

NOTES ON PRICE DISCOVERY WITH PRICE-SETTING AGENTS

[Oft-Told Anecdote] Three people were shipwrecked on an island, a physicist, an engineer, and an economist, along with one can of beans. The physicist said: “Let’s use thermodynamic principles to open the can – we can light a fire and let pressure build up until the can explodes.” The engineer demurred, saying: “Let’s use basic mechanical principles – let’s bang on the can with a rock until it splits open.” The economist, sighing impatiently, then said: “Both your solutions are unpleasantly messy. I propose instead a simple but elegant solution – let’s assume we have a can opener.”

“Like a rapidly growing bush, (economic) theory may sometimes sprout and develop in unhelpful directions, but when pruned with the shears of practical experience it will quickly bear fruit.” Klemperer (2002b, p. 24)

1. Competitive Pricing Theory

The most prominent “can opener” imbedded in standard textbook economic theory is market clearing prices.

Typically, individual seller supply functions are assumed to be constructed from seller responses to the following question: For any given (unit) price for Q , what is the maximum total amount of Q you would be willing to sell?

This question presumes that sellers take prices as given when they decide their sale plans.

Similarly, individual buyer demand functions are assumed to be constructed from buyer responses to the following question: For any given (unit) price for Q , what is the maximum total amount of Q you would be willing to buy?

This question presumes that buyers take prices as given when they decide their purchase plans.

For example, suppose a market consists of N sellers, $n = 1, \dots, N$, and M buyers, $m = 1, \dots, M$.

The *ordinary (or direct) supply function* of seller n :

A schedule that shows, for each (unit) price p for Q , the maximum total amount of Q that seller n would be willing to *sell* at price p , denoted by

$$q_n^s = s_n^o(p) \tag{1}$$

The *ordinary (or direct) demand function* of buyer m :

A schedule that shows, for each (unit) price p for Q , the maximum total amount of Q that buyer m would be willing to *buy* at price p , denoted by

$$q_m^d = d_m^o(p) \tag{2}$$

True total seller ordinary supply function for Q :

A schedule that gives, for each price p for Q , the sum of all individual seller supplies $s_n^o(p)$ at p :

$$q^S = \sum_{n=1}^N s_n^o(p) \equiv S^o(p) \quad (3)$$

True total buyer ordinary demand function for Q :

A schedule that gives, for each price p for Q , the sum of all individual buyer demands $d_m^o(p)$ at p :

$$q^D = \sum_{m=1}^M d_m^o(p) \equiv D^o(p) \quad (4)$$

A competitive market clearing price p^ for Q :*

A solution p^* to the *competitive market clearing equation* $S^o(p) = D^o(p)$. By construction, at any such p^* , $S^o(p^*) = D^o(p^*)$.

Note on terminology:

In market analyses, *inverse* supply and demand functions give price as a function of quantity, whereas *ordinary* supply and demand functions give quantity as a function of price.

The inverse form is commonly used when the stress is on price-setting sellers and buyers who actively make price offers for different considered quantities.

The ordinary form is commonly used when the stress is on “competitive” sellers and buyers who make quantity decisions taking prices as given.

These notes will use the ordinary form of supply and demand functions for easier comparison with competitive pricing models. However, for expositional simplicity, the qualifier “ordinary” and the superscript “o” will be suppressed in the remainder of these notes.

What's wrong with the picture of an economy as depicted by this standard textbook *competitive pricing theory*?

First problem....

The vast majority of transactions in modern decentralized market economies occur in market contexts in which strategically interacting buyers and/or sellers continually adjust their prices in hopes of attracting profitable trades.

Competitive pricing theory does not address whether market clearing prices could ever arise *through the interactions of the market participants themselves*.

Rather, competitive pricing theory requires market prices to be set by a mathematical process external to the modeled economic world: namely, *determine a solution p to the equation $S(p) = D(p)$* .

Competitive pricing theory thus presumes that

- market prices reflect complete and accurate supply and demand information;
- all interactions among market participants take place only through prices;
- market participants take prices as given (i.e., they do not engage in any form of strategic pricing).

A second problem...

In real-world contexts in which some market participants are price takers (e.g., in one-sided posted offer auctions), the price takers generally must engage in some type of search process to discover what price offers have been posted and which price offers would be best to accept.

The resources used up in this search and discovery process are referred to as *transaction costs*.

Standard textbook competitive pricing theory ignores transactions costs.

A third problem...

Many important types of real-world seller-buyer transactions (e.g., financial asset transactions, service contracts) impose future contractual obligations on one or both parties to the transaction.

This can introduce uncertainty into the transaction due to *asymmetric information*, that is, sellers having different information than buyers regarding the ability or even willingness of the transaction participants to carry out their future contractual obligations to each other.

