1

Aggregation of Information in Simple Market Mechanisms: Large Markets

The aim of this chapter is to provide an introduction to the most basic models of information aggregation: static simple market mechanisms where traders do not observe any market statistic before making their decisions. Each agent moves only once, simultaneously with other agents, and can condition his action only on his private information. The issue is whether market outcomes replicate or are close to the situation where agents have symmetric information and share the information in the economy. Examples of such market mechanisms are one-shot auctions and quantity (Cournot) and price (Bertrand) competition markets. In chapters 3–5 we will consider market mechanisms in the rational expectations tradition in which agents can use more complex strategies, conditioning their actions on market statistics. For example, a firm may use a supply function as a strategy, conditioning its output on the market price (chapter 3), or a trader may submit a demand schedule to a centralized stock market mechanism (chapters 4 and 5).

The plan of the chapter is as follows. Section 1.1 introduces the topic of information aggregation, some modeling issues, and an overview of results. Section 1.2 analyzes a large Cournot market with demand uncertainty and asymmetric information. It studies a general model and two examples: linear-normal and isoelastic-lognormal. Section 1.3 examines the welfare properties of price-taking equilibria in Cournot markets with private information. Section 1.4 presents a general smooth market model to examine information aggregation and the value of information. Section 1.5 deals with the case of auctions and section 1.6 introduces endogenous information acquisition.

1.1 Introduction and Overview

1.1.1 Do Markets Aggregate Information Efficiently?

As we stated in the Introduction and Lecture Guide, this has been a contentious issue at least since the debate between Hayek and Lange about the economic viability of socialism.

Hayek's basic idea (1945) is that each trader has some information that can be transmitted economically to others only through the price mechanism and trading. A planner cannot do as well without all that information. We say that the market aggregates information if it replicates the outcome that would be obtained if the agents in the economy shared their private information. The Hayekian hypothesis to check is whether a large (competitive) market aggregates information.

In this chapter we study, as a benchmark, market mechanisms that do not allow traders to condition their actions on prices or other market statistics. In this sense we are making things difficult for the market in terms of information aggregation. In chapter 3 we will allow more complex market mechanisms that allow traders to condition on current prices. Market mechanisms such as Cournot or auctions have the property that when a participant submits his or her trade it can condition only on his private information. For example, in a sealed-bid auction a bidder submits his bid with his private knowledge of the auctioned object but without observing the other bids; in a Cournot market firms put forward their outputs with some private estimate of demand conditions but without observing the market-clearing price or the outputs submitted by other producers.

A first question to ask is whether those simple market mechanisms aggregate information, at least when markets have many participants. We must realize, however, that a large market need not be competitive but rather can be monopolistically competitive. Indeed, firms may be small relative to the market but still retain some market power. If a large market is competitive and it aggregates information, then first-best efficiency follows according to the First Welfare Theorem. This means, in particular, that the full-information Walrasian model may be a good approximation to a large market with dispersed private information. In this case informational and economic efficiency go hand in hand. When a large market is not competitive, informational efficiency will not imply in general economic efficiency. We will discuss in detail the relationship between informational and economic efficiency in chapter 3.

Information aggregation does not obtain, in general, in market mechanisms in which outcomes depend continuously (smoothly) on the actions of the players, like quantity (Cournot) or price (Bertrand) competition markets with product differentiation. In fact, why should a market in which each trader conditions only on its private information be able to replicate the shared-information outcome? We will see that a large Cournot market in a homogeneous product world with a common shock to demand in which each producer receives a private signal about the uncertain demand does not aggregate information in general despite firms being approximately price takers. However, in the same context information may be aggregated if there are constant returns or if uncertainty is of the independent-values type, when the types of traders are independently distributed and, in fact, in the aggregate there is no uncertainty.

In contrast, winner-takes-all markets like auctions or voting mechanisms, in which outcomes do not depend continuously on the actions of players, tend to deliver aggregation of information more easily. Winner-takes-all markets force traders to condition effectively on more information when taking their decisions.

1.1.2 Methodological Issues and Welfare

In the following sections we use the continuum model of a large market. We will examine games with a continuum of agents in which no single one of them can affect the market outcomes. This methodology is in line with the literature on large Cournot markets with complete information (Novshek 1980), and with the view that the continuum model is the appropriate formalization of a competitive economy (Aumann 1964). The advantage of working with the continuum model is that it is very easy to understand the statistical reasons why a competitive market with incomplete information does not aggregate information efficiently in general, and to characterize the equilibrium and its second-best properties. One must check that the equilibria in the continuum economy are the limit of equilibria of finite economies and not artifacts of the continuum specification. This will be done in chapter 2.

The analysis of competitive equilibria in asymmetric-information economies is somewhat underdeveloped. To start with, the very notion of competitive equilibrium needs to be defined in an asymmetric-information environment.[1] If the market aggregates information, then the competitive equilibrium corresponds to the standard concept with full information. Otherwise, we may define the concept of Bayesian (price-taking) equilibrium. This is the situation, e.g., in a Cournot market, where firms' strategies depend on their private information but a firm does not perceive to affect the market price. This will be justified if the firm is very small in relation to the market, that is, in a large market.

Whenever the outcome of a large market, in which agents are price takers, is not first-best efficient, the question arises about what welfare property, if any, it has. The answer is that a price-taking Bayesian equilibrium maximizes expected total surplus subject to the restriction that agents use decentralized strategies (that is, strategies which depend only on the private information of the agents). This welfare benchmark for economies with incomplete information is termed *team efficiency* since the allocation would be the outcome of the decision of a team with a common objective (total surplus) but decentralized strategies (Radner 1962). This means that the large market performs as well as possible, subject to the constraint of using decentralized mechanisms. We will see this result in section 1.3.

1.2 Large Cournot Markets

In this section we will consider a large homogeneous product market in which demand is affected by a random shock and each firm has a private estimate of

[1] See Hellwig (1987) for a discussion of the issue. Progress in the study of competitive markets with asymmetric information has been made, among others, by Harris and Townsend (1981), Prescott and Townsend (1984), and Gale (1996).

the shock and sets an output. This is a very simple and standard framework in which to analyze information aggregation. Quantity setting corresponds to the Cournot model of an oligopolistic market. Here we will have many firms and each will be negligible in relation to the size of the market.

There is a continuum of firms indexed by $i \in [0,1]$.[2] Each firm has a convex, twice continuously differentiable, variable cost function $C(\cdot)$ and no fixed costs. Inverse demand is smooth and downward sloping, and given by $p = P(x,\theta)$, where x is average (per capita) output (and also aggregate output since we have normalized the measure of firms to 1) and θ is a random parameter. Firms are quantity setters.

Firm i receives a private signal s_i, a noisy estimate of θ. The signals received by firms are independently and identically distributed (i.i.d.) given θ and, without loss of generality, are unbiased, i.e., $E[s_i \mid \theta] = \theta$.

We make the *convention* that the strong law of large numbers (SLLN) holds for a continuum of independent random variables with uniformly bounded variances. Suppose that $(q_i)_{i\in[0,1]}$ is a process of independent random variables with means $E[q_i]$ and uniformly bounded variances $\mathrm{var}[q_i]$. Then we let $\int_0^1 q_i \,\mathrm{d}i = \int_0^1 E[q_i]\,\mathrm{d}i$ almost surely (a.s.). This convention will be used, taking as given the usual linearity property of the integral.[3] In particular, here we have that, given θ, the average signal equals $E[s_i \mid \theta] = \theta$ (a.s.).

Section 1.2.1 characterizes the equilibrium with strictly convex costs, section 1.2.2 its welfare properties, section 1.2.3 deals with the constant marginal cost case, and section 1.2.4 provides some examples.

1.2.1 Bayesian Equilibrium

Suppose that $C(\cdot)$ is strictly convex. A production strategy for firm i is a function $X_i(\cdot)$ which associates an output to the signal received. A market equilibrium is a Bayes–Nash (or Bayesian) equilibrium of the game with a continuum of players where the payoff to player i is given by

$$\pi_i = P(x;\theta)x_i - C(x_i), \quad \text{where } x = \int_0^1 X_i(s_i)\,\mathrm{d}i,$$

and the information structure is as described above.[4] At equilibrium, $X_i(s_i)$ maximizes

$$E[\pi_i \mid s_i] = x_i E[P(x;\theta) \mid s_i] - C(x_i)$$

and the firm cannot influence the market price $P(x;\theta)$ because it has no influence on average output x. This is, therefore, a price-taking Bayesian equilibrium. Restricting our attention to strategies with bounded means and uniformly bounded variances across players (this would obviously hold, for example, with

[2] The interval is endowed with the Lebesgue measure.

[3] See section 10.3.1 for a justification of the convention. Equality of random variables has to be understood to hold almost surely. We will not always insist on this in the text.

[4] See section 10.4.2 for an introduction to Bayesian equilibrium.

a bounded output space), the equilibrium must be symmetric. In particular, with $E[X_i(s_i) \mid \theta] < \infty$ and $\text{var}[X_i(s_i) \mid \theta]$ uniformly bounded across players, the random variables $X_i(s_i)$ are independent (given θ) and, according to our convention on a continuum of independent random variables, we have that

$$\int_0^1 X_i(s_i)\,\mathrm{d}i = \int_0^1 E[X_i(s_i) \mid \theta]\,\mathrm{d}i \equiv \tilde{X}(\theta).$$

It follows that the equilibrium must be symmetric, $X_i(s_i) = X(s_i)$ for all i, since the payoff is symmetric, the cost function is strictly convex and identical for all firms, and signals are i.i.d. (given θ). Indeed, for a given average output $\tilde{X}(\theta)$, the best response of a player is unique and identical for all players: $X_i(s_i) = X(s_i)$. Consequently, in equilibrium the random variables $X(s_i)$ will be i.i.d. (given θ) and their average will equal $\tilde{X}(\theta) = E[X(s_i) \mid \theta]$.

An interior (symmetric) equilibrium $X(\cdot)$, that is, one with positive production for almost all signals, is characterized by the equalization of the expected market price, conditional on receiving signal s_i, and marginal production costs:

$$E[P(x;\theta) \mid s_i] = C'(X(s_i)).$$

This characterizes the price-taking Bayesian equilibrium. Firms condition their output on their estimates of demand but they do not perceive, correctly in our continuum economy, any effect of their action on the (expected) market price. Thus the price-taking Bayesian equilibrium coincides with the Bayesian Cournot equilibrium of our large market.

1.2.2 Welfare and Information Aggregation

If firms were to know θ, then a Walrasian (competitive) equilibrium would be attained and the outcome would be first-best efficient. How does the Bayesian market outcome compare with the full-information first-best outcome, where total surplus (per capita) is maximized contingent on the true value of θ?

Given θ and individual production for firm i at x_i, total surplus (per capita) is

$$\text{TS} = \int_0^x P(z;\theta)\,\mathrm{d}z - \int_0^1 C(x_i)\,\mathrm{d}i, \quad \text{where } x = \int_0^1 x_i\,\mathrm{d}i.$$

If all the firms produce the same quantity, $x_i = x$ for all i, we have

$$\text{TS}(x;\theta) = \int_0^x P(z;\theta)\,\mathrm{d}z - C(x).$$

Given strict convexity of costs, first-best production $X^{\text{o}}(\theta)$ is given by the unique x which solves $P(x;\theta) = C'(x)$. If firms were able to pool their private signals, they could condition their production to the average signal, which equals θ a.s., and attain the first-best by producing $X^{\text{o}}(\theta)$. Will a price-taking Bayesian equilibrium, where each firm can condition its production only on its private information, replicate the first-best outcome?

A necessary condition for any (symmetric) production strategy $X(\cdot)$ to be first-best optimal is that, conditional on θ, identical firms produce at the same marginal cost, namely, $C'(X(s_i)) = C'(X^{\text{o}}(\theta))$ (a.s.). However, with increasing marginal costs this can happen only if $X(s_i) = X^{\text{o}}(\theta)$ (a.s.), which boils down to the perfect-information case. Therefore, we should expect a welfare loss at the price-taking Bayesian equilibrium with noisy signals. Proposition 1.1 states the result and a proof follows.

Proposition 1.1 (Vives 1988). *In a large Cournot market where firms have (symmetric) strictly convex costs and receive private noisy signals about an uncertain demand parameter, there is a welfare loss with respect to the full-information first-best outcome.*

Proof. We show that no symmetric production strategy $X(\cdot)$, and therefore no competitive production strategy, can attain the first-best outcome. Let $\tilde{X}(\theta) = \int_0^1 X(s_i)\,\mathrm{d}i$, then expected total surplus (per capita) contingent on θ is given by

$$E[\text{TS} \mid \theta] = \int_0^{\tilde{X}(\theta)} P(z;\theta)\,\mathrm{d}z - E[C(X(s_i)) \mid \theta].$$

We have that

$$\text{TS}(X^{\text{o}}(\theta);\theta) \geqslant \text{TS}(\tilde{X}(\theta);\theta) > E[\text{TS} \mid \theta].$$

The first inequality is true since $X^{\text{o}}(\theta)$ is the first-best, and the second is true since $\text{TS}(\tilde{X}(\theta);\theta) = \int_0^{\tilde{X}(\theta)} P(z;\theta)\,\mathrm{d}z - C(\tilde{X}(\theta))$, the cost function is strictly convex, $\tilde{X}(\theta) = E[X(s_i) \mid \theta]$, and the signals are noisy (which means that, given θ, $X(s_i)$ is still random). Consequently, $C(\tilde{X}(\theta)) < E[C(X(s_i)) \mid \theta]$ according to a strict version of Jensen's inequality (see, for example, Royden 1968, p. 110). □

1.2.3 Constant Marginal Costs

A necessary condition for the market outcome to be first-best optimal is that marginal costs be constant. Firms will then necessarily produce at the same marginal cost. If the information structure is "regular enough," first-best efficiency is achieved in the price-taking limit (Palfrey 1985; Li 1985). The intuition of the result is as follows. Suppose that marginal costs are zero (without loss of generality) and that inverse demand intersects the quantity axis. Firm i will maximize $E[\pi_i \mid s_i] = x_i E[P(x;\theta) \mid s_i]$, where, as before, x is average output. In this constant-returns-to-scale context a necessary condition for an interior equilibrium to exist is that $E[P(x;\theta) \mid s_i] = 0$ for almost all s_i. Looking at symmetric equilibria, $X_i(s_i) = X(s_i)$ for all i, we know that average production $\tilde{X}(\theta)$ given θ is nonrandom, and therefore $E[P(x;\theta) \mid \theta] = P(\tilde{X}(\theta);\theta)$. We will obtain first-best efficiency if the equilibrium condition, $E[P(\tilde{X}(\theta),\theta) \mid s_i] = 0$ for almost all s_i, implies that price equals marginal cost, $P(\tilde{X}(\theta),\theta) = 0$ for almost all θ.

