CHAPTER 11

Auctions

1. *Shading in first-price auctions.* A first-price auction is a selling mechanism in which each potential buyer places a bid, the highest bidder wins the object for sale and pays his bid. A key aspect of bidding in first-price auctions is in deciding on appropriate shading. The tradeoff is a simple one. If the object is worth a value v to you, bidding v results in no profit. Shading your bid will result in a profit if you win, but your chance of winning likely decreases. So, there is a tradeoff between a higher profit conditional on winning, and a lower chance of winning.

To model this tradeoff, denote by b the agent's bid, and by p the maximum bid among other bidders. The agent cannot know p for sure. For now, let us define f as the distribution over the possible maximum bid made by others, and by

$$\phi(\lambda) \equiv \Pr\{v - p > \lambda\} \equiv \int_{p<v-\lambda} f(p)dp$$

the probability that our bidder wins when shading his value v by a fixed amount λ. His expected gain is:

$$\lambda\phi(\lambda).$$

The function ϕ is a decreasing function. The tradeoff is that higher λ increases profit in the event of winning, but lowers the chance of winning. The derivation of the optimal shading is analogous to a standard *monopoly pricing* problem in which $\phi(\lambda)$ is interpreted as a demand function.

2. *Equilibrium shading.* A key aspect of an analysis of auctions is that it describes an equilibrium phenomenon in which each potential buyer solves the tradeoff above, and in which the price p, the distribution f and the function ϕ defined above are determined endogenously. The problem is no longer a

monopoly pricing problem, but an *oligopoly pricing* problem, as the demand ϕ that each bidder faces depends on the behavior of others.[1]

Intuitively, if others shade a lot, one should have little competitive pressure and find shading a relatively safe strategy. On the other hand, if others are aggressive competitors who shade little, substantial shading might lead to very little chance of winning, in which case a more aggressive strategy might be preferable. Equilibrium analysis aims to endogenize the level of shading that each bidder finds optimal, given the behavior of others. Equilibrium shading reflects the degree to which the environment is competitive. It also determines how surplus is shared between the seller and buyers: smaller shading generates higher revenue for the seller and lower expected rents for buyers.

What follows is a simple model that captures the forces involved in determining equilibrium shading.

Define $(v_1, \ldots, v_n)$ as a vector of valuations. Uncertainty is modeled by assuming that the vector is drawn from a (symmetric) distribution.

Bidding is assumed to be characterized for each bidder i by a *uniform* shading level λ_i. That is, bidder i bids

$$b_i = v_i - \lambda_i$$

whenever his value is v_i.[2] Given this behavioral assumption, we can associate, to each vector of shading levels $\lambda = (\lambda_1, \ldots \lambda_n)$ and each bidder i, an expected gain $G_i(\lambda)$:

$$G_i(\lambda) = \lambda_i \Pr(v_i - \lambda_i > \max_{j \neq i} v_j - \lambda_j).$$

Bidders are symmetric, and we look for a symmetric equilibrium. That is, we look for a shading level λ^* such that when all bidders shade by the same amount λ^*, no one finds it profitable to modify his shading λ^*. This implies

$$\lambda^* = \arg\max \lambda_i \phi(\lambda_i - \lambda^*)$$

where

$$\phi(x) = \Pr(v_i - \max_{j \neq i} v_j > x).$$

[1] See Caplin and Nalebuff (1986) for an analysis of price competition in an oligopolistic environment.

[2] We emphasize that the shading level λ_i is applied uniformly to any value v_i in the range of values considered (given the type of object considered). Clearly, acquiring a firm or a used car are different matters, and would likely lead to different shading levels in practice. If we were to look for a shading rule that applies universally, a multiplicative shading (as the one assumed in Chapter 7) would be more plausible. We discuss this further at the end of this chapter.

Looking at first-order conditions, we can characterize the only possible candidate for a symmetric equilibrium λ^*:

$$\lambda^* = \frac{\phi(0)}{|\,\phi'(0)\,|} = \frac{1}{n\,|\,\phi'(0)\,|}$$

3. *The effect of dispersion.* Intuitively, equilibrium shading depends on the proportional change in the winning probability for a bidder when he increases his bid. This proportional change is high (hence shading is small) when there are many bidders (because in this case each has a small chance $\phi(0) = 1/n$ of winning), or when values are not very dispersed (i.e., $|\,\phi'(0)\,|$ is large). This is a standard Bertrand competition effect. When valuations are more dispersed or bidders less numerous, the effect is weaker, and shading is larger.