To reduce this uncertainty typically requires some form of informal or formal contract enforcement process. This might entail the development of trusted relationships (e.g., through reciprocity considerations), periodic monitoring, auditing, performance reviews, and so forth.

Resources used up in order to reduce asymmetric information problems in seller-buyer transactions are referred to as *information costs*.

Standard textbook competitive pricing theory does not explicitly address asymmetric information problems and information costs.

Prices in real world markets might or might not adequately conform to competitive pricing theory – the descriptive accuracy of this theory is an issue that can only be resolved by empirical testing.

A different point is being stressed here, a point forcefully raised by Arrow (1959), among others, and still unresolved today. This point can be summarized as follows.

Standard textbook competitive pricing theory is inconsistent with the commonly accepted idea that the outcomes of economic models should be explicitly grounded in “microfoundations”.

An economic model is said to be grounded in *microfoundations* if its outcomes are the consequences (intended or not) of the activities undertaken by the firms, consumers, and other individual decision makers within the economic model.

In particular, market prices should be determined by the activities of individual decision makers, not by imposed market clearing conditions.

2. Competitive Pricing: 2-Firm 1-Consumer Example

Consider an economy that consists of two price-taking profit-maximizing firms, each producing a distinct consumption good, and a price-taking budget-constrained consumer who obtains utility (happiness) from the consumption of these two goods. Specifically:

- Firm 1 produces a consumption good Q_1 , where the cost of supplying an amount q_1^s of Q_1 is $c_1(q_1^s)$.
- Firm 2 produces a consumption good Q_2 , where the cost of supplying an amount q_2^s of Q_2 is $c_2(q_2^s)$.
- Given any particular unit price p_1 for Q_1 , the objective of Firm 1 is to supply an amount q_1^s of Q_1 that maximizes its profits

$$p_1 q_1^s - c_1(q_1^s) \ . \quad (5)$$

- Given any particular unit price p_2 for Q_2 , the objective of Firm 2 is to supply an amount q_2^s of Q_2 that maximizes its profits

$$p_2 q_2^s - c_2(q_2^s) \ . \quad (6)$$

- The consumer has an income I that he desires to use to demand (buy) amounts q_1^d and q_2^d of Q_1 and Q_2 .
- Given any particular unit prices p_1 and p_2 for Q_1 and Q_2 , the objective of the consumer is to select “feasible” demands q_1^d and q_2^d to maximize his utility of consumption, measured by $U(q_1^d, q_2^d)$.
- *Feasibility* means the consumer’s demands are non-negative and his expenditures do not exceed his income, i.e., his demands satisfy the non-negativity constraints

$$q_1^d \geq 0 \text{ and } q_2^d \geq 0 \tag{7}$$

and the budget constraint

$$p_1 q_1^d + p_2 q_2^d \leq I \text{ .} \tag{8}$$

Price discovery under competitive pricing:

- Assume a fictitious auctioneer AUCT exists for this economy.
- AUCT posts arbitrary unit prices p_1 and p_2 for goods Q_1 and Q_2 .
- Firm 1 submits to AUCT a supply response q_1^s that maximizes its profits, conditional on p_1 .
- Firm 2 submits to AUCT a supply response q_2^s that maximizes its profits, conditional on p_2 .
- The consumer submits to AUCT feasible demand responses q_1^d and q_2^d that maximize his utility, conditional on p_1 and p_2 .
- If supply equals demand for each good, i.e., if $q_1^s = q_1^d$ and $q_2^s = q_2^d$, then AUCT allows trades to take place between Firms 1 and 2 and the consumer at the currently posted prices p_1 and p_2 .
- Otherwise, no trades are permitted to take place. Rather, AUCT posts new unit prices p'_1 and p'_2 for Q_1 and Q_2 , and the process repeats.
- In short, the fictitious auctioneer AUCT keeps posting new unit prices for Q_1 and Q_2 until prices are discovered that clear each market. Price discovery is thus carried out by AUCT, not by the consumers and firms who actually reside within the economy.

3. Strategic Pricing: 2-Firm 1-Consumer Example

Strategic interaction is said to arise between two decision makers A and B if the choices of decision maker B explicitly enter into the choice deliberations of decision maker A because A perceives or expects that B's choices can affect his own outcomes.

Specifically, A asks himself questions of the form: "Given B has done this, what should I do?", and "If I do this, what will B then do?"

As seen in the previous section, no strategic interaction arises among market participants under the assumptions of competitive pricing theory. Market participants are linked through prices and only through prices. Market participants take these prices as given aspects of their decision environments, outside of their control. Consequently, they do not perceive any way in which the decisions of other agents impinge on their own decisions.

