows the instructor to screen the consumer's type at no cost. Finally, the instructor offers a speculative contract only to the consumer type who is "optimistic" about his future needs.

How general are these effects? In Section 4, we characterize the optimal menu of contracts in the class of environments described above. In addition to the standard assumptions (u and v are continuous, c is continuously increasing), we assume that at some ex-post efficient quantity in the state v, u lies weakly below v. The driving school example satisfies this condition. In Sections 3 and 4 we describe additional situations of economic interest that meet this condition as well. However, our motivation for this assumption is primarily methodological: it ensures that all "interesting" effects in our model arise from non-common priors. In the common-prior benchmark, the optimal menu consists of a single contract that maximizes gains from trade and fully extracts the consumer's surplus in each state. In contrast, in our model the optimal menu satisfies the following properties.

Speculative contracts There exist two cutoff beliefs $\theta^* > p$ and $\theta_* < p$, such that optimistic types ($\theta > \theta^*$) are offered a speculative contract; pessimistic types ($\theta < \theta_*$) are offered a speculative contract, too, as long as v lies strictly above u at some point.

Exclusion from speculation The optimal menu does not exclude consumer types in the usual sense of failing to transact with them. Nevertheless, there *is* exclusion, in the sense that the contract offered to consumer types who roughly share the monopolist's prior belief ($\theta \in (\theta^*, \theta_*)$) is the optimal contract under complete information for $\theta = p$. By definition, this is a non-speculative contract. It also has the property that all types evaluate it at the reservation value.

Ex-post (in)efficiency The contracts assigned to types above θ^* induce an efficient outcome in state v and a possibly inefficient outcome in state u. The contracts assigned to types below θ_* induce an efficient outcome in state u and a possibly inefficient outcome in state v (unless u lies weakly above v, in which case types below p are offered the optimal contract under complete information for $\theta = p$). This inefficiency arises from the speculation motive, and would exist also if the monopolist knew the consumer's type. In fact, in some environments the efficiency loss is mitigated by the informational asymmetry.

Thus, in the class of environments studied in Sections 3 and 4 all price discrimination and ex-post inefficiency phenomena are purely artifacts of the consumer's biased beliefs and the betting motive they induce.

Admittedly, the simplifying assumption excludes some environments of economic importance. Therefore, in Section 5 we analyze the model under the assumption that both u and v are increasing concave functions, and that $u' > v'$. This configuration often appears in textbook presentations of price discrimination models. In our context, it complicates the monopolist's considerations. First, the monopolist has a betting motive owing to the fact that the $p \neq \theta$. Second, it faces a standard screening problem: a consumer type $\theta > p$ can earn an informational rent by pretending to be $\theta' = p$. The conflation of these two motives makes the analysis less clear-cut than in the environment studied in Sections 3 and 4. Nevertheless, we provide a characterization of the

optimal menu and show via an example that when u is sufficiently higher than v, our model generates a prediction that is qualitatively distinct from the standard common-prior model.

2. RELATED LITERATURE

Our paper is mostly related to the small "sequential screening" literature described above, which studies non-linear pricing with imperfectly informed consumers. Both this literature and our work study the problem of a monopolist that offers consumers a menu of contingent pricing schedules. In both frameworks, the monopolist's objective is to screen consumers according to their unobservable prior beliefs over their future tastes. However, in contrast to our model, the sequential screening literature assumes common priors as well as common knowledge that the consumer is better informed about his future tastes than the monopolist. We show the importance of this distinction: in some environments, non-common priors are *necessary* for price discrimination to emerge.

Grubb (forthcoming) independently studies optimal contracting in the presence of non-common priors. There are several notable differences between the two papers. First, the two papers study different behavioral biases. We analyze consumers who may be optimistic in the sense that relative to the monopolist, they assign a higher belief to the state with the largest gains from trade. Grubb, on the other hand, studies consumers whose second-period willingness to pay gets a continuum of values, and who are "over-confident" in the sense that their subjective beliefs over second-period valuations are too narrowly concentrated (relative to the actual distribution) around the mean.[1] Second, Grubb's main focus is on the optimal contract under complete information and its implementability by three-part tariffs (in the context of mobile phone services), while we focus on price discrimination in the presence of incomplete information. Third, in our model, the monopolist can perfectly monitor the amount consumed by the consumer (as in the case of mobile phone services). Grubb, on the other hand, studies an environment in which the monopolist has only imperfect monitoring ability and can observe only the amount the consumer has bought but not how much he actually consumed (as, for example, in a market for a perishable good).[2]

Fang and Moscarini (2005) study a principal-agent model with non-common priors, and investigate the implications of non-common priors for the design of wage contracts. Unlike us, they analyze contract design with an informed monopolist, and focus on the signaling aspect of contracts. More specifically, they ask how a monopolist should 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 monopolist, therefore, faces the

---

[1] In particular, Assumprion A* in Grubb (forthcoming), which captures consumers' overconfidence, rules out overoptimism in our sense.

[2] For other examples of real-life speculative contracts, see Bazerman and Gillespie (1999). For further empirical discussion of whether consumers make systematic errors in anticipating their future tastes, see Miravete (2003).

following trade-off. On the one hand, he would like to provide the appropriate mone-
tary 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").

While this paper is concerned with monopolistic screening of a consumer's prior
belief, there have been a few works on competitive screening of priors. Landier and
Thesmar (forthcoming) examine debt contracts that are signed between investors and
entrepreneurs who differ in their degrees of optimism. Assuming a competitive envi-
ronment, in which investors earn zero profits, the authors construct a separating equi-
librium 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 en-
trepreneur 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. Sandroni and Squintani
(2007) modify the Rothschild–Stiglitz insurance market model, to allow for consumers
who are over-optimistic regarding their probability of an accident. They show that al-
though in equilibrium these consumers are under-insured, compulsory insurance need
not be Pareto-improving. Uthemann (2005) shows how to adapt a Hotelling-like model
of competitive price discrimination due to Armstrong and Vickers (2001), in order to
study competitive screening of consumers' prior beliefs regarding the future value of a
taste parameter.

Finally, this paper builds on our own previous work. Eliaz and Spiegler (2006) study
optimal contract design with dynamically inconsistent consumers. A consumer type
is his degree of naivete, modeled as his prior belief that his current preferences will
change in the future. By comparison, the monopolist believes that the consumer's pref-
erences are sure to change. As in the present paper, the monopolist's objective in Eliaz
and Spiegler (2006) is to screen consumers 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 prob-
ability are assigned a contract that serves as a perfect commitment device: it induces
them to choose the action that maximizes their current utility.

Eliaz and Spiegler (2007, 2008, forthcoming) develop further the research agenda
of mechanism design when consumer types consist of their prior beliefs regarding fu-
ture payoffs. In these models, multiple agents hold different priors over an unverifiable
state of nature that affects the outcome of a game they are about to play. In the first
period, the agents negotiate over contracts that define side payments as a function of
the second-period game outcome. Thus, contracts are essentially bets over the second-
period outcome. These papers define a notion of "constrained interim-efficient bets,"
characterize them, and discuss their implementability in terms of the underlying game's
payoff structure. Thus, the mechanism design problem in Eliaz and Spiegler (2007, 2008,
forthcoming) is concerned with interim efficiency, whereas in the present paper the mo-
nopolist's aim is to maximize profit. Moreover, while the present paper assumes that

only the consumer takes actions in the second period, in Eliaz and Spiegler (2007, 2008, forthcoming) all parties to the first-period bets may take actions in the second period.

3. The model

A monopolist offers a consumer the opportunity to choose an action from the set $[0,1]$. The cost of providing an action a is $c(a)$, where c is a continuous, non-decreasing, convex function satisfying $c(0) = 0$. In order to have access to this set of actions, the consumer must sign a contract with the monopolist one period beforehand. If the consumer does not sign a contract with the monopolist, 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 consumer to the monopolist. The monopolist is perfectly able to monitor the consumer's second-period action.

The consumer 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 consumer's preferences are represented by the function $u : [0,1] \to \mathbb{R}$, and in state v they are represented by the function $v : [0,1] \to \mathbb{R}$. We assume that in each state of nature there are non-negative gains from trade, and that the maximal gains from trade in state u are higher than they are in state v, i.e.,

$$\max_a [u(a) - c(a)] > \max_a [v(a) - c(a)] \geq 0.$$

We therefore interpret u as the "good" state. Let e^ω denote an ex-post efficient action in state $\omega \in \{u, v\}$.

The consumer believes that state u occurs with probability θ. Given our interpretation of u as the "good" state, a consumer with a higher θ may be regarded as a consumer with a higher degree of "*optimism.*" Faced with a contract t, the consumer's indirect utility from the contract is

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

The monopolist believes that u occurs with probability p, independently of the consumer's belief. Thus, from the monopolist's point of view, the consumer's beliefs are incorrect. A consumer type $\theta > p$ ($\theta < p$) overestimates (underestimates) the expected maximal gains from trade in the second period. We assume that any difference between the two parties' beliefs is purely due to differences in prior opinion. In Section 6 we discuss the interpretation of this assumption, as well as the implications of allowing p to depend on θ.

A complete characterization of the optimal menu in the above environment turns out to be a difficult task. To simplify our analysis, we restrict attention to environments

in which the utility functions u and v and the cost function c satisfy the additional property

$$v(e^v) \geq u(e^v) \text{ for all } e^v. \tag{*}$$

This property accommodates a number of economically relevant environments. For instance, suppose that the consumer's future preferences have a satiation point. Suppose further that the consumer has no uncertainty about the satiation payoff, but is unsure about the location of the satiation point. Such is the case when the consumer is interested in achieving some goal—passing a test (driving, SAT), quitting smoking, losing weight—but is unsure of the amount of effort (number of classes) required to attain that goal. If costs are sufficiently low, then these environments satisfy property (*).

   Our environment also accommodates situations in which the consumer does not know whether he will like the product at all. That is, he is not sure whether his utility will rise with consumption. A similar situation arises when the consumer is a TV station that buys reruns from major networks and may be uncertain as to whether increasing the dose of reality TV shows at the expense of other genres will cause ratings to rise or fall.

   As we show in the next section, property (*) is methodologically useful in that it precludes discrimination between consumer types when prior beliefs are common (yet privately known by the consumer). This provides us with a sharp benchmark to be compared with the non-common priors case. We examine environments that violate property (*) in Section 5.

4. Analysis

4.1 *Optimal contracts under complete information*

We begin by investigating the problem facing a monopolist who observes the consumer's prior belief θ. Its objective is to design a contract $t_\theta : [0,1] \to \mathbb{R}$ that maximizes expected profit subject to the following constraints: (i) according to the consumer's beliefs, the contract offers him in expectation at least his reservation payoff and (ii) after each state is realized, the consumer chooses the best action for him, given the contract. Formally stated, the monopolist solves the maximization problem

$$\max_{t_\theta}\{p[t_\theta(a_\theta^u) - c(a_\theta^u)] + (1-p)[t_\theta(a_\theta^v) - c(a_\theta^v)]\}$$

subject to the constraint

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

where

$$a_\theta^u \in \underset{a\in[0,1]}{\arg\max}[u(a) - t_\theta(a)] \tag{UI_θ}$$

$$a_\theta^v \in \underset{a\in[0,1]}{\arg\max}[v(a) - t_\theta(a)]. \tag{VI_θ}$$

We refer to a contract t_θ that solves this problem as an optimal contract for type θ under complete information.

The conditions UI_θ and VI_θ represent the fact that a consumer's indirect utility from a contract is determined by the actions he expects to choose in the two states. If the realized state in period 2 is u (v)—an event to which the consumer assigns probability θ $(1-\theta)$—he will choose the optimal action according to the utility function u (v). This is represented by the constraint UI_θ (VI_θ).

It follows that any contract t can be identified with a pair of actions (a_θ^u, a_θ^v). The first action is consistent with u-maximization in the second period, whereas the second action is consistent with v-maximization in the second period. Without loss of generality, we assume that $t(a) = +\infty$ for every $a \notin \{a_\theta^v, a_\theta^u\}$. With slight abuse of terminology, we henceforth refer to the inequality

$$u(a_\theta^u) - t_\theta(a_\theta^u) \geq u(a_\theta^v) - t_\theta(a_\theta^v)$$

as the UI_θ constraint and to the inequality

$$u(a_\theta^u) - t_\theta(a_\theta^u) \geq u(a_\theta^v) - t_\theta(a_\theta^v)$$

as the VI_θ constraint.

