

Price Theory in Economics

Thomas A. Weber*

Forthcoming in: Özer, Ö., Phillips, R. (Eds.) *Oxford Handbook of Pricing Management*, Oxford University Press.

Contents

1	Origin of Value and Prices	1
2	Price-Taking Behavior and Choice	3
2.1	Rational Preferences	3
2.2	Utility Maximization	4
2.3	Comparative Statics	7
2.4	Effects of Uncertainty	9
3	Price Discovery in Markets	10
3.1	Pure Exchange	10
3.2	Competitive Markets	13
3.3	Complete Markets	14
3.4	Incomplete Markets	16
4	Disequilibrium and Price Adjustments	18
4.1	Walrasian Tâtonnement	18
4.2	Information Transmission in Markets	20
5	Externalities and Nonmarket Goods	22
5.1	Externalities	22
5.2	Nonmarket Goods	26
6	Strategic Pricing with Complete Information	28
6.1	Monopoly Pricing	28
6.2	Oligopoly Pricing	30
6.3	Entry Deterrence and Limit Pricing	33
7	Strategic Pricing with Incomplete Information	34
7.1	Screening	34
7.2	Signaling	37
7.3	Mechanism Design	39
7.4	Oligopoly Pricing with Asymmetric Information	42
8	Dynamic Pricing	44
9	Behavioral Anomalies	45
10	Open Issues	47

*Department of Management Science and Engineering, 344 Huang Engineering Center, Stanford University, Stanford, CA 94305-4121, U.S.A. E-mail: webert@stanford.edu.

*And there is all the difference in the world
between paying and being paid.
The act of paying is perhaps
the most uncomfortable infliction (...)
But being paid, —what will compare with it?
The urbane activity with which man receives money
is really marvellous.*
HERMAN MELVILLE, MOBY DICK

1 Origin of Value and Prices

Price theory is concerned with explaining economic activity in terms of the creation and transfer of value, which includes the trade of goods and services between different economic agents. A puzzling question addressed by price theory is, for example: why is water so cheap and diamonds are so expensive, even though water is critical for survival and diamonds are not? In a discussion of this well-known ‘Diamond-Water Paradox,’ Adam Smith (1776) observes that

[t]he word value, it is to be observed, has two different meanings, and sometimes expresses the utility of some particular object, and sometimes the power of purchasing other goods which the possession of that object conveys. The one may be called “value in use;” the other, “value in exchange.” (p. 31)

For him, diamonds and other precious stones derive their value from their relative scarcity and the intensity of labor required to extract them. Labor therefore forms the basic unit of the exchange value of goods (or ‘items’), which determines therefore their ‘real prices.’ The ‘nominal price’ of an item in Smith’s view is connected to the value of the currency used to trade it and might therefore fluctuate. In this labor theory of value the Diamond-Water Paradox is resolved by noting that it is much more difficult, in terms of labor, to acquire one kilogram of diamonds than one kilogram of water.

About a century later, the work of Carl Menger, William Stanley Jevons, and Léon Walras brought a different resolution of the Diamond-Water Paradox, based on marginal utility rather than labor. Menger (1871) points out that the value of an item is intrinsically linked to its utility ‘at the margin.’ While the first units of water are critical for the survival of an individual, the utility for additional units quickly decreases, which explains the difference in the value of water and diamonds. Commenting on the high price of pearls, Jevons (1881) asks “[d]o men dive for pearls because pearls fetch a high price, or do pearls fetch a high price because men must dive in order to get them?” (p. 102), and he concludes that “[t]he labour which is required to get more of a commodity governs the supply of it; the supply determines whether people do or do not want more of it eagerly; and this eagerness of want or demand governs value” (p. 103). Walras (1874/77) links the idea of price to the value of an object in an exchange economy by noting that the market price of a good tends to increase as long as there is a positive excess demand, while

it tends to decrease when there is a positive excess supply. The associated adjustment process is generally referred to as Walrasian *tâtonnement* (“groping”). Due to the mathematical precision of his early presentation of the subject, Walras is generally recognized as the father of general equilibrium theory.¹

To understand the notion of price it is useful to abstract from the concept of money.² In a barter where one person trades a quantity x_1 of good 1 for the quantity x_2 of good 2, the ratio x_1/x_2 corresponds to his price paid for good 2. If apples correspond to good 1 and bananas to good 2, then the ratio of the number of apples paid to the number of bananas obtained in return corresponds to the (average) price of one banana, measured in apples. The currency in this barter economy is denominated in apples, so that the latter is called the *numéraire* good, the price of which is normalized to one.

The rest of this survey, which aims at providing a compact summary of the (sometimes technical) concepts in price theory, is organized as follows. In Section 2, we introduce the concepts of “rational preference” and “utility function” which are standard building blocks of models that attempt to explain choice behavior. We then turn to the frictionless interaction of agents in markets. Section 3 introduces the notion of a Walrasian equilibrium, where supply equals demand and market prices are determined (up to a common multiplicative constant) by the self-interested behavior of market participants. This equilibrium has remarkable efficiency properties, which are summarized by the first and the second fundamental welfare theorems. In markets with uncertainty, as long as any desired future payoff profile can be constructed using portfolios of traded securities, the Arrow-Debreu equilibrium directly extends the notion of a Walrasian equilibrium and inherits all of its efficiency properties. Otherwise, when markets are “incomplete,” as long as agents have “rational expectations” in the sense that they correctly anticipate the formation of prices, the Radner equilibrium may guarantee at least constrained economic efficiency. In Section 4 we consider the possibility of disequilibrium and Walrasian *tâtonnement* as a price-adjustment process in an otherwise stationary economy. Section 5 deals with the problem of “externalities,” where agents’ actions are payoff-relevant to other agents. The presence of externalities in markets tends to destroy the efficiency properties of the Walrasian equilibrium and even threaten its very existence. While in Sections 3 and 4 all agents (including consumers and firms) are assumed to be “price takers,” we consider strategic interactions between agents in Sections 6 and 7, in the presence of complete and incomplete information, respectively. The discussion proceeds from optimal monopoly pricing (which involves the problems of screening, signaling, and, more generally, mechanism design when information is incomplete) to price competition between several firms, either in a level relationship when there are several oligopolists in a market, or as an entry problem, when one incumbent can deter (or encourage) the entrance of other firms into the market. Section 8 deals with dynamic pricing issues, and in Section 9 we mention some of the persistent behavioral irregularities that are not well captured by classical price theory. Finally, Section 10 concludes and provides a number of directions from which further research contributions may be expected.

¹The modern understanding of classical general equilibrium theory is well summarized by Debreu’s (1959) concise axiomatic presentation and by Arrow and Hahn’s (1971) more complete treatise. Mas-Colell (1985) provides an overview from a differentiable viewpoint, and McKenzie (2002) a more recent account of the theory. Friedman (1962/2007), Stigler (1966), and Hirshleifer et al. (2005) present ‘price theory’ at the intermediate level.

²Keynes’ (1936) theory of liquidity gives some reasons for the (perhaps somewhat puzzling) availability of money, which, after all, cannot be directly consumed, but provides a fungible means of compensation in exchange.

2 Price-Taking Behavior and Choice

Normative predictions about what agents do, i.e., about their “choice behavior,” requires some type of model. In Section 2.1, we introduce preferences and the concept of a utility function to represent those preferences. Section 2.2 then presents the classical model of consumer choice in terms of a “utility maximization problem” in an economy where agents take the prices of the available goods as given. In Section 2.3, we examine how choice predictions depend on the given prices or on an agent’s wealth, an analysis which is referred to as “comparative statics.” In reality it is only rarely possible to make a choice which achieves a desired outcome for sure. The effects of uncertainty, discussed in Section 2.4, are therefore important for our understanding of how rational agents behave in an economy. Decision problems faced by two typical agents, named Joe and Melanie, will serve as examples.

2.1 Rational Preferences

An agent’s preferences can be expressed by a partial order over a choice set X .³ For example, consider Joe’s preferences over the choice set $X = \{\text{Apple, Banana, Orange}\}$ when deciding which fruit to pick as a snack. Assuming that he prefers an apple to a banana and a banana to an orange, his preferences on X thus far can be expressed in the form

$$\begin{aligned}\text{Apple} &\succeq \text{Banana}, \\ \text{Banana} &\succeq \text{Orange},\end{aligned}$$

where \succeq denotes “is (weakly) preferred to.” However, these preferences are not complete, since they do not specify Joe’s predilection between an apple and an orange. If Joe prefers an orange to an apple, then the relation

$$\text{Orange} \succeq \text{Apple}$$

completes the specification of Joe’s preference relation \succeq which is then defined for all pairs of elements of X . When Joe is ambivalent about the choice between an apple and a banana, so that he both prefers an apple to a banana (as noted above) *and* a banana to an apple (i.e., $\text{Apple} \succeq \text{Banana}$ and $\text{Banana} \succeq \text{Apple}$ both hold), we say that he is *indifferent* between the two fruits and write $\text{Apple} \sim \text{Banana}$. On the other hand, if $\text{Apple} \succeq \text{Banana}$ holds but $\text{Banana} \succeq \text{Apple}$ is not true, then Joe is clearly not indifferent: he *strictly* prefers an apple to a banana, which is denoted by

$$\text{Apple} \succ \text{Banana}.$$

If the last relation holds true, a problem arises because Joe’s preferences are now such that he strictly prefers an apple to a banana, weakly prefers a banana to an orange, and at the same time weakly prefers an orange to an apple. Thus, Joe would be happy to get an orange in exchange for an apple. Then he would willingly take a banana for his orange, and, finally, pay a small amount of money (or a tiny piece of an apple) to convert his banana back into an apple. This cycle, generated by the intransitivity of his preference relation, leads to difficulties when trying to describe Joe’s behavior as rational.⁴

³Fishburn (1970) and Kreps (1988) provide more in-depth overviews of choice theory.

⁴When aggregating the preferences of a society of at least three rational agents (over at least three items), Arrow’s (1951) seminal ‘impossibility theorem’ states that, no matter what the aggregation procedure may be, these cycles can in general be avoided only by declaring one agent a dictator or impose rational societal preferences

DEFINITION 1 A rational preference \succeq on the choice set X is a binary relationship defined for any pair of elements of X , such that for all $x, y, z \in X$: (i) $x \succeq y$ or $y \succeq x$ (Completeness), (ii) $x \sim x$ (Reflexivity), and (iii) $x \succeq y$ and $y \succeq z$ together imply that $x \succeq z$ (Transitivity).

To make predictions about Joe's choice behavior over complex choice sets, dealing directly with the rational preference relation \succeq is from an analytical point of view unattractive, as it involves many pairwise comparisons. Instead of trying to determine all 'undominated' elements of Joe's choice set, i.e., all elements that are such that no other element is strictly preferred, it would be much simpler if the magnitude of Joe's liking of each possible choice $x \in X$ was encoded as a numerical value of a 'utility function' $u(x)$, so that Joe's most preferred choices also maximize his utility.

DEFINITION 2 A utility function $u : X \rightarrow \mathbb{R}$ represents the rational preference relation \succeq (on X) if for all $x, y \in X$:

$$x \succeq y \quad \text{if and only if} \quad u(x) \geq u(y).$$

It is easy to show that as long as the choice set is finite, there always exists a utility representation of a rational preference relation on X .⁵ In addition, if u represents \succeq on X , then for any increasing function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ the utility function $v : X \rightarrow \mathbb{R}$ with $v(x) = \varphi(u(x))$ for all $x \in X$ also represents \succeq on X .

2.2 Utility Maximization

If Joe has a rational preference relation on X that is represented by the utility function u , in order to predict his choice behavior it is enough to consider solutions of his *utility maximization problem* (UMP),

$$x^* \in \arg \max_{x \in X} u(x). \quad (1)$$

Let us now think of Joe as a consumer with a wealth w that can be spent on a bundle $x = (x_1, \dots, x_L)$ containing nonnegative quantities x_l , $l \in \{1, \dots, L\}$, of the L available consumption

from the outside. For example, if one chooses pairwise majority voting as aggregation procedure, cycles can arise easily, as can be seen in the following well-known voting paradox, which dates back to Condorcet (1785). Consider three agents with rational preferences relations (\succeq_1 , \succeq_2 , and \succeq_3 , respectively) over elements in the choice set $X = \{A, B, C\}$ such that $A \succ_1 B \succ_1 C$, $B \succ_2 C \succ_2 A$, and $C \succ_3 A \succ_3 B$. However, simple majority voting between the different pairs of elements of X yields a societal preference relation \succeq such that $A \succ B$, $B \succ C$, and $C \succ A$ implying the existence of a 'Condorcet cycle,' i.e., by Definition 1 the preference relation \succeq is intransitive and thus not rational.

⁵If X is not finite, it may be possible that no utility representation exists. Consider for example lexicographic preferences defined on the two-dimensional choice set $X = [0, 1] \times [0, 1]$ as follows. Let $(x_1, x_2), (\hat{x}_1, \hat{x}_2) \in X$. Suppose that $(x_1, x_2) \succeq (\hat{x}_1, \hat{x}_2) \stackrel{\text{def}}{\iff} (x_1 > \hat{x}_1) \text{ or } (x_1 = \hat{x}_1 \text{ and } x_2 \geq \hat{x}_2)$. For example, when comparing used Ford Mustang cars, the first attribute might index a car's horsepower and the second its color (measured as proximity to red). An individual with lexicographic preferences would always prefer a car with more horsepower. However, if two cars have the same horsepower, the individual prefers the model with the color that is closer to red. The intuition why there is no utility function representing such a preference relation is that for each fixed value of x_1 , there has to be a finite difference between the utility for $(x_1, 0)$ and $(x_1, 1)$. But there are more than countably many of such differences, which when evaluated with any proposed utility function can be used to show that the utility function must in fact be unbounded on the square, which leads to a contradiction. For more details, see Kreps (1988). In practice, this does not lead to problems, since it is usually possible to discretize the space of attributes, which results in finite (or at least countable) choice sets.

goods in the economy. We assume henceforth that Joe (and any other agent we discuss) strictly prefers more of each good,⁶ so that his utility function is increasing in x . To simplify further, we assume that Joe takes the price p_l of any good l as given, i.e., he is a *price taker*. Given a price vector $p = (p_1, \dots, p_L)$ he therefore maximizes his utility $u(x)$ subject to the constraint that the value of his total consumption, equal to the dot-product $p \cdot x$, does not exceed his total (positive) wealth w . In other words, all feasible consumption bundles lie in his *budget set*

$$B(p, w) = \{x \in \mathbb{R}_+^L : p \cdot x \leq w\}.$$

Joe's so-called *Walrasian demand* correspondence is

$$x(p, w) = \arg \max_{x \in B(p, w)} u(x). \quad (2)$$

Depending on the uniqueness of the solutions to Joe's UMP, the Walrasian demand correspondence may contain more than one value. Optimality conditions for this constrained optimization problem can be obtained by introducing the Lagrangian

$$\mathcal{L}(x, \lambda, \mu; p, w) = u(x) - \lambda(p \cdot x - w) + \mu x,$$

where $\lambda \in \mathbb{R}_+$ is the Lagrange multiplier associated with the inequality constraint $p \cdot x \leq w$ and $\mu = (\mu_1, \dots, \mu_L) \in \mathbb{R}_+^L$ the Lagrange multiplier associated with the nonnegativity constraint $x \geq 0$. The necessary optimality conditions are

$$\frac{\partial \mathcal{L}(x, \lambda, \mu; p, w)}{\partial x_l} = \frac{\partial u(x)}{\partial x_l} - \lambda p_l + \mu_l = 0, \quad l \in \{1, \dots, L\}. \quad (3)$$

Together with the complementary slackness conditions

$$\lambda(p \cdot x - w) = 0 \quad (4)$$

and

$$\mu_l x_l = 0, \quad l \in \{1, \dots, L\}, \quad (5)$$

they can be used to construct Joe's Walrasian demand correspondence. For this, we first note that since Joe's utility function is increasing in x , the budget constraint is binding, i.e., $p \cdot x = w$, at the optimum (Walras' Law). In particular, if Joe consumes positive amounts of all commodities, i.e., if $x \gg 0$,⁷ the complementary slackness condition (5) implies that $\mu_l = 0$ for all $l \in \{1, \dots, L\}$, so that by (3) (with $\lambda > 0$) we obtain that

$$\text{MRS}_{lj}(x) = \frac{\frac{\partial u(x)}{\partial x_l}}{\frac{\partial u(x)}{\partial x_j}} = \frac{p_l}{p_j}, \quad l, j \in \{1, \dots, L\}, \quad (6)$$

where $\text{MRS}_{lj}(x)$ is called the *marginal rate of substitution* between good l and good j (evaluated at x). Condition (6) means that the marginal rate of substitution between any two goods has to be equal to the ratio of the corresponding prices. It is interesting to note that Joe's marginal rate of substitution at the optimum would therefore be the same as Melanie's if she were to solve the same problem, even though her preferences might be very different from Joe's. In other

⁶More precisely, it is sufficient to assume that Joe's preferences are 'locally nonsatiated,' meaning that in the neighborhood of any bundle x there is another bundle \hat{x} (located at an arbitrarily close distance) that Joe strictly prefers. The consequence of this assumption is that Joe values any small increase of his wealth.

⁷We use the following conventions for inequalities involving vectors $x = (x_1, \dots, x_L)$ and $\hat{x} = (\hat{x}_1, \dots, \hat{x}_L)$: (i) $x \leq \hat{x} \Leftrightarrow x_i \leq \hat{x}_i \forall i$; (ii) $x < \hat{x} \Leftrightarrow x \leq \hat{x}$ and $\exists j$ s.t. $x_j < \hat{x}_j$; and (iii) $x \ll \hat{x} \Leftrightarrow x_i < \hat{x}_i \forall i$.

words, *any* consumer chooses his or her Walrasian demand vector such that $L - 1$ independent conditions in Eq. (6) are satisfied (e.g., for $i = 1, j \in \{2, \dots, L\}$). The missing condition for the determination of the Walrasian demand correspondence is given by the budget constraint,

$$p \cdot x = w. \quad (7)$$

As an example, consider the case where Joe has a utility function $u(x) = x_1^{\alpha_1} x_2^{\alpha_2}$ for consuming a bundle $x = (x_1, x_2)$,⁸ where $\alpha_1, \alpha_2 \in (0, 1)$ are constants. Then Eqs. (6) and (7) immediately imply Joe's Walrasian demand,

$$x(p, w) = \left(\frac{\alpha_1 w}{(\alpha_1 + \alpha_2) p_1}, \frac{\alpha_2 w}{(\alpha_1 + \alpha_2) p_2} \right).$$

Note that the Lagrange multiplier λ for the UMP (2) corresponds to the increase in Joe's *indirect utility*,⁹

$$v(p, w) = \max_{x \in B(p, w)} u(x). \quad (8)$$

Indeed, by the well-known envelope theorem (Milgrom and Segal 2002) we have that

$$\frac{\partial v(p, w)}{\partial w} = \frac{\partial \mathcal{L}(x, \lambda, \mu; p, w)}{\partial w} = \lambda. \quad (9)$$

Lagrange multipliers are sometimes also referred to as “shadow prices” of constraints. In this case, λ corresponds to the value of an additional dollar of budget as measured in terms of Joe's utility. Its unit is therefore “utile” per dollar.

Let us now briefly consider the possibility of a *corner solution* to Joe's utility maximization problem, i.e., his Walrasian demand vanishes for (at least) one good l . In that case, $\mu_l > 0$ and therefore

$$\text{MRS}_{lj}(x) = \frac{\lambda p_l - \mu_l}{\lambda p_j} = \frac{p_l}{p_j} \left(1 - \frac{\mu_l}{\lambda} \right) < \frac{p_l}{p_j},$$

provided that $x_j > 0$, i.e., Joe consumes a positive amount of good j .¹⁰ Thus, Joe's marginal utility $\partial u(x)/\partial x_l$ for good l , which is not consumed, is lower relative to his marginal utility for good j that he does consume, when compared to a consumer (e.g., Melanie) who uses positive amounts of both good l and good j . The latter consumer's marginal rate of substitution is equal to the ratio of prices (p_l/p_j), whereas Joe's marginal rate of substitution is strictly less. This means that a consumer may choose to forego consumption of a certain good if its price is too high compared to the price of other goods, so that it becomes ‘too expensive’ to adjust the marginal rate of substitution according to the (interior) optimality condition (6).

⁸This functional form, first proposed by Cobb and Douglas (1928) for identifying ‘production functions’ (cf. Section 3.2) rather than utility functions, is often used for its analytical simplicity and the fact that its parameters can easily be identified using data (e.g., by a linear regression).

⁹Joe's value function $v(p, w)$ is called “indirect utility” because he does usually not have a direct utility for the parameters p and w representing price and wealth, respectively, as they cannot be consumed directly. Yet they influence the problem and v can be thought of as representing a rational preference relation over different values of (p, w) .

¹⁰Clearly, because his utility is increasing and his wealth is positive, he will always demand a positive amount of at least one good.

2.3 Comparative Statics

Samuelson (1941) describes ‘comparative statics’ as the task of examining how a decision variable changes as a function of changes in parameter values. Examining the shifts of model predictions with respect to parameter changes is at the heart of price theory. We illustrate the main comparative statics techniques for the utility maximization problem. Consider, for example, the question of how Joe’s Walrasian demand $x(p, w)$ changes as his wealth w increases. His demand $x_l(p, w)$ for good l is called *normal* if it is nondecreasing in w . If that demand is positive, the first-order condition (3) can be differentiated with respect to the parameter w (using Eq. (9)) to obtain

$$\frac{\partial x(p, w)}{\partial w} = (D^2u(x(p, w)))^{-1} p \frac{\partial^2 v(p, w)}{\partial w^2} = (D^2u(x(p, w)))^{-1} p \frac{\partial \lambda(p, w)}{\partial w}.$$

If u is strongly concave, then $\partial \lambda(p, w)/\partial w \leq 0$,¹¹ i.e., the shadow value of wealth is decreasing as more and more wealth is added, since the individual’s marginal utility for the additional consumption decreases.

REMARK 1 A good that is not normal, i.e., for which demand decreases as wealth rises, is called *inferior*. Typical examples of such goods are frozen foods and bus transportation. In somewhat of a misnomer a normal good is referred to as a *superior* (or *luxury*) good if its consumption increases more than proportionally with wealth. Typical examples are designer apparel and expensive foods such as caviar or smoked salmon. In some cases the consumption of superior goods can drop to zero as price decreases. Luxury goods are often consumed as so-called *positional* goods, the value of which strongly depends on how they compare with the goods owned by others (Hirsch 1976; Frank 1985). \square

The change of the Walrasian demand vector with respect to a change in the price vector is more intricate. It is usually decomposed into a *substitution effect* and a *wealth effect*. For example, if Joe notices a change in the price vector from p to \hat{p} , then the substitution effect of this describes how Joe’s consumption vector changes if enough wealth is either added or subtracted to make sure that Joe’s utility is the same before and after the price change. For example, if his utility level before the price change is U , then his *wealth-compensated* (or *Hicksian*) demand vector $h(p, U)$ is such that it solves the *expenditure minimization problem*

$$h(p, U) \in \arg \min_{x \in \{\xi \in \mathbb{R}_+^L : u(\xi) \geq U\}} \{p \cdot x\}.$$

The Hicksian demand $h(p, U)$ achieves the utility level U at the price p at the lowest possible expenditure. For example, if $U = u(x(p, w))$, where $x(p, w)$ is Joe’s Walrasian demand vector, then $h(p, U) = x(p, w)$. If we denote the *expenditure* necessary to achieve the utility level U by $e(p, U) = p \cdot h(p, U)$, then one can (using the envelope theorem) show that $D_p e(p, U) = h(p, U)$ (Roy’s identity) and from there deduce that for $w = e(p, U)$: $D_p x(p, e(p, U)) = \partial x(p, w)/\partial p + (\partial x(p, w)/\partial w)h(p, U)$, so that one obtains Slutsky’s identity,

$$\underbrace{\frac{\partial x_l(p, w)}{\partial p_j}}_{\text{Price Effect}} = \underbrace{\frac{\partial^2 e(p, u(x(p, w)))}{\partial p_l \partial p_j}}_{\text{Substitution Effect}} - \underbrace{\frac{\partial x_l(p, w)}{\partial w} h_j(p, u(x(p, w)))}_{\text{Wealth Effect}}, \quad l, j \in \{1, \dots, L\}. \quad (10)$$

The (symmetric, negative semi-definite) matrix $S = \left[\frac{\partial^2 e(p, U)}{\partial p_l \partial p_j} \right]$ of substitution effects is referred to as *Slutsky matrix*.

