Plausible theories of behavior

Olivier Compte and Andy Postlewaite

April 2011

Standard game theoretic modelling starts with the definition of a specific game, and it ends with the computation of equilibrium strategies. The general view expressed in this paper is that the equilibrium strategies we derive are often implausible.

The definition of a game generally includes the description of a specific information structure, a specific timing of observations and moves by each player, and specific payoffs associated with the possible moves. Most interactions we analyze involve a rich information structure, in the sense that the set of possible signals or observations that an agent may receive or make is huge. In standard models, this rich information structure translates into a huge strategy space, with each player attempting to pin down the optimal way to adjust his decision as a function of the signals he receives or observations that he makes.

We have several objections to the standard approach.

First, we find implausible that players would *manage* to pin down an optimal response. It could be through introspection. It would then have be based on some a priori knowledge of distributions over observations others make or over signals that others receive, and over payoffs that others get, but where then would this a priori knowledge, if not faith, come from? It could also be through learning. But how can this be achieved if the strategy space is huge? At the very least, the standard approach is silent about that.

Second, we find implausible that players would *attempt* to find it. The situations one faces vary, and rather than trying to look for a strategy that is optimal under a specific circumstance, with no guarantee that it will perform well once the situation changes slightly, one should better attempt to find strategies that perform well in a variety of circumstances.

Finally, as modelers, we run two risks: that of finding 'insights' that would only be valid in this limit of arbitrarily sophisticated players; that of spending much effort on second order effects, and of overlooking first order ones.

Our approach is different, and we shall attempt to illustrate it using various applications. A crucial aspect of our approach is that among the many strategies that a player could use, he only evaluates and compares a limited number of alternative strategies. This subset of alternatives allows the agent to adjust behavior when the circumstances change. Which subset of alternatives he uses will have some importance, as it will affect how the games he faces are eventually played.

Our model can be interpreted as a hierarchical or two-tier model of behavior, whereby not all dimensions of a strategy are endogenized simultaneously: which strategy within the subset is played is subject to quick adjustment, and depends on the particular game being played; which subset is used is subject to slow evolution, and this evolution, if any, is driven by the distribution over the problems he faces.

Once the strategy restrictions are made, analysis becomes relatively easy. One immediate benefit of our approach will thus lie in providing simple models to address issues usually thought to be difficult to handle.

But we aim for more than just a simpler tool. Our model provides a way to deal with less sophisticated agents,[1] and as such, a tool to check whether the insights that we derive in standard models are robust to lesser sophistication. From a methodological perspective, several comments are worth making: for the insights that are robust, the highly sophisticated agent case points to effects that seem second order, and it is not clear that those second order effects should receive special attention, as there is a plentiful of alternative routes to look for second order effects; for the insights that are not robust, we ought to be suspicious of those; finally, the power of the explanations we propose lies in the parsimony of the models we use: how well standard models fare in this respect is not clear.

As said earlier, once the strategy restrictions are made, analysis becomes relatively easy. One challenge will be to motivate and derive appropriate strategy restrictions. Depending on the nature of the strategic interaction, or depending on the nature of the problem that the agent thinks he is facing, different strategy restrictions will be relevant. It is the purpose of the following sections to suggest plausible strategy restrictions for various strategic interactions.

1 Reputation

An agent's reputation refers to the assessment(s) that others' make about that agent's abilities, possibly his moral qualities such as honesty, or about the general way he handles particular situations or behaves in particular circumstances. These assessments are based on past observations, and, because one views that the qualities one attributes to a reputed agent have a good chance to last for a while, these assessments can be used to predict what to expect from an interaction, hence to determine whether such an interaction is worthwhile: one prefers to deal with a reputed agent and avoid interactions with non-reputed ones. This discrimination between reputed and non-reputed agents has one consequence: it creates an incentive to behave according to standards that promotes one's reputation, even if this comes with a short run cost.

We are interested in building a simple model with precisely these ingredients: on one side, an agent (say a seller, or *the* Agent – which we will refer to as player 2) attempting to build or preserve a reputation for good conduct (say, providing quality goods), and concerned about the possibility that he loses his reputation if he does not meet the other side's standard; on the other side, another agent (say a buyer, or the Principal – which we will refer to as player 1) attempting to discriminate between good and bad conduct, and having the option to stop buying for a while if he becomes convinced that the prospects from the interaction are currently not good enough (i.e. that he currently faces a buyer who tends to provide poor quality products).

[1] We will discuss how it compares to some alternative approaches, in particular Crawford.

Before describing a full model of the interaction, we start by describing *the class of decision problem that player 1 typically faces in such interactions.*

1.1 Optimal learning versus Limited learning

In essence, player 1 is facing a two-arm bandit problem. One arm is safe (not interacting) with payoff normalized to 0. The other arm is risky, as the benefit from the interaction is uncertain. And in each stage where he interacts, he receives a signal correlated with the current benefit from the interaction, say he is satisfied ($y = \bar{y}$) or not satisfied ($y = \underline{y}$).

To be more precise, we assume that there is an underlying stochastic process that determines at any given stage the benefit from the risky arm as well as the probability to receive the good signal ($y = \bar{y}$). Formally, there is a state space Θ, where $\theta \in \Theta$ determines the expected benefit v_θ from the risky arm and the probability p_θ to receive a good signal, and there is stochastic process ω over θ that defines the law of motion of θ over time, and which we assume to be Markovian.

Ideally, the agent would like to interact if and only if v_θ is positive. But the agent does not observe θ. Yet the signals y that he observes provide him with information on θ. Following the standard approach, a *learning strategy* σ for the agent determines a decision to interact as a function of the history of observations made so far. The logic of the standard approach would be as follows: a stochastic process ω defines a particular problem that the agent faces; call $\mathcal{P}$ the set of problems that the agent might face, then for any given $\omega \in \mathcal{P}$, one can compute the *optimal learning strategy* σ_ω.

Our view is that it is implausible that the agent would *know* the stochastic process ω that he currently faces. The traditional response is that the difficulty can be bypassed if one assume that the agent can assess which strategy σ performs best. If that can be done, then by picking the strategy that performs best, the agent will behave as if he knew ω. However, the set of possible strategies is huge, and it is implausible all these performance assessments can be done.

Our view is that these performance assessments can only be carried for a limited number of alternative strategies, and our objective below is to suggest such alternatives.

One simple alternative would be to decide to always trust and interact (I), independently of the signal received. This strategy may be represented as a single state automaton that always play I, and we will refer to it as rule r_I. Another alternative would be to always distrust and thus avoid interaction (A), again independently of the signal received. Again this strategy may be represented as a single state automaton, which this time always plays A (not interacting). We refer to this alternative as r_A.

For any given $\omega \in \mathcal{P}$, one can compute the long run value $E_\omega v_\theta$ associated with the alternative r_I. If the set R of alternatives only consisted of these two alternatives (i.e. $R = \{r_I, r_A\}$), then the set of problems $\mathcal{P}$ would be partitioned in essentially two subsets depending on whether $E_\omega v_\theta$ is positive or negative.

In both of the above alternatives, the agent does not use past signals to modify his decision to interact. We now suggest that the agent may be more sophisticated and attempt to track the state θ using the signal y he just received. As an illustration, assume that the agent may be in two possible mental states, normal (N) or upset (S), and that he switches back and forth between those two mental states as a function of the signal he receives:

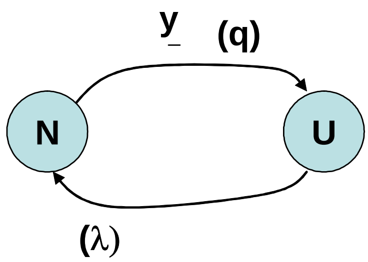

Under that automaton, transition to U requires that $\underline{y}$ is received, and it then occurs with probability q. Transition from U to N occurs with probability λ in any period. Now define the rule $r^{q,\lambda}$ induced by the mental system above that plays I at N and A at U.

For any given pair (q, λ), say (q_0, λ_0), the rule r^{q_0,λ_0} is still a crude one, and it may be far from optimal for specific problems ω. However, there are problems ω for which that rule is an improvement over both r_I and r_N. A collection of rules that the agent might consider using is thus:

$$R = \{r_I, r_N, r^{q_0,\lambda_0}\},$$

and as above, the set of environments/problems can now be partitioned into three subsets. We are not saying that there ought to be only three rules, or that given a three rule constraint, that would be the optimal one to pick.

We are saying that this is a good enough triplet of rules that the agent could use. One could even argue that this particular triplet is optimal for some distribution over problems, so assuming a specific collection of rules does not seem more ad hoc and assuming a specific distribution over problems.

Maybe one point to emphasize is that our approach attempts to separate two issues: which subset of strategies or collection of rules an agent considers using? Which strategy within the subset the agent ought to pick? The first question is related to the distribution over problems that the agent faces, or possibly to some perception of the problem that he is facing. The second question is related to the specific interaction that he is currently facing.

1.2 A reputation model.

So far, we have focused on the decision problem of an agent faced with a bandit problem, and we have proposed a collection of rules that the agent could use for any such problem that he faces. The next step is to apply this collection of rules to the study of a specific interaction.

Stage game. We assume that at any date where the agent interacts (i.e. takes the risky arm), player 2 decides on an effort level $e \in [0,1]$. Effort is costly but it increases the probability that player 1 is satisfied. Whether player is satisfied ($y = \bar{y}$) or dissatisfied ($y = \underline{y}$) is private information to player 1. We assume that

$$p \equiv p_0 + (1 - p_0)\alpha e,$$

where $\alpha \in (0,1)$. In the buyer-seller case, α can be interpreted as a taste parameter that reflects how the seller's product matches with the buyer's taste. We also let $v(p)$denote the expected payoff that player 1 obtains when he has an ex ante probability p of being satisfied, and we assume that $v(p_0) < 0 < v(1)$. It will also be convenient to define p^* as the probability for which the buyer is indifferent between buying and not buying, that is $v(p^*) = 0$. Finally, we denote by $c_\beta(e) = \beta c(e)$ the cost of effort for player 2, where β is a cost parameter. In summary, the stage game will be characterized by two parameters: α, a taste parameter, and β, a cost parameter, and we let $\theta = (\alpha, \beta)$..We shall assume that player 2 observes β, and that player 1 observes neither α nor β.

Dynamic game, strategies and definition of equilibrium.

Turning to the dynamic setting, we shall assume that θ follows a random process taking values in some set Θ, characterized by a transition matrix $\omega : \Theta \to \Theta$.

We shall assume that the agent does not observe the state (nor ω), and that as explained in the previous section, he restricts attention to a small set of rules R_1, independently of the process ω and of the actions taken by player 1.

As for player 2, we also assume that he restricts attention to a small set of rules R_2, independently of the process ω. Specifically, we assume that he picks effort as a function of the cost parameter. In principle, player 2 could attempt to track α by conditioning his effort choice on the current frequency with which he gets to interact with player 1. We first choose to have him focus on the information that is presumably most salient, i.e. the cost parameter. We shall then investigate how enlarging his set of strategies would affect results.

Each pair of rules $(r_1, r_2) \in R_1 \times R_2$ induces a long run distribution over payoffs, to which we may thus associate a value $v_i(r_1, r_2)$ for each player i. A pair of rules (r_1, r_2) is an equilibrium if for each i, player i has no profitable deviation in R_i.[2]

Analysis.