An example will now be given illustrating how strategic interaction can arise between two profit-seeking firms if they are permitted to set their own prices in an attempt to compete for the dollars of a utility-seeking budget constrained consumers.

As in the previous section, consider an economy that consists of two profit-maximizing firms, each producing a distinct consumption good, and a budget-constrained consumer who obtains utility (happiness) from the consumption of these two goods. In contrast to the previous section, however, each firm is now assumed to be a price-

setter rather than a price-*taker*. That is, each firm is free to set its own output price at any level it chooses.

Suppose that the profit obtained by Firm n from the sale of good Q_n is given by

$$p_n q_n - c_n q_n \quad , \quad n = 1, 2, \quad (9)$$

where p_n denotes the (unit) price of good Q_n , q_n denotes the amount sold of good Q_n , and the constant per-unit marginal cost c_n is positive.

Suppose the utility obtained by the consumer from consumption of goods Q_1 and Q_2 can be measured by a utility function of the form

$$U(q_1, q_2) = \log(q_1 - b_1) + \log(q_2 - b_2) \quad , \quad (10)$$

where

- $\log(\cdot)$ denotes the logarithm function (base 10);
- b_1 and b_2 are given nonnegative constants representing “subsistence needs” levels.

Remark: Note that the log function is only defined over positive real numbers — the value of $\log(x)$ approaches minus infinity as x approaches 0 “from the right”. Consequently, b_1 and b_2 actually represent “death” values rather than subsistence levels; consumption of goods Q_1 and Q_2 must be greater than these levels in order for the consumer to survive.

Suppose, also, that the income I of the consumer is a positive exogenously-given constant, and that the consumer takes prices as given. The

utility maximization problem faced by the consumer then takes the following form: Given goods prices p_1 and p_2 , maximize

$$U(q_1, q_2) \tag{11}$$

with respect to the choice of q_1 and q_2 subject to the budget and nonnegativity constraints

$$p_1q_1 + p_2q_2 \leq I ; \tag{12}$$

$$q_1, q_2 \geq 0 . \tag{13}$$

Suppose the prices p_1 and p_2 permit the consumer to purchase Q_1 and Q_2 amounts exceeding b_1 and b_2 with his given income I . That is, suppose the following *survival condition* is satisfied:

$$p_1b_1 + p_2b_2 < I . \tag{14}$$

The solution to the consumer's utility maximization problem (11) then yields the following *demand functions* for q_1 and q_2 :

$$q_1^d = b_1/2 + [I - b_2p_2]/2p_1 = D_1(p_1, p_2) ; \tag{15}$$

$$q_2^d = b_2/2 + [I - b_1p_1]/2p_2 = D_2(p_1, p_2) , \tag{16}$$

where dependence of these demand functions on the exogenous variables b_1 , b_2 , and I has been suppressed for expositional simplicity.

Important Observations:

It follows directly from (15) and (16) that

$$p_n[q_n^d - b_n] = [I - b_1p_1 - b_2p_2]/2 , n = 1, 2$$

Thus, the rule for obtaining optimal demands in the current example takes the following simple form: First meet subsistence needs, then divide any remaining income equally between the two goods. This simple symmetrical form for optimal demand depends critically on the use (10) of a logarithmic utility function that separably weighs each good by the same logarithmic expression applied to goods consumption net of subsistence needs.

Note, also, that q_1^d is independent of p_2 in (15) if and only if $b_2 = 0$, and that q_2^d is independent of p_1 in (16) if and only if $b_1 = 0$. This illustrates the general rule of thumb that the optimal consumer demand for any one good will depend on the prices of all goods. This dependence arises because the demands of a consumer for goods in any one period are simultaneously and jointly constrained by a single budget constraint that requires total expenditures for all goods to equal total income (possibly modified by borrowing or lending).

An exception to this rule arises for the special case of an additive and purely logarithmic utility function, e.g., the utility function in (10) with $b_1 = 0$ and $b_2 = 0$; in this case, as seen in (15) and (16), the demand for each good reduces to being a fixed proportion of income with a proportionality factor that depends only on own price.