¹¹The function u is strongly concave if its Hessian $D^2u(x)$ is negative definite for all x .




Figure 1: Price effect as a result of a price increase for the first good, i.e., a shift from $p = (p_1, p_2)$ to $\hat{p} = (\hat{p}_1, \hat{p}_2)$ (with $\hat{p}_1 > p_1$ and $\hat{p}_2 = p_2$), decomposed into substitution effect and wealth effect using the compensated wealth $\hat{w} = e(\hat{p}, u(x(p, w)))$.

REMARK 2 A *Giffen good* is such that demand for it increases when its own price increases. For example, when the price of bread increases, a poor family might not be able to afford meat any longer and therefore buy even more bread (Marshall 1920). It is necessarily an inferior good with a negative wealth effect that overcompensates the substitution effect, resulting in a positive price effect.¹² A related but somewhat different effect of increased consumption as a result of a price increase is observed for a so-called *Veblen good* (e.g., a luxury car or expensive jewelry). Veblen (1899, p. 75) pointed out that “[c]onspicuous consumption of valuable goods is a means of reputability to the gentleman of leisure.” Thus, similar to positional goods (cf. Remark 1), Veblen goods derive a portion of their utility from the price they cost (compared to other goods), since that implicitly limits their consumption by others. \square

We have seen that situations in which the dependence of decision variables is monotone in model parameters are especially noteworthy. The field of *monotone comparative statics*, which investigates conditions on the primitives of a UMP (or similar problem) that guarantee such monotonicity, is therefore of special relevance for price theory (and economics as a whole). Monotone comparative statics was pioneered by Topkis (1968; 1998), who introduced the use of lattice-theoretic (so-called ‘ordinal’) methods, in particular the notion of supermodularity.¹³ Milgrom and Shannon (1994) provide sufficient and in some sense necessary conditions for the monotonicity of solutions to UMPs in terms of (quasi-)supermodularity of the objective function, as long as

¹²Sørensen (2007) provides simple examples of utility functions which yield Giffen-good effects.

¹³For any two vectors $x, \hat{x} \in \mathbb{R}_+^L$, let $x \wedge \hat{x} = \min\{x, \hat{x}\}$ be their componentwise minimum and $x \vee \hat{x} = \max\{x, \hat{x}\}$ be their componentwise maximum. The choice set $X \subset \mathbb{R}$ is a *lattice* if $x, \hat{x} \in X$ implies that $x \wedge \hat{x} \in X$ and $x \vee \hat{x} \in X$. A function $u : X \rightarrow \mathbb{R}$ is *supermodular* on a lattice X if $u(x \vee \hat{x}) + u(x \wedge \hat{x}) \geq u(x) + u(\hat{x})$ for all $x, \hat{x} \in X$.

the choice set is a lattice.¹⁴ Consider normal goods as an example for monotone comparative statics. One can show that if u is strongly concave and supermodular, then the diagonal elements of the Hessian of u are nonpositive and its off-diagonal elements are nonnegative (Samuelson 1947). Strong concavity and supermodularity of u therefore imply that all goods are normal goods (Chipman 1977). Quah (2007) strengthens these results somewhat using ordinal methods.

REMARK 3 The practical interpretation of supermodularity is in terms of “complementarity.” To see this, consider a world with only two goods. Joe’s utility function $u(x_1, x_2)$ has *increasing differences* if $(\hat{x}_1, \hat{x}_2) \gg (x_1, x_2)$ implies that $u(\hat{x}_1, \hat{x}_2) - u(x_1, \hat{x}_2) \geq u(\hat{x}_1, x_2) - u(x_1, x_2)$. That is, the presence of more of good 2 increases Joe’s utility response to variations in good 1. For example, loudspeakers and a receiver are complementary: without receiver an agent can be almost indifferent about the presence of loudspeakers, whereas when the receiver is present, it makes a big difference if loudspeakers are available or not. A parameterized function is *supermodular* if it has increasing differences with respect to any variable-variable pair and any variable-parameter pair. \square

2.4 Effects of Uncertainty

So far, Joe’s choice problems have not involved any uncertainty. Yet, in many situations, instead of a specific outcome a decision maker can select only an action, which results in a probability distribution over many outcomes $x \in X$, or, equivalently, a random outcome \tilde{x} (or “lottery”) with realizations in X . Thus, given an action set A , the decision maker would like to select an action $a \in A$ so as to maximize the *expected utility* $\text{EU}(a) = E[u(\tilde{x})|a]$. The conditional distributions $F(x|a)$ defined for all $(x, a) \in X \times A$ are part of the primitives for this expected-utility maximization problem,

$$a^* \in \arg \max_{a \in A} \text{EU}(a) = \arg \max_{a \in A} \int_X u(x) dF(x|a).$$

For example, if Joe’s action $a \in A = [0, 1]$ represents the fraction of his wealth w that he can invest in a risky asset that pays a random return \tilde{r} , distributed according to the distribution function $G(r) = P(\tilde{r} \leq r)$, then the distribution of his ex-post wealth $\tilde{x} = w(1 + a\tilde{r})$ conditional on his action a is $F(x|a) = G((\frac{x}{w} - 1)/a)$ for $a \neq 0$ (and $F(x|a = 0) = \delta(x - w)$ and δ is a Dirac distribution¹⁵), so that $\text{EU}(a) = \int_{\mathbb{R}} u(x) dF(x|a) = \int_{\mathbb{R}} u(w(1 + ar)) dG(r)$ for all $a \in A$ by simple substitution. The solution to Joe’s classic *portfolio investment problem* depends on his attitude to risk. Joe’s *absolute risk aversion*, defined by

$$\rho(x) = -u''(x)/u'(x),$$

is a measure of how much he prefers a risk-free outcome to a risky lottery (in a neighborhood of x). For example, in Joe’s portfolio investment problem, there exists an amount $\pi(a)$ such that $u(E[\tilde{x}|a] - \pi(a)) = \text{EU}(a)$, which is called the *risk premium* associated with Joe’s risky

¹⁴Monotone comparative statics under uncertainty is examined by Athey (2002); non-lattice domains are investigated by Quah (2007). The latter is relevant, as even in a basic UMP of the form (2) for more than two commodities the budget set $B(p, w)$ is generally not a lattice. Strulovici and Weber (2008, 2010) provide methods for finding a reparametrization of models that guarantees monotone comparative statics, even though model solutions in the initial parametrization may not be monotone.

¹⁵The (singular) Dirac distribution δ can be defined in terms of a limit, $\delta(x) = \lim_{\varepsilon \rightarrow 0^+} (\varepsilon - |x|)_+ / \varepsilon^2$. It corresponds to a probability density with all its unit mass concentrated at the origin.

payoff lottery as a consequence of his investment action a . If Joe's absolute risk aversion is positive, so is his risk premium. The latter is the difference between the amount he is willing to accept for sure and the actuarially fair value of his investment under action a .

REMARK 4 In economic models, it is often convenient to assume that agents have *constant absolute risk aversion* (CARA) $\rho > 0$, which implies utility functions of the form $u(x) = -\exp(-\rho x)$ for all monetary outcomes $x \in \mathbb{R}$. The advantage is that the agents' risk attitude is then independent of their starting wealth, which insulates models from "wealth effects." Another common assumption, when restricting attention to positive wealth levels (where $x > 0$), is that agents have a *relative risk aversion* $\hat{\rho}(x) = x\rho(x)$ that is constant. Such agents with *constant relative risk aversion* (CRRA) $\hat{\rho} > 0$ have utility functions of the form

$$u(x) = \begin{cases} \ln x, & \text{if } \hat{\rho} = 1, \\ x^{1-\hat{\rho}}/(1-\hat{\rho}), & \text{otherwise.} \end{cases}$$

For more details on decision making under risk, see Pratt (1964) and Gollier (2001). An axiomatic base for expected utility maximization is provided in the pioneering work by Von Neumann and Morgenstern (1944). \square

3 Price Discovery in Markets

So far we have assumed that agents take the prices of all goods as given. Naturally, if goods are bought and sold in a market, prices will depend on the balance between demand and supply. Section 3.1 explains how prices are formed in a pure-exchange economy, in the absence of any uncertainty. Under weak conditions, trade achieves an economically "efficient" outcome in the sense that it maximizes "welfare," i.e., the sum of all agents' utilities. Under some additional assumptions, any efficient outcome can be achieved by trade, provided that a social planner can redistribute the agents' endowments before trading starts. In Section 3.2, we note that these insights carry over to economies where firms producing all available goods are owned and operated by the agents. Section 3.3 shows that these insights also apply in the presence of uncertainty, as long as the market is "complete" in the sense that all contingencies can be priced by bundles of available goods. Section 3.4 provides some details about "incomplete" markets, where this is not possible.

3.1 Pure Exchange

Consider an economy in which N agents can exchange goods in a market at no transaction cost. Each agent i has utility function $u^i(x^i)$, where $x^i \in \mathbb{R}_+^L$ is his consumption bundle, and is endowed with the bundle $\omega^i \in \mathbb{R}_+^L$ (his *endowment*). At the nonnegative price vector p this agent's Walrasian demand $x^i(p, p \cdot \omega^i)$ is obtained by solving a UMP, where $p \cdot \omega^i$ is used as his wealth. The *market-clearing condition* that supply must equal demand yields

$$\sum_{i=1}^N x^i(p, p \cdot \omega^i) = \sum_{i=1}^N \omega^i, \quad (11)$$

and thus L relations (one for each good) that imply the price vector $p = (p^1, \dots, p^L)$ up to a common multiplicative constant, for only $L - 1$ components in Eq. (11) can be independent.

DEFINITION 3 A Walrasian equilibrium (\hat{p}, \hat{x}) in a pure-exchange economy consists of an equilibrium price $\hat{p} \in \mathbb{R}_+^L$ and an equilibrium allocation $\hat{x} = (\hat{x}^1, \dots, \hat{x}^N)$ such that

$$\hat{x}^i \in \arg \max_{x^i \in B(\hat{p}, \hat{p} \cdot \omega^i)} u^i(x^i), \quad i \in \{1, \dots, N\}, \quad (12)$$

and

$$\sum_{i=1}^N (\hat{x}^i - \omega^i) = 0. \quad (13)$$

As an example, let us consider Joe (= agent 1) and Melanie (= agent 2) with identical Cobb-Douglas utility functions $u^1(x_1, x_2) = u^2(x_1, x_2) = x_1^\alpha x_2^{1-\alpha}$, where $\alpha \in (0, 1)$ is a given constant. Suppose that Joe is endowed with the bundle $\omega^1 = (1, 2)$ and Melanie with the bundle $\omega^2 = (2, 1)$. The UMP (12) implies that agent i 's Walrasian demand vector (or "offer curve" (OC) when viewed as a function of price) is

$$x^i(p, p \cdot \omega^i) = \left(\frac{\alpha p \cdot \omega^i}{p_1}, \frac{(1-\alpha)p \cdot \omega^i}{p_2} \right).$$

From Eq. (13) we obtain that $\alpha(p_1 + 2p_2)/p_1 + \alpha(2p_1 + p_2)/p_1 = 3$, so that the ratio of prices becomes $p_1/p_2 = \alpha/(1-\alpha)$. By setting $p_1 = 1$ (which amounts to considering good 1 as the *numéraire*, cf. Section 1) we can therefore immediately determine a unique Walrasian equilibrium (\hat{p}, \hat{x}) , where $\hat{p} = (1, (1-\alpha)/\alpha)$ and $\hat{x}^i = x^i(\hat{p}, \hat{p} \cdot \omega^i)$ for $i \in \{1, 2\}$, so that $\hat{x}^1 = (2-\alpha, 2-\alpha)$ and $\hat{x}^2 = (1+\alpha, 1+\alpha)$. The distribution of resources in a two-agent exchange economy, such as in this example, can be conveniently displayed using the so-called Edgeworth-Bowley box diagram as shown in Figure 2.¹⁶ The figure also shows that the intersection of the agents' offer curves lies on a "contract curve" which contains all "efficient" allocations as explained below.

The beauty of an exchange economy in which price takers interact freely without any transaction cost is that a Walrasian equilibrium allocation cannot be improved upon in the following sense. A feasible allocation $x = (x^1, \dots, x^N)$ is said to be *Pareto-efficient* (relative to the set X of feasible allocations) if there exists no other allocation $\hat{x} = (\hat{x}^1, \dots, \hat{x}^N)$ in X at which all individuals are at least as well off as at x and at least one individual is better off. More precisely, x is Pareto-efficient if for all $\hat{x} \in X$:¹⁷

$$(u^i(x^i) \leq u^i(\hat{x}^i) \forall i) \quad \Rightarrow \quad (u^i(x^i) = u^i(\hat{x}^i) \forall i).$$

Adam Smith (1776, Book IV, Chapter 2) pointed to an "invisible hand" that leads individuals through their self-interest to implement socially optimal outcomes. A key result of price theory is that Walrasian equilibria, even though merely defined as a solution to individual utility maximization problems and a feasibility constraint (supply = demand), produce Pareto-efficient outcomes.

THEOREM 1 (FIRST FUNDAMENTAL WELFARE THEOREM) Any Walrasian equilibrium allocation is Pareto-efficient.¹⁸

¹⁶A simple diagram of this sort was used by Edgeworth (1881, p. 28) to illustrate exchange allocations, and it was later popularized by Bowley (1924).

¹⁷For each Pareto-efficient allocation $x^* \in X$ there exists a vector $\lambda = (\lambda^1, \dots, \lambda^N) \geq 0$ of nonnegative weights such that $x^* \in \arg \max_{x \in X} \left\{ \sum_{i=1}^N \lambda^i u^i(x^i) \right\}$. The last problem can be used to generate all such Pareto-efficient allocations by varying λ .

¹⁸The Pareto-efficiency of a Walrasian equilibrium allocation depends on the fact that consumer preferences are locally nonsatiated (cf. Footnote 6).




Figure 2: Offer curves (OC^1 and OC^2), Walrasian equilibrium, and Pareto-efficient allocations.

The intuition for the Pareto-efficiency of a Walrasian equilibrium allocation is best seen by contradiction. Suppose that there exists an allocation $x = (x^1, x^2)$ that would improve Joe's well-being and leave Melanie at the utility level she enjoys under the Walrasian equilibrium allocation $\hat{x} = (\hat{x}^1, \hat{x}^2)$. By Walras' Law, for Joe allocation x^1 is not affordable under the equilibrium price \hat{p} , so that $\hat{p} \cdot x^1 > \hat{p} \cdot \hat{x}^1$. Furthermore, because Melanie is maximizing her utility in equilibrium, the alternative allocation x^2 cannot leave her with any excess wealth, so that $\hat{p} \cdot x^2 \geq \hat{p} \cdot \hat{x}^2$. Thus, the total value of the alternative allocation, $\hat{p} \cdot (x^1 + x^2)$, is *strictly* greater than the total value of the Walrasian equilibrium allocation, $\hat{p} \cdot (\hat{x}^1 + \hat{x}^2)$. But this contradicts the fact that the total amount of goods in the economy does not depend on the chosen allocation, so that the total value in the economy is constant for a given price. Hence, we have obtained a contradiction, which implies that the Walrasian equilibrium allocation must indeed be Pareto-efficient.

THEOREM 2 (SECOND FUNDAMENTAL WELFARE THEOREM) *In a convex economy it is possible to realize any given Pareto-efficient allocation as a Walrasian equilibrium allocation, after a lump-sum wealth redistribution.*

This important result rests on the assumption that the economy is convex, in the sense that each consumer's sets of preferred goods ("upper contour sets") relative to any feasible endowment point is convex. The latter is necessary to guarantee the existence of an equilibrium price; it is satisfied if all consumers' utility functions are concave.¹⁹ Figure 3 provides an example of how in a nonconvex economy it may not be possible to obtain a Walrasian equilibrium: as a function

¹⁹The proof of the second fundamental welfare theorem relies on the separating hyperplane theorem (stating that there is always a plane that separates two convex sets which have no common interior points). In an economy with production (cf. Section 3.2) the firms' production sets need to also be convex.

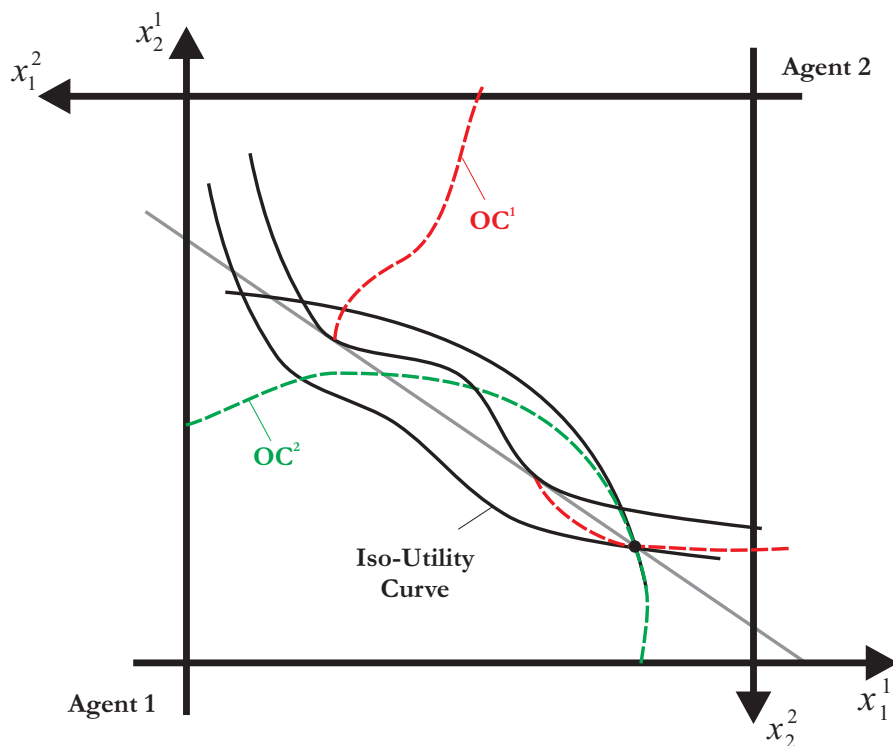


Figure 3: Nonconvex exchange economy without a Walrasian equilibrium.

of price, agent 1 switches discretely in his preference of good 1 and good 2, so that the market cannot clear at any price (except when all its components are infinity, so that agents simply keep their endowments).

The set of Pareto-efficient allocations in an economy is often referred to as the “contract curve” (or Pareto set; cf. Footnote 17 and Figure 2). The subset of Pareto-efficient allocations which present Pareto-improvements over the endowment allocation is called the *core* of the economy. While the first fundamental welfare theorem says that the Walrasian market outcome is in the core of an economy, the second fundamental welfare theorem states that it is possible to rely on markets to implement any outcome in the Pareto set, provided that a lump-sum reallocation of resources takes place before markets are opened.

3.2 Competitive Markets

The setting of the exchange economy in Section 3.1 does not feature any productive activity by firms. Consider M such firms. Each firm m has a (nonempty) production set $Y^m \subset \mathbb{R}^L$ which describes its production choices. For a feasible production choice $y^m = (y_1^m, \dots, y_L^m)$, we say that firm m produces good l if y_l^m is nonnegative; otherwise it uses that good as an input. Thus, since it is generally not possible to produce goods without using any inputs, it is natural to require that $Y \cap \mathbb{R}_+^L \subseteq \{0\}$. This so-called *no-free-lunch* property of the production set Y^m means that if a firm produces goods without any inputs, then it cannot produce a positive amount of anything. Given a price vector p , firm m 's profit is $p \cdot y$. The firms' *profit-maximization problem* is therefore to find

$$y^m(p) \in \arg \max_{y^m \in Y^m} \{p \cdot y\}, \quad m \in \{1, \dots, M\}. \quad (14)$$

In a *private-ownership economy* each firm m is privately held. That is, each agent i owns the nonnegative share ϑ_m^i of firm m , such that

$$\sum_{i=1}^N \vartheta_m^i = 1, \quad m \in \{1, \dots, M\}. \quad (15)$$

DEFINITION 4 A *Walrasian equilibrium* $(\hat{p}, \hat{x}, \hat{y})$ in a *private-ownership economy* is such that all agents maximize utility,

$$\hat{x}^i \in \arg \max_{x^i \in B(\hat{p}, \hat{p} \cdot \omega^i + \sum_{m=1}^M \vartheta_m^i (\hat{p} \cdot \hat{y}^m))} u^i(x^i), \quad i \in \{1, \dots, N\},$$

all firms maximize profits,

$$\hat{y}^m \in \arg \max_{y^m \in Y^m} \{\hat{p} \cdot y\}, \quad m \in \{1, \dots, M\},$$

and the resulting allocation is feasible,

$$\sum_{i=1}^N (\hat{x}^i - \omega^i) = \sum_{m=1}^M \hat{y}^m.$$

The two fundamental welfare theorems continue to hold in the more general setting with production. The existence of a Walrasian equilibrium is guaranteed as long as there are no externalities (cf. Section 5) and the economy is convex (cf. Footnote 19).

3.3 Complete Markets

The agents' consumption choice and the firms' production decisions are generally subject to uncertainty. As in Section 2.4, this may simply mean that agents maximize expected utility and firms expected profits. However, in many situations the agents can use a market mechanism to trade *before* a random state of the world realizes. Arrow (1953) and Debreu (1953) have shown that all the efficiency properties of the Walrasian equilibrium carry over to case with uncertainty, provided that 'enough' assets are available, a notion that will be made precise in the definition of complete markets below (cf. Definition 5). For simplicity we assume that the uncertain state of the world, denoted by the random variable \tilde{s} , can have realizations in the finite state space $S = \{s_1, \dots, s_K\}$. Agent i believes that state s_k occurs with probability μ_k^i .

We now develop a simple two-period model to understand how agents (or "traders") can trade in the face of uncertainty. Consider Joe (= agent 1), who initially owns firm 1, and Melanie (= agent 2), who initially owns firm 2. The state space $S = (s_1, s_2)$ contains only two elements (i.e., $K = 2$). Let V^m be the market value of firm m in period 1. The monetary value of firm m in state s_k is ω_k^m . After trading of firm shares takes place in the first period, Joe and Melanie each hold a portfolio $\vartheta^i = (\vartheta_1^i, \vartheta_2^i)$ of ownership shares in the firms. In the second period all uncertainty realizes and each agent i can make consumption decisions using the wealth $w_k^i(\vartheta^i) = \vartheta_1^i \omega_k^1 + \vartheta_2^i \omega_k^2$ in state s_k , obtaining the indirect utility $v_k^i(w_k^i)$. Hence, in period one, the traders Joe and Melanie, can trade firm shares as follows. Each trader i solves the expected utility maximization problem

$$\hat{\vartheta}^i(V) \in \arg \max_{\vartheta^i \in B((V^1, V^2), V^i)} \left\{ \sum_{k=1}^K \mu_k^i v_k^i(w_k^i(\vartheta^i)) \right\}.$$




Figure 4: Market for contingent claims.

If in addition the market-clearing condition (15) holds, then the resulting equilibrium $(\hat{V}, \hat{\vartheta})$, with $\hat{V} = (\hat{V}^1, \hat{V}^2)$ and $\hat{\vartheta} = (\hat{\vartheta}^1(\hat{V}), \hat{\vartheta}^2(\hat{V}))$, is called a *rational expectations equilibrium*.²⁰ In this equilibrium, Joe and Melanie correctly anticipate the equilibrium market prices of the firm when making their market-share offers, just as in the Walrasian pure-exchange economy discussed in Section 3.1.

