

Cautiousness[‡]

Preliminary and incomplete

Olivier Compte
PSE, Paris

Andrew Postlewaite
University of Pennsylvania

February 2012

Abstract

A standard approach in modeling uncertainty in economic models is to assume that there is a probability distribution over the states of nature and that the decision maker receives a signal that is correlated with the state. In the vast majority of papers in this literature, it is assumed, often implicitly, that the decision maker knows the joint distribution over states and signals. We argue that this assumption is unrealistic, and perhaps more importantly, that this unrealistic assumption can lead directly to unrealistic conclusions. We suggest an alternative approach that does not assume that the decision maker knows the joint distribution, and illustrate the approach in several simple examples.

*Compte: Paris School of Economics, 48 Bd Jourdan, 75014 Paris (e-mail: compte@enpc.fr); Postlewaite: Department of Economics, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 19104-6207 (e-mail: apostlew@econ.upenn.edu). Part of this work was done while Postlewaite visited the Paris School of Economics; their hospitality is gratefully acknowledged. The authors thank the Gould Foundation for financial support, and Postlewaite thanks the National Science Foundation for financial support.

†We would like to thank David Dillenberger, Itzhak Gilboa, David Schmeidler, Peter Wakker and the participants of many seminars at which this work was presented for helpful comments.

1 Introduction

There is a vast literature in economics that models agents' behavior under uncertainty. The overwhelming majority of this literature starts by assuming that the agent has a probability distribution over the underlying unknown states of nature. The agent compares alternatives by their expected utility, calculated using this probability distribution and his utility function. The aim of this paper is to examine the foundation for this modeling strategy and to propose alternatives.

Modeling agents' decision making under uncertainty by expected utility is typically justified by invoking Savage (1954), who showed that intuitively appealing axioms on consistency of choice in the face of uncertainty imply that the decision maker behaves as if he maximized expected utility, with a utility function over certain outcomes and a probability distribution over unknown states of nature. An agent who finds the axioms compelling and faces a choice could use the result to aid in his decision, but might face a problem. Consider the following example.

Example: Tom faces a new, uncertain problem in which he must choose between two alternatives. He understands Savage and, through introspection, determine his utilities for the outcomes in the alternatives, but what probability distribution should he use? Tom considers all past situations that are at all relevant that he has encountered, heard of or thought about in formulating a probability distribution he might use. After availing himself of all possible input, Tom arrives at what he considers the probability distribution that best fits his problem.

With his utility function and probability beliefs at hand, Tom calculates the expected utilities of the two alternatives. But Tom is nervous about committing to the alternative with the higher expected utility. Despite having thought long and hard about what probability distribution he should use, Tom isn't very sure about his probability beliefs; the problem he faces is quite different from anything he has previously encountered, and his beliefs feel like little more than a guess. What bothers Tom most is that he realizes that he might have come to the same beliefs about the state of nature with much more information, and as a result, have been much more confident about his beliefs. His expected utility of the two alternatives is independent of how confident he is about his beliefs, but somehow Tom feels that when he lacks confidence in his beliefs, he should be more cautious in his choice.

Tom's problem illustrates a central point we want to make: modeling an agent making a decision under uncertainty as having a prior over states of nature ignores the fact that the agent typically is less than perfectly confident in his prior. One can, of course require that this uncertainty about what prior he should use be modeled by a prior over possible priors. But there is no reason that the agent should be more confident about his prior over priors than over his initial prior; indeed, it would seem likely he would be less confident. Thus, one can put priors over priors, priors over priors over priors, etc., but at some

point the process stops, and the agent is not completely confident that all the “priors over priors” capture all his uncertainty.

Others have made arguments similar to that above that modeling agents’ uncertainty with a prior over states of nature misses an important aspect of uncertainty and suggested alternatives to the standard model.¹ Schmeidler (1989) proposed that an agent’s beliefs could be represented by a nonadditive probability. Thus, when faced with an urn with black and red balls of unknown proportions, an agent may attach probability .4 to a ball drawn at random to be black, and probability .4 to the ball being red. The sum of the two probabilities being less than 1 can be interpreted as capturing the agent’s residual uncertainty. Agents maximize expected utility of alternatives (using this possibly nonadditive probability) as in the standard model. Gilboa and Schmeidler (1989) propose that when an agent is unsure of precisely what prior to use, the agent considers a *set* of priors. An agent then evaluates an alternative according to the minimum expected utility of the alternative across the priors he considers. These approaches have the attractive feature that they extend Savage: each provides a rule that describes choice behavior, with the rule generated by evaluation via expected utility. Further, the choice rules they generate are derived by weakening Savage’s axioms.

