

Pricing in Matching Markets*

George J. Mailath
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104
gmailath@econ.upenn.edu

Andrew Postlewaite
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104
apostlew@econ.upenn.edu

Larry Samuelson
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706
larrysam@ssc.wisc.edu

March 8, 2006

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 ownership of the default shares—the shares of the surplus owned by the buyer and seller in the absence of transfers—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-0095768, SES-9911219, and SES-0241506) for financial support.

Pricing in Matching Markets

1 Introduction

Prices. Consider three people and the prices that clear the markets for their transactions.¹

- Allen participates in a daily spot market for casual, unskilled labor. A prototypical job would be one in which he is paid a fixed sum to drive a truck to pick up materials at Home Depot. His qualification for this market is that he has a drivers license. The market in which he sells his labor 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 it rejected other students who would have full tuition to attend.
- 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 Allen as *universal* prices. These are the prices that typically show up as heavy horizontal lines on supply-and-demand diagrams in introductory texts. Carol faces *personalized* prices that depend on both her characteristics and those of the firms offering a position.

We say that 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? How can we use these links to design effective markets? This paper addresses these questions, concentrating on uniform and personalized prices.

¹We put aside a variety of reasons why the markets may not clear—perhaps search costs or market power—to concentrate on differences in how prices clear markets.

Default divisions. A match between a buyer and seller creates a surplus. In the absence of any transfers, part of this surplus accrues to each side—the *default share* for that side (possibly negative). Understanding the nature of the surplus and the default division is the key to understanding differences in pricing across markets.

The surplus in Allen’s markets is reasonably modeled as the sum of two terms, a (negative) default share accruing to Allen and depending only on the value of foregone leisure, and a default share accruing to his trading partner that depends on the latter’s characteristics, the value of getting the materials delivered. This separability ensures that there is no issue of *efficient* matching in these markets. It matters that the right people (i.e., high-valuation) trade, but matters not whom they trade with. It is then no surprise that a universal price clears the market.

The surplus in Carol’s markets depends in a complementary fashion upon the agents on both sides of the market. Talented executives are likely to be much 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 transact, but also being sure that they transact with the right partners. We might expect that we 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 default shares in the market. Both Carol and her alma mater receive 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 prestige should she become a Supreme Court Justice, and the increment to its endowment if she donates to the university’s endowment in the future out of gratitude. 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 there before starting her own company. In contrast, Bob’s junior college anticipates no benefit from the transaction with 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.² Prices modify the default distribution of the subsequent surplus and determine investment incentives.

When prices can be personalized, there will exist 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, equilibria still exist. However, *efficient* equilibria exist in this case if and only if the default shares on the side of the market setting prices are independent of the attributes of the agents on the other side of the market—for example, if the side of the market setting prices owns none of the surplus (before transfers), as is the case in Bob’s (but not Carol’s) market.

This 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 if the price-setter’s *default share* exhibits no complementarities. For example, while Bob owns none of the surplus generated by his tax preparations, he may be one of many tax preparers of widely differing abilities. Efficiency likely requires matching highly able tax preparers with clients who have complex finances, so that Bob’s problem is characterized by complementarities in attributes but no complementarities in the buyers’ attributes and the price-setter’s default share. Prices transform default shares into final allocations, and the prices required to support an allocation depend upon the properties of both.³

Why are prices important? Prices guide investments. Workers and firms are matched in markets that take place *after* the parties have made investments that influence their default shares of the surplus. The transac-

²We thus focus on cases with complementarities and the contrast between uniform and personalized pricing. The surplus may exhibit no complementarities without the default share of each agent depending on only that agent’s attribute choice, a possibility not covered in our three opening scenarios but reflected in the fact that hamburgers and pizzas can sell at different prices.

³The nature of a market’s default shares 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.

tion price determines the return to each sides from the match, and hence the value of the pre-match investments. In a world devoid of frictions, personalized prices whereby each seller sets a different price for each potential buyer would be the norm. We will show that in this case, both parties have correct incentives, and invest efficiently. Personalized prices sometimes happen also to be uniform prices, that is, the sellers might charge different prices, but each seller set a single price regardless of the buyer. However, this coincidence would be no more than a curiosity, and there would be little reason to be concerned with how default shares are defined in this case.⁴ The world is not frictionless though. 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 2004.⁵ The going price for identifying whether a high school diploma comes from a legitimate high school is \$100.⁶ There may be substantial savings from posting uniform prices and letting buyers sort themselves (as Bob's clients do), if the default shares are so defined that uniform prices can do this sorting. Alternatively, if the default shares 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.⁷

⁴We thus have no disagreement with Coase (1960)'s observation that property rights would be irrelevant in world without transactions costs.

⁵Expenditures 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.)

⁶"Vetting Those Foreign College Applications," *New York Times*, September 29, 2004, page A21.

⁷For example, Bulow and Levin 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 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

We thus view models of personalized and uniform prices as building blocks for an analysis in which firms first decide whether to acquire the costly technology necessary for personalized pricing. This acquisition decision depends importantly how effectively uniform or personalized prices create subsequent investment and matching incentives. Uniform prices are important because frictions may preclude personalization.

Designing markets. The specification of the default shares 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 surplus in the form of marketable patents from faculty research. Historically, universities have owned these patents, but many universities have unilaterally assigned to professors shares in the revenues from patents stemming from their research. In a similar vein, one could define the default share 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 would then own the future income of students to whom it gives degrees. While it is farfetched to imagine a university owning the entire income stream of a student, a number of countries (including Australia, Sweden and New Zealand) provide income-contingent loans to students that effectively give the lender a share of students' future income (?).⁸

Our results suggest that the appropriate design of the default shares in a market can be valuable. In the absence of constraints, appropriately defined default shares can allow efficient equilibria to be supported by uniform prices, avoiding, or at least ameliorating, the costs of personalized pricing or the costs of inefficient uniform pricing.

Unfortunately, there are often constraints on the design of default shares. 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

(as compared to the \$34,000 stipend they currently receive).

⁸Basketball 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, [http://www.detnews.com/2002/pistons/0204/27/sports-475199.htm/.](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, some of which is assigned to the team.

to measure and collect the increment to income attributable to the university education?⁹ 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 default shares. In the university case, we might instead assign all of the surplus to the student, but again we encounter both measurement and moral hazard issues. 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 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 $[0, 1]$.

The agents make choices in two stages. First, buyers and sellers simultaneously choose attributes. We denote the cost of attribute choice $b \in \mathfrak{R}_+$ to buyer β by $c_B(b, \beta)$, and the cost of attribute choice $s \in \mathfrak{R}_+$ to seller σ by $c_S(s, \sigma)$. Upon entering the market, buyers will be able to observe sellers' attribute choices. However, sellers will be unable to observe buyers' attribute choices unless they have purchased a monitoring technology. A seller who purchases the monitoring technology can determine the attribute choice of *any* buyer. The specification of the cost of this technology plays an important role in the analysis that we discuss in Section 5. Sellers choose whether to be *informed* (i.e., purchase the monitoring technology) at the same time that they choose their attributes.

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 default value* $h_B(b, s)$ and a *seller default value* $h_S(b, s)$.

⁹Measurement 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, or even sung by the Girl Scouts, but collection is impractical. In other cases, legal restrictions may be paramount.

Matching is mediated through prices posted by sellers. These prices may be either positive or negative. Indeed, the distinguishing feature of a seller in this paper is (hereafter) that the seller posts the price that determines the transfer (and not that positive transfers are made to the seller). For example, the illustrations in Section 1 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*.

The price a seller posts is necessarily independent of buyer attribute choices when the seller has not purchased the monitoring technology. We say that prices are *uniform* in this case. If a seller has purchased the monitoring technology, he can post *personalized* prices, that is prices that depend on a buyer's attribute choice. Once prices are posted, each buyer selects a seller. If more than one buyer has selected the same seller, that seller receives a random draw from the set of relevant buyers.

Assumption 1

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

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

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

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

the cost function c_S satisfies analogous conditions.

3. There exists \bar{b} such that for no $b > \bar{b}$, $s \in \mathfrak{R}_+$, $\beta \in [0, 1]$ and $\sigma \in [0, 1]$ is it the case that

$$v(b, s) - c_B(b, \beta) - c_S(s, \sigma) > 0;$$

and there exists an analogous \bar{s} .

4. The default values $h_B : \mathfrak{R}_+ \times \mathfrak{R}_+ \rightarrow \mathfrak{R}$ and $h_S : \mathfrak{R}_+ \times \mathfrak{R}_+ \rightarrow \mathfrak{R}$ are \mathcal{C}^2 and increasing and satisfy

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

and $h_B(b, s)$ is Lipschitz continuous in s on $[0, \bar{b}] \times [0, s^\dagger]$ for some $s^\dagger > \bar{s}$, i.e., there exists Δ such that for all $b \in [0, \bar{b}]$, $\epsilon > 0$ and $s \in [0, s^\dagger - \epsilon]$, $h_B(b, s + \epsilon) - h_B(b, s) < \Delta$.