PROPOSITION 1. *Assume the monopolist knows the consumer's prior belief θ. Then the optimal contract for each consumer type has the following properties.*

(i) *If $p < \theta$, then*

$$a_\theta^v = e^v$$
$$a_\theta^u = \arg\max_a \{[\theta - p][u(a) - v(a)] + p[u(a) - c(a)]\}$$

and

$$t_\theta(a_\theta^u) = v(a_\theta^u) + \theta[u(a_\theta^u) - v(a_\theta^u)] \tag{1}$$
$$t_\theta(a_\theta^v) = v(a_\theta^v) + \theta[u(a_\theta^u) - v(a_\theta^u)]. \tag{2}$$

(ii) *If $p > \theta$, then*

$$a_\theta^u = e^u$$
$$a_\theta^v = \arg\max_a \{[p - \theta][v(a) - u(a)] + [1 - p][v(a) - c(a)]\}$$

and

$$t_\theta(a_\theta^u) = u(a_\theta^u) - (1 - \theta)[u(a_\theta^v) - v(a_\theta^v)] \tag{3}$$
$$t_\theta(a_\theta^v) = u(a_\theta^v) - (1 - \theta)[u(a_\theta^v) - v(a_\theta^v)]. \tag{4}$$

(iii) *If $p = \theta$, then $a_\theta^\omega = e^\omega$ and $t(a_\theta^\omega) = \omega(e^\omega)$ for all $\omega = u, v$.*

The optimal contract under complete information has two interesting features. First, when the monopolist and the consumer hold different prior beliefs, the contract induces the consumer to choose the ex-post efficient action in the state for which the monopolist's prior is higher than the consumer's. That is, if the optimal contract requires some inefficiency, the monopolist prefers to have this inefficiency in the state he deems less likely relative to the consumer. Second, when the monopolist and the consumer share the same prior belief, the optimal contract is ex-post efficient and in each state the monopolist extracts the entire consumer surplus.

Note that the optimal contract under complete information for $\theta = p$ is *independent of θ*. This has the following important implication.

COROLLARY 1. *Suppose the monopolist believes that the consumer's private information (his prior belief θ) is the true objective probability of u. Then it is optimal to offer all consumer types an ex-post efficient contract that extracts the entire consumer surplus in each state.*

This result follows from observing that even if the monopolist could observe the consumer's belief, the optimal contract would not depend on it. Thus, whenever property (*) holds, the common-prior benchmark is especially stark as there is no price discrimination at the stage at which contracts are signed, since all consumer types are offered *exactly the same contract*. Thus, in this environment consumers are offered a menu with several distinct contracts only as a result of the monopolist's attempt to earn speculative gains from (what he considers to be) consumers' biased beliefs.

The proof of Proposition 1 reveals another interesting feature of the optimal contract under complete information: the monopolist believes that a consumer of type θ is better off *rejecting* the complete-information optimal contract, which is aimed at him. In other words, the consumer would not have signed the contract had he shared the monopolist's beliefs. We refer to a contract with this property as a *speculative contract*. Formally, we make the following definition.

DEFINITION 1. A contract t_θ is *speculative* if

$$p[u(a_\theta^u) - t(a_\theta^u)] + (1-p)[v(a_\theta^v) - t(a_\theta^v)] < 0.$$

By this definition, complete-information optimal contracts for $\theta \neq p$ are speculative, whereas the complete-information optimal contract for $\theta = p$ is not speculative. The distinction between speculative and non-speculative contracts plays an important role in the next subsection.

4.2 *Optimal contracts under incomplete information*

The assumption that the monopolist observes the consumer's prior belief is clearly strong. A more realistic assumption is that the monopolist has some subjective probability distribution over the consumer's degree of optimism. Furthermore, in many of the economic environments referred to in the Introduction, we observe that suppliers offer a *menu* of contracts rather than a single contract. This suggests that firms have

incomplete information about the consumer's belief of his future preferences (whether biased or unbiased) and use menus to screen it.

Assume the monopolist does not know the value of θ, but believes that it is distributed over $[0,1]$ according to a continuous, strictly increasing cdf $F(\theta)$. Thus, the consumer's "type" consists of his prior on u. The monopolist's objective is to maximize expected profit. By the revelation principle, a solution to his problem can be obtained via a direct mechanism, in which consumers are asked to report their type, and each reported type ϕ is assigned a contract $t_\phi : [0,1] \to \mathbb{R}$. An optimal menu of contracts $\{t_\theta(a)\}_{\theta \in [0,1]}$ is a solution of the maximization problem

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

subject to the IR_θ, UI_θ, and VI_θ constraints described above, and the additional incentive compatibility constraint, which states that for all $\phi \in [0,1]$, a consumer of type θ cannot be better off by pretending to be of type ϕ and signing the contract assigned to that type. Denote $D_\theta^u \equiv u(a_\theta^u) - t_\theta(a_\theta^u)$ and $D_\theta^v \equiv u(a_\theta^v) - t_\theta(a_\theta^v)$. That is, D_θ^ω is the net payoff the consumer expects to get in state ω. Then the incentive compatibility constraints can be written as

$$\theta D_\theta^u + (1-\theta)D_\theta^v \geq \theta D_\phi^u + (1-\theta)D_\phi^v. \qquad (IC_\theta, \phi)$$

Define $U(\phi, \theta)$ to be the expected payoff of a type θ consumer who pretends to be of type ϕ—that is, $U(\phi, \theta) \equiv \theta D_\phi^u + (1-\theta)D_\phi^v$. 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. We sometimes use the abbreviated notation $t_\theta^\omega \equiv t_\theta(a_\theta^\omega)$ and $c_\theta^\omega \equiv c(a_\theta^\omega)$, for $\omega = u, v$.

The previous subsection introduced the notion of a speculative contract, which turns out to be especially useful in characterizing optimal menus. Our first observation concerns the structure of *non*-speculative contracts in optimal menus. Let t^* denote the optimal contract under complete information for an agent of type $\theta = p$.

LEMMA 1. *Without loss of generality, an optimal menu of contracts includes no non-speculative contract other than t^*.*

This result is based on the observation that the contract t^* induces the reservation value for *all consumer types*. Therefore, including it in the menu of contracts does not add an incentive compatibility constraint. In addition, among all possible non-speculative contracts, this contract generates the highest expected profit for the monopolist. Therefore, if the menu contained other non-speculative contracts, it would be profitable for the monopolist to replace all of them with the optimal contract under complete information for $\theta = p$.

Our next result presents several noteworthy features of optimal menus. First, there exist optimal menus that partition the set of types into three regions. The lowest and highest regions contain those types whose prior beliefs are furthest from the monopolist. These consumers sign contracts that are essentially bets on their second-period

actions. The middle region in between the two cutoffs consists of consumers who more-or-less share the monopolist's beliefs. These consumers are excluded from betting with the monopolist and are therefore assigned non-speculative contracts that fully extract their ex-post consumer surplus in each state.

PROPOSITION 2. *The monopolist can design optimal menus with the following properties. There is a pair of cutoffs, θ_l, θ_h, where $\theta_l < p < \theta_h$, such that*

(i) *all types in the range (θ_l, θ_h) are assigned the optimal non-speculative contract, whereas types in the ranges $[0, \theta_l]$ and $[\theta_h, 1]$ are assigned speculative contracts*

(ii) *the $[0, p]$ problem is independent of the $[p, 1]$ problem*

(iii) *for each $\theta \geq \theta_h$, the action a_θ^v is ex-post efficient, the constraint VI_θ is binding, and $D_u^\theta > 0$ while $D_v^\theta < 0$*

(iv) *for each $\theta \leq \theta_l$, the action a_θ^u is ex-post efficient, the constraint UI_θ is binding, and $D_u^\theta < 0$ while $D_v^\theta > 0$*

(v) *when $u(a) \geq v(a)$ for all a, $\theta_l = 0$.*

Thus, speculative contracts offered to types $\theta < p$ exert no informational externality on types $\theta > p$, and vice versa. The only incentive constraints that the monopolist needs to worry about are those that prevent consumers who are more (less) optimistic than the monopolist to downplay their optimism (pessimism). This simplifies the derivation of an optimal menu, in that it breaks it down into three separate problems: (i) solving for the optimal menu for $\theta \geq \theta_h$, (ii) solving for the optimal menu for $\theta \leq \theta_l$, and (iii) solving for the non-speculative contract for $\theta_l < \theta < \theta_h$. When u lies weakly above v, there is further simplification because the cutoff θ_l becomes irrelevant.

Note that in an optimal menu, speculative contracts preserve the property that the outcome is ex-post efficient in the state the monopolist deems more likely than the consumer. However, Proposition 2 does not characterize the possibly inefficient outcome in the other state. For such a characterization, we need to impose additional assumptions on F. Denote $z = 1 - x$ and $G(z) = 1 - F(1 - x)$, and define

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

Define $\psi_G(z)$ accordingly.

The next result characterizes the actions induced by optimal menus of contracts. To simplify the exposition we characterize the payments induced by this menu in a separate proposition below.

PROPOSITION 3. *(i) If F satisfies the monotone hazard rate condition, then for every $\theta \geq \theta_h$,*

$$a_\theta^u \in \arg\max_a \{[\psi_F(\theta) - p][u(a) - v(a)] + p[u(a) - c(a)]\}, \tag{6}$$

where θ_h is the unique solution of $\psi_F(\theta) = p$.

(ii) If G satisfies the monotone hazard rate condition, then for every $\theta \leq \theta_l$,

$$a_\theta^v \in \arg\max_a \{[\psi_G(1-\theta)-(1-p)][v(a)-u(a)]+[1-p][v(a)-c(a)]\},$$

where θ_l is the unique solution of $\psi_G(1-\theta)=1-p$.

Thus, when F and G satisfy the monotone hazard rate property,[3] the action chosen by the consumer in the state he deems more likely than the monopolist has the following interesting property: it maximizes a weighted sum of the total surplus in that state and the difference between u and v. To interpret this observation, consider a speculative contract signed by a type $\theta > p$. This contract satisfies the VI_θ constraint with equality, and the consumer's indirect utility from the contract may be written as

$$\theta[u(a_\theta^u)-v(a_\theta^u)]+D_\theta^v. \tag{7}$$

Since the monopolist's prior on v is higher than the consumer's, the contract may be viewed as a bet in which the monopolist bets on v and the consumer bets on u. The probability that the consumer wins the bet is θ. From the point of view of the monopolist (who expects to win the bet), the consumer's ex-post utility will be D_θ^v. The consumer accepts the contract because of the extra term $\theta[u(a_\theta^u)-v(a_\theta^u)]$. Therefore, $u(a_\theta^u)-v(a_\theta^u)$ may be viewed as the consumer's "speculative gain."

We now provide an outline of the proof of Proposition 3. By part (ii) of Proposition 2, we can solve two separate optimization problems, one for types below p and another for types above p. Consider the types in $[p,1]$. Since the contracts offered for these types induce an ex-post efficient action in state v (see part (iii) of Proposition 2), the monopolist's objective function may be written as

$$\max_{\theta_h,\{a_\theta^u,t_\theta^u\}_{\theta\in[\theta_h,1]}} \left[\frac{F(\theta_h)-F(p)}{1-F(p)}\right]\{p[u(e^u)-c(e^u)]+(1-p)[v(e^v)-c(e^v)]\}$$
$$+\int_{\theta_h}^1 \{p[t_\theta(a_\theta^u)-c(a_\theta^u)]+(1-p)[t_\theta(e^v)-c(e^v)]\}\left[d\frac{F(\theta)}{1-F(p)}\right].$$

Since the VI_θ constraint of these types is binding (part (ii) of Proposition 2), the monopolist's expected profit from types at or above θ_h equals

$$\int_{\theta_h}^1 \{(1-p)[v(e^v)-c(e^v)-v(a_\theta^u)+c(a_\theta^u)]+t_\theta(a_\theta^u)-c(a_\theta^u)\}\left[d\frac{F(\theta)}{1-F(p)}\right]. \tag{8}$$

Following standard practice in the mechanism design literature, we solve the "relaxed problem" obtained by assuming the solution is incentive compatible, individually rational, and also satisfies the second-period incentive constraints (*UI* and *VI*). This amounts

---

[3]The distribution F satisfies the monotone hazard rate condition if the ratio $f(x)/[1-F(x)]$ is nondecreasing in x. For example, if the density function induced by F is symmetric around $\theta = \frac{1}{2}$, then F satisfies the monotone hazard rate if and only if G satisfies it as well.

to using the integral representation of the IC_θ constraint, together with the fact that the VI_θ constraint is binding, to replace the expression for $t_\theta(a_\theta^u)$ with an expression that depends only on a_θ^u. We then substitute this expression into (8) and use integration by parts to rewrite (8) as

$$\int_{\theta_h}^1 \{[\psi_F(\theta) - p][u(a_\theta^u) - v(a_\theta^u)] + p[u(a_\theta^u) - c(a_\theta^u)]\}\, dF(\theta).$$

Since we have already incorporated the incentive constraints into the objective function, we can maximize the above expression point-by-point, which yields the solution described in (6).

To complete the proof it remains to verify that all the constraints we assumed to be satisfied at the solution are indeed satisfied. Since type θ's expected utility from truth-telling may be written as (7), the incentive compatibility constraints hold if and only if the difference $u(a_\theta^u) - v(a_\theta^u)$ is non-negative at θ_h and non-decreasing in θ (see the Claim in the proof of the proposition). This allows us also to establish that the IR_θ constraint holds for all $\theta > \theta_h$ once we construct type θ_h's contract so that it satisfies this type's individual rationality constraint with equality. The payments made by each type are designed to satisfy the VI_θ constraint with equality (see Proposition 4 below). This implies that the UI_θ constraint is satisfied if and only if $u(a_\theta^u) - v(a_\theta^u) \geq u(e^v) - v(e^v)$. But this inequality follows from property (*) and the result that $u(a_\theta^u) - v(a_\theta^u) \geq 0$ for all $\theta \geq \theta_h$.