In the last example, Joe and Melanie each owned a so-called *asset* (or *security*), the defining characteristic of which is that it entitles the owner to a determined monetary payoff in each state of the world. All that matters for consumption after the conclusion of trade in the first period is how much money an agent has available in each state. Thus, Joe and Melanie could come to the conclusion that instead of trading the firm shares, it would be more appropriate to trade directly in ‘state-contingent claims to wealth.’ If trader i holds a state-contingent claim portfolio $c^i = (c_1^i, \dots, c_K^i)$, then in state s_k that trader obtains the indirect utility $v_k^i(c_k^i)$. In the first period when trading in contingent claims takes place, the price for a claim to wealth in state s_k is p_k . Hence, each trader i has demand

$$\tilde{c}^i \in \arg \max_{c^i \in B(p, p \cdot \omega^i)} \left\{ \sum_{k=1}^K v_k^i(c_k^i) \right\}, \quad i \in \{1, \dots, N\}, \quad (16)$$

that maximizes expected utility. Eq. (16) and the market-clearing relation

$$\sum_{i=1}^N (\hat{c}^i - \omega^i) = 0 \quad (17)$$

together constitute the conditions for a Walrasian equilibrium (\hat{p}, \hat{c}) in the market of contingent claims. Using the insights from Section 2.2, note that the marginal rate of substitution between

²⁰The concept of rational expectations in economics originated with Muth (1961), and the rational expectations equilibrium which is used here with Radner (1972).

claims in state s_k and claims in state s_l equals the ratio of the corresponding market prices at an (interior) equilibrium,

$$\text{MRS}_{kl}(c^i) = \frac{\mu_k^i v_k^i(c_k^i)}{\mu_l^i v_l^i(c_l^i)} = \frac{p_k}{p_l}.$$

This is independent of the agent i , so that after normalizing the market prices to sum to one, the prices p_1, \dots, p_K define a probability distribution over the states of the world, which is referred to as *equivalent martingale measure*. While traders may disagree about their probability assessments for the different states of the world, they are in agreement about the equivalent martingale measure in equilibrium.²¹

It is possible to go back and forth between the market for firm shares and the market for contingent claims if and only if the system of equations

$$\underbrace{\begin{bmatrix} \omega_1^1 & \cdots & \omega_1^M \\ \vdots & & \vdots \\ \omega_K^1 & \cdots & \omega_K^M \end{bmatrix}}_{\Omega} \underbrace{\begin{bmatrix} \vartheta_1^i \\ \vdots \\ \vartheta_M^i \end{bmatrix}}_{\vartheta^i} = \underbrace{\begin{bmatrix} c_1^i \\ \vdots \\ c_K^i \end{bmatrix}}_{c^i}$$

possesses a solution, or, equivalently, if the *asset return matrix* Ω is of rank K .

DEFINITION 5 (i) If the $K \times M$ asset return matrix Ω has rank K , then the market for contingent claims is called *complete*. (ii) A Walrasian equilibrium (p, c) in a complete market for contingent claims, satisfying Eqs. (16) and (17), is called an *Arrow-Debreu equilibrium*.

The concept of completeness can easily be extended to multi-period economies (Debreu 1959), where time is indexed by $t \in \{0, 1, \dots, T\}$. The state of the world \tilde{s}_t at time $t \geq 1$ can depend on the state of the world \tilde{s}_{t-1} at time $t - 1$. As the states of the world successively realize, they plow a path through an *event tree*. In a complete market it is possible to trade claims that are contingent on any possible path in the event tree.

3.4 Incomplete Markets

If the market is not complete (in the sense of Definition 5(i)), the rational expectations equilibrium (also referred to as *Radner equilibrium*) may produce a Pareto-inefficient allocation. The reason is that if there are more states of the world than linearly independent assets, then it is generally not possible for agents to trade contracts that diversify the risks in the economy to the desirable degree. To see this, consider two risk-averse agents in an economy with two or more states, where there are no firms (or assets) that can be traded, so that there is no possibility for mutual insurance. The latter would be the Pareto-efficient outcome in an Arrow-Debreu equilibrium when contingent claims are available.

The consequence of market incompleteness is that agents in the economy cannot perfectly trade contingent claims. One can show that, in an economy with only two periods, a rational expectations equilibrium is “constrained Pareto-efficient,” in the sense that trade in the first period is

²¹Aumann (1976) points out that, in a statistical framework, it is in fact impossible for agents to “agree to disagree” on probability distributions if all the evidence is made available to all agents. Naturally, without such information exchange, the agents’ subjective probabilities may vary significantly.

such that agents obtain a Pareto-efficient allocation in expected utilities, subject to the available assets.²²

The question naturally arises as to how assets should be priced when markets are incomplete. The answer is that depending on the agents' beliefs in the market, a number of different prices are possible. In general, it is reasonable to assume that, as long as transaction costs are negligible, asset prices do not allow for *arbitrage* possibilities, i.e., ways of obtaining a risk-free profit through mere trading of assets. For concreteness, let us assume that there are three assets and that each asset $m \in \{1, 2, 3\}$ is characterized by a return vector $\omega^m = (\omega_1^m, \dots, \omega_K^m)$ which specifies the payoffs for all possible states s_1, \dots, s_K . In addition, suppose that the return vector of the third asset is a linear combination of the return vectors of the first two assets, so that

$$\omega^3 = \phi^1 \omega^1 + \phi^2 \omega^2$$

for some constants ϕ^1 and ϕ^2 . If V^m is the market price of asset m , then clearly we must have that

$$V^3 = \phi^1 V^1 + \phi^2 V^2$$

in order to exclude arbitrage opportunities. In other words, if the state-contingent payoffs of an asset can be replicated by a portfolio of other assets in the economy, then the price of the asset must equal the price of the portfolio.²³ Accordingly, *no-arbitrage pricing* refers to the selection of prices that do not allow for risk-free returns resulting from merely buying and selling available assets.²⁴ But no-arbitrage pricing in an incomplete market alone provides only an upper and a lower bound for the price of an asset. Additional model structure (e.g., provided by general equilibrium assumptions or through the selection of an admissible equivalent martingale measure) is required to pinpoint a particular asset price. Arbitrage pricing theory (Ross 1976; Roll and Ross 1980) postulates that the expected value of assets can be well estimated by a linear combination of fundamental macro-economic factors (e.g., price indices). The sensitivity of each factor is governed by its so-called β -coefficient.

A rational expectations equilibrium in an economy in which the 'fundamentals of the economy,' i.e., the agents' utilities and endowments, do not depend on several states, but consumptions are different across those states, is called a *sunspot equilibrium* (Cass and Shell 1983). The idea is that observable signals that bear no direct effect on the economy may be able to influence prices through traders' expectations. In light of well-known boom-bust phenomena in stock markets, it is needless to point out that traders' expectations are critical in practice for the formation of prices. When market prices are at odds with the *intrinsic value* of an asset (i.e., the value implied by its payoff vector), it is likely that traders are trading in a "speculative bubble" because of self-fulfilling expectations about the further price development. Well-known examples of speculative bubbles include the tulip mania in the Netherlands in 1637 (Garber 1990), the dot-com bubble at the turn of last century (Shiller 2005; Malkiel 2007), and, more recently, the U.S. housing bubble (Sowell 2010). John Maynard Keynes (1936) offered the following comparison:

²²In multi-period (and/or multi-good) economies, rational expectations equilibria are not even guaranteed to be constrained Pareto-efficient (for details see, e.g., Magill and Quinzii (1996)).

²³To see this, assume for example that the market price V^3 is greater than $\phi^1 V^1 + \phi^2 V^2$. Thus, if a trader holds a portfolio with quantities $q^1 = \phi^1 V^3$ of asset 1, $q^2 = \phi^2 V^3$ of asset 2, and $q^3 = -(\phi^1 V^1 + \phi^2 V^2)$ of asset 3, then the value of that portfolio vanishes, as $\sum_{m=1}^3 q^m V^m = 0$. However, provided an equivalent martingale measure p_1, \dots, p_M (cf. Section 3.3) the total return of the portfolio, $\sum_{k=1}^K p_k V^m \omega_k^m q_k = p_k \omega_k^3 (V^3 - (\phi^1 V^1 + \phi^2 V^2))$, is positive for each state s_k .

²⁴Another approach to asset pricing is Luenberger's (2001; 2002) zero-level pricing method, based on the widely used Capital Asset Pricing Model (Markowitz 1952; Tobin 1958; Sharpe 1964). It relies on the geometric projection of the prices of 'similar' traded assets.

“... professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one’s judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees” (p. 156).

In this “beauty contest,” superior returns are awarded to the trader who correctly anticipates market sentiment. For additional details on asset pricing see, e.g., Duffie (1988; 2001).

4 Disequilibrium and Price Adjustments

In practice, we cannot expect markets always to be in equilibrium, especially when agents are free to enter and exit, economic conditions may change over time, and not all market participants possess the same information. Section 4.1 specifies a possible price adjustment process, referred to as “tâtonnement,” that tends to attain a Walrasian equilibrium asymptotically over time. In Section 4.2, we summarize important insights about the strategic use of private information in markets, e.g., the fundamental result that trade might not be possible at all if all agents are rational and some hold extra information.

4.1 Walrasian Tâtonnement

In real life there is no reason to believe that markets always clear. At a given price it may be that there is either excess demand or excess supply, which in turn should lead to an adjustment of prices and/or quantities. While it is relatively simple to agree about the notion of a static Walrasian equilibrium, which is based on self-interested behavior of price-taking agents and firms as well as a market-clearing condition, there are multiple ways of modelling the adjustment dynamics when markets do not clear. Walras (1874/77) was the first to formalize a price-adjustment process by postulating that the change in the price of good $l \in \{1, \dots, L\}$ is proportional to the *excess demand* z_l of good l , where (suppressing all dependencies from entities other than price)

$$z_l(p) = \sum_{i=1}^N (x_l^i(p) - \omega_l^i) - \sum_{m=1}^m y_l^m(p).$$

This leads to an adjustment process, which in continuous time can be described by a system of L differential equations,

$$\dot{p}_l = \kappa_l z_l(p), \quad l \in \{1, \dots, L\}. \quad (18)$$

This process is commonly referred to as a *Walrasian tâtonnement*. The positive constants κ_l determine the speed of the adjustment process for each good l .²⁵ One can show that if the Walrasian equilibrium price vector \hat{p} is unique, and if $\hat{p} \cdot z(p) > 0$ for any p not proportional to \hat{p} , then a solution trajectory $p(t)$ of the Walrasian *tâtonnement* (18) converges to \hat{p} , in the sense that

$$\lim_{t \rightarrow \infty} p(t) = \hat{p}.$$

As an illustration, we continue the example of a two-agent two-good exchange economy with Cobb-Douglas utility functions in Section 3.1. The *tâtonnement* dynamics in Eq. (18) become

$$\begin{cases} \dot{p}_1 &= 3 \kappa_1 (\alpha(p_1 + p_2)/p_1 - 1), \\ \dot{p}_2 &= 3 \kappa_2 ((1 - \alpha)(p_1 + p_2)/p_2 - 1), \end{cases}$$

where $\kappa_1, \kappa_2 > 0$ are appropriate constants. It is easy to see that $p_1/p_2 < \alpha/(1 - \alpha)$ implies that $\dot{p}_1 > 0$ (and that $\dot{p}_2 < 0$) and vice versa. In (p_1, p_2) -space we therefore obtain that a trajectory $p_2(p_1)$ (starting at any given price vector) is described by

$$\frac{dp_2(p_1)}{dp_1} = \frac{3 \kappa_2 ((1 - \alpha)(p_1 + p_2)/p_2 - 1)}{3 \kappa_1 (\alpha(p_1 + p_2)/p_1 - 1)} = -\frac{\kappa_2 p_1}{\kappa_1 p_2}.$$

It follows a concentric circular segment, eventually approaching a ‘separatrix’ defined by $p_1/p_2 = \alpha/(1 - \alpha)$ (cf. Figure 5). We see that when the price of good 1 is too low relative to the price of good 2, then due to the positive excess demand, p_1 will adjust upwards until the excess demand for that good vanishes, which – by the market-clearing condition – implies that the excess demand for the other good also vanishes, so that we have arrived at an equilibrium.²⁶

Walrasian price-adjustment processes have been tested empirically. Joyce (1984) reports experimental results in an environment with consumers and producers that show that for a single unit of a good the *tâtonnement* can produce close to Pareto-efficient prices, and that the process has strong convergence properties. Bronfman et al. (1996) consider the multi-unit case and find that the efficiency properties of the adjustment process depend substantially on how this process is defined. Eaves and Williams (2007) analyze Walrasian *tâtonnement* auctions at the Tokyo Grain Exchange run in 1997/98 and find that price formation is similar to the normative predictions in continuous double auctions.

We note that it is also possible to consider quantity adjustments instead of price adjustments (Marshall 1920). The resulting quantity-adjustment dynamics are sometimes referred to as *Marshallian dynamics* (in contrast to the *Walrasian dynamics* for price adjustments).

²⁵It may be useful to normalize the price vector $p(t)$ such that $p_1^2(t) + \dots + p_L^2(t) \equiv 1$, since then $d(p_1^2(t) + \dots + p_L^2(t))/dt = 2 p(t) \cdot z(p(t)) \equiv 0$ (with $z = (z_1, \dots, z_L)$). Under this normalization any trajectory $p(t)$ remains in the ‘invariant’ set $S = \{p : (p_1)^2 + \dots + (p_L)^2 = 1\}$.

²⁶A Walrasian equilibrium price is locally asymptotically stable if, when starting in a neighborhood of this price, Walrasian *tâtonnement* yields price trajectories that converge toward the equilibrium price. A sufficient condition for local asymptotic stability (i.e., convergence) is that the linearized system $\dot{p} = A(p - \hat{p})$ corresponding to the right-hand side of Eq. (18) around the Walrasian equilibrium price \hat{p} is such that the linear system matrix A has only eigenvalues with negative real parts. In the example, we have that $A = 3\alpha \begin{bmatrix} -\alpha\kappa_1/(1 - \alpha) & \kappa_1 \\ \kappa_2 & -\alpha\kappa_2/(1 - \alpha) \end{bmatrix}$, the eigenvalues of which have negative real parts for all $\kappa_1, \kappa_2 > 0$ and all $\alpha \in (0, 1)$. For more details on the analysis of dynamic systems, see, e.g., Weber (2011).




Figure 5: Walrasian price adjustment process with asymptotic convergence of prices.

4.2 Information Transmission in Markets

It is an economic reality that different market participants are likely to have different information about the assets that are up for trade. Thus, in a market for used cars, sellers may have more information than buyers. For simplicity, let us consider such a market, where cars are either of value 0 or of value 1, but it is impossible for buyers to tell which one is which until after the transaction has taken place. Let $\varphi \in (0, 1)$ be the fraction of sellers who sell “lemons” (i.e., cars of zero value) and let $c \in (0, 1)$ be the opportunity cost of a seller who sells a high-value car. Then, if $1 - \varphi < c$, there is no price p at which buyers would want to buy and high-value sellers would want to sell. Indeed, high-value sellers sell only if $p \geq c$. Lemons sellers simply imitate high-value sellers and also charge p for their cars (they would be willing to sell at any nonnegative price). As a consequence, the buyers stand to obtain the negative expected value $(1 - \varphi) - p$ from buying a car in this market, which implies that they do not buy. Hence, the market for used cars fails, in the sense that only lemons can be traded, if there are too many lemons compared to high-value cars (Akerlof 1970). This self-inflicted disappearance of high-value items from the market is called “adverse selection.”

In the context of Walrasian markets in the absence of nonrational (“noise”) traders, and as long as the way in which traders acquire private information about traded assets is common knowledge, Milgrom and Stokey (1982) show that none of the traders is in a position to profit from the private information, as any attempt to trade will lead to a correct anticipation of market prices. This implies that private information cannot yield a positive return. Since this theoretical no-trade theorem is in sharp contrast to the reality found in most financial markets, where superior information tends to yield positive (though sometimes illegal) returns, the missing ingredient are irrational “noise traders,” willing to trade without concerns about the private information available to other traders (e.g., for institutional reasons).

What information can be communicated in a market? Hayek (1945, p. 526) points out that one should in fact consider the “price system” in a market as “a mechanism for communicating information.” The price system is useful under uncertainty, since (as we have seen in Section 3)

markets generally exist not only for the purpose of allocating resources, but also to provide traders with the possibility of mutual insurance (Hurwicz 1960). To understand the informational role of prices under uncertainty, let us consider a market as in Section 3.4, where N strategic traders (“investors”) and a number of nonstrategic traders (“noise traders”) trade financial securities that have state-contingent payoffs. Each investor i may possess information about the future payoff of these securities in the form of a private signal \tilde{z}^i , which, conditional on the true state of the world, is independent of any other investor j ’s private signal \tilde{z}^j . In the classical model discussed thus far, at any given market price vector p for the available securities, agent i has the Walrasian demand $x^i(p, \omega^i; z^i)$, which – in addition to his endowment ω^i – also depends on the realization z^i of his private signal. However, an investor i who conditions his demand only on the realized price and his private information would ignore the process by which the *other* agents arrive at their demands, which help establish the market price, which in equilibrium must depend on the full information vector $z = (z^1, \dots, z^N)$. For example, if investor i ’s private information leads him to be very optimistic about the market outlook, then at any given price p his demand will be high. Yet, if a very low market price p is observed, investor i obtains the valuable insight that all other investors’ signal realizations must have been rather dim, which means that his information is likely to be an extreme value from a statistical point of view. Grossman and Stiglitz (1980) therefore conclude that in a rational-expectations equilibrium each investor i ’s demand must be of the form $x^i(p, \omega^i; z^i, p(z))$, i.e., it will depend on the way in which the equilibrium price incorporates the available information.²⁷ While in an economy where private information is freely available this may result in the price to fully reveal the entire information available in the market (because, for example, it depends only on an average of the investors’ information; Leland and Pyle 1976), this does not hold when information is costly. Grossman and Stiglitz show that in the presence of a (possibly even small) cost of acquiring private information, prices fail to aggregate all the information in the market, so that a rational expectations equilibrium cannot be ‘informationally efficient.’²⁸

The question naturally arises of how to use private information in an effective way, especially if one can do so over time.²⁹ Kyle (1985) develops a seminal model of insider trading where one informed investor uses his inside information in a measured way over a finite time horizon so as to not be imitated by other rational investors. Another option an informed investor has for using private information is to sell it to other investors. Admati and Pfleiderer (1986) show that it may be best for the informed investor to degrade this information (by adding noise) before

²⁷In an actual financial market (such as the New York Stock Exchange), investors can submit their offer curves in terms of ‘limit orders,’ which specify the number of shares of each given asset that the investor is willing to buy at a given price.

²⁸In the absence of randomness a rational expectations equilibrium may even fail to exist. Indeed, if no investor expends the cost to become informed, then the equilibrium price reveals nothing about the true value of the asset, which produces an incentive for investors to become informed (provided the cost is small enough). On the other hand, if all other investors become informed, then, due to the nonstochastic nature of the underlying values, an uninformed investor could infer the true value of any security from the price, which in turn negates the incentive for the costly information acquisition.

²⁹It is important to realize that while *private* information tends to be desirable in most cases (for exceptions, see Section 7.4), this may not be true for *public* (or “social”) information. Hirshleifer (1971) shows that social information that is provided to all investors in a market might have a negative value because it can destroy the market for mutual insurance. To see this, consider two farmers, one with a crop that grows well in a dry season and the other with a crop that grows well in a wet season. If both farmers are risk averse, then given, say, ex-ante equal chances of either type of season to occur, they have an incentive to write an insurance contract that guarantees part of the proceeds of the farmer with the favorable outcome to the other farmer. Both of these farmers would *ex ante* be strictly worse off if a messenger disclosed the type of season (dry or wet) to them (while, clearly, each farmer retains an incentive to obtain such information privately).

selling it, in order to protect his own trading interests.

REMARK 5 Arrow (1963) realized that there are fundamental difficulties when trying to sell information from an informed party to an uninformed party, perhaps foreshadowing the no-trade theorem discussed earlier. Indeed, if the seller of information just claims to have the information, then a buyer may have no reason to believe the seller.³⁰ The seller therefore may have to ‘prove’ that the information is really available, e.g., by acting on it in a market (taking large positions in a certain stock), to convince potential buyers. This may dissipate a part or all of the value of the available information. On the other hand, if the seller of information transmits the information to a prospective buyer for a ‘free inspection,’ then it may be difficult to prevent that potential buyer from using it, even when the latter later decides not to purchase the information. Arrow (1973) therefore highlights a fundamental “inappropriability” of information.³¹ Clearly, information itself can be securitized, as illustrated by the recent developments in prediction markets (Wolfers and Zitzewitz 2004). Segal (2006) provides a general discussion about the informational requirements for an economic mechanism (such as a market) if its purpose is to implement Pareto-efficient outcomes but where agents possess private knowledge about their preferences (and not the goods). \square

5 Externalities and Nonmarket Goods

Sometimes transactions take place outside of markets. For example, one agent’s action may have an effect on another agent’s utility without any monetary transfer between these agents. In Section 5.1, we see that the absence of a market for such “externalities” between agents can cause markets for other goods to fail. Similarly, the markets for certain goods, such as human organs, public parks, or clean air, may simply not exist. The value of “nonmarket goods,” which can be assessed using the welfare measures of “compensating variation” and “equivalent variation,” is discussed in Section 5.2.

5.1 Externalities

In Section 3 we saw that markets can produce Pareto-efficient outcomes, even though exchange and productive activity is not centrally managed and is pursued entirely by self-interested parties. One key assumption there was that each agent i ’s utility function u^i depends only on his own consumption bundle x^i . However, this may not be appropriate in some situations. For example, if Joe listens to loud music while Melanie tries to study for an upcoming exam, then Joe’s consumption choice has a direct impact on Melanie’s well-being. We say that his action exerts a (direct) *negative externality* on her. Positive externalities also exist, for example when Melanie decides to do her cooking and Joe loves the resulting smell from the kitchen, her action has a

³⁰More recently, ‘zero-knowledge proof’ techniques have been developed for (at least approximately) conveying the fact that information is known without conveying the information itself; see, e.g., Goldreich et al. (1991).

³¹For “information goods” such as software, techniques to augment appropriability (involving the partial transmission of information) have been refined. For example, it is possible to provide potential buyers of a software package with a free version (“cripple ware”) that lacks essential features (such as the capability to save a file) but that effectively demonstrates the basic functionality of the product, without compromising its commercial value.

direct positive impact on his well-being.³²

There are many important practical examples of externalities in the economy, including environmental pollution, technological standards, telecommunication devices, or public goods. For concreteness, let us consider two firms, 1 and 2. Assume that firm 1's production output q (e.g., a certain chemical) – due to unavoidable pollution emissions – makes firm 2's production (e.g., catching fish) more difficult at the margin by requiring a costly pollution-abatement action z (e.g., water sanitation) by firm 2. Suppose that firm 1's profit $\pi^1(q)$ is concave in q and such that $\pi^1(0) = \pi^1(\bar{q}) = 0$ for some $\bar{q} > 0$. Firm 2's profit $\pi^2(q, z)$ is decreasing in q , concave in z , and has (strictly) increasing differences in (q, z) (i.e., $\partial^2\pi^2/\partial q\partial z > 0$), reflecting the fact that its marginal profit $\partial\pi^2/\partial z$ for the abatement action is increasing in the pollution level. Note that while firm 2's payoff depends on firm 1's action, firm 1 is unconcerned about what firm 2 does. Thus, without any outside intervention, firm 1 chooses an output $q^* \in \arg \max_{q \geq 0} \pi^1(q)$ that maximizes its profit. Firm 2, on the other hand, takes q^* as given and finds its optimal abatement response, $z^* \in \arg \max_{z \geq 0} \pi^2(q^*, z)$. In contrast to this, the socially optimal actions, \hat{q} and \hat{z} , are such that they maximize joint payoffs (corresponding to 'social welfare' in this simple model), i.e.,

$$(\hat{q}, \hat{z}) \in \arg \max_{(q, z) \geq 0} \{ \pi^1(q) + \pi^2(q, z) \}.$$

It is easy to show that the socially optimal actions \hat{q} and \hat{z} are at strictly lower levels than the privately optimal actions q^* and t^* : because firm 1 does not perceive the social cost of its actions, the world is over-polluted in this economy; a market for the direct negative externality from firm 1 on firm 2 is missing.

A regulator can intervene and restore efficiency, e.g., by imposing a tax on firm 1's output q which internalizes the social cost of its externality. To accomplish this, first consider the "harm" of firm 1's action, defined as

$$h(q) = \pi^2(q, z^*(0)) - \pi^2(q, z^*(q)),$$

where $z^*(q)$ is firm 2's best abatement action in response to firm 1's choosing an output of q . Then if firm 1 maximizes its profit minus the harm $h(q)$ it causes to firm 1, we obtain an efficient outcome, since

$$\hat{q} \in \arg \max_{q \geq 0} \{ \pi^1(q) - h(q) \} = \arg \max_{q \geq 0} \left\{ \pi^1(q) + \max_{z \geq 0} \pi^2(q, z) \right\}.$$

Hence, by imposing a per-unit excise tax of $\tau = h'(\hat{q})$ (equal to the marginal harm at the socially optimal output \hat{q}) on firm 1, a regulator can implement a Pareto-efficient outcome.³³ Alternatively, firm 1 could simply be required (if necessary through litigation) to pay the total amount $h(q)$ in damages when producing an output of q , following a "Polluter Pays Principle."³⁴

The following seminal result by Coase (1960) states that even without direct government intervention an efficient outcome may result from bargaining between parties, which includes the use of transfer payments.