Given the supermodularity of the surplus function, the restrictions on default values are satisfied by any default division that assigns positive constant proportions of the surplus to the two agents, as well as any division that assigns all of the surplus to the buyer (but not one that assigns all of the surplus to the seller, since then we cannot have $d^2 h(b, s)/dbds > 0$). Remark 3 below explains why we must exclude the latter case. Intuitively, an agent who owns the entire surplus is not an effective choice to post uniform prices. Similarly, the imposition of a Lipschitz requirement only on the buyer's default share reflects asymmetries that arise when sellers post uniform prices.

We can restrict attention to values $b \in [0, \bar{b}]$ and $s \in [0, \bar{s}]$, with Assumption 1.3 ensuring that values outside these intervals will never be optimal. 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, and 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 λ_B and λ_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 $\mathcal{S}' \subset \mathcal{S}$,

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

where λ is Lebesgue measure.

Definition 1 A feasible matching is a bijection $\tilde{b} : \mathcal{S} \rightarrow \mathcal{B}$ that is measure-preserving, i.e., $\lambda_B(\tilde{b}(\mathcal{S}')) = \lambda_S(\mathcal{S}')$ for all Borel $\mathcal{S}' \subset \mathcal{S}$.

The measure-preserving requirement on \tilde{b} ensure 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 , and \tilde{s} denotes its inverse

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 In our model, the surplus generated by a match 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 what an applicant’s accomplishments (chosen attribute), 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 an agent’s attribute choice and his type, at which point our analysis of personalized pricing 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, then attribute choices take on a dual role, directly enhancing the value of a match while also providing signals of types. ? examines such a model. ■

3 The Personalized Pricing Model

In this section, we assume sellers are automatically and costlessly informed and consequently, 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 $p_P(b, s)$.

3.1 Personalized-Price Equilibrium: Formulation

A personalized-price *allocation* is a pair of attribute choice functions, \mathbf{b} and \mathbf{s} , a feasible matching function \tilde{b} , and a personalized price function, $p_P : \mathcal{B} \times \mathcal{S} \rightarrow \mathfrak{R}$. Intuitively, a personalized-price *equilibrium* is a personalized-price allocation from which no agent has an incentive to deviate from the specified behavior. This section makes this precise.

A personalized price function p_P identifies the transfer of a buyer who makes attribute choice $b \in \mathcal{B}$ and matches with a seller attribute choice $s \in \mathcal{S}$, but is silent if either $b \notin \mathcal{B}$ or $s \notin \mathcal{S}$. We accordingly say that a

deviation by a buyer who makes attribute choice $b \in \mathcal{B}$ and matches with a seller with attribute choice $s \in \mathcal{S}$ is *priced*, while a deviation involving attribute choice $b \notin \mathcal{B}$ is unpriced. Similarly, a seller deviation involving an attribute choice $s \in \mathcal{S}$ and price $p_P(\tilde{b}(s), s)$ is said to be priced. A seller deviation to $s \notin \mathcal{S}$ or to an attribute choice $s \in \mathcal{S}$ coupled with a price other than $p_P(\tilde{b}(s), s)$ is said to be unpriced.

3.1.1 Priced Deviations

The payoff to a buyer of type β who chooses attribute $b \in \mathcal{B}$ and seller s , and to a seller of type $\sigma \in \mathcal{S}$ who chooses attribute s , are

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

The first equilibrium condition is that no priced deviation gives the buyer a payoff higher than his equilibrium payoff:

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}(\mathbf{b}(\beta)), \beta) = \sup_{(b,s) \in \mathcal{B} \times \mathcal{S}} \Pi_B(b, s, \beta), \quad \forall \beta \in [0, 1]. \quad (1)$$

Similarly, there must be no profitable priced deviation for the seller:

$$\Pi_S(\mathbf{s}(\sigma), \sigma) = \sup_{s \in \mathcal{S}} \Pi_S(s, \sigma), \quad \forall \sigma \in [0, 1]. \quad (2)$$

3.1.2 Unpriced Deviations

An unpriced deviation requires the seller to propose a match and a price.

Definition 2 *Seller σ has a profitable unpriced deviation if there exists a seller attribute choice s , a buyer attribute choice $b \in \mathcal{B}$, and a price $p \in \mathfrak{R}$ with either $s \notin \mathcal{S}$ or $p \neq p_P(b, s)$, such that*

$$h_B(b, \tilde{s}(b) - p_P(b, \tilde{s}(b))) < h_B(b, s) - p \quad (3)$$

and

$$\Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b, s) + p - c_S(s, \sigma). \quad (4)$$

Hence, the seller must be able to attract a buyer (condition (3)) while making a profit (condition (4)).

There are three types of unpriced seller deviations. First, we may have $s' \notin \mathcal{S}$: the seller chooses an attribute currently not offered in the market

and names the price at which that attribute is to be sold. Second, seller σ could retain the attribute choice $s = \mathbf{s}(\sigma)$ but offer a price $p \neq p_P(\tilde{b}(s), s)$ in order to attract a buyer $b \neq \tilde{b}(s)$. In this case, the cost of attribute choice falls out of (4). Adding the two inequalities, the requirement that there is no p and $b \in \mathcal{B}$ for which (3)–(4) hold is then equivalent to the requirement that the matching \tilde{b} and resulting division of the surplus $v(b, s)$ be *stable*,¹⁰ in the sense that, for all $(b, s) \in \mathcal{B} \times \mathcal{S}$,

$$\begin{aligned} h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) + h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) &\geq h_B(b, s) + h_S(b, s) \\ &= v(b, s). \end{aligned} \quad (5)$$

Finally, a seller σ could choose an attribute $s \in \mathcal{S}$ different than $\mathbf{s}(\sigma)$ and offer a different price than $p_P(\tilde{b}(s), s)$. However, the unprofitability of the priced deviation to s and $p_P(\tilde{b}(s), s)$ (from (2)) and stability imply that such a deviation cannot be profitable.¹¹

An unpriced deviation for a buyer requires an attribute choice $b \in \beta$, for which he will need to propose a price for the transaction, leading to:

Definition 3 *Buyer β has a profitable unpriced deviation if there exists an attribute choice $b' \notin \mathcal{B}$, a price $p \in \mathfrak{R}$, and $s' \in \mathcal{S}$ with*

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

and

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

In other words, the buyer could choose an investment, a transfer and a target seller such that in the resulting transaction, both the buyer and the seller would be better off than by following the behavior prescribed by the proposed equilibrium.

¹⁰This is the standard notion of stability for cooperative games of matching (see Roth and Sotomayer (1990)).

¹¹Suppose not. Then, there exists a seller σ , investments $(b, s) \in \mathcal{B} \times \mathcal{S}$ with $s \neq \mathbf{s}(\sigma)$ and a price p such that (3)–(4) hold. Adding these two inequalities

$$h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) + h_S(\tilde{b}(\mathbf{s}(\sigma)), \mathbf{s}(\sigma)) + p_P(\tilde{b}(\mathbf{s}(\sigma)), \mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma) < h_S(b, s) + h_B(b, s) - c_S(s, \sigma).$$

Applying (2), we have

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

contradicting (5).

3.1.3 Equilibrium

Definition 4 A personalized price equilibrium is a personalized price allocation $(\mathbf{b}, \mathbf{s}, \tilde{b}, p_P)$ satisfying (1) and (2) (i.e., excluding profitable priced deviations) for which no seller or buyer has a profitable unpriced deviation.

Remark 2 The definition appears to treat buyers and sellers asymmetrically (see, in particular, (1) and (2)), but this is illusory. Since information is symmetric, the specification that buyers choose sellers is simply a normalization. The seller effectively has the same option, since he can choose his attribute and price in an attempt to induce a matching with any target buyer. ■

3.2 Personalized-Price Equilibrium: Characterization

The default shares h_B and h_S play no role in the characterization of a personalized price equilibrium or in its existence. Intuitively, personalized prices allow the market to compensate for any idiosyncrasies of the default shares.

Our first result is that a personalized-price equilibrium is “well behaved,” in the sense that more productive agents choose larger attribute choices and match with more productive partners. Let $\tilde{B} : [0, 1] \rightarrow [0, 1]$ identify the type of buyer β matched with seller of type σ . The proof of the following (contained in Section 7) is a straightforward application of the facts that higher types find larger attribute choices more attractive (Assumption 1.2) and that stability requires assortative matching when the surplus function is supermodular (Assumption 1.1).

Lemma 1 Let $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ be a personalized-price equilibrium. Then \mathbf{b} and \mathbf{s} are weakly increasing, \tilde{b} is weakly increasing, and \tilde{B} can be taken to be the identity.