The simple expression that we obtain (characterizing equilibrium shading) relies on the particular restriction on strategies that we assumed. Had we opted for a multiplicative formulation, the mathematics would have been more complex, and dispersion would not be captured by the slope $|\,\phi'(0)\,|$ alone.

Nevertheless, the qualitative statement would be analogous. More dispersion in valuations reduces competition. Both models establish a similar link between value dispersion and competitive forces.

To formalize the effect of dispersion, assume the following structure:

$$v_i = \alpha + d\,\theta_i$$

with all variables drawn from independent distributions. α reflects characteristics of the object affecting the preference of all participants, θ_i reflects a private characteristic, and d is a dispersion parameter. Also define

$$\phi_0(x) \equiv \Pr(\theta_i - \max_{j \neq i} \theta_j > x) \text{ and } \lambda_0^* = \frac{\phi_0(0)}{|\,\phi_0'(0)\,|}.$$

For a fixed d, the condition $v_i - \max_{j \neq i} v_j > x$ is equivalent to

$$\theta_i - \max_{j \neq i} \theta_j > x/d,$$

and we obtain $\phi(x) = \phi_0(x/d)$; hence

$$\lambda^* = d\,\lambda_0^*.$$

In words, the smaller the dispersion parameter, the stronger the competitive forces, and if d is a random variable, we obtain:

$$\lambda^* = \frac{1}{E[1/d]}\lambda_0^*.$$

4. *Objections to standard modeling.* The model departs from standard modeling of auctions in that the strategy is characterized by a single parameter λ, while in standard modeling, one would model agents who optimally adjust bidding, *for each separate realization* v_i.

A possible motivation for the standard model is that, given the distribution from which values are assumed to be drawn, values convey information about the chance of being the highest valuation bidder, and about the dispersion of values below one's own value in the event one wins. Agents should have an opportunity to use that information if it is available.

To illustrate, assume that bidder values are independently drawn on an interval, say $[100, 120]$. For a bidder with value 120, all other bidders have valuations dispersed below his on the whole interval $[100, 120]$. For a bidder with value 101, if he has the highest value, it must be that all others are concentrated between 100 and 101. Competition must thus be much fiercer at the bottom of the interval. The consequence is small shading at the bottom of the interval, higher shading at the top.

An obvious objection is that, in real auctions, one doesn't observe the distributions from which values are drawn. In many situations, it seems difficult to judge, based solely on one's valuation, whether the valuation is high or low compared to others'. When attending an auction for a painting that I like, I might have a personal estimate of $1,000; or possibly $1,100. Whichever figure comes to mind, the specific number seems second order in assessing the chances of being the highest value bidder.

It could of course be that one gets more direct information about one's own chances. If the painting is of my mother or if I am a collector of the artist, I may be quite sure that I'm likely to have the highest value. But, value itself often seems to be a poor instrument. The logic of our simplified model is to disregard this poor instrument altogether. The parameter λ captures a systematic way in which shading occurs (not optimally tailored to each value realization given the particular distribution assumed).

5. *A strategic issue: significant shading.* In bidding in a first-price auction, an agent faces a tradeoff between two types of strategies:

(i) Trying to beat competition by being aggressive, which gives a reasonable chance of winning but a small gain in the event of winning.
(ii) Betting on having a high value compared to others, and shading significantly, which possibly reduces the chance of winning, but generates large gains in the event of winning.

What the equilibrium behavior describes is a smooth resolution of this conflict, in which all bidders find it optimal to settle on the *same* shading level λ^*. An attractive consequence of that smooth resolution is efficiency: the bidder with the highest valuation wins the object.

There are conditions, however, under which no such smooth resolution is possible. When valuations are most likely concentrated, the equilibrium shading level λ^* must be small. Competition is fierce and leaves little rent to bidders. Imagine now that there is a small chance that valuations are dispersed. A bidder may be better off gambling on this being the case and shading significantly: he will only win if in addition he turns out to have a valuation substantially higher than others', but when he wins he gets a substantial gain.