³²In contrast to the direct externalities where the payoff of one agent depends on another agent's action or choice, so-called "pecuniary externalities" act through prices. For example, in a standard exchange economy the fact that one agent demands a lot of good 1 means that the price of that good will increase, exerting a negative pecuniary externality on those agents.

³³This method is commonly referred to as *Pigouvian taxation* (Pigou 1920).

³⁴This principle is also called Extended Polluter Responsibility (EPR) (Lindhqvist 1992).

THEOREM 3 (COASE THEOREM) *If property rights are assigned and there are no informational asymmetries, costless bargaining between agents leads to a Pareto-efficient outcome.*

The intuition for this result becomes clear within the context of our previous example. Assume that firm 1 is assigned the right to produce, regardless of the effect this might have on firm 2. Then firm 2 may offer firm 1 an amount of money, A , to reduce its production (and thus its negative externality). Firm 1 agrees if

$$\pi^1(\hat{q}) + A \geq \pi^1(q^*),$$

and firm 2 has an incentive to do so if

$$\pi^2(\hat{q}, \hat{z}) - A \geq \pi^2(q^*, z^*).$$

Any amount A between $\pi^1(q^*) - \pi^1(\hat{q})$ and $\pi^2(\hat{q}, \hat{z}) - \pi^2(q^*, z^*)$, i.e., when

$$A = \lambda (\pi^1(\hat{q}) + \pi^2(\hat{q}, \hat{z}) - \pi^1(q^*) - \pi^2(q^*, z^*)) + (\pi^1(q^*) - \pi^1(\hat{q})), \quad \lambda \in [0, 1],$$

is acceptable to both parties. Conversely, if firm 2 is assigned the right to a pollution-free environment, then firm 1 can offer firm 2 an amount

$$B = \mu (\pi^1(\hat{q}) + \pi^2(\hat{q}, \hat{z}) - \pi^1(0) - \pi^2(0, z^*(0))) + (\pi^2(0, z^*(0)) - \pi^2(\hat{q}, \hat{z})), \quad \mu \in [0, 1],$$

to produce at the efficient level \hat{q} . The choice of the constant λ (or μ) determines which party ends up with the gains from trade and is therefore subject to negotiation.

Instead of assigning the property rights for pollution (or lack thereof) to one of the two parties, the government may issue marketable pollution permits which confer the right to pollute. If these permits can be traded freely between firm 1 and firm 2, then, provided that there are at least \hat{q} permits issued, the resulting Walrasian equilibrium (cf. Section 3.1) yields a price equal to the marginal harm $h'(\hat{q})$ at the socially optimal output. Hence, efficiency is restored through the creation of a market for the externality.

The preceding discussion shows that without government intervention Adam Smith's "invisible hand" might in fact be absent when there are externalities. Stiglitz (2006) points out that "the reason that the invisible hand often seems invisible is that it is often not there." Indeed, in the presence of externalities the Pareto-efficiency property of Walrasian equilibria generally breaks down. An example is when a good which can be produced at a cost (such as national security or a community radio program) can be consumed by all agents in the economy because they cannot be prevented from doing so. Thus, due to the problem with appropriating rents from this "public" good, the incentive for it is very low, a phenomenon that is often referred to as the "tragedy of the commons" (Hardin 1968). More precisely, a *public good* (originally termed "collective consumption good" by Samuelson (1954)) is a good that is *nonrival* (in the sense that it can be consumed by one agent and is still available for consumption by another agent) and *nonexcludable* (in the sense that it is not possible, at any reasonable effort, to prevent others from using it). Examples include radio waves or a public park. On the other hand, a *private good* is a good that is rival and excludable. All other goods are called *semi-public* (or *semi-private*). In particular, if a semi-public good is nonrival and excludable, it is called a *club good* (e.g., an electronic newspaper subscription or membership in an organization), and if it is rival and nonexcludable it is referred to as a *common good* (e.g., fish or freshwater). Table 1 provides an overview.

Good	<i>Excludable</i>	<i>Nonexcludable</i>
<i>Rival</i>	Private	Common
<i>Nonrival</i>	Club	Public

Table 1: Classification of Goods.

It is possible to extend the notion of Walrasian equilibrium to take into account the externalities generated by the presence of public goods in the economy. Consider N agents and M firms in a private-ownership economy as in Section 3.2, with the only difference that, in addition to L private goods, there are L_G public goods that are privately produced. Each agent i chooses a bundle ξ of the available public goods, and a bundle x^i of the private goods on the market. Each firm m can produce a vector y_G^m of public goods and a vector y^m of private goods, which are feasible if (y_G^m, y^m) is in this firm's production set Y^m . Lindahl (1919) suggested the following generalization of the Walrasian equilibrium (cf. Definition 4), which for our setting can be formulated as follows.

DEFINITION 6 *A Lindahl equilibrium $(\hat{p}_G, \hat{p}, \hat{\xi}, \hat{x}, \hat{y})$ in a private-ownership economy, with personalized prices $\hat{p}_G = (\hat{p}_G^1, \dots, \hat{p}_G^N)$ for the public good, is such that all agents maximize utility,*

$$(\hat{\xi}, \hat{x}^i) \in \arg \max_{(\xi, x^i) \in B((\hat{p}_G^i, \hat{p}), \hat{w}^i)} u^i(\xi, x^i), \quad i \in \{1, \dots, N\},$$

where $\hat{w}^i = \hat{p} \cdot \omega^i + \sum_{m=1}^M \vartheta_m^i (\sum_{i=1}^N \hat{p}_G^i \cdot \hat{y}_G^m + \hat{p} \cdot \hat{y}^m)$, all firms maximize profits,

$$\hat{y}^m \in \arg \max_{(y_G^m, y^m) \in Y^m} \left\{ \sum_{i=1}^N \hat{p}_G^i \cdot y_G^m + \hat{p} \cdot y^m \right\}, \quad m \in \{1, \dots, M\},$$

and the resulting allocation is feasible,

$$\hat{\xi} = \sum_{i=1}^M \hat{y}_G^i \quad \text{and} \quad \sum_{i=1}^n (\hat{x}^i - \omega^i) = \sum_{m=1}^M \hat{y}^m.$$

It can be shown that the Lindahl equilibrium exists (Foley 1970; Roberts 1973) and that it restores the efficiency properties of the Walrasian equilibrium in terms of the two fundamental welfare theorems (Foley 1970).³⁵ Because of personal arbitrage, as well as the difficulty of distinguishing different agents and/or of price discriminating between them, it may be impossible to implement personalized prices, which tends to limit the practical implications of the Lindahl equilibrium.

REMARK 6 Another approach for dealing with implementing efficient outcomes in the presence of externalities comes from game theory rather than general equilibrium theory. Building on insights on menu auctions by Bernheim and Whinston (1986), Prat and Rustichini (2003) examine a setting where each one of M buyers (“principals”) noncooperatively proposes a nonlinear payment schedule to each one of N sellers (“agents”). After these offers are known, each seller i then chooses to offer a bundle (“action”) $x_i^j \in \mathbb{R}_+^L$ for buyer j . Given that buyers’ utility functions can have full externalities, Prat and Rustichini show that there exists an equilibrium which implements an efficient outcome. Weber and Xiong (2007) generalize this finding under minimal assumptions to the fully general setting where both buyers and sellers care about every action taken in the economy. In this situation, prices become nonlinear functions of the quantities,

³⁵We omit the technical assumptions required for a precise statement of these theorems.

effectively extending the traditionally ‘linear’ notion of price (where twice as much of an item typically costs twice as much) in a full-information setting.³⁶ Jackson and Wilkie (2005) consider another generalization of Prat and Rustichini’s results, allowing for side payments between different market participants. \square

REMARK 7 We briefly mention the possibility of central planning as an alternative to resource allocation through markets. While the ‘Chicago School’ in economics with proponents such as Ronald Coase, Frank Knight, Friedrich von Hayek, George Stigler, Gary Becker or Milton Friedman, tends to favor the use of markets for the allocation of resources, the above-mentioned possibility of market failure due to the lack of an invisible hand in the presence of externalities might call for a compromise between markets and government intervention. The main drawback of centrally managed economies stems from the difficulty of collecting and aggregating information at a center (Hayek 1945). Shafarevich (1980) and Stiglitz (1994) provide additional interesting perspectives. \square

REMARK 8 Related to the last remark, when a social planner has full authority, desirable allocations may be implemented using either permits or quotas.³⁷ The question of *what* allocations may be desirable invokes the problem of fairness, which may be addressed by what maximizes the sum of all agents’ utilities (utilitarian solution; Mill 1863), what is best for the economically weakest agents (egalitarian solution; Rawls 1971), a robust intermediate allocation (relatively fair solution; Goel et al. 2009), or by what ensures that no agent would want to swap his allocation with another agent (envy-free solution; Foley 1967). \square

5.2 Nonmarket Goods

For many goods, such as public parks, clean air, the Nobel memorial prize, human organs, or public offices, there are no well-defined markets. Yet, these goods may be of considerable value to some, so that the question begs of how one should determine the value of a *nonmarket good* to a given individual. In contrast to market goods, it is not possible to find an “objective” value by looking up the good’s price. The good has no *a priori* price and its value, as we shall see below, depends on the individual with whom a transaction is to take place and possibly also on the direction of the transaction.³⁸

A general way of thinking about a nonmarket good is in terms of a ‘change of state’ between $s = 0$ and $s = 1$. In state 1 the good is present, whereas in state 0 it is absent. If x denotes a typical consumption bundle of market goods, then Joe’s utility of x in state $s \in \{0, 1\}$ is $u_s(x)$. Thus, given a wealth of $w > 0$ and a price vector p for the bundle of conventional market goods, by Eq. (8) Joe’s indirect utility is

$$v_s(p, w) = \max_{x \in B(p, w)} u_s(x), \quad s \in \{0, 1\}.$$

Hence, if Joe does not have the nonmarket good initially, he would be willing to pay any amount c that leaves him at the reduced wealth $w - c$ at least as well off as he was initially. Hence, the maximum amount Joe is willing to pay, given his initial wealth w , is

$$CV(w) = \sup\{c \in \mathbb{R} : v_1(p, w - c) \geq v_0(p, w)\}. \quad (19)$$

³⁶In Section 7, we see that in the presence of information asymmetries nonlinear price schedules arise naturally.

³⁷In the presence of uncertainty this equivalence may disappear, as noted by Weitzman (1974).

³⁸For a recent survey on nonmarket valuation see Champ et al. (2003).

The welfare measure CV is called Joe's *compensating variation* or *willingness to pay* (WTP) for the change from $s = 0$ to $s = 1$. Conversely, if Joe initially has the nonmarket good, then the smallest amount he is willing to accept, given his initial wealth w , is

$$EV(w) = \inf\{e \in \mathbb{R} : v_0(p, w + e) \geq v_1(p, w)\}. \quad (20)$$

The welfare measure EV is termed *equivalent variation* or *willingness to accept* (WTA) for the change from $s = 1$ to $s = 0$.

As an example, we consider as “the” classical application of these welfare measures, proposed by Hicks (1939), the case where the change of state corresponds to a change in the market price from p to \hat{p} (say, $\hat{p} \ll p$ in case of a price decrease). The equivalent variation measures how much Joe would be willing to give up in wealth in return for a change in the price vector from p to \hat{p} . From Eq. (19) we obtain (dropping any subscripts) that

$$\underbrace{v(\hat{p}, w - CV(w))}_{= v(\hat{p}, e(\hat{p}, U))} = \underbrace{v(p, w)}_{= v(p, e(p, U))},$$

where $e(p, U) = w$ is Joe's expenditure on market goods necessary to attain the utility level $U = u(x(p, w))$ (cf. Section 2.3). Hence, $w - CV(w) = e(p, U) - CV(w)$ must be equal to $e(\hat{p}, U)$, so that

$$CV(w) = e(p, U) - e(\hat{p}, U). \quad (21)$$

In a completely analogous manner one finds that

$$EV(w) = e(p, \hat{U}) - e(\hat{p}, \hat{U}), \quad (22)$$

where $\hat{U} = u(x(\hat{p}, w))$ is the utility level at the new price \hat{p} .³⁹

There is a simple and exact relation between compensating and equivalent variation, provided that both are finite (Weber 2003, 2010), one being a shifted version of the other,

$$EV(w - CV(w)) = CV(w) \quad \text{and} \quad EV(w) = CV(w + EV(w)). \quad (23)$$

REMARK 9 Nonmarket goods are by definition traded using contracts rather than markets. The expense of writing such contracts can be substantial, implying a nonnegligible “transaction cost” (Coase 1937). Williamson (1975; 1985) points out that transaction costs are driven by the specificity and uncertainty attached to the items to be transacted, as well as the bounded rationality and opportunistic behavior of the contracting parties. Transaction cost economics also

³⁹It is interesting to note that the compensating and equivalent variations in Eqs. (21) and (22) depend only on the initial price p and the final price \hat{p} , not on the path from one to the other. This is in contrast to the change in *consumer surplus*, $\Delta CS(w) = \int_p^{\hat{p}} x(p(\gamma)) \cdot d\gamma$, which depends on the path γ taken. Taking classical mechanics as an analogy, the variation in potential energy when moving a rigid body between two points in a (conservative) gravitational field depends only on these two points. Similarly, the expenditure function can be viewed as a ‘potential function’ that measures the variation of welfare for a price change only in terms of the beginning and ending prices. The underlying mathematical justification is that as a consequence of Roy's identity (cf. Section 2.3) the Hicksian demand is a gradient field (generated by the expenditure function), whereas Slutsky's identity (10) implies that (due to the wealth effect) Walrasian demand generally is not a gradient field, i.e., not integrable. (Integrability is characterized by the Frobenius theorem, whose conditions, $\partial x_i(p, w)/\partial p_j = \partial x_j(p, w)/\partial p_i$ for all i, j , are generally not satisfied for a given Walrasian demand.) It is easy to see that $\Delta CS(w)$ always lies between $CV(w)$ and $EV(w)$. Willig (1976) has shown that the difference between the two is likely to be small, so that the variation of consumer surplus (along any path) can often be considered a reasonable approximation for the welfare change. For more general state changes this observation becomes incorrect (see, e.g., Haneman 1991).

recognizes the fact that contracts are necessarily ‘incomplete’ in the sense that not all possible contingencies can be captured and therefore all ‘residual claims’ must be assigned to one party (Hart 1988; Hart and Moore 1990). Transaction cost economics can be used to help explain the boundary of the firm, which is determined by the decision of which goods and services are to be procured from within the firm and which from outside. \square

6 Strategic Pricing with Complete Information

So far, the price or value of a good has been determined in a nonstrategic way, either via an equilibrium in a competitive market with price-taking agents (cf. Section 3.2), or, for nonmarket goods, via indifference using Hicksian welfare measures (cf. Section 5.2). In general, the price can be influenced strategically by agents who have “market power.” Section 6.1 deals with the case where there is one such agent and Section 6.2 with the case where there are several. In Section 6.3, we briefly discuss what a single strategic firm in a market can do to prevent competitors from entering that market.

6.1 Monopoly Pricing

Consider a single firm that can set prices for its products and choose the best element of its production set Y (introduced in Section 3.2). We assume that it is possible to split the components of any feasible production vector $y \in Y$ into *inputs* z and *outputs* (or *products*) q , so that

$$y = (-z, q),$$

where $q, z \geq 0$ are appropriate vectors. Thus, in its production process the firm is able to distinguish clearly between its inputs and its outputs. The problem of finding the firm’s optimal price is typically solved in a *partial equilibrium* setting, where for any price vector p that the firm might choose, it expects to sell to a demand $q = D(p)$.⁴⁰

Given a vector of input prices w , the firm’s *cost* for producing an output q of goods is

$$C(q; w, Y) = \min \{w \cdot z : (-z, q) \in Y\}.$$

This cost corresponds to the minimum expenditure for producing a certain output q . It depends on the market prices of the inputs as well as on the shape of the production set. For convenience, it is customary to omit the parameters w and Y from the cost function $C(q)$. The firm’s *monopoly pricing problem* is to solve

$$\max_p \{p \cdot D(p) - C(D(p))\}. \quad (24)$$

The first-order necessary optimality condition is

$$pD'(p) + D(p) - C'(D(p))D'(p) = 0, \quad (25)$$

⁴⁰In contrast to the *general equilibrium* setting in Section 3 we neglect here the fact that in a closed economy the agents working at the firm have a budget set that depends on the wages that the firm pays, which in turn might stimulate these agents’ demand for the firm’s output. The latter is often referred to as the *Ford effect*, for it is said that Henry Ford used to pay higher wages to his employees in order to stimulate the demand for Ford automobiles (Bishop 1966, p. 658).

so that in the case of a single product we obtain the (*single-product*) *monopoly pricing rule*

$$\frac{p - \text{MC}}{p} = \frac{1}{\varepsilon}, \quad (26)$$

where $\text{MC}(p) = C'(D(p))$ is the firm's marginal cost, and $\varepsilon = -pD'(p)/D(p)$ is the (*own-price*) *demand elasticity* for the good in question. The value on the left-hand side is often referred to as the Lerner index (Lerner 1934). It represents the firm's market power: the relative mark-up a single-product monopolist can charge, i.e., the Lerner index, is equal to the multiplicative inverse of demand elasticity. Note that the pricing rule (26) also implies that the firm chooses to price in the elastic part of the demand curve, i.e., at a point p where the elasticity $\varepsilon \geq 1$.

When there are multiple products, the first-order condition (25) yields a *multi-product monopoly pricing rule* of the form

$$\frac{p_l - \text{MC}_l(p)}{p_l} = \frac{1}{\varepsilon_l} \left(1 - \sum_{j \neq l} \varepsilon_{lj} \left(\frac{p_j - \text{MC}_j(p)}{p_j} \right) \left(\frac{p_j D_j(p)}{p_l D_l(p)} \right) \right),$$

where $\varepsilon_{lj} = -(p_l/D_j(p))\partial D_l(p)/\partial p_j$ is the elasticity of demand for good l with respect to a change in the price of good j . ε_{ll} is the own-price elasticity for good l , whereas ε_{lj} for $l \neq j$ is called a *cross-price elasticity*. When the demand for good l increases with an increase of the price for good j , so that $\varepsilon_{lj} < 0$, then good l and good j are *complements* and the markup for good l tends to be below the single-product monopoly markup (but the firm expects to sell more of both products). Conversely, if the demand for good l decreases with an increase in the price of good j , so that $\varepsilon_{lj} > 0$, then the two goods are *substitutes*, which tends to decrease the price of good l compared to what a single-product monopolist would charge.

An example of complementary goods are the different components of a drum set (e.g., a snare drum and a snare-drum stand). *Perfect complements* are such that one cannot be used without the other, such as a left shoe and a right shoe which are therefore usually sold together in pairs.

Price Discrimination. The practice of selling different units of the same good at different prices to different consumers is called *price discrimination*. Consider, for example, Joe and Melanie's demand for movie tickets. If the seller could somehow know exactly how much each individual would want to pay for a ticket for a given movie (at a given movie theater, at a given time), *perfect* (or *first-degree*) *price discrimination* is possible and the seller is able to extract all surplus from consumers by charging each consumer exactly his willingness to pay. However, the willingness to pay for a good is in many cases private information of potential buyers and therefore not known to a seller, who therefore has to restrict attention to observable characteristics. For example, a movie theater may offer student tickets at a discounted price in order to charge prices based on an observable characteristic (possession of a student ID card).⁴¹ This is called *third-degree price discrimination*. Lastly, if the seller cannot discriminate based on observable characteristics, then it may still obtain a separation of consumers into different groups through product bundling. For example, the movie theater may give quantity discounts by selling larger numbers of movie tickets at a lower per-unit price. Alternately, it could offer discounts on tickets for shows during the day or on weekdays. This lets consumers select the product or bundle they like best. If Joe

⁴¹In some jurisdictions it is illegal for sellers to price discriminate based on certain observable characteristics such as age, gender, race, or certain other characteristics. For example, in the U.S. the Robinson-Patman Act of 1936 (Anti-Justice League Discrimination Act, 15 U.S.C. §13) is a federal law that prohibits price discrimination between equally situated distributors.

is a busy manager, then he might prefer the weekend show at a higher price, whereas Melanie may prefer to see the same movie at the ‘Monday-afternoon special discount price.’ The practice of offering different bundles of (similar) goods by varying a suitable instrument (e.g., quantity, time of delivery) is called *second-degree price discrimination* or *nonlinear pricing* (discussed in Section 7.1).

As an example for third-degree price discrimination, consider a firm which would like to sell a certain good to two consumer groups. The marginal cost of providing the good to a consumer of group i is equal to c_i . At the price p_i group i ’s demand is given by $D_i(p_i)$, so that the firm’s pricing problem (24) becomes

$$\max_{p_1, p_2} \left\{ \sum_{i=1}^2 (p_i - c_i) D_i(p_i) \right\}.$$

Thus, the firm applies the monopoly pricing rule (26) to each group, so that

$$\underbrace{p_i + \frac{D_i(p_i)}{D'_i(p_i)}}_{\text{MR}_i(p_i)} = \underbrace{c_i}_{\text{MC}_i} \quad (27)$$

for all $i \in \{1, 2\}$. Eq. (27) says that a monopolist sets in each independent market a price so as to equalize its marginal cost MC_i and its *marginal revenue* $\text{MR}_i(p_i)$ in that market.⁴²

6.2 Oligopoly Pricing

The presence of other firms tends to erode the market power of any given firm. The reason is that this firm needs to anticipate the actions of all other firms.

In general, we can think of each firm i as having a set A^i of possible actions (or *strategies*) available. The vector $a = (a^1, \dots, a^N) \in A = A^1 \times \dots \times A^N$, which represents the vector of actions of all N firms, is called a *strategy profile*. Each firm i ’s profit, $\pi^i(a)$, depends on its own action a^i and the vector $a^{-i} = (a^1, \dots, a^{i-1}, a^{i+1}, \dots, a^N)$ of all other firms’ actions.

DEFINITION 7 A (*pure-strategy*) *Nash equilibrium of the simultaneous-move game*⁴³ is a strategy profile $\hat{a} = (\hat{a}^1, \dots, \hat{a}^N) \in A$ such that

$$\hat{a}^i \in \arg \max_{a^i \in A^i} \pi^i(a^i, \hat{a}^{-i}), \quad i \in \{1, \dots, N\}. \quad (28)$$

⁴²In general, a firm’s marginal cost MC_i depends on the firm’s output and thus also on its price, just as its marginal revenue does. The marginal revenue is defined as the additional revenue the firm obtains when adding an infinitesimal unit of output. That is, $\text{MR}_i = \partial(p_i D_i) / \partial D_i = p_i + (D_i / D'_i)$, using the product rule of differentiation and the inverse function theorem. We quietly assumed that the firm’s marginal revenue is increasing in price (i.e., decreasing in quantity for all but Giffen goods). This assumption is not always satisfied, neither in theory nor even empirically (Beckman and Smith 1993), but it ensures a convex optimization problem (with concave objective function). We refer to it as the “regular case,” and return to this case in Section 7.3 to avoid technical complications.

⁴³A *game* is a collection of a set of players, a set of strategy profiles, and a set of payoff functions (one for each player, that maps any strategy profile to a real number). A *simultaneous-move game* is a game in which all players select their strategies at the same time. More information on game theory and its applications can be found in Kopalle and Shumsky (this volume).

Depending on if the firms set prices or quantities in their strategic interactions, the game is referred to either as Bertrand pricing game or as Cournot quantity-setting game (Bertrand 1883; Cournot 1838). For simplicity, we restrict attention to the special case in which each firm is producing a single homogeneous product.