Palfrey (1985) shows that if the signal and parameter spaces are finite and the likelihood matrix is of full rank (and demand satisfies some mild regularity conditions), then symmetric interior price-taking equilibria are first-best optimal. The result is easily understood with a two-point support example: θ can be either θ_L or θ_H, $0 < \theta_L \leqslant \theta_H$, with equal prior probability. Firm i may receive a low (s_L) or a high (s_H) signal about θ with likelihood $P(s_H \mid \theta_H) = P(s_L \mid \theta_L) = q$, where $\frac{1}{2} \leqslant q \leqslant 1$. If $q = \frac{1}{2}$, the signal is uninformative; if $q = 1$, it is perfectly informative (see section 10.1.6). Let $\bar{p}$ and $\underline{p}$ be, respectively, the equilibrium prices when demand is high (θ_H) and low (θ_L). Then

$$\begin{bmatrix} E[p \mid s_H] \\ E[p \mid s_L] \end{bmatrix} = \begin{bmatrix} q & 1-q \\ 1-q & q \end{bmatrix} \begin{bmatrix} \bar{p} \\ \underline{p} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

If the likelihood matrix is of full rank, i.e., if $q = \frac{1}{2}$, then $\bar{p} = \underline{p} = 0$ and the symmetric interior equilibrium is efficient.

Remark 1.1. For the result to hold, equilibria have to be interior. For example, suppose that $p = \theta - x$ and that $\theta_L = 0$ and $\theta_H = 1$, then at the price-taking Bayesian equilibrium $X(s_L) = 0$ and $X(s_H) = q/(q^2 + (1-q)^2)$. This means that $\tilde{X}(\theta_H) = X(s_H)q < X^0(\theta_H) = 1$.

1.2.4 Examples

1.2.4.1 Linear-Gaussian Model

Let $p = \theta - \beta x$ and $C(x_i) = mx_i + \frac{1}{2}\lambda x_i^2$, where m is possibly random and $\lambda \geqslant 0$. In this model only the level of $\theta - m$ matters and therefore without loss of generality let $m = 0$.

The joint distribution of random variables is assumed to yield affine conditional expectations. This, when coupled with the linear-quadratic structure of payoffs, yields a unique (and affine in information) Bayesian equilibrium. A leading example of such an information structure is the assumption of joint normality of all random variables. However, there are other pairs of prior distribution and likelihood which do not require unbounded support for the uncertainty and have the affine conditional expectation property.[5]

The random demand intercept θ is distributed according to a prior density with finite variance σ_θ^2 and mean $\bar{\theta}$. Firm i receives a signal s_i such that $s_i = \theta + \varepsilon_i$, where ε_i is a noise term with zero mean, variance $\sigma_{\varepsilon_i}^2$, and with $\mathrm{cov}[\theta, \varepsilon_i] = 0$. Signals can range from perfect ($\sigma_{\varepsilon_i}^2 = 0$ or infinite precision) to pure noise ($\sigma_{\varepsilon_i}^2 = \infty$ or zero precision). The precision of signal s_i is given by $\tau_{\varepsilon_i} = (\sigma_{\varepsilon_i}^2)^{-1}$. We assume that $E[\theta \mid s_i]$ is affine in s_i. All this implies that (see section 10.1)

$$E[\theta \mid s_i] = (1 - \xi_i)\bar{\theta} + \xi_i s_i, \quad \text{where } \xi_i \equiv \tau_{\varepsilon_i}/(\tau_\theta + \tau_{\varepsilon_i}),$$

and

$$E[s_j \mid s_i] = E[\theta \mid s_i], \quad \mathrm{cov}[s_i, s_j] = \mathrm{cov}[s_i, \theta] = \sigma_\theta^2 \quad \text{for all } j \neq i \text{ and all } i.$$

[5] The assumption of normality is very convenient analytically but has the drawback that prices and quantities may take negative values. However, the probability of this phenomenon can be controlled by controlling the parameters of the random variables. Alternatively, we could work with pairs of prior and likelihood functions which admit a bounded support and maintain the crucial property of linear conditional expectations which yields a tractable model. See section 10.2.2.

As τ_{ε_i} ranges from ∞ to 0, ξ_i ranges from 1 to 0. We also assume that the signals received by the firms are identically distributed conditional on θ. Then $\sigma^2_{\varepsilon_i} = \sigma^2_{\varepsilon}$ and $\xi_i = \xi$ for all i.

In equilibrium $X_i(s_i)$ solves

$$\max_{x_i} E[\pi_i \mid s_i] = E[p \mid s_i]x_i - \tfrac{1}{2}\lambda x_i^2,$$

where $p = \theta - \beta x$ and $x = \int_0^1 X_j(s_j)\,\mathrm{d}j$ and where $X_j(s_j)$ is the strategy used by firm j. Under the assumptions on strategies of section 1.2.1, we have that $\int_0^1 X_j(s_j)\,\mathrm{d}j = \tilde{X}(\theta)$. The optimal response of firm i is identical for all firms provided that $\lambda > 0$, $X(s_i) = E[p \mid s_i]/\lambda$, and equilibria will be symmetric. We are looking for a linear (affine) solution of the type $X(s_i) = as_i + c$, where a and c are coefficients to be determined. It follows that $\tilde{X}(\theta) = \int_0^1 X(s_j)\,\mathrm{d}j = a\int_0^1 s_j\,\mathrm{d}j + c = a\theta + c$ using our convention on the SLLN in our continuum economy: $\int_0^1 s_j\,\mathrm{d}j = \theta + \int_0^1 \varepsilon_j\,\mathrm{d}j = \theta$ since $\int_0^1 \varepsilon_j\,\mathrm{d}j = 0$ (a.s.). The first-order condition (FOC) of the program of firm i yields

$$\lambda X(s_i) = E[\theta \mid s_i] - \beta E[\tilde{X}(\theta) \mid s_i].$$

Therefore, we have that

$$\lambda a s_i + \lambda c = E[\theta \mid s_i] - \beta a E[\theta \mid s_i] - \beta c.$$

Since $E[\theta \mid s_i] = (1-\xi)\bar{\theta} + \xi s_i$, we obtain

$$\lambda a s_i + \lambda c = (1 - \beta a)(\xi s_i + (1-\xi)\bar{\theta}) - \beta c.$$

This must hold for (almost) all signals and therefore, solving

$$\lambda a = (1-\beta a)\xi \quad \text{and} \quad \lambda c = (1-\beta a)(1-\xi)\bar{\theta} - \beta c,$$

we obtain

$$a = \frac{\xi}{\lambda + \beta\xi} \quad \text{and} \quad c = \frac{(1-\beta a)(1-\xi)}{\lambda + \beta}\bar{\theta} = \frac{1}{\lambda+\beta}\bar{\theta} - a\bar{\theta}$$

(where in order to obtain the last equality we use the value for a).

In conclusion, we have shown that the unique linear (affine to be precise) function $X(\cdot)$ which satisfies the FOC is given by

$$X(s_i) = a(s_i - \bar{\theta}) + b\bar{\theta}, \quad \text{where } a = \xi/(\lambda + \beta\xi) \text{ and } b = 1/(\lambda + \beta).$$

The linear equilibrium identified can be shown to be in fact the unique equilibrium in the class of strategies with bounded means and uniformly bounded variances (across players). (See remark 1.2 and proposition 7.7 and exercise 7.3 for a proof in a closely related model.)

Average output conditional on θ is given by

$$\int_0^1 X(s_i)\,\mathrm{d}i = \int_0^1 (a(s_i - \bar{\theta}) + b\bar{\theta})\,\mathrm{d}i = a(\theta - \bar{\theta}) + b\bar{\theta}$$

since the average signal conditional on θ equals θ. Note that average output is in fact the same function as the equilibrium strategy $X(\cdot)$ and we may denote it

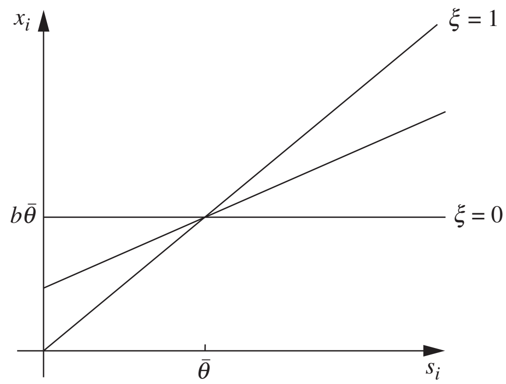

Figure 1.1. Equilibrium strategy of firm i for different values of ξ. The middle line corresponds to $\xi \in (0, 1)$.

by $X(\theta)$. When firms receive no information ($\xi = 0$), then $a = 0$ and production is constant at the level $b\bar{\theta}$; when firms receive perfect information ($\xi = 1$), then $X_i(s_i) = as_i$. As ξ varies from 0 to 1 the slope of the equilibrium strategy, a, increases from 0 to b (see figure 1.1).

When $\lambda = 0$ the equilibrium condition requires that $E[p \mid s_i] = 0$. Restricting our attention to symmetric equilibria, we get that $X(s_i) = s_i/\beta$ provided that $\xi > 0$. Indeed, note that

$$E[p \mid s_i] = E[\theta - \beta x \mid s_i] = E[\theta \mid s_i] - \beta E[X(\theta) \mid s_i] = 0,$$

and therefore

$$E[\theta \mid s_i] - \beta a(E[\theta \mid s_i] - \bar{\theta}) - \beta b\bar{\theta} = (1 - \beta a)E[\theta \mid s_i] + \beta\bar{\theta}(a - b) = 0,$$

which implies that $a = b = 1/\beta$. That is, the above formula is valid on letting $\lambda = 0$, since then $a = b = 1/\beta$. Note that, when the signals are informative ($\xi > 0$), the equilibrium strategies $X(s_i) = s_i/\beta$ do not depend on ξ, the precision of the information. If $\xi = 0$, signals are uninformative and $X(s_i) = \bar{\theta}/\beta$.

In summary, given $\lambda \geqslant 0$ and $\xi \in [0, 1]$ the equilibrium of the continuum economy is given by $X(s_i) = a(s_i - \bar{\theta}) + b\bar{\theta}$ with $a = \xi/(\lambda + \beta\xi)$ ($a = 0$ if $\lambda = \xi = 0$) and $b = 1/(\lambda + \beta)$.

It is worth noting that this equilibrium is the outcome of iterated elimination of strictly dominated strategies if and only if the ratio of the slopes of supply and demand is less than 1 ($\lambda\beta < 1$). This is the familiar cobweb stability condition (see Guesnerie 1992; Heinemann 2004). In this case the equilibrium is the outcome only of the rationality of the players and common knowledge about payoffs and distributions.[6]

Welfare. Assuming that all the firms produce the same quantity $x_i = x$, we have $\mathrm{TS}(x; \theta) = \int_0^x P(z; \theta)\,\mathrm{d}z - C(x)$. Since $P(z; \theta) = \theta - \beta z$ and $C(x) = \frac{1}{2}\lambda x^2$,

[6] See section 10.4.1.1 for the concept of dominated strategy and dominance solvability in games.

we obtain that per capita total surplus with all firms producing output x is given by

$$\mathrm{TS} = \int_0^x (\theta - \beta z)\,\mathrm{d}z - \tfrac{1}{2}\lambda x^2 = \theta x - \left(\frac{\beta + \lambda}{2}\right)x^2.$$

It is immediate then that full-information first-best production is

$$X^{\mathrm{o}}(\theta) = \frac{\theta}{\lambda + \beta} = b\theta,$$

and it equals $X(\theta)$ only if $\xi = 1$ (perfect information) or if $\lambda = 0$ and $\xi > 0$ (constant returns to scale). As ξ goes from 0 to 1, a increases from 0 to b. For $\xi < 1$ and $\lambda > 0$ if θ is high (low) the market underproduces (overproduces) with respect to the full-information first-best. In the constant-returns-to-scale case, average production $X(\theta)$ is independent of the precision of information τ_ε (and of ξ) and the market produces the right amount. Expected total surplus (ETS) at the first-best is

$$\mathrm{ETS}^{\mathrm{o}} = \tfrac{1}{2}bE[\theta^2] = \tfrac{1}{2}b(\sigma_\theta^2 + \bar{\theta}^2).$$

The market ETS can be computed as the sum of per capita consumer surplus $\frac{1}{2}\beta E[(X(\theta))^2]$ and per capita (or firm) profits:

$$E[\pi_i] = \tfrac{1}{2}\lambda E[(X(s_i))^2] = \tfrac{1}{2}\lambda(a^2 \operatorname{var}[s_i] + b^2\bar{\theta}^2).$$

We find that $\mathrm{ETS} = \frac{1}{2}(a\sigma_\theta^2 + b\bar{\theta}^2)$ and therefore the welfare loss $\mathrm{WL} \equiv \mathrm{ETS}^{\mathrm{o}} - \mathrm{ETS}$ equals $\frac{1}{2}(b - a)\sigma_\theta^2$. (See exercise 1.1 for an alternative derivation.)

How does ETS change with variations in the precision of information τ_ε and in the basic uncertainty of demand σ_θ^2?