We illustrate our approach through examples.

[2]Note that we often restrict attention to a finite set of rules, and that an equilibrium in pure strategies is not guaranteed. With finite sets of rules, an equilibrium in mixed strategies always exist. Our preferred interpretation of a mixed strategy equilibrium is as follows. Players alternate between the various rules they consider, and they play each rule r_i for a random length of time that is in expectation increasing, say proportional, to how that rule fares compared to others. The limit case where the coefficient of proportionality is large converge to a (possibly mixed strategy) equilibrium

We start with example 1, in which β remains constant throughout out the interaction.

> **Example 1:** *β is constant and α follows a random process between two values α^H or α^L. We let $\pi = \Pr(\alpha^{t+1} = \alpha^L \mid \alpha^t = \alpha^H)$, and $\pi' = \Pr(\alpha^{t+1} = \alpha^H \mid \alpha^t = \alpha^L)$.*

In equilibrium, player 2 chooses an effort level e^* which we are going to determine. We distinguish several cases.

Case 1: $R_1 = \{r_I, r_A\}$.

Player 1 may only use r_I or r_A, so play is independent of history. It follows that player 2 has no incentives to make effort, and the probability that player 1 is satisfied when interacting is p_0. Since $v(p_0) < 0$, player 1's optimal alternative is r_A. Hence there is a unique equilibrium: $r_1 = r_A$ and $e^* = 0$.

Case 2: $R_1 = \{r_I, r_A, r^{q_0,\lambda_0}\}$.

The above observation is still valid, so $r_1 = r_A$ and $e^* = 0$ remains an equilibrium. We now check whether and when there exists an equilibrium where player 1 uses r^{q_0,λ_0}.

The strategy r^{q_0,λ_0} has two roles.

- First it creates an incentive for player 2 to make an effort, because he is rewarded with a higher frequency of interactions when he makes an effort. Indeed, when he make an effort e, he generates a probability of satisfaction equal to $p(e, \alpha) = p_0 + \alpha e(1 - p_0)$, and a probability of satisfaction p in turn generates a long run frequency of interactions equal to

$$\phi(p) \equiv \frac{\lambda}{\lambda + q(1-p)}$$

when α is constant. When α varies, the expression is more complicated, but still increasing in p hence in the effort e. We can therefore deduce an optimal value e^* for the effort of player 2 when player uses r^{q_0,λ_0}.

- Second, using r^{q_0,λ_0} enables player 1 to track (imperfectly) how his taste currently match with the product produced. For a given effort level e^*, there are two possible values for the probability that player 1 is satisfied:

$$\bar{p} \equiv p_0 + \alpha^H e^*(1 - p_0) \text{ and } \underline{p} \equiv p_0 + \alpha^L e^*(1 - p_0).$$

If $v(\bar{p}) < 0$, then it is always better to avoid interactions: r_A is the best response for player 1. If $v(\underline{p}) > 0$, then it is always better to interact: r_I is the best response for player 1. If $v(\underline{p}) < 0 < v(\bar{p})$, then r^{q_0,λ_0} may be a best response. The benefit of using r^{q_0,λ_0} lies in the fact that it permits player 1 to adjust the frequency of interactions to p. Ideally, the agent would prefer to adjust interactions so that he always interacts if $p > p^*$, and never interacts if $p < p^*$ – both constant rules r_I and r_A would then be dominated. Here, adjustment to the value of p is not perfect and depending on parameter values, r^{q_0,λ_0} may or may not be optimal. It will be optimal for example if $v(\underline{p})$ and $v(\bar{p})$ are sufficiently far away from 0 (this will be the case if e^* is not too small and if α^H

and α^L are sufficiently far apart). If transition probabilities π and π' are very small for example, we have, letting $\rho = \frac{\pi'}{\pi+\pi'}$:

$$v(r^{q_0,\lambda_0}, e^*) = \rho\phi(\bar{p})v(\bar{p}) + (1-\rho)\phi(\underline{p})v(\underline{p})$$

and the equilibrium condition is:

$$\rho\phi(\bar{p})v(\bar{p}) + (1-\rho)\phi(\underline{p})v(\underline{p}) > \max\{0, \rho v(\bar{p}) + (1-\rho)v(\underline{p})\}$$

Comment: In case $v(\underline{p}) > 0$ or $\rho v(\bar{p}) + (1-\rho)v(\underline{p}) > \rho\phi(\bar{p})v(\bar{p}) + (1-\rho)\phi(\underline{p})v(\underline{p}) > 0$, the strategy r^{q_0,λ_0} creates incentives to make effort that are strong enough that the strategy r_I that always interact becomes the best strategy for player 1. This means that there cannot be an equilibrium in pure strategy where interactions occur. However, there may exist an equilibrium in mixed strategy, where player 1 plays r^{q_0,λ_0} with frequency, say η. A lower η implies a smaller marginal effect of effort on the frequency of interactions, hence a reduced incentive to make effort, thereby reducing both $\bar{p}$ and $\underline{p}$. A mixed strategy equilibrium exists if for some effort level e^*, the equality

$$\rho\phi(\bar{p})v(\bar{p}) + (1-\rho)\phi(\underline{p})v(\underline{p}) = \max\{0, \rho v(\bar{p}) + (1-\rho)v(\underline{p})\}$$

holds.

We now move to example 2, in which α remains constant but β varies.

Example 2: *α is constant and β follows a random process between two values β^H or β^L. We let $\pi = \Pr(\beta^{t+1} = \beta^L \mid \beta^t = \beta^H)$, and $\pi' = \Pr(\beta^{t+1} = \beta^H \mid \beta^t = \beta^L)$.*

The difference with the previous example is that player 2 may now adjust his effort level to his cost, so his optimal strategy will be characterized by two numbers e^*_H and e^*_L.

As above the case where $R_1 = \{r_I, r_A\}$ does not allow for reputation formation: there is a unique equilibrium, in which $r_1 = r_A$ and $e^*_H = e^*_L = 0$.

The case where $R_1 = \{r_I, r_A, r^{q_0,\lambda_0}\}$ is the more interesting one: while $(r_1 = r_A,\ e^* = 0)$ remains an equilibrium, there may also be an equilibrium involving r^{q_0,λ_0}. The logic is identical. The strategy r^{q_0,λ_0} has again two roles. It creates an incentive for player 2 to make an effort, and since the cost of effort varies, two distinct values e^*_L and e^*_H for the level of effort result, as a function of β.

These distinct effort levels generate two distinct values for p, $\bar{p} = p(e^*_H, \alpha)$ and $\underline{p} = p(e^*_L, \alpha)$. For player 1, using r^{q_0,λ_0} enables him to track (imperfectly) how the cost hence the probability p varies, and, depending on parameter values, this may turn out to be a better strategy than both r_I and r_A.

Example 3: *Same as example 2, but with* $\beta_L < 0$*: player 2 then strictly prefers to make highest possible effort, that is* $e_L^* = 1$ *at all times.*

This example is in the spirit of the models studied in the literature, and that have introduced behavioral types. We introduce it to make easier the connections with the literature.

Player 2(L) continues to make effort whatever the strategy of player 1. The consequence is that there may no longer be an equilibrium where player 1 always avoid interactions.

In case $R_1 = \{r_I, r_A\}$, player 2 (H) has no incentives to make an effort. Letting $\bar{p} = p(\alpha, 1)$, the value from interacting for player 1 is thus $\rho v(\bar{p}) + (1 - \rho)v(p_0)$, and player 1 has an incentive to play r_I if

$$\rho v(\bar{p}) + (1 - \rho)v(p_0) > 0,$$

that is if ρ is high enough. Of course, ρ has to be quite large for this to be true.

Consider now the case where $R_1 = \{r_I, r_A, r^{q_0,\lambda_0}\}$. The condition under which "not interacting" fails to be an equilibrium is weaker than before, as it is enough that:

$$\max(\rho v(\bar{p}) + (1 - \rho)v(p_0), \rho\phi(\bar{p})v(\bar{p}) + (1 - \rho)\phi(p_0)v(p_0)) > 0$$

The condition is strictly weaker if α is large enough, because then $\bar{p}$ gets close to 1 hence $\phi(\bar{p})$ gets close to 1. When $\alpha = 1$, the condition:

$$\rho v(1) + (1 - \rho)\phi(p_0)v(p_0) > 0$$

ensures that player 1 uses r^{q_0,λ_0} in equilibrium, thereby creating an incentive for player 2(H) to make some effort. For some r^{q_0,λ_0}, $\phi(p_0)$ can be quite small (this is the case when λ_0/q_0 is small), and the above inequality holds even when ρ is small.

This latter observations explains why the presence of a behavioral make it possible to avoid the "no interaction equilibrium". As more rules are available to the agent, it becomes easier to find one for which $\phi(p_0)$ is small and $\phi(\bar{p})$ is large, then explaining why a small uncertainty may swamp out the "no interaction" equilibrium.

Now observe that the incentive to make effort for player 2(H) cannot be too strong, because if it were, and if player $2(H)$ was induced to make some effort e_H^* such that $\underline{p} = p(\alpha, e^*) > p^*$, then there would no longer be incentives for player 1 to use the rule r^{q_0,λ_0} that discriminates between states (H) and (L).

The consequence of the latter observation is that in equilibrium, if ρ is small, then player 2 (H) must have incentives to choose effort close (and slightly below) p^*. This is a (simpler) version of the usual Stackelberg result.

Summary:

The examples above share one key ingredient. There is an underlying state not known to players and that has sufficient permanence. Examples 1 and 2 are meant to suggest that the underlying state need not be confined to a cost type for player 2 (the parameter β), but could also be related to how current circumstances affect how well player 2's effort matches player 1's needs.

The uncertainty about the current state and its permanence creates an incentive to use rules of behavior that condition play on observations correlated with the current state, because the signals can help track the current state. This is precisely what player 1 does, creating an incentives for player 2 to make effort.

We have focused on rules $r^{\lambda,q}$ that alternate between two possible mental states. The approach can be generalized to rules that have more states, each different state representing a different degree of trust in player 2's behavior.

Example 3 is meant to capture the main working hypothesis of the standard literature: the possibility that will small probability player 1 is facing a behavioral type that plays in a specific way, and the conclusion that it induces the rational type (here player 2(H)) to mimic the behavioral type (here, to make higher effort).

Extensions.

Many signals. We have assumed that player 1 could either be satisfied or dissatisfied. One interpretation is that he may only be receiving two signals. Another interpretation consistent with our model is that there is a much reacher set of signals that he may be receiving, say Y, and that he classifies or *partitions* the signals into two categories. Once the set of signals Y is partitioned into two subsets, say $\underline{y}$ and $\bar{y}$, with $Y = \underline{y} \cup \bar{y}$, then our analysis applies almost unchanged. Assuming that there is an underlying distribution $f(y \mid e, \alpha)$ that describes how signals are generated as a function of player 2's effort level and some parameter α, one can define:

$$p(e, \alpha) = \int_{y \in \bar{y}} f(y \mid e, \alpha) dy$$

and apply our analysis using this specific function.