Formally, profits under our tentative equilibrium are equal to $\lambda^*\phi(0)$. This profit must be compared to the gain from shading by a larger amount, say Δ. If the following condition (E) holds:

$$\max_{\Delta} \Delta\phi(\Delta - \lambda^*) > \lambda^*\phi(0), \tag{E}$$

then at least one player prefers to deviate.

Dispersion uncertainty. To better understand circumstances under which (E) might hold, assume that the parameter d takes two values, $\underline{d}$ with probability $1 - p$, and value $\overline{d} > \underline{d}$ with probability p. Then, under minor conditions on ϕ_0, it is sufficient that:

$$\frac{p(1-p)}{4} > \underline{d}/\overline{d}$$

for condition (E) to hold.[3] This means that if p is small, but large compared to relative dispersion $\underline{d}/\overline{d}$, at least one bidder prefers to engage in significant shading, gambling that his valuation is quite high relative to others'.

Intuitively, for any fixed $p > 0$, it is sufficient that $\underline{d}$ is small to drive λ^*, and hence profits, to 0. The reason is that a small change in shading then has a strong effect on the winning probability. And as $\overline{d}$ increases, the gains from the large shading strategy increase.

6. *Revenue comparisons.* We now compare revenues generated by first-price and second-price auctions. Bidding in a first-price auction is not as obvious as it may seem. For the winner, the best strategy is to bid just above the second highest bid. But the winner cannot anticipate what this second highest bid will be. In contrast, in the second-price auction, the winner precisely pays the second highest bid, so he need not adjust his bid to the second highest.

The consequence is that, in the first-price auction, optimal shading depends on a tradeoff between various circumstances; sometimes, it would have been optimal to shade little; and, sometimes, it would have been optimal to shade significantly. Depending on the relative weight of these low- and high-shading

[3] A sufficient condition on ϕ_0 is that for all positive Δ, $|\phi_0'(\Delta)| \leq |\phi_0'(0)|$. Since $E[1/d] \geq (1-p)/\underline{d}$, the tentative equilibrium shading satisfies $\lambda^* \leq \lambda_0^*\underline{d}/(1-p)$.

When the bidder opts for a large shading Δ, he wins with probability at least $p\phi_0(\Delta/\overline{d})$. His expected gain is thus at least equal to $\max_\Delta \Delta p\phi_0(\Delta/\overline{d}) = p\overline{d}\max_\Delta \Delta\phi_0(\Delta)$. Under the condition above on ϕ_0, $\Delta\phi_0(\Delta) \geq \Delta(\phi_0(0) + \Delta\phi_0'(0)) = \phi_0(0)\Delta(1 - \Delta/\lambda_0^*)$, implying that $\max_\Delta \Delta\phi_0(\Delta) \geq \phi_0(0)\lambda_0^*/4$. The optimal deviation thus gives an expected payoff at least equal to $\phi_0(0)p\overline{d}\lambda_0^*/4$. When the latter payoff exceeds $\phi_0(0)\lambda_0^*\underline{d}/(1-p)$, (E) holds.

situations, bidders may end up finding that "somewhat" small shading is optimal, or that "somewhat high" shading is optimal. In the former case, we can expect the first-price auction to generate more revenue, and, in the latter case, we can expect the opposite.

To illustrate, let us return to our low/high dispersion example. If there is a significant chance that values are concentrated (small d), this may be enough to generate fierce competition. In this case, competition is tough even in events where dispersion is large (high d), that is, even in events where rents could potentially be quite high for the buyers. Rents remain small for buyers because they cannot tailor shading to dispersion: shading remains small whether dispersion is small or large, and this is a source of increased revenues for the seller.[4]

Formally, in a second-price auction, the winner, say player i, gets $y = \theta_i - \max_{j\neq i}\theta_j$ in events where y is non-negative. Since y is distributed according to the density $-\phi'(y)$ (by definition of ϕ), a bidder's expected gain, which we denote G^{II}, is therefore:

$$G^{II} = \int_{y\geq 0} -y\phi'(y)dy = \int_{y\geq 0} \phi(y)dy,$$

which can be compared to a bidder's expected gain G^I in the first-price auction: $G^I = \lambda^*\phi(0)$. The following figure illustrates graphically the gains G^I and G^{II} in the case of two bidders ($\phi(0) = 1/2$).