Suppose the market protocol governing market behavior in this econ-

omy is as follows: Firm 1 and Firm 2 simultaneously announce prices p_1 and p_2 , promising to meet any forthcoming demands for their goods from the consumer as long as this can be done with nonnegative profits. A strategic interaction problem then arises for each profit-maximizing firm, because the profits of each firm depend on the price set by the other firm. Specifically, the profit function of Firm 1 takes the form

$$\pi_1(p_1, p_2) = [p_1 - c_1]D_1(p_1, p_2) \quad , \quad (17)$$

and the profit function of Firm 2 takes the form

$$\pi_2(p_1, p_2) = [p_2 - c_2]D_2(p_1, p_2) \quad . \quad (18)$$

Suppose the values taken on by the parameters appearing in the above-described model economy are as follows:

$$b_1 = 1/2; \quad b_2 = 1/2; \quad I = 1; \quad c_1 = 0; \quad c_2 = 0 \quad . \quad (19)$$

Suppose, also, that each firm can set its goods price at only one of two possible values, low L or high H , where

$$L = 1/2 \quad \text{and} \quad H = 3/4 \quad . \quad (20)$$

In this case, the survival condition (14) holds for all possible price choices by the firms. Consequently, the consumer demand functions (15) and (16) reduce to the following particular forms:

$$q_1^d = 1/4 + [2 - p_2]/4p_1 = D_1(p_1, p_2) \quad ; \quad (21)$$

$$q_2^d = 1/4 + [2 - p_1]/4p_2 = D_2(p_1, p_2) \quad . \quad (22)$$

It follows that the consumer demands $D_1(p_1, p_2)$ faced by Firm 1 for all possible settings of the prices p_1 and p_2 are as follows:

$$\begin{aligned} D_1(L, L) &= 1 ; \\ D_1(L, H) &= 7/8 ; \\ D_1(H, L) &= 3/4 ; \\ D_1(H, H) &= 2/3 . \end{aligned}$$

Similarly, the consumer demands $D_2(p_1, p_2)$ faced by Firm 2 for all possible settings of the prices p_1 and p_2 are as follows:

$$\begin{aligned} D_2(L, L) &= 1 ; \\ D_2(L, H) &= 3/4 ; \\ D_2(H, L) &= 7/8 ; \\ D_2(H, H) &= 2/3 . \end{aligned}$$

Using (17) and (18), the profit payoff matrix faced by the two firms then takes the form depicted in Table 1. (For clarity, each profit level has been multiplied by 48 so that profits are represented as whole number rather than as fractions.) The first number in each reported payoff pairing denotes the profit for Firm 1 and the second number denotes the profit for Firm 2.

A *Nash equilibrium* for Firm 1 and Firm 2 is any pair (p_1^*, p_2^*) of pricing strategies such that, given p_1^* , Firm 2 has no incentive to deviate from p_2^* , and given p_2^* , Firm 1 has no incentive to deviate from p_1^* .

As seen in Table 1, there exists a unique Nash equilibrium for the model economy at hand: namely, the pricing strategy pair (H, H) .

		Firm 2	
		L	H
Firm 1	L	(24,24)	(21,27)
	H	(27,21)	(24,24)

Table 1. Profit Payoff Matrix for Firms
(Normalized Values = Multiplied by 48)

In fact, an even stronger property holds for (H,H) . The pricing strategy H constitutes a *dominant* pricing strategy for each firm in the sense that H is the best price for each firm to set in response to *any* price set by its rival.

Moreover, (H,H) is also Pareto efficient, meaning there is no other feasible pricing strategy pair under which both firms would be at least as well off and at least one would be *strictly* better off. Indeed, (L,L) , (L,H) , and (H,L) are Pareto efficient as well. Thus, in the present example, every possible firm strategy combination leads to a Pareto efficient outcome when measured solely in terms of firm profits.

Note that Firm 1 would actually prefer the profit payoff it obtains under the pricing strategy pair (H,L) to the profit payoff it obtains under the pricing strategy pair (H,H) , but it has no power to enforce this outcome: the profit payoff for Firm 2 under (H,L) is worse than

the profit payoff for Firm 2 under (H,H) . And similarly for Firm 2 with regard to (L,H) . However, restricting attention purely to firm profits alone, there is no strategy combination under which *coordination failure* is exhibited, in the sense it is a Pareto-dominated Nash equilibrium.

Note, also, that the high-price outcome (H,H) might appear to be a “price collusion” outcome in which the two firms have conspired to raise the price as high as possible for the consumer. However, as discussed above, the two firms actually could reason their way to this high-price outcome without any explicit communication passing between them. In this case, under current U.S. antitrust law, the high prices would be perfectly legal since a judgement of price collusion requires a demonstration that firms have actively conspired together to fix prices.

Various open questions remain. For example, is (H,H) *always* a dominant (hence Nash) pricing strategy pair for the game at hand, regardless of the particular parameter values? Is there always a Nash equilibrium (even if not a dominant pricing strategy pair), regardless of parameter values? If so, is this Nash equilibrium always Pareto efficient? Could there be a Nash equilibrium which is Pareto dominated, for example, in the sense that there is another feasible pricing strategy pair that yields at least as much profit for both firms and strictly more profit for at least one?

