

Pricing in Matching Markets*

George J. Mailath[†] Andrew Postlewaite[‡] Larry Samuelson[§]

February 14, 2007

Abstract

We study why different markets are cleared by different types of prices—a universal price for all buyers and sellers in some markets, seller-specific prices that are uniform across buyers in others, and personalized prices tailored to both the buyer and the seller in yet others. We link these prices to differences in the premuneration values—the values in the absence of any payments (muneration)—created by the buyer-seller match. The results point to a theory of designing markets to allow effective pricing.

*We thank the National Science Foundation (grants SES-0350969 and SES-0549946) for financial support.

[†]Yale University and University of Pennsylvania, george.mailath@yale.edu

[‡]University of Pennsylvania, apostlew@econ.upenn.edu

[§]University of Wisconsin, LarrySam@ssc.wisc.edu

Pricing in Matching Markets

in 1 Introduction

Prices. Consider three people and the prices that clear their markets:

- Alice works in a daily spot market for casual, unskilled labor. A prototypical job pays her a fixed sum to drive a truck to pick up materials. This market is characterized by a single price, governing the transactions in each buyer/seller pair.
- Carol is a senior executive with an Ivy League degree. She receives job offers at quite different wages from various firms, each of which has made offers to others at wages different than those offered Carol. Her alma mater gave her a half-tuition scholarship while rejecting other students who would have paid full tuition.
- Bob works as a tax preparer, bolstered by a degree from his local junior college. He quotes the same hourly price to all of his clients, though some other tax preparers and accountants charge different prices, just as different junior college and technical schools charge different prices but accept all applicants at those prices.

We refer to the prices faced by Alice as *universal* prices. These are the prices that typically show up in supply-and-demand diagrams in introductory texts. Carol faces *personalized* prices that depend on both her characteristics and those of her trading partner. Bob faces *uniform* prices that depend on the characteristics of the agent posting the price but not those of the agent on the other side of the transaction.

Why do we see universal prices in some markets, uniform prices in others, and personalized prices in yet others? What implications does the type of pricing have for market outcomes? How are these prices linked to market characteristics? This paper addresses these questions, concentrating on uniform and personalized prices.

Premuneration values. An interaction between a buyer and seller entails a cost or benefit to each side, generating a *surplus* if the sum of the costs and benefits is positive. The surplus can be reallocated via a transfer from one side to the other. The *premuneration values* (from the Latin

munerare to give or to pay) are the values to the parties prior to any transfer. Understanding the nature of the surplus and the premuneration values is the key to understanding differences in pricing across markets.

The surplus in Alice's market is reasonably modeled as the sum of two terms, a negative premuneration value for Alice, reflecting her value of foregone leisure, and a premuneration value for her employer reflecting the benefits of getting the materials delivered. This separability ensures that there is no issue of efficient matching in these markets. It matters that the right (i.e., high-valuation) people trade, but matters not with whom they trade. It is then no surprise that a universal price clears the market.

The surpluses in Carol's markets depend in a complementary fashion upon the agents on both sides of the market. Talented executives are likely to be more productive when paired with productive firms than with mediocre firms, and vice versa. Similarly, a good student fares especially well when paired with a good school while the latter is especially effective when working with good students. Clearing such markets requires not only getting the right people to trade, but also being sure that they trade with the right partners. We might then expect to need the precision of personalized prices.

Bob's markets also exhibit complementarity. Even below the Fortune 500 and the Ivy League, there are gains from matching skilled professionals with the right firms and good students with good schools. Why do we see uniform prices in Bob's markets and personalized prices in Carol's, despite complementarities in both cases?

The prices required to achieve a market-clearing allocation depend upon the point of departure provided by the premuneration values. Both Carol and her alma mater own some of the surplus created in the match that gave Carol her education. Carol owns her enhanced earning power, but the university owns the increment to its ranking based on her superb SAT score, the increment to its prestige should she become a Supreme Court Justice, and the increment to its endowment should she become a wealthy donor. In her employment match, Carol's employer owns the revenue her services will generate, but she owns the value of the contacts that she makes before starting her own company. In contrast, Bob's junior college anticipates no benefit from Bob beyond his tuition, while Bob is indifferent over whose taxes he prepares, so long as the client pays.

Our characterization of pricing builds on this distinction. We develop a model in which buyers and sellers invest in attributes in order to enter a market in which they are matched to trade. The surplus created by a match depends supermodularly on the attributes chosen by both agents. Prices transform the premuneration values of the subsequent surplus into a final

allocation.

When prices can be personalized, there exists an equilibrium in which the resulting allocation is efficient, both in the matching and the ex ante investments. When prices are restricted to being uniform, efficient equilibria exist if and only if the premuneration values on the side of the market setting prices are independent of the attributes of the agents on the other side of the market (as is the case in Bob's but not Carol's markets). Equivalently, if (and only if) this constant premuneration value condition fails, price-setting agents strictly prefer to personalize their prices.

Our result is not simply that uniform prices suffice when the *surplus* exhibits no complementarities (and hence the efficient allocation exhibits no matching problem). Instead, it is that an efficient allocation, including both investments and matching, can be supported by uniform prices even when the surplus depends supermodularly on attribute choices and matching is an issue, as long as the price-setter's premuneration value does not depend on the price-setter's match.¹

Why are prices important? In a world devoid of frictions, personalized prices would be the norm. There would be little reason to be concerned with how premuneration values are defined in a frictionless world.² But the world is not frictionless. A seller posting personalized prices must ascertain potential buyers' attributes, a process that can be quite costly. For example, estimates from 11 highly selective liberal arts colleges indicate that they spend about \$3,000 on admissions per matriculating student in 2004.³ The going price for identifying whether a high school diploma comes from a

¹The nature of a market's premuneration and prices will depend upon how broadly we define the market. One could render the questions in this paper moot by defining every market so narrowly that it has only one buyer and seller attribute choice in it, so that universal, uniform and personalized prices necessarily coincide. The appropriate definition of a market depends on the context, being typically broad enough that the distinction between uniform and personalized prices is meaningful.

²Our emphasis on premuneration values thus reflects no disagreement with Coase (1960)'s observation that property rights would be irrelevant in a world without transactions costs.

³Expenditures for the 11 colleges, all but one of which continually appear in the *U. S. News and World Report* top 25 liberal arts colleges, were \$370 per applicant for the 1995-1996 admissions season. Publicly available data on subsequent expenditure growth rates projects an expenditure of \$625 per applicant in the 2004-2005 academic year. The 2002 admission rate for these schools was 34%. Coupling this with an estimated enrollment rate of 60% yields a cost of \$3000 per matriculating student. (Memorandum, Office of Institutional Research and Analysis, University of Pennsylvania, July 1004. We thank Bernie Lentz for his help with these data.)

legitimate high school is \$100.⁴ There may thus be substantial savings from posting uniform prices and letting buyers sort themselves (as Bob's clients do), if the premuneration values are so defined that uniform prices can do this sorting. Alternatively, if the premuneration values are such that uniform prices cannot duplicate the allocation of personalized prices, and if transactions costs or institutional considerations preclude personalized prices, then market outcomes will be inefficient.⁵

Designing markets. The premuneration values in a market can be designed as part of the institutional and legal environment of the market. For example, the match of researchers and universities generates a surplus that includes the value of marketable patents from faculty research. Historically, universities have owned these patents, but we could imagine institutional arrangements that granted them to the faculty. Indeed, the feasibility of such ownership is reflected in the decisions of many universities to unilaterally grant professors shares in the revenues from patents stemming from their research. In a similar vein, one could arrange the premuneration values in a university/student interaction so that the university owns all of the surplus. This would require a somewhat unconventional arrangement in which the university owns the future income of students to whom it gives degrees, but the feasibility of this type of arrangement is reflected in income-contingent loans in a number of countries (including Australia, Sweden and New Zealand) that effectively give the lender a share of students' future income (Johnstone01).

⁴“Vetting Those Foreign College Applications,” *New York Times*, September 29, 2004, page A21.

⁵For example, Bulow and Levin (2004) note that the National Residency Matching Program matching medical residents and hospitals constrains hospitals to make the same offers to all residents. They argue that the primary effect is not inefficient matching but a transfer of surplus to the hospitals. Nicholson (2003) argues, however, that the result is an inefficient allocation of residents to specialties. Medical students who do their residency acquire training that dramatically increases their future earnings. This part of the surplus from the match that is owned by the student is so large in some specialties such as dermatology, general surgery, orthopedic surgery and radiology that if personalized prices were employed, Nicholson argues that medical students would pay hospitals handsomely for the opportunity to do their residency in these specialties (as compared to the \$34,000 stipend they currently receive).

⁶Basketball star Yao Ming (Houston Rockets) has a contract with the China Basketball Association calling for 30% of his NBA earnings to be paid to the Chinese Basketball Association (in which he played prior to joining the Rockets), while another 20% will go to the Chinese government. Similar arrangements hold for Wang Zhizhi (Dallas Mavericks) and Menk Bateer (Denver Nuggets). (See the *Detroit News*, April 26, 2002,

Our results suggest that appropriately designed premuneration values can be valuable, allowing efficient equilibria to be supported by uniform prices, avoiding the costs of personalized pricing or the costs of inefficient uniform pricing. Unfortunately, there are often constraints on the design of premuneration values. If universities owned students' enhanced future income streams, the price in this market would presumably be reversed, with universities paying students to attend in return for owning their future income. But then why would the students exert the effort required to realized this future income? How are we to measure and collect the increment to income attributable to the university education?⁷ Such an arrangement might also require changes in labor laws that preclude involuntary servitude. More generally, laws concerning workplace safety, the (in)ability to surrender legal rights, the division of marital assets and the custody and sale of children may constrain the allocation of premuneration values. Our analysis points to the cost of such constraints or institutional arrangements, in the form of personalization costs or inefficient uniform pricing.

2 The Matching Market

Our model is adapted from ^{CMP01}Cole, Mailath, and Postlewaite (2001). There is a unit measure of buyers whose types are indexed by β and distributed uniformly on $[0, 1]$, and a unit measure of sellers whose types are indexed by σ and distributed uniformly on $[0, 1]$. For ease of reference, the typical buyer is female and typical seller male.

Buyers and sellers have an outside option with payoff zero that precludes participation in the matching process. If they do not take this option, they make choices in two stages. First, buyers and sellers simultaneously choose attributes. We denote the cost of attribute choice $b \in \mathbb{R}_+$ to buyer β by $c_B(b, \beta)$, and the cost of attribute choice $s \in \mathbb{R}_+$ to seller σ by $c_S(s, \sigma)$.

<http://www.detnews.com/2002/pistons/0204/27/sports-475199.htm/>.) We can view the initial match between Yao Ming and his Chinese team as producing a surplus that includes the enhanced value of his earnings as a result of developing his basketball skills. These contracts suggest that the premuneration values could be designed to assign some future earnings to the team.

⁷Measurement and collection both pose difficulties. The University of New Mexico sued a former researcher for rights to patents that he applied for four years after he had left the university, arguing that the patents stemmed from research that he had done before leaving. ("Universities Try to Keep Inventions From Going 'Out the Back Door'", *Chronicle of Higher Education*, May 17, 2002.) In principle, one who owns the rights to a song is entitled to a payment each time the song is played on the radio in a bar or health club, but collection is impractical. In other cases, legal restrictions may be paramount.

Buyers and sellers match in the second stage. A match between a buyer and seller with attribute choices (b, s) produces a total surplus $v(b, s)$. This surplus is the sum of a *buyer premuneration value* $h_B(b, s)$ and *seller premuneration value* $h_S(b, s)$.

Matching and the resulting division of the surplus is mediated through prices. We assume that sellers set prices, which may be either positive or negative. For example, the illustrations in Section [II](#) include cases in which prices were posted by those relinquishing a good (educational services) as well as by those receiving a good (labor services), each of whom would be designated the seller in our model. We are interested in two types of pricing, reflecting differing amounts of information available to sellers. If sellers cannot observe buyers' attribute choices, then the price a seller posts is necessarily independent of these attribute choices. We say that prices are *uniform* in this case. If a seller can observe buyers' attributed choices, then he can post *personalized* prices, that is prices that depend on a buyer's attribute choice. Intuitively, once prices are posted, each buyer selects a seller, with our equilibrium concept requiring that each seller is matched with at most one buyer.

horse

Assumption 1

1. The surplus function $v : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is C^2 , increasing in b and s , and strictly supermodular:

$$\frac{d^2 v(b, s)}{db ds} > 0.$$

2. The cost function $c_B : \mathbb{R}_+ \times [0, 1] \rightarrow \mathbb{R}_+$ is C^2 , strictly increasing and convex in b , with $c_B(0, \beta) = 0 = dc_B(0, \beta)/db$ and

$$\frac{d^2 c_B(b, \beta)}{db d\beta} < 0.$$

The cost function c_S satisfies analogous conditions.

3. There exists $\bar{b}$ such that for all $b > \bar{b}$, $s \in \mathbb{R}_+$, $\beta \in [0, 1]$ and $\sigma \in [0, 1]$,

$$v(b, s) - c_B(b, \beta) - c_S(s, \sigma) < 0.$$

A similar statement applies to sellers with an analogous $\bar{s}$.

4. For every $\beta = \sigma \equiv \phi \in (0, 1]$, there exists $(b, s) \in [0, \bar{b}] \times [0, \bar{s}]$ with $v(b, s) - c_B(b, \phi) - c_S(s, \phi) > 0$.

5. The premuneration values $h_B : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}$ and $h_S : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}$ are C^2 , increasing, and

$$\frac{dh_B^2(b, s)}{dbds} > 0 \quad \text{and} \quad \frac{dh_S^2(b, s)}{dbds} \geq 0.$$

The buyer premuneration value $h_B(b, s)$ is Lipschitz continuous in s on $[0, \bar{b}] \times [0, \bar{s}]$.

Assumption ^{horse} 1.1 restricts attention to supermodular surplus functions, ensuring that matching is important and focussing our attention on uniform and personalized (rather than universal) prices. Assumption ^{horse} 1.2 gives meaning to agents' types by imposing a single-crossing condition on the cost function, so that higher-indexed types have a comparative advantage in choosing large attributes. Assumption ^{horse} 1.3 ensures that the problem is bounded, in the sense that attributes outside the intervals $b \in [0, \bar{b}]$ and $s \in [0, \bar{s}]$ will never be chosen. Assumption ^{horse} 1.4 ensures that the efficient outcome calls for (almost) all agents to make positive attribute choices. The final assumption imposes a single-crossing condition on the premuneration values. Given the supermodularity of the surplus function, these restrictions on premuneration values are satisfied, for example, if $h_B = \theta v$ and $h_S = (1 - \theta)v$ for any constant $\theta \in (0, 1]$. The imposition of strict single crossing for the buyer only and the Lipschitz requirement only on the buyer's premuneration value reflects asymmetries that arise when sellers post uniform prices.

There is always an equilibrium in which every agent chooses the outside option—it does not pay to be the only one in market. We are interested in equilibria in which some agents enter the market, and simplify the analysis by considering only equilibria where everyone enters the market. We hereafter typically omit the possibility of the outside option from our notation (though incorporating its presence in the optimality conditions for equilibrium).

We let $\mathbf{b} : [0, 1] \rightarrow [0, \bar{b}]$ and $\mathbf{s} : [0, 1] \rightarrow [0, \bar{s}]$ be Lebesgue-measurable functions denoting the attributes chosen by buyers and sellers. We assume matching between buyer and seller attribute choices depends only on the distribution of such choices in the market (and not on the specification of which types make which attribute choices). It is an immediate implication of the single-crossing Assumption ^{horse} 1.2 on costs that $\mathbf{b}$ and $\mathbf{s}$ are weakly increasing in such a model. It is conceivable that they are not strictly increasing. Under certain conditions, for example, it may be an equilibrium for every agent to enter them market but choose attribute zero. We say that $\mathbf{b}$ (with $\mathbf{s}$ treated similarly) is *strictly increasing when positive* if $\mathbf{b}(\beta) > 0$ and $\beta' > \beta$

imply $\mathbf{b}(\beta') > \mathbf{b}(\beta)$. We avoid a collection of technical considerations by focussing on equilibria in which $\mathbf{b}$ and $\mathbf{s}$ are strictly increasing when positive.

We denote by $\mathcal{B}$ and $\mathcal{S}$ the closures of the sets of attributes chosen by buyers and sellers respectively,

$$\begin{aligned}\mathcal{B} &= \text{cl}(\mathbf{b}([0, 1])) \\ \mathcal{S} &= \text{cl}(\mathbf{s}([0, 1])).\end{aligned}$$

Let $\lambda_{\mathcal{B}}$ and $\lambda_{\mathcal{S}}$ be the measures induced on $\mathcal{B}$ and $\mathcal{S}$ by the agents' attribute choices. Hence, for Borel sets $\mathcal{B}' \subset \mathcal{B}$ and $\mathcal{S}' \subset \mathcal{S}$,

$$\begin{aligned}\lambda_{\mathcal{B}}(\mathcal{B}') &= \lambda\{\beta \in [0, 1] \mid \mathbf{b}(\beta) \in \mathcal{B}'\} \\ \lambda_{\mathcal{S}}(\mathcal{S}') &= \lambda\{\sigma \in [0, 1] \mid \mathbf{s}(\sigma) \in \mathcal{S}'\},\end{aligned}$$

where λ is Lebesgue measure. We define:

defn-matching

Definition 1 Suppose $\mathbf{b}$ and $\mathbf{s}$ are strictly increasing when positive. Then a feasible matching is a bijection $\tilde{b} : \mathcal{S} \rightarrow \mathcal{B}$ that is measure preserving, i.e., $\lambda_{\mathcal{B}}(\tilde{b}(\mathcal{S}')) = \lambda_{\mathcal{S}}(\mathcal{S}')$ for all Borel $\mathcal{S}' \subset \mathcal{S}$.

The measure-preserving requirement on $\tilde{b}$ ensures that the measure of any set of sellers is equal to the measure of the set of buyers with whom they are matched. Given a feasible matching $\tilde{b}$, $\tilde{b}(s)$ specifies the buyer attribute choice matched to a seller with attribute choice s .

Note that $\mathcal{B}$ and $\mathcal{S}$ are defined as the closures of the sets of attribute choices. This allows us to accommodate the technical complications raised when matching continua of agents characterized by arbitrary attribute choice functions. As a result, however, it is possible that seller σ (with attribute choice $\mathbf{s}(\sigma)$) is matched with a buyer attribute choice b which is chosen by no buyer. We interpret such a seller as matching with a buyer whose attribute choice is arbitrarily close to b , while saying that $\mathbf{s}(\sigma)$ matches with b .

Remark 1 The surplus generated by a match in our model depends only on the attendant attribute choices. Problems in which the attribute chosen is a particular skill, such as the case of the NBA star Yao Ming, fall into this category. In other cases, the surplus might depend on the agents' types as well as attribute choices. Harvard may care not only about an applicant's accomplishments (attribute choice), but also about the applicant's "cost of acquiring" such accomplishments (type). If attribute choices and types can both be observed, then we need only adopt a new definition of "attribute choice" that includes both (the previous notion of) an agent's attribute

choice and his type, at which point our analysis applies. If neither attribute choices nor types can be observed, then such a reformulation of “attribute choice” ensures that our model of uniform pricing applies. If sellers can observe attribute choices but not types, but care about both (or only about types), then attribute choices take on a dual role, directly enhancing the value of a match while also providing signals of types. ^{Hopkins05} Hopkins (2005) and ^{HMS05} Hoppe, Moldovanu, and Sela (2005) examine such models. ■

3 Equilibrium

sect-personalized

3.1 Personalized Pricing

In this section, we assume sellers observe buyers’ attribute choices and post personalized prices. The (possibly negative) price that seller of attribute choice $s \in \mathcal{S}$ receives when selling to a buyer with attribute choice $b \in \mathcal{B}$ is given by $p_P(b, s)$.