That individuals would be able to formulate precise priors is not the only issue we wish to raise. Tom’s problem reveals another central issue that we want to address. To formulate a belief, or a set of beliefs, Tom needs to gather data that he judges relevant. How does data translate into beliefs? The literature essentially adopts one of two perspectives.

The first is the omniscient/Harsanyi perspective. Tom is omniscient in the sense that he understands the connection between the problems that he faces and the data that he uses to form beliefs for each problem that he faces. Based on this fine knowledge and the data that he currently processes, he uses Bayes rule to form correct posteriors beliefs about the actual problem that he faces. Most applications follow this path.

The other is the subjective/Savage perspective. The process that generates beliefs is left unmodelled. Whether Tom has poor or good abilities to formulate beliefs (or sets of beliefs), they become the input that Tom uses in taking decisions: he maximizes subjective expected utility taking at face value his belief (or maxi-mins subjective expected utility taking at face value his set of beliefs). His choices are disconnected from the actual welfare that he derives, which is determined by the actual probability distribution rather than his beliefs. He follows that mechanical rule in all circumstances, not investigating whether other possibly simple ways to use data might increase his welfare.

Between those two polar perspectives, some have tried to relax the omniscient perspective (adding errors to valuations of alternatives (random stimuli/utility models, Block and Marschak (1960)), possibly putting structure on the sources of errors (non-partitional models, Geanakoplos (1989), procedural

¹See, e.g., Gilboa, Postlewaite and Schmeidler (2008, 2009) for a discussion of some of this work.

models, Osborne and Rubinstein (1998)), or to put consistency constraints on beliefs (case based beliefs (Gilboa and Schmeidler (1995)) and analogy based beliefs (Jehiel (2005))). All these attempts, however, take as given that agents would maximize subjective expected utility.

Our approach differs in that we do not assume that the agent would formulate beliefs, nor do we define a definitive rule by which he ought to use beliefs that he might have. Our perspective is that there is an underlying process that defines the connection between the problem faced and the data that the agent uses. We suggest intuitively plausible rules that an agent may employ given the type of data available to him, and we assume that among these plausible rules, the agent knows which are optimal.

Our approach departs from the omniscient perspective in that we assume there is a limit to the number of rules that one compares. Without limitations, the agent's decisions would be the same as if he were omniscient.

Our approach departs from the subjective perspective in that rules are connected to the actual welfare that the agent derives, and in that it does not assume nor produce a definitive rule describing choice behavior in all circumstances. Rather, we demonstrate that rules that incorporate intuitively plausible properties such as caution can, in some circumstances, increase welfare.

Our interest in proposing an alternative way to model uncertainty is not purely methodological. Much of the success of game theory has been achieved using the omniscient perspective. Our view is that this omniscient perspective has shaped our intuitions about important economic problems, often valuable but sometimes not. It is beyond the scope of this paper to review in detail all applications. We gather here a few comments that are developed in companion papers.

Auctions. In a private value first price auction, we expect individuals to shade their bids below their values, with the extent of shading depending on the strength of competition. For the omniscient individual who knows the joint distribution over values, there is a tight relationship between one's value and the strength of competition, and equilibrium analysis provides details on how shading ought to depend on realized value. Our view is that (at least as first approximation) agents are often unable to make much inference about the state of competition based on their value estimate, and that a basic auction model should reflect that inability. Our auction models are complex in ways that sometimes obscure insights rather than illuminate them.

Long run interaction. In interacting with others, we typically use our experience from past encounters to estimate whether continuing to put in effort (or even to interact at all) is worthwhile. For most interactions, these experiences are private, that is, not directly observable to others. These interactions have been analyzed as repeated games with imperfect private monitoring. Although others' experience is not observable, the omniscient individual has a fine understanding of how his effort level translates into others' perceptions of his effort, and it is assumed that this omniscient individual can exploit that understanding as finely as he wishes.