Bertrand Competition. Consider the situation in which there is a unit demand and only two firms which compete on price, so that the firm with the lower price obtains all of the sales. Firm $i \in \{1, 2\}$ therefore faces the demand

$$D^i(p^i, p^j) = \begin{cases} 1, & \text{if } p^i < \min\{p^j, r\}, \\ 1/2, & \text{if } p^i = p^j \leq r, \\ 0, & \text{otherwise,} \end{cases}$$

as a function of its own price p^i and the other firm's price p^j (where $j \in \{1, 2\} \setminus \{i\}$). Firm i 's profit is $\pi^i(p^i, p^j) = (p^i - c)D^i(p^i, p^j)$. If firm j charges a price p^j greater than the firms' common marginal cost c , it is always best for firm i to slightly undercut firm j . If $p^j = c$, then by also charging marginal cost firm i cannot earn a positive profit, so that it becomes indifferent between charging $p^i = c$ or any price that is higher than the other firm's price (and which therefore results in zero sales). The unique Nash equilibrium strategy profile is described by $\hat{p}^i = \hat{p}^j = c$: both firms sell at marginal cost and make no profits at all. The firms are "trapped" in their strategic interaction: they cannot take advantage of technological innovation, i.e., cutting c in half has no effect on their profits (Cabral and Villas-Boas 2005). In particular, both firms charge marginal cost and dissipate all of their profits. None of the firms can make any money. It is interesting to note that the zero-profit outcome persists when we add firms with identical marginal costs, and even if the marginal costs decrease for all firms from c to $\hat{c} < c$. This pessimistic outlook changes somewhat when the marginal costs are different across firms. In that case, there is generally a continuum of Nash equilibria that allow for prices between the most efficient firm's marginal cost and the second-most efficient firm's marginal cost (minus epsilon).

The extreme price competition in a Bertrand oligopoly is softened when firms are selling differentiated products. Hotelling (1929) introduced a by now classical model where two different sellers of the same good are located at two points on a line segment while consumers are distributed on a line segment. Each consumer faces a "transportation cost" to go to one of the two sellers, which allows them to charge a price premium as long as their locations are different. This "spatial" model of product differentiation is "horizontal," since different consumers may have quite different preferences for what is essentially the same good (once it is purchased and brought back home). Shaked and Sutton (1982) allow for vertical product differentiation, which also softens price discrimination.⁴⁴ Another reason why Bertrand oligopolists may obtain positive markups in equilibrium is that firms may have capacity constraints (Osborne and Pitchik 1986), so that the demand needs to be rationed, which can lead to positive rents.

Cournot Competition. If instead of price firm i chooses its production quantity q^i , and at the aggregate quantity $Q(q) = q^1 + \dots + q^N$ the (inverse) market demand is $p(Q) = D^{-1}(Q)$, then its profit is

$$\pi^i(q^i, q^{-i}) = (p(Q(q^i, q^{-i})) - c^i)q^i.$$

Each firm i can then determine its optimal quantity (also referred to as its "best response") as a function of the vector q_{-i} of the other firms' actions. For example, when the inverse demand

⁴⁴Beath and Katsoulacos (1991) give a survey of the economics of product differentiation in a locational setting. Anderson et al. (1992) give a survey in an alternative, random discrete-choice setting.

curve is linear, i.e., when $p(Q) = a - bQ$ for some $a, b > 0$, then the N conditions in Eq. (28) yield firm i 's Nash equilibrium quantity,

$$\hat{q}^i = \frac{a - c^i + \sum_{j=1}^N (c^j - c^i)}{(N + 1)b},$$

provided that all firms produce a positive quantity in equilibrium.⁴⁵ The aggregate industry output becomes

$$\hat{Q} = \frac{N(a - \bar{c})}{(N + 1)b},$$

where \bar{c} is the average marginal cost of active firms in the industry. As the number of firms in the industry increases, the equilibrium price $\hat{p} = p(\hat{Q})$ tends toward the average marginal cost,⁴⁶

$$\hat{p} = \frac{(a/N) + \bar{c}}{1 + (1/N)} \rightarrow \bar{c} \quad (\text{as } N \rightarrow \infty).$$

As in Section 6.1, each firm's market power can be measured in terms of the relative markup it is able to achieve. The more firms in the industry, the closer the market price will approximate the industry's lowest (constant) cost, forcing companies with higher marginal costs to exit. Eventually, the relative markup will therefore approach zero.

A practical measure of the competitiveness (or concentration) of an industry is the Herfindahl-Hirschman index (HHI); it measures industry concentration as the sum of the squares of the firms' market shares $\sigma^i = \hat{q}^i/\hat{Q}$ (Herfindahl 1950; Hirschman 1964),⁴⁷

$$\text{HHI} = \sum_{i=1}^N (\sigma^i)^2 \in [0, 1].$$

REMARK 10 From an empirical point of view (see, e.g., Baker and Bresnahan 1988; 1992) it is useful to identify the relative markup in an industry and compare it to the extreme cases of perfect competition (zero relative markup) on one side, and the monopoly pricing rule in Eq. (26) on the other side. For example, the pricing rule in a symmetric N -firm oligopoly is

$$\frac{p - c}{p} = \frac{\theta}{\varepsilon},$$

where $\theta = \text{HHI} = 1/N$ and ε is the demand elasticity (cf. Section 6.1). In a monopoly $\theta = 1$, and under perfect competition $\theta = 0$. Thus, an empirical identification of θ allows for a full "conjectural variation," as *a priori* it can take on any value between zero and one. \square

REMARK 11 In a Cournot oligopoly firms might be "trapped" in their competitive interaction just as in the Bertrand oligopoly: it may be to all firms' detriment when costs decrease. This

⁴⁵Firms that prefer not to produce anything will be the ones with the highest marginal costs. Thus, after labelling all firms such that $c^1 \leq c^2 \leq \dots \leq c^N$ we obtain that $\hat{q}^i = (a - c^i)/b - (1/(1 + |I|)) \sum_{j \in I} (a - c^j)/b$, where $I = \{i \in \{1, \dots, N\} : (1 + i)(a - c^i) > \sum_{j=1}^i (a - c^j)\}$ is the set of participating firms.

⁴⁶In the symmetric case, when all the firms' marginal costs are the same, i.e., when $c^1 = \dots = c^N$, the equilibrium market price therefore tends toward marginal cost as the number of firms in the market increases.

⁴⁷The HHI is used by the antitrust division of the U.S. Department of Justice. Markets where $\text{HHI} \leq 10\%$ are generally considered as competitive (or unconcentrated). A market for which $\text{HHI} \geq 18\%$ is considered uncompetitive (or highly concentrated). Mergers or acquisitions which stand to change the HHI of an industry by more than 1% tend to raise antitrust concerns.

was pointed out by Seade (1985), who showed that an increase of excise taxes may increase all firms' profits in a Cournot oligopoly. The intuition for this surprising result is that a higher cost may shift firms' output into a less elastic region of the consumers' demand curve, so that the firms' resulting price increase overcompensates them for their increased cost.⁴⁸ In this (somewhat pathological) situation, the consumers therefore bear more than the change in total firm profits, in terms of their welfare losses. \square

6.3 Entry Deterrence and Limit Pricing

In general it is unrealistic to assume that strategic actions in an industry are taken simultaneously. An incumbent firm in an industry can try to protect its monopoly position by discouraging entry by other firms, for example by pricing below the long-run average cost, a practice which is called *limit pricing*.⁴⁹ If some fixed cost is required to enter the industry, a potential entrant may therefore be discouraged. The flip side of this is of course that the mere existence of *potential* entrants tends to limit the pricing power of an incumbent monopolist firm as a function of the required upfront investment. If the latter is close to zero, the markup a monopolist can charge also tends to zero.

Sometimes there are other activities that a monopolist can pursue to keep out competitors, such as advertising and/or building a loyal consumer base. These so-called *rent seeking* activities (Tullock 1980) are generally costly. In principle, the monopolist is willing to spend all of its extra profit (of being a monopolist instead of an oligopolist) pursuing wasteful rent-seeking activities to obtain the extra profit (Posner 1975). Naturally, if one includes the intermediaries providing rent-seeking services, such as search advertising, into the system, then these activities may (even though not fully efficient) not be completely wasteful (Weber and Zheng 2007).

Sometimes the incumbent may resort to threats that *if* a competitor should enter, *then* a price war would be started that would have detrimental effects on the entrant, as they would not be able to recoup any upfront cost (which is already sunk for the incumbent). However, at times such threats may not be credible. For example, if a large incumbent in telecommunications is faced with a small capacity-constrained entrant, it may be better for the monopolist to ignore the entrant, allowing it to obtain a small market share at discount prices while it still retains a lion's share of the market at monopoly prices. The accommodated entrant can then work on building its consumer base and increasing its capacity, and gradually become a more serious threat to the incumbent (Gelman and Salop 1983).

REMARK 12 It is not always in an incumbent firm's best interest to discourage entry by other firms. For example, when other firms' products are complements, their presence would in effect exert a positive externality on the incumbent which may prompt the latter to encourage the entry of such 'complementors' (Economides 1996). For example, a firm that wishes to sell a certain computer operating system may want to encourage the entry of software companies in the market that directly compete with certain of its own products (e.g., browser software) but at the same time render the operating system more valuable to consumers. \square

⁴⁸As an example, one can set $N = 2$, $c^1 = c^2 = c$, and $p(Q) = 100Q^{-3/2}$ in our model, and then consider the increase in marginal cost from $c = 10$ to $c = 20$, which proves to be beneficial for the firms.

⁴⁹It is possible to consider limit pricing also in an oligopolistic setting (Bagwell and Ramey 1981).

7 Strategic Pricing with Incomplete Information

The optimal pricing problem becomes more delicate if there are informational asymmetries between a seller and potential buyers. For example, the buyer's utility function may be unknown to the seller, rendering it impossible to extract all surplus. In the resulting *screening problem*, the seller tries to optimally construct a menu of options for the buyer, so that the buyer through his selection of an option reveals his private information. Or, it may be that the seller has private information about the quality of the good that is being sold, and the resulting *signaling problem* is to convey that information in a credible manner to a potential buyer. When there are several different potential buyers with private information, the seller may be able to solve a *mechanism design problem* so as to extract surplus from the buyers or to implement an efficient allocation of the goods. Lastly, information asymmetries play a role when there are several sellers who may or may not have an incentive to share some of their private information.

7.1 Screening

Consider a seller who would like to sell a bundle of L products $x = (x_1, \dots, x_L)$ (or, equivalently, a single product with L attributes; Lancaster 1966). If the seller charges a price p for this bundle, then the consumer's net utility is

$$u(x, \theta) - p,$$

where $\theta \in \Theta = [\underline{\theta}, \bar{\theta}] \subset \mathbb{R}$ represents the consumer's private information (e.g., the marginal utility) and $u : \mathbb{R}_+^L \times \Theta \rightarrow \mathbb{R}$ is his utility function. The problem the seller has is that each consumer type θ may have different preferences, so that it is generally impossible for him to extract all of a given consumer's surplus, since that surplus is by hypothesis unknown. In order to extract information from the buyer, the seller needs to offer a menu of bundles including a price $p(x)$ for each item on this menu.

The seller's problem is to design an economic *mechanism* $\mathcal{M} = (\hat{\Theta}, \rho)$ that consists of a 'message space' $\hat{\Theta}$ and an 'allocation function' $\rho = (\xi, \tau)$ that maps each element $\hat{\theta}$ of this message space to a product bundle $\xi(\hat{\theta}) \in \mathbb{R}_+^L$ and a price (monetary transfer) $\tau(\hat{\theta}) \in \mathbb{R}$. Given such a mechanism \mathcal{M} , the sequence of events is as follows. The seller proposes $(\hat{\Theta}, \rho)$ to the buyer, meaning that she explains to the buyer that all product choices are coded by elements in the message space (e.g., using tags on items displayed in a store). By sending a message $\hat{\theta} \in \hat{\Theta}$ the buyer effectively points to the item he wants (including the almost always feasible choice of "no item," reflecting the buyer's voluntary participation in the mechanism), and has to pay the transfer $\tau(\hat{\theta})$ to the seller.

The problem of finding an optimal mechanism seems daunting, since there are in principle many equivalent ways to select a message and at least as many ways of defining an allocation function. The following trivial result, generally attributed to Gibbard (1973) and Myerson (1979), considerably simplifies the search for an optimal mechanism.

THEOREM 4 (REVELATION PRINCIPLE) *For any mechanism $\mathcal{M} = (\hat{\Theta}, \rho)$ there exists a 'direct revelation mechanism' $\mathcal{M}_d = (\Theta, \rho_d)$ such that if an agent of type θ finds it optimal to send a message $\hat{\theta}$ under mechanism \mathcal{M} , he finds it optimal to send the fully revealing message θ under \mathcal{M}_d and consequently obtains the identical allocation $\rho_d(\theta) = \rho(\hat{\theta})$.*

It is very easy and instructive to see why this result must hold. Under the mechanism \mathcal{M} agent θ solves

$$\hat{\theta}^*(\theta) \in \arg \max_{\vartheta \in \hat{\Theta}} \{u(\xi(\vartheta), \theta) - \tau(\vartheta)\}. \quad (29)$$

Thus, by setting $\rho_d(\theta) = (\xi(\hat{\theta}^*(\theta)), \tau(\hat{\theta}^*(\theta)))$ the revelation principle follows immediately. The intuition is that the mechanism designer can simulate the buyer's decision problem and change his mechanism to a direct revelation mechanism accordingly.

By using the revelation principle, the seller can – without any loss in generality – consider only direct revelation mechanisms, for which $\hat{\Theta} = \Theta$ and $\hat{\theta}^*(\theta) = \theta$. To simplify the solution of the screening problem, we assume that $u(x, \theta)$ is increasing in θ , and that the Spence-Mirrlees “sorting condition”⁵⁰

$$u_{x\theta}(x, \theta) \geq 0 \quad (30)$$

is satisfied for all (x, θ) . The (direct) mechanism $\mathcal{M} = (\Theta, \rho)$ with $\rho = (\xi, \tau)$ is *implementable*, i.e., is a direct revelation mechanism, if $\hat{\theta}^*(\theta) = \theta$ holds for all $\theta \in \Theta$. The first-order necessary optimality condition for the corresponding ‘incentive-compatibility condition’ (29) is

$$u_x(\xi(\theta), \theta) \xi'(\theta) = \tau'(\theta), \quad \underline{\theta} \leq \theta \leq \bar{\theta}, \quad (31)$$

and the second-order necessary optimality condition is

$$u_{x\theta}(\xi(\theta), \theta) \xi'(\theta) \geq 0, \quad \underline{\theta} \leq \theta \leq \bar{\theta}. \quad (32)$$

The seller's payoff from selling the bundle $x = \xi(\theta)$ to a consumer of type θ at the price $p(x) = \tau(\theta)$ under the implementable mechanism \mathcal{M} is

$$\pi(x, \theta) + p(x).$$

Again, we assume that the seller's payoff satisfies a Spence-Mirrlees sorting condition, so that

$$\pi_{x\theta}(x, \theta) \geq 0$$

for all (x, θ) . Given some beliefs about distribution of the (from her perspective random) con-

⁵⁰Subscripts denote partial derivatives. For example, $u_{x\theta} = \partial^2 u / \partial x \partial \theta$.

sumer type $\tilde{\theta} \in \Theta$ with the cdf $F(\theta) = P(\tilde{\theta} \leq \theta)$, the seller's expected payoff is⁵¹

$$\begin{aligned} \bar{\pi}(\xi, \tau) &= \int_{\underline{\theta}}^{\bar{\theta}} (\pi(\xi(\theta), \theta) + \tau(\theta)) dF(\theta) \\ &= \int_{\underline{\theta}}^{\bar{\theta}} \underbrace{(\pi(\xi(\theta), \theta) + u(\xi(\theta), \theta))}_{\text{Total Surplus}} - \underbrace{\frac{1 - F(\theta)}{F'(\theta)} u_{\theta}(\xi(\theta), \theta)}_{\text{Information Rent}} dF(\theta). \end{aligned} \quad (33)$$

Virtual Surplus ($\equiv S(\xi(\theta), \theta)$)

The previous relation shows that for each type θ the seller would like to maximize a so-called 'virtual surplus' $S(\xi(\theta), \theta)$ which consists of the total surplus $W(\xi(\theta), \theta) = \pi(\xi(\theta), \theta) + u(\xi(\theta), \theta)$ minus a nonnegative 'information rent.' The latter describes the discount that a consumer of type θ obtains compared to perfect price discrimination which the seller could implement under complete information (cf. Section 6.1).

Using basic insights from monotone comparative statics (cf. Section 2.3), it is possible to obtain a characterization of solutions to the seller's expected-profit maximization problem. Indeed, if we assume that π and u are supermodular, the hazard rate $h = F'/(1 - F)$ is nondecreasing, and u_{θ} is submodular, then the virtual surplus $S(x, \theta)$ is supermodular in (x, θ) , so that the pointwise optimal solution $\xi(\theta)$ is nondecreasing in θ , and the second-order condition (32) is automatically satisfied. The first-order condition (31) was used to write the seller's expected profit in terms of virtual surplus (cf. Footnote 51) and is therefore also satisfied. Hence, the optimal solution $\xi(\theta)$ to the seller's screening problem satisfies

$$S_x(\xi(\theta), \theta) = \pi_x(\xi(\theta), \theta) + u_x(\xi(\theta), \theta) - (u_{x\theta}(\xi(\theta), \theta)/h(\theta)) = 0 \quad (34)$$

for all $\theta \in \Theta$ for which $x(\theta) > 0$. Realizing that $\tau(\underline{\theta}) = 0$, we obtain from Eq. (31) that

$$\tau(\theta) = \int_{\underline{\theta}}^{\theta} u_x(\xi(\vartheta), \vartheta) \xi'(\vartheta) d\vartheta. \quad (35)$$

Once (ξ, τ) have been obtained, the seller can use the so-called 'taxation principle' (see, e.g., Rochet (1985)) to determine the optimal nonlinear pricing scheme $p(x)$ for the different bundles $x \in \xi(\Theta)$:

$$p(x) = \begin{cases} \tau(\theta), & \text{if } \exists \theta \in \Theta \text{ s.t. } \xi(\theta) = x, \\ \infty, & \text{otherwise.} \end{cases} \quad (36)$$

⁵¹This identity is obtained, using Eq. (31) and performing twice an integration by parts, as follows:

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \tau(\theta) dF(\theta) &= \int_{\underline{\theta}}^{\bar{\theta}} \left(\int_{\underline{\theta}}^{\theta} \tau'(s) ds \right) dF(\theta) = \int_{\underline{\theta}}^{\bar{\theta}} \left(\int_{\underline{\theta}}^{\theta} u_x(\xi(s), s) \xi'(s) ds \right) dF(\theta) \\ &= \left(\int_{\underline{\theta}}^{\theta} u_x(\xi(s), s) \xi'(s) ds \right) F(\theta) \Big|_{\underline{\theta}}^{\bar{\theta}} - \int_{\underline{\theta}}^{\bar{\theta}} u_x(x(\theta), \theta) \xi'(\theta) F(\theta) d\theta \\ &= \int_{\underline{\theta}}^{\bar{\theta}} (1 - F(\theta)) \left(\frac{du(\xi(\theta), \theta)}{d\theta} - u_{\theta}(\xi(\theta), \theta) \right) d\theta \\ &= (1 - F(\theta)) u(\xi(\theta), \theta) \Big|_{\underline{\theta}}^{\bar{\theta}} + \int_{\underline{\theta}}^{\bar{\theta}} \left(u(\xi(\theta), \theta) - \frac{1 - F(\theta)}{F'(\theta)} u_{\theta}(\xi(\theta), \theta) \right) dF(\theta), \end{aligned}$$

where the firm naturally sets $\xi(\underline{\theta})$ such that $u(\xi(\underline{\theta}), \underline{\theta})$ vanishes.

To understand the properties of the solution to the screening problem, we consider the example where $\pi(x, \theta) = -cx^2/2$ for some $c > 0$, $u(x, \theta) = \theta x$, and $F(\theta) = \theta$ on $\Theta = [0, 1]$, satisfying all of the above assumptions. The virtual surplus is $S(x, \theta) = -cx^2/2 + \theta x - (1 - \theta)x$, so that the optimality condition (34),

$$S_x(x, \theta) = -cx + \theta - (1 - \theta) = 0,$$

is satisfied for $x = \xi(\theta) \equiv (2\theta - 1)_+/c$. From Eq. (35) we therefore obtain

$$\tau(\theta) = \frac{2}{c} \int_0^\theta \mathbf{1}_{\{\vartheta \geq 1/2\}} \vartheta d\vartheta = \frac{(4\theta^2 - 1)_+}{4c}.$$

The taxation principle (36) eliminates the type θ from the expressions for $\xi(\theta)$ and $\tau(\theta)$, returning the optimal nonlinear price,

$$p(x) = \begin{cases} (2 + cx)(x/4), & \text{if } x \in [0, 3/(4c)], \\ \infty, & \text{otherwise.} \end{cases}$$

In this example, we observe a number of regularities that remain valid for other parameterizations. First, the highest consumer type $\theta = \bar{\theta}$ obtains a product bundle $\bar{x} = \xi(\bar{\theta})$ that is efficient, in the sense that \bar{x} maximizes $\pi(x, \bar{\theta}) + u(x, \bar{\theta})$, i.e., the sum of the consumer's and the seller's payoff. In other words, there is 'no distortion at the top.' In general this is true because $S(x, \bar{\theta})$ is precisely the social surplus. Second, the lowest type $\theta = \underline{\theta}$ obtains zero surplus. In the example this is true for all $\theta \in [0, 1/2]$. In general, we have 'full rent extraction at the bottom,' because $u_\theta \geq 0$ would otherwise imply that one could increase the price for all consumers without losing any participation. In the terminology of Section 5.1 the low types through their presence exert a positive externality on the high types. Hence, a (first-order stochastically dominant) shift in the type distribution that increases the relative likelihood of high types vs. low types tends to decrease the information rent they can obtain.

REMARK 13 The screening problem (first examined as such by Stiglitz (1975)) has many applications, such as optimal taxation (Mirrlees 1971), the regulation of a monopolist with an unknown cost parameter (Baron and Myerson 1985), or nonlinear pricing (Mussa and Rosen 1978). The solution to the screening problem without the somewhat restrictive assumptions that guaranteed monotone comparative statics of the solution to Eq. (34) needs the use of optimal control theory (Hamiltonian approach; Guesnerie and Laffont 1984). For $L = 1$, when the solution to Eq. (34) would be decreasing, one needs to implement Mussa and Rosen's 'ironing' procedure, which requires that $\xi(\theta)$ becomes constant over some interval, thus serving those consumer types with identical product bundles (typically referred to as 'bunching'). For $L > 1$, the second-order condition (32) imposes only an average restriction on the slope, which allows avoiding inefficient bunching more easily. An extension of the model to multi-dimensional types θ is nontrivial, as it becomes more difficult to characterize the set of all direct revelation mechanisms over which the seller has to search for an optimal one (Rochet and Stole 2003). \square

7.2 Signaling

In the screening problem, the mechanism operator (referred to as *principal*), who is uninformed, has the initiative and designs the mechanism. If the party with the private information has the initiative, then it faces a *signaling problem*. For example, if Joe tries to sell his car, then it is reasonable to assume that he has important private information about the state of the

car. However, if all he can do is “speak” to the prospective buyer, then (in the absence of any additional guarantees) the latter has no reason to believe what Joe says and can dismiss his sales pitch merely as “cheap talk” (Crawford and Sobel 1982; Farrell and Rabin 1996). However, the seller may be able to convey information when using a costly signal (Spence 1973). More precisely, the cost of the signal has to vary with the private information of its sender, for it to become informative. If each sender ‘type’ sends a different message in such a ‘signaling game,’ then the signaling equilibrium is said to be ‘separating.’ Otherwise it is called ‘pooling’ (or ‘semi-separating,’ if only a subset of types are separated).

Let us discuss a simple signaling game, adapted from Laffont (1989), to highlight the connection between signaling and screening.⁵² More specifically, assume that a seller would like to sell a unit of a product of quality θ known only to him. The seller can offer an observable warranty level of e . Both buyer and seller know that the cost to a type- θ seller of offering the warranty level e is $c(e, \theta)$, which is increasing in e . Suppose further that the cost satisfies the Spence-Mirrlees condition

$$c_{e\theta} \leq 0.$$

That is, the marginal cost c_e of providing a warranty decreases with the quality of the good. Let us now examine the seller’s incentives to report the true quality value θ to the buyer, given that the buyer believes his message.⁵³

If in a truth-telling (i.e., separating) equilibrium the buyer buys a unit amount $\xi(\theta) = 1$ at the price $\tau(\theta)$, then the seller’s profit is

$$\pi(x, p, \theta) = \tau(\theta) - c(e, \theta).$$

The first-order condition for truth-telling is

$$c_e e' = \tau', \tag{37}$$

and the second-order condition

$$c_{e\theta} e' \leq 0. \tag{38}$$

If the quality of the good θ represents the gross utility for the product, then in a separating equilibrium the seller is able to charge a fair price $\tau(\theta) = \theta$, extracting all of the buyer’s surplus. Hence, Eq. (37) yields that

$$e'(\theta) = \frac{1}{c_e(e(\theta), \theta)}.$$

For example, let $\underline{\theta} = 1$, $c(e, \theta) = 1 - \theta e$. Then $c_e = -\theta$, and therefore $e' = -1/\theta$, and $e(\theta) = \ln(\bar{\theta}/\theta) + e(\bar{\theta})$. The notion of a “reactive equilibrium” (Riley 1979) implies that $e(\bar{\theta}) = 0$, and thus $e(\theta) = \ln(\bar{\theta}/\theta)$. There may be many other separating equilibria; there may also be pooling equilibria, in which the seller would offer any quality level at the same price.