A *feasible outcome* is a pair of attribute choice functions $\mathbf{b} : [0, 1] \rightarrow [0, \bar{b}]$ and $\mathbf{s} : [0, 1] \rightarrow [0, \bar{s}]$ that are strictly increasing when positive, and a feasible matching $\tilde{b}$. A *personalized price function* is a function $p_P : \mathcal{B} \times \mathcal{S} \rightarrow \mathbb{R}$. Intuitively, a *personalized-price equilibrium* is a feasible outcome and a personalized price function such that no agent has an incentive to deviate from the behavior specified by the feasible outcome.

Given a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b})$ and a personalized price p_P the payoffs to a buyer β who chooses $b \in \mathcal{B}$ and a seller σ who chooses $s \in \mathcal{S}$ are respectively

$$\begin{aligned} \Pi_B(b, \beta) &\equiv h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) - c_B(b, \beta) \\ \text{and} \quad \Pi_S(s, \sigma) &\equiv h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) - c_S(s, \sigma). \end{aligned}$$

ned seller dev copy(1)

Definition 2 Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$, seller σ has a profitable deviation if either $\Pi_S(\mathbf{s}(\sigma), \sigma) < 0$ or there exists a seller attribute choice $s \in [0, \bar{s}]$, a buyer attribute choice $b \in \mathcal{B}$, and a price $p \in \mathbb{R}$ such that

$$h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) < h_B(b, s) - p \quad (1) \quad \text{costa}$$

$$\Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b, s) + p - c_S(s, \sigma). \quad (2) \quad \text{rica}$$

If $\Pi_S(\mathbf{s}(\sigma), \sigma) < 0$, the outside option is better for the seller. Otherwise, a seller has a profitable deviation if he is able to attract a buyer (condition (1)) while obtaining a higher payoff than had he followed the behavior prescribed ^{costa}

by the given outcome (condition ^{price}(2)). Of particular interest are deviations in which seller σ identifies a seller attribute $s \in \mathcal{S}$, $s \neq \mathbf{s}(\sigma)$, that is, an attribute that exists in the market. Such an attribute is matched with a buyer attribute $\tilde{b}(s)$ at price $p_P(\tilde{b}(s), s)$. If matching with that buyer attribute at the market price yields a higher net payoff to seller σ , then the seller has a profitable deviation:

Lemma 1 *If there exists $\sigma \in [0, 1]$ and $s \in \mathcal{S}$ such that $\Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) - c_S(s, \sigma)$, then seller σ has a profitable deviation.*

The proof is immediate. Since attribute choice s and buyer $\tilde{b}(s)$ at price $p_P(\tilde{b}(s), s)$ make seller σ strictly better off, so will the same attribute choice and buyer at price $p_P(\tilde{b}(s), s) - \epsilon$. The latter also makes the buyer strictly better off, yielding a profitable deviation. Hence, the absence of profitable deviations ensures that no seller can envy the trade of any other seller. The definition of profitable deviation goes beyond that, ensuring the sellers also do not desire transactions that would be attractive to a buyer but are not currently offered in the proposed outcome.

We can define an analogous notion for buyers.

Definition 3 *Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$, buyer β has a profitable deviation if either $\Pi_B(\mathbf{b}(\beta), \beta) < 0$ or there exists an attribute choice b , a price $p \in \mathbb{R}$, and $s \in \mathcal{S}$ with*

$$\Pi_B(\mathbf{b}(\beta), \beta) < h_B(b, s) - p - c_B(b, \beta)$$

and

$$h_S(\tilde{b}(s), s) - p_P(\tilde{b}(s), s) < h_S(b, s) + p.$$

In other words, there is a buyer attribute, a transfer, and a target seller such that in the resulting transaction, both the buyer and the seller would be better off than had they followed the behavior prescribed by the proposed equilibrium.

Analogous to the case with sellers, no buyer can envy the transaction proposed for any other buyer: a buyer who strictly prefers the transaction available to another buyer can acquire that buyer's attribute and agree to pay a slightly higher price to the given seller, making both better off than at their proposed transactions. Also as in the seller case, if there is a buyer attribute that has not been chosen by any other buyer that would be attractive to a seller at some price and yields the buyer a higher payoff, the buyer has a profitable deviation.

We then have:

Definition 4 A personalized price equilibrium is a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b})$ and a personalized price function p_P such that no seller or buyer has a profitable deviation.

sect-uniform

3.2 Uniform Pricing

We now consider the case in which sellers do not have the information necessary to set personalized prices. In the absence of this information, a seller can post a price that depends on his own attribute choice, but not on the buyer's attribute choice, so that the price function is *uniform-price function* $p_U : \mathcal{S} \rightarrow \mathbb{R}$.

Given a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b})$ and a uniform price p_U the payoffs to a buyer β who chooses $b \in \mathcal{B}$ and a seller σ who chooses $s \in \mathcal{S}$ are as before:

$$\begin{aligned} \Pi_B(b, \beta) &\equiv h_B(b, \tilde{s}(b)) - p_U(b, \tilde{s}(b)) - c_B(b, \beta) \\ \text{and} \quad \Pi_S(s, \sigma) &\equiv h_S(\tilde{b}(s), s) + p_U(\tilde{b}(s), s) - c_S(s, \sigma). \end{aligned}$$

As in the case of personalized prices, we want to capture the idea that buyers and sellers cannot profitably deviate. Because sellers cannot observe buyers' attributes, however, a seller cannot choose an attribute s and a price p targeted at a particular buyer attribute b . Instead, the attribute s and price p may attract a range of buyer attributes. The following definition embodies this constraint.

Definition 5 Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_U)$, an uninformed seller σ has a profitable deviation if either $\Pi_S(\mathbf{s}(\sigma), \sigma) < 0$ or there exists s' and a price $p \in \mathbb{R}$ such that there exists $b' \in \mathcal{B}$ with

$$h_B(b', \tilde{s}(b')) - p_U(\tilde{s}(b')) < h_B(b', s') - p,$$

and for all $b'' \in \mathcal{B}$,

$$h_B(b'', \tilde{s}(b'')) - p_U(\tilde{s}(b'')) < h_B(b'', s') - p \Rightarrow \Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b'', s') + p - c_S(s', \sigma).$$

Here, we require that when a seller chooses an attribute-price pair (s, p) , his net payoff should be higher than had he followed the prescribed behavior, for *all* buyers who would be attracted to the transaction (s, p) (cf. Remark [2](#)).

As in the case of personalized prices, if there are no profitable deviations, then no seller can envy the transaction proposed for any other seller. However, the argument is somewhat more complicated in the case of uniform

pricing. With personalized pricing, any seller who envied another seller's transaction could simply offer a slightly lower price to the target buyer, making both the given seller and the target buyer better off and yielding a profitable deviation, with the ability to personalize prices ensuring that the seller need not be concerned with other buyers. An uninformed seller cannot target a given buyer in this way. A seller who mimicked the attribute selection of another seller while undercutting his price would typically attract not only the buyer matched with the target seller but also buyers with lower attributes, potentially making the seller worse off. Nonetheless, it can be shown that if a seller envies the transaction of any other seller, there will be *some* profitable deviation for the former. Section A.1 proves:

wabash **Lemma 2** *Consider a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b})$ and uniform price p_U . If there is a seller σ and $s' \in \mathcal{S}$ with $\Pi_S(\mathbf{s}(\sigma), \sigma) < \Pi_S(s', \sigma)$, then seller σ has a profitable deviation.*

The definition of profitable deviations for buyers is as before:

Definition 6 *Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$, buyer β has a profitable deviation if either $\Pi_B(\mathbf{b}(\beta), \beta) < 0$ or there exists an attribute choice b , a price $p \in \mathbb{R}$, and $s \in \mathcal{S}$ with*

$$\Pi_B(\mathbf{b}(\beta), \beta) < h_B(b, s) - p - c_B(b, \beta)$$

and

$$h_S(\tilde{b}(s), s) - p_P(\tilde{b}(s), s) < h_S(b, s) + p.$$

Definition 7 *A uniform-price equilibrium is a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b})$ and a uniform price function p_U such that no buyer or seller has a profitable deviation.*

color **Remark 2** A seller is defined to have a profitable deviation under uniform pricing only if he is better off when matched with *any* buyer who is attracted to the deviation. Why make sellers so pessimistic? One could alternatively think of requiring only that the seller be better off under a random draw from the set of attracted buyers, or given the buyer from this set most preferred by the seller. Allowing the seller to select his most preferred buyer essentially restores the ability to personalize prices (Lemma 5 below makes this connection precise), so that some pessimism is essential to uniform pricing. A more pessimistic formulation makes seller deviations less attractive and hence expands the set of uniform-price equilibria. Our key results (Propositions 1 and 2) establish conditions under which personalized-price

and uniform-price equilibria coincide, and hence are rendered more powerful by such a permissive definition of the latter. ■

4 Examples

fortran

4.1 The Market

risk

Let the surplus and cost functions be:

$$v(b, s) = bs, \quad c_B(b, \beta) = \frac{b^3}{3\beta}, \quad c_S(s, \sigma) = \frac{s^3}{3\sigma}.$$

The premuneration values assign a fixed share of the surplus, $\theta \in [0, 1]$, to the buyer:

$$h_B(b, s) = \theta bs \quad \text{and} \quad h_S(b, s) = (1 - \theta)bs.$$

4.2 Efficiency

Efficiency requires that for each matched pair β and σ , attribute choices b and s solve:

$$\max_{b,s} bs - \frac{b^3}{3\beta} - \frac{s^3}{3\sigma},$$

giving first-order conditions

$$\begin{aligned} s - \frac{b^2}{\beta} &= 0 \\ b - \frac{s^2}{\sigma} &= 0. \end{aligned}$$

Setting $\sigma = \beta$ and hence $s = b$, we have the solution

$$\mathbf{b}(\beta) = \beta, \quad \mathbf{s}(\sigma) = \sigma.$$

4.3 Personalized Pricing

emu

We next show that

$$p_P(b, s) = \frac{s^2}{2} - (1 - \theta)bs \tag{3} \quad \text{ostrich}$$

is a personalized-price function associated with an efficient equilibrium outcome for our example. Note that for any seller attribute s , the price that a seller would receive in a match with a buyer with attribute b is decreasing in

b. It is not surprising that this might occur since the seller is getting $(1 - \theta)$ of the surplus and the surplus is increasing in the buyer's attribute choice.

Given the price function ^{ostrich} (3), buyer β solves

$$\max_{b,s} \theta bs - \frac{s^2}{2} + (1 - \theta)bs - \frac{b^3}{3\beta} = \max_{b,s} bs - \frac{s^2}{2} - \frac{b^3}{3\beta} \Rightarrow \mathbf{b}(\beta) = \beta.$$

Hence, $\tilde{b}(s) = s$. Seller σ then chooses attribute s to solve

$$\max_s (1 - \theta)\tilde{b}(s)s + p_P(\tilde{b}(s), s) - c_S(s, \sigma) = \frac{s^2}{2} - \frac{s^3}{3\sigma} \Rightarrow \mathbf{s}(\sigma) = \sigma.$$

Equilibrium payoffs to the seller and buyer are

$$\begin{aligned} \frac{(\mathbf{s}(\sigma))^2}{2} - \frac{(\mathbf{s}(\sigma))^3}{3\sigma} &= \frac{\sigma^2}{2} - \frac{\sigma^3}{3\sigma} = \frac{1}{6}\sigma^2 \\ \frac{(\mathbf{b}(\beta))^2}{2} - \frac{(\mathbf{b}(\beta))^3}{3\beta} &= \frac{\beta^2}{2} - \frac{\beta^3}{3\beta} = \frac{1}{6}\beta^2. \end{aligned}$$

Note that the agents' attribute choices are independent of the buyer's premuneration value, determined by θ . The agents' equilibrium utilities are also independent of the premuneration values. It does not matter who "owns" the technology that combines buyer and seller attribute choices into surplus when there is a competitive market with personalized prices for the attributes.

4.4 Uniform Pricing May Duplicate Personalized Pricing

Let $\theta = 1$, and hence $h_S(b, s) = 0$. Consider a uniform price function $p_U(s)$ that attaches to each seller attribute choice s the equilibrium price the seller receives in the personalized price equilibrium we have just constructed, or

$$p_U(s) = p_P(\tilde{b}(s), s) = \frac{s^2}{2} - (1 - \theta)\tilde{b}(s)s = \frac{s^2}{2}.$$

Section ^{screech} 4.5 shows that this pricing function is part of a uniform-pricing equilibrium that duplicates the efficient outcome of the personalized-pricing equilibrium, and indeed that would exhibit no profitable deviations even if prices could be personalized. In this case, the ability to personalize prices is irrelevant. Proposition ^{feaster} 11 below shows that this coincidence is a general implication of the property $dh_S(b, s)/db = 0$.

screech

4.5 Uniform Pricing Need Not Match Personalized Pricing

We calculate a uniform-price equilibrium for our example. The buyer's problem is now

$$\max_{b,s} \theta bs - p_U(s) - \frac{b^3}{3\beta},$$

for first-order conditions

$$\begin{aligned} \theta s - \frac{b^2}{\beta} &= 0, \\ \theta b - p'_U(s) &= 0. \end{aligned}$$

The seller's objective is

$$\max_s (1 - \theta) \tilde{b}(s)s + p_U(s) - \frac{s^3}{3\sigma},$$

for a first-order condition

$$(1 - \theta)[\tilde{b}'(s)s + \tilde{b}(s)] + p'_U(s) - \frac{s^2}{\sigma} = 0.$$

We now conjecture that the equilibrium attribute choice functions are given by

$$\mathbf{b}(\beta) = A\beta \tag{4} \quad \boxed{\mathbf{bfirst}}$$

$$\mathbf{s}(\sigma) = B\sigma. \tag{5} \quad \boxed{\mathbf{sfirst}}$$

If so, and assuming that, in equilibrium, a buyer of type β matches with seller of type $\sigma = \beta$, we have, for any matched pair of b and s values, $bB = sA$. Using this, we can rewrite the second buyer first-order condition as $\theta \frac{A}{B}s - p'(s) = 0$ and solve for the price function

$$p_U(s) = \frac{\theta A}{2B}s^2.$$

Similarly, we can rewrite the first buyer first-order condition as $\theta \frac{B}{A}b - \frac{b^2}{\beta} = 0$ and solve for

$$b = \theta \frac{B}{A}\beta. \tag{6} \quad \boxed{\mathbf{bsecond}}$$

Turning to the seller, we can write the first-order condition as $2(1 - \theta) \frac{A}{B}s + \theta \frac{A}{B}s - \frac{s^2}{\sigma} = 0$ and solve for

$$s = (2 - \theta) \frac{A}{B}\sigma. \tag{7} \quad \boxed{\mathbf{ssecond}}$$

Combining $\frac{\text{bfirst}}{(4)}$ with $\frac{\text{bsecond}}{(6)}$ and $\frac{\text{sfirst}}{(5)}$ with $\frac{\text{ssecond}}{(7)}$, we can solve for $A = \theta^{\frac{2}{3}}(2 - \theta)^{\frac{1}{3}}$ and $B = \theta^{\frac{1}{3}}(2 - \theta)^{\frac{2}{3}}$, and hence

$$\mathbf{b}(\beta) = \theta^{\frac{2}{3}}(2 - \theta)^{\frac{1}{3}}\beta$$

$$\mathbf{s}(\sigma) = \theta^{\frac{1}{3}}(2 - \theta)^{\frac{2}{3}}\sigma$$

$$p_U(s) = \frac{\theta}{2} \left(\frac{\theta}{2 - \theta} \right)^{1/3} s^2$$

$$\tilde{b}(s) = \left(\frac{\theta}{2 - \theta} \right)^{1/3} s.$$

When $\theta = 1$, we have the efficient solution, with buyers and sellers behaving symmetrically. When $\theta < 1$, so that the seller's premuneration value is positive, we have $A/B = ((\theta/(2 - \theta))^{1/3} < 1$. This implies that buyers now choose smaller attributes than do sellers, with buyers of attribute choice level b matching with values of $s > b$. Differentiating $B = \theta^{\frac{1}{3}}(2 - \theta)^{\frac{2}{3}}$, as θ falls below one, seller attribute choices initially increase, and ultimately decrease to zero as θ goes to 0. Differentiating $A = \theta^{\frac{2}{3}}(2 - \theta)^{\frac{1}{3}}$, we see that as θ falls below one, so do buyers' attribute choices, again falling to zero as θ goes to zero.

We can confirm that for $\theta < 1$, this uniform-price equilibrium is not a personalized-price equilibrium. To do this, it suffices to choose an attribute choice level x , and consider a buyer who makes attribute choice $b = x$ and a seller who makes attribute choice $s = x$ (and hence is not matched with the buyer in question), and then to show that their payoffs sum to less than x^2 (ensuring that the buyer and seller could do better matching with each other). This ensures that choosing attribute x and setting a (personalized) price that will attract the buyer constitutes a profitable deviation for the seller. This condition is:

$$\begin{aligned} & (1 - \theta)\tilde{b}(x)x + p_U(x) + \theta\tilde{s}(x)x - p_U(\tilde{s}(x)) \\ = & (1 - \theta) \left(\frac{\theta}{1 - \theta} \right)^{\frac{1}{3}} x^2 + \frac{\theta}{2} \left(\frac{\theta}{1 - \theta} \right)^{\frac{1}{3}} x^2 + \theta \left(\frac{2 - \theta}{\theta} \right) x^2 - \frac{\theta}{2} \left(\frac{\theta}{2 - \theta} \right)^{\frac{1}{3}} \left(\frac{2 - \theta}{\theta} \right)^{\frac{2}{3}} x^2 \\ = & \frac{1}{2} \left[(2 - \theta)^{\frac{2}{3}} \theta^{\frac{1}{3}} + (2 - \theta)^{\frac{1}{3}} \theta^{\frac{2}{3}} \right] x^2 \\ < & x^2. \end{aligned}$$

The equilibrium payoffs for the seller in the uniform-price equilibrium are given by

$$\begin{aligned}
& (1 - \theta)\tilde{b}(s(\sigma))s(\sigma) + p_U(s(\sigma)) - \frac{(s(\sigma))^3}{3\sigma} \\
= & (1 - \theta) \left(\frac{\theta}{2 - \theta} \right)^{\frac{1}{3}} \left[\theta^{\frac{1}{3}}(2 - \theta)^{\frac{2}{3}}\sigma \right]^2 + \frac{\theta}{2} \left(\frac{\theta}{2 - \theta} \right)^{\frac{1}{3}} \left[\theta^{\frac{1}{3}}(2 - \theta)^{\frac{2}{3}}\sigma \right]^2 - \frac{\left(\theta^{\frac{1}{3}}(2 - \theta)^{\frac{2}{3}}\sigma \right)^3}{3\sigma} \\
= & \frac{1}{6}\theta(2 - \theta)^2\sigma^2.
\end{aligned}$$

When $\theta = 1$, this duplicates the payoff from the personalized price equilibrium. For values of θ not too much smaller than 1, the seller earns a higher payoff under the uniform price equilibrium.