There may be many personalized prices that support the choices in a given personalized price equilibrium. In particular, most of the pairs $(b, s) \in \mathcal{B} \times \mathcal{S}$ correspond to matches that do not form in equilibrium, and there may be a range of prices for which such transactions are unprofitable. We will focus on the prices described in the following lemma.

Lemma 2 Suppose $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ is a personalized price equilibrium. Then, $(\mathbf{b}, \mathbf{s}, \hat{p}_P, \tilde{b})$ is also a personalized price equilibrium, where $\hat{p}_P(b, s)$ is the

reservation price for a seller with attribute choice s to match with b ,

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

Proof. The construction of $\hat{p}_P(b, s)$ as the reservation price of seller attribute choice s for buyer b ensures that sellers have no profitable deviations, priced or unpriced, under price \hat{p}_P . If, under \hat{p}_P , buyer b now strictly prefers to buy seller attribute choice $s' \neq \tilde{s}(b)$, then a price of $\hat{p}_P(b, s') + \varepsilon$ for the pair (b, s') vitiates the stability of $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$. ■

We can use the fact that the equilibrium type matching \tilde{B} will be the identity to 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 equilibria are *constrained efficient* in the sense that no matched pair of agents can increase their net surplus without both agents deviating to attribute choices outside the sets \mathcal{B} and \mathcal{S} . This is shown in the following lemma, proven (in Section 7) 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.

Lemma 3 *Suppose $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ is a personalized price equilibrium. Then, for all $\phi \in [0, 1]$, $b \in \mathcal{B}$, $s \in \mathcal{S}$ and all $b' \in [0, \tilde{b}]$, $s' \in [0, \tilde{s}]$,*

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

This result does not ensure that a personalized-price equilibrium 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.

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

Definition 5 *An allocation $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ is a complete personalized-price*

equilibrium if $p_P : [0, \bar{b}] \times [0, \bar{s}] \rightarrow \Re$ and for all β and all σ ,

$$\begin{aligned}\Pi_B(\mathbf{b}(\beta), \tilde{\mathbf{s}}(\mathbf{b}(\beta)), \beta) &= \sup_{(b,s) \in [0, \bar{b}] \times [0, \bar{s}]} \Pi_B(b, s, \beta) \quad \forall \beta \in [0, 1] \\ \Pi_S(\mathbf{s}(\sigma), \sigma) &= \sup_{(b,s) \in [0, \bar{b}] \times [0, \bar{s}]} \Pi_S(s, \sigma) \quad \forall \sigma \in [0, 1] \\ h_B(b, \tilde{\mathbf{s}}(b)) - p_P(b, \tilde{\mathbf{s}}(b)) + h_S(\tilde{\mathbf{b}}(s), s) + p_P(\tilde{\mathbf{b}}(s), s) &\geq h_B(b, s) + h_S(b, s).\end{aligned}$$

The first two conditions again stipulate that there be no profitable priced deviations. Because p_P specifies prices for every possible pair of characteristics, the only unpriced deviations are those in which seller σ chooses $s \in [0, \bar{s}]$ (possibly equal to $\mathbf{s}(\sigma)$, but possibly not) and chooses a price other than that specified by $p_P(b, s)$. The third (stability) condition ensures that these deviations are not profitable.

Lemma 4 *A complete personalized-price equilibrium is efficient.*

Proof. Suppose $(\mathbf{b}, \mathbf{s}, p_P, \tilde{\mathbf{b}})$ 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

$$v(b, s) - c_B(b, \phi) - c_S(s, \phi) > v(\mathbf{b}(\phi), \mathbf{s}(\phi)) - c_B(\mathbf{b}(\phi), \phi) - c_S(\mathbf{s}(\phi), \phi). \quad (6)$$

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

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

Adding these two contradicts (6). ■

We say that a personalized price allocation $(\mathbf{b}, \mathbf{s}, p_P, \tilde{\mathbf{b}})$ can be supported by a *uniform rationing price* if

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

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

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

Proof. Let $(\mathbf{b}, \mathbf{s}, p_U, \tilde{b})$ be a personalized price equilibrium 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 $\mathbf{b}(\beta)$ to match with a seller with $s < \tilde{s}(\mathbf{b}(\beta))$ (and hence $\tilde{b}(s) < \mathbf{b}(\beta)$). 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. ■

3.3 Example: Efficient Equilibrium

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 default values assign a fixed share of the surplus, $\theta \in [0, 1]$, to the buyers:

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

Efficiency requires that for each value of β and σ , attribute choices b and s must 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 (by symmetry) $s = b$, we have the solution

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

This characterizes an efficient allocation.

We next show that

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

is a personalized-price function corresponding to an efficient equilibrium for our example.¹² Given this price function, the seller attribute choice optimally chosen by a buyer of attribute choice b is

$$\theta bs - p_P(b, s) = \theta bs - \left\{ \frac{s^2}{2} - (1 - \theta)bs \right\} = bs - \frac{s^2}{2},$$

giving $\tilde{s}(b) = b$. Consequently, buyer β makes attribute choice b to maximize

$$\theta b \tilde{s}(b) - p_P(b, \tilde{s}(b)) - c_B(b, \beta) = \frac{b^2}{2} - \frac{b^3}{3\beta} \Rightarrow \mathbf{b}(\beta) = \beta.$$

Given the price, seller σ makes attribute choice s to maximize

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

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 share of the surplus that goes to the buyer, θ . This is not surprising, since the markets are complete and the efficient attribute choices are independent of the distribution of the surplus. The agents' equilibrium utilities are also independent of the default division of surplus. It does not matter who "owns" the technology that combines buyer and seller attribute choices into surplus when there is a competitive market for the attributes.

¹²Note that for any seller investment s , the price that a seller would get in a match with a buyer with investment b is decreasing in b . It is not surprising that this might occur since the seller is getting $(1 - \theta)$ of the surplus, which is increasing in the buyer's investment. In fact, this is the pricing function, identified in Lemma 2, that makes sellers indifferent over the buyer with whom they are matched, since

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

3.4 Personalized-Price Equilibrium: Existence

One route to existence is to note that a personalized price equilibrium is equivalent to Cole, Mailath, and Postlewaite (2001)'s *ex post contracting equilibrium*, and then to note that Cole, Mailath, and Postlewaite (2001) establish conditions for the existence of an *ex post contracting equilibria* exist. We take an alternative route here, building on the relationship between personalized-price and uniform-price equilibria (and hence postpone the proof until after examining uniform-price equilibria).

Proposition 1 *There exists an efficient personalized-price equilibrium.*

Proof. See appendix. ■

4 The Uniform Price Model

We now consider the case where sellers are constrained *not* to 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 $p_U : \mathcal{S} \rightarrow \mathfrak{R}$. In contrast to personalized pricing, default buyer and seller values play an important role in the determination of attribute choices under uniform pricing.

4.1 Uniform-Price Equilibrium: Formulation

A uniform-price allocation consists of a pair of attribute choice functions \mathbf{b} and \mathbf{s} , a feasible matching function \tilde{b} , and a price function, $p_U : \mathcal{S} \rightarrow \mathfrak{R}$. A uniform-price equilibrium is a uniform-price allocation from which no agent has an incentive to deviate. As with personalized pricing, we analyze these incentives in several steps.

We begin with the analogs of (1)–(2) for priced deviations. A buyer β who makes attribute choice $b \in \mathcal{B}$ and pays $p_U(s)$ to match with a seller with attribute choice $s \in \mathcal{S}$, and a seller σ who chooses attribute $s \in \mathcal{S}$ to match with buyer $\tilde{b}(s)$ at price $p_U(s)$, receive respective payoffs of

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

The requirement that there be no profitable priced deviations is then

$$\Pi_B(\mathbf{b}(\beta), \tilde{\mathbf{s}}(\mathbf{b}(\beta)), \beta) = \sup_{(b,s) \in \mathfrak{R}_+ \times \mathcal{S}} \Pi_B(b, s, \beta), \quad \forall \beta \in [0, 1] \quad (7)$$

$$\Pi_S(\mathbf{s}(\sigma), \sigma) = \sup_{s \in \mathcal{S}} \Pi_S(s, \sigma), \quad \forall \sigma \in [0, 1]. \quad (8)$$

The definitions of profitable unpriced seller deviations must incorporate the sellers' lack of information about buyer attribute choices. Consider a seller deviation to an attribute choice $s \notin \mathcal{S}$ or to a price inconsistent with p_U . In evaluating the profitability of a deviation, the seller must have a conjecture as to the attribute choice of the buyer that the deviation would attract. Assuming buyers are rational, the seller can rule out any buyer attribute choice b' for which the implied payoff is less than the buyer's current gross payoff, $h_B(b', \tilde{\mathbf{s}}(b')) - p_U(\tilde{\mathbf{s}}(b'))$. Since the seller must induce a deviation by the buyer, we assume that the seller rules out any buyer-attribute choice indifferent between the implied payoff and the putative equilibrium gross payoff. Since we are interested in the scope of behavior with uniform pricing, we impose the most permissive equilibrium conditions consistent with rationality: sellers only depart from the proposed equilibrium behavior if they *always* benefit from transacting with a buyer who is attracted by their deviation¹³ (and, of course, they must attract at least one buyer):

Definition 6 *An uninformed seller σ has a profitable unpriced deviation if there exists s' and a price $p \in \mathfrak{R}$, with either $s' \notin \mathcal{S}$ or $p \neq p_U(s')$, such that there exists $b' \in \mathcal{B}$ with*

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

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

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

Since the buyer attribute choice is not observed by the seller, there are no unpriced deviations by buyers.