One would hope that improvements in the precision of information (τ_ε) would increase ETS (reducing the welfare loss). This is indeed the case since ETS increases with a, and a is in turn increasing with τ_ε. Per capita expected consumer surplus (ECS) increases with τ_ε since consumer surplus is a convex function of average output ($\mathrm{CS} = \frac{1}{2}\beta x^2$) and increases in ξ increase the slope of $X(\cdot)$ and make it more variable. It can be checked that expected profits may increase or decrease with τ_ε. Increasing τ_ε increases the sensitivity of output to the signal received by the firm, and this tends to increase expected profits, but it also decreases the variance of the signal, and this tends to depress expected profits. (A fuller explanation of the drivers of the comparative statics of expected profits is given in section 1.4.3.) However, the effect on the welfare of the consumers always dominates the profit effect in the total surplus computation.

With $\lambda > 0$, increasing the basic uncertainty of demand increases ETS (by increasing both ECS and expected profits) but $\mathrm{ETS}^{\mathrm{o}}$ increases by more and the welfare loss increases with[7] σ_θ^2; increasing the precision of information

[7] To show the result notice that

$$\frac{\partial\,\mathrm{WL}}{\partial\sigma_\theta^2} = \frac{1}{2b}(b^2 - a^2).$$

increases ETS and ETS$^{\mathrm{o}}$ stays constant and, therefore the welfare loss decreases with τ_ε. With $\lambda = 0$, WL $= 0$ provided that $\tau_\varepsilon > 0$ (but with $\tau_\varepsilon = 0$, WL increases with σ_θ^2).

Similar results hold when we replace quadratic costs by capacity limits. Indeed, a capacity limit can be interpreted as a very steep marginal cost schedule. This can be formally checked in a model with constant marginal costs and capacity limits with a finite support information structure (see exercise 1.2). The qualitative properties of the capacity model are identical to those of the model analyzed in this section.

1.2.4.2 Isoelastic-Lognormal Model

Let $p = \mathrm{e}^\theta x^{-\beta}$, $\beta > 0$, and the information structure be exactly as in section 1.2.4.1, where θ and the error terms of the signals are jointly normally distributed. Firms have constant elasticity cost functions given by

$$C(x_i) = (1+\lambda)^{-1} x_i^{1+\lambda}, \quad \lambda \geqslant 0.$$

We find an equilibrium in log-linear strategies. This specification avoids the unpalatable feature of the linear-normal model where outputs and prices may become negative.

The FOC for profit maximization yields $E[p \mid s_i] = x_i^\lambda$ and, as in the general model, the equilibrium must be symmetric when the cost function is strictly convex (i.e., for $\lambda > 0$). It can be easily checked that there is a unique symmetric equilibrium in which the price is a log-linear function of θ. For this we use the fact that, if z is normally distributed $N(\mu, \sigma^2)$ and r is a constant, then $E[\mathrm{e}^{rz}] = \mathrm{e}^{r\mu + r^2\sigma^2/2}$ (see section 10.2.4). Postulate an equilibrium of the form $X(s_i) = \mathrm{e}^{as_i+b}$. Using our convention about the average of a continuum of independent random variables (which implies that $\int_0^1 \mathrm{e}^{\varepsilon_i}\,\mathrm{d}i = \int_0^1 E[\mathrm{e}^{\varepsilon_i}]\,\mathrm{d}i$) and the properties of lognormal distributions, namely that if $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$, then $E[\mathrm{e}^{\varepsilon_i}] = \mathrm{e}^{\sigma_\varepsilon^2/2}$, we obtain

$$\int_0^1 \mathrm{e}^{s_i}\,\mathrm{d}i = \int_0^1 \mathrm{e}^{\theta+\varepsilon_i} = \mathrm{e}^\theta \int_0^1 \mathrm{e}^{\varepsilon_i}\,\mathrm{d}i = \mathrm{e}^\theta \mathrm{e}^{\sigma_\varepsilon^2/2} \quad \text{(a.s.)}.$$

This implies that

$$\tilde{X}(\theta) = \int_0^1 X(s_i)\,\mathrm{d}i = \int_0^1 \mathrm{e}^{as_i+b}\,\mathrm{d}i = \mathrm{e}^{a\theta+b+a^2\sigma_\varepsilon^2/2},$$

$$p = \mathrm{e}^{(1-\beta a)\theta - \beta(b + a^2\sigma_\varepsilon^2/2)}.$$

From the FOC $E[p \mid s_i] = x^\lambda$ we obtain

$$E[\exp\{(1-\beta a)\theta - \beta(b + \tfrac{1}{2}a^2\sigma_\varepsilon^2)\} \mid s_i] = \mathrm{e}^{\lambda a s_i + \lambda b}.$$

Given that $\theta \mid s_i \sim N(\bar{\theta}(1-\xi) + \xi s_i, \sigma_\theta^2(1-\xi))$, we have that the left-hand side

---

Therefore, $\partial\,\mathrm{WL}/\partial\sigma_\theta^2 > 0$ if and only if $b > a$. This is the case if $\lambda > 0$ and $\xi < 1$ or if $\lambda = 0$ and $\xi = 0$ (no information).

equals

$$\exp\{(1-\beta a)(\bar{\theta}(1-\xi)+\xi s_i)+\tfrac{1}{2}(1-\beta a)^2(1-\xi)\sigma_\theta^2-\beta(b+\tfrac{1}{2}a^2\sigma_\varepsilon^2)\}.$$

By identifying coefficients on s_i we obtain $a = \xi/(\lambda+\beta\xi)$. Doing the same with the constant term, substituting for the value of a, and using the fact that $\xi = \tau_\varepsilon/(\tau_\varepsilon+\tau_\theta) = \sigma_\theta^2/(\sigma_\theta^2+\sigma_\varepsilon^2)$ and therefore $\sigma_\varepsilon^2 = ((1-\xi)/\xi)\sigma_\theta^2$, and simplifying, we obtain

$$b = \frac{1-\xi}{\lambda+\beta}\left[\left(\frac{\lambda}{\lambda+\beta\xi}\right)\bar{\theta}+\left(\frac{\sigma_\theta^2}{2(\lambda+\beta\xi)^2}\right)[\lambda^2-\beta\xi]\right].$$

In conclusion, the equilibrium is given by

$$X(s_i) = \mathrm{e}^{as_i+b},$$

where

$$a = \frac{\xi}{\lambda+\beta\xi} \quad \text{and} \quad b = \frac{1-\xi}{\lambda+\beta}\left(\frac{\lambda}{\lambda+\beta\xi}\bar{\theta}+\frac{\lambda^2-\beta\xi}{(\lambda+\beta\xi)^2}\frac{\sigma_\theta^2}{2}\right).$$

Note that the full-information output (corresponding to the case $\xi = 1$ and $\sigma_\varepsilon^2 = 0$) is $X^{\mathrm{o}}(\theta) = \mathrm{e}^{\theta/(\lambda+\beta)}$.

The output of a firm is more sensitive to its signal the better the information is (a increases and b decreases with ξ) and $E[\pi_i]$ increase or decrease with ξ depending on whether demand is elastic or inelastic (β smaller or larger than 1). The same forces as in the linear-normal model are present here. Profits are a convex function of output; increasing ξ increases the responsiveness of output to the signal, this induces more output variation and is good for expected profits. However, as in section 1.2.3, the decreased variance of the signal works in the opposite direction, decreasing expected profits. With constant returns to scale the (full-information) competitive outcome is obtained and then $E[\pi_i]$ are independent of ξ.

1.2.5 Summary

The main learning points of the section are the following.

- A large market, even inducing price-taking behavior, may fail to be first-best efficient because of a lack of information aggregation.
- This is the case for a large Cournot market with common-value uncertainty and decreasing returns because firms with different demand assessments will not produce at the same marginal cost.
- With constant returns to scale a large Cournot market will aggregate information under regularity conditions. In this case the market attains the first-best aggregate output and this is all that matters for welfare purposes.

1.3 Welfare in Large Cournot Markets with Asymmetric Information

The first-best full-information outcome is too stringent a benchmark of comparison for the market outcome. The reason goes back to the idea of Hayek that the private information of agents may not be easily communicable and therefore that there is no omniscient center that knows the realization of the uncertainty θ. The welfare benchmark must therefore be a decentralized one where agents use strategies which are measurable in their information. This is the approach pioneered in Vives (1988). In our market we will say that an allocation is *team-efficient* if it maximizes expected total surplus with decentralized strategies. A team is a group of people with a common objective. Team allocations with private information have been studied by Radner (1962). This will be equivalent to the solution to the problem of a planner who wants to maximize expected total surplus and can control the action (strategy) of an agent but can make it contingent only on the private information of the agent and not on the information of other agents.

We show below that the allocation of a price-taking equilibrium in a Cournot market is team-efficient, where each firm follows a production rule, contingent on its private information, with the common objective of maximizing expected total surplus. That is, the price-taking market solves the team problem with expected total surplus as an objective function. This provides a general welfare characterization of price-taking equilibrium in a Cournot market with private information allowing for a general information structure. A price-taking equilibrium obtains in the case of a continuum of firms as in section 1.2.

Consider an n-firm Cournot market. Firm i, $i = 1, \dots, n$, has smooth convex costs, $C(x_i; \theta_i)$, where x_i is its output. Inverse demand is given by $P(x; \theta_0)$, a smooth and downward-sloping function of total output $x = \sum_{j=1}^{n} x_j$. Suppose that firm i receives a private signal vector s_i about the (potentially) random parameters (θ_0, θ_i). At a price-taking Bayesian equilibrium, firm i maximizes its expected profits (conditional on receiving signal s_i):

$$E[\pi_i \mid s_i] = E[P(x; \theta_0) \mid s_i]x_i - E[C(x_i; \theta_i) \mid s_i],$$

without taking into account the influence of its output on the market price $P(x; \theta_0)$. That is, for an interior equilibrium we find that (for almost all s_i) the expected price equals the expected marginal cost:

$$E[P(x; \theta_0) \mid s_i] = E[\mathrm{MC}(x_i; \theta_i) \mid s_i].$$

The following result provides an analogue to the First Welfare Theorem for price-taking Bayesian equilibria. We say that firms use decentralized strategies if each firm can choose its output as a function only of its own signal.

Proposition 1.2 (Vives 1988). *In a smooth Cournot private-information environment, price-taking Bayesian equilibria maximize expected total surplus (ETS) subject to the use of decentralized production strategies.*

Proof. The maximization of ETS subject to decentralized production strategies is the problem of a team whose members have as common objective

$$\text{ETS} = E\left[\int_0^x P(z;\theta_0)\,\mathrm{d}z - \sum_{j=1}^{n} C(X_j(s_j);\theta_j)\right],$$

with strategies $X_j(s_j)$, $j = 1,\dots,n$, where $x = \sum_{j=1}^{n} X_j(s_j)$. Under our assumptions the optimal decision rules are determined (for interior solutions) by the set of FOCs $E[\partial TS/\partial x_i \mid s_i] = 0$ or, equivalently, $E[P(x;\theta_0) \mid s_i] = E[\text{MC}(x_i;\theta_i) \mid s_i]$. Indeed, a set of decision rules are optimal if and only if they are person-by-person optimal given that the team function is concave and differentiable (Radner 1962, theorem 1). The conditions are fulfilled in our case. Now, price-taking firm i will maximize

$$E[\pi_i \mid s_i] = E[P(x;\theta_0) \mid s_i]x_i - E[C(x_i;\theta_i) \mid s_i],$$

yielding an FOC,

$$E[P(x;\theta_0) \mid s_i] = E[\text{MC}(x_i;\theta_i) \mid s_i].$$

These conditions are sufficient given our assumptions and therefore the solutions to both problems coincide. □

The result also applies to our continuum economy. Consider the payoff for player $i \in [0,1]$, $\pi_i = P(x;\theta)x_i - C_i(x_i)$, where $P(\cdot\,;\theta)$ is smooth and downward sloping and $C_i(\cdot)$ is smooth and strictly convex. We have that $\text{TS} = \int_0^x P(z;\theta)\,\mathrm{d}z - \int_0^1 C_i(x_i)\,\mathrm{d}i$, where $x = \int_0^1 x_i\,\mathrm{d}i$. The team problem is to find strategies $(X_i(\cdot))_{i\in[0,1]}$, with bounded means $E[X_i(s_i) \mid \theta] < \infty$ and uniformly bounded variances $\text{var}[X_i(s_i) \mid \theta]$ across players, that maximize ETS. We have that

$$\int_0^1 X_i(s_i)\,\mathrm{d}i = \int_0^1 E[X_i(s_i) \mid \theta]\,\mathrm{d}i \equiv \tilde{X}(\theta)$$

and

$$\text{ETS} = E\left[\int_0^{\tilde{X}(\theta)} P(z;\theta)\,\mathrm{d}z - \int_0^1 C_i(X_i(s_i))\,\mathrm{d}i\right].$$

The Bayesian equilibrium implements the team-efficient solution.

Remark 1.2. The fact that equilibria can also be obtained as the outcome of the optimization of a strictly concave welfare function can be used to show uniqueness of the equilibrium. For example, if the welfare function is quadratic and strictly concave, as is TS in the linear-normal model of section 1.2.4.1, the team solution is unique and therefore the Bayesian equilibrium is also unique (in the class of strategies with bounded means and uniformly bounded variances across players). The insight is more general and applies whenever the outcome of a game can be replicated optimizing an appropriate potential function. We will use it to show uniqueness of a Bayesian equilibrium in a game with a finite number of players in chapter 2 (see proposition 2.1 and its proof).

Remark 1.3. Price-taking Bayesian equilibria will differ from Bayesian Cournot equilibria in a market with a finite number of firms. This is so because at a Bayesian Cournot equilibrium a firms takes into account the impact of its output on the market price. However, as the economy increases in size (increasing at the same time the number of firms and consumers, perhaps because of lowering of entry costs) Bayesian Cournot equilibria converge to price-taking Bayesian equilibria, as we will see in section 2.2. At the limit economy with a continuum of firms they coincide. Proposition 1.2 applies similarly to the Bayesian equilibrium of the limit continuum economy. A price-taking Bayesian equilibrium is efficient as long as only decentralized strategies can be used.

Remark 1.4. According to proposition 1.2, in a Cournot environment there is no room for a planner to improve on market performance taking as given decentralized decision making. This contrasts with environments where agents can condition on prices and therefore there are potential informational externalities with prices as a public signal. Chapter 3 will deal with the welfare analysis of such models in the rational expectations tradition.