This omniscient perspective forces us to focus on how omniscient individuals

might manage to cooperate. It is not clear that these help us understand cooperation in real repeated relationships, that is, when agents are not omniscient.

Information aggregation. Aggregating information from different sources is central to many of our models. A decision has to be taken, each agent has an opinion about the appropriate decision, and these possibly distinct opinions must be aggregated. This is a difficult problem in practice, but the omniscient individual knows the joint distribution over appropriate decisions and opinions, and he can easily translate opinions into decisions. The omniscient perspective thus makes information aggregation a trivial issue, and it has led us to focus exclusively on incentive issues (agents' incentives to report opinions).

This emphasis on incentive issues leads us to overlook information aggregation issues and to focus on incentive issues among *omniscient individuals*. For example, in perhaps the most widely used models of information transmission, Crawford and Sobel (1982), the omniscient receiver is able to precisely interpret the sender's message and to extract information from it. The sender's equilibrium behavior must be carefully calibrated to circumvent the receiver's extraordinary ability. Focusing on exactly how this can be done (i.e., the equilibrium construction) may distract us from an important aspect of information transmission: biased senders have more incentive to bias their message, which likely results in their message having less chance to be taken into account at all.

We turn next to a discussion of the standard approach.

2 The standard approach

Modelling a decision problem starts with a set of *alternatives* or possible decisions, and a set of *states* that capture all relevant aspects of the decision environment. If one cares about the quantity of oil in an oil field, the state could be the number equal to that quantity. If one considers a choice between two alternatives, the state could be a pair of numbers, each measuring the intensity of the preference for one of the alternatives.

Formally, we define A as the set of alternatives, and S as the set of states. A state s summarizes everything that the agent needs to know to properly evaluate which alternative is best. *Preferences over alternatives* are then characterized by a utility function $u(a, s)$ that specifies a welfare level for each $a \in A$ and $s \in S$.

The next step consists of modelling the *uncertainty* that the agent faces. This is done by assuming that the agent does not know the state, but that he has some imperfect knowledge of it. Formally, it is assumed that the agent observes a signal $z \in Z$ correlated with the true state s . That signal may be interpreted as the agent's (imperfect) perception of the true state. The joint distribution over state and signal is denoted ω . Then $\omega \in \Delta(Z \times S)$ and we will refer to ω as a possible environment that the agent may be facing, and to $\mathcal{P} \subset \Delta(Z \times S)$ as a set of possible environments.

Finally, rational decision making consists of assuming that for any signal

z that he might receive, the agent takes the alternative that maximizes his expected welfare. Denoting by $r^*(z)$ that optimal decision, we have:

$$r^*(z) \equiv \arg \max_a E_\omega[u(a, s) \mid z].$$

Alternatively, one may define the set \overline{R} consisting of all possible decision rules $r(\cdot)$ that map Z into A . For each rule $r \in \overline{R}$, one can define the expected utility (or performance)

$$v_\omega(r) = E_\omega u(r(z), s).$$

The optimal decision rule r^* solves:

$$r^* = \arg \max_{r \in \overline{R}} v_\omega(r).$$

Example 1 *An estimation problem*

An agent has to determine a number that is complicated to evaluate (a quantity of oil in an oil field, a number of coins in a jar, etc.). The agent makes a report, and he is penalized as a function of the distance between the report that he makes and the true number (e.g., quadratic preferences).

Formally, we choose $u(a, s) = -(s - a)^2$ to describe preferences, and assume that

$$z = s + \varepsilon$$

where s and ε are drawn from independent normal distributions centered on s_0 and 0 respectively:

$$s = s_0 + \eta$$

with $\eta \sim \mathcal{N}(0, \sigma_\eta^2)$ and $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$.

The optimal report then consists of choosing

$$r^*(z) = E_\omega[s \mid z] = z - \rho(z - s_0) \text{ where } \rho = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\eta^2}.$$

In this example, the parameter z can be interpreted as a noisy estimate of the true state, and the optimal rule above can be interpreted as that the estimate should be taken with caution and regressed to the mean (s_0).