Definition 7 *A uniform-price equilibrium is an allocation $\{\mathbf{b}, \mathbf{s}, p_U, \tilde{\mathbf{b}}\}$, satisfying (7) and (8), and such that no seller has a profitable unpriced deviation. A complete uniform-price equilibrium is a uniform-price equilibrium for which $p_U : [0, \bar{s}] \rightarrow \mathfrak{R}$.*

¹³We discuss this definition in the last section.

4.2 Uniform-Price Equilibrium: Characterization

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

$$0 = \frac{dh_B(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \quad (9)$$

$$0 = \frac{dh_B(b, s)}{ds} - \frac{dp_U(s)}{ds}, \quad (10)$$

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

$$0 = \frac{dh_S(\tilde{b}(s), s)}{db} \frac{d\tilde{b}(s)}{ds} + \frac{dh_S(\tilde{b}(s), s)}{ds} + \frac{dp_U(s)}{ds} - \frac{dc_S(s, \sigma)}{ds}. \quad (11)$$

Using (10) to eliminate $dp_U(s)/ds$ in (11) and then using the identity $v(b, s) = h_B(b, s) + h_S(b, s)$ in (9) and (11), these three first-order conditions can be reduced to

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

>From Lemma 3, establishing the constrained efficiency of a personalized-price equilibrium, we know that a personalized-price equilibrium must be characterized by the first-order conditions:

$$\begin{aligned} 0 &= \frac{dv(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \\ 0 &= \frac{dv(b, s)}{ds} - \frac{dc_S(s, \sigma)}{ds}. \end{aligned} \quad (12)$$

Comparing these, it is immediate that the solution to the first order conditions for the personalized price equilibrium maximization problems will be a solution for the first order conditions for the uniform price equilibrium problem if $\frac{dh_S(b, s)}{db} = 0$, that is, if each seller's default share 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.

Proposition 2 *A personalized price equilibrium allocation can be achieved in a uniform price equilibrium if the sellers' default shares do not depend on the level of the buyer's investment.*

Proof. Let $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ be a personalized price equilibrium. From Lemma 2, $(\mathbf{b}, \mathbf{s}, \hat{p}_P, \tilde{b})$ is a personalized-price equilibrium, with

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

If $h_S(b, s)$ does not depend on b , then neither does \hat{p}_P , ensuring that $(\mathbf{b}, \mathbf{s}, p_U, \hat{b})$ for $p_U = \hat{p}_P$ 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 all pairs (b, s) that are not matched in equilibrium.¹⁴

Proposition 3 *Let $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ be a personalized price equilibrium for which setting $p_U(s) = p_P(\tilde{b}(s), s)$ yields a uniform-price equilibrium. Then for all $s \in \mathcal{S}$,*

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

Proof. It follows from (9)–(12) (without any differentiability assumptions beyond those placed on the primitives of the model in Assumption 1), that if $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$ 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 $\frac{dh_S(\tilde{b}(s), s)}{db} = 0$. ■

4.3 Example: Coincident Personalized-Price and Uniform-Price Equilibria

Let $\theta = 1$ in our example from Section 3.3. Then $h_S(b, s) = 0$, and Proposition 2 ensures that the personalized-price equilibrium we have constructed is

¹⁴ Analogously, the single-crossing condition is essentially necessary for a separating equilibrium in a signaling model.

a uniform-price equilibrium. In confirmation, we note that the equilibrium pricing function is given by (when $\theta = 1$)

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

and hence is independent of the buyer's type.

4.4 Example: Divergent Personalized-Price and Uniform-Price Equilibria

We calculate a uniform-price equilibrium for our example, under the assumption that $\theta \leq 1$. The buyer's problem is now

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

for first-order conditions of

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

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 of

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

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

$$\mathbf{b}(\beta) = A\beta \tag{13}$$

$$\mathbf{s}(\sigma) = B\sigma. \tag{14}$$

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 buyer's second 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

$$b = \theta \frac{B}{A} \beta. \quad (15)$$

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

$$s = (2 - \theta) \frac{A}{B} \sigma. \quad (16)$$

Combining (13) with (15) and (14) with (16), 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 symmetric behavior from buyers and sellers. When $\theta < 1$, so that the seller owns some of the pre-transfer surplus, 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 choice levels initially increase, and ultimately decrease to zero as $\theta = 0$. 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 pricing equilibrium is not a personalized pricing equilibrium. To do this, it suffices to choose an attribute choice level x , 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 payoff is less than x^2 (ensuring that the buyer and seller could do better matching with each other). This ensures that we have a violation of the stability condition (5) and hence that

there exists a profitable (unpriced) deviation. 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.

Remark 3 Notice that if $\theta = 0$, so that the seller owns all of the surplus, then the equilibrium collapsed into the trivial equilibrium in which no trade occurs. (Assumption 1.4 excludes this case by assuming that $dh_B^2(b, s)/dbds > 0$.) In this case, buyers receive all of their payoff from 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 price. Because sellers cannot condition prices on buyer attribute choice, every buyer will choose $\mathbf{b} = 0$ in equilibrium. In this case, buyers would be a more appropriate choice to post prices. ■

4.5 Uniform-Price Equilibria: Existence

A fixed-point argument (in Section 7) allows us to establish:

Proposition 4 *A complete uniform price equilibrium exists.*

5 Endogenous Pricing

We turn now to the full model in which the first round calls for buyers and sellers to choose their attributes and for sellers to simultaneously choose whether to obtain the monitoring technology. We refer to this as the *endogenous* model to reflect the fact that sellers choose between personalized prices and uniform prices.

5.1 The Monitoring Decision

As before, we denote by \mathcal{B} and \mathcal{S} the closure of the sets of attributes chosen by buyers and sellers respectively. We denote the set of informed sellers by \mathcal{I} and the closure of the set of attributes chosen by informed sellers by $\mathcal{S}(\mathcal{I})$. Similarly, the set of uninformed sellers is \mathcal{U} and the closure of the set of attributes chosen by uninformed sellers is $\mathcal{S}(\mathcal{U})$. It is possible that $\mathcal{S}(\mathcal{I}) \cap \mathcal{S}(\mathcal{U}) \neq \emptyset$ 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 sufficiently well-behaved that $\lambda(\mathcal{S}(\mathcal{I}) \cap \mathcal{S}(\mathcal{U})) = 0$ (we return to this below). This is 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 $\tilde{b}_i : \mathcal{S}(\mathcal{I}) \rightarrow \mathcal{B}$ is a one-to-one measure-preserving function describing the match of an informed seller with attribute choice $s \in \mathcal{S}(\mathcal{I})$, and its inverse on $\tilde{b}_i(\mathcal{S}(\mathcal{I}))$ is denoted \tilde{s}_i . The function $\tilde{b}_u : \mathcal{S}(\mathcal{U}) \rightarrow \mathcal{B}$ is a one-to-one measure-preserving function describing the match of an uninformed seller with attribute choice $s \in \mathcal{S}(\mathcal{U})$ and its inverse on $\tilde{b}_u(\mathcal{S}(\mathcal{U}))$ is denoted \tilde{s}_u . The pair $(\tilde{b}_i, \tilde{b}_u)$ is a feasible matching if, in addition, $\tilde{b}_i(\mathcal{S}(\mathcal{I})) \cup \tilde{b}_u(\mathcal{S}(\mathcal{U})) = \mathcal{B}$. Since only sellers who become

informed can condition their price on buyers' investments, there will be two price functions: $p_P : \mathcal{B} \times \mathcal{S}(\mathcal{I}) \rightarrow \mathfrak{R}$ and $p_U : \mathcal{S}(\mathcal{U}) \rightarrow \mathfrak{R}$. 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.

An allocation is thus a collection $\{\mathbf{b}, \mathbf{s}, \mathcal{I}, p_P, p_U, \tilde{b}\}$, where \tilde{b} is a feasible matching. Let \mathcal{A} denote an allocation. We let the cost of the monitoring technology to a seller of type σ be written as $K(\sigma, \mathcal{A})$. We thus entertain the possibility that the cost of the technology may depend upon the allocation. At one extreme, the cost may consist entirely of a fixed cost, so that $K(\alpha, \mathcal{A}) = K$. 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.