PROPOSITION 4. *(i) If F satisfies the monotone hazard rate condition, then in state u the cutoff type $\theta = \theta_h$ makes the payment*

$$t_{\theta_h}(a_{\theta_h}^u) = p[u(e^u) + c(a_{\theta_h}^u) - c(e^u)] + (1-p)v(a_{\theta_h}^u)$$

and any higher type makes the payment

$$t_\theta(a_\theta^u) = v(a_\theta^u) + \theta[u(a_\theta^u) - v(a_\theta^u)] - \int_{\theta_h}^\theta [u(a_x^u) - v(a_x^u)]\, dx.$$

The payment made by each type $\theta \geq \theta_h$ in state v is given by

$$t_\theta(e^v) = t_\theta(a_\theta^u) - v(a_\theta^u) + v(e^v).$$

(ii) If G satisfies the monotone hazard rate condition, then in state v the cutoff type $\theta = \theta_l$ makes the payment

$$t_{\theta_l}(a_{\theta_l}^v) = pu(a_{\theta_l}^v) + (1-p)[v(e^v) + c(a_{\theta_l}^v) - c(e^v)]$$

and any higher type makes the payment

$$t_\theta(a_\theta^v) = v(a_\theta^v) + \theta[u(a_x^v) - v(a_x^v)] + \int_\theta^{\theta^l} [u(a_x^v) - v(a_x^v)]\, dx.$$

The payment made by each type $\theta \leq \theta_l$ in state u is given by

$$t_\theta(e^u) = t_\theta(a_\theta^v) - u(a_\theta^v) + u(e^u).$$

To understand how these transfers are constructed, consider the types in $[\theta_h, 1]$. The threshold type θ_h has the property that the monopolist is indifferent between offering him a speculative and a non-speculative contract. By Lemma 1, the non-speculative contract is ex-post efficient, hence the monopolist's expected profit from this contract is

$$p[u(e^u) - c(e^u)] + (1-p)[v(e^v) - c(e^v)].$$

The monopolist's expected profit from a speculative contract offered to type θ is given by the integrand in expression (8). It follows that the payment $t_\theta(a_{\theta_h}^u)$ equates these two amounts. The payment made by each type $\theta > \theta_h$ in state u is computed by equating $U(\theta, \theta)$, as given in (7), with the integral representation of the IC_θ constraint,

$$U(\theta, \theta) = \int_{\theta_h}^{\theta} [u(a_x^u) - v(a_x^u)] \, dx,$$

assuming the VI_θ constraint is binding (which allows us to substitute $v(a_\theta^u) - t_\theta(a_\theta^u)$ for $v(e^v) - t_\theta(e^v)$). We then verify that indeed the VI_θ constraint is binding by simply letting $t_\theta(e^v) = t_\theta(a_\theta^u) - v(a_\theta^u) + v(e^v)$. The transfers made by types in $[0, \theta_l]$ are constructed in an analogous way.

Our final result in this section shows that an optimal menu almost always contains a speculative contract.

PROPOSITION 5. *The optimal menu always contains a non-speculative contract. It contains no speculative contracts if and only if $p = 1$ and $u(a) \geq v(a)$ for all a.*

The following examples illustrate the type of contracts that can be found in an optimal menu. These examples also illustrate some of the applications that fit the current environment. 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 illuminate pricing schemes that we observe in reality.

EXAMPLE 1 (Flat-rate versus variable-rate price schemes). A commonly observed menu of pricing schemes offers unlimited consumption at a flat rate, side-by-side with a variable-rate scheme that charges according to consumption. Menus of this kind are found in the telecom industry, where firms often offer a choice between "unlimited calling plans" and plans that condition the per-minute rate on the number of minutes used. Similarly, DVD rental stores offer "unlimited plans" as well as "limited plans." Coffee shops allow customers to purchase a prepaid debit card that entitles them to a certain number of cups, side-by-side with the usual on-the-spot orders. Finally, as in the example presented in the Introduction, driving schools and learning centers that offer prep courses for SAT/GMAT/TOEFL sometimes offer students a choice between a package

that provides a basic course and an option to take additional lessons at an extra charge, and a package in which the student can take as many classes as he wants.

We propose to interpret such a menu as a tool for screening consumers according to their prior beliefs over their future satiation point (where there is no uncertainty about the satiation *payoff*). How many lessons will the consumer need in order to pass a test? Will the consumer need a lot of airtime to communicate with friends and colleagues on the mobile phone, or will short conversations suffice? How many DVDs will he need to watch to satiate his taste for films? A variable-rate contract may be viewed as a bet, where the consumer "wins" if he manages to consume only a small amount and the monopolist "wins" if the consumer ends up consuming a large amount. In contrast, an unlimited consumption contract has no speculative component, because the payment the consumer makes is equal to his satiation payoff and therefore independent of his level of consumption. Consumers who believe that their satiation point is likely to be low prefer a speculative, variable-rate contract, whereas consumers who believe that their satiation point is likely to be high opt for the non-speculative, unlimited-consumption contract.

We illustrate this idea with an extension of the driving school example of the Introduction. Consider a student who has several weeks to prepare for his first College Board exam. This student contemplates purchasing a package from a learning center. Having never taken such an exam before, the student does not know how many lessons he will need to pass the exam. Formally, let a denote the amount of preparation the student acquires (where $a = 1$ means that a consumer takes every available lesson before the exam). The cost to the learning center is given by $c(a) = \frac{1}{4}a$. Let u and v represent the student's willingness to pay in the "low" and "high" states. Specifically, let $v(a) = a$ and $u(a) = \min(2a, 1)$. Let $F(\theta) = \theta$ and $p = \frac{1}{2}$.

From the learning center's point of view, a student of type $\theta > \frac{1}{2}$ ($\theta < \frac{1}{2}$) underestimates (overestimates) the amount of prep work he will require. Since u lies weakly above v, Proposition 2 implies that the learning center will offer a non-speculative contract to types below $\theta_h > \frac{1}{2}$. Such a contract can be implemented by a package that offers an unlimited number of lessons until the exam, for a flat fee of 1. Turning to the speculative contracts in the menu, Proposition 2 implies that $a_\theta^v = 1$ for all $\theta \geq \theta_h$. It can be shown that this proposition also implies that

$$a_\theta^u \in \underset{a \in [0,1]}{\arg\max} [(2\theta - 1)u(a) + (\tfrac{5}{4} - 2\theta)a].$$

Solving this optimization problem yields $a_\theta^u = \frac{1}{2}$, independently of θ. Thus, all types above θ_h end up taking an ex-post efficient action in both states. This in turn implies that the optimal menu contains a single speculative contract, denoted t^s, in addition to the optimal non-speculative contract.

Recall that VI_θ is binding for every $\theta \geq \theta_h$ and that IR_θ is binding for θ_h. Moreover, at the cutoff θ_h, the monopolist is indifferent between the speculative and non-speculative contracts. It follows that $\theta_h = \frac{3}{4}$. The binding constraints IR_{θ_h} and VI_{θ_h} then imply $t^s(1) = \frac{11}{8}$ and $t^s(\frac{1}{2}) = \frac{7}{8}$. The speculative contract t^s can be implemented by a two-part tariff: $t(0) = 0$ and $t(a) = \frac{3}{8} + a$ for $a > 0$.

In this example, flat-rate, unlimited-consumption pricing schemes are non-speculative contracts, while limited plans with extra charges for additional consumption are speculative contracts aimed at consumers who optimistically underestimate their satiation quantity. However, this classification is sensitive to the specification of u and v. For instance, let $u(a) = 2a$ and $v(a) = \min(2a, 1)$. The interpretation of this alternative specification is that in state v, the consumer reaches satiation at $a = \frac{1}{2}$, whereas in state u he is willing to pay for additional units. In this case, there is an optimal menu consisting of a variable-rate scheme, $t(a) = 2a$ for all a, as well as a flat-rate scheme, $t(a) = \frac{7}{4}$ for all a. The former is a non-speculative contract selected by types $\theta \leq \frac{3}{4}$, whereas the latter is a speculative contract selected by types $\theta > \frac{3}{4}$. All consumer types end up taking ex-post efficient actions in both states ($a^u = 1$, $a^v = \frac{1}{2}$). Thus, the classification of flat-rate and variable-rate contracts as speculative or non-speculative depends on the fine details of u and v. ◊

EXAMPLE 2 (Speculative contracts and inefficiency). In many situations, buyers contract with a supplier prior to knowing the exact specifications of the product or service they require, because these depend on the future realization of a state of nature. Bilateral contracts of this nature have been widely studied in the literature. Most often, the focus has been on the hold-up problem that may arise and its various remedies (see Tirole 1999 and the references therein). However, the manner in which these contingent contracts allow for speculation (see Bazerman and Gillespie 1999) has been largely overlooked by theorists. Our next example captures a contracting situation of this kind. We employ this example to demonstrate the subtle effect of speculative contracts on ex-post efficiency.

A seller can provide (at zero cost) a product of 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 functional forms

$$u(a) = 3 - (a - \tfrac{1}{2})^2$$
$$v(a) = 2 - (a - \tfrac{3}{4})^2.$$

That is, the buyer's ideal variety is $a = \frac{1}{2}$ in state u and $a = \frac{3}{4}$ in state v, and, moreover, state u is characterized by higher demand. Let $F(\theta) = \theta$ and $p = \frac{1}{2}$.

For brevity, we focus on the contracts that the seller designs for types $\theta \in [\frac{1}{2}, 1]$. The non-speculative contract for types $\frac{1}{2} \leq \theta < \theta_h$ extracts the entire consumer surplus in each state. Let us turn to the speculative contracts for types $\theta \geq \theta_h$. We first note that for each of these types, a^v_θ can be set to $\frac{3}{4}$. From the proof of Proposition 2, it follows that to compute a^u_θ for $\theta \geq \theta^h$ we need to solve

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

yielding

$$a^u_\theta = \tfrac{5}{4} - \theta.$$

Finally, from the fact that the monopolist is indifferent between betting and not betting with the cutoff type θ_h, it follows that $\theta_h = \frac{3}{4}$.

Note that $a_\theta^u \leq \arg\max u$ for all $\theta \geq \frac{3}{4}$. 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' prior beliefs become more polarized, the contract they sign becomes more inefficient ex-post (in state u).

If the seller could observe the buyer's type θ, he would assign to any buyer type $\theta > \frac{1}{2}$ a speculative contract that induces $a_\theta^v = \frac{3}{4}$ and

$$a_\theta^u = \arg\max[(\theta - p)(u - v) + pu] = \frac{3}{4} - \frac{1}{2}\theta.$$

Compare this with our result that when the seller does not observe θ,

$$a_\theta^u = \frac{1}{2} \text{ for } \theta \in [\frac{1}{2}, \frac{3}{4})$$

and

$$a_\theta^u = \frac{5}{4} - \theta \text{ for } \theta > \frac{3}{4}.$$

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.[4]

Note that in contrast to previous examples, the optimal menu here displays fine discrimination among optimistic types. Specifically, there is a continuum of speculative contracts. To characterize these contracts for $\theta > \frac{3}{4}$, recall that VI_θ is binding for these types. Therefore,

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

This means that the higher the buyer's prior on u, the larger the difference between the payments he makes in the two states.   ◊

5. An alternative environment

Our analysis in the previous section highlighted the necessity of non-common priors in generating price discrimination at the signing stage. The key assumption behind this result is property (*). Although this property can accommodate a number of situations of economic interest, it does rule out environments that are commonly studied in the literature. In particular, for the canonical textbook environment in which both u and v are increasing, concave functions, property (*) requires that there be no positive gains from trade in state v. In this section we characterize the optimal menu for this canonical environment in the absence of property (*). Let $u(0) = v(0) = 0$, $u'(a) > v'(a)$ for all a, and u'', $v'' < 0$. In addition, assume that F satisfies the monotone hazard rate condition.

---

[4]These welfare implications are not general, but a consequence of certain features of the payoff structure: (i) $v(a) \equiv u(a - d)$, where d 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).

5.1 *Optimal contracts under complete information*

We begin by considering the optimal contract when the monopolist faces a consumer whose type θ is known. Proposition 1 characterizes this contract for the three possible case, $\theta > p$, $\theta < p$, and $\theta = p$. It is straightforward to show that the optimal contracts under complete information for types $\theta \neq p$ in the current environment are precisely the same as those characterized by Proposition 1. Property (*) is applied in the original proof only to verify that the UI_θ and VI_θ constraints are satisfied. In the current environment, this follows from the assumption that u lies above v, and that $u(a) - v(a)$ increases with a.

Property (*) also plays a key role in showing that when $\theta = p$, the monopolist offers the consumer an ex-post efficient contract that extracts the entire ex-post surplus in each state. In the present environment, the optimal contract under complete information continues to be ex-post efficient. However, since property (*) does not hold, the UI_θ and VI_θ constraints would be violated if the consumer were required to surrender his entire ex-post surplus in each state. Because the contract t_θ extracts the consumer's entire *ex-ante* surplus, the transfers $t_\theta(e^u)$ and $t_\theta(e^v)$ lie on the line

$$\theta\, t_\theta(e^u) + (1 - \theta) t_\theta(e^v) = \theta\, u(e^u) + (1 - \theta) v(e^v).$$

The UI_θ and VI_θ constraints further imply that these transfers must lie between the point given by (1)–(2) and the point given by (3)–(4).