Different partitions would generate different functions $p(.,.)$, thereby affecting the possibility to generate a reputation mechanism in equilibrium. But qualitatively the same conclusions would remain, the joint possibility that player 1 finds optimal to use a sophisticated strategy and that player 2 has incentives to make some effort.

Now which partition is reasonable is a legitimate question, as was the question asking which subset of strategies should the agent consider. Endogenizing which partitions would arise based on optimality criteria would have some theoretical appeal. However, as we have already emphasized, our view is that it is implausible that an agent would be looking for a partition that would be optimal in the space of all possible partitions (let alone the fact that optimality would be specific to the particular distribution over signals chosen).

Rather our view is that the agent has a perception of the problem he is facing, most plausibly an erroneous one, and that this perception could guide the partition he consciously or unconsciously selects. As an illustration, we might assume that he has a perception that signals are either generated according to $\widetilde{f}(y \mid e_H)$, or according to $\widetilde{f}(y \mid e_L)$, and that the (perceived) likelihood ratio

$$\widetilde{l}(y) = \frac{\widetilde{f}(y \mid e_H)}{\widetilde{f}(y \mid e_L)}$$

then determines, by comparing it to some threshold l^*, whether he is satisfied or unsatisfied:

$$\bar{y} \equiv \{y, \widetilde{l}(y) > l^*\} \text{ and } \underline{y} \equiv \{y, \widetilde{l}(y) < l^*\}.$$

More sophisticated strategies. One could think of several ways by which a rule of behavior could be more sophisticated than $r^{\lambda,q}$. There could be more states. There could also be more categories than "satisfied" or "unsatisfied", for example various degrees of satisfaction, generating different transition across states.

More signals (continued). We have assumed that if he does not interact, player 1 receives no signals. It is conceivable however that he would receive signals from other sources. This would for example be the case if player 2 were also interacting with other players, and if, in addition, he were to receive signals $z \in \{\bar{z}, \underline{z}\}$ correlated with other player's satisfaction.[3] A natural extension of $r^{\lambda,q}$ would be to consider the automaton:

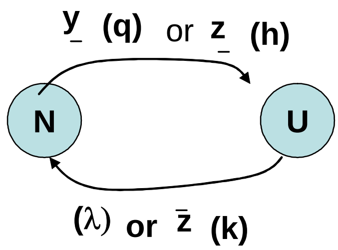

and have player 1 play I at N and A at U. In comparison with $r^{\lambda,q}$, the above automaton is characterized by three parameters $(\lambda, q, h, k))$, and it provides player 1 with an additional chance (k) to transit back to N, in the event good news $(\bar{y})$ has been received, and with an additional "chance" (h) to transit from N to U in the event $\underline{y}$ is received.[4]

To illustrate the kind of insight one could get, let us discuss two cases: $R_1 = \{r_I, r_A, r^{q_0,\lambda_0}\}$ and $R_2 = \{r_I, r_A, r^{q_0,\lambda_0,h_0,k_0}\}$.

If player 1 interacts very infrequently with player 2, then under R_1, we should not expect r^{q_0,λ_0} to be particularly useful. Now under R_2, if satisfaction for other players is poorly correlated with player 1's satisfaction, then the rule $r^{q_0,\lambda_0,h_0,k_0}$ will induce a behavior that is not that different from a mixed strategy uncorrelated with the true state, and for most parameters, that rule will not be a best response for player 1. No reputation should emerge in this case.

[3] Think here of many player 1s, all interacting relatively unfrequently.

[4] That is, at the end of each period, $\Pr(U \to N \mid \underline{y}) = \lambda$ and $\Pr(U \to N \mid \bar{y}) = \lambda + (1-\lambda)k$.

In other words, the possibility that player 1 uses rule $r^{q_0,\lambda_0,h_0,k_0}$ may be useful in context where player 1 interacts infrequently, and where other players' tastes are correlated with player 1's.

Discussion and comparison with the literature.

Standard approaches to reputation start with the possibility that (with small probability) player 2 follows some pre-determined behavior (a crazy, or say honest, type). The exercise then consists in endogenizing the behavior of the rational (or non-crazy type) player 2, as well as player 1's behavior and beliefs about player 2's honestly. One can interpret these beliefs as a measure of player 2's reputation. Deriving equilibrium behavior and beliefs is a relatively simple task when behavior is observed perfectly, as reputation completely vanishes as soon as observed behavior conflicts with that of an honest type. That task is much more difficult when observations are noisy (and private), and results generally take the form of a bound on equilibrium payoffs for arbitrarily patient players, with little qualitative insights on how past observations concur to generate a good (or bad) reputation.[5]

The simplicity of our approach lies in the fact that for the possible rules considered, there are only few "belief states" in which player 1 can be. Under r_I or r_A, there is a single "belief state", and thus no hope for reputation effects. Under r^{q_0,λ_0}, there are two beliefs states, as player 1 thinks either highly or poorly of player 2, with changes from one belief state to the other being fixed exogenously as a function of the signal received. With these two belief states, one can hope for reputation effects, as player 2 can try to influence player 1's belief state through his actions. However there is no guarantee that player 1 will wish to use that rule, and our analysis is precisely aimed at finding the set of parameters for which this will be the case.

2 Cooperation

Through the course of a relationship or partnership, incentives to invest vary: at times we may be upset over the way the partnership goes, with poor expectations for sustained cooperation and little hope that making effort would increase prospects; at other times, we feel good about prospects, ready to invest again in the relationship and worried that decreasing effort would undermine cooperation. We are interested in building a simple model that explains conditions under which cooperation (i.e. high effort level on both sides) can be sustained.

A partnership bears some resemblance with the strategic interaction just studied. The basic stage game is essentially identical. A player chooses an effort level, and this affects whether his partner will be satisfied or dissatisfied. There is a difference however, as moral hazard is two sided, and sustained cooperation requires that both sides make an effort. In the rest of this Section, we shall assume that the stage game is played repeatedly, with one player choosing an

[5] Standard approaches also have a hard time dealing with settings where behavioral uncertainty does not stem from the existence of crazy types but from cost or preference uncertainty.

effort level, next the other player choosing an effort level. And so on indefinitely.[6]

From a given player's perspective, an objective may be to attempt to *discriminate* between instances where effort is worthwhile, and instances where it is not, and discrimination can be based on the signals that the player gets (i.e. whether he is satisfied or not). Discrimination is key, because as in the reputation model, if such an attempt take place, this tends to create incentives for the other side to make some effort.

But this cannot be the full story: if for the other side, the incentives to make effort do not vary over time, and if effort therefore remains unchanged, there won't be incentives to discriminate in the first place. The objective of what follows is to explain how this can be achieved.

As before, we start by describing the class of decision problem that each player typically faces in such interactions.

2.1 The decision problem

In partnership games of the kind described above, the decision problem is more complex than that studied before. Even if we limit ourselves to two possible effort levels $e \in \{0, 1\}$, so that again we can think of a two arm bandit problem (one arm for each effort level), a key difference is that the benefits from taking a particular arm may now depend on his own previously chosen effort levels. Formally, there is an underlying state space Θ, where $\theta \in \Theta$ determines an expected benefit v_θ and the probability p_θ to receive a good signal, and there is stochastic process ω (assumed to be Markovian) that defines a transition probabilities over states *as a function of one's current effort level.*

Following the standard approach, a *strategy* σ would determine an effort decision e as a function of the history of observations made so far, but as before, we find implausible that an optimal strategy could be found.

Rather, we will focus on few rules or alternatives that an agent might consider using.

One simple alternative would be to decide to always make effort, independently of the signal received. This strategy may be represented as a single state automaton that always play $e = 1$, and we will refer to it as rule r_C. Another alternative would be to never make effort, independently of the signal received. Again this strategy may be represented as a single state automaton, which this time always plays $e = 0$. We refer to this alternative as r_D.

In both of the above alternatives, the agent does not use past signals to modify his decision to interact. We now suggest that the agent may be more sophisticated and attempt to track the state θ using the signal y he just received. As an illustration, assume that the agent may be in two possible mental states, normal (N) or upset (S), and that he switches back and forth between those

[6] Another difference with the reputation game is that stopping the interation for a while is no longer an option.

two mental states as a function of the signal he receives:

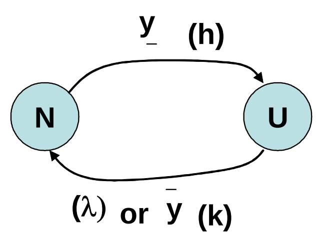

Under that automaton, transition to U requires that $\underline{y}$ is received, and it then occurs with probability h. Transition from U to N occurs with probability λ in any period, with an additional chance to transit back to N in case $y = \bar{y}$ has been received.[7] Now define the rule $r^{h,k,\lambda}$ induced by the mental system above that plays $e = 1$ at N and $e = 0$ at U.

For any given triplet (h_0, k_0, λ_0), the rule r^{h_0,k_0,λ_0} is still a crude one, and it may be far from optimal for a specific problem ω. However, there are problems ω for which that rule is an improvement over both r_C and r_D. A collection of rules that the agent might consider using is thus:

$$R = \{r_C, r_D, r^{h_0,k_0,\lambda_0}\}.$$

2.2 A partnership game.

Given the class of problems that the agent faces, we have proposed a collection of rules that the agent could use for any such problem that he faces. The next step is to apply this collection of rules to the study of a specific partnership model.

Stage game. There are two kinds of stage game, depending on who has the option to make an effort. In stage game 1, player 1 decides on an effort level $e_1 \in \{0, 1\}$. Effort is costly but it increases the probability $p(e)$ that player 2 is satisfied. Whether player is satisfied $(y = \bar{y})$ or dissatisfied $(y = \underline{y})$ is private information to player 1. We let $p_0 \equiv p(0)$ and $p_1 \equiv p(1)$. We also let $v(p)$denote the expected payoff that a player obtains when he has an ex ante probability p of being satisfied, let $v_0 \equiv v(0)$ and $v_1 \equiv v(1)$ and we assume that $v_0 < v_1$. Finally we denote by c the cost of effort. Stage game 2 is defined similarly, with player 2 making effort.

In summary, the stage game is characterized by the vector of parameters $\xi = (p_0, p_1, v_0, v_1, c)$. For simplicity we shall assume that these parameters do not vary overtime, but in principle they could.

Dynamic game, strategies and definition of equilibrium.

The stage game described above is repeated, with each player taking turn in getting the option to make an effort, and as before, we assume that each player restricts attention to a small set of rules $R_i = \{r_C, r_D, r^{h_0,k_0,\lambda_0}\}$.

[7] That is, at the end of each period, $\Pr(U \to N \mid \underline{y}) = \lambda$ and $\Pr(U \to N \mid \bar{y}) = \lambda + (1 - \lambda)k$.

Each pair of rules $(r_1, r_2) \in R_1 \times R_2$ induces a long run distribution over payoffs, to which we may thus associate a value $v_i(r_1, r_2)$ for each player i. A pair of rules (r_1, r_2) is an equilibrium if for each i, player i has no profitable deviation in R_i.

Result.