5.2 Equilibrium

Given a set of informed sellers, \mathcal{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)

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

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

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}_i(\mathbf{b}(\beta)), \beta) = \sup_{(b, s) \in \mathcal{B} \times \mathcal{S}} \Pi_B(b, s, \beta) \quad (17)$$

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

$$\Pi_B(\mathbf{b}(\beta), \tilde{s}_u(\mathbf{b}(\beta)), \beta) = \sup_{(b, s) \in \mathcal{B} \times \mathcal{S}} \Pi_B(b, s, \beta). \quad (18)$$

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

We turn now to the priced deviations by the seller. In keeping with our discussion above on the matching, a choice of $s \in \mathcal{S}(\mathcal{I})$ is then the joint decision to become an informed seller and a choice of s , while a choice of $s \in \mathcal{S}(\mathcal{U})$ 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,

$$\Pi_S(\mathbf{s}(\sigma), \sigma) = \sup_{s \in \mathcal{S}} \Pi_S(s, \sigma), \quad \forall \sigma \in [0, 1], \quad (19)$$

and

$$\sigma \in \mathcal{I} \iff \mathbf{s}(\sigma) \in \mathcal{S}(\mathcal{I}). \quad (20)$$

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 2 and 6:

Definition 8 *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 $p \in \mathfrak{R}$, with either $s' \notin \mathcal{S}(\mathcal{I})$ 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 $p \in \mathfrak{R}$, with either $s' \notin \mathcal{S}(\mathcal{U})$ 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' ,

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

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 9 *Buyer β has a profitable unpriced deviation (with an informed seller) if there exists an attribute choice $b' \notin \mathcal{B}$, a price $p \in \mathfrak{R}$, and $s' \in \mathcal{S}(\mathcal{I})$ 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 10 *An endogenous-price equilibrium is an allocation satisfying (17), (18), (19), and (20), and such that no seller or buyer has a profitable unpriced deviation.*

5.3 Optimal Uniform Prices

We can readily provide conditions under which the equilibrium of the endogenous model exhibits uniform pricing.¹⁵

¹⁵Proposition 4 established the existence of uniform price equilibria (and, via Proposition 1, 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.

Proposition 5

5.1 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 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 default shares 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 4 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 uniform price equilibrium $(b(\beta), s(\sigma), p_U(s))$ is not a personalized-price equilibrium outcome. 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

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

and

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

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. ■

Remark 5 One might then 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 existing of a personalized price equilibrium. ■

5.4 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

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

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 transfer equilibrium precludes their from 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.¹⁶

6 Discussion

Moral Hazard. It is worth discussing the default shares 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' default surplus should be independent of the buyer to whom they are matched. An obvious question is that if this allocation of the default shares is optimal, why don't individuals simply contract to make this the case? 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 landlords. Suppose that sellers' default shares are set at zero and uniform pricing prevails. The parties will match, but the 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 default share, and it will typically not leave the seller indifferent over buyers with whom he might match.

Uniform pricing definition. In the definition of a uniform price equilibrium, a seller was assumed to have a profitable unpriced deviation only if he

¹⁶Let $\bar{c} = \lim_{\sigma \uparrow \psi} c_S(\mathbf{s}(\sigma), \sigma)$ and $\underline{c} = \lim_{\sigma \downarrow \psi} c_S(\mathbf{s}(\sigma), \sigma)$. Then we need

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

benefited by matching with *any* buyer who was attracted to the seller following the deviation. Alternatively, one might think that a seller could identify a deviation that would attract multiple buyers and choose randomly among the pool of such buyers. Stiglitz and Weiss (1981) show that pricing strategies of this type can, in fact, be optimal for banks making loans to potential borrowers of differing, unobservable creditability. There is a difference here however. In Stiglitz and Weiss, potential borrowers who are willing to pay high interest rates are those with projects that have high variance returns. Thus, when a bank raises its interest rate, the pool of potential borrowers shrinks, but it is the borrowers who are most attractive to the lender that disappear. This is not the case in our framework. We have assumed that higher attribute buyers have a strictly higher default value of matching with a particular seller than lower attribute buyers do. Thus if the seller offered a price that attracted multiple buyers, the high attribute buyers in the pool have higher default surplus than do the low attribute buyers, and consequently, the high attribute buyers in the pool would have a higher payoff trading with the seller than in the proposed equilibrium. Thus, if a seller who attracted multiple buyers raised his price, the lowest attribute buyers would no longer find the seller attractive. This does not mean that a seller should necessarily raise his price to discourage low attribute buyers however. While the high default surplus potential buyers gain the most from matching with a given seller, they typically will also have the most attractive alternatives.

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.¹⁷ To avoid this, Harvard Business School students successfully lobbied for policies that prohibit students’ divulging their grades to potential employers, and Wharton student government adopted a policy banning the release of grades.¹⁸ While we cannot be certain that it is always the case, we point out that in our example above, 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.

Who should set prices? Our analysis makes clear that when both buyers’ and sellers’ default surpluses depend on the attribute of the person they match with, uniform pricing will lead to inefficient investment regardless of which side is assigned the price-setting role. We can see from the uniform price calculations in our running example that it is not irrelevant who sets prices. For a matched pair $\beta = \sigma = \phi$, the buyers’ and sellers’ payoffs in that example are $[\frac{1}{6}\theta(2 - \theta)^2\phi^2]$ and $[\frac{1}{6}\theta^2(2 - \theta)\phi^2]$ respectively, and the sum is $\frac{1}{3}\theta(1 - \theta)\phi^2$. If we were to reverse the roles of buyer and seller, the share of the default surplus owned by the buyer would be $(1 - \theta)$ instead of θ . The surpluses to the two agents after the reversal of the buyer-seller roles are $\frac{1}{6}(1 - \theta)(1 + \theta)^2\phi^2$ and $\frac{1}{6}(1 - \theta)(1 + \theta)^2\phi^2$, and the total surplus is $\frac{1}{3}(1 - \theta)^2\phi^2$. It is straightforward to verify that the total surplus for any matched pair is greater when the price setter is the agent who has a smaller share of the surplus, that is, the buyer when $\theta < 1/2$ and the seller when $\theta > 1/2$. Reversing the roles of the buyer and seller when the buyer has a

¹⁷See Grossman (1981), Milgrom (1981) or Okuno-Fujiwara and Suzumura (1990), e.g., for analyses of this.

¹⁸Ostrovsky and Schwarz (2005) investigate the optimal amount of information to disclose from the students’ perspective.

smaller share of the surplus than the seller ($\theta < 1/2$) increases the combined surplus of each matched pair, but the reversal doesn't result in a Pareto improvement. The seller in our example ???

We emphasize that we are only demonstrating the optimality of assigning the seller role to the agent with the smaller default surplus share for this example. Examining this question more generally is an important question for future research.

Do actual markets use personalized prices? Young and Burke (2001) 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 $3/5-2/5$.¹⁹ 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.

Real estate markets have a similar regularity: brokers typically charge 6% independent of the characteristics of the house to be sold or of market

¹⁹Bardhan (1976), Bardhan and Ashok (1980), Bardhan (1981), Bardhan (1984) and Bardhan and Rudra (1986), document similar patterns for village economies in India and Africa.

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.²⁰ If such improbable regularities coincide with nontrivial investments by one side or another, inefficient investments are likely.

7 Appendix

Proof of Lemma 1. Suppose $\bar{\beta} > \underline{\beta}$ but $\mathbf{b}(\bar{\beta}) < \mathbf{b}(\underline{\beta})$. The optimality of $\mathbf{b}(\underline{\beta})$ implies

$$\Pi(\mathbf{b}(\underline{\beta}), \tilde{s}(\mathbf{b}(\underline{\beta})), \underline{\beta}) - c_{\beta}(\mathbf{b}(\underline{\beta}), \underline{\beta}) \geq \Pi(\mathbf{b}(\bar{\beta}), \tilde{s}(\mathbf{b}(\bar{\beta})), \underline{\beta}) - c_{\beta}(\mathbf{b}(\bar{\beta}), \underline{\beta}).$$

The single-crossing property of the cost function and $\bar{\beta} > \underline{\beta}$ now imply

$$\Pi(\mathbf{b}(\underline{\beta}), \tilde{s}(\mathbf{b}(\underline{\beta})), \bar{\beta}) - c_{\beta}(\mathbf{b}(\underline{\beta}), \bar{\beta}) > \Pi(\mathbf{b}(\bar{\beta}), \tilde{s}(\mathbf{b}(\bar{\beta})), \bar{\beta}) - c_{\beta}(\mathbf{b}(\bar{\beta}), \bar{\beta}),$$

contradicting the optimality of $\mathbf{b}(\bar{\beta})$. A similar argument holds for \mathbf{s} .