Similarly, the buyer's payoff is

$$\begin{aligned}
& \theta\tilde{s}(b(\beta)) - p_U(\tilde{s}(b(\beta))) - \frac{(b(\beta))^3}{3\beta} \\
= & \theta \left(\frac{2 - \theta}{\theta} \right)^{\frac{1}{3}} \left[\theta^{\frac{2}{3}}(2 - \theta)^{\frac{1}{3}}\beta \right]^2 - \frac{\theta}{2} \left(\frac{\theta}{2 - \theta} \right)^{\frac{1}{3}} \left(\frac{2 - \theta}{\theta} \right)^{\frac{2}{3}} \left[\theta^{\frac{2}{3}}(2 - \theta)^{\frac{1}{3}}\beta \right]^2 - \frac{\left(\theta^{\frac{2}{3}}(2 - \theta)^{\frac{1}{3}}\beta \right)^3}{3\beta} \\
= & \frac{1}{6}\theta^2(2 - \theta)\beta^2.
\end{aligned}$$

This payoff is always smaller under the uniform than personalized price equilibrium.

favor

Remark 3 When $\theta = 0$, so the seller owns all of the surplus, the equilibrium collapses into the trivial equilibrium in which no trade occurs. In this case, a buyer's payoff is solely the price p_U , which will have to be negative in order to bring buyers into the market, and buyers will choose the seller posting the smallest ("largest negative") price. Because sellers cannot condition prices on buyer attribute choice, every buyer will choose $b = 0$ in equilibrium. Similarly, when θ is positive but small, the equilibrium is markedly inefficient, featuring tiny attribute choices. This is an indication that the wrong side of the market is setting prices—sellers are appropriate price-setters for large θ , but buyers would be a more efficient choice to post prices when θ is small. ■

5 Allocating Buyers and Sellers

How do personalized and uniform prices allocate buyers and sellers? Let $(\mathbf{b}, \mathbf{s}, \tilde{b})$ be a feasible outcome. Define $\tilde{B} : [0, 1] \rightarrow [0, 1]$ by letting $\tilde{B}(\beta) = \mathbf{s}^{-1}(\tilde{b}(\mathbf{b}(\beta)))$. Thus, $\tilde{B}$ maps each buyer type into the seller type with whom that buyer matches.

5.1 Assortative Matching

Our first result, proven in Section ^{pend1}A.2, is that equilibrium requires assortative matching.

hand

Lemma 3

^{hand}3.1 Let $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$ be a personalized-price equilibrium. Then $\tilde{b}$ is strictly increasing and $\tilde{B}$ can be taken to be the identity.

^{hand}3.2 Let $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_U)$ be a uniform-price equilibrium. Then $\tilde{b}$ is strictly increasing and $\tilde{B}$ can be taken to be the identity.

5.2 Premuneration Values

Since personalized prices allow the market to compensate for any particularities of premuneration values, premuneration values play no role in the characterization of a personalized price equilibrium or in its existence. The following is a straightforward calculation:

Lemma 4 Let $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$ be a personalized-price equilibrium, with premuneration values $h_B(b, s)$ and $h_S(b, s)$. Then $(\mathbf{b}, \mathbf{s}, \tilde{b}, p'_P)$ is a personalized-price equilibrium, with premuneration values $h'_B(b, s)$ and $h'_S(b, s)$, where

$$p'_P(b, s) = p_P(b, s) + h'_B(b, s) - h_B(b, s) = [p(b, s) + h_S(b, s) - h'_S(b, s)].$$

Prices can thus be adjusted to “undo” any changes in premuneration values. Premuneration values would be irrelevant if we cared only about personalized price equilibria. Section ^{fortran}4 makes it clear that the same is *not* the case for uniform pricing equilibrium. In particular, setting $\theta = 1$ in the specification of premuneration values for the example gives an efficient equilibrium outcome that cannot be achieved for any value $\theta < 1$.

5.3 Rationing

Personalized prices allow a seller to accept some buyers while excluding others who would be willing to pay the same price. The buyers excluded by

the seller have lower attribute choices than the seller's equilibrium match. In particular, we say that a personalized price specification $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ can be supported by a *uniform rationing price* if

$$p_P(b, s) = p_P(\tilde{b}(s), s) \quad \forall b \geq \tilde{b}(s),$$

and we call a personalized price equilibrium outcome that can be supported by a uniform rationing price a *uniform rationing equilibrium outcome*. In this case, we can think of seller whose attribute choice is s as setting a uniform price $p(s) = p(\tilde{b}(s), s)$, but then excluding any buyers $b < \tilde{b}(s)$.

ration

Lemma 5 *Any personalized price equilibrium outcome is a uniform rationing equilibrium outcome.*

Proof. Let $(\mathbf{b}, \mathbf{s}, \tilde{b})$ be a personalized price equilibrium outcome and consider its associated uniform rationing price. The conditions for the latter to be a personalized price equilibrium are implied by the former, with the exception that there may now be profitable priced deviations by a seller β with attribute choice $b(\beta)$ to match with a seller with $s < \tilde{s}(b(\beta))$ (and hence $\tilde{b}(s) < b(\beta)$). But since $h_S(b, s)$ is increasing in b , the seller in question would welcome such a match. Hence, if this match is a profitable priced deviation in the uniform rationing equilibrium, it is a profitable unpriced deviation in the personalized price equilibrium, a contradiction. ■

In contrast, if a uniform-price equilibrium outcome is not also a personalized-price equilibrium outcome, then it must be that some seller would like to lower his price in order to attract better buyers, but is deterred from doing so by the specter of less desirable buyers:

Lemma 6 *Suppose $(\mathbf{b}, \mathbf{s}, \tilde{b})$ is a uniform-price equilibrium outcome and there is no p_P for which $(\mathbf{b}, \mathbf{s}, \tilde{b}; p_P)$ is a personalized-price equilibrium. Then there must exist a seller s and buyers $\underline{b}$ and $\bar{b}$ and a price $p < p_U(s)$ such that*

$$\begin{aligned} h_S(\bar{b}, s) + p &> h_S(\tilde{b}(s), s) + p_U(s) \\ h_B(\bar{b}, s) - p &> h_S(\bar{b}, \tilde{s}(\bar{b})) + p_U(\tilde{s}(\bar{b})) \\ h_S(\underline{b}, s) + p &< h_S(\tilde{b}(s), s) + p_U(s) \\ h_B(\underline{b}, s) - p &> h_S(\underline{b}, \tilde{s}(\underline{b})) + p_U(\tilde{s}(\underline{b})). \end{aligned}$$

Hence, the seller could profitably lower his price enough to attract buyer $\bar{b}$, but unfortunately instead attracts the unprofitable buyer $\underline{b}$.

Finally, Section ^{fortran}4 provided an example in which personalized-price equilibrium outcomes and uniform-price equilibrium outcomes coincide. In this case, the personalized-price power to exclude buyers is unnecessary—buyers sort themselves among sellers just as sellers would have them do.

6 Efficiency

Since we can take the equilibrium type matching $\tilde{B}$ in either a uniform-price or personalized-price equilibrium to be the identity, we define the ex ante surplus for buyer and seller types $\beta = \sigma = \phi \in [0, 1]$ as

$$W(b, s, \phi) = v(b, s) - c_B(b, \phi) - c_S(s, \phi).$$

An efficient choice of attributes maximizes $W(b, s, \phi)$ for (almost) all ϕ .

Personalized-price equilibrium outcomes are constrained efficient in the sense that no matched or unmatched pair of agents can increase its net surplus without both agents deviating to attribute choices outside the sets $\mathcal{B}$ and $\mathcal{S}$. Section ^{pend2}(A.3) proves the following by showing that if these constrained efficiency conditions fail, then the attribute choices involved in the failure can be exploited to construct a profitable deviation.

prop-constrained eff

Lemma 7 *Suppose $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$ is a personalized price equilibrium. Then, for all $\phi \in [0, 1]$, $b \in \mathcal{B}$, $s \in \mathcal{S}$ and all $b' \in [0, \bar{b}]$, $s' \in [0, \bar{s}]$,*

$$\begin{aligned} W(b, s', \phi) &\leq W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi) \\ \text{and} \quad W(b', s, \phi) &\leq W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi). \end{aligned}$$

Section ^{screech}4.5 makes it clear that the same is typically not true of uniform-price equilibria.

This result does not ensure that a personalized-price equilibrium outcome is efficient. The possibility remains that $W(b, s, \phi)$ may be maximized by values $b \notin \mathcal{B}$ and $s \notin \mathcal{S}$. In this sense, the inefficiency is the result of a coordination failure. For example, consider the surplus function $v(b, s) = bs$; here, it is an equilibrium that all agents choose attribute 0. In contrast, uniform-price equilibria in general do not satisfy even constrained efficiency.

We view the possible inefficiency of a personalized pricing equilibrium as reflecting incomplete markets. If there are “enough” prices, a personalized pricing equilibrium is efficient:

slushpump

Definition 8 *The feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b})$ and personalized-price p_P is a complete personalized-price equilibrium if there is an extension of p_P to $[0, \bar{b}] \times [0, \bar{s}]$ (also denoted p_P) such that for all β and all σ ,*

$$\begin{aligned} h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) - c_B(b, \beta) &\geq \sup_{(b,s) \in [0, \bar{b}] \times [0, \bar{s}]} h_B(b, s) - p_P(b, s) - c_B(b, \beta) \\ h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) - c_S(s, \sigma) &\geq \sup_{(b,s) \in [0, \bar{b}] \times [0, \bar{s}]} h_S(b, s) + p_P(b, s) - c_S(s, \sigma) \geq 0. \end{aligned}$$

As the names suggest, every complete personalized-price equilibrium outcome is indeed a personalized-price equilibrium outcome. Section A.4 ^{spread} proves:

tulip

Lemma 8

^{tulip}(8.1) *Every complete personalized-price equilibrium outcome is a personalized-price equilibrium outcome.*

^{tulip}(8.2) *A complete personalized-price equilibrium outcome is efficient.*

Remark 4 We could similarly define a complete uniform-price equilibrium by requiring a price for all seller attributes in $[0, \bar{s}]$, while expanding to $[0, \bar{s}]$ the set of seller attribute choices over which the buyer optimizes. It is immediate from the definition that a complete uniform-price equilibrium is a uniform-price equilibrium, and apparent from Section 4.5 ^{screech} that a complete uniform-price equilibrium need not be efficient. ■

7 When is Personalization Redundant?

When can a personalized-price equilibrium outcome be supported by uniform prices? Or, alternatively, when can matching when sellers are uninformed achieve outcomes attainable as equilibrium outcomes when they are informed?

We begin with some intuition, appropriate when equilibrium is characterized by first-order conditions. Fix a uniform price equilibrium, including the uniform-price function p_U . The first order conditions implied for the buyer's choice of attribute b and matching attribute choice s in a uniform-price equilibrium are

$$0 = \frac{dh_B(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \tag{8} \quad \text{buddy}$$

$$0 = \frac{dh_B(b, s)}{ds} - \frac{dp_U(s)}{ds}, \tag{9} \quad \text{water}$$

while the seller's first-order condition for choosing s is

$$0 = \frac{dh_S(\tilde{b}(s), s)}{db} \frac{d\tilde{b}(s)}{ds} + \frac{dh_S(\tilde{b}(s), s)}{ds} + \frac{dp_U(s)}{ds} - \frac{dc_S(s, \sigma)}{ds}. \quad (10) \quad \boxed{\text{boy}}$$

Using (9) to eliminate $\frac{dp_U(s)}{ds}$ in (10) and then using the identity $v(b, s) = h_B(b, s) + h_S(b, s)$ in (8) and (10), these three first-order conditions can be reduced to

$$\begin{aligned} 0 &= \frac{dv(b, s)}{db} - \frac{dh_S(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \\ 0 &= \frac{dh_S(b, s)}{db} \frac{d\tilde{b}(s)}{ds} + \frac{dv(b, s)}{ds} - \frac{dc_S(s, \sigma)}{ds}. \end{aligned}$$

From Lemma prop-constrained eff 7, establishing the constrained efficiency of a personalized-price equilibrium outcomes, we know that a personalized-price equilibrium must be characterized by the first-order conditions:

$$\begin{aligned} 0 &= \frac{dv(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \\ 0 &= \frac{dv(b, s)}{ds} - \frac{dc_S(s, \sigma)}{ds}. \end{aligned} \quad (11) \quad \boxed{\text{holly}}$$

Comparing these, it is immediate that the solution to the first order conditions for the personalized price equilibrium will be a solution for the first order conditions for the uniform price equilibrium if $\frac{dh_S(b, s)}{db} = 0$, that is, if each seller's premuneration value of the surplus is independent of the attribute choice of the buyer with whom the seller is matched. This argument is summarized in the following proposition.

feaster **Proposition 1** *A personalized price equilibrium outcome can be achieved in a uniform price equilibrium if the sellers' premuneration values do not depend on the buyer's attribute.*

Proof. Let $(\mathbf{b}, \mathbf{s}, \tilde{b}; p_P)$ be a personalized price equilibrium. Then $(b, s, \tilde{b}, \hat{p}_P)$ is also a personalized-price equilibrium, where $\hat{p}_P(\tilde{b}(s), s) = p_P(\tilde{b}(s), s)$ for all $s \in \mathcal{S}$ and

$$\hat{p}_P(b, s) = p_P(\tilde{b}(s), s) + h_S(\tilde{b}(s), s) - h_S(b, s)$$

for all $(b, s) \in \mathcal{B} \times \mathcal{S}$ satisfying $b \neq \tilde{b}(s)$. Intuitively, $\hat{p}_P(b, s)$ is the reservation price for a seller with attribute choice s to match with buyer attribute

$b \neq \tilde{b}(s)$.⁸ In particular, the construction of $\hat{p}_P(b, s)$ as the reservation price of seller attribute choice s for buyer attribute b ensures that sellers have no profitable deviations under price $\hat{p}_P$. If, under $\hat{p}_P$, a buyer with attribute b now strictly prefers to buy seller attribute choice $s' \neq \tilde{s}(b)$, then for sufficiently small $\varepsilon > 0$ the price $\hat{p}_P(b, s') + \varepsilon$ allows us to construct a profitable deviation in $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$, a contradiction.

The argument is now completed by noting that if $h_S(b, s)$ does not depend on b , then neither does $\hat{p}_P$, ensuring that $(b, s, \hat{b}; p_U)$ for $p_U(s) = \hat{p}_P(\cdot, s)$ is a uniform-price equilibrium. ■

The constancy of $h_S(b, s)$ in b is also essentially necessary for personalized price equilibria to be achieved via uniform pricing. The “essentially” here is that this constancy need not hold for pairs (b, s) that are not matched in equilibrium.⁹

refeaster

Proposition 2 *Let $(\mathbf{b}, \mathbf{s}, \tilde{b}; p_P)$ be a personalized price equilibrium and that the outcome $(\mathbf{b}, \mathbf{s}, \tilde{b})$ can be supported as a uniform price equilibrium outcome with price $p_U(s) = p_P(\tilde{b}(s), s)$. Then for all $s \in \mathcal{S}$,*

$$\frac{dh_S(\tilde{b}(s), s)}{db} = 0.$$

Proof. It follows from (8)–(11) (without any differentiability assumptions beyond those placed on the primitives of the model in Assumption I), that if $(\mathbf{b}, \mathbf{s}, \tilde{b}; p_P)$ is a personalized price equilibrium that can be supported by uniform prices, then

$$\frac{dh_B(\tilde{b}(s), s)}{db} = \frac{dv(\tilde{b}(s), s)}{db},$$

implying $\frac{dh_S(\tilde{b}(s), s)}{db} = 0$. ■

⁸ For example, the pricing function (3) constructed in Section 4.3 attaches to each pair (b, s) the seller’s reservation value for the buyer, since

$$p_P(b, s) - p_P(b', s) = (1 - \theta)(b' - b)s = h_S(b', s) - h_S(b, s).$$

⁹ Analogously, the single-crossing condition is essentially necessary for a separating equilibrium in a signaling model.

8 Existence of Equilibrium

8.1 Uniform Price Equilibria

For any $\beta = \sigma$, let $\hat{\mathbf{b}}(\beta)$ and $\hat{\mathbf{s}}(\sigma)$ be efficient, i.e.,

$$(\hat{\mathbf{b}}(\beta), \hat{\mathbf{s}}(\sigma)) \in \arg \max_{[0, \bar{b}] \times [0, \bar{s}]} h_B(b, s) + h_S(b, s) - c_B(b, \beta) - c_S(s, \sigma).$$

A fixed-point argument (in Section ^{pend4}A.5) allows us to establish:

Proposition 3

^{prop-uniform exist}
(3.1) ^{prop-uniform exist} Let there exist $(b, s) \in [0, \bar{b}] \times [0, \bar{s}]$ with $h_B(b, s) + h_S(b, s) - c_B(b, 1) - c_S(s, 1) > 0$. Then there exists a uniform-price equilibrium in which some buyers and some sellers make strictly positive attribute choices.

^{prop-uniform exist}
(3.2) ^{prop-uniform exist} Suppose that for all $\phi \in (0, 1]$, $h_B(\hat{\mathbf{b}}(\phi), \hat{\mathbf{s}}(\phi)) + h_S(0, \hat{\mathbf{s}}(\phi)) - c_B(\hat{\mathbf{b}}(\phi), \phi) - c_S(\hat{\mathbf{s}}(\phi), \phi) > 0$. Then there exists a uniform-price equilibrium with $\mathbf{s}(\sigma), \mathbf{b}(\beta) > 0$ for $\sigma, \beta > 0$.

The condition that $h_B(\hat{\mathbf{b}}(\phi), \hat{\mathbf{s}}(\phi)) + h_S(0, \hat{\mathbf{s}}(\phi)) - c_B(\hat{\mathbf{b}}(\phi), \phi) - c_S(\hat{\mathbf{s}}(\phi), \phi) > 0$ is satisfied if $h_S(b, s)$ is independent of b (in which case personalized and uniform pricing correspond). It can fail if $dh_S(b, s)/dS$ is large (e.g., when θ is small in Section ^{screech}4.5). In such cases, buyers ^{favor} are the appropriate side of the market to be setting prices (cf. Remark 3). Our sufficient conditions for existence thus require that prices be set on the appropriate side of the market. In particular, uniform pricing equilibria are inefficient when $h_S(b, s)$ depends on b . If this dependence is too pronounced, they may be so inefficient as to preclude trade.

8.2 Personalized-Price Equilibrium

One route to existence is to note that a personalized price equilibrium is equivalent to ^{CMP01}Cole, Mailath, and Postlewaite (2001)'s ex post contracting equilibrium, and then to note that ^{CMP01}Cole, Mailath, and Postlewaite (2001) establish conditions for the existence of an ex post contracting equilibria. We take an alternative route here, building on the relationship between personalized-price and uniform-price equilibria.

^{prop-personalized exist} **Proposition 4** *There exists an efficient personalized-price equilibrium.*

Proof. Suppose first that $h_S(b, s) = 0$ and hence $h_B(b, s) = v(b, s)$ for all pairs (b, s) . Proposition 3 ^{prop-uniform exist} ensures that there exists a complete uniform-price equilibrium. Proposition 1 ^{feaster} ensures that there is a corresponding complete

personalized price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}; p_P)$, which Lemma ^{tulip}8 ensures is efficient. Then setting

$$p'_P(b, s) = p_P(b, s) - h_S(b, s) = p_P(b, s) + h_B(b, s) - v(b, s)$$

gives a complete (and hence efficient) personalized price equilibrium for the market in question. ■

sect-general pricing

9 Endogenous Pricing

Given that personalized prices allow sellers to do everything they can do with uniform prices and more, we would ⁱⁿexpect personalized prices everywhere if they were not costly. Section ⁱⁿ1 explains why we think personalization is costly. Uniform prices may thus appear either because personalization is prohibitively costly or unnecessary (cf. Proposition ^{feaster}11). In general, we would expect sellers to be able to decide whether to personalize, balancing the costs and benefits of doing so. This section explores this trade-off.