> **Proposition (Compte and Postlewaite 2010):** (r_D, r_D) *is always an equilibrium. But there also exist a non-trivial set of parameter values for which sustained cooperation (in which each player follows rule* r^{h_0,k_0,λ_0}*) is an equilibrium. The condition* $h_0 < k_0$ *is a necessary condition for existence of such an equilibrium.*

In essence, the reason why a cooperative equilibrium exists is as follows. The rule r^{h_0,k_0,λ_0} induces play that depends on the history of signal received. In contrast to the reputation example, uncertainty about the stage game parameters ξ is not be needed to provide players with incentives to use such a rule. The mere fact that, say player 2 uses rule r^{h_0,k_0,λ_0} creates endogenous strategic uncertainty, and at least for a non trivial subset of parameter values h_0, k_0, λ_0, this fact alone may be sufficient to make optimal the choice of rule r^{h_0,k_0,λ_0}. by player 1.

It is not guaranteed however that r^{h_0,k_0,λ_0} is always an equilibrium. If h_0 is too small, then incentives to play r_D will be strong. If k_0 is smaller than h_0, then recoordination on cooperation is difficult, which makes the alternative r_C an attractive one because it avoids spending much time in phases where no player makes effort.

3 Cautiousness

Limiting the strategy sets has proved helpful in providing simple intuitions and insights on otherwise complex problems. We now wish to illustrate the consequences of **not limiting** the strategy sets. Beyond the technical difficulties that this creates, we bring up to two issues: (i) it prevents us from separating between first order and second effects; (ii) it leads us to restrict attention or formulate informational assumptions in ways that ensure the model can be solved, with the consequence that it may sometimes shape our intuition in special ways. This section mostly attempts to illustrate the first issue.

We are interested in building a simple model of *cautiousness*. We have in mind that the decision problems we face vary a great deal, and that when faced with a specific one, an agent is uncertain about the appropriate decision that he ought to take. He may come up with some estimate of the appropriate decision, but inevitably, and despite being aware of it, his estimate will incorporate errors.

How then should he behave? Common sense would call for *cautiousness*, that is, not taking the estimate at face value. What cautiousness means varies across problems. In one dimensional problems in which overshooting the decision is more costly than undershooting it, caution would mean shading one's estimate.

And how cautious he ought to be should presumably depends on the magnitude of the errors that he tends to make. We are interested in building a simple model that captures these intuitions.

3.1 A simple class of decision problems.

For each decision problem that the agent may face, there is an appropriate decision that he ought to take. We denote by $s \in \mathcal{R}$ that decision and refer to it as the state. The agent however does not know the appropriate decision. Rather, he forms a noisy estimate of the state, which we denote x. We assume that

$$x = s + \alpha\varepsilon.$$

where ε is an error term (independent of the state) drawn according to some distribution h, $\alpha \geq 0$ parameterizes the magnitude of the error and is drawn from g (independently of s and ε), and s is drawn according to some distribution f. We let $\omega = (f, g, h)$ summarize the environment that the agent faces, and $\mathcal{P}$ the set of possible environments. As usual, we assume that the agent does not know the environment ω that is he facing.

We denote by $u(a, s)$ the utility that the agent gets from taking action a when the appropriate action is s. We define the difference $u(s, s) - u(a, s)$ as the loss function and assume below that it depends only the difference $z = a - s$. We let $L(z)$ denote the loss function.

Plausible alternatives. One (naive) alternative that the agent could follow would be to take his estimate at face value:

$$r_0(x) \equiv x$$

Other alternatives would have him pick an action that distorts, upward or downward, his initial estimate. We define rule r_γ as:

$$r_\gamma(x) \equiv x - \gamma$$

We examine three cases:

Case 1: $R_1 = \{r_0, r_{\gamma_0}, r_{-\gamma_0}\}$.

In that case, the set R_1 is a relatively simple set of rules. The agent is unsure as too whether his estimate is correct or not, and he learns from comparing the relative performance of these rules whether he ought to be cautious, and distort the decision he takes upward or downward. By doing so, and despite not knowing the distribution over problems that he face nor the distribution over the errors that he makes, he partitions the set of environments $\mathcal{P}$ into three classes, as a function of whether $E_\varepsilon L(\alpha\varepsilon)$, $E_\varepsilon L(\alpha\varepsilon - \gamma)$ or $E_\varepsilon L(\alpha\varepsilon + \gamma)$ is largest.

Case 2: $R = \{r_\gamma\}_{\gamma \in R}$

This is a more sophisticated set of rules. The agent obtains $E_\varepsilon L(\gamma + \alpha\varepsilon))$ by following rule γ, so the optimal choice of γ is the value γ^* that solves

$$E_\varepsilon L'(\alpha\varepsilon + \gamma^*) = 0.$$

For example, with a quadratic loss function, the cost from overshooting or undershooting the optimal decision is the same, and not surprisingly, one obtains $\gamma^* = 0$. With preferences that would generate a cost from overshooting the decision, for example:

$$\begin{aligned} L(z) &= -z^2 \text{ if } z \leq 0 \\ &= -z^2 - \beta z \text{ if } z > 0, \end{aligned}$$

the optimal decision is to choose γ^* that solves:

$$\gamma^* = \frac{\beta}{2} \Pr(\alpha\varepsilon > \gamma^*),$$

that is, shading one's estimate, with shading being stronger when $\alpha\varepsilon$ is more dispersed or when the cost parameter β is larger.

Case 3: no restrictions

This case corresponds to a standard model of decision under uncertainty, with no limits put on the set of rules that the agent considers: any function $r(x)$ is a possible alternative for the agent (case 3). In general, computing the optimal rule is difficult. To compute it in closed form, we turn to a special case, assuming that the loss function is quadratic, that s and ε drawn from normal distributions, say $\mathcal{N}(s_0, \sigma)$ and $\mathcal{N}(0, \sigma_\varepsilon)$, and that α is constant. The optimal rule is then

$$r^*(x) \equiv x - \gamma^*(x)$$

where

$$\gamma^*(x) \equiv E[\alpha\varepsilon \mid x] = -\frac{\alpha^2\sigma_\varepsilon^2}{\alpha^2\sigma_\varepsilon^2 + \sigma^2}(x - s_0).$$

In other words, the agent ought to exercise caution and not take his estimate at face value. Being aware that his estimate may be mistaken, he ought to correct it by $\gamma^*(x)$, taking into account the joint distribution over problems and errors (this is regression to the mean, s_0).

3.2 Discussion.

Tracking relevant aspects of ω. A central assumption of our approach is that the agent does not know the specific environment he is facing ω. Going from case 1 to case 3, one sees how enlarging the set of rules considered translates into a fine dependence of γ^* on the specific environment ω. Despite not knowing ω, the set of rules that he evaluates allows him to *track* some aspects of ω that are relevant to his decision problems. This tracking of ω is crude in case 1, but it nevertheless permits to understand under which circumstances he should (or should not) exert caution. Case 2 gives a more complete picture, as one can derive the magnitude of the caution he should exert as a function of the dispersion of his errors. Case 3 goes one step further, and our view is that this step is questionable. The agent behaves as if he knew ω in fine details. In particular, he behaves as if he was able to make inferences about the precision of his signals, or on the bias of his estimate, just based on the actual value of the estimate he gets. Our view is that it not plausible that an agent would do that, and that in most applications, this is at best a second order effect.

Signals about ω. We do not claim that an agent is unable to make some assessment about the accuracy if his estimate. Depending on the type of problem he faces, or the difficulty that he experiences in just coming up with an estimate, an agent could have an assessment of the precision of his estimate that differs across problems.

For example, if we are given one second to evaluate the square root of 4307, then whatever the estimate we come up with, we will probably be less confident about that estimate than if we are asked the square root of 10020. Besides, the value of the estimate we get will probably not be that relevant to assessing its accuracy.

Consistent with the example above, our view however is that this assessment of accuracy should be modelled as a distinct signal, for example a signal $\rho \in \{h, l\}$ correlated with α (say with h more likely when α is larger).

A rule would now be a function of x and ρ. For example one could define $r_{\gamma,\mu}$ as

$$r_{\gamma,\mu}(x, L) = x - \gamma \text{ and } r_{\gamma,\mu}(x, H) = x - \mu$$

and then

$$R = \{r_{\gamma,\mu}\}_{(\gamma;\mu)\in\{0,\gamma_0,-\gamma_0\}^2}$$

thereby allowing the agent to assess when it is useful to take into account his own perception of whether his estimate is accurate.

Diffuse priors. To conclude, let us mention some of the difficulties we point out with case 3 can be bypassed if one assumes a diffuse prior over states. Under that assumption, observing a particular estimate x does not provide information about the estimation error (sign and size) one makes, and the corrective term $\gamma^*(x)$ corresponds to that which obtains in the limit where σ^2 is infinite. Assuming a diffuse prior can thus be viewed as a device that enable the modeler to abstract from details which he thinks are irrelevant to the decision problem or strategic interaction being considered. We will get back to this when we discuss auctions.

3.3 Extension to choice problems.

We shall restrict attention to a class of choice problems between two alternatives, one of which, labelled a, can be easily evaluated by the agent, the other, labelled b, being not so easy to evaluate (by the agent).[8] We associate to each alternative $k = \{a, b\}$ a state s_k with the understanding that taking alternative k yields utility $u(s_k)$ to the agent.[9] If the agent knew the state $s = (s_a, s_b)$, the optimal decision would consist in choosing the alternative $k \in \{a, b\}$ for which s_k is largest.

[8] Extension to more general cases is easy.

[9] If k is lottery, s_k would thus correspond to the certainty equivalent associated with k.

However, the agent knows s_a but does not know s_b. Rather, he forms a noisy estimate of the state s_b.[10] Accordingly, we let $x = (x_a, x_b)$, and assume that $x_a = s_a$ and

$$x_b = s_b + \alpha_b \varepsilon_b.$$

where as before ε_b is an error term (independent of the state) drawn according to some distribution h, $\alpha_b \geq 0$ parameterizes the magnitude of the error and is drawn from g (independently of s and ε_b), and s is drawn according to some distribution f. We let $\omega = (f, g, h)$ summarize the environment that the agent faces, and $\mathcal{P}$ the set of possible environments. As usual, we assume that the agent does not know the environment ω that is he facing.

Plausible alternatives.

One (naive) alternative that the agent could follow would be to take his estimate at face value and pick the alternative that maximize x_k :

$$r_0(x) \equiv \arg\max_k x_k$$

Other alternatives would have him first distort, upward or downward, his estimate x_b by γ. Letting $\gamma_a = 0$ and $\gamma_b = \gamma$, we thus define rule r_γ as:

$$r_\gamma(x) \equiv \arg\max_k x_k - \gamma_k$$

To the agent, the performance of a rule r_γ is given by the expected utility he obtains under ω when he follows r_γ.[11] We denote is $v(r_\gamma)$:

$$v(r_\gamma) = E_\omega u(r_\gamma(x))$$

We examine the case where $R_1 = \{r_0, r_\gamma\}$.

The environment ω defines the size of the errors that the agent is making and the agent's risk aversion affects how costly these errors are. With sufficiently large errors and risk aversion, rule r_γ performs better, and the agent thus learns to be cautious of his estimate x_b. Otherwise, the agent prefers to take his estimate at face value.

Cautiousness vs ambiguity aversion.