In summary, we have claimed that an appropriate welfare benchmark for our incomplete-information economy is team efficiency, where expected total surplus is maximized subject to the use of decentralized strategies, and we have shown that a price-taking equilibrium is team-efficient. This means in particular that a *large* Cournot market, where firms are effectively price takers, is team-efficient.

1.4 Information Aggregation in Smooth Large Markets

A large market need not aggregate information, that is, it need not replicate the shared-information outcome, as we have seen in the Cournot case in section 1.2. In fact, we will see in this section that a large market aggregates information only in very particular circumstances other than the independent-values case. Furthermore, large markets may be competitive (price-taking) or monopolistically competitive. In the latter case, each firm produces a differentiated commodity, is negligible in the sense that its actions alone do not influence the profits of any other firm, and has some monopoly power (this is the Chamberlinian large group case).[8] The monopolistically competitive market is not efficient even with complete information. That is, even if the market were to aggregate information, there would be a welfare loss in relation to the full-information first-best allocation. A question arises about the value of information in this situation.

In this section we present a quadratic payoff game for which monopolistic competition is a leading example. The monopolistically competitive market

[8] See Vives (1990 and chapter 6 in 1999).

will be a large market linear-quadratic model with normally distributed random variables that encompasses Cournot or Bertrand competition with product differentiation and uncertainty of common- or private-value type.

More generally, consider a game among a continuum of players where each player has a (symmetric) smooth payoff function $\pi(y_i, \tilde{y}; \theta_i)$ with y_i the action of player i, $\tilde{y}$ a vector of statistics (e.g., mean and variance) that characterizes the distribution of the actions of players, and θ_i a possibly idiosyncratic payoff-relevant random parameter. Suppose that player i receives a signal s_i about the parameter θ_i. As before a strategy for player i is a measurable function $Y_i(\cdot)$ from the signal space to the action space of the player. A set of strategies $(Y_i(\cdot))_{i\in[0,1]}$ forms a Bayesian equilibrium if for any player (almost surely)

$$Y_i(s_i) \in \arg\max_{z_i} E[\pi(z_i, \tilde{y}; \theta_i) \mid s_i],$$

where $\tilde{y}$ is the vector of statistics that characterizes the equilibrium distribution of the actions of players. As in section 1.2 player i when optimizing takes as given the equilibrium statistics since his action cannot influence them. A linear-quadratic-Gaussian specification of the game is presented in the following section. We in turn analyze information aggregation and perform a welfare analysis.

1.4.1 A Linear-Quadratic-Gaussian Model

Consider a quadratic profit function model with a continuum of players. The payoff to player i is

$$\pi(y_i, y; \theta_i) = \theta_i y_i - \tfrac{1}{2}\omega_1 y_i^2 - \omega_2 y y_i,$$

where $(-\partial^2\pi/(\partial y_i)^2) = \omega_1 > 0$ ensures strict concavity of π with respect to the firm's action y_i (in the real line), and where $y = \int_0^1 y_j \,\mathrm{d}j$ denotes the average action. Actions are strategic complements (substitutes) if $\partial^2\pi/\partial y_i \partial y = -\omega_2 > (<)\ 0$.[9] Under complete information the best response of a player to the aggregate action y is

$$y_i = \frac{\theta_i - \omega_2 y}{\omega_1}.$$

Since the slope of the best response is

$$\kappa \equiv \frac{\partial^2\pi/\partial y_i \partial y}{-\partial^2\pi/(\partial y_i)^2} = -\frac{\omega_2}{\omega_1},$$

this quotient provides a natural measure of the degree of strategic complementarity or substitutability. A game is of strategic complementarities if the best response for each player increases with the actions of rivals. We assume that the slope $\kappa < 1$ and therefore $\omega_1 + \omega_2 > 0$ always. This allows for games of strategic substitutability and of strategic complementarities with a bounded degree of complementarity. An example of a strategic complementarity ($\omega_2 < 0$) game is

[9] Section 10.4.1.2 provides a brief introduction to games of strategic complementarities.

provided by an adoption externalities or investment complementarities game in which y_i is the adoption or investment effort with return $(\theta_i - \omega_2 y)y_i$ and cost $\frac{1}{2}\omega_1 y_i^2$.[10] The return to adoption or investment increases with the aggregate effort y.

We can obtain Bertrand and Cournot competition, with linear demands and constant unit costs (equal to zero without loss of generality), by choosing parameters appropriately. The case of Cournot competition with quadratic production costs can also be accommodated. The payoffs are in line with the following demand system with random intercepts (x and p are, respectively, average quantity and price):

$$p_i = \alpha_i - (1-\delta)x_i - \delta x \quad \text{with } \delta \in [0,1],$$
$$x_i = \beta_i - (1+\gamma)p_i + \gamma p \quad \text{with } \gamma \geqslant 0,$$

where (p_i, x_i) is the price–output pair of firm i. We can obtain the second equation by inverting the first and letting $\beta_i = (\alpha_i - \delta\tilde{\alpha})/(1-\delta)$ and $\gamma = \delta/(1-\delta)$ with $\tilde{\alpha} = \int_0^1 \alpha_i \, \mathrm{d}i$. When $\delta = 0$, the firms are isolated monopolies, and when $\delta = 1$, they are perfect competitors because then the product is homogeneous and firms are price takers. The parameter δ represents the degree of product differentiation and when $0 < \delta < 1$ the market is monopolistically competitive. In the quantity competition (Cournot) case, let $\theta_i = \alpha_i$, $\frac{1}{2}\omega_1 = 1 - \delta$, and $\omega_2 = \delta$. In the price competition (Bertrand) case, let $\theta_i = \beta_i$, $\frac{1}{2}\omega_1 = 1 + \gamma$, and $\omega_2 = -\gamma$. In the Cournot case with homogeneous product and increasing marginal costs (as in section 1.2) we set $\omega_1 = \lambda$, $\omega_2 = \beta$, and $\theta_i = \theta$ to obtain $\pi_i = (\theta - \beta x)x_i - \frac{1}{2}\lambda x_i^2$. Note that in the Cournot (Bertrand) $\omega_2 > 0$ ($\omega_2 < 0$) case actions are strategic substitutes (complements).

Assume that the information structure is symmetric and given as follows.[11] Each pair of parameters (θ_i, θ_j) is jointly normally distributed with $E[\theta_i] = \bar{\theta}$, $\operatorname{var}[\theta_i] = \sigma_\theta^2$, and $\operatorname{cov}[\theta_i, \theta_j] = \varsigma\sigma_\theta^2$ for $j \neq i$, $0 \leqslant \varsigma \leqslant 1$. Agent i receives a signal $s_i = \theta_i + \varepsilon_i$, where $\theta_i \sim N(\bar{\theta}, \sigma_\theta^2)$, $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$, and $\operatorname{cov}[\varepsilon_i, \varepsilon_j] = 0$ for $j \neq i$. The error terms of the signals are also independent of the θ parameter. The precision of signal s_i is given by $\tau_{\varepsilon_i} = (\sigma_\varepsilon^2)^{-1}$. As before we let $\xi = \tau_\varepsilon/(\tau_\theta + \tau_\varepsilon)$.

Our information structure encompasses the cases of "common value" and of "private values." For $\varsigma = 1$ the θ parameters are perfectly correlated and we are in a *common-values* model. When signals are perfect, $\sigma_{\varepsilon_i}^2 = 0$ for all i, and $0 < \varsigma < 1$, we will say we are in a *private-values* model. Agents receive idiosyncratic shocks, which are imperfectly correlated, and each agent observes his shock with no measurement error. When $\varsigma = 0$, the parameters are independent, and we are in an *independent-values* model.

It is not difficult to see (see section 10.2.3) that

$$E[\theta_i \mid s_i] = \xi s_i + (1-\xi)\bar{\theta} \quad \text{and} \quad E[s_j \mid s_i] = E[\theta_j \mid s_i] = \xi\varsigma s_i + (1-\xi\varsigma)\bar{\theta}.$$

[10] Models in this vein have been presented by, among many others, Diamond (1982), Bryant (1983), Dybvig and Spatt (1983), and Matsuyama (1995).

[11] See section 10.2.3 for results on the family of information structures to which the one presented here belongs.

When signals are perfect, $\xi = 1$, $E[\theta_i \mid s_i] = s_i$, and $E[\theta_j \mid s_i] = \varsigma s_i + (1-\varsigma)\bar{\theta}$. When they are not informative, $\xi = 0$ and $E[\theta_i \mid s_i] = E[\theta_j \mid s_i] = \bar{\theta}$.

We can also derive the relationship between θ_i, s_i, and the average parameter $\tilde{\theta} = \int_0^1 \theta_j \, \mathrm{d}j$. The average parameter $\tilde{\theta}$ is normally distributed with mean $\bar{\theta}$ and variance equal to $\varsigma\sigma_\theta^2$. This is in accordance with the finite-dimensional analogue for the average of a collection of symmetrically correlated random variables (see section 10.2.3). We have that $E[\theta_i \mid \tilde{\theta}] = \tilde{\theta}$, $E[\tilde{\theta} \mid \theta_i] = E[\theta_j \mid \theta_i] = \varsigma\theta_i + (1-\varsigma)\bar{\theta}$, $E[\tilde{\theta} \mid s_i] = E[\theta_j \mid s_i]$, and

$$E[\theta_i \mid \tilde{\theta}, s_i] = (1-d)\tilde{\theta} + ds_i,$$

where $d = (\sigma_\theta^2(1-\varsigma))/(\sigma_\theta^2(1-\varsigma) + \sigma_\varepsilon^2) = (\tau_\varepsilon(1-\varsigma))/(\tau_\varepsilon(1-\varsigma) + \tau_\theta) = (1-\varsigma)/(\xi^{-1}-\varsigma)$. If signals are perfect, then $d = 1$ and $E[\theta_i \mid \tilde{\theta}, s_i] = s_i$. If signals are useless or correlation is perfect ($\varsigma = 1$), then $d = 0$ and $E[\theta_i \mid \tilde{\theta}, s_i] = \tilde{\theta}$. If both signals and correlation are perfect, then $E[\theta_i \mid \tilde{\theta}, s_i] = \tilde{\theta} = s_i$ (a.s.).

There is a unique symmetric Bayesian equilibrium in linear strategies in the incomplete-information game. This is also the case if agents share information. Observe that $\tilde{s} = \int_0^1 \theta_i \, \mathrm{d}i + \int_0^1 \varepsilon_i \, \mathrm{d}i = \tilde{\theta}$, since $\int_0^1 \varepsilon_i \, \mathrm{d}i = 0$ according to our convention on the average of i.i.d. random variables as the error terms ε_i are uncorrelated and have mean zero. Agents who share information need only know the average signal $\tilde{s}$ (on top of their signal). Indeed, given the information structure and linear equilibrium, $(s_i, \tilde{s})$ is a sufficient statistic[12] in the estimation of θ_i by agent i, that is, to estimate θ_i with the pooled information available, firm i need only look at $(s_i, \tilde{s})$. Furthermore, with $\tilde{s}$ $(= \tilde{\theta})$ the firm can predict with certainty the aggregate action y since y depends only on $\tilde{\theta}$ in a linear equilibrium. The following proposition characterizes the equilibrium. The proof is in the appendix.

Proposition 1.3 (Vives 1990). *There is a unique symmetric linear Bayesian equilibrium in the case of both private and shared information. The equilibrium strategy of agent i in the private-information case is given by* $Y(s_i) = a(s_i - \bar{\theta}) + b\bar{\theta}$, *where* $a = \xi/(\omega_1 + \omega_2\varsigma\xi)$, $\xi = \tau_\varepsilon/(\tau_\varepsilon + \tau_\theta)$, *and* $b = 1/(\omega_1 + \omega_2)$. *In the shared-information case, it is given by* $Z(s_i, \tilde{\theta}) = \hat{a}(s_i - \tilde{\theta}) + b\tilde{\theta}$, *where* $\hat{a} = d/\omega_1$ *and* $d = (1-\varsigma)/(\xi^{-1} - \varsigma)$.

Note that $\omega_1 + \omega_2\varsigma\xi > 0$ since $\omega_1 > 0$, $\omega_1 + \omega_2 > 0$, and $0 < \varsigma\xi < 1$. Therefore, a and $\hat{a}$ are nonnegative and b is positive.

Remark 1.5. As in sections 1.2.1 and 1.3 there cannot be asymmetric equilibria as long as we restrict our attention to equilibria with bounded means and uniformly bounded variances (across players). It can be checked that the equilibrium can be obtained by optimizing a strictly concave potential (or team) function which delivers a unique solution.

[12] See section 10.1.4 for the statistical concept of sufficiency. Intuitively, an aggregate signal about a parameter is a sufficient statistic for the information available if the posterior distribution given the aggregate signal is the same as the one given all the individual signals.

1.4.2 When Does a Large Market Aggregate Information?

The linear-quadratic model has a certainty-equivalence property. The expected value of an individual (as well as average) action, with either pooling or not pooling of information, equals $b\bar{\theta}$, which is the action agents would choose if they did not have any information. However, equilibria are not independent of pooling arrangements except in particular circumstances. The exceptions are if $a = \hat{a}$ and $\tilde{\theta} = \bar{\theta}$ (a.s.) or if $b = a = \hat{a}$. The first case obtains if there is no correlation between parameters, $\varsigma = 0$ (independent values) and the second obtains if $\omega_1(1-\xi) = \omega_2\xi(1-\varsigma)$. This latter equality obviously holds in the perfect-information case, $\varsigma = \xi = 1$. It cannot hold if $\omega_2 < 0$ (strategic complements), or for ω_1 and ω_2 different from zero, in the common-value case ($\varsigma = 1$, $\xi \in (0,1)$), or in the private values case ($\xi = 1$, $\varsigma \in (0,1)$).

In summary, the large market aggregates information if and only if $\varsigma(\omega_1(1-\xi) - \omega_2\xi(1-\varsigma)) = 0$.