9.1 The Monitoring Decision

We must first expand the model by assuming that, simultaneously with their attribute choices, sellers also decide whether to become *informed*, in which case they can set personalized prices, or *uninformed*, in which case they must set uniform prices. We denote the set of informed sellers by I and the closure of the set of attributes chosen by informed sellers by $\mathcal{S}(I)$. Similarly, the set of uninformed sellers is U and the closure of the set of attributes chosen by uninformed sellers is $\mathcal{S}(U)$. It is possible that $\mathcal{S}(I) \cap \mathcal{S}(U) \neq \emptyset$ since we have taken the closure of the set of chosen attributes. We present a formulation of feasible matchings, and hence equilibrium, only for the case in which the set of informed sellers is such that $\lambda(\mathcal{S}(I) \cap \mathcal{S}(U)) = 0$ (we return to this below). The resulting notions are well defined, as no unilateral deviation can disrupt such a condition.

In describing the matching, we must distinguish between attributes chosen by informed and uninformed sellers. The function $\tilde{b}_i : \mathcal{S}(I) \rightarrow \mathcal{B}$ is a one-to-one measure-preserving function describing the match of an informed seller with attribute choice $s \in \mathcal{S}(I)$, and its inverse on $\tilde{b}_i(\mathcal{S}(I))$ is denoted $\tilde{s}_i$. The function $\tilde{b}_u : \mathcal{S}(U) \rightarrow \mathcal{B}$ is a one-to-one measure-preserving function describing the match of an uninformed seller with attribute choice $s \in \mathcal{S}(U)$ and its inverse on $\tilde{b}_u(\mathcal{S}(U))$ is denoted $\tilde{s}_u$. The pair $(\tilde{b}_i, \tilde{b}_u)$ is a feasible matching if, in addition, $\tilde{b}_i(\mathcal{S}(I)) \cup \tilde{b}_u(\mathcal{S}(U)) = \mathcal{B}$. Since only sellers who become

informed can condition their price on buyers' investments, there will be two price functions: $p_P : \mathcal{B} \times \mathcal{S}(I) \rightarrow \mathbb{R}$ and $p_U : \mathcal{S}(U) \rightarrow \mathbb{R}$. The first function is the price an informed seller with attribute choice s charges for buyer of attribute choice b , and the latter the price set by an uninformed seller of attribute choice s for any buyer.

A feasible outcome in this case is thus a collection $\{\mathbf{b}, \mathbf{s}, I, \tilde{b}\}$, and p_P and p_U are associated prices that respect sellers' decisions to be informed or not. Let $\mathcal{A}$ denote a specification. We denote the cost of the monitoring technology to a seller of type σ by $K(\sigma, \mathcal{A})$. We thus leave open the possibility that the cost of the technology may depend upon a seller's type. At one extreme, the cost may consist entirely of a fixed cost, so that $K(\sigma, \mathcal{A}) = K$. Alternatively, the cost may be decreasing in σ , so that types who find attributes cheaper also find the monitoring technology cheaper, but may otherwise be independent of $\mathcal{A}$. These formulations would be applicable if, for example, the primary cost to a college of monitoring applicants is the fixed cost of creating an admissions department. However, the cost of monitoring may depend upon how many buyers would like to purchase at the observed price $p(\tilde{b}(s), s)$, but are excluded. This is consistent with an interpretation of personalized prices as uniform rationing prices.

An *equilibrium with endogenous pricing* is a feasible outcome and prices p_P, p_U , $\{\mathbf{b}, \mathbf{s}, I, \tilde{b}; p_P, p_U\}$, with the properties that first period (investment and monitoring) choices are optimal, the second period match is feasible, second-period prices respect sellers' monitoring choices (i.e., informed seller personalize while uninformed sellers set uniform prices), and there are no profitable second-period deviations. The ideas are familiar from the definitions of personalized and uniform prices but the details requires some care. The formalities are presented in Section A.6.

9.2 Optimal Uniform Prices

It is clear that the equilibrium of the endogenous model will exhibit uniform pricing if the gains from personalization are small and the cost large.¹⁰

¹⁰~~Proposition 3 established the existence of uniform price equilibria (and, via Proposition 4, personalized price equilibria) by relying crucially on the ability to restrict attention to behavior that is monotonic in type. A similar approach to a general existence result for the endogenous-price model is precluded by the fact that we have no reason to expect the decision to become informed to be monotonic in type. The difficulty is not that existence is expected to fail, but that we can expect equilibria to look quite different, and to be established via different techniques, in different instances of the model, particularly different formulations of the monitoring cost.~~

tumor

Proposition 5

5.1 ^{tumor} Let $h_S(b, s)$ be independent of b . Then it is an equilibrium of the endogenous model for each seller to choose not to obtain the monitoring technology, coupled with a uniform price equilibrium.

5.2 ^{tumor} Fix the functions v , c_B and c_S , and let $\{h_B^n, h_S^n\}_{n=1}^\infty$ be default-share functions converging uniformly to a limit in which h_S is constant in b . Then for every $\tilde{K} > 0$, there is an N such that if the monitoring-cost is bounded below by $\tilde{K}$, then there exists an equilibrium of the general model in which no seller buys the monitoring technology, for all $n > N$.

The first statement reiterates our basic conclusion, that markets can be (efficiently) cleared by uniform prices when premuneration values are appropriately defined. The next statement notes that this is not a “razor-edge” result. The benefit to a seller to obtain the monitoring technology is to be able to discriminate among the potential buyers with whom he might transact. But if the differences in the size of $h_S(b, s)$ across potential buyers is sufficiently small, the benefits from acquiring the technology will be less than paying the cost to obtain the technology, and consequently it will be an equilibrium in the endogenous model for no seller to purchase and to set a uniform price.

Remark 5 Turning this around, if the monitoring cost K is sufficiently small and $h_S(b, s)$ is not independent of b , then we will not have completely uniform pricing. In particular, suppose the outcome of a uniform price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}; p_U(s))$ cannot be supported in a personalized-price equilibrium. Then there must exist a seller who can use personalized prices to construct a profitable deviation, i.e., a seller σ , attribute choice s and price $p_U < p_U(s)$ and buyer β with attribute choice b such that

$$h_S(b, s) + p_U - c_S(s, \sigma) > h_S(\tilde{b}(\sigma), \mathbf{s}(\sigma)) + p_U(s(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma)$$

and

$$h_B(b, s) - p_U(s) - c_B(b, \beta) \geq h_B(\mathbf{b}(\beta), \tilde{s}(b)) - p_U(\tilde{s}(b)) - c_B(\mathbf{b}(\beta), \beta).$$

Hence, if seller s had the ability to do so, he would post a price that would accept buyer b at price p_U , but exclude some buyers. It is the inability to exclude such buyers that deters the seller from posting price p_U in the uniform price equilibrium.¹¹

¹¹As pointed out in the previous footnote, we do not have general conditions for existence of equilibria in the endogenous case.

■

Remark 6 One might conjecture that if $h_S(b, s)$ is not independent of b , then we will have a personalized price equilibrium if K is sufficiently small, without further assumptions on how it is small. However, this is not the case. Let $K(\sigma, \mathcal{A}) = K > 0$. Then the endogenous model does not have an equilibrium in which all sellers acquire the monitoring technology. The lowest type of seller attribute choice necessarily matches with the lowest buyer investment, and hence has no buyers to exclude. It then cannot be in this seller's best interests to pay to acquire the monitoring technology. For values $K > 0$, this applies to an interval of lowest-type sellers, which precludes the existence of a personalized price equilibrium. ■

9.3 Endogenously Mixed Uniform and Personalized Prices

Continuing with our running example, we illustrate an equilibrium with a mixture of uniform and personalized prices. We assume that K is independent of $\mathcal{A}$ but decreasing in σ , with $K(1) = 0$, and consider the class of cost functions αK for $\alpha > 0$.

Let the sum of the payoffs to a buyer and seller of types $\beta = \sigma = \phi$ in the personalized and uniform price equilibria be denoted by $v_P(\phi, \phi)$ and $v_U(\phi, \phi)$. Then in our example,

$$\begin{aligned} v_P(\phi, \phi) &= \frac{1}{3}\phi^2 \\ v_U(\phi, \phi) &= \frac{1}{6}[\theta(2 - \theta)^2 + \theta^2(2 - \theta)]\phi^2 \\ &= \frac{1}{3}\theta(2 - \theta)\phi^2. \end{aligned}$$

Let ψ satisfy

$$v^P(\psi, \psi) - v^U(\psi, \psi) = K(\psi).$$

A match between two agents of type ψ is then the switch-point at which the efficiency of the personalized-pricing equilibrium just suffices to warrant paying the cost κ of the technology. Agents with types below ψ will not purchase the monitoring technology and will play as in the uniform-pricing equilibrium. Agents above ψ will purchase the monitoring technology, and

will play as in the personalized-pricing equilibrium, with the exception that the price will now be given by

$$p_P(b, s) = \frac{s^2}{2} - (1 - \theta)bs + \Delta.$$

The constant Δ affects none of the incentives in the personalized-pricing equilibrium. It is chosen to equalize the payoffs of the marginal seller $\sigma = \psi$ in the two equilibria. This is the required condition for this seller to be indifferent between buying and not buying the monitoring technology. We have

$$\begin{aligned}\Delta &= \frac{1}{3}\psi^2 - \frac{1}{3}\theta(2 - \theta)\psi^2 \\ &= \frac{1}{3}(1 - \theta)^2\psi^2 \\ &> 0.\end{aligned}$$

Hence, the division of the surplus is pushed in the seller's favor, compared to the personalized-pricing equilibrium, in response to seller $\sigma = \psi$'s outside option of saving the cost of the monitoring technology by entering the uniform-pricing segment of the market.

The seller's attribute choice drops as σ crosses ψ while the buyer's jumps up. The price jumps down:

$$\begin{aligned}p_P(\mathbf{b}(\psi), \mathbf{s}(\psi)) &= (\theta - \frac{1}{2})\psi^2 + \Delta \\ &< \frac{\theta}{2} \left(\frac{\theta}{2 - \theta} \right)^{\frac{1}{3}} \theta^{\frac{2}{3}} (2 - \theta)^{\frac{4}{3}} \psi^2 \\ &= p_U(\mathbf{b}(\psi), \mathbf{s}(\psi)).\end{aligned}$$

The inequality is equivalent to

$$(\theta - \frac{1}{2}) + \frac{1}{3}(1 - \theta)^2 < \frac{\theta}{2} \left(\frac{\theta}{2 - \theta} \right)^{\frac{1}{3}} \theta^{\frac{2}{3}} (2 - \theta)^{\frac{4}{3}}$$

which is readily verified numerically. At the switch point ψ , the marginal buyer thus trades off a high-attribute choice seller and a high price (just below ψ) against a relatively low-attribute choice low-price seller (above ψ). As the seller moves across ψ , the seller is able to pay less for higher-attribute

choice buyers, at the cost of buying the transfer-setting technology. Notice that some buyers below ψ would like to buy sellers above ψ , at the observed transfers, without increasing their investments. However, the equilibrium precludes their doing so.

The only optimality condition that is not obvious in this formulation concerns the behavior of sellers near the critical type ψ . Seller ψ is indifferent between acquiring and not acquiring the monitoring technology, which may initially appear to suffice for optimality. However, we have noted that the seller attribute choice falls at type ψ . If K is independent of σ , then sellers' types enter their payoffs only through the cost function c_S . Given the single-crossing property satisfied by c_S , we could then conclude that the equilibrium seller attribute choice must be increasing in type, ensuring that the proposed strategies are not an equilibrium. Seller ψ to can be indifferent between a large attribute choice coupled with uniform pricing and a small attribute choice coupled personalized pricing, without seller $\psi - \varepsilon$ for small ε strictly preferring the latter (disrupting the equilibrium) only if seller ψ has a cost advantage in purchasing the monitoring technology, i.e., only if $K(\sigma)$ declines sufficiently rapidly in σ , i.e., if $dK(\psi)/d\sigma$ is sufficiently negative. This will be the case, and we will have an equilibrium, for all α sufficiently large.¹²

10 Discussion

out

Moral Hazard. It is worth discussing the premuneration values of the surplus as they play a central role in our analysis. Our main result is that a necessary and sufficient condition for efficient investments is that sellers' premuneration values should be independent of the buyer to whom they are matched. An obvious question is that if this allocation of the premuneration values is optimal, why doesn't society specify property rights to make this so? We will expand on the moral hazard issue that we touched upon briefly above.

Consider a standard moral hazard-matching problem in which a set of heterogeneous risk averse tenants will be matched with a set of heterogeneous risk neutral landlords. Suppose that sellers' premuneration values are set at zero and uniform pricing prevails. The parties will match, but the

¹²Let $\bar{c} = \lim_{\sigma \uparrow \psi} c_S(\mathbf{s}(\sigma), \sigma)$ and $\underline{c} = \lim_{\sigma \downarrow \psi} c_S(\mathbf{s}(\sigma), \sigma)$. Then we need

$$\frac{d(\bar{c} - \underline{c})}{d\sigma} > \frac{dK(\psi)}{d\sigma}$$

division of the surplus that leaves the risk averse tenant bearing all risk is inefficient. If the parties then recontract to an efficient outcome, the landlord will bear some risk, that is, he will have a share of the surplus. But if the parties foresee this outcome, it is in essence the premuneration value, and it will typically not leave the seller indifferent over buyers with whom he might match.

Menus of seller offers. It is straightforward to see that when buyers' attributes are not observable to sellers, sellers cannot benefit by offering menus of contracts in our model. Each seller will ultimately match with a single buyer, and there is no underlying uncertainty in our model. Thus, a seller can perfectly forecast what contract will be chosen in equilibrium and could offer that contract alone.

Who pays to reveal attributes? Firms resort to uniform pricing equilibria when they do not have the information necessary for personalized pricing. It is not intrinsically costly to generate the information; individuals are assumed to know their attributes in our model. Rather, it is costly to firms to get this information because individuals have an incentive to distort firms' estimates of their attribute if it is to be used for personalized pricing. There are, however, alternatives to firms acquiring a monitoring technology to assess buyers' attributes. Universities typically require students to take SAT exams that, at least partially, reveal the attribute of interest to universities. Even if universities did not require students to take the exams, it is unlikely that it would be an equilibrium for all students to not take the exams. If all students were not taking the exam and uniform pricing prevailed, students of high ability would find it in their interest to take an exam, if one existed, to certify their attribute. Thus a more general model would include a richer set of technologies by which either buyers or sellers could make attributes known to all participants.

Who benefits if attributes are observable? The previous paragraph suggests that if the cost to buyers of certifying their attribute is not too high, the uncertainty might "unravel": high attribute buyers reveal themselves, making it optimal for the highest attribute buyers in the remaining pool to reveal their attribute, and so on until all buyers' attributes are known.¹³ To avoid this, Harvard Business School students successfully lobbied for policies that prohibit students' divulging their grades to potential employers,

¹³See [G. Milgrom \(1981\)](#) or [OkunoPostlewaiteSuzumura](#), e.g., for analyses of this.

and Wharton student government adopted a policy banning the release of grades.¹⁴ While we cannot be certain that it is always the case, we point out that in our example above, buyers are worse off when information about their attributes is not known than when it is. It is interesting that this holds for any non-zero share the buyers may have of the surplus, and for all students, not just the low attribute students. That even the lowest attribute students would be made worse off if buyer attribute information were suppressed is due to the altered incentives to invest when attributes will be known.

Do actual markets use personalized prices? ^{YoungBurke01} document that of nearly a thousand Illinois sharecropping contracts, over 80% split the crop equally between tenant and landlord, independent of the quality of the soil or the crop being raised. Of the remaining contracts virtually all specify shares of $2/3$ - $1/3$ or $3/5$ - $2/5$.¹⁵ Looking at the sharecropping problem through the lens of our model, we would think of the landlord as the seller and the tenant as the buyer. The attributes of the seller would be a measure of the crop that could be produced on the land with given inputs, and the landlord's investment in attribute would include building access roads. Tenants' attribute would be the equipment they own and their agricultural abilities. The share of the crop offered by landlord to tenants is then a personalized price. Each landlord can be thought of as offering half the crop to the tenant, with a lower bound on the tenant attribute he was willing to accept. What is surprising about this personalized price is the regularity by which in each matched pair, the price is half the output. Given that both the attributes of the tenants and the landlords must vary substantially from region to region, it is surprising that the personalized price function we analyze above would systematically take this form. Young and Burke present an alternative explanation to our personalized prices that lays out a dynamic model that gives rise to the norm-based shares that are observed. However, if the equilibrium price that mediates transactions between matched pairs is norm-based, and hence diverges from our equilibrium personalized price, there must be inefficient investments. The inefficiency can take the form of excess investment on one side in order to match with a more desirable partner (for example, tenants overinvesting in equipment or landlords making excessive land improvements), or too little investment if partners are not particularly desirable at the given price.

¹⁴~~OstrovskySchwarz05~~ investigate the optimal amount of information to disclose from the students' perspective.

¹⁵~~BaBaranPillai00BhSchan86~~ ^{YoungBurke01} document similar patterns for village economies in India and Africa.

Real estate markets have a similar regularity: brokers typically charge 6% independent of the characteristics of the house to be sold or of market conditions. Here, one would think of the sellers as being real estate brokers whose attributes include the ability to bring potential buyers to a listed house and the buyers as being homeowners who wish to sell, whose attributes are the value of their house. As with landlord-tenant matching, one can think of the sales commissions as a personalized price with constraints on acceptable buyer attributes. As with sharecropping contracts, it is highly unlikely that the equilibrium personalized price would exhibit this kind of regularity across markets and market conditions; collusion is a more likely explanation. Again, if the observed transaction prices diverge from the equilibrium personalized price, there will be inefficient investments.

These are but two of a number of instances in which we observe equilibrium prices that exhibit regularities that are improbable if markets were governed by personalized pricing.¹⁶ If such improbable regularities coincide with nontrivial investments by one side or another, inefficient investments are likely.

A Appendix

pend

pend0

A.1 Proof of Lemma ^{wabash}2

Proof: Suppose seller σ has a priced deviation. By definition, there exist s' such that $h_S(\tilde{b}(s'), s') + p_U(s') - c_S(s', \sigma) > h_S(\tilde{b}(s), s) + p_U(s) - c_S(s, \sigma)$.

Case I: $\mathbf{b}$ is discontinuous at $\beta = \mathbf{b}^{-1}(\tilde{b}(s'))$. In this case, $\mathbf{b}$ jumps up at β , and consequently $\sup_{b < \mathbf{b}(\beta)} h_B(b, s) < h_B(\mathbf{b}(s'), s)$. Suppose now that seller σ chooses an attribute s'' slightly larger than s' and chooses a price p' slightly higher than $p_U(s')$, with the changes such that $h_B(\tilde{b}(s'), s') - p_U(s') = h_B(\tilde{b}(s'), s'') - p'$. That is, the increase in the premuneration value to the buyer who chooses seller attribute s' is exactly offset by the increase in price from $p_U(s')$ to p' . If the seller's changes from s' to s'' and from $p_U(s')$ to p' are sufficiently small,

$$h_S(\tilde{b}(s'), s'') + p' - c_S(s'', \sigma) > h_S(\tilde{b}(s), s) + p_U(s) - c_S(s, \sigma).$$

That is, the seller is better off choosing attribute s'' and price p' than by playing the putative equilibrium *if* he could still be assured matching with

¹⁶A further example would be lawyers, who charge standard contingent fees in all personal injury cases ranging from 33 1/3% to 50% depending on the jurisdiction.

buyer $\tilde{b}(s')$. We note that $\sup_{b < \beta} h_S(b, s'') < h_S(\tilde{b}(s'), s'')$, and since the change in seller's attribute and price left indifferent, all buyers with lower attributes are strictly worse off choosing seller attribute s'' at price p' than at the putative equilibrium. Thus, they will still be strictly worse off even if seller σ chooses attribute s'' and a price $p'' = p' - \varepsilon$ for sufficiently small ε . Thus, seller σ choosing attribute s'' and price p'' will attract buyer $\tilde{b}(s')$ but no buyer with lower attribute. The combination may attract buyers with higher attributes than $\tilde{b}(s')$, but the seller will be strictly better off for all buyers who are attracted to attribute s'' and price p'' . Summarizing, if there is a profitable priced deviation and the buyer attribute choice function is discontinuous at the buyer attribute, there is a (generalized) profitable deviation.