Suppose that the matching \tilde{b} is not weakly increasing, so that there exists $\bar{s} > \underline{s}$ with $\underline{b} \equiv \tilde{b}(\bar{s}) < \tilde{b}(\underline{s}) \equiv \bar{b}$. The stability condition (5) gives

$$\begin{aligned} h_B(\bar{b}, \underline{s}) - p_P(\bar{b}, \underline{s}) + h_S(\underline{b}, \bar{s}) + p_P(\underline{b}, \bar{s}) &\geq h_B(\bar{b}, \bar{s}) + h_S(\bar{b}, \bar{s}) \\ h_B(\underline{b}, \bar{s}) - p_P(\underline{b}, \bar{s}) + h_S(\bar{b}, \underline{s}) + p_P(\bar{b}, \underline{s}) &\geq h_B(\underline{b}, \underline{s}) + h_S(\underline{b}, \underline{s}). \end{aligned}$$

Adding these two inequalities, we get a contradiction to the supermodularity of the surplus function v .

Finally, we note that since \mathbf{b} , \mathbf{s} and \tilde{b} are increasing, so can we take \tilde{B} to be increasing.²¹ Then the measure-preserving requirement on \tilde{b} ensures that \tilde{B} is the identity. \blacksquare

²⁰ A 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.

²¹ Note that \tilde{B} may not be uniquely determined if \mathbf{b} or \mathbf{s} is not strictly increasing.

Proof of Lemma 3. Suppose there exists $b \in \mathcal{B}$ and $s' \in [0, \bar{s}]$ such that $W(b, s', \phi) > W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi)$. We will argue that 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 $W(b, s', \phi) > W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi) + 2\varepsilon$,

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

where the last inequality is an implication of (1). Now, the payoff to the seller from this deviation is then

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

and so the unpriced deviation is profitable.

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

Proof of Proposition 1. Suppose first that $h_S(b, s) = 0$ and hence $h_B(b, s) = v(b, s)$ for all pairs (b, s) . Proposition 4 ensures that there exists a complete uniform-price equilibrium. Proposition 2 ensures that there is a corresponding complete personalized price equilibrium $(\mathbf{b}, \mathbf{s}, p_P, \tilde{b})$, which Lemma 4 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 of the market in question. \blacksquare

Proof of Proposition 4. Let $P = v(\bar{b}, s^\dagger)$. Then P is sufficiently large that no buyer would be willing to purchase any seller attribute choice $s \in [0, s^\dagger]$ 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 1.4, so that for all $\epsilon > 0$, $s \in [0, s^\dagger - \epsilon]$, and $b \in [0, \bar{b}]$, we have $h_B(b, s + \epsilon) - h_B(b, s) < \Delta\epsilon$. As a result, given a choice between seller s and seller $s + \epsilon$ at a price higher by $\Delta\epsilon$, buyers would always choose the former. Equilibrium prices will thus never increase at a rate faster than Δ . Let $\bar{\Delta} \geq 2\Delta$ be set so that $2P/\bar{\Delta} < s^\dagger - \bar{s}$.

We now consider a sequence of games, indexed by n . Each game n has three players, consisting of a buyer, a seller, and a price-setter. We let -1 denote a decision to not participate in the market (which is distinct from a decision to participate, but with investment 0).

The buyer chooses two functions, $\mathbf{b} : [0, 1] \rightarrow \{-1\} \cup [0, \bar{b}]$ and $\tilde{s} : [0, 1] \rightarrow \{-1\} \cup [0, s^\dagger]$, identifying a buyer attribute choice and seller attribute choice with whom to match, each as a function of the buyer's type, with the restrictions that (i) $\tilde{s}(\beta) = -1$ if and only if $\mathbf{b}(\beta) = -1$, (ii) for any $\epsilon > 0$ and $\beta \in [0, 1 - \epsilon]$ with $\mathbf{b}(\beta) \geq 0$, $\epsilon/n \leq \mathbf{b}(\beta + \epsilon) - \mathbf{b}(\beta) \leq n\epsilon$, and (iii) for any $\epsilon > 0$ and $\beta \in [0, 1 - \epsilon]$ with $\tilde{s}(\beta) \geq 0$, $\epsilon/n \leq \tilde{s}(\beta + \epsilon) - \tilde{s}(\beta) \leq n\epsilon$. The seller chooses a function $\mathbf{s} : [0, 1] \rightarrow \{-1\} \cup [0, s^\dagger]$ with $\epsilon/n \leq \mathbf{s}(\sigma + \epsilon) - \mathbf{s}(\sigma) \leq n\epsilon$ for any ϵ and $\sigma \in [0, 1 - \epsilon]$ with $\mathbf{s}(\sigma) \geq 0$. The price-setter chooses an increasing function $p_U : [0, s^\dagger] \rightarrow [-P, P]$ with the property that for all $\epsilon > 0$ and $s \in [0, s^\dagger - \epsilon]$, $p_U(s + \epsilon) - p_U(s) < \bar{\Delta}\epsilon$, and with $p_U(s^\dagger) = P$. Hence, we restrict the price setter so that the set of prices offered in the market is of the form $[p, P]$ for some $p \in [-P, P]$. The presence of high prices will be useful when arguing that an equilibrium of the game we examine corresponds to a uniform-price equilibrium of the market. The assumption that the price-setter prices the interval $[0, s^\dagger]$ rather than $[0, \bar{s}]$ ensures that this constraint can be satisfied without disrupting the properties of the equilibrium.

Let

$$\begin{aligned} G_B(b) &= \lambda\{\beta : \mathbf{b}(\beta) \leq b\} \\ F_B(s) &= \lambda\{\beta : \tilde{s}(\beta) \leq s\} \\ F_S(s) &= \lambda\{\sigma : \mathbf{s}(\sigma) \leq s\}. \end{aligned}$$

We delete the superscript n identifying the game (and hence the functions \mathbf{b}^n , \tilde{s}^n and \mathbf{s}^n) in which these functions are defined. Each of the functions $G_B : \{-1\} \cup [0, s^\dagger] \rightarrow [0, 1]$ and $F : \{-1\} \cup [0, s^\dagger] \rightarrow [0, 1]$ is increasing and continuous from the right. Let $\hat{b}(s)$ be such that $G_B(\hat{b}(s)) = F_S(s)$ for all $s \in [0, s^\dagger]$. Our interpretation will be that a seller of type s matches with a buyer randomly drawn from the interval $[\lim_{s' \uparrow s} \hat{b}(s'), \hat{b}(s)]$ according to the distribution induced by $G_B(s)$. Notice that this interval is degenerate

in each game n , so that the random assignment feature of this definition is germane only in game ∞ . Here, it will allow us to address cases where a positive measure of buyers choose different attributes b but choose the same seller attribute s .

The buyer maximizes (again, omitting the superscript n)

$$\int_0^1 (h_B(\mathbf{b}(\beta), \tilde{s}(\beta)) - p_U(\tilde{s}(\beta)) - c_B(\mathbf{b}(\beta), \beta)) d\beta.$$

The seller maximizes

$$\int_0^1 h_S(\hat{b}(\mathbf{s}(\sigma)), \mathbf{s}(\sigma)) + p_U(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma) d\sigma.$$

The price-setter maximizes

$$\int_0^{\bar{s}} \varphi(p_U(s)) p_U(s) [F_B(s) - F_S(s)]_+ + \varphi(p_U(s)) [F_B(s) - F_S(s)]_- ds,$$

where $[x]_+ = \max\{x, 0\}$, $[x]_- = \min\{x, 0\}$, and φ is a strictly increasing strictly concave function from $[-P, P]$ into \mathfrak{R} .

We refer to the n th element of this sequence of games as game n . We refer to the game with no restrictions of the form $1/n \leq f(x + \epsilon) - f(x) \leq n$ on the buyer and seller as game ∞ . Note that the restrictions on the price-setter remain fixed across games.

We now examine a product space of functions. We let Υ_b be the set of pairs of functions $\mathbf{b} : [0, 1] \rightarrow \{-1\} \cup [0, \bar{b}]$ and $\tilde{s} : [0, 1] \rightarrow \{-1\} \cup [0, s^\dagger]$, with the properties that (i) $\tilde{s}(\beta) = -1$ if and only if $\mathbf{b}(\beta) = -1$, (ii) for any $\epsilon > 0$ and $\beta \in [0, 1 - \epsilon]$ with $\mathbf{b}(\beta) \geq 0$, $\epsilon/n \leq \mathbf{b}(\beta + \epsilon) - \mathbf{b}(\beta) \leq n\epsilon$, and (iii) for any $\epsilon > 0$ and $\beta \in [0, 1 - \epsilon]$ with $\tilde{s}(\beta) \geq 0$, $\epsilon/n \leq \tilde{s}(\beta + \epsilon) - \tilde{s}(\beta) \leq n\epsilon$. Let a norm for an element (\mathbf{b}, \tilde{s}) be given by the sum of the L^1 norms on the functions \mathbf{b} and \tilde{s} . We let Υ_s be the set of functions $\mathbf{s} : [0, 1] \rightarrow \{-1\} \cup [0, s^\dagger]$ with $\epsilon/n \leq \mathbf{s}(\sigma + \epsilon) - \mathbf{s}(\sigma) \leq n\epsilon$ for any ϵ and $\sigma \in [0, 1 - \epsilon]$ with $\mathbf{s}(\sigma) \geq 0$. Endow this set with the L^1 norm. Let Υ_p be the set of increasing functions $p_U : [0, s^\dagger] \rightarrow [-P, P]$ with the properties that $p_U(s^\dagger) = P$ and for all $\epsilon > 0$ and $s \in [0, s^\dagger - \epsilon]$, $p_U(s + \epsilon) - p_U(s) < \bar{\Delta}\epsilon$, again endowed with the L^1 norm. Let $\Upsilon = \Upsilon_b \times \Upsilon_s \times \Upsilon_p$, with norm given by the sum of the three constituent norms. Then Υ is a compact metric space.²²