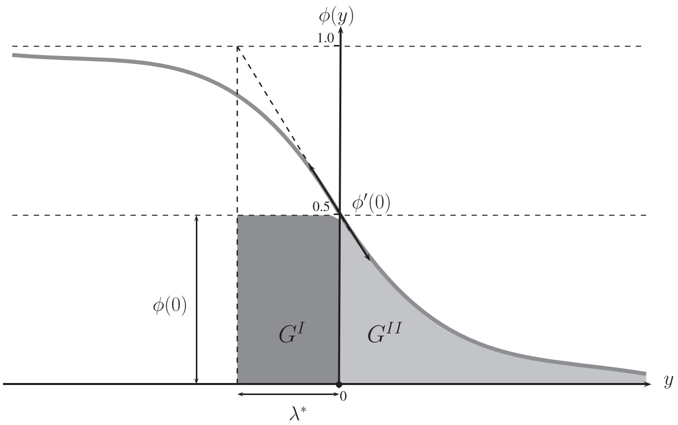

[4] The opposite argument works, of course, when the chance of a large dispersion drives shading to high levels. There are events where dispersion turns out to be small and yet buyers get large rents.

Call ρ the ratio between G^{II} and G^{I}:

$$\rho = G^{II}/G^{I},$$

and ρ_0 this ratio when the dispersion parameter d is equal to 1.[5] Since the allocation does not change across formats, the seller's revenue is highest when the bidder's expected gain is smallest. So *the first-price auction generates more revenue than the second-price if and only if* $\rho > 1$, and we have:

$$\rho > 1 \Leftrightarrow \int_{y\geq 0} \frac{\phi(y)}{\phi(0)} dy > \frac{\phi(0)}{\mid \phi'(0) \mid}. \tag{P}$$

Our main observation is that *dispersion uncertainty* makes it easier to satisfy inequality (P). Indeed, we have:[6]

$$\rho = \kappa\rho_0 \text{ where } \kappa = E(1/d) * Ed,$$

and the coefficient κ is equal to 1 if d is certain, larger than 1 otherwise. So for a fixed distribution ϕ_0, dispersion uncertainty increases ρ.

Said differently, the first-price auction is better for the seller when the "demand" function *combines* high concentration and substantial dispersion: concentration (i.e., high $\mid \phi'(0) \mid$) implies a strong Bertrand competition effect, hence low gains for buyers in the first-price auction, while substantial dispersion implies high rents in the second-price auction. Dispersion uncertainty makes it easier to satisfy these two conditions.[7]

7. *Misperceptions and misadjustments.* In an ascending-price auction, the best strategy is obvious – dropping from the auction when the price reaches one's value. Bidding optimally in a first-price auction is less obvious as it depends on the degree to which the environment is competitive, on the number of bidders, and on the dispersion of values.

Said differently, we have characterized equilibrium shading λ^*, and in so doing, we have proposed a model that links properties of the distribution over values to behavior. Distributions, however, are not meant to be observable by participants. The model is a shortcut that leaves unmodeled the process by which people conform to equilibrium behavior: we can only hope that agents eventually figure out that behaving in this way is optimal, and, for example,

5 $\rho_0 = \int_{y\geq 0} \frac{\phi_0(y)}{\phi_0(0)} dy/\lambda_0^*$.

6 This is because for any given d, $\phi(y) = \phi_0(y/d)$.

7 Note that dispersion uncertainty may also generate incentives for large shading (strategic issue). It can be checked, however, that (P) may hold without generating such incentives. The technical reason is that $\int_{y\geq 0} \phi(y)dy > \max_y y\phi(y)$.

that they will shade less when there are more bidders or when values are less dispersed.

So, although the model does not explain or describe the agents' thought process, it suggests features of the environment that agents should consider when deciding how to bid, such as the number of bidders or the dispersion of values (i.e., $\phi'(0)$).

Standard models thus portray agents as *adjusting perfectly and instantaneously* to these features of the environment. This is a useful modeling simplification; but in a world in which one cannot presume such an immediate and perfect adjustment, strategies are bound to diverge from the "correct" equilibrium outcome. Sources of discrepancies are numerous, as an agent may possibly rely on poorly informative or irrelevant signals about the environment without realizing their strength or lack of relevance.

Our aim, in what follows, is to describe the consequences of these misperceptions or misadjustments.