3 Implicit informational assumptions

The basic model and standard methods of analysis outlined above come with implicit informational assumptions. We explain what these assumptions are and discuss their plausibility and implications.

Optimization and information.

As explained above, modelling uncertainty proceeds by defining a joint distribution (ω) over signals (z) and states (s). Poorly informed agents can then be modelled as agents for whom z is poorly correlated with s .

From the agent's perspective however, following the optimal decision rule r^* requires *more* knowledge/information than just the mere observation of some signal z . Our models are generally silent about how the agent manages to follow the optimal rule. But whether this is done through computation or learning, it amounts to detailed knowledge of the joint distribution ω . In the estimation problem above for example, it requires the agent to behave as if he knew the conditional expectations $E_\omega[s | z]$.²

So, although it is not necessary that the agent knows ω to determine the optimal decision rule, the traditional view is that the agent knows ω . He can then use ω to form a correct belief $\beta_z \in \Delta(S)$ about the true state for each signal z that he may receive:³

$$\beta_z = \omega(\cdot | z),$$

and next use that belief to derive the optimal rule, that is:

$$\begin{aligned} r^*(z) &= a^*[\beta_z] \text{ where for any } \beta \in \Delta(S), \\ a^*[\beta] &= \arg \max_a E_\beta u(a, s). \end{aligned}$$

In other words, the initial objective is to model agents who are incompletely informed about the characteristics of the specific problem they face, and this is done by assuming that an agent does not observe s but rather, a signal z that may differ from s . For the purpose of evaluating the performance of a given rule r , the modeler defines the actual distribution ω that the agent faces. *A priori, we may not intend to assume that the agent knows that distribution ω , but the assumption that the agent can determine the optimal rule amounts to it.*

Plausibility.

Our first concern is that at least descriptively, this informational assumption does not seem plausible:

(i) If the agent already finds it difficult to determine the state s in S , it is presumably even more difficult for the agent to determine the correct ω in \mathcal{P} .

(ii) We end up modelling agents who know with great *precision* the distribution over the problems they might be facing (that is, the s they are facing) and the distribution over the mistakes they might be making. However, most beliefs are at best crude or vague. In the words of Savage (p58):

The postulate of personal probability imply that I can determine, to any degree of accuracy whatsoever, the probability (for me) that the next president will be a Democrat. Now it is manifest that I cannot determine that number with great accuracy, but only roughly.

²This statement assumes quadratic preferences. In general, what the agent would need to know depends on his preferences.

³For that reason, the standard approach generally identifies the signal that the agent receives with his belief:

$$z \equiv \beta_z.$$

Nature and robustness of insights.

Models need not be realistic. However we ask that a model provide relevant insights as to how agents might behave. Our second concern is that the results and insights that we find may be finely tuned to the joint distribution ω , and eventually *driven by* the implicit assumption that the agent knows that distribution.

In the estimation problem above (Illustration 1), the relevant aspects of ω consist of two parameters, the mean value s_0 and the noise ratio ρ ,⁴ and optimal behavior calls for regressing to the mean as follows:

$$r^*(z, \rho, s_0) = (1 - \rho)z + \rho s_0.$$

From an agent's perspective, there is probably a rough sense in which if an estimate looks extreme, then one should be cautious in taking it at face value. Choosing a more conservative figure that would look closer to "normal" might then seem appropriate. In many applications however, there doesn't exist a "normal" estimate. If one has to estimate a quantity of oil in an oil field, there is nothing to make most estimates seem abnormal or extreme. If I am asked to make a point estimate for a number of coins in a jar, there is nothing that would make me think the estimate is extreme. That estimate just happens to be mine. Of course, if confronted with other people's estimates and seeing that mine lies far above all others, some reconsideration of my own estimate will probably seem advisable.

In other words, regressing to the mean may make sense when an appropriate reference point is available, but not otherwise. And even if one is available, the agent's perception is likely to fall into one of few categories: "more or less normal", "somewhat high", "somewhat low", "extremely high" or "extremely low", without his being able to say more about precisely *how* extreme an estimate seems to be.