And what about the welfare of the consumer in all of this? In particular, what utility level would be achieved by the *consumer* under each of the four possible strategy combinations for the two

firms? Given definition (10) for the consumer's utility function, as well as the calculated consumer demands, it can be shown that the consumer's achieved utility level under each possible combination of firm strategies is as depicted in Table 2.

		Firm 2	
		L	H
Firm 1	L	-28.8	-69.6
	H	-69.6	-74.70

Table 2. Consumer Welfare (Attained Utility Level)
Under Each Firm Strategy Combination
(Normalized Values = Multiplied by 48)

Finally, define *total net surplus* for this economy to be the sum of attained consumer utility AND total attained firm profits (where for simplicity it is assumed that utility is in dollar terms so that utility and profits can be added together). Which of the four possible strategy combinations for the two firms would achieve the highest SOCIAL welfare in the sense that TOTAL net surplus is maximized? As shown in Table 3, the answer is (L,L) .

Consequently, although the Nash equilibrium strategy combination

		Firm 2	
		L	H
Firm 1	L	19.2	-21.16
	H	-21.6	-26.7

Table 3. Social Welfare (Total Net Surplus)
Under Each Firm Strategy Combination
(Normalized Values = Multiplied by 48)

(H,H) is not Pareto-dominated in terms of firm profits, it does not result in the highest possible social welfare (total net surplus).

Interestingly, if the firms could be induced (or forced) to commit to (L,L) , they would achieve the same profits as at (H,H) yet social welfare would also be maximized. In the absence of any commitment or enforcement mechanism, however, how likely is it that the firms in this example – free to choose their own prices and caring only about maximizing their own profits – will actually end up choosing the strategy combination (L,L) that maximizes social welfare?

Final Cautionary Note:

How robust are the above SOCIAL welfare findings to changes in the consumer's utility function?

For example, suppose that the NATURAL logarithm is used instead of the BASE 10 logarithm? Or suppose that the utility function is simply multiplied by a positive constant?

Both of these utility transformations are *ordinal* in the sense that they preserve the ORDER of consumer preferences: that is, if a consumption bundle A is preferred to another consumption bundle B under the original utility function, then A is still preferred to B under the transformed utility function. However, the transformation changes the INTENSITY with which A is preferred to B.

By construction, an ordinal transformation of the consumer's utility function does not affect the consumer's preference order over consumption bundles. However, it can have a dramatic effect on social welfare outcomes when these outcomes are measured in terms of total net surplus, as above. This is because total net surplus requires an adding up of surplus across individual firms and consumers (i.e. it requires inter-agent welfare comparisons), so the precise way individual firm and consumer surplus measures are constructed can substantially affect the resulting social welfare conclusions for the economy as a whole.

For example, suppose under alternative parameter settings that the two firms achieve strictly greater total profits under (H,H) than under (L,L) whereas the consumer achieves a higher utility level (10)

under (L,L) than under (H,H). Suppose the resulting total net surplus, i.e., the sum of firm profits and consumer utility (assumed to be in dollar units), is higher under (L,L) than under (H,H), as in the example above. If you then multiply the consumer’s utility function (10) by a sufficiently small constant $k > 0$ — a “positive linear transformation” that leaves unaffected the representation of the comparative preferences of the consumer (what is preferred to what) — you can make the resulting total net surplus higher at (H,H) than at (L,L) because the gain in firm profits now outweighs the loss in consumer utility.

4. More General Observations on Price Discovery

The robustness of competitive pricing theory to changes in its basic assumptions can be tested by considering three basic questions.

Question 1: How might strategic interaction become important if firms set prices for their inputs and outputs?

As illustrated in the previous section, firms’ actions become strategically linked together if they understand and exploit the fact that the demand for their outputs and the distribution of labor and capital rental services across firms depend on the prices they set as well as on the prices set by other firms. For example, firm X might be able to bid away labor services from firm Y by offering a higher wage than the wage set by firm Y ; and similarly with regard to attracting an increased supply of capital services and an increased output demand. Consequently, the attraction and retention of service suppliers and

customers now involves a careful consideration of the pricing strategies of other firms.

Question 2: How might expectations and learning rules become important if firms set prices for their inputs and outputs?

Realistically, firms would not have costless access to complete and correct information regarding the supply and demand functions they face for inputs and outputs, information that is critical for the price-setting process. In this case, firms would face a “dual control” problem at each point in time in the sense that each firm would have two potentially conflicting objectives:

Information exploitation: Set prices in an attempt to ensure that total profits are as high as possible, conditional on the firm’s current information regarding the supply and demand functions it faces.

Information exploration: Set prices in an attempt to learn more about the supply and demand functions faced by the firm, so that *future* profits can be increased even if these learning efforts lead to lower current profits.