²²It 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 allows us to show that Υ is

We next fix a value n and show that in game n , the payoff functions are continuous, given this topology and the strategy sets in game n . We first note that for increasing, bounded functions on a compact set, L^1 convergence implies convergence almost everywhere.²³

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 $c_B(\mathbf{b}(\beta), \beta)$. The first and the third of these follows from the continuity of h_B and c_B (from Assumption 1), while for the remaining case it suffices to note that game n is defined so that p_U is continuous

We next note that the functions F_B and F_S converge almost everywhere. This follows from the facts that \mathbf{s} and \tilde{s} are increasing, and are continuous and strictly increasing when nonnegative. For the price setter, continuity then follows from arguments analogous to those applied to the buyer, since we have convergence almost everywhere of $p_U(s)[F_B(s) - F_S(s)]$.

For the seller, continuity follows from similar arguments if we can show that $\hat{b}(\mathbf{s}(\sigma))$ converges almost everywhere. Fix a value σ at which $\mathbf{s}^*(\sigma)$ is continuous (i.e., any value at which \mathbf{s}^* does not make a jump from -1 to 0). Then $F_S^n(s)$ converges to a value $F_S^*(s)$. Let $\beta' = F_S^*(s)$. Then $\hat{b}(\mathbf{s}^n(\sigma))$ converges to $\mathbf{b}^*(\beta')$ as long as \mathbf{b}^* does not jump from -1 to 0 at β' , and hence converges almost everywhere.

Glicksberg (1952) shows that this game has a Nash equilibrium, though this equilibrium may be mixed.

We now select one such fixed point for each game in our sequence, de-

sequentially compact. In particular, given a sequence $\{(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}^n, p_U^n)\}$, we can choose a subsequence 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 $\{(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}^n, p_U^n)\}$ 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 value β (for example) at which $\lim_{\mathcal{Q} \ni \beta' \uparrow \beta} \mathbf{b}^\infty(\beta') = \lim_{\mathcal{Q} \ni \beta' \downarrow \beta} \mathbf{b}^\infty(\beta')$, and hence almost everywhere, which suffices (for bounded functions) for L^1 convergence.

²³Suppose 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 $\epsilon > 0$, all $z \in [x, y]$ and for all sufficiently large n , we have $f_n(z) \geq f(y) + \epsilon \geq f(z) + \epsilon$. This in turn ensures that $\int |f_n(z) - f(z)| dz > (y - x)\epsilon$, precluding the L^1 convergence of $\{f_n\}_{n=1}^\infty$ to f .

noted by $(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}^n, p_U^n)$. By taking an appropriate subsequence (and re-defining to avoid additional notation), we can assume that $\{(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}^n, p_U^n)\}_{n=1}^\infty$ converges to a limit $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$. We remember here that the equilibrium strategies might be mixed, without introducing new notation to capture this possibility.

The argument now continues in four steps.

[STEP 1] The profile $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$ is an equilibrium of game ∞ . We present the argument for pure strategies, and for the seller (the least obvious case). The extension to mixtures is an analogous argument is a similar manipulation of expected payoffs.

Let Π_s^n denote the seller's payoff in game n . Then we need to show that

$$\Pi_s^\infty(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*) \geq \Pi_s^\infty(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}, p_U^*)$$

for any alternative \mathbf{s} . It suffices to show that for any game n ,

$$\Pi_s^n(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}^n, p_U^n) \geq \Pi_s^n(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}, p_U^n) - \epsilon_n,$$

where $\{\epsilon_n\}_{n=1}^\infty$ is a sequence of positive values converging to zero, and that

$$\begin{aligned} \Pi_s^n(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}^n, p_U^n) &\rightarrow \Pi_s^\infty(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*) \\ \Pi_s^n(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}, p_U^n) &\rightarrow \Pi_s^\infty(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}, p_U^*), \end{aligned}$$

for any sequence \mathbf{s}^n converging to a limit \mathbf{S}^* and for any strategy \mathbf{s} .

We first consider the convergence results. It suffices to prove the first result, which subsumes the second. The argument is reminiscent of the argument that payoff functions in game n are continuous, in that we prove that the constituents of the integrand in the seller's payoff converge almost everywhere and then appeal to the dominated convergence theorem.

Because c_B is continuous, it is immediate that $c_S(\mathbf{s}(\sigma))$ converges almost everywhere if $\mathbf{s}(\sigma)$ does. Similarly, because p_U^n converges to continuous limit, $p_U(\mathbf{s}(\sigma))$ converges almost everywhere. (Here, we use our continuity restriction on the function p_U .) Hence, consider $h_S(\hat{b}(\mathbf{s}(\sigma), \mathbf{s}(\sigma)))$. We again note that the function h_S is continuous and \mathbf{s} converges almost everywhere. However, we must cope with the fact that in game ∞ (though in no game n) the function \hat{b} and its implied matching may attach many values of b , comprising a set of positive measure, to a single value of s , chosen by a set of seller types of matching measure. To address this, define, for any increasing \mathbf{b} and \mathbf{s} , the function given by

$$\check{b}(\sigma) = \mathbf{b}(\sigma).$$

We view this as matching seller σ , who chooses attribute $\mathbf{s}(\sigma)$, with the attribute chosen by buyer $\beta = \text{sigma}$. We now note that in any game,

$$\int_0^1 h_S(\hat{b}(\mathbf{s}(\sigma)), \mathbf{s}(\sigma)) + p_U(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma) d\sigma = \int_0^1 h_S(\check{b}(\sigma), \mathbf{s}(\sigma)) + p_U(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma) d\sigma$$

(here, we reap the advantage of working with a single buyer and seller with payoffs defined as integrals over types). It thus suffices to show that this payoff function converges as $(\mathbf{b}^n, \tilde{s}^n, \mathbf{s}^n, p_U^n) \rightarrow (\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$. But the convergence almost everywhere of \mathbf{s} , \mathbf{b} , and p_U , coupled with the continuity of h_S , c_S , p_U^n and p_U^∞ , suffices for the result.²⁴

We next need to show that for any $\eta > 0$, there is a sufficiently large N such that the restrictions placed on the seller's strategy in game n (i.e., that $\epsilon/n \leq \mathbf{s}(\sigma + \epsilon) - \mathbf{s}(\sigma) \leq n\epsilon$) cause no more than η loss, compared to an unrestricted best response. To do this, we first note that because the seller's cost function satisfies single-crossing, the optimal function \mathbf{s} in the absence of restrictions, in any game n or game ∞ , is increasing. This in turn implies that, for any \mathbf{s} and in game n , there exists a strategy \mathbf{s}^n that is within $1/n$ of \mathbf{s} for all s on a set of measure $1 - \bar{s}/n$. On the residual set of measure \bar{s}/n , the payoff loss from playing a best response is perhaps very large, but bounded. On the remaining set, the restriction on the slope of the payoff function ensures that the payoff loss contributed by the term $p_U(s)$ is at most Δ/n . Coupled with the uniform continuity of h_S and c_S (being continuous on compact sets), this suffices for the result.

[STEP 2] The profile $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$ is pure. (Notice that Glicksberg's theorem does not imply this.) The concavity of φ ensures that the price setter will never choose a mixture in any game n . In particular, if the price setter is mixing over a collection of functions $\{p_U(\theta)\}$ for some θ drawn from an index set Θ , with each function satisfying the constraint that it be no steeper than $\bar{\Delta}$, then the expected value similarly satisfies this constraint, and yields a higher payoff.

Consider the buyer (the case of the seller is analogous). Suppose that \mathbf{b}^* and/or \tilde{s} are mixed. This mixture induces the pair of functions G_B and F_B , describing the distributions of buyer attributes b and seller attributes \tilde{s} chosen by the buyer.

Now define a pair of increasing functions $\mathbf{b}' : [0, 1] \rightarrow \{-1\} \cup [0, \bar{b}]$ and

²⁴Why not work with the function \check{b} rather than \hat{b} throughout? Because it does not suffice to show that an equilibrium of the game is a uniform price equilibrium of the market.