In line with our previous analysis in the common-value case (section 1.2.1), the equality holds with constant marginal costs, $\omega_1 = \lambda = 0$ (Palfrey 1985), but it does not hold with strictly convex costs, $\lambda > 0$, provided signals are not perfect, $\xi < 1$ (Vives 1988). With independent values the information aggregation result should not be surprising because in the limit there is no aggregate uncertainty. In equilibrium a firm can predict with certainty the average action in the market. In fact, what seems surprising is that there are *any* circumstances under which information aggregation obtains when aggregate uncertainty (coupled with imperfect signals) remains in the limit.

The knife-edge information aggregation result of the common-value case with constant marginal costs extends to other parameter configurations along the curve $\omega_1(1-\xi) - \omega_2\xi(1-\varsigma) = 0$. In general, in our linear-normal model the equilibrium strategy of firm i in the shared-information regime depends both on its private signal s_i and on the average signal, which in the limit equals the average parameter: $\tilde{s} = \tilde{\theta}$. Parameter configurations along the line $\omega_1(1-\xi) - \omega_2\xi(1-\varsigma) = 0$ have the property that the equilibrium strategy does not depend on $\tilde{\theta}$ in the limit. This may seem surprising, since when $\varsigma > 0$ the average market output $\tilde{y}$ is random and firms can predict it exactly knowing the average parameter $\tilde{\theta}$. How can an agent not respond to the information contained in $\tilde{\theta}$? Notice first that $\tilde{\theta}$ gives information about the average action $Y(\tilde{\theta})$ and about θ_i, since the θ_i parameters are correlated. In the Cournot model (with actions being strategic substitutes) a high $\tilde{\theta}$ is good news for agent i since this means that θ_i is likely to be high but it is bad news for agent i at the same time since it means that $Y(\tilde{\theta})$ is also likely to be high, which tends to lower the payoff of the agent. For the specific parameter configurations $\omega_1(1-\xi) - \omega_2\xi(1-\varsigma) = 0$ it so happens that the two forces exactly balance out and it is optimal not to respond to changes in $\tilde{\theta}$. It is worth noting, however, that this could not happen if the actions of the firms were strategic complements $\omega_2 < 0$ (say, prices in a differentiated product market with constant marginal costs) since then a low

$\tilde{\theta}$ is bad news on two counts: it means that the demand intercept β_i and the average price are likely to be low and then the agent wants to set a low price. In this situation it is never optimal not to respond to $\tilde{\theta}$.

In summary, a large smooth market aggregates information only in particular circumstances. An important example is when types are independent because then there is no aggregate uncertainty. Another relevant case is when there is Cournot competition, common-value uncertainty, and constant returns. In general, however, we should not expect to have information aggregation when there is aggregate uncertainty in a smooth market.

1.4.3 Welfare Analysis and the Value of Information

Suppose that we are in the (usual) case that the market does not aggregate information. Will more precise private information for all players always be to the benefit of each one of them? The answer turns out to depend on whether competition is of the strategic substitutes or complements type. Let us consider the common-value case.

In the common-value case expected profits increase with a uniform increase in the precision of the signal τ_ε with strategic complements (Bertrand case) and may increase or decrease with τ_ε with strategic substitutes (Cournot case). In any case expected profits increase with prior uncertainty σ_θ^2 (Vives (1990) and exercise 1.4). However, expected profits always increase with a uniform increase in the precision of the signals for a given signal correlation ($\text{cov}[s_i, s_j]/\text{var}[s_i]$), and for given signal precisions, expected profits increase (decrease) with increased correlation of signals with strategic complements (substitutes) (see section 8.3.1 in Vives 1999). With strategic complements (substitutes) an increased (decreased) signal correlation is good for expected profits. The reason is that with strategic complements (substitutes) best responses are upward (downward) sloping. Once the correlation of the signals received by players is controlled for, increasing the precision of signals is always good.

The consequences for welfare in the monopolistic competition model are as follows. Expected consumer and total surplus increase (decrease) with the precision of private information in the Cournot (Bertrand) case. The intuition for these results is derived from the form of the consumer surplus function. Expected consumer surplus as a function of quantities is a convex combination of the variance of individual and average output, and expected consumer surplus as a function of prices increases with the variance of individual prices and decreases with the variance of the average price and with the covariance of the demand shock and the individual price. (See exercise 1.4.)

It is worth noting again that, in contrast to the competitive economy in section 1.2, even if firms were to have complete information there would be a welfare loss in the market due to monopolistic behavior. Furthermore, under incomplete information, the monopolistically competitive market is not team-efficient. That is, the market allocation does not maximize expected total surplus subject to decentralized strategies since the team-efficient allocation would imply that expected prices are equal to marginal costs.

Angeletos and Pavan (2007a) provide a characterization of the equilibrium and efficient use of information as well as the social value of information in a closely related model that generalizes the quadratic payoff structure in the common-value case including the dispersion (standard deviation) of the actions of the players. This allows encompassing other applications such as beauty contest games where the payoff to a player depends also on the distance between the action of the player and the actions of other players (Morris and Shin 2002). Angeletos and Pavan find that stronger complementarity increases the sensitivity of equilibrium actions to public information, raising aggregate volatility while stronger substitutability increases the sensitivity to private information, raising the cross-sectional dispersion of actions.

Consider a symmetric game with a continuum of players in which the payoff to player i is $\pi(y_i, y, \sigma_y; \theta)$ with y_i his action, $y = \int_0^1 y_i \,\mathrm{d}i$ the average action, $\sigma_y = (\int_0^1 (y_i - y)^2 \,\mathrm{d}i)^{1/2}$ the standard deviation of the action distribution, and θ a common payoff-relevant random parameter. Assume that $\pi(y_i, y, \sigma_y; \theta) = u(y_i, y; \theta) + \frac{1}{2}\nu\sigma_y^2$ with u quadratic and

$$\frac{\partial^2 u}{(\partial y_i)^2} < 0, \quad \frac{\partial^2 u}{(\partial y_i)^2} + 2\frac{\partial^2 u}{\partial y_i \partial y} + \frac{\partial^2 u}{(\partial y)^2} < 0, \quad \frac{\partial^2 u}{(\partial y_i)^2} + \nu < 0$$

to ensure global concavity (note that $\nu = \partial^2\pi/(\partial\sigma_y)^2$). It is also assumed that the slope of the best response of a player is less than 1:

$$\kappa \equiv \frac{\partial^2\pi/\partial y_i \partial y}{-\partial^2\pi/(\partial y_i)^2} < 1.$$

Each player receives a private and a public signal about θ and all variables are jointly normally distributed.

Angeletos and Pavan characterize the linear equilibrium of the game. The equilibrium can be shown to be unique as in proposition 1.3 and remark 1.5.[13] They show that the equilibrium strategy is a convex combination of the expectation of the full-information equilibrium allocation and the expectation of aggregate activity (both conditional on the information set of a player) where the weight given to the latter is precisely κ. This is called the "equilibrium degree of coordination" by Angeletos and Pavan but we will call it the degree of complementarity. When $\kappa = 0$ the equilibrium strategy is just the best predictor of the full-information equilibrium allocation. This means that the weights, respectively, to public and private information, in the equilibrium strategy are just the Bayesian weights. When $\kappa > 0$ there is strategic complementarity and players then weigh more public information, and when $\kappa < 0$ there is strategic substitutability and players then weigh less public information. The reason is that, with strategic complementarity, when a player wants to align his action with the average action, to predict the average action it is better to weigh more

[13] In the paper it is claimed to be the unique one, at least if $\kappa \in (-1, 1)$.

public information, while, with strategic substitutability, when a player wants to differentiate his action from the average action, it is better to weigh more private information. An increase in κ decreases the dispersion of equilibrium activity ($\mathrm{d}\,\mathrm{var}[y_i - y]/\mathrm{d}\kappa < 0$) and increases the nonfundamental volatility as measured by $\mathrm{d}\,\mathrm{var}[y - y^{\mathrm{f}}]/\mathrm{d}\kappa > 0$, where y^{f} is the full-information equilibrium allocation.

Angeletos and Pavan also perform a welfare analysis using team efficiency as a benchmark. This is defined in relation to the aggregate welfare of the players in the game. An allocation is then team-efficient if it maximizes the expected utility of a representative player using a (symmetric) decentralized strategy. We introduced this benchmark in section 1.3 using as team function total surplus in a partial equilibrium market. Both definitions are closely connected. Indeed, with the present definition this would imply, in the context of the competitive linear-normal economy of section 1.2.4.1, for example, to take consumers with quasilinear quadratic preferences as players who at the same time own the firms. We would then see that the price-taking Bayesian allocation is also team-efficient under the present definition. At the team-efficient solution both payoff externalities and the efficient use of information, with no communication, are taken into account. A team-efficient allocation is then uniquely characterized by a convex combination of the expectation of the (full-information) first-best allocation and the expectation of aggregate activity (both conditional on the information set of a player) where the weight given to the latter is

$$\kappa^{\mathrm{e}} \equiv 1 - \frac{\partial^2\pi/(\partial y_i)^2 + 2(\partial^2\pi/\partial y_i \partial y) + \partial^2\pi/(\partial y)^2}{\partial^2\pi/(\partial y_i)^2 + \partial^2\pi/(\partial \sigma_y)^2}.$$

In the quadratic environment it then follows that the efficient allocation for the original game can be replicated by the equilibrium of a fictitious game with the same information structure, in which the full-information equilibrium equals the first-best allocation of the original economy and where the degree of complementarity is precisely κ^{e}. This is the fictitious degree of complementarity that guarantees efficiency under incomplete information once we have efficiency under complete information. Some interesting results follow.

In an efficient economy (that is, one which is efficient under complete information and for which $\kappa = \kappa^{\mathrm{e}}$), the social value of public versus private information increases with κ (we have that $(\partial E[\pi]/\partial\tau_\eta)/(\partial E[\pi]/\partial\tau_\varepsilon) = \tau_\varepsilon/(1-\kappa)\tau_\eta$, where τ_ε (τ_η) is the precision of the private (public) information).

If we take an economy which is efficient under complete information, then when $\kappa > \kappa^{\mathrm{e}}$ ($\kappa < \kappa^{\mathrm{e}}$) there is overreaction (underreaction) to public information and excessive nonfundamental volatility (cross-sectional dispersion).

An example is provided by the beauty contest model of Morris and Shin (2002). The idea of the beauty contest goes back to Keynes and views the financial market as a contest where investors try to guess what other investors will do instead of trying to assess the fundamentals (see the Introduction and Lecture

Guide and section 8.3). In this example,

$$\pi(y_i, y, \sigma_y; \theta) = -(1-r)(y_i - \theta)^2 - r\left(L_i - \int_0^1 L_i \, \mathrm{d}i\right)$$

with $r \in (0,1)$ and $L_i = \int_0^1 (y_i - y_j)^2 \, \mathrm{d}j = (y_i - y)^2 + \sigma_y^2$. Note that $\int_0^1 L_i \, \mathrm{d}i = 2\sigma_y^2$.

Players try to get close to θ but derive a private value from taking an action close to others. However, socially the latter attempt is a waste since $\int_0^1 \pi_i \, \mathrm{d}i = -(1-r)\int_0^1 (y_i - \theta)^2 \, \mathrm{d}i$. This economy is efficient under complete information,

$$\kappa = \frac{\partial^2 \pi / \partial y_i \partial y}{-\partial^2 \pi / (\partial y_i)^2} = \frac{2r}{2} = r \quad \text{and} \quad \kappa^{\mathrm{e}} = 0$$

(since $\partial^2 \pi / (\partial y)^2 = -2r$ and $\partial^2 \pi / (\partial \sigma_y)^2 = 2r$). Therefore, there is overreaction to public information and excessive nonfundamental volatility. In fact, Morris and Shin (2002) show that welfare may decrease with the precision of public information. This fact has prompted a debate on the desirability of transparency of central bank policy. The point is that the disclosures of a central bank may reduce welfare whenever the beauty contest analogy applies to financial markets (see Morris and Shin 2005; Hellwig 2005; Svensson 2006; Woodford 2005). We will come back to this issue in chapter 3 and in the models of pure informational externalities in chapter 6.[14]

Another example is provided by new Keynesian business cycle models (e.g., Woodford 2002; Hellwig 2005), where the economy is efficient under complete information. However, in this case welfare increases with public information (Hellwig 2005). The reason is that in those models the externality with the dispersion in relative prices is negative ($\partial^2 \pi / (\partial \sigma_y)^2 < 0$), due to imperfect substitutability across goods in a monopolistic competition model, and this leads to a higher κ^{e} with $\kappa < \kappa^{\mathrm{e}}$.

Angeletos and Pavan (2007a) also look at the monopolistic competition model developed in this section, which is not efficient under complete information, and derive some comparative statics results with respect to information parameters. For example, expected profits increase with the precision of public information when competition is of the strategic complements variety. In this case (as in the quadratic model of section 1.4.1) we have that $\kappa^{\mathrm{e}} = 2\kappa$.[15]

Angeletos and Pavan (2007b) look at Pigouvian corrective tax policy in a similar environment and show that if the government can set marginal tax rates contingent on aggregate activity the (decentralized) efficient allocation can be

---

[14] Calvó-Armengol and de Martí (2007) consider a game with n players with payoffs similar to the beauty contest example and model the signals received as the outcome of communication in a network.

[15] There is in fact an apparent contradiction of their corollary 10 (Cournot competition), in which they claim that expected profits increase with the precision of private information, with the comparative statics result presented above that expected profits may increase or decrease with this precision (see exercise 1.4). This is due to the fact that corollary 10 only holds under the parameter restriction $\kappa > -1$ and this fact is not stated in the corollary.

implemented. The basic idea is to introduce taxes that control payoff externalities, so that the complete information equilibrium coincides with the first best allocation, and make agents perceive the fictitious degree of complementarity κ^e instead of the true κ, so that there is a socially efficient use of information. Tax progressivity is a crucial instrument for the first objective and the sensitivity of marginal taxes to aggregate activity to the second. It is found, for example, that in economies which are inefficient only under incomplete information and for which $\kappa > \kappa^e$, marginal optimal taxes are increasing in aggregate activity. In this way the perceived degree of complementarity is reduced.