The situation is further complicated by the fact that the supplies and demands for a firm’s inputs and outputs depend not only on its own prices but also on the prices set by other firms. Indeed, past prices will also affect the firm’s current supplies and demands to the extent that these past prices affect current consumer budget constraints and search behavior.

Also, firms can now offer different wage rates to observationally equivalent workers and different goods prices to observationally distinct consumers for units of the same type of good. Consequently, it is highly unrealistic to assume that consumers can costlessly acquire complete and correct information regarding the wage rates and goods prices they face. Rather, consumers would presumably have to undertake some form of sequential search to learn about the current distribution of wage rates and goods prices.

Presumably, however, this search would involve opportunity costs for consumers in terms of delayed consumption and foregone wages. Consequently, consumers might decide to sample only a small fraction of the available wage offers and goods prices and then to accept the highest wage offer and lowest goods prices found to date instead of carrying out a complete global search of all possibilities. The rule by which a consumer decides to stop sequential sampling is called a *sequential stopping rule* in the statistical decision theory literature. When consumers use stopping rules, a nondegenerate distribution of wage rates can exist in “equilibrium” for a single type of labor, and a nondegenerate distribution of prices can exist in “equilibrium” for a single type of good. Moreover, consumers deciding to supply labor to or buy goods from a firm in some time period may simply decide to stick with this firm in future time periods without engaging in more search (habit, brand effects,...). These considerations can further complicate the strategic price-setting rivalry among firms.

Question 3: How might bankruptcy rules, rationing rules, and

inventory management become important if firms set prices for their inputs and outputs?

Markets are no longer guaranteed to clear at the levels of *planned* supplies and demands, since wage rates and goods prices might be “wrong.” Some type of formally or informally established rationing rule is needed to determine who gets what if the planned demand for a good happens to exceed its supply. If this situation arises frequently, a firm might want to institute an inventory plan so that excess demand can be satisfied out of inventory.

In the reverse case of supply exceeding demand, unintended inventories arise, and the firm would presumably want to take this possibility into consideration when making its price and quantity decisions. Even with an inventory plan, the possibility of bankruptcy arises for a firm if it cannot sell all it produces; for the firm might then be unable to fulfill its obligations with respect to wage and capital rental payments. If bankruptcy occurs, some type of formally or informally established rule is needed to determine how the assets that the firm still possesses get divided among its various creditors.

5. Brief Literature Review on Price-Setting Agents

Although a *definitive* theory regarding the working of decentralized market economies with price-setting agents is currently lacking, there is a large and growing body of empirical, analytical, and experimental evidence that bears on this issue.

Empirical studies generally confirm that the markup of price over

the cost of production is a key focus of many firms. Increases in input costs tend to be passed through to output prices either in whole or in part as firms attempt to maintain their markups. For example, Blinder et al. (1998) surveyed 200 U.S. firms between April 1990 and March 1992 regarding factors they considered to be relevant to their price-setting decisions. About half of the surveyed firms responded that they changed prices at most once a year, primarily in response to input cost changes. The firms also cited competitive pressures from rival firms, material and labor costs of changing prices, risk of antagonizing customers, and loss of managerial time as additional factors important for their pricing decisions.

Willis (2003) examines how structural changes in the U.S. economy over the past two decades – spurred on, in particular, by advances in information technology – might have influenced the price-setting behavior of firms and the behavior of aggregate inflation. For example, as reported in a *Wall Street Journal* article (September 18, 2002), new pricing methods facilitated by improved information technology include the creation of company pricing teams that track regional trends in the prices of competitors as well as international surveys of customers to determine their willingness to pay for an item.

Regarding wage setting, it is a puzzling empirical fact in the U.S. that recessions increase unemployment and retard the average growth rate of nominal wages and salaries but that cuts in nominal wages and salaries are rarely observed. In an attempt to understand this puzzle, Bewley (1999) conducted extensive interviews with over 300 business people, labor leaders, counselors of the unemployed, and business

consultants in the northeast U.S. during the early 1990s. He finds that social considerations outside of conventional economic theory strongly influence the wage-setting process. Employers want workers to identify with company goals and to cooperate with co-workers and supervisors. What restrains employers from cutting worker pay even during recessions is the strong belief that pay cuts hurt worker morale and increase labor turnover. The advantage of layoffs over pay reductions is that they “get the misery out the door.” Based on these findings, Bewley argues that firms should be viewed as communities, not simply as profit maximizers, because considerations such as trust, enthusiasm, and commitment are needed to make firms function well.

Bewley (1999) also provides a comprehensive survey of existing theories of the wage-setting process. He notes that his empirical findings support none of these theories, apart from those which emphasize the impact of pay cuts on morale. He concludes that “despite the considerable resources required, it is well worth the effort to get out of the (academic) office and face economic reality rather than invent it.”