Case II: $\mathbf{b}$ is continuous at $\beta = \mathbf{b}^{-1}(\tilde{b}(s'))$. As in case I, suppose the seller chooses attribute s'' and price p' that are respectively higher than s' and $p_U(s')$ that leave buyer β indifferent, that is, $h_B(\tilde{b}(s'), s') - p_U(s') = h_B(\tilde{b}(s'), s'') - p'$. $h_B(\cdot, s')$ is a continuous strictly increasing function, so for any ε there exists b' and $\delta > 0$ such that $\tilde{b}(s') - b' < \varepsilon$ and $h_B(\tilde{b}(s'), s') - h_B(b, s') > \delta$ for all $b < b'$. Thus if seller σ lowers the price from p' to p'' by less than δ , buyer attribute $\tilde{b}(s')$ will be strictly attracted, but no buyer with attribute less than b' will be attracted. Since the seller is strictly better off attracting buyer $\tilde{b}(s')$ at price $p_U(s')$, he will be strictly better off choosing s'' and p'' if the increase in the attribute is sufficiently small and the changes are such that the attributes of the buyers attracted to s'' and p'' are sufficiently close to $\tilde{b}(s')$. There may be buyers with attributes greater than $\tilde{b}(s')$ who are attracted as well, but the seller is clearly better off with those buyers as well.

A.2 Proof of Lemma ^{hand}3

pend1

Suppose $\tilde{b}$ is not weakly increasing, so that there exists $s_1 < s_2$ with $b_1 = \tilde{b}(s_2) < \tilde{b}(s_1) \equiv b_2$. Adding conditions (I) ^{Costa} and (2) ^{Rica}, the requirement that seller choosing s_2 not have a profitable deviation that maintains attribute s_2 but matches with buyer attribute b_2 instead of b_1 holds only if

$$h_B(b_2, s_1) - p_P(b_2, s_1) + h_S(b_1, s_2) + p_P(b_1, s_2) \geq h_B(b_2, s_2) + h_S(b_2, s_2).$$

Similarly, if the seller choosing s_1 is to not have a profitable deviation maintaining s_1 but matching with b_1 , we must have

$$h_B(b_1, s_2) - p_P(b_1, s_2) + h_S(b_2, s_1) + p_P(b_2, s_1) \geq h_B(b_1, s_1) + h_S(b_1, s_1).$$

Adding these two inequalities, we get a contradiction to the supermodularity of the surplus function v . The function $\tilde{b}$ is thus increasing.

It follows immediately from the single-crossing Assumption ^{horse} 1.5 that the matching $\tilde{b}$ corresponding to a uniform-pricing equilibrium must be increasing.

In each case, the requirement that $\tilde{b}$ be measure preserving, coupled with the facts that $\mathbf{b}$ and $\mathbf{s}$ are strictly increasing when positive, ensures that $\tilde{b}$ is also strictly increasing.

Finally, we note that since b , s and $\tilde{b}$ are strictly increasing when positive, the measure-preserving requirement on $\tilde{b}$ ensures that $\tilde{B}$ can be taken to be the identity.

pend2

A.3 Proof of Lemma ^{prop-constrained eff} 7

Suppose there exists ϕ , $b \in \mathcal{B}$ and $s' \in [0, \bar{s}]$ such that $W(b, s', \phi) > W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi)$. We argue this implies that the seller has a profitable unpriced deviation, and hence, that $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ is not a personalized-price equilibrium.

Let $\varepsilon = [W(b, s', \phi) - W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi)]/3 > 0$ and set $p = h_B(b, s') - h_B(b, \tilde{s}(b)) + p_P(b, \tilde{s}(b)) - \varepsilon$. Note that the seller of type $\sigma = \phi$ can induce a buyer with attribute choice b to buy from him by choosing s' and offering a price p . Since $W(b, s', \phi) > W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi) + 2\varepsilon$,

$$\begin{aligned} v(b, s') - h_B(b, \tilde{s}(b)) + p_B(b, \tilde{s}(b)) - c_S(s', \sigma) \\ &> \Pi_S(\mathbf{s}(\phi), \phi) + h_B(\mathbf{b}(\phi), \mathbf{s}(\phi)) \\ &\quad - p_B(\mathbf{b}(\phi), \mathbf{s}(\phi)) - c_B(\mathbf{b}(\phi), \phi) \\ &\quad - [h_B(b, \tilde{s}(b)) - p_B(b, \tilde{s}(b)) - c_B(b, \phi)] + 2\varepsilon \\ &\geq \Pi_S(\mathbf{s}(\phi), \phi) + 2\varepsilon, \end{aligned} \tag{12}$$

giants

where the last inequality is an implication of ^{buyer IC eq} (12). Now, the payoff to the seller from this deviation is then

$$\begin{aligned} h_S(b, s') + p - c_S(s', \phi) &= h_S(b, s') + h_B(b, s') - h_B(b, \tilde{s}(b)) \\ &\quad + p_P(b, \tilde{s}(b)) - \varepsilon - c_S(s', \phi) \\ &= v(b, s') - h_B(b, \tilde{s}(b)) + p_P(b, \tilde{s}(b)) - c_S(s', \phi) - \varepsilon \\ &> \Pi_S(\mathbf{s}(\phi), \phi) + \varepsilon, \end{aligned}$$

where the inequality follows from ^{giants} (12). Hence, the unpriced deviation is profitable.

The alternative possibility, $s \in \mathcal{S}$ and $b' \in \mathbb{R}_+ \setminus \mathcal{B}$ satisfying $W(b', s, \beta) > W(b(\phi), s(\phi), \phi)$, implies a profitable unpriced deviation for the buyer by an identical argument.

spread

A.4 Proof of Lemma ^{tulip}8

Fix a complete personalized price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$. We need only verify that there are no profitable unpriced deviations from the specification $(\mathbf{b}, \mathbf{s}, \tilde{b}, p'_P)$, where p'_P is restricted to $\mathcal{B} \times \mathcal{S}$. Suppose the buyer has a profitable unpriced deviation, so there exists a type β and an attribute choice $b' \notin \mathcal{B}$, a price $p \in \mathbb{R}$, and $s' \in \mathcal{S}$ with

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}(\mathbf{b}(\beta)), \beta) < h_B(b', s') - p - c_B(b', \beta)$$

and

$$h_S(\tilde{b}(s'), s') - p_P(\tilde{b}(s'), s') < h_S(b', s') + p.$$

Since $s' \in \mathcal{S}$, there is $\sigma \in [0, 1]$ for which we can subtract $c_S(s, \sigma)$ from each side of the second expression to get

$$\begin{aligned} \Pi_B(\mathbf{b}(\beta), \tilde{s}(\mathbf{b}(\beta)), \beta) &< h_B(b', s') - p - c_B(b', \beta) \\ \Pi_S(\mathbf{s}(\sigma), \sigma) &< h_S(b', s') + p - c_S(s', \sigma). \end{aligned}$$

But then there is no $p_P(b', s')$ for which the equalities in Definition ^{slushpump}8 can be satisfied. Unpriced deviations on the part of the seller are analogous.

Suppose $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ is a complete personalized price equilibrium that is not efficient. The inefficiency implies that there exist buyer and seller types $\beta = \sigma = \phi$ and attribute choices b and s for which

$$v(b, s) - c_B(b, \phi) - c_S(s, \phi) > v(\mathbf{b}(\phi), \mathbf{s}(\phi)) - c_B(\mathbf{b}(\phi), \phi) - c_S(\mathbf{s}(\phi), \phi). \quad (13) \quad \text{sofa}$$

Because $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ is a complete personalized price equilibrium, we have

$$\begin{aligned} h_B(\mathbf{b}(\phi), \mathbf{s}(\phi)) - p_P(\mathbf{b}(\phi), \mathbf{s}(\phi)) - c_B(\mathbf{b}(\phi), \phi) \\ \geq h_B(b, s) - p_P(b, s) - c_B(b, \phi) \end{aligned}$$

and

$$\begin{aligned} h_S(\mathbf{b}(\phi), \mathbf{s}(\phi)) + p_P(\mathbf{b}(\phi), \mathbf{s}(\phi)) - c_S(\mathbf{s}(\phi), \phi) \\ \geq h_S(b, s) + p_P(b, s) - c_S(s, \phi). \end{aligned}$$

Adding these two contradicts ^{sofa}(13).

pend4

A.5 Proof of Proposition ^{prop-uniform exist}3.

Since uniform price equilibria involve two sets of qualitatively different incentive constraints (one set corresponding to priced deviations, and the other to unpriced deviations), the existence proof is involved and indirect. We would like to construct a game Γ whose equilibria induce uniform price equilibria. However, the obvious such game Γ is itself difficult to handle, so we work with an approximating double sequence of games $\Gamma^{n,k}$. We verify that each $\Gamma^{n,k}$ has an equilibrium, and then take limits, first n to ∞ , and then k to ∞ , and show that the limiting strategy profile induces a uniform price equilibrium. Loosely, the n index allows us to accommodate (in the limit) the possibility of jumps in the attribute choice functions (precluded in game $\Gamma^{n,k}$), while the k index (in the limit) ensures that unpriced deviations are unprofitable.

A.5.1 Preliminaries

Let $P = \max\{h_B(\bar{b}, \bar{s}), h_S(\bar{b}, \bar{s})\}$. Then P is sufficiently large that no buyer would be willing to purchase any seller attribute choice $s \in [0, \bar{s}]$ at a price exceeding P , nor would any seller be willing to sell to a buyer $b \in [0, \bar{b}]$ at price less than $-P$. We can thus limit prices to the interval $[-P, P]$.

Let Δ be the Lipschitz constant from Assumption ^{horse}I.5, so that for all $\varepsilon > 0$, $s \in [0, \bar{s} - \varepsilon]$, and $b \in [0, \bar{b}]$, we have $h_B(b, s + \varepsilon) - h_B(b, s) < \Delta\varepsilon$. As a result, given a choice between seller s and seller $s + \varepsilon$ at a price higher by $\Delta\varepsilon$, buyers would always choose the former. Equilibrium prices will thus never need to increase at a rate faster than Δ .

A.5.2 The game $\Gamma^{n,k}$

Each game $\Gamma^{n,k}$ has three players, consisting of a buyer, a seller, and a price-setter.

Strategy spaces. We begin by defining the strategy spaces for $\Gamma^{n,k}$. The strategy spaces are only a function of n , with k only affecting the seller's payoffs.

The buyer chooses a pair of functions, $(\mathbf{b}, \tilde{s})$, where $\mathbf{b} : [0, 1] \rightarrow [0, \bar{b}]$ specifies a buyer attribute choice and $\tilde{s} : [0, 1] \rightarrow [0, \bar{s}]$ a seller attribute with which to match, each as a function of the buyer's type, with the restrictions that

$$(\beta' - \beta)/n \leq \mathbf{b}(\beta') - \mathbf{b}(\beta) \leq n(\beta' - \beta) \quad (14) \quad \boxed{\text{I}}$$

and

$$(\beta' - \beta)/n \leq \tilde{s}(\beta') - \tilde{s}(\beta) \leq n(\beta' - \beta) \quad (15) \quad \boxed{\text{II}}$$

for any $\beta < \beta' \in [0, 1]$. We denote the set of pairs of functions $(\mathbf{b}, \tilde{s})$ satisfying (II4) and (II5) normed by the sum of the L^1 norms on the component functions by Υ_B^n , and the “limit” set where (II4) and (II5) need not hold for any n by Υ_B^∞ . As usual, we do not distinguish between functions that agree almost everywhere (this is only relevant in Υ_B^∞).

The seller chooses an increasing function $\mathbf{s}$, where $\mathbf{s} : [0, 1] \rightarrow [0, \bar{s}]$ specifies a seller attribute choice as a function of seller’s type, satisfying

$$(\sigma' - \sigma)/n \leq \mathbf{s}(\sigma') - \mathbf{s}(\sigma) \leq n(\sigma' - \sigma) \quad (16) \quad \boxed{\text{III}}$$

for any $\sigma < \sigma' \in [0, 1]$.

We denote the set of functions $\mathbf{s}$ satisfying (III6) with the L^1 norm by Υ_S^n , and the “limit” set where (III6) need not hold for any n by Υ_S^∞ .

The price setter chooses an increasing function $p_U : [0, \bar{s}] \rightarrow [-P, P]$ satisfying

$$p_U(s') - p_U(s) < 2\Delta(s' - s) \quad (17) \quad \boxed{\text{IV}}$$

for all $s < s' \in [0, \bar{s}]$. Denote the set of increasing functions p_U satisfying (IV7), again endowed with the L^1 norm, by Υ_P (note that Υ_P is not indexed by n). Every function in Υ_P is continuous; indeed the collection Υ_P is equicontinuous.

The sets $\Upsilon^n \equiv \Upsilon_B^n \times \Upsilon_S^n \times \Upsilon_P$ and $\Upsilon^\infty \equiv \Upsilon_B^\infty \times \Upsilon_S^\infty \times \Upsilon_P$, when normed by the sum of the three constituent norms, are compact metric spaces.¹⁷

Payoffs. The buyer’s payoff from $(\mathbf{b}, \tilde{s}) \in \Upsilon_B^n$, when the price-setter has chosen $p_U \in \Upsilon_P$ is

$$\int (h_B(\mathbf{b}(\beta), \tilde{s}(\beta)) - p_U(\tilde{s}(\beta)) - c_B(\mathbf{b}(\beta), \beta)) d\beta. \quad (18) \quad \boxed{\text{boccanegra}}$$

Note that the buyer’s payoff is independent of seller behavior.

¹⁷It suffices for this conclusion to show that Υ is sequentially compact, since sequential compactness is equivalent to compactness for metric spaces (Dunford and Schwartz (1988, p. 20). An argument analogous to that of Helly’s theorem (Billingsley (1986, P. 345)) shows Υ is sequentially compact. In particular, given a sequence $\{(\mathbf{b}^m, \tilde{s}^m, \mathbf{s}^m, p_U^m)\}$, we can choose a subsequence along which each function converges at every rational value in its domain to a limit $\{(\mathbf{b}^\infty, \tilde{s}^\infty, \mathbf{s}^\infty, p_U^\infty)\}$. Because each function in the sequence $\{(\mathbf{b}^m, \tilde{s}^m, \mathbf{s}^m, p_U^m)\}$ is increasing, so must be each limiting function $\{(\mathbf{b}^\infty, \tilde{s}^\infty, \mathbf{s}^\infty, p_U^\infty)\}$. This ensures convergence at every continuity point of the limit functions, and hence almost everywhere, sufficing (for bounded functions) for L^1 convergence.

The price-setter's payoff from $p_U \in \Upsilon_P$, when the buyer and seller have chosen $(\mathbf{b}, \tilde{s}, \mathbf{s}) \in \Upsilon_B^n \times \Upsilon_S^n$ is

$$\int_0^{\tilde{s}} (\varphi_+(p_U(s))[F_B(s) - F_S(s)]_+ + \varphi_-(p_U(s))[F_B(s) - F_S(s)]_-) ds, \quad (19)$$

simon2

where $[x]_+ = \max\{x, 0\}$, $[x]_- = \min\{x, 0\}$, φ_+ is an increasing strictly concave function from $[-P, P]$ into $\mathbb{R}$ with slope bounded away from zero and infinity, φ_- an increasing strictly convex function from $[-P, P]$ into $\mathbb{R}$ with slope bounded away from zero and infinity, and

$$\begin{aligned} F_B(s) &\equiv \lambda\{\beta \mid \tilde{s}(\beta) \leq s\} \\ F_S(s) &\equiv \lambda\{\sigma \mid \mathbf{s}(\sigma) \leq s\}. \end{aligned}$$

Hence, the price-setter has an incentive to raise the price of seller attribute choices in excess demand and lower the price of seller attribute choices in excess supply.

The specification of the seller's payoff is complicated by the need to incorporate incentives arising from the possibility of unpriced deviations. Given an attribute choice s , price p , and price function p_U , set

$$B(s, p, p_U) \equiv \left\{ b \in [0, \bar{b}] \mid h_B(b, s) - p \geq \max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\} \right\}.$$

Hence, $B(s, p, p_U)$ is the set of buyer attribute choices that find attribute choice s at price p (weakly) more attractive than any attribute $s' \in [0, \bar{s}]$ at price $p_U(s')$. Note that for all s and $p_U \in \Upsilon_P$, since there is no a priori restriction on p , $B(s, p, p_U)$ is nonempty for sufficiently low p (possibly requiring $p < -P$, e.g., if $p_U \equiv -P$), and it is empty if $p > p_U(s)$.

mahler

Claim 1

(I.1) If $B(s, p, p_U) \neq \emptyset$, then $B(s, p, p_U) = [b_1, b_2]$ with $b_1 \leq b_2$.

(I.2) For fixed s and p_U , let $\bar{p}(s, p_U) \equiv \max\{p \mid B(s, p, p_U) \neq \emptyset\}$ and write $[b_1(p), b_2(p)]$ for $B(s, p, p_U)$ when $p \leq \bar{p}(s, p_U)$. Denote the set of discontinuity points in the domain of $b_j(p)$ by $\mathcal{D}_j(s, p_U)$. The set $\{s \mid \mathcal{D}_j(s, p_U) \neq \emptyset\}$ has zero Lebesgue measure.

(I.3) Suppose $\{(s^\ell, p^\ell, p_U^\ell)\}_\ell$ is a sequence converging to (s, p, p_U) with $\emptyset \neq B(s^\ell, p^\ell, p_U^\ell) \equiv [b_1^\ell, b_2^\ell]$. Then $B(s, p, p_U) \neq \emptyset$, and so $B(s, p, p_U) = [b_1, b_2]$, where

$$b_1 \leq \liminf_\ell b_1^\ell \leq \limsup_\ell b_2^\ell \leq b_2.$$

^{mahler}
(I.4) Moreover, if $p \notin \mathcal{D}_j(s, p_U) \cup \{\bar{p}(s, p_U)\}$, then $b_j = \lim_\ell b_j^\ell$.

Proof.

1.

^{mahler}
(I.1) Suppose $b_1, b_2 \in B(s, p, p_U)$ with $b_1 < b_2$, and $\hat{b} \notin B(s, p, p_U)$ for some $\hat{b} \in (b_1, b_2)$. Then there exists $\hat{s} \in [0, \bar{s}]$ such that

$$h_B(\hat{b}, s) - p < h_B(\hat{b}, \hat{s}) - p_U(\hat{s}).$$

If $\hat{s} > s$, then assumption ^{horse} I.4 implies

$$\begin{aligned} h_B(b_2, \hat{s}) - h_B(b_2, s) &\geq h_B(\hat{b}, \hat{s}) - h_B(\hat{b}, s) \\ &> p_U(\hat{s}) - p, \end{aligned}$$

contradicting $b_2 \in B(s, p, p_U)$. Similarly, $\hat{s} < s$ contradicts $b_1 \in B(s, p, p_U)$, and so $\hat{s} = s$. But $b_2 \in B(s, p, p_U)$ then implies $p_U(\hat{s}) \geq p$ while $\hat{b} \notin B(s, p, p_U)$ implies $p_U(\hat{s}) < p$, the final contradiction, and so $\hat{b} \in B(s, p, p_U)$. It is immediate that $B(s, p, p_U)$ is closed.

^{mahler}
(I.2) Since $B(s, p', p_U) \supset B(s, p, p_U)$ for $p' < p$, $b_1(p)$ and $b_2(p)$ are monotonic functions of p , and so are continuous except at a countable number of points.