Whenever r_γ is optimal with $\gamma > 0$, one could conclude that the agent is subject to ambiguity aversion. This effect however just reflects the effect of the agent's uncertainty as to the accuracy of his estimate, and it results from a standard notion of risk aversion combined with the noisiness of the estimation. Two risk averse agents with *different* dispersion of estimation errors but otherwise *identical* value estimate y and attitude toward risk would behave differently. So there could be differences in behavior from these two agents,

[10] An alternative assumption would be to assume that he forms an estimate of the difference $z = s_a - s_b$, as it may sometimes be easier to evaluate differences than each alternative separately.

[11] This is not a computation that the agent performs. It should be interpreted as the average utility that the agent experiences when confronted with ω for a long enough time.

despite the fact that they have the same preferences and the same value estimate y. Differences in behavior would be explained by differences in the strength or accuracy of their estimates.

Illustration.

Examples meant to illustrate ambiguity aversion are usually made in a context of choice between lotteries. The analysis applies to choice over lotteries, but it also applies to choices over sure outcomes, as the following example illustrates.

> *Option a:* you get $\$1M$ (for sure)
>
> *Option b*: you get x times $\$1M$ (for sure), where x is computed in the following way.
>
> Define p_n to be the n^{th} prime number, and set $x = 2^{\frac{p_{n+1}-p_n}{\ln p_n}}$ for $n = 1003$.

Which option would you take?

For people who know Riemann's hypothesis or the prime number theorem, that is, asymptotically, there are $\lg n$ numbers less than n that are prime, option b might look like a good bet, so that a reasonable estimate of x is 2. Of course, for people who also happen to know the twin prime conjecture, which states that there are infinitely many pairs of prime numbers whose difference is equal to 2, option b might look a bit risky, unless they happen to be able to compute p_{1003} fast.

So option a and b are two sure bets that do not call for the same decision as n varies. Yet our guess is that most people would have answered in the same way to that question, whatever the number $n > 1003$ would have be chosen to describe option B, and that they might be more inclined to choose B if they get to pick n (above 1003).

3.4 Summary

The underlying model of behavior proposed can be summarized as followed. There is an underlying state $s \in S$ that defines all what the agent would need to know to take an appropriate decision. The decision maker gets data that we summarize by a vector $z \in Z$. The process by which data is generated can be summarized by a joint distribution ω over (z, s). That process however is not known to the agent. He may form some ideas about the process that generates it or the accuracy of the estimates he forms, but this is formalized as a signal that is included in the description of z (see for example (x, ρ) in the discussion section). Note that the data the agent gets may include the fact that he feels unable to evaluate some alternatives, or to compare it with others.

Next the agent is endowed with a number of rules R. A rule $r \in R$ maps each $z \in Z$ to a decision a. Defining $u(a, s)$ as the utility that the agent derives from taking a (possibly inappropriate) decision a when the state is s, one can define the performance $v(r)$ associated with a rule r as $E_\omega u(r(z), s)$. We assume that the agent is able to pick the rule that has maximum performance:

$$r^* = \arg\max_{r \in R} v(r)$$

Looking at simple sets of rules, we easily obtain the insight that estimation errors combined with risk aversion is a source of caution. Caution may take different forms, depending on the class of problem the agent faces. In section 3.1, with larger costs from overshooting the caution will mean shading one's decision. In Section 3.2, caution means tilting the decision in favor of the one for which we know we are not making estimation errors. It also implies that more caution should be expected from less sophisticated agents (for whom estimation errors are likely to be larger), or when alternatives are more complex to evaluate or compare. (more abstract reasoning required, more complex computations, compound lotteries...).

Looking at simple sets of rules is *not crucial* however. Similar insights would obtain even if we did not restrict the set of rules. But they would be mixed up with another one: as in case 3 in section 3.1., the optimal rule would also take the joint distribution over (z, s) in a fine way, with the agent implicitly modifying his estimating to take into account the distribution over problems that he faces.

Beyond simplicity, the model of decision that we propose has an important difference with standard decision models. In standard models, modelling uncertainty proceeds exactly as above, with a joint distribution (ω) over signals (z) and states (s). Poorly informed agent are then modelled as agents for whom z is poorly correlated with s. These agents are poorly informed of s but yet they have (or end up behaving as if they had) a precise knowledge of ω. If the agent already finds is hard to figure what the state s is, it is presumably even harder for the agent to figure what ω is within $\mathcal{P}$, it seems odd to restrict our attention to cases where the agent would know ω for sure.

4 Information transmission

When uncertain about what decision we should take, we may seek or get advice or opinion from an external source. As a decision maker, we cannot be sure however that the advice we get is reliable, and we have to decide whether, when and how that advice should be taken into account. As an expert, we may have a vested interest in the decision about to be taken, or we may wish that our advice is followed, and these considerations may affect the advice we give.

These two strategic considerations are linked, as the extent to which an expert *bias* his advice should affect the extent to which the agent *trusts* the advice.

The seminal work on this issue is due to Crawford and Sobel, and the form of the equilibria that they obtain has shaped our intuition about how expert's information is transmitted in this type of strategic environment (i.e. partition equilibria).

We provide here a simple model of information transmission that accounts for the strategic issues described above. Our objective is twofold: (i) we wish to illustrate how our approach can be applied to information transmission; (ii) we wish to point out how *tractability motives* in the standard approach leads to special informational assumptions, that in turn shapes our intuition in a special way.

Before describing the strategic interaction, we start by describing *the class of decision problem that the decision maker typically faces in such interactions.*

4.1 The decision problem

We consider a class of decision problems where (i) the agent is uncertain about the appropriate decision he ought to take, and (ii) the agent receives additional information from a possibly unreliable source. The agent's initial uncertainty will be modelled exactly as in previous Section, as a noisy estimate of an underlying state. The additional information will be modelled as another possibly noisy estimate of that same underlying state.

We now turn to the formal model.

Preferences. For each decision problem that the agent may face, there is an appropriate decision that he ought to take. We denote by $s \in \mathcal{R}$ that decision and refer to it as the state, and choosing a decision a different than s involves a cost, which we assume to be quadratic: $u(a, s) = -(a - s)^2$.

Initial uncertainty. We assume that the agent does not know the appropriate decision. Rather, he forms a noisy estimate of the state, which we denote x. We assume that

$$x = s + \alpha\varepsilon.$$

where ε is an error term (independent of the state) drawn according to some distribution h, $\alpha \geq 0$ parameterizes the magnitude of the error and is drawn from g (independently of s and ε). We assume that s is drawn from a diffuse prior on R.

Additional information. We assume that the agent obtains another estimate, say y, of the appropriate decision. Such an estimate might come from an advice given by some other agent, or from taking a fresh look at the problem. For the sake of illustration, assume that

$$y = s + \alpha'\varepsilon'$$

where ε' is drawn from the same distribution h as before (independently of s and ε), and where $\alpha' \geq 0$ is drawn from R^+ according to some distribution g'.

Finally, we let $\omega = (g, g', h)$ summarize the environment that the agent faces, and $\mathcal{P}$ the set of possible environments. As usual, we assume that the agent does not know the environment ω that is he facing.

In general the agent ought to take this new estimate into account, because both x and y are informative of the true state. Exactly how he ought to take it in account may not be an easy task though. A fully rational agent who would know ω would have to compute the function

$$\phi(z) = E[\varepsilon \mid \alpha\varepsilon - \alpha'\varepsilon' = z].$$

The optimal decision would then consist in choosing

$$a^* = x - \phi(x - y)$$

Our view is that it is not plausible that a decision maker would know ϕ nor is it plausible, even in this relatively simple strategic environment (only one additional estimate y, no correlation across errors, both errors are unbiased) that the decision maker would learn to play optimally, that is, that he would find out the function $a^*(x, y)$ that would maximize welfare among all possible functions – a hopeless task.

Rather, we find it more plausible that an agent would try to look for an optimal strategy among a limited number of alternatives, and our objective below is to suggest such alternatives.

A minimal set of alternatives.

A simple rule of thumb that the agent might use would be to ignore the additional estimate, hence choose $a = x$. Another one would be to trust it, hence choose $a = y$. . We shall refer to these alternatives (or rules) as I (ignore) and T (trust). We denote by $v(r)$ the value associated with an alternative $r(\in \{I, T\})$. Letting $\sigma^2 = E\varepsilon^2$, we have:

$$\begin{aligned} v(I) &= -E\alpha^2\ \sigma^2, \text{ and} \\ v(T) &= -E\alpha'^2\ \sigma^2 \end{aligned}$$

Equipped with this simple set of possible decision rules, the agent thus partitions the set of environments into two subsets:

$$\mathcal{P}_T = \{\omega \mid E\alpha'^2 < E\alpha^2\} \text{ and } \mathcal{P}_I = \{\omega \mid E\alpha'^2 > E\alpha^2\}.$$

That is, he ends up trusting the second estimate, or ignoring it, depending on how $E\alpha^2$ compares to $E\alpha'^2$.[12]

A larger set of alternatives.

Rather than being either fully trusting or fully distrusting, the agent might try to *discriminate* between reliable and poorly reliable signals. Since a large value for the difference $\mid y - x \mid$ may be an indication that y is poorly informative, one plausible rule of behavior is that the agent ignores the signal if $\mid y - x \mid$ is above a threshold, say k, and that he trusts the estimate y otherwise:[13] We shall call r_k this threshold rule.

[12] One possible interpretation is that x coincides with first hand information on a particular case, while y corresponds to aggregate data on similar cases. The agent ends up trusting more x than y when y is noisier estimate of each specific case. This could account for the difficulty that agents apparently face in taking into account aggregate data.

[13] What would determine the agent's threashold is somewhat arbitrary. For an agent who has a precise idea of the distribution h, these threashold could be parameterized by the probability that $\mid y - x \mid$ be consistent with y being a perfect estimate, i.e. the value of $\Pr\{\mid \varepsilon \mid \geq k\}$.

For a given $k_0 > 0$, the rule r_{k_0} is still a crude rule, but it may improve the agent's welfare over both $I(= r_0)$ and $T(= r_\infty)$. A collection of rule that the agent might consider using is thus

$$R = \{r_k\}_{k \in \{0, k_0, \infty\}},$$

and as above, for any ω, one can compute the value $v(r_k)$ associated with rule or alternative k. The set of environment can thus be partitioned into three sets $\{\mathcal{P}_k\}_{k \in \{0,k_0,\infty\}}$ depending on which rule is optimal.

4.2 Strategic information transmission

We consider the usual strategic environment in which one agent makes an observation $y_0 = s$ and sends a signal y to the decision maker, and in which the agent is biased compared to the decision maker: his preference are quadratic with a bliss point at $s + b$.

The decision maker has an initial estimate $x = s + \varepsilon$ of the true state, and he receives the signal y sent by the agent. Based on these two estimates x and y, he take a decision. The set of rules R that he considers is defined as above, and we will illustrate our approach by considering two cases: $R = \{I, T\}$ or $R = \{I, T, r_{k_0}\}$.

From the perspective of the sender, a plausible rule is that he bias his estimate y_0 by a constant amount Δ when he makes a report, thus choosing $y = m_\Delta(y_0)$ where

$$m_\Delta(y_0) = y_0 + \Delta.$$

We will study the case where the set of rules that the agent considers is:

$$M = \{m_\Delta\}_{\Delta \in \mathcal{R}}.$$

A rule profile then consists of a pair $(m_\Delta, r) \in M \times R$. To each such pair, we may associate values for each player, and for any given set of rules M and R, we may define an equilibrium accordingly.