The theoretical industrial organization (IO) literature on price setting largely focuses on monopolistic (single seller) and oligopolistic (multiple seller) markets with price-setting firms. See, for example, Pepall et al. (1999) and Varian (1992, Chapters 14–16). A wide variety of interesting topics have been addressed, including product tie-ins and price discrimination, predatory pricing, limit pricing, collusion, mergers, and research and development.

In the theoretical IO literature, however, analytical tractability

considerations often require the imposition of strong restrictions on the beliefs, behaviors, and interaction patterns of the firms. For example, in the well-studied *Bertrand model*, firms producing an identical good simultaneously choose prices in a one-shot game. It is assumed that all firms believe that the firm with the lowest price will capture all of the market, and that each firm will choose a price that is a “best response” to the anticipated price choices of its rival firms.

The strong restrictions on beliefs, behaviors, and interaction patterns imposed in the theoretical IO literature can be problematic from the viewpoint of obtaining realistic microfoundations, since sellers in actual real-world markets typically have a good deal more autonomy and flexibility.

In addition, buyers need not be as passive as the theoretical IO literature typically presumes. Rather, buyers in some market contexts can coordinate their quantity demands and/or their price offers to form countervailing buyer market power against seller market power, a tactic that can be very effective. See, for example, Engle-Warnick and Ruffle (2002) and Nicolaisen et al. (2001).

The theoretical literature focusing specifically on auctions is also quite extensive; see Klemperer (2000,2002a). Although some work has been done on double auctions, much of the literature to date has focused on the more tractable case of one-sided posted offer auctions.

The basic type of auction considered in this theoretical literature is a seller posted-offer auction consisting of a seller with a single item to sell and a fixed set of expected-payoff maximizing buyers who in-

independently submit bids for the seller's item during the course of a single trading period. The buyers typically are assumed to have either private valuations (reservation prices) for the item known with certainty, private valuations for the item drawn from a common probability distribution, or a common but unknown valuation for the item that must be separately deduced by each buyer from private information.

The key issues emphasized in this theoretical auction literature to date have been the effects of relaxing one of more assumptions of this basic model, e.g., though the introduction of information correlation, aversion to risk, or budget constraints. Attention is generally focused on the existence and uniqueness of equilibria in the form of various Nash equilibrium refinements. As Klemperer (2002a,b) cautions, although this work is important and interesting, what really matters for *practical* auction design are more traditional IO concerns: preventing collusion, predatory pricing, and entry-detering behavior.

There is also an extensive body of industrial organization laboratory research work with human subjects focusing in part on pricing behaviors in markets; see, for example, Holt (1995) and Smith (1989,1994). Particular attention has been focused on the behavior of prices in one-sided and double auctions. Although results overall are mixed, in experiments in which the market structure is kept fairly simple — e.g., a single market with a unique market clearing point — double auction pricing mechanisms have been regularly observed to result in convergence to a market clearing outcome. Thus, competitive pricing theory does have predictive content in certain experimental settings

even when its behavioral assumptions (in particular, price taking) are known to be false.

For a seminal paper investigating research and development effort under two types of industrial structures (Cournot and Bertrand) by means of systematic and comprehensive computational experiments, see Quirnbach (1993).

A growing number of researchers are now studying a wide variety of industrial structures with price-setting agents by means of agent-based computational economics (ACE) experiments; see Tesfatsion (2002). One key issue addressed is how imperfectly informed buyers and sellers co-adapt and co-evolve their market behaviors over time under alternative institutional arrangements. A more ambitious goal is to understand how market behaviors and institutions co-evolve *together* over time.

Some ACE researchers have assumed that buyers and/or sellers set their price offers in accordance with simple learning algorithms that incrementally adjust these price offers on the basis of past profit earnings. An example is the *derivative follower (DF)* algorithm used by Greenwald and Kephart (1999) for the design of shopbots and pricebots on the Internet. The DF algorithm is a local search algorithm that starts from a user-specified price level and then incrementally changes price in the same direction as long as profits increase. Whenever a drop in profits occurs, the direction of price change is reversed.

Other researchers [e.g., Nicolaisen et al. (2001)] have implemented reinforcement learning algorithms in which the probability of a price

offer being selected by a buyer or seller at any given time is proportional to its average relative profitability in comparison with other price offers selected in the past. Various alternative specifications of learning representations for buyers and sellers will be taken up more carefully in later lectures.