REMARK 14 In the signaling model presented here, the warranty offered has no effect on the buyer’s utility whatsoever. The only important feature is that the marginal cost of providing

⁵²The classic setting, explored by Spence (1973), deals with signaling in a labor market. A full analysis of signaling games is complicated by a generic equilibrium multiplicity, even in the simplest models. To reduce the number of equilibria, in view of sharpening the model predictions, a number of ‘equilibrium refinements,’ such as the ‘intuitive criterion’ by Cho and Kreps (1987), have been proposed, with mixed success.

⁵³If the seller does not tell the truth under these circumstances, the buyer, of course, will not trust the seller in equilibrium. We omit the details for the precise specification of the equilibrium.

warranty depends monotonically on the types. This insight can be transferred to other application domains such as advertising. The latter – while still just a “conspicuous expenditure” – can become an informative signal about product quality if the firms’ costs satisfy some type of relaxed Spence-Mirrlees condition (Kihlstrom and Riordan 1984). With multiple time periods to consider, even product prices can become signals (Bagwell and Riordan 1991), since a low-quality firm with a high price would face a much larger drop in sales than a high-quality firm, as the qualities become known to consumers. See Bagwell (2007) for an overview of the related economics of advertising. Kaya and Özer (this volume) discuss signaling (and screening) in the context of supply-chain contracting. \square

7.3 Mechanism Design

In the screening problem an uninformed party can design an economic mechanism to extract information from an informed party and implement its payoff-maximizing allocation of resources.⁵⁴ An economic mechanism can generally involve more than one agent, in addition to the mechanism operator (principal). As a leading example, consider Vickrey’s (1961) celebrated second-price auction. In that mechanism N bidders (agents) jointly submit their bids to the auctioneer (principal), who awards a single object to the highest bidder at a price equal to the second-highest submitted bid. It is easy to see that if every bidder $i \in \{1, \dots, N\}$ has a private value $\theta^i \in \Theta \subset \mathbb{R}_+$ for the object, it is best for that bidder to submit his true value as bid b_i . This leads to a unique equilibrium “bidding function” $\beta : \Theta \rightarrow \mathbb{R}$ which describes each player i ’s bid, $b_i = \beta(\theta_i) \equiv \theta_i$, as a function of his type θ_i , independent of the players’ beliefs.⁵⁵

We illustrate the design of a mechanism⁵⁶ with the example of an optimal auction (Myerson 1981) that, on the one hand, parallels third-degree price discrimination (cf. Section 6.1), and, on the other hand, employs the same methods as our solution to the screening problem. Consider a seller who would like to sell a single indivisible item to either Joe (= agent 1) or Melanie (= agent 2) so as to maximize his expected revenues. As in the second-price auction, we assume that both Joe’s and Melanie’s value for the item is private. Suppose that this seller thinks that Joe’s valuation θ_1 for the item is distributed with the cdf F_1 on the support $\Theta_1 = [\underline{\theta}_1, \bar{\theta}_1]$, while Melanie’s valuation θ_2 is distributed with the cdf F_2 on the support $\Theta_2 = [\underline{\theta}_2, \bar{\theta}_2]$. Suppose further that the seller’s opportunity cost for selling the item to agent i is $c_i \geq 0$.

Consider first the problem of how the seller would price to the two individuals so as to maximize expected profit from either of them, just as in third-degree price discrimination. At a price $p_i \in [0, 100i]$, the demand for the seller’s good by individual i is $D_i(p_i) = 1 - F_i(p_i)$. Hence, the monopoly pricing rule (27) in Section 6.1 yields that the seller’s optimal price r_i to agent i is

⁵⁴This allocation is referred to as ‘second-best,’ because it is subject to information asymmetries and therefore inferior in terms of the principal’s expected payoff under the ‘first-best’ allocation when information is symmetric. The latter is the case if *either* the agent is ex-ante uninformed about his own type (so that the principal can extract all his expected surplus by an upfront fixed fee) *or* both parties are informed about the agent’s type (so that the principal can practice first-degree price discrimination; cf. Section 6.1).

⁵⁵Indeed, if – provided that all other bidders use the equilibrium bidding function β – bidder i chooses a bid b_i that is strictly less than his value θ_i , then it may be possible that another bidder j bids $\beta(\theta_j) \in (b_i, \theta_i)$, so that bidder i loses the object to a bidder with a lower valuation. If, on the other hand, bidder i places a bid $b_i > \theta_i$, then there may be a bidder k such that $b_k \in (\theta_i, b_i)$ forcing bidder i to pay more than his private value θ_i for the object. Hence, it is optimal for bidder i to set $b_i = \beta(\theta_i) = \theta_i$.

⁵⁶The *design* aspect of mechanisms for resource allocation was highlighted by Hurwicz (1973).




Figure 6: Allocations under the optimal selling mechanism.

such that marginal revenue equals marginal cost, i.e.,⁵⁷

$$\text{MR}_i(r_i) = r_i + \frac{D_i(r_i)}{D'_i(r_i)} = r_i - \frac{1 - F(r_i)}{F'_i(r_i)} = c_i. \quad (39)$$

Let us now consider a mechanism-design approach for the optimal pricing problem. The revelation principle (Theorem 4) extends to the context of “Bayesian implementation,” i.e., situations with multiple agents, in the sense that it is possible for the principal to restrict attention to ‘direct revelation mechanisms’ of the form $\mathcal{M} = (\Theta, \rho)$, where $\Theta = \Theta_1 \times \Theta_2$ is the type space, and $\rho = (\xi_1, \xi_2; \tau_1, \tau_2)$ the allocation function, with $\xi_i : \Theta \rightarrow [0, 1]$ the probability that agent i obtains the item and $\tau_i : \Theta \rightarrow \mathbb{R}$ the transfer that this agent pays to the principal.

Given that agent j reports the true value θ_j , agent i finds it optimal to also tell the truth if

$$\theta_i \in \arg \max_{\hat{\theta}_i \in [\underline{\theta}_i, \bar{\theta}_i]} E \left[\theta_i \xi_i(\hat{\theta}_i, \tilde{\theta}_j) - \tau_i(\hat{\theta}_i, \tilde{\theta}_j) \mid \hat{\theta}_i, \theta_i \right] = \arg \max_{\hat{\theta}_i \in [\underline{\theta}_i, \bar{\theta}_i]} \left\{ \theta_i \bar{\xi}_i(\hat{\theta}_i) - \bar{\tau}_i(\hat{\theta}_i) \right\},$$

for all $(\theta_1, \theta_2) \in \Theta$, where $\bar{\xi}_i(\hat{\theta}_i) = E[\xi_i(\hat{\theta}_i, \tilde{\theta}_j) \mid \hat{\theta}_i]$ is agent i ’s expected probability of obtaining the item when sending the message $\hat{\theta}_i$ and $\bar{\tau}_i(\hat{\theta}_i) = E[\tau_i(\hat{\theta}_i, \tilde{\theta}_j) \mid \hat{\theta}_i]$ is the resulting expected payment by that agent. Completely analogous to the implementability conditions (31) and (32) in the screening problem, we obtain

$$\theta_i \bar{\xi}'_i(\theta_i) = \bar{\tau}'_i(\theta_i), \quad \theta_i \in [\underline{\theta}_i, \bar{\theta}_i], \quad i \in \{1, \dots, N\}, \quad (40)$$

and

$$\bar{\xi}'_i(\theta_i) \geq 0, \quad \theta_i \in [\underline{\theta}_i, \bar{\theta}_i], \quad i \in \{1, \dots, N\}, \quad (41)$$

⁵⁷As in Section 6.1 (cf. Footnote 42) we assume that we are in the “regular case,” i.e., that the marginal revenue is increasing in price. If this assumption is not satisfied, an ironing procedure analogous to the one mentioned in Remark 13 needs to be performed (Bulow and Roberts 1989).

as implementability conditions for the principal's direct revelation mechanism.⁵⁸ The seller maximizes his profits by solving

$$\max_{\xi, \tau} \sum_{i=1}^2 \int_{\theta_i}^{\bar{\theta}_i} (\bar{\tau}_i(\theta_i) - c_i \bar{\xi}_i(\theta_i)) dF_i(\theta_i) = \max_{\xi} \sum_{i=1}^2 \int_{\theta_1}^{\bar{\theta}_1} \int_{\theta_2}^{\bar{\theta}_2} (\text{MR}_i(\theta_i) - c_i) \xi_i(\theta_1, \theta_2) dF_1(\theta_1) dF_2(\theta_2),$$

subject to $0 \leq \xi_1(\theta_1, \theta_2) + \xi_2(\theta_1, \theta_2) \leq 1$.⁵⁹ The optimal solution can be found by simply 'looking hard' at the last expression. For a given type realization (θ_i, θ_j) , the principal finds it best to place all probability mass on the agent with the largest difference between marginal revenue and marginal cost, as long as that difference is nonnegative. If it is, then it is better not to sell the item at all. Hence,

$$\xi_i^*(\theta_i, \theta_j) = \begin{cases} \mathbf{1}_{\{\text{MR}_i(\theta_i) - c_i > \text{MR}_j(\theta_j) - c_j\}} + (\mathbf{1}_{\{\text{MR}_i(\theta_i) - c_i = \text{MR}_j(\theta_j) - c_j\}}/2), & \text{if } \text{MR}_i(\theta_i) \geq c_i, \\ 0, & \text{otherwise.} \end{cases}$$

To obtain a concrete result, let us consider the case where Joe's value is uniformly distributed on $\Theta_1 = [0, 100]$ and Melanie's value is uniformly distributed on $\Theta_2 = [0, 200]$, so that $F_i(\theta_i) = \theta_i/(100i)$. Furthermore, let $c_1 = c_2 = 20$. One can see that condition (41) is satisfied and that, in addition, condition (40) holds if we set⁶⁰

$$\begin{aligned} \tau_i^*(\theta_i, \theta_j) &= \mathbf{1}_{\{\text{MR}_i(\theta_i) \geq c_i\}} \int_{\theta_i}^{\theta_i} s \frac{\partial \xi_i^*(s, \theta_j)}{\partial s} ds = \mathbf{1}_{\{\text{MR}_i(\theta_i) \geq c_i\}} \int_{\theta_i}^{\theta_i} s \delta(s - \theta_j - 50(i - j)) ds \\ &= \begin{cases} \theta_j + 50(i - j), & \text{if } \theta_i \geq \max\{(c/2) + 50i, \theta_j + 50(i - j)\}, \\ 0, & \text{otherwise,} \end{cases} \end{aligned}$$

where $\delta(\cdot)$ is the Dirac impulse distribution, which appears naturally as the (generalized) derivative of the indicator function (cf. Footnote 15). In other words, the optimal selling mechanism for the principal is an auction where the winner is the agent with the highest marginal profit $(\text{MR}_i - c_i)$, and the transfer is determined by the valuation of the losing bidder and the difference in their bids. In this asymmetric second-price auction, the optimal reserve price r_i for agent i is chosen such that $\text{MR}_i(r_i) = c_i$, i.e., $r_i = (c_i/2) + 50i$, as already computed for the third-degree price discrimination in Eq. (39) above: $r_1 = (c_1/2) + 50 = 60$ and $r_2 = (c_2/2) + 100 = 110$. Using the optimal auction mechanism, the firm obtains an expected profit of $E[\max\{0, \text{MR}_1(\theta_1) - c_1, \text{MR}_2(\theta_2) - c_2\}] = 2E[\max\{0, \theta_1 - 60, \theta_2 - 110\}] = 50.3\bar{6}$. The reason that the optimal auction seems to lead to a lower expected revenue than third-degree price

⁵⁸Differentiating Eq. (40) with respect to θ_i we obtain $\xi_i'(\theta_i) + \theta_i \bar{\xi}_i''(\theta_i) = \bar{\tau}_i''(\theta_i)$. On the other hand, the second-order optimality condition is $\theta_i \bar{\xi}_i''(\theta_i) - \bar{\tau}_i''(\theta_i) \leq 0$, so that (by combining the last two relations) we obtain inequality (41).

⁵⁹The last identity is obtained, using Eq. (40) and performing twice an integration by parts, as follows:

$$\begin{aligned} \int_{\theta_i}^{\bar{\theta}_i} \bar{\tau}_i(\theta_i) dF_i(\theta_i) &= \int_{\theta_i}^{\bar{\theta}_i} \left(\int_{\theta_i}^{\theta_i} s \bar{\xi}_i'(s) ds \right) dF_i(\theta_i) = \left(\int_{\theta_i}^{\theta_i} s \bar{\xi}_i'(s) ds \right) F_i(\theta_i) \Big|_{\theta_i}^{\bar{\theta}_i} - \int_{\theta_i}^{\bar{\theta}_i} \theta_i \bar{\xi}_i'(\theta_i) F_i(\theta_i) d\theta_i \\ &= \int_{\theta_i}^{\bar{\theta}_i} (1 - F_i(\theta_i)) \theta_i \bar{\xi}_i'(\theta_i) d\theta_i = (1 - F_i(\theta_i)) \theta_i \bar{\xi}_i(\theta_i) \Big|_{\theta_i}^{\bar{\theta}_i} - \int_{\theta_i}^{\bar{\theta}_i} (1 - F_i(\theta_i) - \theta_i F_i'(\theta_i)) \bar{\xi}_i(\theta_i) d\theta_i \\ &= -\theta_i \bar{\xi}_i(\theta_i) + \int_{\theta_i}^{\bar{\theta}_i} \left(\theta_i - \frac{1 - F_i(\theta_i)}{F_i'(\theta_i)} \right) \bar{\xi}_i(\theta_i) F_i'(\theta_i) d\theta_i = \int_{\theta_i}^{\bar{\theta}_i} \text{MR}_i(\theta_i) \bar{\xi}_i(\theta_i) dF_i(\theta_i), \end{aligned}$$

where the firm naturally sets $\bar{\xi}_i(\theta_i) = \bar{\tau}_i(\theta_i) = 0$. The derivations are analogous to the ones in Footnote 51.

⁶⁰This is not the only way to satisfy Eq. (40). For more details on Bayesian implementation, see Palfrey and Srivastava (1993) as well as Weber and Bapna (2008).

discrimination is that the latter allows for items to be simultaneously sold to high-value types of both buyers (“markets”), whereas the former allocates one unit to either one of the two buyers (“agents”). The allocations resulting from the (asymmetric) optimal auction, for any $(\theta_1, \theta_2) \in \Theta$, are shown in Figure 6.

REMARK 15 The interpretation of the design of an optimal auction in terms of a third-degree monopoly pricing problem stems from Bulow and Roberts (1989). While Vickrey’s second-price auction leads to a Pareto-efficient allocation *ex post* (i.e., the item ends up with the individual that has the highest value for it), Myerson’s optimal auction may lead to an inefficient allocation, since it is possible that no transaction takes place even though an agent values the item above the seller’s marginal cost.⁶¹ This illustrates a fundamental conflict between efficient allocation and surplus extraction when information is asymmetric. Note that the expected-revenue difference between the Myerson and Vickrey auctions is small in the following sense: the expected revenues from a Vickrey auction with $N + 1$ ex-ante symmetric bidders cannot be less than the expected revenue for an optimal auction with N ex-ante symmetric bidders (Bulow and Klemperer 1996). For additional details on auction theory see, e.g., Milgrom (2004). In many practical applications the principal also has private information, which complicates the mechanism design problem (Myerson 1983; Maskin and Tirole 1992). \square

7.4 Oligopoly Pricing with Asymmetric Information

The most general strategic pricing problems arise when multiple sellers interact with multiple buyers in the presence of asymmetric information. An interesting case is the situation where the sellers observe different private signals about their demands and/or their own costs which they can choose to exchange with other sellers or not. Raith (1996) provides a model of information sharing when the sellers’ payoffs are quadratic in their action variables, which encompasses the standard specifications of Cournot and Bertrand oligopolies with differentiated goods. A seller’s incentive to reveal private information depends on the correlation of the signals (i.e., if seller i observes a high cost or a high demand, how likely is it that seller j has a similar observation?) and on the type of strategic interaction (i.e., does an increase in seller i ’s action tend to increase (‘strategic complements’) or decrease (‘strategic substitutes’) seller j ’s action?). For example, in a Cournot duopoly with substitute products, increasing firm i ’s output tends to decrease firm j ’s output, so that their actions are strategic substitutes. The firms’ incentives to share information then increase as the correlation in their signals decreases. With very correlated signals, firm i does not expect to increase its payoffs by forcing its strategy to be even more negatively correlated with firm j ’s after information revelation than before. As signals are less correlated and the firms’ products gain in complementarity, their strategic actions become more positively correlated *ex ante*, so that information sharing can profitably increase the precision of how the firms can take advantage of this implicit collusion against the consumers.

It is noteworthy that in an oligopoly setting it may be better for a seller to actually have less precise information if that fact can be made reliably known to others (Gal-Or 1986). Indeed, if in a Cournot oligopoly with substitute products a seller i does not know that uncertain market demand (common to all sellers) is likely to be low, then that seller will produce a larger output

⁶¹The multi-unit generalization of the second-price auction is the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey 1961; Clarke 1971; Groves 1973). Ausubel and Milgrom (2006) describe why the VCG mechanism, despite its many attractive features (such as truthful bidding strategies which are incentive-compatible in dominant-strategies, and ex-post efficient allocations) it is rarely used in practice.

than sellers who are aware of the demand forecast, especially if those sellers know that seller i is uninformed. Thus, not having demand information helps a seller to credibly act “crazy” and thus to commit to an otherwise unreasonable action (an informed seller would never rationally opt to produce such a large output). Thus, while in single-person decision problems it is always better to have more information (with a partial information order induced by statistical sufficiency; Blackwell 1951), in strategic settings the value of additional information may well be negative.⁶²

Another way for multiple buyers and sellers to interact is through intermediaries, who may be able to alleviate some of the problems that arise from asymmetric information. One can argue that most economic transactions are ‘intermediated,’ in the sense that a buyer who resells his item effectively becomes an intermediary between the party that sold the item and the party that finally obtains the item. Thus, retailers are intermediaries, wholesalers are intermediaries, banks are intermediaries, and so on. A necessary condition for a buyer and a seller to prefer an intermediated exchange to a direct exchange is that the fee charged by the intermediary does not exceed the cost of direct exchange for either of the two transacting parties. Thus, for the intermediary to have a viable business proposition it is necessary that the intermediation cost for a given transaction is less than the expected transaction cost the parties face in a direct exchange. As an example consider the market for used cars discussed at the beginning of Section 4.2. We saw there that if a buyer’s expected value $1 - \varphi$ (where φ is the proportion of zero-value lemons in the market) is below the high-value seller’s opportunity cost c , i.e., if $1 - \varphi < c$, then the market fails, due to adverse selection. If an intermediary can observe and ‘certify’ the quality of a seller’s car at the intermediation cost κ , then the intermediary can charge the buyer a ‘retail price’ $R \in \{R_L, R_H\}$ (with $R_L = 0$ and $R_H = 1$) and pass on to the seller a ‘wholesale price’ $W \in \{W_L, W_H\}$ (with $W_L = 0$ and $W_H = c$), contingent on the observed value of the used car, such that all buyers and sellers are happy to use the intermediary. The intermediary is viable if and only if its expected revenues outweigh the intermediation cost, i.e., if and only if

$$P(\text{Used Car is of High Value})(R_H - W_H) = (1 - \varphi)(1 - c) \geq \kappa.$$

Spulber (1999) provides an extensive discussion of market intermediation and different ways in which intermediaries can create value by effectively lowering the transaction cost in the market. More recently, the theory of intermediation has evolved into a “theory of two-sided markets,” which emphasizes the intermediary’s profit objective and the possibility of competition between intermediaries or “platforms” (Rochet and Tirole 2003).⁶³

It is also possible to explicitly consider the competition between different economic mechanism designs as in Section 7.3 (Biais et al. 2000), with one of the main technical difficulties being that the revelation principle (Theorem 4) ceases to hold in that environment. For example, in the case of nonlinear pricing this leads to so-called “catalog games” (Monteiro and Page 2008).

REMARK 16 The source of incomplete information in a product market is in many cases directly related to a costly search process that limits the consumers’ ability to compare different products.

⁶²Weber and Croson (2004) show that this may be true even in a bilateral setting where an information seller sells information to an agent with limited liability. The fact that too much information might induce the agent to take overly risky actions, and thus perhaps be unable to pay the seller a pre-agreed ex-post fee, is enough motivation for the seller to want to degrade the information before selling it, knowing well that its price decreases but that the increased likelihood of obtaining payment outweighs this loss.

⁶³Rochet and Tirole (2005) point out that a “market is two-sided if the platform can affect the volume of transactions by charging more to one side of the market and reducing the price paid by the other side by an equal amount; in other words, the price structure matters, and platforms must design it so as to bring both sides on board.”

Weitzman (1979) points out that in an optimal sequential search a consumer should follow ‘Pandora’s rule,’ trading off the cost of an incremental search effort against the expected benefit from that extra search over the best alternative already available. A particularly striking consequence of even the smallest search cost in a product market is the following paradox by Diamond (1971): if N symmetric firms are Bertrand-competing on price in a market for a homogenous good and consumers can visit one (random) firm for free and from then on have to incur a small search cost $\varepsilon > 0$ to inspect the price charged by another firm, then instead of the Bertrand outcome where all firms charge marginal cost that would obtain for $\varepsilon = 0$ (cf. Section 6.2), one finds that all firms charge monopoly price. This rather surprising result will not obtain if consumers are heterogeneous in terms of their search costs (tourists vs. natives) or their preferences (bankers vs. cheap skates). Stiglitz (1989) provides a summary of the consequences of imperfect information in the product market related to search. See also Vives (1999) for a survey of oligopoly pricing under incomplete information.

8 Dynamic Pricing

For a seller with market power, the ability to adjust his prices over time is both a blessing and a curse.⁶⁴ On the one hand, the flexibility of being able to adjust prices allows the seller at each time t to incorporate all available information into his pricing decision. On the other hand, the fact that a seller has the flexibility to change prices at any time in the future removes his ability to commit to any particular price today. The following conjecture by Coase (1972) makes this last point more precise.⁶⁵

THEOREM 5 (COASE CONJECTURE) *Consider a monopolist who sells a durable (i.e., infinitely lived) good to consumers with possibly heterogeneous valuations at times $t \in \{0, \Delta, 2\Delta, \dots\}$ at the marginal cost $c \geq 0$. As $\Delta \rightarrow 0+$ the monopolist’s optimal price $p_0(\Delta)$ at time $t = 0$ tends to his marginal cost, i.e., $\lim_{\Delta \rightarrow 0+} p_0(\Delta) = c$.*

When the length Δ of the time period for which a monopolist can guarantee a constant price decreases, the price it can charge drops. The intuition is that the monopolist at time $t = 0$ is competing with a copy of its own product that is sold at time $t = \Delta$. Clearly, when viewed from the present, that copy is not quite as good as the product now. However, that ‘quality’ difference vanishes when Δ tends to zero. Thus, as $\Delta \rightarrow 0+$ arbitrarily many copies of virtually the same product will be available in any fixed time interval, so that the resulting perfect competition must drive the monopolist’s price down to marginal cost.

The Coase problem, which arises from the lack of commitment power, can be ameliorated by renting the product instead of selling it, or by making binding promises about future production (e.g., by issuing a “limited edition”). Perishable products also increase the monopolist’s commitment power (e.g., when selling fresh milk), as well as adjustment costs for price changes (e.g., due to the necessity of printing a new product catalogue). The ability to commit to a ‘price path’ from the present into the future is a valuable asset for a seller; the question of commitment as a result of the consumers’ option of intertemporal arbitrage (i.e., they can choose between buying

⁶⁴For more details on dynamic pricing, see Aviv and Vulcano (this volume).