1.5 Auctions and Voting

In this section we examine information aggregation properties of auction and voting mechanisms. We will see how auctions aggregate information under less restrictive conditions than smooth market mechanisms, such as a Cournot market. The better information aggregation properties of auction and voting mechanisms are explained by their winner-takes-all feature. This property implies that a bidder, while submitting his bid, or a voter, while casting his vote, has to think, respectively, of the implications that winning or being pivotal for the election outcome have in terms of the signals that other players may have received. That is, the bidder, effectively, has to condition on the information that winning conveys. In section 1.5.1 we examine common-value auctions and in section 1.5.2 we draw some connections with voting mechanisms.

1.5.1 Information Aggregation in Common-Value Auctions

Do auction markets aggregate information efficiently? This question got an affirmative answer in the studies of Wilson (1977) and Milgrom (1979, 1981). Wilson (1977) considers a first-price auction (the bidder with the highest bid wins the object and pays his bid) where buyers have some private information about the common value θ of the good to be sold and shows that, under certain regularity conditions, as the number of bidders tends to ∞ the maximum bid is almost surely equal to the true value. Milgrom (1979) obtains a necessary and sufficient condition on the information structure so that convergence to the true value is in probability.[16] The (second price) Vickrey auction has the same type of limiting properties (Milgrom 1981).

The conditions required for obtaining aggregation of information are relatively strong. In equilibrium, winning the object means that the other $n - 1$ bidders have received worse signals about the value of the object (the "winner's curse" since this means that the winner may have overestimated the value of the object). Consider the following illustration of the winner's curse in a common-value auction. An oil tract is auctioned and bidders have private estimates (signals) of its value. If bidders were to bid naively, each one just on the basis of

[16] See section 10.3.1 for an account of the different convergence concepts for random variables.

his private estimate, the winner would have overbid. The reason is that winning means that you have bid above the other bidders. To avoid the winner's curse, each bidder has to condition his bid not only on his private signal but also on the information that winning the auction conveys (Milgrom and Weber 1982). Now, for a high bid to be optimal, when $n - 1$ is large, the signal of the bidder must be quite strong. In particular, it must be the case that for any value θ there is a signal for which, conditional on winning, the bidder puts very small probability on a value lower than θ.[17]

Pesendorfer and Swinkels (1997) show that information aggregation is obtained under less stringent information conditions provided that supply also becomes large as the number of bidders grows (see also Kremer 2002). Pesendorfer and Swinkels consider an auction of k identical objects of unknown value. The k highest bidders obtain an object and pay the $(k+1)$th bid. Each bidder receives a signal of bounded precision (more precisely, the posterior of a bidder after receiving a signal has full support, with density bounded away from zero and infinity) before submitting his bid. This contrasts with the setup in Wilson and Milgrom. Pesendorfer and Swinkels characterize the unique symmetric Bayesian equilibrium. They show that a necessary and sufficient condition for the equilibrium price to converge in probability to the true value is that both the number of objects sold k and the number of bidders who do not receive an object $n - k$ go to ∞. The driving force behind the convergence result is that a loser's curse is added (losing means that at most $n - k$ of the other buyers bid below your own bid) to the usual winner's curse (winning means that at most k of the other buyers bid above your own bid). When both k and $n - k$ tend to ∞, then signals need not be very strong for the equilibrium price to converge to the value. The reason is that a bid on the brink of winning or losing conveys precise information on the true value. Furthermore, for the result to hold supply need not grow in proportion to demand (it may grow less quickly) and only a vanishing fraction of bidders may be informed.

1.5.2 Voting

Voting aggregates information in ways similar to auctions. Consider a two-candidate election in which voters have private information about a common-value characteristic θ of the alternatives. For example, the common unknown value may be the quality of a public good about which voters have different assessments. The connection between auctions and elections is that a voter must condition his beliefs about θ on the event that his vote is pivotal (that is, that his vote can change the outcome of the election) in the same way that a bidder must condition his bid on the event that he wins the auction (with the highest bid).

---

[17] Milgrom's (1981) necessary and sufficient condition for information aggregation with a finite set of values is that for any $\theta < \theta'$ and any constant K there is a signal s' which yields a likelihood ratio on θ' versus θ of at least K. (See section 10.1.5 for an explanation of the likelihood ratio.)

Feddersen and Pesendorfer (1997) analyze an election with two alternatives A and Q in which the payoff to a voter depends on his type, the (one-dimensional) state of nature, and the winning alternative. Voting is costless. Each voter knows his type (types are i.i.d.) and receives a private signal about the uncertain state of nature. The alternative Q wins the election if it gets at least a fraction of votes q. Feddersen and Pesendorfer characterize symmetric Bayes–Nash equilibria of the voting game in which no voter uses a weakly dominated strategy. In equilibrium there are types who always vote for A, types who always vote for Q, and the rest cast their vote depending on the signal received (i.e., they take an “informative action”).

Feddersen and Pesendorfer show that elections aggregate information. That is, the alternative chosen under shared information obtains with probability close to 1 in a large election with private information. We say that an election is close if in equilibrium the winning candidate obtains a fraction of votes very close to the winning percentage. Feddersen and Pesendorfer show that although the fraction of voters who take an informative action tends to 0 as the number of voters grows unboundedly, large elections are almost always very close. As a consequence, elections are decided by those who take an informative action. However, information aggregation is not obtained in general if the distribution from which preference types are drawn is uncertain. An interesting related insight is that uninformed voters may prefer to abstain rather than vote because of the analogue to the winner’s curse in the context of the election (Feddersen and Pesendorfer 1996). This is so because uninformed voters may abstain to maximize the probability that informed voters are pivotal and decide the outcome of the election.

In summary, when the market mechanism has the winner-takes-all feature, as in auctions or voting mechanisms, information aggregation seems to be facilitated. The reason is that these mechanisms force agents to take into consideration the informational implications of their winning the auction or being pivotal in the election.

1.6 Endogenous Information Acquisition

In this section we examine the implications of costly information acquisition for information aggregation in the context of the Cournot market of section 1.2.[18] We confirm that with endogenous information acquisition a large Cournot market with decreasing returns does not aggregate information and does not attain first-best efficiency. Furthermore, the welfare loss increases with the cost of information acquisition. The market, however, is still second-best optimal when firms can only use decentralized production strategies. The case of constant returns to scale is more subtle.

[18] Milgrom (1981) and Milgrom and Weber (1982) consider information acquisition in auctions.

After setting up the model we deal first with the decreasing-returns-to-scale case (section 1.6.1) and then with the constant-returns-to-scale case (section 1.6.2).

Consider the model with a continuum of firms, uncertain demand, and linear-normal specification of the example in section 1.2.4.1: $p = \theta - \beta x$ and $C(x_i) = \frac{1}{2}\lambda x_i^2$. The timing of events is as follows. In the first stage firms contract information (a sample) of a certain precision about the unknown parameter θ. At the second stage firms receive their private signals and compete in quantities. The precision of information of firm i, τ_{ε_i}, is proportional to the sample size the firm obtains.[19] The cost of obtaining information with precision τ_{ε_i} is $c\tau_{\varepsilon_i}$, where $c > 0$. It will be more convenient to work with ξ_i ($\xi_i \equiv \tau_{\varepsilon_i}/(\tau_\theta + \tau_{\varepsilon_i})$) and, with some abuse of language, we will speak of ξ_i as the precision of information of firm i. The cost of obtaining information ξ_i will then be given by $\varphi(\xi_i) = (c/\sigma_\theta^2)(\xi_i/(1-\xi_i))$. Purchase of zero precision ($\xi_i = 0$) is costless and to purchase perfect information ($\xi_i = 1$) is infinitely costly. The parameters ξ_i are common knowledge at the second stage (in a competitive market, though, a firm needs to know only the average precision).

An alternative would be to consider the case where firms purchase precision simultaneously with their output choice or, equivalently, that the precision acquisition by a firm is not observable by the rivals. In this case a strategy for firm i is a pair $(\xi_i, X_i(\cdot))$ determining the precision purchased and the output strategy. In our continuum economy it does not matter whether the precisions purchased are observable or not. Note first that if $(\xi_i, X_i(\cdot))_{i\in[0,1]}$ is a Nash equilibrium of the one-shot game, then necessarily $(X_i(\cdot))_{i\in[0,1]}$ is a (Bayes) Nash equilibrium of the market game for given $(\xi_i)_{i\in[0,1]}$. Let ξ denote the average information precision in the market. Then equilibrium strategies and expected profits of firm i at the market stage are easily seen to depend only on ξ_i and ξ. This means that there is no strategic effect derived from the information purchases at the first stage. The reason is that a single firm must take the average ξ as given. It follows that equilibrium precisions, both at the one-shot and two-stage games, will solve

$$\operatorname*{Max}_{\xi_i} \Pi_i(\xi_i, \xi) = E[\pi_i] - \varphi(\xi_i).$$

1.6.1 Decreasing Returns to Scale

Assuming that $\lambda > 0$ (increasing marginal costs), we can find a price-taking Bayesian equilibrium of this (continuum) economy as in section 1.2 for given $(\xi_i)_{i\in[0,1]}$. Let us postulate $X_i(s_i) = a_i(s_i - \bar{\theta}) + b\bar{\theta}$ as a candidate linear equilibrium strategy for firm i.

[19] The model is taken from Vives (1988). Li et al. (1987) model information acquisition similarly.

We then have that

$$\int_0^1 X_i(s_i)\,\mathrm{d}i = \int_0^1 (a_i(s_i - \bar{\theta}) + b\bar{\theta})\,\mathrm{d}i$$
$$= a(\theta - \bar{\theta}) + \int_0^1 a_i\varepsilon_i\,\mathrm{d}i + b\bar{\theta} = a(\theta - \bar{\theta}) + b\bar{\theta},$$

where $a = \int_0^1 a_j\,\mathrm{d}j$, and, assuming that $\mathrm{var}[a_i\varepsilon_i]$ is uniformly bounded across firms and given that $E[a_i\varepsilon_i] = 0$, we have $\int_0^1 a_i\varepsilon_i\,\mathrm{d}i = 0$ according to our convention on the integral.

We can use the FOC $E[p \mid s_i] = \lambda X_i(s_i)$, where $E[p \mid s_i] = E[\theta \mid s_i] - \beta E[\tilde{X}(\theta) \mid s_i]$ and $E[\tilde{X}(\theta) \mid s_i] = a(E[\theta \mid s_i] - \bar{\theta}) + b\bar{\theta}$ to solve for the parameters a_i and obtain $a_i = \xi_i/(\lambda + \beta\xi)$ and $b = 1/(\lambda + \beta)$. Note that $\mathrm{var}[a_i\varepsilon_i] = \xi_i(\lambda + \beta\xi)^{-2}(\tau_{\varepsilon_i} + \tau_\theta)^{-1}$, which is bounded between 0 and 1. Using the FOC it is immediate that

$$E[\pi_i] = \tfrac{1}{2}\lambda E[(X_i(s_i))^2] = \frac{\lambda}{2}\left(b^2\bar{\theta}^2 + \frac{\xi_i\sigma_\theta^2}{(\lambda + \beta\xi)^2}\right).$$

Expected profits for firm i increase with ξ_i and decrease with ξ. The total payoff for firm i is

$$\Pi_i(\xi_i, \xi) = \frac{\lambda}{2}\left(\frac{\xi_i}{(\lambda + \beta\xi)^2}\sigma_\theta^2 + \frac{1}{(\lambda + \beta)^2}\bar{\theta}^2\right) - \frac{c}{\sigma_\theta^2}\frac{\xi_i}{1 - \xi_i}.$$

The marginal benefit to a firm to acquire information decreases with the average amount of information purchased by the other firms ξ. This means that ξ_i and ξ are strategic substitutes in the payoff of firm i. More information in the market reduces the incentive for any firm to do research. A firm wants information to estimate the market price $p = \theta - \beta x$. When firms have better information (ξ is high) the market price varies less since producers match better production decisions with the changing demand. Consequently, for high ξ an individual firm has less incentive to acquire information.

The FOC (which is sufficient) is given by

$$\frac{\partial \Pi_i}{\partial \xi_i} = \frac{\lambda\sigma_\theta^2}{(\lambda + \beta\xi)^2} - \frac{c}{\sigma_\theta^2}\frac{1}{(1 - \xi_i)^2} \leqslant 0.$$

It will hold with equality if $\xi_i > 0$. Since the equilibrium will be symmetric, we let $\xi_i = \xi$ and solve for ξ. We obtain

$$\xi^* = \max\left\{0, \frac{\sigma_\theta^2 - \sqrt{2c\lambda}}{\beta\sqrt{2c/\lambda} + \sigma_\theta^2}\right\},$$

which can be written in terms of the precision of the signal as

$$\tau_\varepsilon^* = \max\left\{0, \frac{1 - \tau_\theta\sqrt{2c\lambda}}{\lambda + \beta}\sqrt{\frac{\lambda}{2c}}\right\}.$$

It is immediate that the precision purchased is monotone in cost c. If $c = 0$, firms obtain perfect information ($\tau_\varepsilon^* = \infty$); as c increases, τ_ε^* decreases monotonically until c is so high that no information is purchased. Furthermore, τ_ε^*

increases with σ_θ^2. More uncertainty (or less prior public knowledge τ_θ) induces the firms to acquire more information. The relationship between the slope of marginal costs (λ) and τ_ε^* is not monotonic: for both small and large λ, expected profits at the market stage are low and consequently there is low expenditure on research, and τ_ε^* peaks at intermediate values of λ. It increases with λ for low values of λ and decreases for high values.

Differentiated products. The results generalize to differentiated products. This is easily seen by noting that the model examined is formally equivalent to a linear demand differentiated product market with constant marginal costs (equal to zero for simplicity) as studied in section 1.4. The results obtained apply, therefore, to the case of Cournot competition with payoff for firm i given by $\pi_i = (\theta - \frac{1}{2}\lambda x_i - \beta x)x_i$, and to the case of Bertrand competition with payoff $\pi_i = (\theta - \frac{1}{2}\lambda p_i - \beta p)p_i$. An interesting difference is that with price competition and differentiated products, ξ_i and ξ are strategic complements in the payoff of firm i. The reason is that, as other firms are now better informed, the intercept of the residual demand of firm i, $\theta - \beta\tilde{p}(\theta)$, is more variable because $\beta < 0$ and the average price $\tilde{p}(\theta)$ is more responsive to θ.[20]

Welfare. The first-best outcome with costly acquisition of information is just the full-information first-best. This can best be understood by considering finite-economy approximations to the large market. It is verified in section 2.3 that, as finite economies grow large, in the sense that the numbers of consumers and firms grow as the market is replicated, the optimal expenditure on information converges to zero in per capita terms and the precision of the aggregate signal tends to ∞. That is, when a centralized planner determines the purchase of an aggregate signal and sets output to maximize expected total surplus, then as the market grows large less and less information is purchased per firm but at the same time the aggregate precision grows unboundedly.