5.2 *Optimal contracts under incomplete information*

We now turn to the monopolist's problem when consumer types are unobserved. The following proposition characterizes optimal menus of contracts in this case.

PROPOSITION 6. *The monopolist can design optimal menus with the following properties. There exists a pair of cutoffs, $\theta^* \in [0,1]$ and $\bar{\theta} \in [\theta^*, 1)$, such that*

   (i) *if $\theta^* > 0$, then for all $\theta < \theta^*$, $a_\theta^u = e^u$, $t_\theta(e^u) = u(e^u)$, $a_\theta^v = 0$, and $t_\theta(0) = 0$*

   (ii) *if $\bar{\theta} > \theta^*$, then for all $\theta \in [\theta^*, \bar{\theta})$, the induced actions are*

$$a_\theta^u = a_\theta^v$$
$$= \underset{a \in [0,1]}{\arg\max}\left\{[\psi(\theta) - p][u(a) - v(a)] + p[u(a) - c(a)] + (1-p)[v(a) - c(a)]\right\} \quad (9)$$

and the payments for these actions are given by

$$t_{\theta^*}(a_{\theta^*}^u) = p[u(e^u) - c(e^u) + c(a_{\theta^*})] - (1-p)[v(e^v) - c(e^v) - v(a_{\theta^*})]$$
$$t_\theta(a_\theta^u) = v(a_\theta^u) + \theta[u(a_\theta^u) - v(a_\theta^u)] - \int_{\theta^*}^{\theta} [u(a_x^u) - v(a_x^u)]\, dx \quad (10)$$

and $t_\theta(a_\theta^v) = t_\theta(a_\theta^u)$

(iii) for all $\theta \in [\bar{\theta}, 1]$, the induced actions are $a^v_\theta = e^v$ and

$$a^u_\theta \in \underset{a \in [0,1]}{\arg\max}\{[\psi(\theta) - p][u(a) - v(a)] + p[u(a) - c(a)]\} \tag{11}$$

and the payments for these transfers are given by (10) and $t_\theta(e^v) = t_\theta(a^u_\theta) - v(a^u_\theta) + v(e^v)$.

Thus, the optimal menu partitions the set of types into at most three regions. The types in the lowest region are assigned a contract that essentially excludes them from transacting with the monopolist in the low state v. This contract guarantees *all* types their reservation payoff. Consumer types in the middle region are assigned contracts that commit them to the same action in both states. This action may be different for each type in this region, and it may be inefficient in both states. Finally, types in the highest range are assigned contracts that induce the efficient action in state v and a possibly inefficient action in state u. Thus, the contracts assigned to these types retain the feature of the complete-information optimal contracts that efficiency is maintained in the state the monopolist deems more likely than the consumer.

The absence of property (*) implies that some of the arguments we employed in proving Proposition 2 are not valid in the present environment. Consequently, the proof of Proposition 6 is somewhat more involved than the proof of Proposition 2. In what follows we give a brief outline of the proof.

One important implication of relaxing property (*) is that the monopolist cannot offer a non-speculative contract that satisfies ex-post efficiency and fully extracts consumer surplus in each state. The reason is that such a contract would violate one of the second-period incentive constraints (*UI* or *VI*). We therefore need to characterize the optimal contract among those that fully extract consumer surplus in each state. Claims 1–5 in the proof establish that this contract is ex-post efficient in state u and induces no-trade in state v. Consumer types in the range $[0, \theta^*)$ opt for this contract.

Another consequence of relaxing property (*) is that the UI_θ constraint is *not* implied by the fact that the IC_θ constraint is satisfied and the VI_θ constraint is binding. To see why, note that the UI_θ constraint holds if

$$u(a^u_\theta) - u(a^v_\theta) \geq t_\theta(a^u_\theta) - t_\theta(a^v_\theta).$$

If the VI_θ constraint is binding, this inequality may be rewritten as

$$u(a^u_\theta) - v(a^u_\theta) \geq u(a^v_\theta) - v(a^v_\theta). \tag{12}$$

If we were to solve the monopolist's relaxed problem, assuming the solution necessarily satisfies the UI_θ constraint, then we would obtain (by the same arguments as in Proposition 2) that $a^v_\theta = e^v$. Recall from our discussion of Proposition 2 that the IC_θ constraint implies that the left-hand side of (12) is non-negative. Since $a^v_\theta = e^v$, property (*) implies that the right-hand side of (12) is negative. Hence, the UI_θ constraint is satisfied. However, once we relax property (*), then it may very well be the case that $u(e^v) > v(e^v)$,

such that the right-hand side of (12) is positive. Hence we can no longer ignore the UI_θ constraint when maximizing the monopolist's expected profit.

To tackle this problem, we partition the set of types above θ^* into two subsets: the set of types for whom $u(a_\theta^u) - v(a_\theta^u) \geq u(e^v) - v(e^v)$ (we denote this set by Θ^+), and the set of types for whom this inequality does not hold (we denote this set by Θ^-). If Θ^- is empty, then the monopolist's problem is essentially the same as the problem of designing the optimal menu for types above p in an environment in which property (∗) holds. The question is, what is the optimal menu when Θ^- is not empty?

To answer this question, we conjecture that the problem of characterizing contracts in Θ^+ is independent of the problem of characterizing contracts in Θ^-. This means that each type in Θ^+ is assigned a speculative contract that maximizes the monopolist's expected profit, assuming all the necessary constraints are satisfied. The solution to this maximization problem is that $a_\theta^v = e^v$ and a_θ^u is given by (11). In addition, each type in Θ^- is assigned a speculative contract that maximizes the monopolist's expected profit *subject to the constraint* that both UI_θ and VI_θ are binding. This, together with our assumption on u and v, implies that types in Θ^- must choose precisely the same action in each of the states.

The most involved part of the proof is to verify that the above conjecture is correct. The main step is to show that the incentive compatibility constraints hold. In particular, we need to show that no type in Θ^+ has an incentive to mimic a type in Θ^-, and vice versa. To achieve this, we show that (i) any pair of types, $\theta \in \Theta^+$, $\theta' \in \Theta^-$ satisfies $\theta' > \theta$ and $a_\theta^u > e^v > a_{\theta'}^u$, and (ii) $u(a_\theta^u) - v(a_\theta^u)$ is non-decreasing on $\Theta^- \cup \Theta^+ = [\theta^*, 1]$. Since

$$U(\phi, \theta) = \theta[u(a_\phi^u) - v(a_\phi^u)] + [v(a_\phi^v) - t_\phi(a_\theta^v)],$$

(i) and (ii) together allow us to establish incentive-compatibility (see Claims 6 and 7 in the proof).

The following example illustrates some aspects of the characterization of optimal menus. In particular, it demonstrates that if the consumer's marginal utilities in the two states are sufficiently wide apart, the optimal menu in the common-prior benchmark is a singleton, whereas in our model it contains multiple contracts. The restrictiveness of the example allows us to provide a simple characterization of the optimal menu, including a closed solution for the cutoffs.

PROPOSITION 7. *Let* $u \equiv kv$, *where* $k > 2$; *let* $F(\theta) = \theta$; $p > \frac{1}{3}$; *and* $c(a) = 0$ *for all* a. *Then, there exists an optimal menu of contracts with the following properties.*

(i) $\theta^* = \bar{\theta} = (k(p+1) - 2)/(2k - 2)$.

(ii) *For every* $\theta < \theta^*$, $a_\theta^u = 1$, $t_\theta(1) = u(1)$, $a_\theta^v = 0$, *and* $t_\theta(1) = 0$.

(iii) *For every* $\theta > \theta^*$, $a_\theta^u = a_\theta^v = 1$.

(iv) *In contrast, suppose the monopolist believes that the consumer's private information (his prior belief* θ*) is the true objective probability of* u. *Then, the optimal menu is a singleton, given by (ii) above.*

This example is simple in the sense that the cutoffs θ^* and $\bar{\theta}$ coincide. The simplification is due to the feature that $u-c$, $v-c$, and $u-v$ are all strictly increasing functions. The characterization of optimal menus in the common-prior and non-common-prior cases shares several features with this example. In particular, in both cases low types are offered a contract that induces no trade in state v and efficient trade in state u, and extracts the consumer's entire surplus. The crucial difference is that in our model, the monopolist evaluates the expected revenue from a consumer of type θ who accepts this contract at $pu(1)$, whereas in the common-prior case, the monopolist evaluates this contract at $\theta u(1)$. This in turn implies that in the common-prior case, all consumer types are offered this contract, whereas in the case of non-common priors the optimal menu exhibits price discrimination. This qualitative distinction holds for any p as long as there exists some sufficiently large $k > 1$ such that $u(a) \geq kv(a)$ for all $a > 0$.

6. Concluding remarks

We have argued that menus of non-linear pricing schemes in monopolistic environments can be usefully interpreted as a consequence of the monopolist's attempt to screen the consumer's prior belief regarding his future willingness to pay. In particular, we have demonstrated that in certain environments, price discrimination of this sort emerges only when the monopolist and the consumer have different priors. Furthermore, in some of our examples the consumer's actions are independent of the pricing scheme he selects from the menu. In these cases, incentive provision cannot be the explanation for price discrimination. Rather, some of the pricing schemes on the menu are bets aimed at consumers whose prior beliefs are sufficiently different from the monopolist's.

6.1 *The interpretation of non-common priors*

We motivated the assumption of non-common priors by the idea that consumers may have incorrect beliefs due to market inexperience or inherent biases (such as underestimating the level of consumption required for satiation). Indeed, there is a rich psychological literature on people's limited ability to forecast future tastes (for a comprehensive and popular exposition, see Gilbert 2006).

However, we wish to re-emphasize that for a consistent interpretation of the model, it is not enough to assume that consumers have biased prior beliefs; one has to assume in addition that consumers do not believe that the monopolist is better informed—otherwise, they would regard the menu as a signal of the monopolist's prior and use it to update their beliefs. One reason consumers may hold this belief is that they are unaware of their bias, and therefore (erroneously) believe that no other party could be better informed about their future tastes. Alternatively, consumers may be aware of a general tendency to hold biased beliefs, yet they may be (erroneously) confident that they themselves do not suffer from this problem.

Of course, the judgment that consumer types $\theta \neq p$ hold incorrect beliefs is not essential. One could simply assume that the two parties hold different priors, without

making any assumption as to which of the two parties, if any, is right. Under that interpretation, the monopolist and the consumer simply "agree to disagree"—it is common knowledge between them that they hold different priors. However, we find the original interpretation more interesting, partly because it generates non-trivial predictions. For example, if we believe that more experienced consumers are less likely to hold biased beliefs, then our model implies that inexperienced consumers will tend to select speculative contracts on the menu while experienced consumers will tend to select non-speculative contracts.

6.2 *The independence between the two parties' priors*

The assumption that the monopolist's prior p is independent of the consumer's type θ is quite restrictive. Even if the monopolist believes that the consumer's prior belief is biased, why should he rule out the possibility that the consumer is better informed about his future needs? A more satisfying assumption is that when the consumer's type is θ, the monopolist believes that the consumer's second-period utility function will be u with probability $p(\theta)$. This generalization subsumes the common-prior case and our model as special cases, namely $p(\theta) = \theta$ and $p(\theta) = p$, respectively.

It can be shown that the analysis of Section 4 is robust to this generalization, under the assumption that $\theta - p(\theta)$ is a strictly increasing function that attains negative (positive) values for low (high) θ. Let θ^0 denote the unique solution to the equation $p(\theta) = \theta$. Then, the characterization of optimal menus is qualitatively the same as in the model of Section 3, where θ^0 plays the role of p. The task of characterizing optimal menus under more general classes of $p(\theta)$ is left for future work.

6.3 *More than two states*

We have restricted our analysis to an environment in which there are only two possible states of nature. A natural extension of the model is to allow for more than just two possible utility functions in the second period. Since the consumer type in our model is given by his prior beliefs, such an extension requires us to consider multi-dimensional types. Multidimensional mechanism-design introduces a number of technical issues that are secondary to the main ideas of the current paper, and we therefore leave this extension for future research.

APPENDIX

PROOF OF PROPOSITION 1. Since the monopolist knows the consumer's type, it is straightforward to show that the IR_θ constraint is binding. We may therefore express the monopolist's objective function as

$$\theta[u(a_\theta^u) - c(a_\theta^u)] + (1-\theta)[v(a_\theta^v) - c(a_\theta^v)] + \delta[t(a_\theta^v) - t(a_\theta^u) + c(a_\theta^u) - c(a_\theta^v)], \quad (13)$$

where $\delta \equiv p - \theta$.

Suppose $\delta < 0$. Then to maximize the above expression subject to the UI_θ and VI_θ constraints, the VI_θ constraint must bind. By substituting $v(a_\theta^v) - v(a_\theta^u)$ for $t(a_\theta^v) - t(a_\theta^u)$

in (13) and solving for a_θ^u and a_θ^v that maximize the resulting expression, we obtain the values given in the statement of the proposition.