Example 1: $R = \{I, T\}$

There exists an equilibrium with no information transmission where Δ is large and $r = I$. So long as $b^2 < \sigma^2$ there also exists an equilibrium where the agent chooses $\Delta = b$ and $r = T$. In words, *information can be transmitted so long as the agent's bias is not too large compared to the noisiness of the decision maker's signal.*

Example 2: $R = \{I, T\}$, sender observes a noisy signal $y = s + \alpha\varepsilon'$.

There exists an equilibrium with information transmission so long as $\alpha^2\sigma^2 + b^2 < \sigma^2$. Information can thus be transmitted is the agent's bias is not too large, but the constraint is tighter because the agent's signal is not that good.

Example 3: $R = \{I, T, r_{k_0}\}$

For a range of k_0, there exist equilibria with information transmission where the decision maker chooses r_{k_0} and the sender chooses some $\Delta \in (0, b)$. In equilibrium, the decision maker ignores the agent's advice when the realization of ε is too low (because then the difference $y-x=\Delta-\varepsilon$ is large), and the agent chooses a bias lower than b because of the risk that his signal will be ignored.

This trade-off seems realistic: when we give some advice, we would like to influence the decision in a particular way, yet we often feel that if our recommendation is too extreme, it won't be taken into account. When used in equilibrium, rule r_{k_0}, which attempts to discriminate between reliable and poorly reliable signals, creates an incentive device that disciplines the agent (to some extent).

Example 4: $R=\{I,T,r_{k_0}\}$ and $y=s+\alpha\varepsilon'$

For a range of k_0, there also exist equilibria with information transmission where the decision maker chooses r_{k_0}. Compared to the previous example, the sender chooses a less cautious strategy $\Delta'>\Delta$. The reason is that he has lesser control over the chance that his advice will be ignored.

One implication of this equilibrium is that more poorly informed senders bias their reports to a greater extent.

4.3 Comparison with standard models.

We model uncertainty as a noisy signal x correlated with an underlying state s, where $x=s+\alpha\varepsilon$, with s drawn from a diffuse prior. A more poorly informed agent corresponds to one for which the estimation error has larger variance.

In Crawford Sobel (CS), there is an underlying distribution, say f, over states, and the decision maker is uninformed in the sense that he receives no signals about the state. In principle, one could assume that the decision maker receives a signal x correlated with the state. Within the standard approach, this is a hard problem to solve (because the equilibrium strategy of the decision maker can be a complex function of his signal x), and much attention has been devoted to the case where the decision maker receives no signal.

These modelling choices have important consequences on the nature of informational asymmetries, on the structure of the equilibria that are found, and on the nature of communication that takes place.

Information asymmetries. In our approach, the agent is uncertain about the estimation error that the agent makes, while in CS, both players know f so the agent knows the exact way in which the decision maker is uncertain about the true state. We find the latter to be a peculiar assumption, as the following jar example illustrates:

> **Jar example**. The state s refers the amount of money contained in a jar, and a decision maker has τ seconds to provide an estimate x of the amount of money contained in a jar, assuming a quadratic loss function for the decision maker when x differs from s.

In this example, it seems reasonable to expect that lower values of τ will induce noisier estimates from the decision maker. Taking the perspective of an agent who would know s, we can speculate that lower values of τ will give rise to noisier estimates. However, reducing τ should not make it easier to predict the decision marker's estimate.

In CS, the prediction goes the other way around. When the decision maker is "completely" uninformed, the agent has no difficulty predicting his decision. The reason is that the decision maker is uninformed about the actual state, but perfectly informed about the distribution over states.

Structure of equilibria and nature of communication.

In CS, equilibria take the form of partition equilibria: in equilibrium, the agent sends one among a bounded number messages, and the decision maker interpret the message according to some common grid of interpretation that *pre-exist* to communication, and that is tailored to the particular distribution over states and to the particular preferences of the agent and the decision maker.

Two features of these equilibria are implausible.

One feature is that the players would come up to the interaction with a (common) grid of interpretation that pre-exists the interaction. Coming up with an a priori notion of what high, low, very high or very low means requires some reference point that the current state can be compared to, or some knowledge of the distribution over problems that we are facing. In CS, common knowledge of the distribution over state provides such a common reference point. There are many instances however where we attempt to communicate our expertise without knowing much about the distribution over problems of similar nature that we might face, or when there is no obvious reference point to which the current state can be compared to. When reporting information on the volume of oil present in an oil field, there may not be an a priori notion of what high very high, low or very low means, just because the types of oil field we face vary and do not share the same characteristics. Yet meaningful communication is possible, even if we believe that the expert has a small vested interest in our bidding high for the oil field.

Another implausible feature is that the grid is finely tailored to the specific strategic environment (distribution over states and preferences) that players face. Rather common sense would suggest that interpretations follow general rules, and that the strength of interpretation precisely comes from the fact that they are useful even when the environment changes.

One could argue that the model should not be interpreted too literally, and that it is not meant to be an actual description of how players behave in such interactions. It would just capture the idea that the agent ought to be ambiguous when communicating. We have two comments:

(i) That conclusion relies both on the agent's perfect prediction of the decision maker's decision, and on the decision maker's extreme ability to decode the sender's message; We already commented on the first aspect (see the Jar example), and on the second one, one can think of many reasons as to why the decision maker's ability to decode the sender's message would be limited, an obvious one being that he is unsure about the environment ω that he is facing. Our model captures that later aspect by assuming that decision maker only attempts to compare few strategies (which of course limits his ability to decode the sender's message);

(ii) Even when the decision maker has a high ability to decode the agent's message, the prediction that ambiguity would take the particular form of a partition equilibrium is questionable. If one assumes that the agent only compares two alternatives, say $M = \{m_0, m_\Delta\}$, there exist equilibria where the agent randomizes between these two alternatives.

4.4 Extensions

One possible extension would be to consider a class of decision problems in which the decision maker receives signals from two (possibly unreliable) sources (in addition to own estimate x). Such an extension would then permit us to deal with strategic situations where the decision maker has hired two experts.

For example, we could assume that in addition to x, the decision maker receives signals y_1 and y_2. In the spirit of the alternative rules defined above, plausible rules could be that: (i) the agent ignores both signals; (ii) the agent trusts the signal that is closest to his own signal; (iii) the agent trust the signal that is closest to his own signal only if the two signals are not too different.

The following strategic situations, involving two senders, could then be analyzed:

- two senders having same bias,
- two senders, same bias but noisy estimate y_1 and y_2,
- two senders, opposite bias,
- two senders opposite bias and noisy estimates,

Other extensions involving only one sender could also be analyzed:

- one sender with multiple decision makers,
- one sender with a population of decision maker and a majority rule to take a decision,
- one sender who has a cost of sending a message m that differs from y ($c(m, y) = \gamma \mid m - y \mid$)

5 Auctions

When participating in a first price auction, one has to calculate how to shade one's bid. We expect that shading will be small when competition is tough, and more substantial when competition is weaker. Ideally, we would like a simple theory that captures what is meant by competition being tough or weak, and that allows us to understand how bidding, welfare and seller revenue vary across formats (first price and second price) as one introduces asymmetries between bidders, signals about bidder's comparative advantage, and uncertainty about own valuation.

We provide here a simple model of auctions. Our objective is again twofold: (i) we wish to illustrate how our approach can be applied to auctions; (ii) we wish to point out how the standard approach has shaped our intuition in special ways.

The change in perspective that we propose can be summarized as follows:

> I might go to an auction at Sotheby's, see a painting and decide that my value for a painting is \$1100, but have little idea of where my valuation stands in relation to the valuations of other bidders. That is, I might think it equally likely that my valuation is the highest, second highest, or lowest; in other words, my valuation gives me little guidance in predicting my rank in the valuations. Had my value of the painting been \$1200 rather than \$1100, would my prediction be different?

We shall take the perspective that it is hard for an individual to sort out whether the basis of his higher value was common to all bidders or was idiosyncratic to him, and that a plausible model of auctions should reflect that difficulty.

This is a *change* in perspective because standard auction theory implicitly poses no limits on the agent's ability disentangle common and idiosyncratic elements.

The problem is similar to that explored in Section 3 (Cautiousness). In that Section, the agent gets an estimate of the value of an alternative; say $x = s + \varepsilon$. Ideally, the agent would like to know his estimation error ε, but we should not expect him to be able to make precise inferences about the size or sign of ε based on the estimate x he forms. There too, we took the perspective that a plausible model of decision making should reflect that difficulty, contrasting with standard models of decision making that implicitly allow agents to use their knowledge of the joint distribution over problems and errors to make inferences about their current estimation errors.

5.1 The bidder's decision problem

For each auction problem that the agent may face, we assume the agent has a valuation v_i for the object, and that he gets a possibly noisy estimate of that valuation which we denote x_i. We assume that

$$v_i = s + \theta_i$$

and

$$x_i = v_i + \alpha_i \varepsilon_i.$$

s represents a component common to all bidders' values and θ_i represents an idiosyncratic component of i's value (drawn from g independently of s). ε_i is an error term (independent of v_i) drawn according to some distribution h, and $\alpha_i \geq 0$ parameterizes the magnitude of the error.

In a first price auction, the agent bids a price which we denote b_i, and he wins $v_i - b_i$ when b_i exceeds some price p, and 0 otherwise. Letting ρ denote the joint distribution over states and prices, we can characterize by $\omega = (\rho, g, h, \alpha_i)$ the environment that the agent faces.

Plausible bidding rules. The data that the agent gets is x_i, and we consider family of bidding rules based on x_i. For any γ, we define rule r_γ as:

$$r_\gamma(x) \equiv x - \gamma.$$

Defining

$$\phi(\gamma) \equiv \Pr(x_i - \gamma > p),$$

the expected payoff that the agent gets from using rule r_γ is:

$$v(r_\gamma) = \gamma\phi(\gamma)$$

If R is the set of rules that the agent considers, he thus picks the rule maximizes $v(r_\gamma)$, that is he picks:

$$r_{\gamma^*} = \arg\max_{r_\gamma \in R} v(r_\gamma)$$

Throughout most of this section, we shall assume that $R = \{r_\gamma\}_{\gamma \in \mathcal{R}}$ (but a coarser set of rules could be considered).

Using signals about ω. Standard analysis would not restrict the set of bid functions, and we would look for a bid function $r^*(x)$ that fully exploits the environment ω. We do not claim that an agent is unable to make some assessment related to ω, for example about the size of $v_i - p$. Depending on the type of competitors he faces, or because he also forms a (noisy) estimate of others' valuations, he might come up with the idea that he is in a strong or weak position.

Following our discussion in Section 3, our view is that this assessment of ω should be modelled as a distinct signal, for example a signal $\rho \in \{h, l\}$ correlated with $v_i - p$ (say with h more likely when α is larger).