⁶⁵The Coase conjecture was proved by Stokey (1981), Bulow (1982), and Gul et al. (1986) in varying degrees of generality.

now or later) must be at the center of any dynamic pricing strategy, at least in an environment where information is fairly complete. We note that the Coase problem is not significant in situations when consumers are nonstrategic (i.e., not willing or able to wait) or when goods respond to a current demand (e.g., for electricity). An intertemporal aspect is then introduced into a firm's pricing strategy mainly through its own cost structure, production technology, or changes in boundary conditions.⁶⁶

Under incomplete information, buyers and sellers need to learn about their economic environment, which generally includes exogenous factors (such as technological possibilities, macroeconomic conditions, and consumer preferences) or persistent endogenous information (such as past experience with products, reputations, and recent observable actions by market participants). All participants' beliefs may change from one time period to the next, based on their respective new information. At each time t , any agent trades off the value of additional costly information (important for improving decisions in the next period) against the benefit of using the available resources to directly augment current payoffs through non-informational actions. The difficulties of finding an optimal balance between exploration (i.e., the acquisition of new information) and exploitation (i.e., the accumulation of benefit) can be illustrated well with the so-called multi-armed bandit problem. In the simplest version of this problem a single agent can choose to pull one of several arms of a slot machine. Each arm pays a random reward with an unknown, stationary probability distribution. The agent weighs the current benefit of pulling a relatively well-known arm to gain a fairly certain reward against the possibility of experimenting with a new arm that could pay even better rewards. Gittins and Jones (1974) provide a surprisingly simple solution to what seems like a difficult dynamic optimization problem which for some time had a reputation of being unsolvable from an analytic point of view. Their solution shows that the problem can be decomposed so as to perform a separate computation of a *Gittins index* for each arm based on the available information for that arm. The Gittins index corresponds to the retirement reward (as a per-period payment in perpetuity) that would make the agent indifferent between stopping the experimentation with that arm or obtaining the retirement perpetuity instead. Given the vector of Gittins indices at time t , it is optimal for the agent to choose the arm with the highest Gittins index.

The fact that the Gittins indices can be computed separately for each arm allows for a use of this method in mechanism design, for example, in the construction of a mechanism that allocates an object of unknown value in different rounds to different bidders (Bapna and Weber 2006; Bergemann and Valimäki 2006). Based on their own observations, bidders can then submit bids which are related to the computation of their own Gittins index, the latter representing in effect the private types, the collection of which determines the socially optimal allocation at each point in time.

9 Behavioral Anomalies

A discussion of price theory could not be complete without the 'disclaimer' that the self-interested *economic man* (or *homo economicus*) that appeared in all of the rational models in this survey

⁶⁶For example, a firm may need to consider a production lead time, the replenishment of inventories, a limited-capacity infrastructure (requiring 'peak-load pricing'), or demand seasonality. Many of these issues have been considered in the operations management and public pricing literature; they are beyond the scope of this survey.

is not exactly what is encountered in reality.⁶⁷ Individuals are generally subject to decision biases which may be psychological in nature or due to limited cognitive abilities (or both). The resulting “bounded rationality” leads human decision makers to suboptimal decisions and to what Simon (1956) terms “satisficing” (rather than optimizing) behavior, which means that individuals tend to stop the search for better decision alternatives as soon as a decision has been found that promises a payoff which satisfies a certain aspiration level.

The behavioral anomalies encountered in practice or often too significant to be ignored in models, so that descriptive and normative models of bounded rationality increasingly influence and shape the evolution of economic theory (Thaler 1994; Rubinstein 1998; Kahneman and Tversky 2000; Camerer 2006). As an example for how decision making can be biased by ‘mental accounting’ (Thaler 1999), consider the following experiment by Tversky and Kahneman (1981, p. 457), in which individuals were asked the following.

Imagine that you are about to purchase a jacket for \$125 (resp. \$15) and a calculator for \$15 (resp. \$125). The calculator salesman informs you that the calculator you wish to buy is on sale for \$10 (resp. \$120) at the other branch of the store, located 20 minutes drive away. Would you make the trip to the other store?

The empirical observation in this experiment suggests that when individuals stand to save \$5 from an original amount of \$15 they are willing to travel to another store, but not if the original amount is \$125. Thus, the same total amount of \$5 tends to be aggregated into one mental account, which contains the total expense. This explains the decreasing marginal utility for a \$5 savings in the base price of the item. The latter is explained by the empirical regularity that real individuals’ utility functions for money tend to be convex for losses (implying ‘loss aversion’) and concave for gains (implying ‘risk aversion’; cf. Section 2.4) (Kahneman and Tversky 1979). An individual with a utility function that, relative to the current wealth, exhibits both loss aversion and risk aversion, tends to avoid actions (i.e., lotteries) that change his *status quo*.

Related literature on the so-called *endowment effect* points to an important *caveat* when trying to identify a ‘behavioral anomaly’ which may require the adjustment of economic theory. In a well-known experiment, Knetsch (1989) observes systematic differences between WTA and WTP when the welfare change is induced by the transfer of a nonmarket good. These empirical results, repeated in a variety of settings by other authors, suggest that a consumer’s WTA tends to be larger than her WTP, a behavioral pattern that is generally referred to as the “endowment effect” or “status-quo bias.” In Knetsch’s experiment, one group of randomly selected subjects was given a coffee mug and another group of subjects a candy bar. No matter who was given the mug, the individuals’ average WTAs tended to be much larger than their average WTPs, which has been commonly viewed as supporting evidence for the presence of an endowment effect. Yet, as the normative relation between WTA and WTP (i.e., the Hicksian welfare measures EV and CV) in Eq. (23) shows, it is

$$\text{WTA}(w - \text{WTP}(w)) = \text{WTP}(w) \quad \text{and} \quad \text{WTA}(w) = \text{WTP}(w + \text{WTA}(w)),$$

so that an identity between WTA and WTP is not actually required, even from a normative point of view, unless one makes sure that individuals are wealth-compensated as the previous

⁶⁷The term “economic man” dates back to Ingram (1888) who criticizes the work of Mill (1836) which – in Ingram’s words – “dealt not with real but with imaginary men – ‘economic men’ (...) conceived as simply ‘money-making animals’” (p. 218); see also the discussion by Persky (1995). Behavioral issues in pricing are discussed further by Özer and Zheng (this volume).

identities indicate. More recently, Plott and Zeiler (2005) note (without alluding to the wealth-compensation issue) that the observed differences between WTA and WTP might disappear altogether if one controls for all possible biases in experimental settings, so that – as these authors point out – the endowment effect may be a mere artifact produced by observations under imperfect laboratory conditions. The controversy surrounding the endowment effect is symptomatic of the difficulties in separating implications of available normative models from behavioral anomalies that would require an extension of these models.⁶⁸

10 Open Issues

The discussion in this survey has shown that the strategic and nonstrategic pricing of resources is complicated by the presence of externalities, asymmetric information, and behavioral anomalies. The following three axes will be instrumental in the further development of price theory.

First, as the growth of networks around us increases the ‘connectedness’ of individuals, on the one hand new markets are created, but on the other hand the potential for direct externalities through actions taken by participants in these networks also increases. New economic theory is needed for the pricing and transfer of resources in networks which considers the informational, transactional, and behavioral aspects specific to the relevant networks.

A second fundamental driver for a change in the way price theory continues to develop is given by our transition to a ‘data-rich society,’ in the sense that firms (and individuals) are increasingly able to tap a variety of data sources for informing their decisions about the buying, selling, and pricing of resources. Consumers whose behavior can be tracked, firms whose products become customizable by informational components, and intermediaries that increase market transparency (up to a deliberate obfuscation point), imply reduced information rents. The countervailing trend is that a data-rich society allows for experimentation on an unprecedented scale. The resulting optimization of decisions increases the rents from information. It therefore becomes more critical to connect economic models directly to data, blurring the boundaries of model identification and solution.

Lastly, with the advent of a data-rich society there remains the issue of “long-tails,” i.e., the fundamental impossibility to have well-defined beliefs about all relevant aspects of economic behavior. Low-probability, high-consequence events require decisions to be “robust,” i.e., valid under a range of possible scenarios.⁶⁹ One important aspect of robustness is an implementation of satisficing (cf. Section 9), where a simple expected-payoff objective is replaced by a ‘robustness criterion,’ such as a conservative worst-case value or a less stringent competitive-ratio criterion.

⁶⁸Weber (2010) shows that the endowment effect is unsurprising from a normative viewpoint, i.e., it is not necessarily a behavioral anomaly. For example, a “normative endowment effect” arises automatically if the welfare measures WTA and WTP are increasing in wealth.

⁶⁹The second-price auction discussed in Section 7.3 is an example of a robust mechanism in the following sense. All bidders’ equilibrium strategies (cf. Footnote 55) are completely independent of their beliefs about the distribution of other bidders’ valuations. In general, this is not case. For example, if the auction were such that a winning bidder has to pay his bid (termed a first-price auction), then the equilibrium bid strategies become very sensitive to the beliefs about other players’ types and to the players’ common knowledge about these beliefs. This is to the extent that with asymmetric beliefs the first-price auction can yield inefficient allocations (although it is revenue-equivalent to the second-price auction in the symmetric case). The second-price auction, therefore, is in accordance with the so-called ‘Wilson doctrine’ (Wilson 1987), which calls for the design of distribution-free mechanisms. Auction pricing is discussed further by Steinberg (this volume).

Sample-sparse estimation is likely to become an important part of data-driven economic decision making under uncertainty (Chehrazi and Weber 2010).

References

- ADMATI, A.R., PFLEIDERER, P. (1986) "A Monopolistic Market for Information," *Journal of Economic Theory*, Vol. 39, No. 2, pp. 400–438.
- AKERLOF, G.A. (1970) "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, Vol. 84, No. 3, pp. 488–500.
- ANDERSON, S., DE PALMA, A., THISSE, J.-F. (1992) *Discrete-Choice Theory of Product Differentiation*, MIT Press, Cambridge, MA.
- ARROW, K.J. (1951) *Social Choice and Individual Values*, Wiley, New York, NY. [Second edition published in 1963 by Cowles Foundation, Monograph 12, Yale University Press, New Haven, CT.]
- ARROW, K.J. (1953) "Le Rôle des Valeurs Boursières pour la Répartition la Meilleure des Risques," *Économétrie, Colloques Internationaux du CNRS*, XL, pp. 41–48. [English translation, "The Role of Securities in the Optimal Allocation of Risk," published in *Review of Economic Studies*, Vol. 31, No. 2, pp. 91–96. The paper was originally read at the *Colloque sur les Fondements et Applications de la Théorie du Risque* on May 13, 1952, in Paris, France.]
- ARROW, K.J. (1963) "Uncertainty and the Welfare Economics of Medical Care," *American Economic Review*, Vol. 53, No. 5, pp. 941–973.
- ARROW, K.J., HAHN, F.J. (1971) *General Competitive Analysis*, Holden-Day, San Francisco, CA.
- ARROW, K.J. (1973) "Information and Economic Behavior," Lecture Presented to the Federation of Swedish Industries, Stockholm, Sweden. Reprinted in: *Collected Papers of Kenneth J. Arrow, Vol. 4*, Belknap Harvard, Cambridge, MA, in 1984, pp. 136–152.
- ATHEY, S. (2002) "Monotone Comparative Statics Under Uncertainty," *Quarterly Journal of Economics*, Vol. 117, No. 1, pp. 187–223.
- AUMANN, R.J. (1976) "Agreeing to Disagree," *Annals of Statistics*, Vol. 4, No. 6, pp. 1236–1239.
- AUSUBEL, L., MILGROM, P.R. (2006) "The Lovely But Lonely Vickrey Auction," in: CRAMTON, P., SHOHAM, Y., STEINBERG, R. (Eds.) *Combinatorial Auctions*, MIT Press, Cambridge, MA, pp. 17–40.
- BAGWELL, K., RAMEY, G. (1991) "Oligopoly Limit Pricing," *RAND Journal of Economics*, Vol. 22, No. 2, pp. 155–172.
- BAGWELL, K., RIORDAN, M.H. (1991) "High and Declining Prices Signal Product Quality," *American Economic Review*, Vol. 81, No. 1, pp. 224–239.
- BAGWELL, K. (2007) "The Economic Analysis of Advertising," in: Armstrong, M., Porter, R. (Eds.) *Handbook of Industrial Organization, Vol. 3*, Elsevier, Amsterdam, NL, pp. 1701–1844.
- BAKER, J.B., BRESNAHAN, T.F. (1988) "Estimating the Residual Demand Curve Facing a Single Firm," *International Journal of Industrial Organization*, Vol. 6, No. 3, pp. 283–300.
- BAKER, J.B., BRESNAHAN, T.F. (1992) "Empirical Methods of Identifying and Measuring Market Power," *Antitrust Law Journal*, Vol. 61, pp. 3–16.
- BAPNA, A., WEBER, T.A. (2006) "Efficient Dynamic Allocation with Uncertain Valuations," Technical Report 2006-03-31, Department of Management Science and Engineering, Stanford University, Stanford, CA.
- BARON, D.P., MYERSON, R.B. (1982) "Regulating a Monopolist with Unknown Costs," *Econometrica*, Vol. 50, No. 4, pp. 911–930.
- BAUMOL, W.J., BRADFORD, D.F. (1970) "Optimal Departures from Marginal Cost Pricing," *American Economic Review*, Vol. 60, No. 3, pp. 265–283.
- BEATH, J., KATSOUACOS, Y. (1991) *The Economic Theory of Product Differentiation*, Cambridge University Press, New York, NY.
- BECKMAN, S.R., SMITH, W.J. (1993) "Positively Sloping Marginal Revenue, CES Utility and Subsistence Requirements," *Southern Economic Journal*, Vol. 60, No. 2, pp. 297–303.
- BERGEMANN, D., VÄLIMÄKI, J. (2006) "Efficient Dynamic Auctions," Cowles Foundation Discussion Paper No. 1584, Yale University, New Haven, CT.

- BERNHEIM, B.D., WHINSTON, M.D. (1986) "Menu Auctions, Resource Allocations, and Economic Influence," *Quarterly Journal of Economics*, Vol. 101, No. 1, pp. 1–31.
- BERTRAND, J.L.F. (1883) "Théorie Mathématique de la Richesse Sociale, Recherches sur les Principes Mathématiques de la Théorie des Richesses," *Journal des Savants*, Vol. 67, pp. 499–508. [Review of the books by Walras (1874/77) and Cournot (1838).]
- BIAIS, B., MARTIMORT, D., ROCHET, J.-C. (2000) "Competing Mechanisms in a Common Value Environment," *Econometrica*, Vol. 68, No. 4, pp. 799–838.
- BISHOP, R.L. (1966) "Monopoly Under General Equilibrium: Comment," *Quarterly Journal of Economics*, Vol. 80, No. 4, pp. 652–659.
- BLACKWELL, D. (1951) "Comparison of Experiments," Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, pp. 93–102.
- BOWLEY, A.L. (1924) *The Mathematical Groundwork of Economics*, Oxford University Press, Oxford, UK.
- BRONFMAN, C., MCCABE, K., PORTER, D., RASSENTI, S., SMITH, V. (1996) "An Experimental Examination of the Walrasian Tâtonnement Mechanism," *RAND Journal of Economics*, Vol. 27, No. 4, pp. 681–699.
- BULOW, J. (1982) "Durable Goods Monopolists," *Journal of Political Economy*, Vol. 90, No. 2, pp. 314–332.
- BULOW, J., KLEMPERER, P. (1996) "Auctions vs. Negotiations," *American Economic Review*, Vol. 86, No. 1, pp. 180–194.
- BULOW, J., ROBERTS, J. (1989) "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, Vol. 97, No. 5, pp. 1060–1090.
- CABRAL, L.M.B., VILLAS-BOAS, M. (2005) "Bertrand Supertraps," *Management Science*, Vol. 51, No. 4, pp. 599–613.
- CAMERER, C.F. (2006) "Behavioral Economics," in: Blundel, R., Newey, W.K., Persson, T. (Eds.) *Advances in Economics and Econometrics, Vol. II*, Cambridge University Press, Cambridge, UK, pp. 181–214.
- CASS, D., SHELL, K. (1983) "Do Sunspots Matter?" *Journal of Political Economy*, Vol. 91(2), pp. 193–227.
- CHAMP, P.A., BOYLE, K.J., BROWN, T.J. (2003) *A Primer on Nonmarket Valuation*, Kluwer, Dordrecht, NL.
- CHEHRAZI, N., WEBER, T.A. (2010) "Monotone Approximation of Decision Problems," *Operations Research*, Vol. 58, No. 4/2, pp. 1158–1177.
- CHIPMAN, J. (1977) "An Empirical Implication of Auspitz-Lieben-Edgeworth-Pareto Complementarity," *Journal of Economic Theory*, Vol. 14, No. 1, pp. 228–231.
- CHO, I.-K., KREPS, D.M. (1987) "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, Vol. 102, No. 2, pp. 179–222.
- CLARKE, E.H. (1971) "Multipart Pricing of Public Goods," *Public Choice*, Vol. 2, No. 1, pp. 19–33.
- COASE, R.H. (1937) "The Nature of the Firm," *Economica*, Vol. 4, No. 16, pp. 386–405.
- COASE, R.H. (1960) "The Problem of Social Cost," *Journal of Law and Economics*, Vol. 3, No. 2, pp. 1–44.
- COASE, R.H. (1972) "Durability and Monopoly," *Journal of Law and Economics*, Vol. 15, No. 1, pp. 143–49.
- COBB, C.W., DOUGLAS, P.H. (1928) "A Theory of Production," *American Economic Review*, Vol. 18, No. 1, pp. 139–165.
- CONDORCET, M.J.A.N. DE CARITAT, MARQUIS DE (1785) "Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix (Essay on the Application of Analysis to the Probability of Majority Decisions)," Imprimerie Royale, Paris, France. [Reprinted by American Mathematical Society, Providence, RI, in 1972.]
- COURNOT, A. (1838) *Recherches sur les Principes Mathématiques de la Théorie des Richesses*, Hachette, Paris, France.
- CRAWFORD, V.P., SOBEL J. (1982) "Strategic Information Transmission," *Econometrica*, Vol. 50, No. 6, pp. 1431–1451.
- CREMER, H., THISSE, J.-F. (1991) "Location Models of Horizontal Differentiation: a Special Case of Vertical Differentiation Models," *Journal of Industrial Economics*, Vol. 39, No. 4, pp. 383–390.
- DEBREU, G. (1953) "Une Économie de l'Incertain," Working Paper, Électricité de France, Paris, France. [Published 1960, *Économie Appliquée*, Vol. 13, No. 1, pp. 111–116.]
- DEBREU, G. (1959) *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*, Cowles Foundation, Monograph 17, Yale University Press, New Haven, CT.
- DENECKERE, R.J., MCAFEE, R.P. (1996) "Damaged Goods," *Journal of Economics and Management Strategy*, Vol. 5, No. 2, pp. 149–174.

- DIAMOND, P.A. (1971) "A Model of Price Adjustment," *Journal of Economic Theory*, Vol. 3, No. 2, pp. 156–168.
- DUFFIE, D. (1988) *Security Markets*, Academic Press, New York, NY.
- DUFFIE, D. (2001) *Dynamic Asset Pricing Theory*, Third Edition, Princeton University Press, Princeton, NJ.
- EAVES, J., WILLIAMS, J. (2007) "Walrasian Tâtonnement Auctions on the Tokyo Grain Exchange," *Review of Financial Studies*, Vol. 20, No. 4, pp. 1183–1218.
- ECONOMIDES, N. (1996) "Network Externalities, Complementarities, and Invitations to Enter," *European Journal of Political Economy*, Vol. 12, No. 2, pp. 211–233.
- EDGEWORTH, F.Y. (1881) *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*, C. Kegan Paul, London, UK.
- FARRELL, J., RABIN, M. (1996) "Cheap Talk," *Journal of Economic Perspectives*, Vol. 10, No. 3, pp. 103–118.
- FISHBURN, P.C. (1970) *Utility Theory for Decision Making*, Wiley, New York, NY.
- FOLEY, D.K. (1967) "Resource Allocation and the Public Sector," *Yale Economic Essays*, Vol. 7, pp. 45–98.
- FOLEY, D.K. (1970) "Lindahl's Solution and the Core of an Economy with Public Goods," *Econometrica*, Vol. 38, No. 1, pp. 66–72.
- FRANK, R.H. (1985) "The Demand for Unobservable and Other Nonpositional Goods," *American Economic Review*, Vol. 75, No. 1, pp. 101–116.
- FRIEDMAN, M. (1962/2007) *Price Theory*, Aldine Transaction, London, UK.
- GAL-OR, E. (1988) "The Advantages of Imprecise Information," *RAND Journal of Economics*, Vol. 19, No. 2, pp. 266–275.
- GARBER, P.M. (1990) "Famous First Bubbles," *Journal of Economic Perspectives*, Vol. 4, No. 2, pp. 35–54.
- GELMAN, J.R., SALOP, S.C. (1983) "Judo Economics: Capacity Limitation and Coupon Competition," *Bell Journal of Economics*, Vol. 14, No. 2, pp. 315–332.
- GIBBARD, A. (1973) "Manipulation of Voting Schemes," *Econometrica*, Vol. 41, No. 4, pp. 587–601.
- GITTINS, J.C., JONES, D.M. (1974) "A Dynamic Allocation Index for the Sequential Design of Experiments," in: Gani, J., Sarkadi, K., Vincze, I. (Eds.) *Progress in Statistics*, North-Holland, Amsterdam, NL, pp. 241–266.
- GOEL, A., MEYERSON, A., WEBER, T.A. (2009) "Fair Welfare Maximization," *Economic Theory*, Vol. 41, No. 3, pp. 465–494.
- GOLDREICH, O., MICALI, S., WIGDERSON, A. (1991) "Proofs that Yield Nothing But Their Validity or All Languages in NP Have Zero-Knowledge Proof Systems," *Journal of the Association for Computing Machinery*, Vol. 38, No. 1, pp. 691–729.
- GOLLIER, C. (2001) *The Economics of Risk and Time*, MIT Press, Cambridge, MA.
- GROSSMAN, S.J. (1989) *The Informational Role of Prices*, MIT Press, Cambridge, MA.
- GROSSMAN, S.J., STIGLITZ, J.E. (1980) "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, Vol. 70, No. 3, pp. 393–408.
- GROVES, T. (1973) "Incentives in Teams," *Econometrica*, Vol. 41, No. 4, pp. 617–631.
- GUESNERIE, R., LAFFONT, J.-J. (1984) "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm," *Journal of Public Economics*, Vol. 25, No. 3, pp. 329–369.
- GUL, F., SONNENSCHN, H., WILSON, R.B. (1986) "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, Vol. 39, No. 1, pp. 155–190.
- HANEMANN, W.M. (1991) "Willingness to Pay and Willingness to Accept: How Much Can They Differ?" *American Economic Review*, Vol. 81, No. 3, pp. 635–647.
- HARDIN, G. (1968) "The Tragedy of the Commons," *Science*, Vol. 162, No. 3859, pp. 1243–1248.
- HART, O.D. (1988) "Incomplete Contracts and the Theory of the Firm," *Journal of Law, Economics and Organization*, Vol. 4, No. 1, pp. 119–139.
- HART, O.D., MOORE, J. (1990) "Property Rights and the Nature of the Firm," *Journal of Political Economy*, Vol. 98, No. 6, pp. 1119–1158.
- HAYEK, F.A. (1945) "The Use of Knowledge in Society," *American Economic Review*, Vol. 35, No. 4, pp. 519–530.