The computational experiments conducted by ACE researchers suggest that convergence to a market clearing outcome can sometimes be ensured if the learning processes of agents take particular forms. Another observed possibility, however, is that markets can tend toward states exhibiting *coordination failure*. That is, markets can tend toward states which are Nash equilibria, so that no individual market participant perceives any incentive to change their current price offer unilaterally, but which are also Pareto dominated, meaning that a better state could feasibly be obtained if the current market participants would only undertake coordinated changes in their current price offers.

6. Concluding Remarks

Real-world decentralized market economies have coevolved an enormous variety of market institutions (legal systems, credit systems, bankruptcy rules,...) that help to prevent market breakdowns. Given these supporting institutions, it does appear that many decentralized market economies are able to achieve, on average, a rough balancing of supplies and demands for goods, services, and financial assets despite the fact that prices are being set by individual buyers and

sellers without any overall centralized control.

More generally, many decentralized market economies are able to stay approximately coordinated over time, in the sense that inflation and unemployment do not get wildly out of control and a positive GDP growth rate is sustained on average. The accomplishment of such feats is incontrovertible evidence that economists have a lot left to learn about the market institutions underlying the determination of prices in decentralized market economies.

References

- [1] Kenneth Arrow (1959), “Towards a Theory of Price Adjustment,” in M. Abramovitz (ed.), *Allocation of Economic Resources*, Stanford University Press, Stanford.
- [2] Truman Bewley (1999), *Why Wages Don't Fall During a Recession*, Harvard University Press, Cambridge, MA.
- [3] Alan Blinder, Elie Canetti, David Lebow, and Jeremy Rudd (1998), *Asking About Prices*, Russell Sage Foundation, New York, N.Y.
- [4] Jim Engle-Warnick and Bradley J. Ruffle (2002), “Buyer Countervailing Power versus Monopoly Power: Evidence from Experimental Posted-Offer Markets,” Working Paper, Ben Gurion University, Israel, May. <https://econ.bgu.ac.il/papers/154.pdf>
- [5] Amy M. Greenwald and Jeffrey O. Kephart (1999), “Shopbots and Pricebots,” pp. 506–511 in *Proceedings of the 16th Joint Conference on Artificial Intelligence (IJCA-99)*, August.
- [6] Charles A. Holt (1995), “Industrial Organization: A Survey of Laboratory Research,” pp. 349–443 in John H. Kagel and Alvin E. Roth, *Handbook of Experimental Economics*, Princeton University Press, Princeton, N.J.
- [7] Paul Klemperer (2000), “Auction Theory: A Guide to the Literature,” pp. 3–62 in Paul Klemperer, *The Eco-*

- conomic Theory of Auctions*, Edward Elgar, Cheltenham, UK.
<https://www.paulklempere.org>
- [8] Paul Klemperer (2002a), “What Really Matters in Auction Design,” *Journal of Economic Perspectives*, Vol. 16, No. 1, 169–189. <https://www.paulklempere.org>
- [9] Paul Klemperer (2002b), “Using and Abusing Economic Theory,” Alfred Marshall Lecture to the European Economic Association, December 2002. <https://www.paulklempere.org/>
- [10] James Nicolaisen, Valentin Petrov, and Leigh Tesfatsion (2001), “Market Power and Efficiency in a Computational Electricity Market with Discriminatory Double-Auction Pricing,” *IEEE Transactions on Evolutionary Computation*, Vol. 5, No. 5, 437–441. <https://www2.econ.iastate.edu/tesfatsi/mpeieee.pdf>
- [11] Lynne Pepall, Daniel J. Richards, and George Norman (1999), *Industrial Organization: Contemporary Theory and Practice*, South-Western College Publishing, New York, N.Y.
- [12] Herman C. Quirnbach (1993), “R&D: Competition, Risk, and Performance,” *Rand Journal of Economics* 24(2), Summer, 157-197.
- [13] Vernon L. Smith (1989), “Theory, Experiments, and Economics,” *Journal of Economic Perspectives* Vol. 3, 151–169.
- [14] Vernon L. Smith (1994), “Economics in the Lab,” *Journal of Economic Perspectives* Vol. 8, Winter, 113–131.

- [15] Tesfatsion, L. (2002), “Agent-Based Computational Economics: Growing Economies from the Bottom Up,” *Artificial Life*, Vol. 8(1), 55–82. <https://www2.econ.iastate.edu/tesfatsi/acealife.pdf>
- [16] Hal R. Varian (1992), *Microeconomic Analysis*, Third Edition, W. W. Norton and Company, New York, N.Y.
- [17] Jonathan L. Willi (2003), “Implications of Structural Changes in the U.S. Economy for Pricing Behavior and Inflation Dynamics,” *Economic Review*, Federal Reserve Bank of Kansas City, Vol. 88, No. 1, First Quarter, 5–26.