8. *Noisy shading and misperceptions.* A simple way to model misperceptions/misadjustments is to assume noisy shading, with agents having only imperfect control over shading. A mechanical consequence of noise is that we are bound to lose on efficiency, hence on the surplus to be shared. Another consequence, which we analyze below, is that noise modifies incentives to shade, possibly altering the way that surplus is shared between the seller and the buyers.

Formally, we model misperceptions by assuming bid strategies that take the following form:

$$b_i = v_i - \lambda_i + \varepsilon_i. \tag{11.1}$$

The parameter ε_i is meant to be a (small)[8] noise parameter, which we take to be centered on 0, and that prevents each bidder i from perfectly adjusting shading to the underlying distribution over values. We keep, however, the assumption that bidder i controls λ_i and tries to find the optimal *target shading*. This ensures that despite his errors, his behavior is driven by relevant welfare comparisons, in expectation over the mistakes that he might make.

[8] We keep ε_i small to ensure that shading remains positive despite the noise. There are many ways to introduce errors, and other ways to correct for errors that we could imagine. For example, we could assume: $b_i = v_i - \lambda_i/(1+\eta_i)$ where η_i is a positive random variable. This way of modeling noise and strategies would introduce a more direct motive for shading less (and possibly other comparative statics with respect to noise), as smaller shading is a way to reduce mistakes. Our objective however is not to be exhaustive about possible strategic effects, but to illustrate two different forces that may affect bidding.

An alternative interpretation of the model is that the agent has access only to a noisy estimate of value, denoted z_i:

$$z_i = v_i + \varepsilon_i$$

and that bidding takes the form assumed earlier:

$$b_i = z_i - \lambda_i.$$

Of course, we are not suggesting that misadjustments to the environment and misperceptions of one's own value are always equivalent. We just point out that (11.1) encompasses both interpretations.

9. *Strategic consequences.* Misperceptions can have two effects: they may increase the dispersion of bids, and may generate a selection bias (as a bidder with a positive noise term ε_i has a greater chance of being selected). Higher dispersion weakens competition. The selection bias is potentially favorable to sellers, though bidders may try to compensate for it by increasing shading.

Given our behavioral assumption, we can compute the gain $G_i(\lambda_i, \lambda)$ that bidder i obtains when he shades by λ_i and others shade by λ. To do this, we define

$$y_i = v_i + \varepsilon_i - \max_{j \neq i}(v_j + \varepsilon_j)$$

as the margin by which player i wins when all use the same bid strategy. Next, define

$$\phi_\varepsilon(y) \equiv \Pr(y_i > y); \quad \psi(y) \equiv E[\varepsilon_i \mid y_i > y] \quad \text{and} \quad \Psi(y) = E[\varepsilon_i \mid y_i = y].$$

$\psi(y)$ (and $\Psi(y)$) are thus the expected error that bidder i makes when he wins by a margin at least equal to y (exactly equal to y). Note that $\psi(y)$ is positive and increasing because bidder i tends to win more often when ε_i is positive. This is a selection bias, and the selection bias is stronger when you are winning by a larger margin.

If other bidders shade by λ, we have:

$$G_i(\lambda_i) = (\lambda_i - \psi(\lambda_i - \lambda))\phi_\varepsilon(\lambda_i - \lambda).$$

A symmetric equilibrium shading λ_ε^* must thus satisfy:

$$\lambda_\varepsilon^* + \psi(0) = \frac{\phi_\varepsilon(0)}{-\phi'_\varepsilon(0)}(1 - \psi'(0)),$$

or equivalently:[9]

$$\lambda^*_\varepsilon = \frac{\phi_\varepsilon(0)}{-\phi'_\varepsilon(0)} + \Psi(0).$$

This formula captures the two strategic consequences of noise on incentives. Noise creates greater dispersion in bids, which typically increases shading (because $-\phi'_\varepsilon(0)$ typically decreases). It also generates a selection bias of size $\psi(0)$, which bidders only partially offset by shading an additional $\Psi(0)$.[10] For bidders, the consequence for profits is:

$$G_i(\lambda^*_\varepsilon) = \phi_\varepsilon(0)\left[\frac{\phi_\varepsilon(0)}{-\phi'_\varepsilon(0)} + \Psi(0) - \psi(0)\right],$$

where we see that the partial offsetting of the selection bias may reduce profits ($\Psi(0) < \psi(0)$); while the dispersion of bids generally makes the environment less competitive (smaller $-\phi'_\varepsilon(0)$), hence conducive to higher profit.