It remains to verify that the UI_θ constraint holds at the solution for a_θ^u and a_θ^v, i.e., that

$$u(a_\theta^u) - v(a_\theta^u) \geq u(e^v) - v(e^v).$$

By property (*), $u(e^v) - v(e^v) \leq 0$. Suppose that the optimal contract under complete information to type θ, denoted t_θ, induces an action $a_\theta^u = a^*$ satisfying $u(a^*) - v(a^*) < 0$. Suppose the monopolist offered the consumer a different contract, t_θ', that satisfies $t_\theta'(a^*) = \infty$, $t_\theta'(e^u) = t_\theta(a^*)$, and $t_\theta'(e^v) = t_\theta(e^v)$. By construction, t_θ' satisfies the IR_θ and VI_θ constraints (both constraints are binding). By the assumption that $\max(u - c) > \max(v - c)$, $u(e^u) - v(e^u) > 0$. Since the VI_θ constraint is binding, it follows that t_θ' also satisfies the UI_θ constraint.

To reach a contradiction, we need to show that t_θ' generates a higher expected profit than does t_θ. Because t_θ' and t_θ induce the same action and generate the same profit in state v, it suffices to show that t_θ' generates a higher profit in state u. Define

$$W(a) \equiv p[u(a) - c(a)] - \delta[u(a) - v(a)].$$

Since both the IR_θ and VI_θ constraints are binding, the monopolist's objective function may be expressed as

$$W(a_\theta^u) + (1 - \theta)[v(a_\theta^v) - c(a_\theta^v)].$$

By the definition of e^u, $u(e^u) - c(e^u) \geq u(a^*) - c(a^*)$. By property (*),

$$u(e^u) - v(e^u) > 0 \geq u(a^*) - v(a^*).$$

Hence, since $\delta < 0$ we have $W(e^u) > W(a^*)$, in contradiction to our assumption that t_θ is an optimal contract under complete information.

The case of $\delta > 0$ is proven in a similar manner.

Suppose next that $\delta = 0$. Then the monopolist's objective function, given by (13), is maximized at the ex-post efficient actions, e^u and e^v. By property (*), the monopolist can extract the entire ex-post consumer surplus in each state without violating the UI_θ or VI_θ constraints.   □

PROOF OF LEMMA 1.   Assume an optimal menu includes a set of non-speculative contracts T (which may be a singleton) that generate an expected profit for the monopolist that is lower than his expected profit from the following contract t^*, the optimal contract under complete information for $\theta = p$:

$$t^*(a) = \begin{cases} u(e^u) & \text{if } a = e^u \\ v(e^v) & \text{if } a = e^v \\ \infty & \text{if } a \notin \{e^u, e^v\}. \end{cases}$$

This contract has two important features: it yields zero expected payoff to all types, and it satisfies the UI and VI constraints of any type who chooses it. Let Θ^T denote the set of types whose most preferred contract in the menu is in T.

Consider amending the original menu by replacing all the contracts in T with t^*. Since the original menu is assumed to be optimal, every speculative contract in it must satisfy the *UI* and *VI* constraints. Hence, every contract in the new menu satisfies these constraints.

We wish to show that there is a way to assign a contract in the new menu to each consumer type, such that (i) no consumer type is assigned a contract with a negative expected payoff, (ii) the assigned contract is weakly preferred to any other contract in the menu, and (iii) the new menu generates a higher expected profit than the original menu.

By our construction of t^*, this is straightforward if we can assign t^* to every type in Θ^T, while assigning all other types their most preferred contract in the original menu. Suppose there is a positive measure of types in Θ^T who strictly prefer a contract $t \notin T$ to the contract t^*. If t is non-speculative, then by the definition of T, the expected profit from t is larger or equal to the expected profit from t^*. If t is speculative, then its expected profit is at least as high as the expected profit from t^* (otherwise the monopolist is able to increase his expected profit by replacing the original menu with one that consists of t^* and all the original speculative contracts except for the least profitable one).

Suppose next that there is a positive measure of types outside Θ^T whose most preferred contract in the new menu is t^*. Then each of these types must have obtained a strictly negative expected payoff from his most preferred contract in the original menu, contradicting the *IR* constraint. $\qquad\square$

PROOF OF PROPOSITION 2.   *Proof of (i).* Note first that by IR_p, type $\theta = p$ cannot be assigned a speculative contract. Our proof relies on the following claim.

CLAIM. *If an optimal menu assigns a speculative contract to a type $\theta \neq p$, then $D_\theta^u > 0$ and $D_\theta^v < 0$ when $\theta > p$, while $D_\theta^u < 0$ and $D_\theta^v > 0$ when $\theta < p$.*

PROOF.   Suppose $\theta > p$. Let $x \equiv \theta - p$. Then by Definition 1,

$$(\theta - x)D_\theta^u + (1 - \theta + x)D_\theta^v < 0.$$

By IR_θ,

$$\theta D_\theta^u + (1 - \theta)D_\theta^v \geq 0.$$

Because $\theta > x > 0$, these two inequalities imply that $D_\theta^u > 0$ and $D_\theta^v < 0$. A similar argument applies for $\theta < p$. $\qquad\triangleleft$

We now show that if a type $\theta > p$ is assigned a speculative contract, then a higher type is also assigned a speculative contract. Assume the optimal menu assigns a speculative contract to type θ and a non-speculative contract to some type $\phi > \theta$. Then by Lemma 1, $U(\phi, \phi) = 0$. By $IC_{\phi, \theta}$, a consumer of type $\phi > \theta$ satisfies $U(\phi, \phi) \geq U(\theta, \phi)$. By definition, $U(\theta, \phi) = \phi(D_\theta^u - D_\theta^v) + D_\theta^v$. Since, by assumption, θ is assigned a speculative contract, the Claim implies that $D_\theta^u - D_\theta^v > 0$ (recall that $\theta > p$). This, in turn, implies that $U(\theta, \phi) > U(\theta, \theta)$. Since $U(\theta, \theta) \geq 0$, we have reached a contradiction. A similar argument applies for types below p.

Proof of (ii). Assume the monopolist 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 $[p, 1]$ (respectively $[0, p]$) weakly prefers his assigned contract to every contract in T^- (respectively T^+).

By part (i) and Lemma 1, there exist $\theta^l \in [p, 1]$ and $\theta^h \in [0, p]$ such that $U(\theta, \theta) = 0$ for all $\theta \in (\theta_l, \theta_h)$. By the Claim, $D_\theta^v > D_\theta^u$ for all $\theta \le \theta_l$, while $D_\theta^u > D_\theta^v$ for all $\theta \ge \theta_h$. In addition, the contract assigned to every $\theta \notin (\theta_l, \theta_h)$ violates IR_p. It follows that for every pair of types θ, ϕ such that $\theta \le p$ and $\phi \ge \theta_h$ we have $U(\phi, \theta) < 0$. Similarly, $U(\phi', \theta') < 0$ for every pair of types θ', ϕ' with $\theta' \ge p$ and $\phi' \le \theta_l$.

By assumption, T^+ (respectively T^-) maximizes the monopolist's expected profit conditional on $\theta \in [p, 1]$ (respectively $\theta \in [0, 1]$). Hence the union of T^+ and T^- maximizes the *unconditional* expected profit of the monopolist. In addition, T^+ satisfies the *IR, IC, UI,* and *VI* 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 (iii). Consider some type $\theta \ge \theta_h$. By part (ii), this type is assigned a speculative contract. Hence, by the Claim, $D_\theta^u > 0$ and $D_\theta^v < 0$. Let $\Theta^* = \{\theta \ge \theta_h : a_\theta^v \ne e^v\}$. Assume the optimal menu has the property that Θ^* is non empty. Consider modifying the original menu by changing each contract t_θ to a new contract t_θ' that differs from the original contract only in two actions, a_θ^v and e^v, such that $t_\theta'(a_\theta^v) = \infty$ and $t_\theta'(e^v)$ satisfies

$$v(e^v) - t_\theta'(e^v) = v(a_\theta^v) - t_\theta(a_\theta^v). \tag{14}$$

Because the consumer's net payoff in state v under the modified contract is the same as it is under the original contact, all the *IR* and *IC* constraints, as well as VI_θ, continue to hold. To see that UI_θ is also satisfied, note that because $\theta > p$ and t_θ is a speculative contract, it follows from the Claim that $D_\theta^u > 0$ and $D_\theta^v < 0$. Hence by (14), $u(a_\theta^u) - t(a_\theta^u) \ge v(e^v) - t_\theta'(e^v)$. By property (*), this inequality implies UI_θ. Note also that conditional on the consumer's type being θ, the contract t_θ' generates a higher expected profit to the monopolist. This follows from the observation that (14), together with our assumption on a_θ^v, implies $t_\theta'(e^v) - c(e^v) > t_\theta(a_\theta^v) - c(a_\theta^v)$.

It remains to show that the VI_θ constraint is binding. Note that by the Claim, $D_\theta^u - D_\theta^u > 0$. In what follows, we adopt Krishna's (2002, 63–66) derivation of incentive compatibility for direct mechanisms. Define $q(\theta) \equiv D_\theta^u - D_\theta^v$ and $m(\theta) \equiv -D_\theta^v$. The optimal menu is incentive compatible if for all types θ and ϕ,

$$V(\theta) \equiv \theta q(\theta) - m(\theta) \ge \phi q(\theta) - m(\phi).$$

By the Claim, $q(\theta) \ge 0$ for all $\theta \ge \theta_h$ (while $q(\theta) = 0$ for all $\theta \in (\theta_l, \theta_h)$ and $q(\theta) \le 0$ for all $\theta \le \theta_l$). Hence, the left-hand side of the above inequality is an affine function of the true value θ. Incentive compatibility implies that for all $\theta \ge \theta^*$,

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

I.e., $V(\theta)$ is a maximum of a family of affine functions, and hence is convex on $[\theta^*, 1]$ (since $D_\theta^u = D_\theta^v = 0$ for all types in $[0, \theta^*)$, $V(\theta) = 0$ for all these types).

From the standard argument it follows that for all $\theta > \theta^*$,

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

By part (i), $V(\theta_h) = 0$, and so we obtain $U(\theta, \theta) = \int_{\theta_h}^{\theta} q(x) \, dx$ for $\theta \geq \theta^*$.

Assume there is an optimal menu in which the VI_θ constraint does not bind for a positive measure of types above θ_h. Consider the types $\theta \geq \theta_h$. Let δ_θ denote the slack in the VI_θ constraint of type θ. Then

$$t_\theta(a_\theta^v) - t_\theta(a_\theta^u) = v(a_\theta^v) - v(a_\theta^u) - \delta_\theta.$$

By the previous paragraph, if type $\theta \geq \theta_h$ is assigned a speculative contract, which is incentive compatible, then

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

Assume the optimal menu satisfies $\delta_\theta > 0$ for some positive measure of types in $[\theta_h, 1]$. Consider amending the menu by changing only $t_\theta(a_\theta^v)$ for all types $\theta \geq \theta_h$ such that the new transfer is equal to $t_\theta(a_\theta^v) + \delta_\theta$, making the VI_θ constraint bind for all these types. Clearly, this change does not violate the UI_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. Since this new menu only increases the monopolist's revenue without changing his costs (the new menu induces exactly the same actions for each type as does the original menu), it increases the expected profit, in contradiction to our assumption that the original menu was optimal.

Proof of (iv).   The proof is essentially the same as that for part (iii).   □

PROOF OF PROPOSITION 3.   We provide a detailed proof for the distribution of types F, conditional on $\theta \in [p, 1]$. An optimal menu for types distributed on $[0, p)$ is derived by essentially the same argument.

Assume that F satisfies the monotone hazard rate condition. Proposition 2 implies the following. First, we may restrict attention to menus that induce $a_\theta^v = e^v$ for $\theta \geq p$. Second, the monopolist's problem may be written as

$$\max_{\theta_h, \{a_\theta^u, t_\theta^u\}_{\theta \in [\theta_h, 1]}} \left\{ \left[\frac{F(\theta_h) - F(p)}{1 - F(p)} \right] \{ p[u(e^u) - c(e^u)] + (1-p)[v(e^v) - c(e^v)] \} \right.$$

$$\left. + \int_{\theta_h}^{1} \{ p[t_\theta(a_\theta^u) - c(a_\theta^u)] + (1-p)[t_\theta(e^v) - c(e^v)] \} \left[d \frac{F(\theta)}{1 - F(p)} \right] \right\} \tag{16}$$

subject to the IR_θ, $IC_{\theta,\phi}$, UI_θ, and VI_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 these constraints are satisfied (moreover, we assume the VI_θ constraint is binding for all $\theta \geq \theta_h$), and then checking that this is true at the solution we obtain.