$\tilde{s}' : [0, 1] \rightarrow \{-1\} \cup [0, s^\dagger]$ by

$$\begin{aligned}\mathbf{b}'(\beta) &= \inf\{b : G_B(b) \geq \beta\} \\ \tilde{s}'(\beta) &= \inf\{\tilde{s} : G_B(\tilde{s}) \geq \beta\}.\end{aligned}$$

These functions gives the same distribution of b and \tilde{s} in the market, but feature positive assortivity between buyer types and attribute choices and between buyer and seller attribute choices, both of which increase the buyer's payoff. Hence, the buyer cannot mix.

[STEP 3] The profile $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$ balances the market in the sense that $F_B^*(s) = F_S^*(s)$ for all s . Since F_B^* and F_S^* are continuous from the right, it suffices to show that they agree almost everywhere. We first argue that $F_B^*(s) - F_S^*(s) \leq 0$ almost everywhere. Suppose this is not the case, so there exists $s < s^\dagger$ with $F_B^*(s) - F_S^*(s) = \epsilon > 0$ and with s a continuity point of $F_B^* - F_S^*$. Then there exists s_1 and s_2 with $s \in [s_1, s_2)$, and $F_B^*(s) - F_S^*(s) \geq \epsilon/2$ on $[s_1, s_2]$ and either $s_1 = 0$ or, for every $\eta > 0$, a value $s_\eta \in [s_1 - \eta, s_1)$ with $F_B^*(s_\eta) < \epsilon/2$. We consider the case in which $s_1 > 0$, with the case $s_1 = 0$ being a straightforward simplification.

Since $F_B^*(s) - F_S^*(s) > 0$ on $[s_1, s_2]$, the price setter must be setting prices as large as possible on this interval, and hence there must be $s' \in [s_1, s_2]$ with the property that $p_U^*(s) = \inf_{s \geq s_2} p_U(s)$ for $s \in [s', s_2]$ and with $dp_U(s)/ds = \bar{\Delta}$ on $[s_1, s']$. Hence, prices increase at the maximum rate possible until hitting a constraint, after which they are fixed at the constrained value. (Notice that one of these regions (maximum increase or constraint) may be degenerate.) As a result, $\tilde{s}([0, 1]) \cup [s_1, s_2] \subset \{s_1, s_2\}$, 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 will demand only attribute choice s_2 from this set, while the price of a seller attribute choice increases so sharply on $[s_1, s']$ that from this set buyers will demand only s_1 .) If buyers are not choosing s_1 , this is a contradiction to the facts that $F_B^*(s) - F_S^*(s) = \epsilon$ and either $s_1 = 0$ or, for every $\eta > 0$, there exists a value a $s_\eta \in [s_1 - \eta, s_1)$ with $F_B^*(s_\eta) < \epsilon/2$. Hence, some buyers must be choosing s_1 .

If buyers are to choose s_1 , there must exist values of s arbitrarily close to s_1 whose prices are at least $p_U(s_1) - \Delta(s_1 - s)$. Otherwise, Assumption 1.4 ensures that no buyer chooses s_1 . Hence, for some $\eta > 0$, which we can take to be as small as desired, the function p_U on $[s_1 - \eta, s']$ is at least

$$P'_U = \max\{p_U(s_1 - \eta), p_U(s_1) + (s - s_1)\bar{\Delta}\}.$$

Now let $p_U(s, \zeta)$ be a function defined on $[s_1 - \eta, s']$ with $p_U(s, 0) = p_U(s)$ and with $p_U(s, \zeta) = \max\{p_U(s), p_U(s_1) + (s - s_1)\bar{\Delta}\}$ if this value is smaller than

$p_U(s')$, and equal to $p_U(s')$ otherwise. Hence, as ζ increases, we increase the function p_U uniformly on the interval s_1, s' , subject to the constraint that this not push the price beyond $p_U(s')$ and that prices below s_1 are increased so as to preserve the slope constraint on p_U .²⁵ A lower bound on the effect of such a price change on the price-setter's payoff is given by taking the derivative of

$$\int_{s_1-\eta}^{s_1} \varphi(p_U(s, \zeta))(-1)ds + \int_{s_1}^{s'} \varphi(p_U(s, \zeta))(F_B(s) - F_S(s))ds.$$

in ζ . The derivative in ζ , evaluated at $\zeta = 0$, is given by

$$\int_{s_1}^{s'} \frac{d\varphi(p_U(s, \zeta))}{dp_U(s, \zeta)}(F_B(s) - F_S(s))ds > 0,$$

a contradiction to the optimality of the price-setter's behavior.

We conclude that $F_B^*(s) - F_S^*(s) \leq 0$ for almost all s . It remains to show 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 $p_U = -P$ for $s < \hat{s}$, $F_B^*(s) - F_S^*(s) < 0$ for a positive-measure subset of $[0, \hat{s}]$, and $F_B^*(s) - F_S^*(s) = 0$ for almost all $s > \hat{s}$. But then no seller would choose attributes in $[0, \hat{S})$, a contradiction.

[STEP 4] The profile $(\mathbf{b}^*, \tilde{s}^*, \mathbf{s}^*, p_U^*)$ corresponds to a uniform price equilibrium of the matching market. 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.

For the buyer, the absence of profitable deviations in the game is immediate, since these are the only possible type of deviation. We then note that any profitable priced deviation in the matching market would imply the existence of a profitable buyer deviation in the game. The same argument applies to priced deviations for the seller. For unpriced deviations, we first note that such a deviation can never entail a lower price. The lower price will always suffice to attract the seller's current match, 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

²⁵The observation that there exist values of s arbitrarily close to s_1 whose prices are at least $p_U(s_1) - \Delta(s_1 - s)$, plays a role here, ensuring that this price adjustment affects no prices lower than $s_1 - \eta$.

all such prices are already in the game, being attached to larger sellers. If such sellers already exist in the matching market, the seller in question will attract no buyers. If they do not exist in the matching market, it must be that no buyers are selecting them in the game. The fact that buyers are not doing so ensures that they will not select the deviating seller in the matching market, again making the deviation unprofitable. ■

References

- BARDHAN, P. K. (1976): “Variations in Extent and Forms of Agricultural Tenancy – I: Analysis of Indian Data Across Regions and Over Time,” *Economic and Political Weekly*, pp. 1505–1546.
- (1984): *Land Labor, and Rural Poverty: Essays in Development Economics*. Columbia University Press, New York.
- BARDHAN, P. K., AND R. ASHOK (1980): “Terms and Conditions of Sharecropping Contracts: An Analysis of Village Survey Data in India,” *Journal of Development Studies*, 16(3), 287–302.
- BARDHAN, P. K., AND A. RUDRA (1986): “Labour Mobility and the Boundaries of the Village Moral Economy,” *Journal of Peasant Studies*, 13, 90–115.
- BARDHAN, P. K. A. R. (1981): “Terms and Conditions of Labor Contracts in Agriculture: Results of a Survey in West Bengal, 1979,” *Oxford Bulletin of Economics and Statistics*, 43(1), 89–111.
- BULOW, J., AND J. LEVIN (2004): “Matching and Price Competition,” Mimeo, Stanford University.
- COASE, R. (1960): “The Problem of Social Cost,” *Journal of Law and Economics*, 2, 1–40.
- COLE, H. L., G. J. MAILATH, AND A. POSTLEWAITE (2001): “Efficient Non-Contractible Investments in Large Economies,” *Journal of Economic Theory*, 101, 333–373.
- DUNFORD, N., AND J. T. SCHWARTZ (1988): *Linear Operators Part I: General Theory*. John Wiley and Sons, New York, Wiley Classics Library Edition.

- GLICKSBERG, I. L. (1952): “A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points,” *Proceedings of the American Mathematical Society*, 3(1), 170–174.
- GROSSMAN, S. (1981): “The Informational Role of Warranties and Private Disclosure about Product Quality,” *Journal of Law and Economics*, 24, 461–483.
- MILGROM, P. R. (1981): “An Axiomatic Characterization of Common Knowledge,” *Econometrica*, 49, 215–218.
- NICHOLSON, S. (2003): “Barriers to Entering Medical Specialties,” Working Paper 9649, National Bureau of Economic Research.
- OKUNO-FUJIWARA, MASAHIRO, A. P., AND K. SUZUMURA (1990): “Strategic Information Revelation,” *Review of Economic Studies*, 57, 25–47.
- OSTROVSKY, M., AND M. SCHWARZ (2005): “Equilibrium Information Disclosure and Unraveling,” *mimeo*.
- ROTH, A., AND M. A. O. SOTOMAYER (1990): *Two-Sided Matching*. Cambridge University Press, Cambridge.
- STIGLITZ, J., AND A. WEISS (1981): “Credit Rationing in Markets with Imperfect Information,” *American Economic Review*, 71, 393–410.
- YOUNG, H. P., AND M. A. BURKE (2001): “Competition and Custom in Economic Contracts: A Case Study of Illinois Agriculture,” *American Economic Review*, 91, 559–573.