For normal distributions and two bidders, total welfare and total profits can be characterized in closed form. Assume v_i and ε_i are normal distributions with mean and variance respectively equal to (v_0, η^2) and $(0, \sigma^2)$. Let $\rho = \sigma^2/\eta^2$. We provide below the exact formula for total welfare (W) and total profits (G):[11]

$$W = v_0 + \frac{\eta}{\sqrt{\pi(1+\rho)}} \text{ and } G = \frac{\eta}{\sqrt{\pi}}\left(\frac{1}{\sqrt{1+\rho}} + \sqrt{1+\rho}(\pi - \sqrt{2}\right).$$

Thus, as ρ increases above 0, welfare decreases (because the object is no longer necessarily allocated to the highest value bidder), but total profits for buyers nevertheless increase. The reason for this is that, due to noise, competition is weaker (smaller $-\phi'_\varepsilon(0)$), and this effect dominates the fact that the selection bias is only partially offset ($\Psi(0) < \psi(0)$). The immediate consequence is that the seller is worse off.

To summarize, a bidder is more likely to win when optimistic about his own valuation. This is a selection bias, which creates a gap between the

9 This follows because $\Psi = \frac{(\psi\phi)'}{\phi'} = \psi + \psi' \frac{\phi}{\phi'}$.

10 Intuitively, the reason that bidders only partially offset the selection bias is that at the optimal shading, bidders implicitly consider what happens when they win by a zero margin; while, on average, they win by a larger-than-zero margin. They, thus, face a selection bias of magnitude $\psi(0)$ (larger than $\Psi(0)$).

11 To simplify notation, denote by $h = -\phi'_\varepsilon$ the density over y_i. Note that $y_i \sim \mathcal{N}(0, \sigma_y^2)$ with $\sigma_y = \sqrt{2}\sqrt{\sigma^2 + \eta^2}$. Define $M = E[z_i - v_0 \mid z_i > z_j]$. We observe that $W = v_0 + M - \psi(0)$ and that by symmetry,

$$M = \frac{1}{2}E[z_i - z_j \mid z_i - z_j > 0] = \int_{y_i > 0} y_i h(y_i) dy_i = E\max(0, y_i) = \sigma_y/\sqrt{2\pi}.$$

Also observe that $\Psi(y) = E[\varepsilon_i \mid y_i = y] = \alpha y$, with $\alpha = \frac{\sigma^2}{2\sigma^2 + 2\eta^2}$, hence $\psi(0) = 2\alpha \int_{y_i > 0} y_i h(y_i) dy_i$. To obtain G, we observe that $G = W - S$ where S is the seller's gain. To compute S, we observe that $S = M - \gamma^*_\varepsilon$ and, since $y_i \sim \mathcal{N}(0, \sigma_y^2)$, $\lambda^*_\varepsilon = \sqrt{\pi}\sigma_y$.

estimated value and the realized value conditional on winning. This gap is called the winner's curse, and profit-maximizing bidders ought to bid in a way that takes into account this selection bias. The analysis above highlights that in equilibrium, the offsetting of the selection bias is incomplete, which is detrimental to bidders. But, it also highlights that the dispersion in bids that misperceptions create may sufficiently weaken competition that bidders may nevertheless benefit from the noisier environment. It is not increased information that is a source of rents, but increased dispersion.

Further Comments

Independence and correlation. *A standard division in studying auctions is whether values are* independent or correlated. *Technically the distinction means that in the former case, values should be thought of as being drawn from independent distributions, while in the latter case, one should think of a random parameter affecting all values.*

The distinction seems awkward. The object for sale is the same for every potential buyer, it has characteristics that each one can observe. These characteristics may not be valued in the same way, but to any potential buyer, the value of a car surely differs substantially from the value he attaches to a plant. This justified our representation

$$v_i = \alpha + d\,\theta_i,$$

where α is a shift parameter that captures the "typical" value of the object for sale.

So there should be no question that values are correlated across auctions. The idea that the analyst can nevertheless focus on the case where values are independent stems from the fact that, if bidders know α, d and the structure of the model, they can infer the value of the independently drawn characteristic $\theta_i \equiv (v_i - \alpha)/d$, and use it as an instrument to derive optimal/equilibrium bidding $b^*(v_i) = v_i - \lambda^*(\theta_i)$. From this perspective, the independent value is the more cognitively demanding, as it requires knowing α and d, and determining how to condition bids on three variables (v_i, α and d).