Suppose $p \in \mathcal{D}_1(s, p_U)$, and let $b_1^+ = \lim_{p' \searrow p} b_1(p')$. Since b_1 is left-continuous, $b_1(p) < b_1^+$. Then for all $b \in [b_1(p), b_1^+]$,

$$h_B(b, s) - p = \max_{s' \in [0, \bar{s}]} h_B(b, s') - p_U(s'). \quad (20)$$

dvorak string qtt

From the envelope theorem (^{MilgromSegal02} Milgrom and Segal, 2002, theorem 2), this implies for all $b \in (b_1(p), b_1^+)$,

$$\frac{\partial h_B(b, s)}{\partial b} = \frac{\partial h_B(b, s'(b))}{\partial b},$$

where $s'(b) \in \arg \max_{s' \in [0, \bar{s}]} h_B(b, s') - p_U(s')$. Assumption ^{horse} I.5 then implies $s = s'(b)$ for all $b \in (b_1(p), b_1^+)$, and so $p = p_U(s)$.

Moreover, for all $s'' > s$,

$$h_B(b_1^+, s'') - h_B(b_1^+, s) \leq p_U(s'') - p_U(s)$$

so that

$$\frac{\partial h_B(b_1^+, s)}{\partial s} \leq \liminf_{s'' > s} \frac{p_U(s'') - p_U(s)}{s'' - s},$$

while for all $s' < s$,

$$p_U(s) - p_U(s') \leq h_B(b_1, s) - h_B(b_1, s'),$$

so that

$$\limsup_{s' < s} \frac{p_U(s) - p_U(s')}{s - s'} \leq \frac{\partial h_B(b_1, s)}{\partial s},$$

Consequently, since

$$\frac{\partial h_B(b_1, s)}{\partial s} < \frac{\partial h_B(b_1^+, s)}{\partial s}$$

the price function p_U cannot be differentiable at s . Finally, since $\hat{p}_U$ is a monotonic function, it is differentiable almost everywhere (Royden, 1988, theorem 5.3), and hence $\{s \mid \mathcal{D}_1(s, p_U) \neq \emptyset\}$ has zero Lebesgue measure. A similar argument shows that $\{s \mid \mathcal{D}_2(s, p_U) \neq \emptyset\}$ has zero Lebesgue measure.

^{mahler}(I.3) Suppose $\{(s^\ell, p^\ell, p_U^\ell)\}_\ell$ is a sequence converging to (s, p, p_U) , and let $\{b^\ell\}$ be a sequence of attributes with $b^\ell \in B(s^\ell, p^\ell, p_U^\ell)$ for all ℓ . Without loss of generality, we assume $\{b^\ell\}$ is a convergent sequence with limit b . Since

$$h_B(b^\ell, s^\ell) - p^\ell \geq \max_{s' \in [0, \bar{s}]} \{h_B(b^\ell, s') - p_U^\ell(s')\}, \quad \forall \ell,$$

taking limits gives

$$h_B(b, s) - p \geq \max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\},$$

and so $b \in B(s, p, p_U)$.

^{mahler}(I.4) Consider b_2 and suppose $p \notin \mathcal{D}_2(s, p_U) \cup \{\bar{p}(s, p_U)\}$. Hence, $b_2 = b_2^+ = \lim_{p' \searrow p} b_2(p')$. Consider $b \in (b_1^+, b_2)$. For $p' > p$ sufficiently close to p , we have $b \in B(s, p', p_U)$, and so

$$h_B(b, s) - p > \max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\}.$$

Consequently, for ℓ sufficiently large,

$$h_B(b, s^\ell) - p^\ell > \max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\},$$

i.e., $b \in B(s^\ell, p^\ell, p_U^\ell)$. This implies that $b_2^\ell(p^\ell) > b$, and hence $\liminf b_2^\ell \geq b$. Since this holds for all $b \in (b_1^+, b_2)$ and the claim ^{mahler}I.3 has established $\limsup_\ell b_2^\ell \leq b_2$, we have $\lim_\ell b_2^\ell = b_2$. The argument for b_1 is an obvious modification of this argument.

■

Our specification of seller payoffs captures in a smooth way the idea that he attracts buyer attributes in $B(s, p, p_U)$. Fix (s, p, p_U) and suppose $\lambda(\{\beta \mid \mathbf{b}(\beta) \in B(s, p, p_U)\}) > 0$. Since $\mathbf{b}$ is strictly increasing and continuous, it then follows from claim I that $\mathbf{b}([0, 1]) \cap B(s, p, p_U) = [b'_1, b'_2]$ for some $0 \leq b'_1 < b'_2 \leq \bar{b}$. We next define a measure on $[0, \bar{b}]$ to reflect the idea that in the limit (as k gets large), the seller attracts the worst buyers in $[b'_1, b'_2]$. Set

$$f^k(b) = \frac{ke^{-k(b-b'_1)}}{1 - e^{-k(b'_2-b'_1)}}.$$

The two critical properties of this function are that $\int_{b'_1}^{b'_2} f^k(b)db = 1$ for all k , and

$$\lim_{k \rightarrow \infty} \int_{b'_1}^{b'_2} g(b)f^k(b)db = g(b'_1) \quad (21) \quad \boxed{\text{Symph 2}}$$

for all continuous g . Define the measure ψ^k on $[0, \bar{b}]$ by

$$\psi^k(A) = \int_{\mathbf{b}(\beta) \in A \cap [b'_1, b'_2]} f^k(\mathbf{b}(\beta)) d\beta \quad \text{for all measurable } A$$

and the associated expected payoff to the seller from $(s, p, \mathbf{b}, p_U)$ by

$$H^k(s, p, \mathbf{b}, p_U) \equiv \int h_S(b, s) \psi^k(db) + p. \quad (22) \quad \boxed{\text{cello}}$$

This function depends upon p_U and $\mathbf{b}$ through ψ^k 's dependence on $B(s, p, p_U)$ and $\mathbf{b}$. Moreover, ψ^k inherits any decorations or superscripts from p_U and $\mathbf{b}$. For later reference, note that for fixed s , $\mathbf{b}$, and p_U , the function $H^k(s, p, \mathbf{b}, p_U)$ is continuous from the left in p (since $\mathbf{b}$ satisfies (I4) and both $b_1(p)$ and $b_2(p)$, defined just before claim I, are left-continuous).

We set $\tilde{P}(s, \mathbf{b}, p_U) \equiv \{p \mid \lambda(\{\beta \mid \mathbf{b}(\beta) \in B(s, p, p_U)\}) > 0\}$, and noting that this set is nonempty, define

$$\bar{H}^k(s, \mathbf{b}, p_U) \equiv \max \left\{ \sup_{p \in \tilde{P}(s, \mathbf{b}, p_U)} H^k(s, p, \mathbf{b}, p_U), h_S(0, s) + p_U(s) \right\}. \quad (23) \quad \boxed{\text{mice1}}$$

Taking the maximum over $\sup_{p \in \tilde{P}(s, \mathbf{b}, p_U)} H^k(s, p, \mathbf{b}, p_U)$ and $h_S(0, s) + p_U(s)$ effectively assumes that the seller can always sell attribute choice s at the posted price $p_U(s)$, though perhaps only attracting buyer attribute choice 0. This ensures that the price setter indeed acts like a price setter, allowing buyers to purchase any s at $p_U(s)$ and sellers to sell at that price. Notice

that if $\tilde{P}(s, \mathbf{b}, p_U)$ contains $p_U(s) - \epsilon$ for all $\epsilon > 0$, then the first term in $(23)^{\text{mice1}}$ will be the maximum.

The seller's payoff from $\mathbf{s} \in \Upsilon_S^n$ when the buyer and price-setter have chosen $(\mathbf{b}, \tilde{s}, p_U) \in \Upsilon_B^n \times \Upsilon_P$ is then

$$\int \left(\bar{H}^k(\mathbf{s}(\sigma), \mathbf{b}, p_U) - c_S(\mathbf{s}(\sigma), \sigma) \right) d\sigma. \quad (24) \quad \boxed{\text{mice}}$$

A.5.3 Equilibrium in game $\Gamma^{n,k}$

Our next task is to show that each game $\Gamma^{n,k}$ has a Nash equilibrium, and that the price-setter plays a pure strategy in any such equilibrium. To do this, we first note that the price-setter's payoff is strictly concave in p_U , though we have not shown that the buyer's and sellers's payoffs are even quasiconcave. If the payoffs functions in game $\Gamma^{n,k}$ are continuous, then Glicksberg's fixed point theorem, applied to the game where we allow the buyer and seller to randomize, yields a Nash equilibrium in which the buyer and seller may randomize, but the price-setter does not.

It thus remains to show:

Claim 2 *The buyer, price-setter and seller payoff functions given by $(18)^{\text{bocca}}, (19)^{\text{bocca}}$ and $(24)^{\text{mice}}$, are continuous functions of $(\mathbf{b}, \mathbf{s}, \tilde{s}, p_U)$ on Υ^n .*

Proof. We first note that for increasing, bounded functions on a compact set, L^1 convergence implies convergence almost everywhere.¹⁸

Consider first the buyer. The functions $\mathbf{b}$, $\tilde{s}$, and p_U are bounded functions on compact sets, and hence the absolute value of each of these functions is dominated by an integrable function (e.g., the constant function equal to the relevant upper bound). The continuity of the buyer's payoff then follows immediately from Lebesgue's dominated convergence theorem, if we can show that the convergence of $\mathbf{b}$, p_U , and $\tilde{s}$ in the L^1 norm (and hence almost everywhere) implies the convergence almost everywhere of $h_B(\mathbf{b}(\beta), \tilde{s}(\beta))$, $p_U(\tilde{s}(\beta))$ and $c_B(\mathbf{b}(\beta), \beta)$ (note that we are talking about sequences of functions within a given game $\Gamma^{n,k}$). The first and the third of these follows from

¹⁸Suppose f_n converges in L^1 norm to an increasing function f without converging almost everywhere. Then since f is discontinuous on a set of measure zero, there exists (for example) a continuity point x of f with $g(x) \equiv \lim f_n(x) > f(x)$ (with the case $\lim f_n(x) < f(x)$ analogous). The continuity of f then ensures that for some point $y > x$, some $\varepsilon > 0$, all $z \in [x, y]$ and for all sufficiently large n , we have $f_n(z) \geq f(y) + \varepsilon \geq f(z) + \varepsilon$. This in turn ensures that $\int |f_n(z) - f(z)| dz > (y - x)\varepsilon$, precluding the L^1 convergence of $\{f_n\}_{n=1}^\infty$ to f .

the continuity of h_B and c_B (from Assumption $\overline{\text{horse}}_{\text{II}}$), while for the remaining case it suffices to note that the collection Υ_P is equicontinuous.

Consider now the price-setter. Suppose $\mathbf{s}^\ell$ converges in L^1 , and so almost everywhere, to $\mathbf{s}$. Then F_S^ℓ converges weakly to F_S (and so a.e.).¹⁹ Similarly, if $\tilde{s}^\ell$ converges in L^1 to $\tilde{s}$, then F_B^ℓ converges a.e. to F_B . Continuity for the price-setter's payoff then follows from arguments analogous to those applied to the buyer, since we have convergence almost everywhere of $p_U(s)[F_B(s) - F_S(s)]$.

Finally, we turn to the seller, where the proof of continuity is more involved. It suffices to argue that $\bar{H}^k(s, \mathbf{b}, p_U)$ is continuous in $(s, \mathbf{b}, p_U)$ for almost all s (since $\tilde{s}$ is irrelevant in the determination of the seller's payoff and the continuity with respect to $\mathbf{s}$ is then obvious, at which point again appeal to Lebesgue's dominated convergence theorem).

Fix a point $(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$ and a sequence $(s^\ell, \mathbf{b}^\ell, p_U^\ell)$ converging to it. Since we need continuity for almost all $s \in [0, \bar{s}]$, we can assume $\mathcal{D}_1(\hat{s}, \hat{p}_U) \cup \mathcal{D}_2(\hat{s}, \hat{p}_U) = \emptyset$ (or equivalently that $\hat{p}_U$ is differentiable at $\hat{s}$ (cf. the proof of claim $\overline{\text{mahler}}_{\text{I.2}}$)).

We show $\lim_{\ell \rightarrow \infty} \bar{H}^k(s^\ell, \mathbf{b}^\ell, p_U^\ell) = \bar{H}^k(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$. Notice that $\bar{H}^k(s, \mathbf{b}, p_U)$ is the maximum of two terms. It suffices to show that

$$\begin{aligned} \lim_{\ell \rightarrow \infty} \sup_{p \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)} H^k(s^\ell, p, \mathbf{b}^\ell, p_U^\ell) &= \sup_{p \in \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)} H^k(\hat{s}, p, \hat{\mathbf{b}}, \hat{p}_U) \\ \lim_{\ell \rightarrow \infty} h_S(0, s^\ell) + p_U^\ell(s^\ell) &= h_S(0, \hat{s}) + \hat{p}_U(\hat{s}). \end{aligned}$$

The second is immediate from the continuity of h_S and the differentiability of $\hat{p}_U$ at $\hat{s}$. We accordingly turn to the first. To conserve on notation, let $\sup_{p \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)} H^k(s^\ell, p, \mathbf{b}^\ell, p_U^\ell) \equiv \bar{\bar{H}}^k(s^\ell, \mathbf{b}^\ell, p_U^\ell)$.

We now first show that

$$\liminf_{\ell \rightarrow \infty} \bar{\bar{H}}^k(s^\ell, \mathbf{b}^\ell, p_U^\ell) \geq \bar{\bar{H}}^k(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U). \quad (25) \quad \boxed{\text{tristan}}$$

For all $\varepsilon > 0$ there exists $\hat{p} \in \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$ such that

$$H^k(\hat{s}, \hat{p}, \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon/2 \geq \bar{\bar{H}}^k(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U).$$

¹⁹Fix $\varepsilon > 0$. By Egoroff's theorem ($\overline{\text{Royden88}}$ (Royden, 1988, p.73), $\mathbf{s}^\ell$ converges uniformly to $\mathbf{s}$ on a set E of measure at least $1 - \varepsilon$. Suppose s is a continuity point of F_S . There then exists $\delta > 0$ such that $|F_S(s) - F_S(s')| < \varepsilon$ for all $|s - s'| \leq \delta$. There exists ℓ' such that, for all $\sigma \in E$, for all $\ell > \ell'$, $|\mathbf{s}^\ell(\sigma) - \mathbf{s}(\sigma)| < \delta$. Consequently, $F_S^\ell(s) \leq F_S(s + \delta) + \varepsilon$ and $F_S(s - \delta) - \varepsilon \leq F_S^\ell(s)$, and so $|F_S^\ell(s) - F_S(s)| < 2\varepsilon$. Hence, F_S^ℓ converges weakly to F_S .

Since $H^k(\hat{s}, p, \hat{\mathbf{b}}, \hat{p}_U)$ is continuous from the left in p , there exists $\hat{p}' < \bar{p}(\hat{s}, \hat{p}_U)$ with $\hat{p}' \leq \hat{p}$ satisfying

$$|H^k(\hat{s}, \hat{p}, \hat{\mathbf{b}}, \hat{p}_U) - H^k(\hat{s}, \hat{p}', \hat{\mathbf{b}}, \hat{p}_U)| < \varepsilon/2,$$

and so

$$H^k(\hat{s}, \hat{p}', \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon \geq \bar{H}^k(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U).$$

Since $\hat{p}' \notin \mathcal{D}_1(\hat{s}, \hat{p}_U) \cup \mathcal{D}_2(\hat{s}, \hat{p}_U) \cup \{\bar{p}(\hat{s}, \hat{p}_U)\}$ and $\hat{\mathbf{b}}$ satisfies (II.4), for sufficiently large ℓ , $\hat{p}' \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)$, and so (applying claim I.3) ^{mahler}

$$\lim_{\ell \rightarrow \infty} H^k(s^\ell, \hat{p}', \mathbf{b}^\ell, p_U^\ell) = H^k(\hat{s}, \hat{p}', \hat{\mathbf{b}}, \hat{p}_U).$$

Hence,

$$\liminf_{\ell \rightarrow \infty} \bar{H}^k(s^\ell, \mathbf{b}^\ell, p_U^\ell) + \varepsilon \geq \bar{H}^k(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U), \quad \forall \varepsilon > 0,$$

yielding (25) ^{tristan}.

We now argue that

$$\bar{H}^k(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U) \geq \limsup_{\ell \rightarrow \infty} \bar{H}^k(s^\ell, \mathbf{b}^\ell, p_U^\ell), \quad (26) \quad \boxed{\text{isolde}}$$

which with (25) ^{tristan} gives continuity.

Fix $\varepsilon > 0$. For each ℓ , there exists $p^\ell \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)$ such that

$$H^k(s^\ell, p^\ell, \mathbf{b}^\ell, p_U^\ell) + \varepsilon \geq \bar{H}^k(s^\ell, \mathbf{b}^\ell, p_U^\ell).$$

Without loss of generality, we can assume $\{p^\ell\}_\ell$ is a convergent sequence, with limit $\hat{p}$. Suppose first that $\hat{p} \in \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$. If $\hat{p} \neq \{\bar{p}(\hat{s}, \hat{p}_U)\}$, it is immediate that

$$H^k(\hat{s}, \hat{p}, \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon \geq \limsup_{\ell \rightarrow \infty} \bar{H}^k(s^\ell, \mathbf{b}^\ell, p_U^\ell), \quad (27) \quad \boxed{\text{prager}}$$

which (since it holds for all ε) implies (26) ^{isolde}.

Suppose now that $\hat{p} \notin \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$ or $p = \bar{p}(\hat{s}, \hat{p}_U)$. Since $\hat{p}_U$ is differentiable at $\hat{s}$, there cannot be a nondegenerate interval of buyer attributes indifferent between $(\hat{s}, \hat{p})$ and the unconstrained optimal seller attribute under $\hat{p}_U$. This implies $\hat{\mathbf{b}}([0, 1]) \cap B(\hat{s}, \hat{p}, \hat{p}_U) = \{\hat{b}\}$ for some $\hat{b}$, and so

$$\bar{H}^k(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U) \geq h_S(\hat{b}, \hat{s}) + \hat{p}.$$

From claim I.3 ^{mahler},

$$\lim_{\ell \rightarrow \infty} H^k(s^\ell, p^\ell, \mathbf{b}^\ell, p_U^\ell) + \varepsilon = h_S(\hat{b}, \hat{s}) + \hat{p} + \varepsilon,$$

and so

$$\bar{H}^k(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon \geq \limsup_{\ell \rightarrow \infty} \bar{H}^k(s^\ell, \mathbf{b}^\ell, p_U^\ell),$$

which (since it holds for all $\varepsilon > 0$) implies (26) ^{isolde}. ■

A.5.4 The limit $n \rightarrow \infty$

We continue to fix k and now examine the limit as $n \rightarrow \infty$. In particular, let $\{(\xi_B^{n,k}, \xi_S^{n,k}, p_U^{n,k})\}_n \subset \Delta(\Upsilon_B^\infty) \times \Delta(\Upsilon_S^\infty) \times \Upsilon_P$ be a sequence of Nash equilibria of the games $\Gamma^{n,k}$. Without loss of generality (since the relevant spaces are sequentially compact), we may assume that both the sequence of equilibria converges to some limit $(\xi_B^k, \xi_S^k, p_U^k)$, and that each players' payoffs also converge.