From section 1.2.4.1 we know that in the continuum economy the first-best per capita expected total surplus with no cost of information acquisition is given by $\mathrm{ETS}^0 = (\frac{1}{2}b)E[\theta^2]$, where $b = 1/(\lambda + \beta)$. With decreasing returns to scale, a competitive market always falls short of this first-best level unless the cost of information is zero. If the competitive market spends a positive per capita amount on information, it cannot attain first-best efficiency, which involves zero average expenditure on information acquisition. If the market does not buy any information, then the precision of signals is zero and again an inefficient outcome obtains. The welfare loss is given by $\frac{1}{2}(b - a)\sigma_\theta^2 + c\tau_\varepsilon^*$, where $a = \xi^*/(\lambda + \beta\xi^*)$. This is immediate from section 1.2.4.1 adding the cost of information. It is always positive unless $c = 0$, in which case $b = a$.

Nevertheless, with decreasing returns the market is team-efficient: expected total surplus is maximized given that firms can base their decisions only on their private information. We know from proposition 1.2 that, contingent on ξ, the market works like a team which maximizes expected gross surplus (gross of

[20] See exercise 8.15 in Vives (1999) for the duopoly case.

the cost of information) and attains $\text{EGS}(\xi) = \frac{1}{2}(a\sigma_\theta^2 + b\bar{\theta}^2)$, where $a = \xi/(\lambda + \beta\xi)$. The team maximizes $\text{EGS}(\xi) - \varphi(\xi)$. The FOC of the optimization problem, which is also sufficient for a maximum to obtain, is exactly as in the market solution. The reason is that at the first stage there is no private information and with negligible single players each firm has the right (second-best) incentives to purchase information.

Proposition 1.4. *With endogenous information acquisition a competitive market with increasing marginal costs works like a team which chooses expenditures on information acquisition and decentralized production rules to maximize total expected surplus.*

As a corollary to the proposition we obtain that the welfare loss with respect to the first-best increases with the cost of information. This is easily seen since with increasing marginal costs, the competitive market acts as if it were solving the team program

$$\operatorname*{Max}_{\tau_\varepsilon} \phi(\tau_\varepsilon, c) = \text{EGS}(\tau_\varepsilon) - c\tau_\varepsilon.$$

Therefore, using the envelope condition,

$$\frac{\mathrm{d}\phi(\tau_\varepsilon^*, c)}{\mathrm{d}c} = \frac{\partial\phi}{\partial c}(\tau_\varepsilon^*, c) = -\tau_\varepsilon^*,$$

and the net expected total surplus of the market decreases with the cost of information (as long as $\tau_\varepsilon^* > 0$). The following proposition summarizes the market performance with respect to the first-best with increasing marginal costs.

Proposition 1.5. *With endogenous and costly information acquisition a competitive market with decreasing returns always falls short of first-best efficiency. Furthermore, the welfare loss increases with the cost of information as long as the market expenditure on information acquisition is positive.*

1.6.2 Constant Returns to Scale

In the constant-returns-to-scale case ($\lambda = 0$), no equilibrium exists if information acquisition is costly. The argument is as follows. Notice first that in equilibrium it must be the case that no firm makes any profit at the market stage when prices and quantities are set. With constant returns to scale, equilibrium at the second stage implies that the expected value of the market price conditional on the signal received by any firm is nonpositive. Otherwise the firm would expand indefinitely. Therefore, firms cannot purchase any information in equilibrium since this would imply negative profits. Furthermore, zero expenditures on information acquisition are not consistent with equilibrium either. If no firm purchases information, there are enormous incentives for any firm to get some information and make unbounded profits under constant returns to scale. With no firm acquiring information the average output in the market is nonrandom and any firm with some information could make unbounded profits by shutting down operations if the expected value of the market price conditional on the received signal is nonpositive and producing an unbounded amount otherwise.

The nonexistence of equilibrium argument does not depend on the linear-normal specification and we will encounter it later when we deal with the Grossman–Stiglitz paradox in section 4.2.2. The nonexistence result is also reminiscent of Wilson's analysis of informational economies of scale (Wilson 1974). He finds that the compounding of constant returns in physical production with information acquisition yields unbounded returns in models where the choice of the decision variable applies in the same way to the production of all units of output. Obviously, this implies nonexistence of a competitive equilibrium. The fact that no equilibrium exists with endogenous information acquisition under constant returns is particular to the continuum model. With a finite number of firms an equilibrium exists in the two-stage model and it approaches the first-best efficient outcome when the market grows large. The analysis is deferred to section 2.3.

1.6.3 Summary

In a large Cournot market with uncertain demand where firms acquire information, we can draw the following conclusions.

- With decreasing returns there is a welfare loss with respect to the first-best outcome and it increases with the cost of information. Yet the market is team-efficient and maximizes expected total surplus subject to the use of decentralized production strategies.
- With constant returns there is no welfare loss and the market acquires the right amount of information.
- Information acquisition decisions by firms are strategic substitutes (if competition were to be à la Bertrand in a differentiated product environment, then they would be strategic complements).

1.7 Summary

In this chapter we have examined how simple market mechanisms—such as Cournot or auctions, in which agents do not have the opportunity to condition on market statistics—aggregate information. The market aggregates information if it replicates the outcome of competition when all agents pool their information in the economy. The chapter has introduced information aggregation in standard partial equilibrium models and some of the basic tools for the analysis:

- The idealization of a large market as a continuum of players.
- Price-taking (Bayesian) equilibria in markets with incomplete information.

- The linear-normal model and the computation and characterization of linear equilibria.
- The concept of team efficiency as a welfare benchmark for an incomplete-information economy.

The main results are as follows.

- Except in the independent-values case or in particular circumstances, smooth market mechanisms, such as Cournot with homogeneous product or Bertrand with differentiated products, do not aggregate information as the market grows large.
- Auctions, like voting mechanisms, tend to aggregate information under less restrictive conditions. This is because of their winner-takes-all feature which forces bidders to condition effectively on the information that winning conveys.
- A large homogeneous product Cournot market under private information will not be first-best efficient, except possibly under constant returns to scale. However, it is team-efficient, that is, it maximizes total expected surplus under the constraint of using of decentralized strategies. The property also holds with costly information acquisition.
- The strategies of agents in incomplete-information economies, and the relative weights to private and public information in particular, depend on the degree of strategic complementarity or substitutability of payoffs. This may explain, for example, patterns of under- or overreaction to public information with respect to the team-efficient benchmark.

1.8 Appendix

Proof of proposition 1.3 (see p. 32). We first derive the linear equilibrium in the case of private information, where each player conditions on his signal only. Player i maximizes over y_i:

$$E[\pi(y_i, y; \theta_i) \mid s_i] = E[\theta_i \mid s_i]y_i - \tfrac{1}{2}\omega_1 y_i^2 - \omega_2 E[y \mid s_i]y_i,$$

where y is the random equilibrium average action y. We obtain the following FOC,

$$E[\theta_i \mid s_i] - \omega_1 y_i - \omega_2 E[y \mid s_i] = 0,$$

and, since $\omega_1 > 0$, the best response of player i is

$$y_i = \frac{E[\theta_i \mid s_i] - \omega_2 E[y \mid s_i]}{\omega_1}.$$

We are looking for a linear symmetric equilibrium of the form $Y(s_i) = as_i + \hat{b}$, where the parameters a and $\hat{b}$ need to be determined. It follows that

$$y = \int_0^1 Y(s_j)\,\mathrm{d}j = a\int_0^1 \theta_j\,\mathrm{d}j + a\int_0^1 \varepsilon_j\,\mathrm{d}j + \hat{b} = a\tilde{\theta} + \hat{b}$$

since according to our convention $\int_0^1 \varepsilon_j \,\mathrm{d}j = 0$ (a.s.). We know that $E[\theta_i \mid s_i] = \xi s_i + (1-\xi)\bar{\theta}$ and $E[\tilde{\theta} \mid s_i] = E[\theta_j \mid s_i] = \xi\varsigma s_i + (1-\xi\varsigma)\bar{\theta}$.

It must therefore hold that

$$as_i + \hat{b} = \frac{\xi s_i + (1-\xi)\bar{\theta} - \omega_2(a(\xi\varsigma s_i + (1-\xi\varsigma)\bar{\theta}) + \hat{b})}{\omega_1}.$$

Identifying coefficients we obtain

$$\omega_1 a = \xi - \omega_2 \xi\varsigma a \quad \text{and} \quad \omega_1 \hat{b} = (1-\xi)\bar{\theta} - \omega_2(1-\xi\varsigma)\bar{\theta}a - \omega_2\hat{b}$$

and, since $\omega_1 + \omega_2 > 0$,

$$a = \frac{\xi}{\omega_1 + \omega_2\xi\varsigma} \quad \text{and} \quad \hat{b} = \left(\frac{1}{\omega_1+\omega_2} - a\right)\bar{\theta}.$$

It follows that linear symmetric Bayesian equilibrium strategies in the private-information case are given by

$$Y(s_i) = as_i + \left(\frac{1}{\omega_1+\omega_2} - a\right)\bar{\theta} = a(s_i - \bar{\theta}) + b\bar{\theta},$$

where a and b are given as in the proposition. We have that $E[y(s_i)] = b\bar{\theta}$ since $E[s_i - \bar{\theta}] = 0$.

Next we derive the symmetric linear Bayesian equilibrium strategies in the shared-information case, where each firm conditions both on its signal and the average parameter $\tilde{\theta}$ (since $\tilde{s} = \tilde{\theta}$). The FOC is given by

$$E[\theta_i \mid s_i, \tilde{\theta}] - \omega_1 y_i - \omega_2 E[y \mid s_i, \tilde{\theta}] = 0,$$

which, since $\omega_1 > 0$, implies that

$$y_i = \frac{E[\theta_i \mid s_i, \tilde{\theta}] - \omega_2 E[y \mid s_i, \tilde{\theta}]}{\omega_1}.$$

We are looking for a linear symmetric equilibrium of the form $Z(s_i, \tilde{\theta}) = \hat{a}s_i + \hat{b}\tilde{\theta} + \hat{c}$, where the parameters $\hat{a}$, $\hat{b}$, and $\hat{c}$ need to be determined. With hindsight we let $\hat{c} = 0$. We therefore have that $z = \int_0^1 Z(s_i, \tilde{\theta})\,\mathrm{d}i = \hat{a}\tilde{\theta} + \hat{b}\tilde{\theta} = (\hat{a} + \hat{b})\tilde{\theta}$ since $\tilde{s} = \int_0^1 s_i \,\mathrm{d}i = \tilde{\theta}$.

It must therefore hold that

$$\hat{a}s_i + \hat{b}\tilde{\theta} = \frac{E[\theta_i \mid s_i, \tilde{\theta}] - \omega_2 z}{\omega_1}.$$

From the properties of conditional expectations with normal distributions we have that

$$E[\theta_i \mid s_i, \tilde{\theta}] = (1-d)\tilde{\theta} + ds_i,$$

where $d = (1-\varsigma)/(\xi^{-1} - \varsigma)$. Then

$$\hat{a}s_i + \hat{b}\tilde{\theta} = \frac{(1-d)\tilde{\theta} + ds_i - \omega_2(\hat{a}+\hat{b})\tilde{\theta}}{\omega_1},$$

from which $\hat{a} = d/\omega_1$ and $\omega_1 \hat{b} = (1-d) - \omega_2(\hat{a} + \hat{b})$, and therefore, since $\omega_1 + \omega_2 > 0$,

$$\hat{b} = \frac{(1-d) - \omega_2 \hat{a}}{\omega_1 + \omega_2} = \frac{1}{\omega_1 + \omega_2} - \hat{a} = b - \hat{a}.$$

In conclusion, linear symmetric Bayesian equilibrium strategies in the shared-information case are given by

$$Z(s_i, \tilde{\theta}) = \hat{a}(s_i - \tilde{\theta}) + b\tilde{\theta},$$

where $\hat{a}$ and b are as in the proposition. We also have that $E[Z(s_i, \tilde{\theta})] = b\bar{\theta}$ since $E[s_i - \tilde{\theta}] = 0$. □

1.9 Exercises

1.1 (*expected surplus in the linear-normal model*). Consider the linear model of section 1.2.4.1. Show that expected gross surplus (per capita) at the market solution is given by $E[\int_0^{\tilde{X}(\theta)} P(z;\theta)\,\mathrm{d}z] = a\sigma_\theta^2 + b\bar{\theta}^2 - \frac{1}{2}\beta(a^2\sigma_\theta^2 + b^2\bar{\theta}^2)$ and expected cost by $E[C(X_i(s_i))] = \frac{1}{2}\lambda[a^2\sigma_\theta^2/\xi + b^2\bar{\theta}^2]$. Conclude that ETS $= \frac{1}{2}(a\sigma_\theta^2 + b\bar{\theta}^2)$. Show also that per capita expected consumer surplus

$$E\bigg[\int_0^{\tilde{X}(\theta)} P(z;\theta)\,\mathrm{d}z\bigg] - P(\tilde{X}(\theta),\theta)\tilde{X}(\theta)$$

is given by $\frac{1}{2}\beta E[(\tilde{X}(\theta))^2]$.

Hint. The first part is a straight computation using the equilibrium expressions. For the second you do not need to know the form of $\tilde{X}(\theta)$.