Nevertheless, whether bidders condition behavior on v_i or on (v_i, α, d), both formulations are quite demanding cognitively. In both formulations, bidders may exploit the structure of the model and, for example, adjust their shading strategy using information on the difference $D = v_i - \max_{j \neq i} v_j$ conveyed by v_i (or v_i, α and d).

What we have proposed is a drastic simplification in which v_i is used in bidding, but where the information about D potentially provided by v_i is not. This provides a more parsimonious auction model. It also opens the path to intermediate models in which bidders get "some" information about D, in the form of coarse signals about rank that they could condition behavior on, *without necessarily embodying the level of sophistication assumed in standard models.*[12]

Dispersion rents or information rents? *The literature often refers to* "information rents" *to describe the gains that an "informed" player gets. In auctions, a more appropriate terminology might be* "dispersion rents": *bid dispersion generates rents, and private information in standard models generates rents insofar as it creates bid dispersion. Our model illustrates that poor information (i.e., noisier estimates) may translate into higher bid dispersion, and that improving an agent's information (i.e., less noisy estimates) may induce less bid dispersion, hence smaller rents.*[13]

This effect of noise on bid dispersion would not hold in a standard model – the opposite would actually be true. Noisier estimates would translate into less dispersed posteriors (by a regression to the mean effect), and therefore greater competition when symmetry is assumed. The latter conclusion, however, is (in our view) an artifact of the standard model, and of the implicit assumption that agents know (or behave as if they knew) all distributions: as noise increases, value estimates decrease in importance and more weight is put on priors,[14] *as explained in Chapter 10.*

Common values, interdependence and estimation errors. *In modeling auctions, the distinction between private and common values is often seen as a key dividing line. In common or interdependent value auctions, the bids of others' reveal information about one's own valuation, and, in bidding, a rational bidder ought to take into account those inferences. An omniscient bidder will indeed find this advice useful. To most bidders, however, the precise ways in which preferences are interdependent are likely obscure, and the appropriate inference likely out of reach.*

From a less sophisticated bidder's perspective, a more useful dividing line may be whether he is subject to estimation errors *or not. If a bidder is subject to estimation errors, he faces selection bias: he is more likely to win when the error is positive. This selection bias has been identified first by Capen et al. (1971), and the optimal response to it is caution.*[15] *This phenomenon is not specific to auctions or the presence of interdependencies in valuations: it may arise in any decision problem where an agent compares an alternative that is easy to evaluate (not buying) to one that is more difficult to evaluate hence subject to estimation error.*[16]

[12] This is the path taken in Compte and Postlewaite (2012). In this paper, we investigate whether and when some rank related signals promote or diminish competition.
[13] Note that the motive invoked here as to why poor information hurts the seller is different from that invoked in Milgrom and Weber (1982), which relies on affiliation.
[14] That logic is pursued in Ganuza (2004) for example.
[15] See Compte (2001), which examines the effect of increasing the number of bidders on this selection bias, in the context of the second-price auction.
[16] See Compte and Postlewaite (2012) and Chapter 21.

Now, the level of caution depends on context, and indeed, the degree of interdependence then matters. If idiosyncratic components are less dispersed (small d), which can be interpreted as values being more interdependent, the estimation errors carry more weight and caution should increase.

Efficiency in first price. *In a symmetric environment with independent values, an equilibrium in monotonic strategies is guaranteed to exist, implying an efficient allocation of the object.*[17] *With our focus on a restricted set of strategies, this is no longer guaranteed, and we explained in Section 6 when it may fail to exist.*[18] *Since our model may be interpreted as one in which players are unable to draw precise inference about rank given their values, as in a correlated value model, our model also gives hints as to when there may be nonexistence of an equilibrium in monotonic strategies in a standard correlated value model.*

In any event, one reason that first-price auctions are likely to generate inefficient outcomes is that symmetry in perceptions of the environment is likely to fail, hence differences in shading levels are likely to arise, generating inefficient allocations. Whether these inefficiencies are large in practice remains, however, an important applied question.