We now examine the limit $(\xi_B^k, \xi_S^k, p_U^k)$. Intuitively, we would like to think of this profile as the equilibrium of a “limit game.” However, the definition of this limit game is not straightforward. When defining the payoffs in game $\Gamma^{n,k}$, especially the sellers' payoffs, we have relied heavily on the strategies $\mathbf{b}, \mathbf{s}$ that have properties (such as being strictly increasing and continuous) that need not carry over to their limits. In establishing properties of $(\xi_B^k, \xi_S^k, p_U^k)$, we will accordingly typically begin our argument in the limit, and then pass back to the approximating equilibrium profile $(\xi_B^{n,k}, \xi_S^{n,k}, p_U^{n,k})$ to obtain a contradiction. The latter step of the argument is notationally cumbersome, and we do not always make the approximation explicit. It is accordingly helpful to keep in mind what it means for $(\xi_B^{n,k}, \xi_S^{n,k}, p_U^{n,k})$ to converge to the pure profile $(\mathbf{b}^k, \tilde{\mathbf{s}}^k, \mathbf{s}^k, p_U^k)$: For all k and for all $\varepsilon > 0$ there exists n' such that for all $n \geq n'$,

$$\begin{aligned} \xi_B^{n,k} \left(\{(\mathbf{b}, \tilde{\mathbf{s}}) \in \Upsilon_B^n \mid \int |\mathbf{b}(\beta) - \mathbf{b}^k(\beta)| d\beta < \varepsilon, \int |\tilde{\mathbf{s}}(\beta) - \tilde{\mathbf{s}}^k(\beta)| d\beta < \varepsilon\} \right) \\ \geq 1 - \varepsilon, \end{aligned}$$

$$\xi_S^{n,k} \left(\{\mathbf{s} \in \Upsilon_S^n \mid \int |\mathbf{s}(\sigma) - \mathbf{s}^k(\sigma)| d\sigma < \varepsilon\} \right) \geq 1 - \varepsilon,$$

and

$$\int |p_U^{n,k}(s) - p_U^k(s)| ds < \varepsilon.$$

Note that while the seller is best responding to $\xi_B^{n,k}$ in choosing $\mathbf{s}$, the choice of p implicit in (23) is made after $(\mathbf{b}, \tilde{\mathbf{s}})$ is realized.

Claim 3 *The limit profile $(\xi_B^k, \xi_S^k, p_U^k)$ is pure, which we denote $(\mathbf{b}^k, \tilde{\mathbf{s}}^k, \mathbf{s}^k, p_U^k)$.*

Proof. Consider the buyer (the case of the seller is analogous). Let $\xi_{B,b}^k$ and $\xi_{B,\tilde{\mathbf{s}}}^k$ denote the marginal measures induced on choices b and $\tilde{\mathbf{s}}$.

Suppose the buyer's strategy is not pure. Then define a pair of increasing functions $\mathbf{b}' : [0, 1] \rightarrow \{-1\} \cup [0, \bar{b}]$ and $\tilde{\mathbf{s}}' : [0, 1] \rightarrow \{-1\} \cup [0, \bar{\mathbf{s}}]$ by

$$\begin{aligned} \mathbf{b}'(\beta) &= \inf \{b \mid \xi_{B,b}^k(b) \geq \beta\} \\ \text{and} \quad \tilde{\mathbf{s}}'(\beta) &= \inf \{\tilde{\mathbf{s}} \mid \xi_{B,\tilde{\mathbf{s}}}^k(\tilde{\mathbf{s}}) \geq \beta\}. \end{aligned}$$

These functions give the same distribution of b and $\tilde{s}$ in the market, but feature positive assortivity between the buyer's types and attribute choice, and between the buyer's attribute choice the seller attribute with which the buyer matches, both of which increase the buyer's payoff. Hence, this pure strategy strictly increases the buyer's payoff. It then follows from straightforward continuity arguments that for sufficiently large n , i.e., for a game in which the slope requirements on the buyer's strategy are sufficiently weak and the equilibrium profile $(\xi_B^{n,k}, \xi_S^{n,k}, p_U^{n,k})$ is sufficiently close to $(\xi_B^k, \xi_S^k, p_U^k)$, there is a pure strategy sufficiently close to $\mathbf{b}'$ and $\tilde{s}'$ as to give the buyer a payoff higher than his supposed equilibrium payoff in $\Gamma^{n,k}$, a contradiction. Hence, the buyer cannot mix. ■

simon

Claim 4 *The profile $(\mathbf{b}^k, \tilde{s}^k, \mathbf{s}^k, p_U^k)$ balances the market, i.e., $F_B^k(s) = F_S^k(s)$ for all s .*

Proof. Since F_B^k and F_S^k are continuous from the right, it suffices to show that they agree almost everywhere. We first argue that $F_B^k(s) - F_S^k(s) \leq 0$ almost everywhere. Suppose this is not the case, so there exists $s < \bar{s}$ with $F_B^k(s) - F_S^k(s) = \varepsilon > 0$ and with s a continuity point of $F_B^k - F_S^k$. Then there exists s_1 and s_2 with $s \in [s_1, s_2)$, and $F_B^k(s) - F_S^k(s) \geq \varepsilon/2$ on $[s_1, s_2]$ and either $s_1 = 0$ or, for every $\eta > 0$, there is a value $s_\eta \in [s_1 - \eta, s_1)$ with $F_B^k(s_\eta) - F_S^k(s_\eta) < \varepsilon/2$ (note that $F_B^k(s_\eta) - F_S^k(s_\eta)$ may be negative, and so is bounded below by -1). We consider the case in which $s_1 > 0$ and $p_U^k(s_1) < p_U^k(s_2)$, with the remaining cases a straightforward simplification.

Since $F_B^k(s) - F_S^k(s) > 0$ on $[s_1, s_2]$, for fixed $p_U^k(s_1)$ and $p_U^k(s_2)$, the price setter must be setting prices as large as possible on this interval. If not, there is a price function $\hat{p}_U \in \Upsilon_P$ with $\hat{p}_U(s) \geq p_U^k(s)$ for all s and $\hat{p}_U(s) > p_U^k(s)$ for some s yielding strictly higher payoffs to the price-setter than p_U^k in $\Gamma^{n,k}$ for sufficiently large n , when the buyer and seller choose $(\xi_B^{n,k}, \xi_S^{n,k})$. But this contradicts the equilibrium property of $(\xi_B^{n,k}, \xi_S^{n,k}, p_U^{n,k})$.

Hence, there exists $s' \in [s_1, s_2]$ such that $dp_U(s)/ds = 2\Delta$ on (s_1, s') and $p_U^k(s) = p_U^k(s_2)$ for $s \in [s', s_2]$. That is, prices increase at the maximum rate possible until hitting $p_U^k(s_2)$ (with $s' = s_2$ possible, but since $p_U^k(s_1) < p_U^k(s_2)$, we have $s_1 < s'$). Consequently, $\tilde{s}([0, 1]) \cup [s_1, s_2] \subset \{s_1, s_2\}$, i.e., buyers demand only seller attribute choices s_1 and s_2 from this interval. [Since all seller attribute choices in $[s', s_2]$ command the same price, buyers demand only attribute choice s_2 from this set, while the price of a seller attribute choice increases sufficiently quickly on $[s_1, s']$ that from this set buyers demand only s_1 .]

Since for every $\eta > 0$, there exists $s_\eta \in [s_1 - \eta, s_1)$ with $F_B^k(s_1) - F_B^k(s_\eta) < \epsilon/2$, the buyer must choose s_1 for some buyer types. But if the buyer is to choose s_1 , there is a range of seller attributes just below s_1 with prices that are not too low, that is, there exists $\eta' > 0$ such that

$$p_U^k(s) > p_U^k(s_1) - \Delta(s_1 - s)$$

for all $s \in [s_1 - \eta', s_1)$. Consider now the price function $\hat{p}_U^\eta \in \Upsilon_P$ given by

$$\hat{p}_U^\eta(s) = \begin{cases} p_U^k(s), & \text{if } s \geq s', \\ \min\{p_U^k(s_1 - \eta) + 2\Delta(s - s_1 + \eta), p_U^k(s')\}, & \text{if } s \in (s_1 - \eta, s'), \\ p_U^k(s), & \text{if } s \leq s_1 - \eta, \end{cases}$$

and note that $p_U^0 = p_U^k$. The price-setter's payoff from choosing $p_U^\eta \in \Upsilon_P$ (less the payoff from p_U^k) is bounded below by

$$\begin{aligned} & \int_{s_1}^{s'} \left(\varphi_+(p_U^\eta(s)) - \varphi_+(p_U^k(s)) \right) [\epsilon/2]_+ ds \\ & + \int_{s_1 - \eta}^{s_1} \left(\varphi_-(p_U^\eta(s)) - \varphi_-(p_U^k(s)) \right) [-1]_- ds. \end{aligned}$$

For η close to 0, $p_U^\eta(s) < p_U^k(s')$ for all $s \in (s_1, s_1 + (s' - s_1)/2)$, and so for small η , the above expression is no smaller than

$$\begin{aligned} & \int_{s_1}^{s_1 + (s' - s_1)/2} \left\{ \varphi_+(p_U^k(s_1) + \Delta\eta + 2\Delta(s - s_1)) \right. \\ & \quad \left. - \varphi_+(p_U^k(s_1) + 2\Delta(s - s_1)) \right\} [\epsilon/2]_+ ds \\ & + \int_{s_1 - \eta}^{s_1} \left(\varphi_-(p_U^k(s_1 - \eta) + 2\Delta\eta) - \varphi_-(p_U^k(s_1 - \eta)) \right) [-1]_- ds. \end{aligned}$$

Since the first integral is of the same order of magnitude as η , while the second is lower order, for sufficiently small η , the lower bound is strictly positive, implying the price-setter has a profitable deviation (in $\Gamma^{n,k}$ for large n), a contradiction.

We conclude that $F_B^k(s) - F_S^k(s) \leq 0$ for almost all s . It remains to argue that it is not negative on a set of positive measure. Suppose it is. Then there must exist a seller characteristic $\hat{s} > 0$ such that $p_U(s) = -P$ for $s < \hat{s}$, $F_B^k(s) - F_S^k(s) < 0$ for a positive-measure subset of $[0, \hat{s}]$, and $F_B^k(s) - F_S^k(s) = 0$ for almost all $s > \hat{s}$. But then no seller would choose attributes in $[0, \hat{s})$, a contradiction. ■

A.5.5 The limit $k \rightarrow \infty$

Without loss of generality (since the relevant spaces are sequentially compact), we may assume that both the sequence $\{(\mathbf{b}^k, \tilde{s}^k, \mathbf{s}^k, p_U^k)\}_k \subset \Upsilon_B^\infty \times \Upsilon_S^\infty \times \Upsilon_P$ converges to some limit $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$, and that each players' payoffs also converge. Since we have not proved uniform convergence in k as $n \rightarrow \infty$, we cannot assert the existence of a subsequence of $\{(\xi_B^{n,k}, \xi_S^{n,k}, p_U^{n,k})\}_{n,k}$ converging to $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$. However, it suffices for us that for all $\varepsilon > 0$ there exists K such that for all $k \geq K$, there exists n_k , such that for all $n \geq n_k$, the profile $(\xi_B^{n,k}, \xi_S^{n,k}, p_U^{n,k})$ is within ε of $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$, with a similar claim holding for the associated payoffs.

We now seek a characterization of the seller's payoffs. We know that $\mathbf{b}^*$ and $\tilde{s}^*$ are increasing, and the market-clearing result of Claim 4^{Simon} carries over to the limit $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$. Intuitively, we would like to use these facts to conclude that there is positive assortative matching, and indeed that a seller of type σ matches with a buyer of type $\beta = \sigma$. However, these properties may not hold if $\mathbf{b}$ and $\tilde{s}$ are not strictly increasing (properties we need to establish in any case for the existence of a uniform price equilibrium). Moreover, even if we had such matching, the specification of the seller's payoffs given by (24)^{rice} leaves open the possibility that the (gross) payoff to a seller of type σ choosing attribute s may not be given by $h_S(\mathbf{b}(\sigma), s) + p_U(s)$. Hence, the buyers that sellers are implicitly choosing in their payoff calculations may not duplicate those whose seller choices balance the market.

Our first step in addressing these issues is to show that the buyer's limiting attribute choice function is indeed strictly increasing:

rhyme

Claim 5 *The function $\mathbf{b}^*$ is strictly increasing when nonzero.*

Proof. By construction, $\mathbf{b}^*$ is weakly increasing. We show that $\beta'' > \beta'$ and $\mathbf{b}^*(\beta') > 0$ imply $\mathbf{b}^*(\beta'') > \mathbf{b}^*(\beta')$. Suppose to the contrary that $b = \mathbf{b}^*(\beta) > 0$ for two distinct values of β .

Define $\beta_1 \equiv \inf\{\beta \mid \mathbf{b}^*(\beta) = b\}$, $\beta_2 \equiv \sup\{\beta \mid \mathbf{b}^*(\beta) = b\}$, and $\bar{\beta} = (\beta_1 + \beta_2)/2$. We assume $0 < \beta_1$ and $\beta_2 < 1$ (if equality holds in either case, then the argument is modified in the obvious manner). We now define a new attribute choice function (as a function of a parameter $\eta > 0$) that is strictly increasing on a neighborhood of $[\beta_1, \beta_2]$ and agrees with $\mathbf{b}^*$ outside that neighborhood. First, define

$$\begin{aligned} \beta_1^\eta &\equiv \inf\{\beta \leq \beta_1 \mid \mathbf{b}^*(\beta) \geq b + \eta(\beta - \bar{\beta})\} \\ \text{and} \quad \beta_2^\eta &\equiv \sup\{\beta \geq \beta_2 \mid \mathbf{b}^*(\beta) \leq b + \eta(\beta - \bar{\beta})\}. \end{aligned}$$

Note that as $\eta \rightarrow 0$, $\beta_j^\eta \rightarrow \beta_j$ for $j = 1, 2$. Finally, define

$$\mathbf{b}^\eta(\beta) = \begin{cases} \mathbf{b}^*(\beta), & \text{if } \beta > \beta_2^\eta, \\ b + \eta(\beta - \bar{\beta}), & \text{if } \beta \in [\beta_1^\eta, \beta_2^\eta], \\ \mathbf{b}^*(\beta), & \text{if } \beta < \beta_1^\eta. \end{cases}$$

The payoffs to the buyer under $\mathbf{b}^\eta$ less that under $\mathbf{b}^*$ is

$$\int_{\beta_1^\eta}^{\beta_2^\eta} h_B(\mathbf{b}^\eta(\beta), \tilde{s}(\beta)) - h_B(\mathbf{b}^*(\beta), \tilde{s}(\beta)) - [c_B(\mathbf{b}^\eta(\beta), \beta) - c_B(\mathbf{b}^*(\beta), \beta)] d\beta. \quad (28) \quad \boxed{\text{rachmaninov}}$$

Now,

$$\begin{aligned} & \int_{\beta_1}^{\beta_2} [c_B(\mathbf{b}^\eta(\beta), \beta) - c_B(\mathbf{b}^*(\beta), \beta)] d\beta \\ &= \int_{\beta_1}^{\beta_2} \left[\frac{\partial c_B(b, \beta)}{\partial b} \eta(\beta - \bar{\beta}) + o(\eta) \right] d\beta \\ &= \eta \int_0^{(\beta_2 - \beta_1)/2} \left[\frac{\partial c_B(b, \bar{\beta} + x)}{\partial b} - \frac{\partial c_B(b, \bar{\beta} - x)}{\partial b} \right] x dx + o(\eta). \end{aligned}$$

From assumption ^{horse}1.2, the integrand is strictly negative, and so the integral is strictly negative and independent of η . Since $\tilde{s}$ is increasing, a similar argument applied to the difference in the premuneration values shows that

$$\begin{aligned} & \int_{\beta_1}^{\beta_2} h_B(\mathbf{b}^\eta(\beta), \tilde{s}(\beta)) - h_B(\mathbf{b}^*(\beta), \tilde{s}(\beta)) - [c_B(\mathbf{b}^\eta(\beta), \beta) - c_B(\mathbf{b}^*(\beta), \beta)] d\beta \\ & \geq \eta \int_0^{(\beta_2 - \beta_1)/2} \left[\frac{\partial c_B(b, \bar{\beta} - x)}{\partial b} - \frac{\partial c_B(b, \bar{\beta} + x)}{\partial b} \right] x d\beta + o(\eta). \end{aligned}$$

It remains to argue that the contribution to ^{bachmaninov}(28) from the intervals $[\beta_1^\eta, \beta_1)$ and $(\beta_2, \beta_2^\eta]$ is of order $o(\eta)$. But this is immediate, since $|\mathbf{b}^\eta(\beta) - \mathbf{b}^*(\beta)| \leq \eta$ and $\beta_j^\eta \rightarrow \beta_j$ as $\eta \rightarrow 0$ (for $j = 1, 2$). Hence, for $\eta > 0$ sufficiently small, $\mathbf{b}^\eta$ gives the buyer a strictly higher payoff under ^{boccanegra}(18) than $\mathbf{b}^*$. But, then by a now familiar argument, the buyer has a profitable deviation in $\Gamma^{n,k}$ for sufficiently large k , and then n , a contradiction. So $\mathbf{b}^*$ is strictly increasing when nonzero. ■

Given an attribute choice s , let $[\beta_1(s), \beta_2(s)]$ be the closure of the set of buyers for whom $\tilde{s}^*(\beta) = s$. Because $\tilde{s}$ is increasing, this set is an interval (though possibly degenerate) if it is nonempty. Let $b_1(s) = \lim_{\beta \downarrow \beta_1(s)} \mathbf{b}(s)$ and $b_2(s) = \lim_{\beta \uparrow \beta_2(s)} \mathbf{b}(s)$.

Claim 6

$$\begin{aligned} \lim_k \lim_n \int \int \bar{H}^k(\mathbf{s}^{n,k}(\sigma), \mathbf{b}^{n,k}, p_U^{n,k}) - c_S(\mathbf{s}^{n,k}(\sigma), \sigma) d\sigma d\xi^{n,k} \\ \leq \int h_S(b_1(\mathbf{s}^*(\sigma)), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)) - c_S(\mathbf{s}^*(\sigma), \sigma) d\sigma. \end{aligned}$$

Proof. Fix $s = \mathbf{s}^*(\sigma)$ for some seller type σ . For fixed k and n , any $p > p_U^{n,k}(s)$ gives an empty $B(s, p, p_U^{n,k})$, and so a payoff $h_S(0, s) + p_U(s) - c_S(s, \sigma) \leq h_S(b_1(\mathbf{s}^*(\sigma)), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)) - c_S(\mathbf{s}^*(\sigma), \sigma)$. For $p < p_U^{n,k}(s)$, for large n , there is a positive measure of buyers who find s attractive (since $\tilde{s}^k(\sigma)$ is optimal for buyer σ under p_U^k). As k becomes arbitrarily large, the average value of the attracted attributes under ψ^k converges to the value of the worst type (see (21)), which is at best $b_1(s)$, and so in the limit the payoff cannot exceed $h_S(b_1(s), \mathbf{s}(\sigma)) + p_U^*(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma)$. ■

We now argue that $\tilde{s}$ and $\mathbf{s}$ are strictly increasing. From claim ^{Simon}4, $\tilde{s}(x) = \mathbf{s}(x)$ for almost all $x \in [0, 1]$, and so $\tilde{s}^*$ is strictly increasing if $\mathbf{s}$ is. We examine the latter.