The traditional answer to these comments is that if one thinks that the agent is uncertain about some aspect of ω , then this uncertainty should be modelled explicitly.

Let us examine that path with illustration 1, assuming that the mean value is actually uncertain, and that the agent receives a noisy signal of it.

Example 2 *Estimation problem continued*

We build on the above example and now assume that

$$s = s_0 + \eta \text{ and } s_0 = \theta_0 + \gamma$$

where η and γ are independent normal distributions. Thus for a given realization of γ , the state s is distributed as $\mathcal{N}(s_0, \sigma^2)$, but s_0 is now a random variable.⁵ As before, the agent receives a signal correlated with s , which we now denote x :

$$x = s + \varepsilon.$$

⁴This relies on the structure of the model assumed.

⁵So in effect, s is now a normal distribution centered in θ_0 rather than s_0 .

But we also assume that he receives a signal y correlated with the "realized mean value" s_0 .

$$y = s_0 + \varepsilon_0$$

Defining $z = (x, y)$ and ω as the joint distribution over (z, s) , the optimal report consists as before of choosing

$$r^*(z) = E_\omega[s \mid z].$$

The simplicity of this formula hides the fact that the rule now depends not only on x and y , but also on the numbers $\theta_0, \sigma_\varepsilon, \sigma_{\varepsilon_0}, \sigma_\eta, \sigma_\gamma$ that parameterize the distribution ω . By trying to capture the uncertainty about the joint distribution ω of illustration 1, we have just pushed up one level the issue we had wished to address (i.e., s_0 is not known with precision), replacing it with another even more implausible assumption, that all numbers $\theta_0, \sigma_\varepsilon, \sigma_{\varepsilon_0}, \sigma_\eta, \sigma_\gamma$ are known with precision.

A comparative statics exercise.

To illustrate further our concern about insights being driven by the implicit knowledge of ω , we do a comparative statics exercise. We consider a population of identical agents, each facing the same estimation problem as that described in Example 1.

It may be useful to think of agents all facing the same jar filled with coins, and being asked the number of coins in the jar. We assume that by varying the length of time an agent can spend studying the jar, the precision of their estimate varies. We model this by assuming that the noise term ε has higher variance when the time spent is smaller.

The prediction of this model is that as the variance of the noise term increases, the agents' reports should put more weight on s_0 , lowering the variance of the estimates. For any two agents with estimates z_1 and z_2 respectively, and letting $r_i = r^*(z_i)$, we have:

$$E[r_1 - r_2]^2 = (1 - \rho)^2 E[z_1 - z_2]^2.$$

This implies that the variance of the reports vanishes as the estimates become noisier:⁶

$$\lim_{\sigma_\varepsilon \nearrow} E[r_1 - r_2]^2 \searrow 0$$

In other words, as agents have less time to make their estimates, their opinions converge to one another (in the sense of vanishing variance). This is of course an (unfortunate) artifact of the model and of the implicit informational assumption that they all know ω (or behave as if they knew ω).

⁶This is because $2(\frac{\sigma_\eta^2}{\sigma_\varepsilon^2 + \sigma_\eta^2})^2 \sigma_\varepsilon^2$ tends to 0 as σ_ε increases.

4 A different approach

The standard approach embodies strong informational assumptions because there are implicitly no limits on the agent’s ability to compare alternative rules. We propose a theory in which weaker informational assumptions are made by limiting the agent’s ability to evaluate or compare alternatives.

Main assumption

We follow the steps of the standard approach. The state s characterizes a particular decision problem that the agent faces,⁷ while z characterizes the agent’s perception of that decision problem, or the data that the agent considers relevant to the particular decision problem he faces. We shall refer to the pair (z, s) as a particular situation that the agent faces.

Our aim is a model in which decisions are based on the data/perception z , and in which decisions are driven by payoff considerations.

The primitives of our model consist of (i) D , a range of situations over which the model applies; (ii) R , a set of rules that apply to situations in D , where each rule maps the agent’s perception to a decision; (iii) $\omega \in \Delta(D)$, the distribution over situations in D that the agent faces.