- HERFINDAHL, O.C. (1950) *Concentration in the U.S. Steel Industry*, Doctoral Dissertation, Department of Economics, Columbia University, New York, NY.
- HICKS, J.R. (1939) *Value and Capital*, Clarendon Press, Oxford, UK.
- HIRSHLEIFER, J. (1971) "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, Vol. 61, No. 4, pp. 561–574.
- HIRSHLEIFER, J., GLAZER, A., HIRSHLEIFER, D. (2005) *Price Theory and Applications*, Seventh Edition, Cambridge University Press, Cambridge, UK.
- HIRSCH, F. (1976) *Social Limits to Growth*, Harvard University Press, Cambridge, MA.
- HIRSCHMAN, A.O. (1964) "The Paternity of an Index," *American Economic Review*, Vol. 54, No. 5, pp. 761–762.
- HOTELLING, H. (1929) "Stability in Competition," *Economic Journal*, Vol. 39, No. 153, pp. 41–57.
- HURWICZ, L. (1960) "Optimality and Informational Efficiency in Resource Allocation Processes," in: Arrow, K.J., Karlin, S. (Eds.) *Mathematical Methods in the Social Sciences*, Stanford University Press, Stanford, CA, pp. 27–46.
- HURWICZ, L. (1973) "The Design of Mechanisms for Resource Allocation," *American Economic Review*, Vol. 63, No. 2, pp. 1–30.
- HURWICZ, L., REITER, S. (2006) *Designing Economic Mechanisms*, Cambridge University Press, Cambridge, UK.
- INGRAM, J.K. (1888) *A History of Political Economy*, Macmillan, New York, NY.
- JACKSON, M.O., WILKIE, S. (2005) "Endogenous Games and Mechanisms: Side Payments Among Players," *Review of Economic Studies*, Vol. 72, No. 2, pp. 543–566.
- JEVONS, W.S. (1881) *Political Economy*, Macmillan, London, UK.
- JOYCE, P. (1984) "The Walrasian Tâtonnement Mechanism and Information," *RAND Journal of Economics*, Vol. 15, No. 3, pp. 416–25.
- KAHNEMAN, D., TVERSKY, A. (1979) "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, Vol. 47, No. 2, pp. 263–291.
- KAHNEMAN, D., TVERSKY, A. (EDS.) (2000) *Choices, Values, and Frames*, Cambridge University Press, Cambridge, UK.
- KEYNES, J.M. (1936) *The General Theory of Employment, Interest and Money*, Macmillan, London, UK.
- KIHLSTROM, R.E., RIORDAN, M.H. (1984) "Advertising as a Signal," *Journal of Political Economy*, Vol. 92, No. 3, pp. 427–450.
- KNETSCH, J.L. (1989) "The Endowment Effect and Evidence of Nonreversible Indifference Curves," *American Economic Review*, Vol. 79, No. 5, pp. 1277–1284.
- KREPS, D. (1988) *Notes on the Theory of Choice*, Westview Press, Boulder, CO.
- KYLE, A.S. (1985) "Continuous Auctions and Insider Trading," *Econometrica*, Vol. 53, No. 6, pp. 1315–1335.
- LAFFONT, J.-J. (1989) *The Economics of Uncertainty and Information*, MIT Press, Cambridge, MA.
- LANCASTER, K.J. (1966) "A New Approach to Consumer Theory," *Journal of Political Economy*, Vol. 74, No. 2, pp. 132–157.
- LELAND, H.E., PYLE, D.H. (1977) "Informational Asymmetries, Financial Structure, and Financial Intermediation," *Journal of Finance*, Vol. 32, No. 2, pp. 371–387.
- LERNER, A.P. (1934) "The Concept of Monopoly and the Measurement of Monopoly Power," *Review of Economic Studies*, Vol. 1, No. 3, pp. 157–175.
- LINDAHL, E. (1919) *Die Gerechtigkeit der Besteuerung: eine Analyse der Steuerprinzipien auf der Grundlage der Grenznutzentheorie*, Hakan Ohlssons, Lund, Sweden. [Partially reprinted in 1958 as "Just Taxation: A Positive Solution," in: Musgrave, R.A., Peacock, A.T. (Eds.) *Classics in the Theory of Public Finance*, St. Martin's Press, New York, NY, pp. 168–176.]
- LINDHQVIST, T. (1992) "Extended Producer Responsibility," in: Lindhqvist, T. *Extended Producer Responsibility as a Strategy to Promote Cleaner Products*, Seminar Proceedings, Department of Industrial Environmental Economics, Lund University, Lund, Sweden.
- LUENBERGER, D.G. (2001) "Projection Pricing," *Journal of Optimization Theory and Applications*, Vol. 109, No. 1, pp. 1–25.

- LUENBERGER, D.G. (2002) "A Correlation Pricing Formula," *Journal of Economic Dynamics and Control*, Vol. 26, No. 7-8, pp. 1113–1126.
- MAGILL, M., QUINZII, M. (1996) *Theory of Incomplete Markets*, MIT Press, Cambridge, MA.
- MALKIEL, B.G. (2007) *A Random Walk Down Wall Street*, Ninth Edition, Norton, New York, NY.
- MARKOWITZ, H. (1952) "Portfolio Selection," *Journal of Finance*, Vol. 7, No. 1, pp. 77–91.
- MARSHALL, A. (1920) *Principles of Economics*, Macmillan, London, UK.
- MAS-COLELL, A. (1985) *The Theory of General Economic Equilibrium*, Econometric Society Monographs, Cambridge University Press, Cambridge, UK.
- MAS-COLELL, A., WHINSTON, M.D., GREEN, J.R. (1995) *Microeconomic Theory*, Oxford University Press, Oxford, UK.
- MASKIN, E., TIROLE, J. (1992) "The Principal-Agent Relationship with an Informed Principal, II: Common Values," *Econometrica*, Vol. 60, No. 1, pp. 1–42.
- MCKENZIE, L.W. (2002) *Classical General Equilibrium Theory*, MIT Press, Cambridge, MA.
- MENGER, C. (1871) *Grundsätze der Volkswirtschaftslehre: Erster, Allgemeiner Theil*, W. Braumüller, Vienna, Austria.
- MILGROM, P.R. (2004) *Putting Auction Theory to Work*, Cambridge University Press, Cambridge, UK.
- MILGROM, P., SEGAL, I. (2002) "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, Vol. 70, No. 2, pp. 583–601.
- MILGROM, P.R., SHANNON, C. (1994) "Monotone Comparative Statics," *Econometrica*, Vol. 62, No. 1, pp. 157–180.
- MILGROM, P.R., STOKEY, N. (1982) "Information, Trade, and Common Knowledge," *Journal of Economic Theory*, Vol. 26, No. 1, pp. 17–27.
- MILL, J.S. (1836) "On the Definition of Political Economy and on the Method of Investigation Proper to It," *London and Westminster Review*, October Issue. [Reprinted as "Essay V" in: Mill, J.S. (1844) *Essays on Some Unsettled Questions of Political Economy*, Parker, London, UK.]
- MILL, J.S. (1863) *Utilitarianism*, Routledge, London, UK.
- MIRRLEES, J.A. (1971) "An Exploration in the Theory of Optimal Income Taxation," *Review of Economic Studies*, Vol. 38, No. 2, pp. 175–208.
- MONTEIRO, P.K., PAGE, JR., F.H. (2008) "Catalog Competition and Nash Equilibrium in Nonlinear Pricing Games," *Economic Theory*, Vol. 34, No. 3, pp. 503–524.
- MUSSA, M., ROSEN, S. (1978) "Monopoly and Product Quality," *Journal of Economic Theory*, Vol. 18, No. 2, pp. 301–317.
- MUTH, J.F. (1961) "Rational Expectations and the Theory of Price Movements," *Econometrica*, Vol. 29, No. 3, pp. 315–335.
- MYERSON, R.B. (1979) "Incentive Compatibility and the Bargaining Problem," *Econometrica*, Vol. 47, No. 1, pp. 61–74.
- MYERSON, R.B. (1981) "Optimal Auction Design," *Mathematics of Operations Research*, Vol. 6, No. 1, pp. 58–73.
- MYERSON, R.B. (1983) "Mechanism Design by an Informed Principal," *Econometrica*, Vol. 51, No. 6, pp. 1767–1798.
- OSBORNE, M.J., PITCHIK, C. (1986) "Price Competition in a Capacity-Constrained Duopoly," *Journal of Economic Theory*, Vol. 38, No. 2, pp. 238–260.
- PALFREY, T.R., SRIVASTAVA, S. (1993) *Bayesian Implementation*, Harwood Academic Publishers, Chur, Switzerland.
- PERSKY, J. (1995) "The Ethology of Homo Economicus," *Journal of Economic Perspectives*, Vol. 9, No. 2, pp. 221–231.
- PIGOU, A.C. (1920) *The Economics of Welfare*, Macmillan, New York, NY.
- PLOTT, C.R., ZEILER, K. (2005) "The Willingness to Pay–Willingness to Accept Gap, the 'Endowment Effect,' Subject Misconceptions, and Experimental Procedures for Eliciting Valuations," *American Economic Review*, Vol. 95, No. 3, pp. 530–545.
- POSNER, R.A. (1975) "The Social Costs of Monopoly and Regulation," *Journal of Political Economy*, Vol. 83, No. 4, pp. 807–828.
- PRAT, A., RUSTICHINI, A. (2003) "Games Played Through Agents," *Econometrica*, Vol. 71, No. 4, pp. 989–1026.
- PRATT, J. (1964) "Risk Aversion in the Small and in the Large," *Econometrica*, Vol. 32, No. 1/2, pp. 122–136.

- QUAH, J. K.-H. (2007) "The Comparative Statics of Constrained Optimization Problems," *Econometrica*, Vol. 75, No. 2, pp. 401–431.
- RADNER, R. (1972) "Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets," *Econometrica*, Vol. 40, No. 2, pp. 289–303.
- RAWLS, J. (1971) *A Theory of Justice*, Harvard University Press, Cambridge, MA. [Revised edition published in 1999.]
- ROBERTS, D.J. (1973) "Existence of Lindahl Equilibrium with a Measure Space of Consumers," *Journal of Economic Theory*, Vol. 6, No. 4, pp. 355–381.
- ROCHET, J.-C., STOLE, L.A. (2003) "The Economics of Multidimensional Screening," in: Dewatripont, M., Hansen, L.-P., Turnovsky, S.J. (Eds.), *Advances in Economics and Econometrics: Theory and Applications: Eighth World Congress*, Vol. I, Cambridge University Press, New York, NY, pp. 150–197.
- ROLL, R., ROSS, S.A. (1980) "An Empirical Investigation of the Arbitrage Pricing Theory," *Journal of Finance*, Vol. 35, No. 5, pp. 1073–1103.
- ROSS, S.A. (1976) "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, Vol. 13, No. 3, pp. 341–360.
- RAITH, M. (1996) "A General Model of Information Sharing in Oligopoly," *Journal of Economic Theory*, Vol. 71, No. 1, pp. 260–288.
- RAMSEY, F.P. (1927) "A Contribution to the Theory of Taxation," *Economic Journal*, Vol. 37, No. 145, pp. 47–61.
- RILEY, J.G. (1979) "Informational Equilibrium," *Econometrica*, Vol. 47, No. 2, pp. 331–359.
- ROCHET, J.-C. (1985) "The Taxation Principle and Multitime Hamilton-Jacobi Equations," *Journal of Mathematical Economics*, Vol. 14, No. 2, pp. 113–128.
- ROCHET, J.-C., TIROLE, J. (2003) "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, Vol. 1, No. 4, pp. 990–1029.
- ROCHET, J.-C., TIROLE, J. (2005) "Two-Sided Markets: A Progress Report," Working Paper, Institut d'Économie Industrielle, Toulouse, France.
- RUBINSTEIN, A. (1998) *Modeling Bounded Rationality*, MIT Press, Cambridge, MA.
- SAMUELSON, P.A. (1941) "The Stability of Equilibrium: Comparative Statics and Dynamics," *Econometrica*, Vol. 9, No. 2, pp. 97–120.
- SAMUELSON, P.A. (1947) *Foundations of Economic Analysis*, Harvard University Press, Cambridge, MA.
- SAMUELSON, P.A. (1954) "The Pure Theory of Public Expenditure," *Review of Economics and Statistics*, Vol. 36, No. 4, pp. 387–389.
- SEADE, J. (1985) "Profitable Cost Increases and the Shifting of Taxation: Equilibrium Responses of Markets in Oligopoly," Warwick Economic Research Paper No. 260, Department of Economics, University of Warwick, Warwick, UK.
- SEGAL, I. (2006) "Communication in Economic Mechanism," in: Blundel, R., Newey, W.K., Persson, T. (Eds.) *Advances in Economics and Econometrics, Vol. I*, Cambridge University Press, Cambridge, UK, pp. 222–268.
- SHAFAREVICH, I.R. (1980) *The Socialist Phenomenon*, Harper and Row, New York, NY.
- SHAKED, A., SUTTON, J. (1982), "Relaxing Price Competition through Product Differentiation," *Review of Economic Studies*, Vol. 49, No. 1, pp. 3–13.
- SHARPE, W.F. (1964) "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance*, Vol. 19, No. 3, pp. 425–442.
- SHILLER, R.J. (2005) *Irrational Exuberance*, Second Edition, Princeton University Press, Princeton, NJ.
- SIMON, H.A. (1956) "Rational Choice and the Structure of the Environment," *Psychological Review*, Vol. 63(2), pp. 129–138.
- SMITH, A. (1776) *The Wealth of Nations*, Printed for W. Strahan and T. Cadell, London, UK. [Reprinted in 1994 by Modern Library, New York, NY.]
- SØRENSEN, P.N. (2007) "Simple Utility Functions with Giffen Demand," *Economic Theory*, Vol. 31, No. 2, pp. 367–370.
- SOWELL, T. (2010) *The Housing Boom and Bust (Revised Edition)*, Basic Books, New York, NY.
- SPENCE, A.M. (1973) "Job Market Signaling," *Quarterly Journal of Economics*, Vol. 87(3), pp. 355–374.

- SPULBER, D.F. (1999) *Market Microstructure: Intermediaries and the Theory of the Firm*, Cambridge University Press, Cambridge, UK.
- STIGLER, G.J. (1966) *The Theory of Price*, Macmillan, New York, NY.
- STIGLITZ, J.E. (1975) "The Theory of 'Screening,' Education, and the Distribution of Income," *American Economic Review*, Vol. 65, No. 3, pp. 283–300.
- STIGLITZ, J.E. (1989) "Imperfect Information in the Product Market," in: Schmalensee, R., Willig, R.D. (Eds.) *Handbook of Industrial Organization*, Elsevier, Amsterdam, NL, pp. 769–847.
- STIGLITZ, J.E. (1994) *Whither Socialism?* MIT Press, Cambridge, MA.
- STIGLITZ, J.E. (2006) *Making Globalization Work*, Norton, New York, NY.
- STOKEY, N.L. (1979) "Intertemporal Price Discrimination," *Quarterly Journal of Economics*, Vol. 93, No. 3, pp. 355–371.
- STOKEY, N.L. (1981) "Rational Expectations and Durable Goods Pricing," *Bell Journal of Economics*, Vol. 12, No. 1, pp. 112–128.
- STRULOVICI, B.H., WEBER, T.A. (2008) "Monotone Comparative Statics: Geometric Approach," *Journal of Optimization Theory and Applications*, Vol. 137, No. 3, pp. 641–673.
- STRULOVICI, B.H., WEBER, T.A. (2010) "Generalized Monotonicity Analysis," *Economic Theory*, Vol. 43, No. 3, pp. 377–406.
- THALER, R.H. (1994) *The Winner's Curse: Paradoxes and Anomalies of Economic Life*, Princeton University Press, Princeton, NJ.
- THALER, R.H. (1999) "Mental Accounting Matters," *Journal of Behavioral Decision Making*, Vol. 12(3), pp. 183–206.
- TOBIN, J. (1958) "Liquidity Preference as Behavior Towards Risk," *Review of Economic Studies*, Vol. 25, No. 2, pp. 65–86.
- TOPKIS, D.M. (1968) *Ordered Optimal Solutions*, Doctoral Dissertation, Stanford University, Stanford, CA.
- TOPKIS, D.M. (1998) *Supermodularity and Complementarity*, Princeton University Press, Princeton, NJ.
- TULLOCK, G. (1980) "Efficient Rent-Seeking," in: Buchanan, J.M., Tollison, R.D., Tullock, G. (Eds.) *Toward a Theory of the Rent-Seeking Society*, Texas A&M University Press, College Station, TX, pp. 97–112.
- TVERSKY, A., KAHNEMAN, D. (1981) "The Framing of Decisions and the Psychology of Choice," *Science*, Vol. 211, No. 4481, pp. 453–458.
- VEBLEN, T. (1899) *The Theory of the Leisure Class: An Economic Study of Institutions*, Macmillan, New York, NY. [Reprinted by Penguin, New York, NY, in 1994.]
- VICKREY, W. (1961) "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, Vol. 16, No. 1, pp. 8–37.
- VIVES, X. (1999) *Oligopoly Pricing*, MIT Press, Cambridge, MA.
- VON NEUMANN, J., MORGENSTERN, O. (1944) *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, NJ.
- VON NEUMANN, J. (1928) "Zur Theorie der Gesellschaftsspiele," *Mathematische Annalen*, Vol. 100(1), pp. 295–320.
- WALRAS, L. (1874/77) *Éléments d'Économie Politique Pure*, Corbaz, Lausanne, Switzerland. [English translation: *Elements of Pure Economics*, Allen and Unwin, London, UK, 1954.]
- WEBER, T.A. (2003) "An Exact Relation between Willingness to Accept and Willingness to Pay," *Economics Letters*, Vol. 80, No. 3, pp. 311–315.
- WEBER, T.A., BAPNA, A. (2008) "Bayesian Incentive Compatible Parametrization of Mechanisms," *Journal of Mathematical Economics*, Vol. 44, No. 3–4, pp. 394–403.
- WEBER, T.A., CROSON, D.C. (2004) "Selling Less Information for More: Garbling with Benefits," *Economics Letters*, Vol. 83, No. 2, pp. 165–171.
- WEBER, T.A., XIONG, H. (2007) "Efficient Contract Design in Multi-Principal Multi-Agent Supply Chains," *Proceedings of the 18th Annual Conference of the Production and Operations Management Society (POMS)*, Dallas, TX.
- WEBER, T.A., ZHENG, Z. (2007) "A Model of Search Intermediaries and Paid Referrals," *Information Systems Research*, Vol. 18, No. 4, pp. 414–436.
- WEBER, T.A. (2010) "Hicksian Welfare Measures and the Normative Endowment Effect," *American Economic Journal: Microeconomics*, Vol. 2, No. 4, pp. 171–194.

- WEBER, T.A. (2011) *Optimal Control Theory with Applications in Economics*, MIT Press, Cambridge, MA.
- WEITZMAN, M.L. (1974) "Prices vs. Quantities," *Review of Economic Studies*, Vol. 41, No. 4, pp. 477–491.
- WEITZMAN, M.L. (1979) "Optimal Search for the Best Alternative," *Econometrica*, Vol. 47, No. 3, pp. 641–654.
- WILLIAMSON, O.D. (1975) *Markets and Hierarchies: Analysis and Antitrust Implications*, Free Press, New York, NY.
- WILLIAMSON, O.D. (1985) *The Economic Institutions of Capitalism*, Free Press, New York, NY.
- WILLIG, R.D. (1976) "Consumer's Surplus Without Apology," *American Economic Review*, Vol. 66, No. 4, pp. 589–597.
- WILSON, R.B. (1987) "Bidding," in: Eatwell, J., Milgate, M., Newman, P. (Eds.) *The New Palgrave: A Dictionary of Economics, Vol. 1*, Macmillan, London, UK, pp. 238–242.
- WILSON, R. (1993) *Nonlinear Pricing*, Oxford University Press, Oxford, UK.
- WOLFERS, J., ZITZEWITZ, E. (2004) "Prediction Markets," *Journal of Economic Perspectives*, Vol. 18, No. 2, pp. 107–126.

Index

- Adverse selection, 20
- Arbitrage, 17
- Arrow impossibility theorem, 3
- Asset, 15
- Asset return matrix, 16
- Bertrand
 - Competition, 31
- Best-response (correspondence), 31
- Bounded rationality, 46
- Cheap talk, 38
- Club good, 24
- Coase conjecture, 44
- Coase theorem, 24
- Common good, 24
- Comparative statics, 7
- Compensating variation, *see* Welfare measures
- Complement (good), 29
- Conjectural variation, 32
- Conspicuous consumption, 8
- Consumer surplus, 27
- Contract curve, *see* Pareto set
- Core (of an economy), 13
- Cost (function), 28
- Demand
 - Excess, 18
 - Hicksian (wealth-compensated), 7
 - Walrasian, 5
- Demand elasticity, 29
- Diamond-Water Paradox, 1
- Elasticity
 - Cross-price, 29
 - Own-price, 29
- Endowment effect, 46
 - Normative, 47
- Equilibrium
 - Arrow-Debreu, 2
 - General, 28
 - Lindahl, 25
 - Nash, 30
 - Partial, 28
 - Rational expectations (Radner), 15, 16
 - Reactive, 38
 - Sunspot, 17
 - Walrasian, 2
- Equivalent martingale measure, 16
- Equivalent variation, *see* Welfare measures
- Event tree, 16
- Expected utility maximization, 9
- Expenditure function, 7
- Expenditure minimization problem, 7
- Extended polluter responsibility (EPR), 23
- Externalities, 22
- Ford effect, 28
- Fundamental welfare theorem
 - First, 11
 - Second, 12
- Game, 30
 - Simultaneous-move, 30
- Giffen good, 8
- Gittins index, 45
- Good
 - Classification, 25
 - Club, 24
 - Common, 24
 - Complement, 29
 - Giffen, 8
 - Inferior, 7
 - Nonexcludable, 24
 - Nonmarket, 22, 26
 - Nonrival, 24
 - Normal, 7
 - Numéraire, 2, 11
 - Perfect complements, 29
 - Positional, 7
 - Private, 24
 - Public, 24
 - Substitute, 29
 - Superior, 7
 - Veblen, 8
- Hicksian demand, 7
- Impossibility theorem, 3
- Increasing differences, 9
- Indifference, 3
- Inferior good, 7

- Information
 - Goods, 22
 - Inappropriability of, 22
 - Private, 21
 - Public, 21
- Invisible hand, 11
- Lattice, 8
- Lemons market, *see* Adverse selection
- Limit pricing, 33
- Luxury good, 7
- Marginal
 - Rate of substitution, 5
 - Revenue, 30
- Market
 - Two-sided, 43
- Market power, 28
- Market-clearing condition, 10
- Marshallian dynamics, 19
- Mechanism, 34
 - Bayesian implementation, 40
 - Design, 34, 39
 - Wilson doctrine, 47
 - Implementable, 35
- Monopoly pricing problem, 28
- Monopoly pricing rule
 - Multi-product, 29
 - Single-product, 29
- Monotone comparative statics, 8
- No-arbitrage pricing, 17
- Noise traders, 20, 21
- Nonexcludable good, 24
- Nonlinear pricing, 30
- Nonmarket good, 22, 26
- Nonrival good, 24
- Normal good, 7
- Numéraire good, 2, 11
- Pareto efficiency, 11
 - Constrained, 16
- Pareto set, 11, 13
- Perfect complements, 29
- Pigouvian taxation, 23
- Platform, 43
- Polluter Pays Principle, 23
- Portfolio investment problem, 9
- Positional good, 7
- Preference, 3
 - Locally nonsatiated, 5
 - Rational, 4
 - Representation, 4
- Price
 - Nominal, 1
- Price discrimination, 29
 - First-degree, 29
 - Second-degree (nonlinear pricing), 30
 - Third-degree, 29
- Price-taking behavior, 3, 5
- Pricing
 - Limit, 33
 - No-arbitrage, 17
- Private good, 24
- Private-ownership economy, 14
- Product differentiation
 - Horizontal, 31
- Production
 - Input, 28
 - Output, 28
- Production function
 - No-free-lunch property, 13
- Profit-maximization problem, 13
- Public good, 24
- Rent seeking, 33
- Revelation principle, 34
- Risk aversion
 - Absolute, 9
 - Constant absolute (CARA), 10
 - Constant relative (CRRA), 10
 - Relative, 10
- Risk premium, 9
- Roy's identity, 7
- Satisficing, *see* Bounded rationality
- Screening, *see* Nonlinear pricing
- Security, 15
- Signaling, 34, 37
- Slutsky matrix, 7
- Slutsky's identity, 7
- Sorting condition, 35
- Speculative bubble, 17
- Spence-Mirrlees sorting condition, 35
- Status-quo bias, *see* Endowment effect
- Strategy
 - Profile, 30
 - Set, 30
- Substitute (good), 29

- Substitution effect, 7
- Superior good, 7
- Supermodularity, 8, 9
- Taxation
 - Pigouvian, 23
- Transaction cost, 27
- Upper contour set, 12
- Utility
 - At the margin, 1
 - Expected, 9
 - Function, 4
 - Indirect, 6
- Utility maximization problem (UMP), 3, 4
 - Corner solution, 6
 - Under uncertainty, *see* Expected utility maximization
- Value, 1
 - In exchange, 1
 - In use, 1
 - Intrinsic, 17
- Veblen good, 8
- Walras' law, 5
- Walrasian
 - Tâtonnement*, 18, 19
 - Demand (correspondence), 5
 - Dynamics (for price adjustments), 19
 - Equilibrium, 2
- Wealth effect, 7
- Wealth-compensated demand, 7
- Welfare measures, 22
 - Compensating variation, 27
 - Equivalent variation, 27
- Willingness to accept (WTA), 27
- Willingness to pay (WTP), 27
- Wilson doctrine, 47