Claim 7 *The function $\mathbf{s}$ is strictly increasing when nonzero.*

Proof. Suppose to the contrary there is a strictly positive constant $\hat{s}$ and a nondegenerate interval (σ_1, σ_2) with $\mathbf{s}((\sigma_1, \sigma_2)) = \hat{s}$. Then because the market balances (Claim ^{Simon}4 which carries over to the limit $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$) and the buyer attribute-choice function is strictly increasing (Claim ^{Rhyme}5), the attribute choice $\hat{s}$ must be “matched” with buyer attributes from a set $[b_1, b_2]$ i.e. $b_1(\hat{s}) = \lim_{\beta \downarrow \beta_1(\hat{s})} \mathbf{b}(\hat{s}) < b_2(\hat{s}) = \lim_{\beta \uparrow \beta_2(\hat{s})} \mathbf{b}(\hat{s})$. Then from Claim ^{Symph 3.5}6, we have:

$$\begin{aligned} \lim_k \lim_n \int \int_{\sigma_1}^{\sigma_2} \bar{H}^k(\mathbf{s}^{n,k}(\sigma), \mathbf{b}^{n,k}, p_U^{n,k}) - c_S(\mathbf{s}^{n,k}(\sigma), \sigma) d\sigma d\xi^{n,k} \\ \leq \int_{\sigma_1}^{\sigma_2} h_S(b_1, \hat{s}) + p_U^*(\hat{s}) - c_S(\hat{s}, \sigma) d\sigma. \end{aligned}$$

Now let $\sigma(\eta) = \inf\{\sigma : \mathbf{s}(\sigma) \geq \hat{s} + \eta\}$. Notice that $\lim_{\eta \rightarrow 0} \sigma(\eta) = \sigma_2$. Now let the seller set attribute choice function

$$\mathbf{s}'(s) = \begin{cases} \mathbf{s}(\sigma) & \text{if } \sigma \notin (\sigma_1, \sigma(\eta)) \\ \hat{s} + \eta & \text{if } \sigma \in (\sigma_1, \sigma(\eta)). \end{cases}$$

Notice that $\mathbf{s}'$ is weakly increasing. A seller choosing attribute $\hat{s} + \eta$ sets price $\hat{p} > p_U(\hat{s})$ satisfying

$$\hat{p} = \sup\{p : B(\mathbf{s}^*, p, p_U^*) \neq \emptyset\}.$$

The price $\hat{p}$ is at least as high as that the value p' satisfying $h_B(b_2, \hat{s}) - p_U(\hat{s}) = h_B(b_2, \hat{s}_\eta) - p'$, where p' is calculated so that buyer b_2 will prefer attribute choice $\hat{s} + \eta$ at price $\hat{p}$ to the buyer's current choice of $\hat{s}$ at price $p_U^*(\hat{s})$. By setting the price $\hat{p}$ for attribute choice $\hat{s} + \eta$, the seller then ensures that attribute choice $\hat{s} + \eta$ is chosen by a buyer at least as high as b_2 (the single-crossing condition on buyer premoneration values ensures that no smaller buyers will choose $\hat{s} + \eta$). In addition, any buyer who currently chooses a seller $s > \hat{s} + \eta$ will continue to prefer that seller. As a result, the switch to attribute choice function $\mathbf{s}'$ increases the seller's payoff by at least

$$\begin{aligned} & \int_{\sigma_1}^{\sigma_2} (h_S(b_2, \hat{s} + \eta) + \hat{p}) d\sigma - \int_{\sigma_1}^{\sigma_2} (h_S(b_1, \hat{s}) + p_U^*(\hat{s})) d\sigma - \int_{\sigma_1}^{\sigma(\eta)} (c_S(\hat{s} + \eta, \sigma) - c_S(\mathbf{s}^*(\sigma), \sigma)) d\sigma \\ & > (\sigma_2 - \sigma_1)[h_S(b_2, \hat{s} + \eta) - h_S(b_1, \hat{s})] - (\sigma(\eta) - \sigma_1)[c_S(\hat{s} + \eta, \sigma_1) - c_S(\hat{s}, \sigma_1)]. \end{aligned}$$

The first term in the second line is bounded away from zero as η approaches zero, while the second approaches zero as does η , ensuring that there is some $\eta > 0$ for which the payoff differences is positive. Intuitively, each seller in the interval (σ_1, σ_2) experiences a discontinuous increase in expected buyer (at a higher price) when increasing her attribute choice, while sellers in the interval $\sigma_1, \sigma(\eta)$ experience a continuous increase in cost. The attribute choice function $\mathbf{s}'$ increases the seller's payoff for sufficiently small η , yielding the result. $\blacksquare$

With the various functions strictly increasing, it is now straightforward to show:

Symph 4

Claim 8

$$\begin{aligned} & \lim_k \lim_n \int \int \bar{H}^k(\mathbf{s}(\sigma), \mathbf{b}, p_U^{n,k}) - c_S(\mathbf{s}(\sigma), \sigma) d\sigma d\xi^{n,k} \\ & = \int h_S(\mathbf{b}^*(\sigma), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)) - c_S(\mathbf{s}^*(\sigma), \sigma) d\sigma \\ & = \int \max_s h_S(\mathbf{b}^*(\sigma), s) + p_U^*(s) - c_S(s, \sigma) d\sigma. \end{aligned}$$

A.5.6 Uniform Price Equilibria

We finally argue that the profile $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$ induces a uniform price equilibrium of the matching market with identical attribute choices (but perhaps a shift in the price function). First, we note that the equilibrium functions $\mathbf{b}^*$, $\tilde{s}^*$, and $\mathbf{s}^*$ are increasing. We can then take them to be continuous from the right, since doing so requires adjusting at most a countable set of values, which leaves expected payoffs unaffected. This in turn ensures that if almost all buyers and sellers have no profitable deviation, then (via a continuity argument) none do.

The first task is to show that equilibrium payoffs are nonnegative, so that agents would not prefer to be out of the market. Suppose $\{\xi_B^{n,k}, \xi_S^{n,k}, p_U^{n,k}\}_{n,k}$ is the double sequence that converges to $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$. We have

$$\begin{aligned} h_B(0, 0) - p_U^*(0) &= h_B(0, 0) - p_U^*(0) - c_B(0, \beta) \leq h_B(\mathbf{b}(\beta), \tilde{s}(\beta)) - p_U^*(\tilde{s}(\beta)) - c_B(\mathbf{b}(\beta), \tilde{s}(\beta)), \\ h_S(0, 0) + p_U^*(0) &= h_S(0, 0) + p_U^*(0) - c_S(0, \sigma) \leq h_S(\tilde{b}(s(\sigma)), \mathbf{s}(\sigma)) + p_U^*(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma) \end{aligned} \quad \begin{matrix} (29) \\ (30) \end{matrix}$$

Let

$$P = h_B(0, 0) - p_U^*(0) \geq -h_S(0, 0) - p_U^*(0)$$

and replace the price function p^* with $p^* + P$. Both $\xi_B^{n,k}$ and $\xi_S^{n,k}$ remain best responses given price $p^* + P$ and markets still clear in the limit of $n \rightarrow \infty$. Moreover, replacing p^* with $p^* + P$ in (29)–(30) gives

$$\begin{aligned} 0 &\leq h_B(\mathbf{b}(\beta), \tilde{s}(\beta)) - p_U^*(\tilde{s}(\beta)) - c_B(\mathbf{b}(\beta), \tilde{s}(\beta)) \\ 0 &\leq h_S(\tilde{b}(s(\sigma)), \mathbf{s}(\sigma)) + p_U^*(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma) \end{aligned}$$

and hence nonnegative payoffs.

It is immediate from the formulation of the buyer's payoffs in the game and from [Symph 4](#) that neither buyer nor seller has a profitable priced deviation for a positive measure set of types.

For unpriced deviations, consider first deviations to values of s' that are chosen by some other sellers. We first note that such a deviation can never entail a lower price than that which would make the deviation priced. The lower price will always suffice to attract the current match of seller s' , which combines with the pessimism built into the seller's evaluation of unpriced deviations to ensure that they are not optimal. For higher prices, we need only note that since s' is available at the going price, the deviating seller in question will attract no buyers.

Now consider deviations to values s that are chosen by no sellers. Then s is contained in an interval of unchosen attribute choices; let the closure of

the largest such interval be denoted by $[s', s'']$. Assume that attribute choices s' and s'' are purchased by buyers b' and b'' respectively. (The argument is easily modified to cover the case in which one of these buyers does not exist, i.e., in which s is either larger or smaller than every chose attribute.) If there is a profitable unpriced seller deviation to s , it must involve a lower price than $p_U(s)$ and attract buyer b'' . (It is clear that higher prices will attract no buyers. The price schedule p_U already offers the seller the chance to sell s at price $p_U(s)$ to a buyer at least as good as b' , so there is no hope in offering a lower price to attract b' .) Suppose there exists p such that

$$h_B(b'', s) - p \geq h_B(b'', s'') - p_U(s'')$$

and a seller who strictly prefers choosing s and p to the sellers' equilibrium outcome. Then there must exist a price p'' such that the seller strictly prefers s and p'' , and such that

$$h_B(b'', s) - p'' > h_B(b'', s'') - p_U(s'').$$

But then for sufficiently large k and then large n , the set $B(s, p'', p_U^{n,k})$ must include a positive measure of buyers with types above b'' , and an arbitrarily small measure of smaller buyers. The seller must then evaluate this payoff as if it could attract buyer b'' , which is a contradiction.

A.5.7 Nontriviality

We now show that under the conditions of the proposition, the profile $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$ is nontrivial. First, suppose there exists $(b, s) \in [0, \bar{b}] \times [0, \bar{s}]$ with

$$h_B(b, s) + h_S(0, s) - c_B(b, 1) - c_S(s, 1) > 0. \quad (31)$$

dugout

We suppose further that the equilibrium is trivial, so that $\mathbf{b}^*$ and $\mathbf{s}^*$ are identically zero, and derive a contradiction. Then there is no agent for whom it is profitable to trade at price p_U , so that for all $(b, s) \in [0, \bar{b}] \times [0, \bar{s}]$,

$$\begin{aligned} h_S(0, s) + p_U(s) - c_S(s, 1) &\leq 0 \\ h_B(b, s) - p_U(s) - c_B(b, 1) &\leq 0, \end{aligned}$$

where we focus on agents $\beta = 1 = \sigma$ since they are the most likely to want to trade. Notice that we are using here the maximum that appears in the building block (23) for the specification of the seller's payoff, and which effectively allows the seller to sell any attribute choice $s \in [0, \bar{s}]$ at price

$p_U(s)$, assuming in the process that he can attract at least a zero-attribute buyer. For these two inequalities to hold, it must be that

$$h_B(b, s) + h_S(0, s) \leq c_B(b, 1) + c_S(s, 1),$$

contradicting ^{dugout}(31).

Suppose that for all $\phi \in (0, 1]$, there exist attribute choices $b(\phi)$ and $s(\phi)$ such that

$$h_B(b(\phi), s(\phi)) + h_S(0, s(\phi)) - c_B(b(\phi), \phi) - c_S(s(\phi), \phi) > 0. \quad (32) \quad \boxed{\text{pancake}}$$

Suppose that not all agents participate in the market. Because equilibrium attribute choice functions are increasing when nonzero, and because the market for attributes clears in equilibrium, there must be intervals of types $[0, \beta']$ and $[0, \sigma']$ with $\beta' = \sigma'$ who do not participate.²⁰ Choose $\beta \in [0, \beta']$ and $\sigma = \beta \equiv (\phi)$. If neither agent is to participate, it must be that

$$\begin{aligned} h_S(0, s(\phi)) + p_U(s(\phi)) - c_S(s(\phi), \phi) &\leq 0 \\ h_B(b(\phi), s(\phi)) - p_U(s(\phi)) - c_B(b(\phi), \phi) &\leq 0. \end{aligned}$$

For these two inequalities to hold, it must be that

$$h_S(0, \mathbf{s}^*(\sigma)) + h_B(\mathbf{b}^*(\beta), \mathbf{s}^*(\sigma)) \leq c_B(\mathbf{b}^*(\beta), \beta) + c_S(\mathbf{s}^*(\sigma), \sigma),$$

contradicting ^{pancake}(32). ■

A.6 Equilibrium with Endogenous Monitoring

pend5

Given a set of informed sellers, I , and a pair of price functions (p_P, p_U) , a buyer β who makes attribute choice $b \in \mathcal{B}$ and a seller attribute choice $s \in \mathcal{S}$ has payoff (where we distinguish between informed and uninformed sellers)

$$\Pi_B(b, s, \beta) \equiv \begin{cases} h_B(b, s) - p_P(b, s) - c_B(b, \beta), & \text{if } s \in \mathcal{S}(I) \text{ and} \\ & \text{the seller is informed,} \\ h_B(b, s) - p_U(s) - c_B(b, \beta), & \text{if } s \in \mathcal{S}(U) \text{ and} \\ & \text{the seller is uninformed.} \end{cases}$$

²⁰This observation is important to the proof, as it allows us to conclude not just that some sellers and some buyers do not participate, but that there is a buyer/seller pair who do not participate and who can produce a positive surplus when matched.

The first condition is that there be no profitable priced buyer deviations, i.e., $\forall \beta \in [0, 1]$, if $b(\beta) \in \tilde{b}_i(\mathcal{S}(I))$, then

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}_i(\mathbf{b}(\beta)), \beta) = \sup_{(b,s) \in \mathcal{B} \times \mathcal{S}} \Pi_B(b, s, \beta) \quad (33) \quad \boxed{\text{buyer IC gen 1}}$$

and if $b(\beta) \in \tilde{b}_u(\mathcal{S}(U))$, then

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}_u(\mathbf{b}(\beta)), \beta) = \sup_{(b,s) \in \mathcal{B} \times \mathcal{S}} \Pi_B(b, s, \beta). \quad (34) \quad \boxed{\text{buyer IC gen 2}}$$

Note that if $b(\beta) \in \tilde{b}_i(\mathcal{S}(I)) \cap \tilde{b}_u(\mathcal{S}(U))$, buyer β is indifferent between the informed seller of attribute choice $\tilde{s}_i(b(\beta))$ and the uninformed seller of attribute choice $\tilde{s}_u(b(\beta))$.

We turn now to the priced deviations by the seller. In keeping with our discussion above on the matching, a choice of $s \in \mathcal{S}(I)$ is then the joint decision to become an informed seller and a choice of s , while a choice of $s \in \mathcal{S}(U)$ is the joint decision not to become an informed seller and a choice of s . Given a matching $\tilde{b}$ and a pair of price functions (p_P, p_U) , a seller who chooses an attribute choice $s \in \mathcal{S}$ receives a payoff of

$$\Pi_S(s, \sigma) \equiv \begin{cases} h_S(\tilde{b}_i(s), s) + p_P(\tilde{b}_i(s), s) - c_S(s, \sigma) - K(\sigma, \mathcal{A}), & \text{if } s \in \mathcal{S}(\mathcal{I}) \text{ and } \\ & \sigma \text{ is informed,} \\ h_B(\tilde{b}_u(s), s) - p_U(s) - c_B(b, \beta), & \text{if } s \in \mathcal{S}(\mathcal{U}) \text{ and } \\ & \sigma \text{ is uninformed.} \end{cases}$$

The condition that there be no profitable priced seller deviations has two parts,

$$\Pi_S(\mathbf{s}(\sigma), \sigma) = \sup_{s \in \mathcal{S}} \Pi_S(s, \sigma), \quad \forall \sigma \in [0, 1], \quad (35) \quad \boxed{\text{seller IC gen 1}}$$

and

$$\sigma \in \mathcal{I} \iff \mathbf{s}(\sigma) \in \mathcal{S}(\mathcal{I}). \quad (36) \quad \boxed{\text{seller IC gen 2}}$$

Since now sellers choose whether to become informed, a profitable unpriced deviation for a seller can take two forms, one in which he chooses to become informed and the other in which he chooses to remain uninformed. This leads us to the following definition, which should be compared with Definitions ?? and ??:

Definition 9 *A seller σ has a profitable unpriced deviation as an informed seller if there exists an attribute choice s' , an available buyer attribute choice $b \in \mathcal{B}$ and a price $p \in \mathbb{R}$, with either $s' \notin \mathcal{S}(I)$ or $p \neq p_P(b, s')$, such that*

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}_i(\mathbf{b}(\beta)), \beta) < h_B(b, s') - p - c_B(b, \beta)$$

and

$$\Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b', s') + p - c_S(s', \sigma) - K(\sigma, \mathcal{A}).$$

A seller σ has a profitable unpriced deviation as an uninformed seller if there exists s' and a price $p \in \mathbb{R}$, with either $s' \notin \mathcal{S}(U)$ or $p \neq p_U(s')$, such that there exists $\mathbf{b}(\beta)$ with

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}_i(\mathbf{b}(\beta)), \beta) < h_B(\mathbf{b}(\beta), s') - p - c_B(\mathbf{b}(\beta), \beta)$$

and for all β' ,

$$\text{if } \Pi_{B'}(\mathbf{b}(\beta'), \tilde{s}_i(\mathbf{b}(\beta')), \beta') < h_B(\mathbf{b}(\beta'), s') - p - c_B(\mathbf{b}(\beta'), \beta'),$$

$$\text{then } \Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(\mathbf{b}(\beta'), s') + p - c_S(s', \sigma).$$

A seller σ has a profitable unpriced deviation if he has a profitable deviation as either an informed or as an uninformed seller.

The first part of the definition corresponds to the case in which the seller chooses an attribute and incurs the cost of becoming informed, K , and hence can target the buyer with whom he transacts. The second part corresponds to the seller choosing not to become informed, and only considers a deviation profitable when, as in the definition of uniform price equilibria, he benefits from transacting with all buyers who would be attracted to this investment-transfer proposal.

The only potential profitable deviations for buyers are deviations with an informed seller as in personalized price equilibrium:

Definition 10 Buyer β has a profitable unpriced deviation (with an informed seller) if there exists an attribute choice $b' \notin \mathcal{B}$, a price $p \in \mathbb{R}$, and $s' \in \mathcal{S}(I)$ with

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}_i(\mathbf{b}(\beta)), \beta) < h_B(b', s') - p - c_B(b', \beta)$$

and

$$h_S(\tilde{b}_i(s'), s') - p_P(\tilde{b}_i(s'), s') < h_S(b', s') + p.$$

We now define an endogenous price equilibrium.

Definition 11 An endogenous-price equilibrium is a specification satisfying buyer IC gen 1, seller IC gen 1, buyer IC gen 2, seller IC gen 2, (33), (34), (35), and (36), and such that no seller or buyer has a profitable unpriced deviation.

References

- Billingsley1986** BILLINGSLEY, P. (1986): *Probability and Measure*. John Wiley and Sons, New York.
- BandL03** BULOW, J., AND J. LEVIN (2004): “Matching and Price Competition,” Mimeo, Stanford University.
- Coase60** COASE, R. (1960): “The Problem of Social Cost,” *Journal of Law and Economics*, 2, 1–40.
- CMP01** COLE, H. L., G. J. MAILATH, AND A. POSTLEWAITE (2001): “Efficient Non-Contractible Investments in Large Economies,” *Journal of Economic Theory*, 101, 333–373.
- DandS88** DUNFORD, N., AND J. T. SCHWARTZ (1988): *Linear Operators Part I: General Theory*. John Wiley and Sons, New York, Wiley Classics Library Edition.
- Hopkins05** HOPKINS, E. (2005): “Job Market Signaling of Relative Position, or Becker Married to Spence,” Discussion paper, University of Edinburgh.
- HMS05** HOPPE, H. C., B. MOLDOVANU, AND A. SELA (2005): “The Theory of Assortative Matching Based on Costly Signals,” Mimeo, University of Bonn.
- MilgromSegal02** MILGROM, P., AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70(2), 583–601.
- Milgrom81** MILGROM, P. R. (1981): “An Axiomatic Characterization of Common Knowledge,” *Econometrica*, 49, 215–218.
- Nicholson03** NICHOLSON, S. (2003): “Barriers to Entering Medical Specialties,” Working Paper 9649, National Bureau of Economic Research.
- Royden88** ROYDEN, H. L. (1988): *Real Analysis*. Prentice Hall, Englewood Cliffs, NJ, 3rd edn.