A rule would now be a function of x and ρ. For example, one could define $r_{\gamma,\mu}$ as

$$r_{\gamma,\mu}(x, L) = x - \gamma \text{ and } r_{\gamma,\mu}(x, H) = x - \mu$$

and then

$$R = \{r_{\gamma,\mu}\}_{(\gamma;\mu) \in \mathcal{R}^2}$$

thereby allowing the agent to assess when it is useful to take into account his own perception of strength.

Optimal bidding rule.

Returning to the case where only x is observed and $R = \{r_\gamma\}_{\gamma \in \mathcal{R}}$, the optimal bidding rule is characterized by a single parameter γ^* that measures the extent of shading, and that is given by the following first order condition:

$$\gamma = \frac{\phi(\gamma)}{-\phi'(\gamma)}.$$

Shading is larger when the agent has higher chances of winning, or when the distribution over $x_i - p$ is more dispersed (because then $\mid \phi' \mid$ is then smaller).

Noisy estimates and the winner's curse.

When estimates are noisy (positive α), the bidder runs the risk of getting the object only because $\alpha\varepsilon_i$ was positive and high). The bidder should thus exert some caution, which means *shading* one's bid (increasing γ).

In other words, *noisy estimates* make a bidder subject to *a winner's curse*, and an optimal reaction to that possibility is cautiousness – or shading.

5.2 The strategic interaction

We now assume that there are n bidders and that each considers a set of bidding rules $R_i = \{r_\gamma\}_{\gamma\in\mathcal{R}}$. For any $\gamma = (\gamma_1, ..., \gamma_n)$, each rule profile $r_\gamma = (r_{\gamma_1}, ..., r_{\gamma_n})$ induces an expected payoff which we denote by $v_i(\gamma_1, ..., \gamma_n)$. We have:

$$v_i(\gamma) = \gamma_i \varphi(\gamma) \text{ where } \varphi(\gamma) \equiv \Pr_\omega(x_i - \gamma_i \geq \max_{k\neq i} x_k - \gamma_k)).$$

An equilibrium is then defined in the usual way. We illustrate the approach with some equilibrium computation.

Case 1. *Symmetric case, no noise (*$\alpha_i = 0$*).*

We look for a symmetric equilibrium in which all bidders pick rule r_{γ^*}. Define

$$\phi(y) = \Pr(\theta_i - \max_{j\neq i} \theta_j > y).$$

We have

$$v_i(\gamma_i, \gamma^*) = \gamma_i \phi(\gamma_i - \gamma^*)$$

implying that γ^* solves:

$$\gamma^* = \frac{\phi(0)}{-\phi'(0)} = \frac{1}{-\phi'(0)n}.$$

In other words, equilibrium shading is driven by the number of bidders (through the expected chance of winning) and the dispersion of idiosyncratic terms.

Case 2. *Symmetric case, no noise (*$\alpha_i = 0$*)*, information about ranking

Case 3. *Asymmetric case (player 1 has a comparative advantage:* $v_i = s + \theta_i + \mu$*)*

Case 4. *Symmetric case, noise (*$\alpha_i > 0$*).*

5.3 Discussion.

In standard models, auctions are hard problems because one attempts to derive an optimal bid function within the whole set of possible functions. In addition to being hard problems, auction models assume that bidders can tailor their bid functions to the specific environment (i.e., the joint distribution over values and estimates, say ω) they face. Our view is that it is implausible that bidders would do that.

Beyond complexity and implausibility, we wish to argue to these modelling choices have either made it difficult to get intuition about the problems, or that they have shaped intuition in questionable ways.

Parsimony.

In standard models, bidders come to an auction knowing the joint distribution over values and estimates of their value, or they behave as if they knew the distribution. Inevitably, this means that each distinct value estimate comes with a distinct assessment of how their valuation compares to others, and in the standard approach it also means that the agent will fully exploit these fine assessments.

We have two objections, related to the two roles that we think theory should accomplish: help us shape our intuition about relevant strategic issues; help us predict how players behave in particular interactions.

To accomplish the first role, one would like a theory that enables us to focus on first order effects, and this may be more easily accomplished in a more parsimonious model that does not rely on highly sophisticated and correct assessments of the environment one is facing.

To accomplish the second goal, one would like a theory that provides a *compelling picture* of how bidders behave. Bidders get information about their valuations, and possibly some information about the specific environment ω they are facing: how then is bidding accomplished? How is information used and/or combined? If one views (as we do) that signals about the environment have greater chances of being noisier than signals about own valuation, then we should not be surprised that agents disregard signals about the environment, and we should study models that do not rely or hinge on a specific environment or the precise knowledge of that environment.

The common/private values distinction.

The literature has used the private/common value distinction as a central dividing line in the study of auctions (and other mechanism design related issues). With this division, the private value case is perceived as a *simple* environment, while the common value case is perceived as a more elaborate strategic situation. We argue below that this perception is misleading: from a bidder's standpoint (as opposed to the modeler's standpoint), an alternative dividing line may instead be *whether or not he gets an estimate related to his rank.*

Consider an extension of our model in which in addition to x_i, bidders each receive a (possibly noisy) estimate of s, that we denote y_i:

$$y_i = s + \beta_i \xi_i,$$

where ξ_i is a noise term drawn independently of s.

The standard approach would ask bidders to look for an optimal rule $r(x_i, y_i)$ within the set of all possible functions. To the modeler the limit case $\beta_i = 0$ is simpler because he can use the structural assumptions to figure that there is an equilibrium where bidders use rules that are function of the difference $x_i - y_i$. The private value case precisely corresponds to that case (with $\alpha_i = 0$).

As a participant in the auction however, determining how to use and/or combine the two signals x_i and y_i is a more complex task than just determining how to use a single signal x_i. It does not necessarily mean that y_i will be ignored, but we would like a plausible way that y_i could be taken into account. The following example suggests one possibility.

A class of plausible rules: The agent partitions signals into two classes ($\rho_i \in \{h, l\}$), and views the event $x_i > y_i$ as a good signal about own rank ($\rho_i = h$) and the event $x_i < y_i$ as bad signal about rank ($\rho_i = l$). He then adjusts shading (γ_l or γ_h) to the signal $\rho_i = h, l$ that they receive. Though agents are unsure of the reliability of y_i, their optimal strategy (and the difference $\gamma_h - \gamma_l$) will reflect how reliable signal y_i is.

How does adding signal y_i modify the strategic interaction? From the agent's perspective, the essential difference with our basic model (case 1 above) is that he now gets a (noisy) signal about rank that he may wish to exploit, and the class of rules proposed above suggest a way he can do that, thereby turning our basic model (case 1) into a slightly more elaborate one (case 2 above).

In comparison, the standard private value model requires substantially greater sophistication from the agent,[14] as it assumes a fine understanding of how variations in $y_i - x_i$ translate into variations in ranking and value differences.

Revenue ranking

How do auction formats (first and second price) compare? Revenue comparisons often begin with the following observation. Second price auctions do not look like good mechanisms because they seem to give rents to winners. In first price auctions however, bidders shade their bids in comparison to what they would do in a second price auction, so the effect on revenues is unclear.

Revenue equivalence results have been obtained in the independent private value case, and these results can be seen as a nice illustration of the above phenomenon. It is not clear however that these results can be viewed as more than an illustration. The complexity of the approach makes it difficult to get economic intuition for these results, and the sophistication required by the agents makes their relevance unclear. The same comment applies to more recent revenue ranking results in correlated value settings.

In contrast, under our approach, everything that is relevant in the comparison between the two auction formats is contained in the function ϕ (that is derived from the distribution over the idiosyncratic elements) and revenue comparisons can be made based on the shape of that function: first price auctions will yield greater (expected) revenue than second price auctions when ϕ exhibits low dispersion at 0 (implying small shading) and a reasonably fat tail (implying that substantial rents are given away in the second price auction).

Noisy/perfect estimates distinction

In terms of sophistication, we have seen that a relevant distinction is whether the agent receives (and uses) information about rank. In terms of bidding behavior, a relevant distinction is whether the agent receives noisy or perfect estimates of valuations. As case 4 illustrates, noisy estimates induce cautious behavior, and caution translates into increased shading, for fear that one tends to win more often when valuation is overestimated.

[14] This is in addition to assuming that x_i and y_i are perfect estimates (i.e. $\alpha_i = \beta_i = 0$).

This effect is greater with more competition as the costs induced by over and under estimation are asymmetric: underestimation is costly because you may lose the object, but the cost is small when competition is fierce because winning; in contrast, overestimation can be very costly.

In other words, estimation errors can lead to a winner's curse, and caution is an appropriate response. As we pointed out in Section 5.1., this same phenomenon arises even in a single bidder environment.

In comparison with the standard literature, our model provides intuition about the winner's curse in a different way. In auctions with common values, our usual understanding is that winning conveys information about the *common value part* of one's valuation, and that to avoid the winners' curse, bidders ought to take into account that information. This is not true in our setting. Winning conveys information about s because

$$E[s \mid x_i] \neq E[s \mid x_i,\ x_i > \max x_j].$$

Yet that information is not used in bidding because $x_i = v_i + \alpha_i \varepsilon_i$ so

$$E[v_i \mid x_i,\ x_i > \max x_j] = x_i - E[\alpha_i \varepsilon_i \mid x_i > \max x_j].$$

This explains why in our setting, bidders in equilibrium only take into account the information about the *estimation error* conveyed by winning (rather than the information about the common component s).

Existence issues and asymmetric settings

As for revenue ranking, standard auction models provide little economic intuition as to when existence of pure strategy equilibria fails, and few predictions about the effect of asymmetries. These issues are more easily addressed in our setting.

6 Beliefs

As decision makers, beliefs seem to matter. We forge an opinion about the individuals with whom we are about to deal with, we get a conviction that a person is guilty or innocent, though we may also express doubts, we get a sense that a situation or an action is risky, thinking that there is a non trivial chance that an unfortunate event arises, and we also get a sense of whether our answers to a test were correct, or not.

How are these opinions forged? How are they used in decision making?

We forge opinions as a function of the observations we make, and our analysis of reputation provides an example in this respect. In that example, we focused on observations related to the satisfaction experienced in the interaction. But other types of signals could be used, related for example to the behavior that one observes. And although we do not observe the exact circumstances that led to the behavior we observe, we often draw inferences from that observation, and theses inferences in turn contribute to forging our opinions.

Once formed, these opinions can be used as a guide to decision making, as one may *condition* behavior on the opinion one gets. To the extent these opinions do not end up being independant of the decision problem being faced, they should help take better decisions.

We provide here a model of decision making where decisions are indeed shaped or guided by the opinions we form, and where the opinions we form are a product of the observations and inferences that we make.

Our objective is twofold: (i) we wish to illustrate how our approach can be applied to more general decision problems; (ii) we also wish to point out that the standard approach fails to provide a compelling theory of behavior.

We start by describing *a class of decision problems that involves information aggregation.*

6.1 Decisions problems with information aggregation

We consider a class of decision problems in which an individual processes a stream of informative signals before taking a decision.

Decision problem. For each problem considered, there are two possible states, $\theta = 1, 2$. The decision is a choice between two alternatives, $a \in \{1, 2\}$. Taking action $a = \theta$ is the appropriate action when the state is θ, and it yields payoff v_θ. Taking an inappropriate action yields 0.