Additive versus multiplicative shading. *Technically, one difference between additive and multiplicative shading is that the latter mechanically incorporates into bidding some information about rank that the former does not: in the multiplicative version, shading is stronger for higher value realizations. For some auction formats such as the first-price auction, the qualitative effects are similar under both assumptions. For other auction formats such as the all-pay auction, an additive shading assumption could drastically modify the analysis and make existence of an equilibrium an issue even with few bidders. To us, however, failure of existence is not a weakness of the additive formulation, but an illustration of what is necessary to generate stable predictions under that particular auction format.*

Suggestions for further research/applications. Through strategy restrictions, one obtains a more parsimonious treatment of auctions, with only one, or few, dimensions of behavior being endogenized. This, in turn, may help deal with problems that are technically difficult to address within the standard framework, or at least simplify their analysis. In the spirit of Chapter 7, it may also help assess the robustness of existing

[17] See Milgrom and Weber (1982).

[18] In essence the reason is similar to why competition for differentiated goods may not generate a pure strategy equilibrium. See Caplin and Nalebuff (1986).

models in which players' knowledge of their rank in the distribution plays a key role.

Specific topics and relevant references: (i) auctions in which buyers get other signals beyond their own valuation (Fang and Morris (2006)); (ii) auctions in which multiple units or slots are sold;[19] (iii) dynamic selling problems in which the seller faces a sequence of buyers (Lauermann and Wolinski (2016)); (iv) contests in which one or several prizes are offered (Moldovanu and Sela (2001)),[20] and more generally, matching problems;[21] (v) auctions in which the seller chooses a reserve price strategically, as a function of his perception of the buyers' valuations.[22]

References

Capen, E. C., Clapp, R. B. and Campbell, W. M. (1971). Competitive bidding in high risk situations. *Journal of Petroleum Technology*, 23, 641–653.

Caplin, A. and Nalebuff, B. (1986). Multi-dimensional product differentiation and price competition. *Oxford Economic Papers*, 38, 129–145.

Compte, O. (2001). The winner's curse with independent private values, mimeo.

Compte, O. and Postlewaite, A. (2012a). Simple auctions, mimeo.

Compte, O. and Postlewaite, A. (2012b). Cautiousness, mimeo.

Fang, H. and Morris, S. (2006). Multidimensional private value auctions. *Journal of Economic Theory*, 126, 1–30.

Ganuza, J.-J. (2004). Ignorance promotes competition: an auction model with endogenous private valuations. *The RAND Journal of Economics*, 35(3), 583–598.

[19] One could for example study the performance of a product-mix auction (Klemperer (2010)), as compared with a sequential auction of each good. More generally, many selling procedures are possible, each allowing bidders to express their preferences to varying degrees, and each having different rules for selecting winners and transfers. Comparing the performance and efficiency properties of these formats is a challenge which can be easily surmounted if the set of bidding strategies examined is reduced (e.g., by assuming linear shading across all units).

[20] These games typically have the structure of an all-pay auction, as that analyzed in Chapter 7.

[21] One part of the matching literature deals with the matching of students to various schools. In that literature, students behave as if they had perfect assessments of their chance of being admitted to each school. One might reconsider these models, by assuming either that agents get only noisy signals of their position (in the queue for schools), or that the intensity of their preferences does not give them accurate information about how that intensity compares to others'. A second segment of the matching literature considers matching of buyers and sellers to understand how surplus in matching will be shared. One can think of such models as one side of the market suggesting shares to agents on the other side. Equilibria of such games implicitly assume that each agent's suggested shares are precisely calibrated to his relative ranking within his side and to characteristics of the other side. Strategy restrictions could model situations where agents are less able to discern their relative ranking on the basis of their own characteristics, and less able to precisely target each agent on the other side.

[22] In standard terms, these are called "informed principal problems" and they are notoriously difficult to solve. From our perspective, we are only adding a single strategic variable. See Chapter 17.

Klemperer, P. (2010). The product-mix auction: a new auction design for differentiated goods. *Journal of the European Economic Association*, 8, 526–536.
Lauermann, S. and Wolinsky, A. (2016). Search with adverse selection. *Econometrica*, 84, 243–315.
Milgrom, P. and Weber, R. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50, 443–459.
Moldovanu, B. and Sela, A. (2001). The optimal allocation of prizes in contests. *The American Economic Review*, 91, 542–558.