For the optimal menu, the VI_θ constraint must bind for all $\theta \geq \theta_h$, hence

$$t_\theta(a_\theta^u) = t_\theta(a_\theta^v) + v(a_\theta^u) - v(e^v). \tag{17}$$

This allows us to simplify the objective function of the monopolist rewriting (16) as

$$\int_{\theta_h}^1 \{(1-p)[t_\theta(e^v) - c(e^v) - t_\theta(a_\theta^u) + c(a_\theta^u)] + t_\theta(a_\theta^u) - c(a_\theta^u)\}\left[d\frac{F(\theta)}{1-F(p)}\right]$$

$$= \int_{\theta_h}^1 \{(1-p)[v(e^v) - c(e^v) - v(a_\theta^u) + c(a_\theta^u)] + t_\theta(a_\theta^u) - c(a_\theta^u)\}\left[d\frac{F(\theta)}{1-F(p)}\right].$$

By (15)—given that VI_θ is binding—and since $U(\theta,\theta) = \theta q(\theta) - m(\theta)$,

$$t_\theta(a_\theta^v) = \theta q(\theta) + v(e^v) - \int_{\theta_h}^\theta q(x)\,dx. \tag{18}$$

Substituting this expression into (17), we obtain

$$t_\theta(a_\theta^u) = \theta q(\theta) - \int_{\theta_h}^\theta q(x)\,dx + v(a_\theta^u). \tag{19}$$

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

$$\max_{\theta_h, \{a_\theta^u\}_{\theta \in [\theta_h,1]}} \left\{ [F(\theta_h) - F(p)]p[u(e^u) - c(e^u)] \right.$$

$$\left. + \int_{\theta_h}^1 \left\{ \theta q(\theta) - \int_{\theta_h}^\theta q(x)\,dx + p[v(a_\theta^u) - c(a_\theta^u)] \right\} dF(\theta) \right\}$$

Note that we have taken the constant $(1-p)[v(e^v) - c(e^v)]$ out of the objective function. By interchanging the order of integration in

$$\int_{\theta_h}^1 \int_{\theta_h}^\theta q(x)\,dx$$

we can rewrite this expression as

$$\int_{\theta_h}^1 [1 - F(\theta)]q(\theta)\,d\theta.$$

We can thus replace

$$\int_{\theta_h}^1 [\theta q(\theta) - \int_{\theta_h}^\theta q(x)\,dx]\,dF(\theta)$$

with

$$\int_{\theta_h}^1 \psi_F(\theta) q(\theta)\,dF(\theta),$$

where $\psi_F(\theta)$ is defined in (5). Because VI_θ is binding we can replace $q(\theta)$ with $u(a_\theta^u) - v(a_\theta^u)$ to obtain the objective function

$$\max_{\theta_h, \{a_\theta^u\}_{\theta \in [\theta_h,1]}} \left\{ [F(\theta_h) - F(p)] p[u(e^u) - c(e^u)] \right. \tag{20}$$
$$\left. + \int_{\theta_h}^1 \{\psi_F(\theta)[u(a_\theta^u) - v(a_\theta^u)] + p[v(a_\theta^u) - c(a_\theta^u)]\}\,dF(\theta) \right\}$$

Because F satisfies the monotone hazard rate property, it follows that the optimal a_θ^u for each $\theta \in [\theta_h, 1]$ is given by (6).

The optimal cutoff θ_h is the value of θ that solves a first-order condition, which represents the monopolist's indifference between the optimal non-speculative contract described in Lemma 1 and the second-best contract. That is, θ_h satisfies

$$p[u(e^u) - c(e^u)] = \max_a \{\psi_F(\theta)[u(a) - v(a)] + p[v(a) - c(a)]\}.$$

By the monotone hazard rate property, $\psi_F(\theta)$ is increasing in θ. Therefore there is a unique value of θ_h that solves this equation, given by $\psi_F(\theta_h) = p$.

It remains to verify that the menu we constructed indeed satisfies all the necessary constraints. We begin by verifying that for type θ_h both the IR_{θ_h} and the VI_{θ_h} constraints are binding. We do this by letting $t_{\theta_h}(a_{\theta_h}^v)$ and $t_{\theta_h}(a_{\theta_h}^u)$ be the solutions to the system of equations

$$\theta_h[u(a_{\theta_h}^u) - t_{\theta_h}(a_{\theta_h}^u)] + (1 - \theta_h)[v(e^v) - t_{\theta_h}(e^v)] = 0$$
$$v(e^v) - t_{\theta_h}(e^v) = v(a_{\theta_h}^u) - t_{\theta_h}(a_{\theta_h}^u).$$

Similarly, we verify that the VI_θ constraint of all types $\theta > \theta_h$ is binding by setting $t_\theta(a_\theta^v)$ and $t_\theta(a_\theta^u)$ according to equations (18) and (19).

In verifying the remaining constraints, we rely on the following claim, which establishes two important properties of $q(\theta)$.

CLAIM. *$q(\theta)$ is non-decreasing, and $q(\theta_h) \geq 0$.*

PROOF.   Since, by construction, the VI_θ constraint is binding for all $\theta \geq \theta_h$, $q(\theta)$ is non-decreasing for these types if and only if $u(a_\theta^u) - v(a_\theta^u)$ is non-decreasing for all $\theta \geq \theta_h$. Consider a pair of types ϕ, θ such that $\phi > \theta$. By construction,

$$\psi_F(\phi)[u(a_\phi^u) - v(a_\phi^u)] + p[v(a_\phi^u) - c(a_\phi^u)] \geq \psi_F(\phi)[u(a_\theta^u) - v(a_\theta^u)] + p[v(a_\theta^u) - c(a_\theta^u)]$$

and

$$\psi_F(\theta)[u(a_\theta^u) - v(a_\theta^u)] + p[v(a_\theta^u) - c(a_\theta^u)] \geq \psi_F(\theta)[u(a_\phi^u) - v(a_\phi^u)] + p[v(a_\phi^u) - c(a_\phi^u)]$$

Adding these two inequalities and cancelling common terms yields

$$[\psi_F(\phi) - \psi_F(\theta)]\{[u(a_\phi^u) - v(a_\phi^u)] - [u(a_\theta^u) - v(a_\theta^u)]\} \geq 0.$$

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

Let us now show that $q(\theta_h) \geq 0$. Recall that we determined θ_h as the unique solution of $\psi_F(\theta) = p$. Together with (6), this means that $a_{\theta_h}^u = e^u$. By assumption, $u(e^u) - v(e^u) > 0$, hence $q(\theta_h) \geq 0$.   ◁

We now verify that the $IC_{\theta,\phi}$ constraints are satisfied for every distinct pair of types $\theta, \phi \in [\theta_h, 1]$. Since, by the Claim, $q(\theta) \geq 0$ for all $\theta \geq \theta_h$, incentive compatibility is equivalent to the requirement that $q(\theta)$ is non-decreasing on $[\theta_h, 1]$ (see Krishna 2002, 68), which is the case, according to the Claim.

Next, we verify that the UI_θ constraint is satisfied for all $\theta \in [\theta_h, 1]$. Since $a_\theta^v = e^v$ and since the VI_θ constraint is binding by construction, the UI_θ constraint is satisfied if and only if $u(a_\theta^u) - v(a_\theta^u) \geq u(e^v) - v(e^v)$. By property (*), $u(e^v) - v(e^v) < 0$ and by the Claim, $u(a_\theta^u) - v(a_\theta^u) \geq 0$ for all $\theta \geq \theta_h$.

Finally, we verify that the IR_θ constraint is satisfied for all $\theta \in [\theta_h, 1]$. We have already shown, by construction, that this constraint is binding for $\theta = \theta_h$, i.e., that $U(\theta_h, \theta_h) = 0$. By the Claim, $U(\theta_h, \theta) > U(\theta_h, \theta_h)$ for all $\theta > \theta_h$ and by IC_{θ,θ_h}, $U(\theta, \theta) \geq U(\theta_h, \theta)$.   □

PROOF OF PROPOSITION 4.   At type θ_h the monopolist is indifferent between offering a speculative contract and a non-speculative contract. By Lemma 1, the monopolist's expected profit from a non-speculative contract is

$$p[u(e^u) - c(e^u)] + (1 - p)[v(e^v) - c(e^v)].$$

From the proof of Proposition 3 it follows that the monopolist's expected profit from a speculative contract offered to a type θ is given by

$$(1 - p)[v(e^v) - c(e^v) - v(a_{\theta_h}^u) + c(a_{\theta_h}^u)] + t_{\theta_h}(a_{\theta_h}^u) - c(a_{\theta_h}^u)$$

(see (16)). The transfer $t_{\theta_h}(a_{\theta_h}^u)$ is then derived by equating these two amounts. The transfer $t_\theta(a_\theta^u)$ for each $\theta > \theta_h$ is computed according to (19) from Proposition 3. For each $\theta \geq \theta_h$ we let $t_\theta(e^v) = t_\theta(a_\theta^u) - v(a_\theta^u) + v(e^v)$, thereby guaranteeing that the VI_θ constraint binds for these types. The transfers made by types in $[0, \theta_l]$ are constructed in an analogous way.   □

PROOF OF PROPOSITION 5.   Assume $p < 1$. Note first that an optimal menu cannot contain only speculative contracts since this would violate the IR constraints of types in the neighborhood of p. Suppose the optimal menu includes *no* speculative contracts. Then

by Lemma 1, we may restrict attention to the optimal contract under complete informa-
tion for type $\theta = p$, which guarantees each type his reservation payoff. Consider some
$0 < \varepsilon < u(e^u) - v(e^u)$ that satisfies

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

Then

$$p[u(e^u) - \varepsilon] + (1-p)[v(e^v) + u(e^u) - v(e^u) - \varepsilon] > pu(e^u) + (1-p)v(e^v). \qquad (21)$$

In addition, there exists $\bar{\theta} < 1$ sufficiently close to 1 that satisfies

$$\bar{\theta}\varepsilon - (1 - \bar{\theta})[u(e^u) - v(e^u) - \varepsilon] = 0. \qquad (22)$$

Suppose that the monopolist added to the original menu the following contract t:

$$t(a) = \begin{cases} u(e^u) - \varepsilon & \text{if } a = e^u \\ v(e^v) + u(e^u) - v(e^u) - \varepsilon & \text{if } a = e^v \\ \infty & \text{if } a \notin \{e^u, e^v\}. \end{cases}$$

By (22), all types $\theta > \bar{\theta}$ strictly prefer t to t^*, while the opposite is true for types below $\bar{\theta}$.
By (21) and the continuity of F, the expected profit from the new menu is strictly higher
than from the original menu, which is a contradiction.

   Assume next that $p = 1$. Suppose v lies strictly above u at some $a \in [0,1]$. Then we
can apply essentially the same argument as the one given in the previous paragraph to
show that the optimal contract must contain at least one speculative contract. Suppose
next that $u(a) \geq v(a)$ for all a. When the monopolist knows the consumer's type, his
objective is to solve the maximization problem

$$\max_{t^u, t^v, a^u, a^v} t^u - c(a^u)$$

subject to

$$\theta[u(a^u) - t^u] + (1 - \theta)[v(a^v) - t^v] \geq 0 \qquad (IR)$$
$$u(a^u) - t^u \geq u(a^v) - t^v \qquad (UR)$$
$$v(a^v) - t^v \geq v(a^u) - t^u. \qquad (VR)$$

Since $u(a) \geq v(a)$ for every a, the *UI* constraint implies $u(a^u) - t^u \geq v(a^v) - t^v$. By the
IR constraint, $t^u \leq u(a^u)$. Therefore, the monopolist's objective function is bounded
from above by $u(a^u) - c(a^u)$. It follows that the non-speculative contract, which extracts
$u(e^u)$ in state u and $v(e^v)$ in state v, maximizes the monopolist's profit, regardless of the
consumer's type. Therefore, the optimal menu contains no speculative contracts. □

PROOF OF PROPOSITION 6.   We proceed by proving a series of claims.

CLAIM 1. *Any contract t_θ in an optimal menu satisfies $D_\theta^u \geq \max(0, D_\theta^v)$.*

PROOF.   First, note that by IR_θ, there exists no t_θ with $D_\theta^u \leq 0$ and $D_\theta^v \leq 0$, where one of these inequalities is strict. Suppose $D_\theta^u \leq 0$ and $D_\theta^v > 0$ for some contract t_θ . Then, $u(a_\theta^u) - t_\theta(a_\theta^u) \leq 0$. Hence, by UI_θ, $u(a_\theta^v) - t_\theta(a_\theta^v) \leq 0$. Since $u(a) \geq v(a)$ for all a, $v(a_\theta^v) - t_{\theta^*}(a_\theta^v) \leq 0$, contradicting our initial assumption that $D_\theta^v > 0$. It follows that $D_\theta^u \geq 0$ for every θ. Suppose $D_\theta^u < D_\theta^v$ for some contract t_θ. Then, $u(a_\theta^u) - t_\theta(a_\theta^u) < v(a_\theta^v) - t_\theta(a_\theta^v) \leq u(a_\theta^v) - t_\theta(a_\theta^v)$, contradicting UI_θ.   ◁

CLAIM 2. *(i) There exists a non-empty subset of types $\Theta^* \subseteq [0,1]$ such that for every $\theta \in \Theta^*$ the IR_θ constraint binds.*

(ii) Let $\theta^* \equiv \max_\theta \Theta^*$. If $\theta^* > 0$, then $D_\theta^u = D_\theta^v = 0$ for all $\theta < \theta^*$.