The signals received. The agent does not know the state, and before deciding which action to take, the agent receives a random number of signals each correlated with the true state. Formally the set of possible signals for this problem is denoted X, and a signal is denoted x. Signals are received over time, at date $t = 0, 1, ...$, and the decision must be taken at some random date $\tau \geq 1$ drawn from some distribution g, and we denote by $\mathbf{x} = (x_0, ..., x_\tau)$ the stream of signals received. The probability that the true state is θ is π_θ, and given state θ, the signals are drawn independently of one another, each from some distribution $f(. \mid \theta)$. The environment that the agent faces is thus summarized by $\omega_0 = (v, \pi, X, f, g)$.

From a standard perspective, the agent's objective is to use the signals he gets optimally, that is, his task is to find the function $r(\mathbf{x})$ that maximizes expected utility accross all possible functions. One cannot imagine that an agent could perform this task directly, but standard decision theory suggest an indirect way: for any given ω_0, the agent ought to use Bayes rule to compute the belief $\beta_\theta(\mathbf{x}) \equiv \Pr_{\omega_0}(\theta \mid \mathbf{x})$, and then choose the action that maximize expected utility given β.

As already emphasized, our view is that either way is highly implausible: the agent cannot know ω_0, and he cannot evaluate the performance of each possible rule $r(.)$. Rather our view that these performance assessments can only be carried for a limited number of rules, and our objective below is to suggest plausible rules.

Inferences

The signals x are specific the problem considered, and we look for rules that apply accross problems. So rules that take $\mathbf{x}$ as an argument are not good candidates. Rather we shall assume that the agent use signals to form inferences about which state holds. We shall later see how inferences can plausibly be used to form opinions.

We assume that following each signal received x, the agent makes some inference $\widetilde{\theta} \in \{1, 2\}$ as to whether the signal is evidence of state 1 or 2, and that he also forms an estimate $\widetilde{l}$ of the strenght of the evidence. In other words, he gets another signal

$$\widetilde{z} = (\widetilde{\theta}, \widetilde{l})$$

which we assume to be correlated with ω_0 and x.

To fix ideas, we describe two cases, one in which the agent makes perfect inferences, and one in which the agent makes noisy inferences. First define the state $\theta(x)$ that has highest likelihood given x, namely:

$$\theta(x) = \arg\max_{\theta} f(x \mid \theta),$$

and the likelihood ratio $l(x)$ defined by:

$$l(x) = \frac{f(x \mid \theta = \theta(x))}{f(x \mid \theta \neq \theta(x))}.$$

The state $\theta(x)$ is the state for which signal x provides support, and the likelihood ratio $l(x)$ provides a measure of the strength of the evidence in favor of $\theta(x)$. We let $z(x) = (\theta(x), l(x))$

Perfect inferences correspond to the case where $\widetilde{z}(x) = z(x)$. An example of noisy inferences is:

$$\widetilde{\theta}(x) = \theta(x) \text{ and } \widetilde{l} - 1 = \widetilde{\mu}(x)(l(x) - 1)$$

where $\widetilde{\mu}(x)$ is a random variable. More generally, we let h denote the process by which inferences are generated, and we let

$$\omega = (\omega_0, h)$$

denote the environment that the agent faces.

Mental systems and plausible rules.

The agent makes a number of inferences $(\widetilde{z}_0, ..., \widetilde{z}_\tau)$ and we wish to define plausible rules that the agent could follow to take a decision. A naive plausible rule would consist in ignoring these inferences, and choose an action a independently of the inferences made.

A more sophisticated class of plausible rule would be to assume that the agent has a limited number of *states of mind*, and that each inference the agent makes (possibly) triggers a change in his state of mind. Formally a state of mind is denoted $s \in S$, where S is a finite set. For any signal x received, changes in state of mind depend on the inference $\widetilde{z}$. We denote by T the transition function:

$$s' = T(s, \widetilde{z}).$$

The case where S is a singleton corresponds to the naive case. The case with more states of mind correspond to that described informally in introduction, where opinions are forged as a function inferences made.

We shall refer to the triplet

$$\mu = (S, T, s_0)$$

as a mental system: it consists of a set of states S, transitions T, and an initial state s_0. Given μ, a strategy $\sigma \in A^S$ can be defined as a function that maps mental states to actions in $A = \{1, 2\}$, and it generates a plausible rule $r_{\mu,\sigma}$ that maps $\widetilde{z}$ to $\{1, 2\}$. We denote by R_μ the family of plausible rules that are induced in this way:

$$R_\mu = \{r_{\mu,\sigma}\}_{\sigma \in A^S}$$

Example 1: (a simple mental system) The agent may be in one of three states of mind $\{s_0, s_1, s_{-1}\}$. His initial state is s_0. Define A_0^+ as the event $\{\widetilde{\theta} = 1\}$ (evidence in favor of state $\theta = 1$), and A_0^- as the event $\{\widetilde{\theta} = 2\}$ (evidence in favor of $\theta = 2$). Transitions are defined as follows:

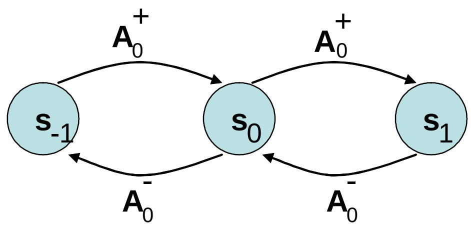

Figure 2: Transition function

Casual beliefs and biased beliefs

As the above figure illustrates, if the agent finds himself in state s_1 when he is called upon to make his decision, there may be many histories that have led to his being in state s_1: the mental system μ is simply a device that generates a particular pooling of the histories that the agent faces when making a decision.

The state the agent finds himself in upon taking a decision reflects the agent's current opinion or *casual belief.* In example 1 above state s_1 reflects the agent casual belief that state $\theta = 1$ holds, while state s_{-1} reflects the agent's casual belief that state state $\theta = 2$ holds. The belief is casual in the sense that no particular posterior probability can be assigned to a particular state, as this probability would depend on the environment ω that the agent actually faces (which the agent does not know).

Given θ and ω, a mental system μ leads to a probability distribution over the state that the agent is in upon taking a decision, and we denote by $\phi_{\theta,\omega,\mu}$

that distribution. Calling S_1 the set of mental states casually associated with state $\theta = 1$ (and similarly for state 2), an agent is said to hold *correct casual beliefs* when $s \in S_k$ whenever $\theta = k$. *A mental system is said to be unbiased* if the agent has more chance of holding correct casual belief, that is:

$$\phi_{1,\omega,\mu}(S_1) > \phi_{1,\omega,\mu}(S_2) \text{ and } \phi_{2,\omega,\mu}(S_2) > \phi_{2,\omega,\mu}(S_1).$$

Welfare.

As for welfare, for any given (ω, μ),we can compute the expected utility associated with each rule r, that is $u_{\omega,\mu}(r)$, and let

$$v_\omega(R_\mu) = \max_{r \in R_\mu} u_{\omega,\mu}(r)$$

6.2 Analysis

From the agent's perspective, a mental system gives rise to a particular opinion or casual belief. Two questions can be asked. Is a mental system helpful? Does it give rise to biased beliefs?

Welfare. A mental system μ (with at least two states) provides a signal correlated with the true state, so a mental system cannot hurt and it sometimes help take better decisions: it improves welfare over the naive set of rules. As the number of states increase, or as the accuracy of the inferences improves, welfare may improve further. But if there are more states, this comes at the cost of more difficulties in evaluating alternative rules.

When do we expect unbiased casual beliefs?

In what follows, we discuss the case where each signal x consists of a sample of observations.

Large sample size.

If with high probability the sample size is large enough, then

$$p_1 \equiv \Pr(\theta(x) = 1 \mid \theta = 1) > 1/2 \text{ and } p_2 \equiv \Pr(\theta(x) = 2 \mid \theta = 2) > 1/2 \qquad (1)$$

So to the extent that inferences are not too noisy (i.e. $\widetilde{\theta}(x)$ strongly correlated with $\theta(x)$), even a simple mental system such as that in example 1 will induce unbiased casual beliefs.

If inferences are noisy in the sense that $\widetilde{l}(x)$ is poorly correlated $l(x)$, then a more sophisticated mental system may lead to both poorer performance and biased beliefs

Small sample size.

Then a mental system will still improve welfare, but there is no guarantee that inequalities (1) both hold, and even the simple mental system in example 1 then leads to biased beliefs

More than two states.

All our analysis extends to cases where there are more than two (true) states. Define inferences as before, with $\theta(x) \equiv \arg\max_\theta f(x \mid \theta)$ and assume the agent makes correct inferences on the true state, i.e. $\widetilde{\theta}(x) = \theta(x)$. Then, if sample size is small, there are simple problems for which

$$p_k = \Pr(\theta(x) = k \mid \theta = k) = 0 \text{ for some } k,$$

which implies that the agent never finds evidence in favor of state k.

6.3 Evolution of mental systems

For any mental system μ, and any given ω, we have defined the value associated with the set of rules R_μ, and we denoted it $v_\omega(R_\mu)$. Another mental system μ' will in general induce another set of rule $R_{\mu'}$, and for a specific ω, that could induce a higher welfare. However the task of tailoring μ (say, the transitions) to a specific environment ω so as to maximize welfare seems misguided, as the purpose of these mental system is to enable the agent to deal with a whole range of environment.

Nevertheless, if there are modifications of the mental systems that increase welfare accross many problems, then we may expect that evolution selects such modifications. The main insight of Compte and Postlewaite (2010) is that ignoring weak evidence is a modification that indeed improves welfare (across many problems).

The formal result, which we describe below assuming perfect inferences, can described as follows. For any mental system μ, we consider the modification of μ in which all weak inferences (that is, inferences $\widetilde{z} = (\widetilde{\theta}, \widetilde{l})$ such that $\widetilde{l} < 1 + \beta$) are ignored. We denote by μ^β that modified mental system, and denote by A the event where the two mental systems induce different transitions. We consider distributions over environments $\kappa \in \Delta(\mathcal{P})$ that put weight on all realizations of (v_1, v_2) in $[0,1]^2$. Our main result in Compte and Postlewaite (2010) is that for any such distribution κ,

$$E_\kappa v_\omega(R_{\mu^\beta}) - E_\kappa v_\omega(R_{\mu^0}) \geq a \Pr_\kappa\{A\}[E_\kappa v_\omega(R_{\mu^0}) - E_\kappa v_\omega(R_0)]$$

for some constant a that depends only μ^0, where R_0 is the set of naive rules (single mental state)

6.4 Ignoring weak evidence and direction of bias

The previous section suggests a reason as to why agents may ignore weak evidence. What are the consequences of ignoring weak evidence? To fix ideas we consider the simple mental system of example 1.

A consequence is that with small sample size, it may be that most of the evidence in favor of say $\theta = 2$ is weak, and as a result casual beliefs will be biased towards $\theta = 1$.

Formally, for a given β, consider

$$\mathcal{P}_2^\beta \equiv \{\omega, \Pr_\omega(\widetilde{l} > 1 + \beta \mid \theta = 2) > 0 \text{ and } \Pr_\omega(\widetilde{\theta}(x) = 2, \widetilde{l} > 1 + \beta \mid \theta = 2) = 0\},$$

Any $\omega \in \mathcal{P}_2$ leads to biased casual beliefs in favor of state 1.