** **1.2** (*information aggregation in a Cournot market with capacity constraints and a finite support information structure*). Consider a market with inverse demand $p = \theta - \beta x$, with firms receiving private signals about the uncertain θ and competing in quantities. Firm i has a capacity of production k_i. This enables the firm to produce at zero cost up to the capacity limit. We will assume here a finite support information structure where θ can take two values, θ_{H} and θ_{L}, $\theta_{\mathrm{H}} \geqslant \theta_{\mathrm{L}} > 0$ with equal prior probability. Let $\mu = \frac{1}{2}(\theta_{\mathrm{H}} + \theta_{\mathrm{L}})$. Firm i may receive a low (s_{L}) or high (s_{H}) signal about θ with a likelihood $\mathrm{P}(s_{\mathrm{H}} \mid \theta_{\mathrm{H}}) = \mathrm{P}(s_{\mathrm{L}} \mid s_{\mathrm{L}}) = \ell$, where $\frac{1}{2} \leqslant \ell \leqslant 1$. If $\ell = \frac{1}{2}$, the signal is uninformative; if $\ell = 1$, it is perfectly informative. Signals received by firms are i.i.d. conditional on θ. With these assumptions:

(i) Check that $E[\theta \mid s_{\mathrm{H}}] = \ell\theta_{\mathrm{H}} + (1-\ell)\theta_{\mathrm{L}}$, $E[\theta \mid s_{\mathrm{L}}] = (1-\ell)\theta_{\mathrm{H}} + \ell\theta_{\mathrm{L}}$, $\mathrm{P}(s_{\mathrm{H},j} \mid s_{\mathrm{H},i}) = \ell^2 + (1-\ell)^2$, and $\mathrm{P}(s_{\mathrm{H},j} \mid s_{\mathrm{L},i}) = 2\ell^2(1-\ell)$ with $j \neq i$.

(ii) Let $\varsigma \equiv \ell^2 + (1-\ell)^2$, $\Delta = \theta_{\mathrm{H}} - \theta_{\mathrm{L}}$, $\ell \in [\frac{1}{2}, 1]$, and assume that $\ell\theta_{\mathrm{L}} > (1-\ell)\theta_{\mathrm{H}}$. Show that when all firms have a common capacity k there is a unique symmetric Bayesian equilibrium in which each firm produces according to $X(s_{\mathrm{H}}) = \bar{y} \equiv \omega\theta_{\mathrm{H}} - (\omega - 1)\theta_{\mathrm{L}}$ and $X(s_{\mathrm{L}}) = \underline{y} \equiv \omega\theta_{\mathrm{L}} - (\omega - 1)\theta_{\mathrm{H}}$ if $k \geqslant \bar{y}$, where $\omega = \ell/(2\ell - 1)$ if $\ell > \frac{1}{2}$ and $\omega = \frac{1}{2}$ if $\ell = \frac{1}{2}$; $X(s_H) = k$ and $X(s_L) = \underline{z} \equiv (E[\theta \mid s] - (1-\varsigma)k)/\varsigma$ if $\bar{y} > k > E[\theta \mid s]$ and $X(s_H) = X(s_L) = k$ otherwise.

(iii) Show that in equilibrium the average output $\tilde{X}(\theta)$ equals θ if $k \geqslant \bar{y}$, $q(\theta)$ if $\bar{y} > k > E[\theta \mid s]$ (where $q(\theta_H) = (k(\varsigma + \ell - 1) + E[\theta \mid s_L](1-\ell))/\varsigma$ and $q(\theta_L) = (\ell E[\theta \mid s_L] - (\ell - \varsigma)k)/\varsigma$), and k otherwise, provided the signals are informative ($\ell > \frac{1}{2}$). When does the market aggregate information?

Solution. Follow similar steps as in section 1.2.4.1. See Vives (1986).

** **1.3** (*investment in flexibility under uncertainty and private information*). Consider the same market as in exercise 1.2 but now each firm has an opportunity in a first stage to invest in capacity. Firm i may purchase a capacity k_i at the cost ck_i ($c > 0$). Assume that $\ell\theta_L > (1-\ell)\theta_H$.

(i) Show that there is a unique symmetric subgame-perfect equilibrium of the two-stage investment-quantity setting game.[21] The equilibrium capacity k^* is given by

$$k^* = \max\left\{0, \bar{\theta} - c, \bar{\theta} - c + \frac{\Delta(\ell - \frac{1}{2}) - c}{2\varsigma - 1}\right\}.$$

(ii) Show that if $c < \bar{c} \equiv \Delta(\ell - \frac{1}{2})$ the equilibrium investment in k^* increases with a mean-preserving spread of demand (i.e., with Δ) and increases or decreases with the precision of the information (i.e., ℓ) according to whether c is larger or smaller than $\tilde{c} \equiv (1 - 4\ell(1-\ell))\Delta/(4(2\ell - 1))$. If $c > \bar{c}$, then k^* is independent of Δ and ℓ. Interpret the results in terms of the effect of uncertainty in investment in flexibility. Distinguish between increases in prior uncertainty Δ and increases in the variability of beliefs ℓ.

(iii) *Welfare.* Show that $k^* = \arg\max\{\text{EGS}(k) - ck\}$, where $\text{EGS}(k)$ is the expected gross total surplus with capacity k.

(iv) Show that there is a welfare loss with respect to the full-information first-best unless the cost of capacity c is zero or high enough ($c \geqslant \frac{1}{2}\Delta$). The welfare loss is decreasing with the precision of information if $\Delta(\ell - \frac{1}{2}) > c > 0$ and independent of the latter if $\frac{1}{2}\Delta > c > \Delta(\ell - \frac{1}{2})$. Interpret the results.

Solution. See Vives (1986).

1.4 (*welfare in the monopolistic competition model with a common value*). The linear demand system for differentiated products of section 1.4.1 can be obtained from the optimizing behavior of a representative consumer who maximizes the quadratic utility

$$U = \int_0^1 \alpha_i x_i \, \mathrm{d}i - \frac{1}{2}\left(\delta x^2 + (1-\delta)\int_0^1 x_i^2 \, \mathrm{d}i\right)$$

[21] A subgame-perfect equilibrium requires that for any investment decision at the first stage, a Nash equilibrium in outputs obtains at the second stage. This rules out incredible threats. Nash equilibria only require optimizing behavior along the equilibrium path.

minus expenditure $\int_0^1 p_i x_i \,\mathrm{d}i$, where $x = \int_0^1 x_i \,\mathrm{d}i$. Consumer surplus (CS) in terms of quantities is given by $\frac{1}{2}(\delta x^2 + (1-\delta)\int_0^1 x_i^2 \,\mathrm{d}i)$ and in terms of prices by

$$\frac{1}{2}\Big((1-\delta)\int_0^1 \beta_i^2 \,\mathrm{d}i + \delta\tilde{\beta}^2 + (1+\gamma)\int_0^1 p_i^2 \,\mathrm{d}i - \gamma p^2 - 2\int_0^1 \beta_i p_i \,\mathrm{d}i\Big),$$

where $\tilde{\beta} = \int_0^1 \beta_i \,\mathrm{d}i$ and, as before, $\beta_i = (\alpha_i - \delta\tilde{\alpha})/(1-\delta)$ and $\gamma = \delta/(1-\delta)$. Note that $\tilde{\beta} = \tilde{\alpha}$.

Since costs are assumed to be zero, U represents total surplus (TS). Consider the common-value case ($\varsigma = 1$) and show that

(i) expected consumer surplus in the Cournot case $\mathrm{ECS_C}$ increases and in the Bertrand case $\mathrm{ECS_B}$ decreases with the precision of information τ_ε;

(ii) expected profits increase with τ_ε in the Bertrand case ($E[\pi_\mathrm{B}]$) and may increase or decrease with τ_ε in the Cournot case ($E[\pi_\mathrm{C}]$); expected profits in both cases increase with prior uncertainty σ_θ^2;

(iii) expected total surplus in the Cournot case $\mathrm{ETS_C}$ increases and in the Bertrand case $\mathrm{ETS_B}$ decreases with the precision of information τ_ε.

Interpret the results in terms of the impact on the variability of individual and aggregate strategies and potential covariance of the uncertain parameters with strategies. What would be the impact on welfare if firms were to share information?

Solution. Let $\xi = \tau_\varepsilon/(\tau_\theta + \tau_\varepsilon)$. From the equilibrium strategies in proposition 1.3 and the expressions for CS obtain that

$$\frac{\partial\,\mathrm{ECS_C}}{\partial\xi} = \frac{2 - 4\delta + 2\delta^2 + 3\delta(1-\delta)\xi}{(2 - 2\delta + \delta\xi)^3}\frac{\sigma_\theta^2}{2} > 0$$

and

$$\frac{\partial\,\mathrm{ECS_B}}{\partial\xi} = -\frac{(12-\xi)\gamma + (6+\xi)\gamma^2 + 6}{(2 + 2\gamma - \gamma\xi)^3}\frac{\sigma_\theta^2}{2} < 0.$$

Note that in equilibrium $E[\pi_\mathrm{C}] = (1-\delta)E[x_i^2]$ and $E[\pi_\mathrm{B}] = (1+\gamma)E[p_i^2]$ and derive the results for profits from the equilibrium strategies in proposition 1.3. Furthermore, obtain that

$$\frac{\partial\,\mathrm{ETS_C}}{\partial\xi} = \frac{6 - 12\delta + 6\delta^2 + \delta(1-\delta)\xi}{(2 - 2\delta + \delta\xi)^3}\frac{\sigma_\theta^2}{2} > 0$$

and

$$\frac{\partial\,\mathrm{ETS_B}}{\partial\xi} = \frac{(4-3\xi)\gamma + (2-\xi)\gamma^2 + 2}{(2 - 2\gamma - \gamma\xi)^3}\frac{\sigma_\theta^2}{2} < 0.$$

For the effect of information sharing note that it is equivalent to letting $\xi = 1$. (See Vives (1990) for more details.)

References

Angeletos, G. M., and A. Pavan. 2007a. Efficient use of information and social value of information. *Econometrica* 75:1103–42.

——. 2007b. Policy with dispersed information. Mimeo.

Aumann, R. 1964. Markets with a continuum of traders. *Econometrica* 32:39–50.

Bryant, J. 1983. A simple rational expectations Keynes-type model. *Quarterly Journal of Economics* 98:525–28.

Calvó-Armengol, A., and J. de Martí. 2007. On optimal communication networks. Mimeo.

Diamond, P. 1982. Aggregate demand management in search equilibrium. *Journal of Political Economy* 90:881–94.

Dybvig, P. H., and C. S. Spatt. 1983. Adoption externalities as public goods. *Journal of Public Economics* 20:231–47.

Feddersen, T., and W. Pesendorfer. 1996. The swing voter's curse. *American Economic Review* 86:408–24.

——. 1997. Voting behavior and information aggregation in elections with private information. *Econometrica* 65:1029–58.

Gale, D. 1996. Equilibria and Pareto optima of markets with adverse selection. *Economic Theory* 7:207–35.

Guesnerie, R. 1992. An exploration of the eductive justifications of the rational-expectations hypothesis. *American Economic Review* 82:1254–78.

Harris, M., and R. Townsend. 1981. Resource allocation under asymmetric information. *Econometrica* 49:33–64.

Hayek, F. 1945. The use of knowledge in society. *American Economic Review* 35:519–30.

Heinemann, M. 2004. Are rational expectations equilibria with private information eductively stable? *Journal of Economics* 82:169–94.

Hellwig, C. 2005. Heterogeneous information and the welfare effects of public information disclosures. UCLA mimeo.

Hellwig, M. F. 1987. Some recent developments in the theory of competition in markets with adverse selection. *European Economic Review* 31:319–25.

Kremer, I. 2002. Information aggregation in common value auctions. *Econometrica* 70:1675–82.

Li, L. 1985. Cournot oligopoly with information sharing. *RAND Journal of Economics* 16:521–36.

Li, L., R. McKelvey, and T. Page. 1987. Optimal research for Cournot oligopolists. *Journal of Economic Theory* 42:140–66.

Matsuyama, K. 1995. Complementarities and cumulative processes in models of monopolistic competition. *Journal of Economic Literature* 33:701–29.

Milgrom, P. R. 1979. A convergence theorem for competitive bidding and differential information. *Econometrica* 47:679–88.

——. 1981. Rational expectations, information acquisition, and competitive bidding. *Econometrica* 49:921–43.

Milgrom, P. R., and R. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* 50:1089–122.

Morris, S., and H. Shin. 2002. The social value of public information. *American Economic Review* 92:1521–34.

——. 2005. Central bank transparency and the signal value of prices. *Brookings Papers on Economic Activity* 2:1–66.

Novshek, W. 1980. Cournot equilibrium with free entry. *Review of Economic Studies* 47:473–86.

Palfrey, T. 1985. Uncertainty resolution private information aggregation, and the Cournot competitive limit. *Review of Economic Studies* 52:69–83.

Pesendorfer, W., and J. Swinkels. 1997. The loser's curse and information aggregation in common value auctions. *Econometrica* 65:1247–81.

Prescott, E., and R. Townsend. 1984. Pareto optima and competitive equilibria with adverse selection and moral hazard. *Econometrica* 52:21–45.

Radner, R. 1962. Team decision problems. *Annals of Mathematical Statistics* 33:857–81.

Royden, H. L. 1968. *Real Analysis*. New York: Macmillan.

Svensson, L. 2006. Social value of public information. Comment: Morris and Shin (2002), Is actually pro-transparency, not con. *American Economic Review* 96:448–52.

Vives, X. 1986. Investment in flexibility in a competitive market with incomplete information. CARESS Working Paper 86-08, University of Pennsylvania.

——. 1988. Aggregation of information in large Cournot markets. *Econometrica* 56:851–76.

——. 1990. Trade associations, disclosure rules, incentives to share information and welfare. *RAND Journal of Economics* 21:409–30.

——. 1999. *Oligopoly Pricing: Old Ideas and New Tools*. Boston, MA: MIT Press.

Wilson, R. 1975. Informational economies of scale. *Bell Journal of Economics* 6:184–95.

——. 1977. A bidding model of perfect competition. *Review of Economic Studies* 44:511–18.

Woodford, M. 2002. Imperfect common knowledge and the effects of monetary policy. In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps* (ed. P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford). Princeton University Press.

——. 2005. Central bank communication and policy effectiveness. *Proceedings*, Federal Reserve Bank of Kansas City, August, pp. 399–474.