(iii) If $\theta^* < 1$, then for any pair of types (θ, θ') satisfying $\theta > \theta' \geq \theta^*$, $D_{\theta'}^u - D_{\theta'}^v > 0$ implies $D_\theta^u - D_\theta^v > 0$.

PROOF.   To prove (i), suppose the IR_θ constraint did not bind for any of the types. Then the monopolist could strictly increase his profit by raising $t_\theta(a_\theta^u)$ and $t_\theta(a_\theta^v)$ by an arbitrarily small $\varepsilon > 0$ for all types.

To prove (ii), suppose $\theta^* > 0$ but there is some type $\theta < \theta^*$ for whom it is not true that $D_\theta^u = D_\theta^v = 0$. Note first that if $D_\theta^u = D_\theta^v \neq 0$, then by Claim 1, $D_\theta^u = D_\theta^v > 0$. But this contradicts the definition of Θ^*. Suppose next that $D_{\theta^*}^u > D_{\theta^*}^v$. By the $IC_{\theta^*, \theta}$ constraint, $\theta^*(D_\theta^u - D_\theta^v) + D_\theta^v \leq 0$. But since $\theta < \theta^*$ and $D_\theta^u - D_\theta^v > 0$, $\theta(D_\theta^u - D_\theta^v) + D_\theta^v < \theta^*(D_\theta^u - D_\theta^v) + D_\theta^v$, in violation of IR_θ.

Finally, to prove (iii), assume that $\theta^* < 1$ yet $D_{\theta'}^u - D_{\theta'}^v > 0$ and $D_\theta^u - D_\theta^v \leq 0$ for some $\theta > \theta' \geq \theta^*$. By Claim 1, $D_\theta^u = D_\theta^v = 0$. By the $IR_{\theta'}$ constraint, $\theta'(D_{\theta'}^u - D_{\theta'}^v) + D_{\theta'}^v \geq 0$. Because $\theta > \theta'$ and $D_{\theta'}^u - D_{\theta'}^v > 0$, $\theta(D_{\theta'}^u - D_{\theta'}^v) + D_{\theta'}^v > \theta'(D_{\theta'}^u - D_{\theta'}^v) + D_{\theta'}^v$, in violation of $IC_{\theta, \theta'}$.   ◁

CLAIM 3. *The VI_θ constraint is binding for all types $\theta \geq \theta^*$.*

The proof is essentially the same as the proof in Proposition 3 that the VI_θ constraint is binding for all types $\theta \geq \theta_h$, and therefore omitted.

CLAIM 4. *An optimal menu has the property that for all types $\theta < \theta^*$, $a_\theta^u = e^u$, $t_\theta(e^u) = u(e^u)$, $a_\theta^v = 0$, and $t_\theta(0) = 0$.*

PROOF.   By Claim 2, $D_\theta^u = D_\theta^v = 0$ for all types $\theta \leq \theta^*$. Since $v(a) < u(a)$ for all $a > 0$, UI_θ implies $a_\theta^v = 0$ and $t_\theta(a_\theta^v) = 0$ for all types $\theta \leq \theta^*$. This means that the monopolist's profit from each $\theta \leq \theta^*$ is $p[u(a_\theta^u) - c(a_\theta^u)]$. Therefore, the optimal contract for each $\theta \leq \theta^*$ must satisfy $a_\theta^u = e^u$.   ◁

CLAIM 5. *Suppose the optimal menu is incentive compatible and individually rational and satisfies the VI_θ constraint of each type with equality. Then the contract assigned to each $\theta \in [\theta^*, 1]$ induces an action pair (a_θ^u, a_θ^v) that maximizes*

$$[\psi(\theta) - p][u(a_\theta^u) - v(a_\theta^u)] + p[u(a_\theta^u) - c(a_\theta^u)] + (1-p)[v(a_\theta^v) - c(a_\theta^v)] \qquad (23)$$

subject to the UI_θ constraint.

PROOF. The proof applies the same arguments as those used in the proof of Proposition 3 to obtain the objective function (20). The only difference is that here we include the term $(1-p)[v(a_\theta^v) - c(a_\theta^v)]$ inside the objective function because in contrast to the environment of Section 3, we cannot prove directly that $a_\theta^v = e^v$.   ◁

Before we proceed to the next sequence of claims, we introduce the following notation. Define Θ^+ as the subset of types in $[\theta^*, 1]$ for whom $u(a_\theta^u) - v(a_\theta^u) \geq u(e^v) - v(e^v)$, where a_θ^u is given by (11). Let $\Theta^- = [\theta^*, 1]/\Theta^+$. Note that if $\theta^* < 1$, then at least one of the two sets, Θ^+ or Θ^-, is non-empty.

CLAIM 6. *Suppose both Θ^+ and Θ^- are non-empty. Assume that an optimal menu assigns to each type in Θ^+ the action pair (a_θ^u, a_θ^v), where a_θ^u is given by (11) and $a_\theta^v = e^v$. Assume further that the menu assigns to each type in Θ^- the action pair (a_θ^u, a_θ^v) given by (9). Then*

(i) *for any pair of types $\theta \in \Theta^+$ and $\theta' \in \Theta^-$, $\theta' > \theta$ and $a_\theta^u > e^v > a_{\theta'}^u$,*

(ii) *$u(a_\theta^u) - v(a_\theta^u)$ is non-decreasing on $\Theta^- \cup \Theta^+ = [\theta^*, 1]$.*

PROOF. (i) We first claim that Θ^+ includes all the types for whom $\psi(\theta) \geq 0$. To see why, consider a type θ with $\psi(\theta) \geq 0$. By our assumptions on u and v, the difference $u(a) - v(a)$ is strictly increasing and reaches a maximum at $a = 1$. Suppose $a_\theta^u \leq e^v$. Since $\psi(\theta) \geq 0$, $(\partial/\partial a)[u(a) - v(a)] > 0$, and $(\partial/\partial a)[v(a) - c(a)] > 0$, if we were to increase a by an infinitesimal amount, we would strictly increase the objective function (23), a contradiction. It follows that under an optimal menu, $a_\theta^u > e^v$. But this means that $u(a_\theta^u) - v(a_\theta^u) \geq u(e^v) - v(e^v)$.

From the previous paragraph, it follows that for every $\theta \in \Theta^-$, $\psi(\theta) < 0$. We now show that $a_{\theta'}^u < e^v$ for all $\theta' \in \Theta^-$. Suppose, to the contrary, that $a_{\theta'}^u \geq e^v$ for some type $\theta' \in \Theta^-$. Since $\psi(\theta') < 0$, $(\partial/\partial a)[u(a) - v(a)] > 0$, and $(\partial/\partial a)[v(a) - c(a)] \geq 0$, if we were to decrease a by an infinitesimal amount, we would strictly increase the objective function

$$\psi(\theta')[u(a_{\theta'}^u) - v(a_{\theta'}^u)] + [v(a_{\theta'}^u) - c(a_{\theta'}^u)], \qquad (24)$$

a contradiction. It follows that under an optimal menu, $a_{\theta'}^u \leq e^v$.

Consider a pair of types $\theta \in \Theta^+$ and $\theta' \in \Theta^-$. Since $\psi(\theta') < 0$, $\theta' < \theta$ for every $\theta \in \Theta^+$ with $\psi(\theta) \geq 0$. Suppose there exists $\theta \in \Theta^+$ with $\psi(\theta) < 0$ such that $\theta < \theta'$ for some $\theta' \in \Theta^-$. Let

$$b_{\theta'}^u \in \underset{a}{\arg\max}\{\psi(\theta')[u(a_{\theta'}^u) - v(a_{\theta'}^u)] + p[v(a_{\theta'}^u) - c(a_{\theta'}^u)]\}.$$

The same argument we used to show that for $\theta' \in \Theta^-$ expression (24) is maximized by $a^u_{\theta'} < e^v$ implies that $b^u_\theta < e^v$. By assumption, $\theta \in \Theta^+$ implies that a^u_θ is given by (11). Hence

$$\psi(\theta)[u(a^u_\theta) - v(a^u_\theta)] + p[v(a^u_\theta) - c(a^u_\theta)] \geq \psi(\theta)[u(b^u_{\theta'}) - v(b^u_{\theta'})] + p[v(b^u_{\theta'}) - c(b^u_{\theta'})]. \quad (25)$$

As was shown in the first paragraph of the proof, $a^u_\theta \geq e^v$, and therefore $u(a^u_\theta) - v(a^u_\theta) > u(b^u_{\theta'}) - v(b^u_{\theta'}) \geq 0$. Since, by assumption, $\psi(\theta) < 0$ and $p \geq 0$, inequality (25) holds only if $v(a^u_\theta) - c(a^u_\theta) > v(b^u_{\theta'}) - c(b^u_{\theta'})$. By the definition of $b^u_{\theta'}$,

$$\psi(\theta')[u(b^u_{\theta'}) - v(b^u_{\theta'})] + p[v(b^u_{\theta'}) - c(b^u_{\theta'})] \geq \psi(\theta')[u(a^u_\theta) - v(a^u_\theta)] + p[v(a^u_\theta) - c(a^u_\theta)]. \quad (26)$$

Adding inequalities (25) and (26), and cancelling common terms yields

$$[\psi(\theta') - \psi(\theta)]\{[u(b^u_{\theta'}) - v(b^u_{\theta'})] - [u(a^u_\theta) - v(a^u_\theta)]\} \geq 0.$$

Since $\psi(\theta') > \psi(\theta)$ and $u(b^u_{\theta'}) - v(b^u_{\theta'}) < u(a^u_\theta) - v(a^u_\theta)$, the left hand side of the above inequality is strictly negative, a contradiction. It follows that $\theta > \theta'$ for every $\theta \in \Theta^+$ and $\theta' \in \Theta^-$.

   (ii) We first show that $u(a^u_\theta) - v(a^u_\theta)$ is non-decreasing on Θ^+. Since the difference $u(a) - v(a)$ is strictly increasing in a, it suffices to show that a^u_θ is non-decreasing on Θ^+. Consider a pair of types $\phi, \theta \in \Theta^+$ such that $\phi > \theta$. By construction,

$$\psi(\phi)[u(a^u_\phi) - v(a^u_\phi)] + p[v(a^u_\phi) - c(a^u_\phi)] \geq \psi(\phi)[u(a^u_\theta) - v(a^u_\theta)] + p[v(a^u_\theta) - c(a^u_\theta)]$$

and

$$\psi(\theta)[u(a^u_\theta) - v(a^u_\theta)] + p[v(a^u_\theta) - c(a^u_\theta)] \geq \psi(\theta)[u(a^u_\phi) - v(a^u_\phi)] + p[v(a^u_\phi) - c(a^u_\phi)].$$

Adding these two inequalities and cancelling common terms yields

$$[\psi(\phi) - \psi(\theta)]\{[u(a^u_\phi) - v(a^u_\phi)] - [u(a^u_\theta) - v(a^u_\theta)]\} \geq 0.$$

Since $\phi > \theta$ and ψ is an increasing function, $\psi(\phi) > \psi(\theta)$. Hence, for the above inequality to hold, it must be that $u(a^u_\phi) - v(a^u_\phi) \geq u(a^u_\theta) - v(a^u_\theta)$. In a similar manner we can show that $u(a^u_\theta) - v(a^u_\theta)$ is non-decreasing on Θ^-.   ◁

CLAIM 7. *The optimal menu assigns each type in Θ^+ the action pair (a^u_θ, a^v_θ) given by $a^v_\theta = e^v$ and (11), and if Θ^- is not empty it assigns each type in this set the action pair (a^u_θ, a^v_θ) given by (9).*

PROOF.   First, note that because $\psi(1) = 1$ and ψ is an increasing function, $\Theta^+ \neq \emptyset$. Since the action pair (a^u_θ, a^v_θ) given by $a^v_\theta = e^v$ and (11) is the solution to the *unconstrained* maximization of (23), we need to show that it satisfies all the necessary constraints. We do this first for all $\theta \in \Theta^+ \setminus \{\theta^*\}$, and then later treat the case of $\theta^* \in \Theta^+$ separately.

   We begin with the VI_θ constraint, which we assumed to be binding for each $\theta \geq \theta^*$. To verify that indeed this is true for all $\theta \in \Theta^+$, we simply set $t_\theta(a^v_\theta)$ and $t_\theta(a^u_\theta)$ according to the equations (18) and (19).

Next we verify that the UI_θ constraint is satisfied for all $\theta \in \Theta^+$. Since the VI_θ constraint is binding for these types, $t_\theta(a_\theta^u) - t_\theta(a_\theta^v) = v(a_\theta^u) - v(a_\theta^v)$. This means that the UI_θ constraint is satisfied if and only if $u(a_\theta^u) - v(a_\theta^u) \geq u(a_\theta^v) - v(a_\theta^v)$. By the definition of Θ^+, this must be true for every type in this set.

We now verify that for every pair $\theta, \theta' \in \Theta^+$, the $IC_{\theta,\theta'}$ constraint is satisfied. Incentive compatibility between any two types in some subset of $[0, 1]$ is equivalent to the requirement that $q(\theta)$ is non-decreasing on this subset. Since VI_θ binds for all $\theta \in \Theta^+$, $q(\theta)$ is non-decreasing for these types if and only if $u(a_\theta^u) - v(a_\theta^u)$ is non-decreasing on Θ^+. But this follows from part (ii) of Claim 6.

There are two remaining constraints that we need to verify: first, that each $\theta \in \Theta^+$ satisfies the IR_θ constraint, and second, that for each $\theta \in \Theta^+$ and each $\theta' \in [0,1] \setminus \Theta^+$, the $IC_{\theta,\theta'}$ constraint is satisfied. In addition, we need to characterize the contract assigned to the threshold type θ^*. To do all this, we distinguish between two cases.

Case 1: $\theta^* \in \Theta^+$. We need to show that the IR_θ and $IC_{\theta,\phi}$ constraints are satisfied for every $\theta, \phi \in [\theta^*, 1]$. The proof is essentially the same as the proof that these constraints are satisfied for every $\theta, \phi \in [\theta_h, 1]$ in Proposition 3 (with θ^* playing the role of θ_h).

Case 2: $\theta^* \notin \Theta^+$. What this case means is that for $\theta = \theta^*$, a solution to the unconstrained maximization of (23), which satisfies the VI_{θ^*} constraint with equality, necessarily violates the UI_{θ^*} constraint. This means that in order to solve for $(a_{\theta^*}^u, a_{\theta^*}^v)$, we must solve the *constrained* optimization problem of maximizing (23) subject to the constraint that both the UI_{θ^*} and VI_{θ^*} constraints are binding.

By Claim 6, the difference $u(a_\theta^u) - v(a_\theta^u)$, evaluated at a_θ^u given by (11), is positive and non-decreasing. This means that there cannot be a pair of types $\phi > \theta$ such that for a_ϕ^u, a_θ^u defined by (11), $u(a_\theta^u) - v(a_\theta^u) > 0$ while $u(a_\phi^u) - v(a_\phi^u) < 0$. This means that there exists a unique threshold type $\overline{\theta} \in (\theta^*, 1)$ such that $\Theta^+ = [\overline{\theta}, 1]$ and $\Theta^- = [\theta^*, \overline{\theta})$.

Suppose the UI_θ constraint binds for some type θ. Then

$$t_\theta(a_\theta^u) - t_\theta(a_\theta^v) = u(a_\theta^u) - u(a_\theta^v). \tag{27}$$

Since the VI_θ constraint must also bind,

$$t_\theta(a_\theta^u) - t_\theta(a_\theta^v) = v(a_\theta^u) - v(a_\theta^v).$$

It follows that (27) holds if and only if

$$u(a_\theta^v) - v(a_\theta^v) = u(a_\theta^u) - v(a_\theta^u).$$

But since by assumption, the difference $u(a) - v(a)$ is strictly increasing in a, it follows that $a_\theta^u = a_\theta^v$. Hence, a_θ^u and a_θ^v are given by (9).

It follows that for every $\theta \in \Theta^-$, a_θ^u and a_θ^v are given by (9). Define $t_{\theta^*}(a_{\theta^*}^u)$ and $t_{\theta^*}(a_{\theta^*}^v)$ to be the solutions to the pair of equations defined by the requirement

that IR_{θ^*} and VI_{θ^*} are binding. By construction, the IR_θ constraint of type θ^* is binding, while for any $\theta > \theta^*$, this constraint is satisfied if the IC_{θ,θ^*} constraint is satisfied.

We now verify that for every pair of types $\theta, \theta' \in \Theta^-$, the $IC_{\theta,\theta'}$ constraint is satisfied. As before, we achieve this by showing that $q(\theta)$ is non-decreasing on Θ^-. Since the VI_θ constraint is binding for all $\theta \in \Theta^-$, this is equivalent to showing that $u(a_\theta^u) - v(a_\theta^u)$ is non-decreasing on Θ^-. But this follows from part (ii) of Claim 6.

Next, we verify that for every $\theta' \in \Theta^-$ and $\theta'' \in \Theta^+$, the $IC_{\theta',\theta''}$ and $IC_{\theta'',\theta'}$ constraints are satisfied. As explained before, to achieve this it suffices to show that $q(\theta)$ is non-decreasing on $[\theta^*, 1]$. Since we have already shown that the VI_θ constraint is binding for all $\theta \in [\theta^*, 1]$, this is equivalent to showing that $u(a_\theta^u) - v(a_\theta^u)$ is non-decreasing on $[\theta^*, 1]$. Since we have also shown that $u(a_\theta^u) - v(a_\theta^u)$ is non-decreasing on both Θ^+ and Θ^-, and that $\theta'' > \theta'$ for any $(\theta', \theta'') \in \Theta^- \times \Theta^+$, it remains to be shown that $u(a_{\theta''}^u) - v(a_{\theta''}^u) \geq u(a_{\theta'}^u) - v(a_{\theta'}^u)$. Because the difference $u(a) - v(a)$ is increasing for all a, this is equivalent to showing that $a_{\theta''}^u > a_{\theta'}^u$. But this follows from part (i) of Claim 6.

We have therefore shown that the contract assigned to all types at or above the cutoff θ^* satisfy all the required constraints. We now verify that all the necessary constraints are satisfied for types below the cutoff. These types are assigned a contract that guarantees them a zero indirect utility (hence, IR_θ trivially holds for all $\theta < \theta^*$). In addition, for each $\theta < \theta^*$, $a_\theta^u = e^u$, $a_\theta^v = 0$, and $D_\theta^u = D_\theta^v = 0$. It follows that for each of these types, the UI_θ constraint binds, and because $v(e^u) < u(e^u)$, the VI_θ constraint is satisfied. Finally, we need to verify that for every $\theta < \theta^* \leq \phi$, the $IC_{\theta,\phi}$ constraints are satisfied. To do so, note that since IR_θ^* binds and since we verified above that $IC_{\theta^*,\phi}$ holds,

$$\theta^* q(\phi) + D_\phi^v \leq 0$$

for every $\phi \geq \theta^*$. Since $q(\phi) \geq 0$, this means that for every $\theta < \theta^*$,

$$\theta q(\phi) + D_\phi^v \leq 0.$$

It follows that the $IC_{\theta,\phi}$ constraints are satisfied.

Finally, we show that if $\theta^* \in \Theta^-$, then either $\theta^* = 0$ or there exists a unique cutoff type θ^* at which the monopolist is indifferent between assigning the contract given by Claim 4 and a contract assigned to types in Θ^-. By definition, θ^* is a type with the property that every lower type is assigned the contract described in Claim 4. Let

$$h(\theta) \equiv \max_{a \in [0,1]} \{\psi(\theta)[u(a) - v(a)] + [v(a) - c(a)]\} - p \max_{a \in [0,1]} [u(a) - c(a)].$$

Then, $\theta^* = 0$ (respectively, $\theta^* \notin \Theta^-$) whenever $h(\theta) \geq 0$ (respectively, $h(\theta) < 0$) for every $\theta \in [0, \bar{\theta})$. If, however, there exists a solution $\theta^* \in \Theta^-$ to the equation $h(\theta) = 0$, then we need to show that this solution is unique.

Since $\psi(\theta)$ is non-decreasing and u lies above v, $h(\theta)$ is non-decreasing. Since $\psi(0) < 0$, $\psi(0) = 1$, and $\psi(0)$ is non-decreasing, there exists a unique solution $\theta^0 \in [0,1]$ to the equation $\psi(\theta^0) = 0$. Recall that $\psi(\theta) < 0$ for all $\theta^* \in \Theta$. It follows that if $h(\theta^0) < 0$, then $h(\theta) < 0$ for every $\theta < \theta^0$, implying that $\theta^* \notin \Theta^-$. If $h(0) > 0$, then this remains true for every $\theta \in [0, \bar{\theta})$, implying that $\theta^* = 1$. If, however, $h(0) < 0$ while $h(\theta^0) > 0$, then since $h(\theta)$ is non-decreasing, there exists a unique solution $\theta^* \in \Theta^-$ to the equation $h(\theta) = 0$.   ◁

To complete the proof of the proposition, it remains to derive the consumers' payments in each state. This is essentially the same as in the proof of Proposition 4, hence we do not provide a separate derivation.   ☐

PROOF OF PROPOSITION 7.   We begin by computing the optimal contract for types below the cutoff θ^*. From Proposition 6, it follows that $a_\theta^u = 1$, $t_\theta(1) = u(1)$, $a_\theta^v = 0$, and $t_\theta(0) = 0$. Hence the monopolist's profit from each type who signs this contract is $p(\theta)u(1) = p(\theta)kv(1)$. Let us turn to the contracts for types above θ^*. Assume that $\theta^* \geq \frac{1}{2} - p/(2(k-1))$. Then, expressions (9) and (11) imply $a_\theta^u = a_\theta^v = 1$ for all $\theta \geq \theta^*$. In particular, this means that $\bar{\theta} = \theta^*$. The monopolist's revenue from types $\theta > \theta^*$ is $\theta^* u(1) + (1 - \theta^*)v(1)$, while its revenue from each $\theta < \theta^*$ is $pu(1)$. The monopolist's expected revenue is thus

$$[\theta^* \cdot pk + (1-\theta^*)(\theta^*(k-1)+1)]v(1).$$

The optimal cutoff is thus given by

$$\theta^* = \frac{k(p+1)-2}{2k-2},$$

which is greater than $\frac{1}{2} - p/(2(k-1))$, since $p \geq \frac{1}{3}$.

Let us turn to the case of common priors. It can be shown that since $\arg\max(u-c) = \arg\max(v-c) = \arg\max(u-v) = 1$, the characterization of the optimal menu shares some features with the optimal menu under non-common priors. In particular, there exists a cutoff θ^{**} such that for every $\theta < \theta^{**}$, $a_\theta^u = 1$ and $a_\theta^v = 0$, whereas for every $\theta > \theta^{**}$, $a_\theta^u = a_\theta^v = 1$. This means that all types below θ^{**} are assigned the contract $t(1) = u(1)$, $t(0) = 0$ (and $t(a) = \infty$ for all other a), while all types above θ^{**} are assigned the contract $t(1) = \theta^{**} u(1) + (1-\theta^{**})v(1)$ (and $t(a) = \infty$ for all other a). Thus, the monopolist's expected revenue is

$$\int_0^{\theta^{**}} \theta k v(1)\, d\theta + [(1-\theta^{**})(\theta^{**}(k-1)+1)]v(1).$$

The optimal cutoff is thus $\theta^{**} = 1$. This means that all types are offered the same contract: $t(1) = u(1)$, $t(0) = 0$ (and $t(a) = \infty$ for all other a).   ☐

REFERENCES

Armstrong, Mark (1996), "Nonlinear pricing with imperfectly informed consumers." Unpublished paper, Nuffield College, Oxford. [459]

Armstrong, Mark and John Vickers (2001), "Competitive price discrimination." *Rand Journal of Economics*, 32, 579–605. [464]

Baron, David P. and David Besanko (1984), "Regulation and information in a continuing relationship." *Information Economics and Policy*, 1, 267–302. [459]

Bazerman, Max H. and James J. Gillespie (1999), "Betting on the future: The virtues of contingent contracts." *Harvard Business Review*, 155–160. [463, 475]

Courty, Pascal and Hao Li (2000), "Sequential screening." *Review of Economic Studies*, 67, 697–717. [459]

Eliaz, Kfir and Ran Spiegler (2006), "Contracting with diversely naive agents." *Review of Economic Studies*, 73, 689–714. [464]

Eliaz, Kfir and Ran Spiegler (2007), "A mechanism-design approach to speculative trade." *Econometrica*, 75, 875–884. [464, 465]

Eliaz, Kfir and Ran Spiegler (2008), "Optimal speculative trade among large traders." *Review of Economic Design*, 12, 45–74. [464, 465]

Eliaz, Kfir and Ran Spiegler (forthcoming), "Bargaining over bets." *Games and Economic Behavior*. [464, 465]

Fang, Hanming and Giuseppe Moscarini (2005), "Morale hazard." *Journal of Monetary Economics*, 52, 749–777. [463]

Gilbert, Daniel T. (2006), *Stumbling on Happiness*. Knopf, New York. [480]

Grubb, M. (forthcoming), "Selling to overconfident consumers." *American Economic Review*. [463]

Krishna, Vijay (2002), *Auction Theory*. Academic Press, San Diego. [484, 486, 488]

Landier, Augustin and David Thesmar (forthcoming), "Contracting with optimistic entrepreneurs: Theory and evidence." *Review of Financial Studies*. [464]

Miravete, Eugenio J. (2002), "Estimating demand for local telephone service with asymmetric information and optional calling plans." *Review of Economic Studies*, 69, 943–971. [459]

Miravete, Eugenio J. (2003), "Choosing the wrong calling plan? Ignorance and learning." *American Economic Review*, 93, 297–310. [459, 463]

Sandroni, Alvaro and Francesco Squintani (2007), "Overconfidence, insurance, and paternalism." *American Economic Review*, 97, 1994–2004. [464]

Tirole, Jean (1999), "Incomplete contracts: Where do we stand?" *Econometrica*, 67, 741–781. [475]

Uthemann, Andreas (2005), "Competitive screening of consumers with non-common priors." Unpublished paper, University College London. [464]