Each distribution $\omega \in \Delta(D)$ and rule $r \in R$ induces a joint distribution over decisions and states, from which we derive the performance $v_\omega(r)$ as before. We let $r_\omega^* = \arg \max_{r \in R} v_\omega(r)$. Our main assumption is:

A1: The set of considered rules R consists of a limited set of rules, i.e. $R \subset \overline{R}$.

Although the agent does not know ω , he identifies (or learns) which rule r_ω^* is optimal in R , and he follows it.

The standard approach corresponds to the case where no restrictions are imposed on R , i.e., $R = \overline{R}$. A1 thus corresponds to a weakening of the standard assumption.

A1 requires that the agent identifies which rule is optimal despite not knowing ω . We are agnostic about how the agent does it, though learning is a natural candidate. We emphasize that *how* the agent has come to know which rule is optimal is outside our model.

A1 limits the agent’s ability to compare rules. However, A1 assumes a perfect ability to discriminate across rules in R . This assumption could be weakened, assuming that agents can only imperfectly evaluate which rule is best.⁸ The qualitative features of behavior that we care about are robust to the agent not knowing *exactly* which of the two rules is optimal.

⁷The state s may consist of a specification of preferences over sure outcomes, and of the outcome induced by each alternative. Or it may consist of a specification of preferences over objective lotteries, and of the objective lotteries associated with each alternative.

⁸Denoting by r^* the rule that the agent perceives as best, and denoting by $\mu_\omega(r)$ the probability that r is perceived as best under ω , we could assume:

A2: $\mu_\omega(r) = \phi_r(\{v_\omega(r)\}_{r \in R})$ where for any vector $v = \{v^r\}_{r \in R}$, $\phi(v) = \{\phi_r(v)\}_{r \in R} \in \Delta(R)$. For example, the function ϕ could be a logit function ($\phi_r(v) = \exp \beta v^r / \sum_r \exp \beta v^r$).

Being able to identify which rule is optimal in R is precisely what we interpret as an implicit informational assumption. This assumption implicitly assumes that if a rule is added to the set the agent compares, the agent knows more.

From an (omniscient) outsider's perspective, A1 is tantamount to the agent behaving *as if he had information on the distribution* ω . With no restrictions on R , the agent behaves as if he knew ω and then followed expected utility maximization. With restrictions, the informational assumption is weakened. For example, if R consists of two rules, say

$$R = \{r_0, r_1\}$$

then one can define a partition of $\Delta(D)$ into two subsets \mathcal{P}_{r_0} and \mathcal{P}_{r_1} depending on which rule is better.⁹ When, say $r_\omega^* = r_0$, the agent understands that the optimal rule is r_0 and he follows it. So he behaves as if he knew that $\omega \in \mathcal{P}_{r_0}$, *without being able to exploit this information any further*.¹⁰ We shall sometimes refer to $\{\omega \in \mathcal{P}_{r_0}\}$ as a (crude) representation of the uncertainty that the agent faces, or implicit belief, with the understanding that the agent believes that the class of situations calls for rule r_0 . We do not have in mind however that the agent can formulate with precision that ω lies in subset \mathcal{P}_{r_0} , only the omniscient outsider can. In particular, we emphasize that one can think of agents understanding which rule to follow in a given circumstance without understanding probability at all.

(D, R, ω) is a primitive, hence the *analyst's* choice, in the same way that ω is in the standard approach. We do not suggest that there is a unique set of rules that the analyst should consider. Different sets (say R or R') could in principle lead to different insights about behavior, and the set of rules the analyst chooses will be governed by the questions of interest in a given problem. This is analogous to practice in the standard approach. When analyzing auctions, one could argue that for any auction of interest, there are both private value and common value components. Nevertheless, it is common to focus attention on one or the other components, with different insights obtained for the two.

We do not propose a cookbook for choosing the set of rules R , as there is no general recipe for writing models. It is the analyst's contribution to find sets of rules that look *plausible* and insights that appear *relevant* or *useful*?. Relevance may stem from insights being robust to modifications of R . A rule's plausibility may stem from the combination of various considerations: (i) at least for some plausible ω , that rule remains optimal even when R is enlarged; (ii) it captures computational constraints, or other cognitive limitations; (iii) it applies to a variety of situations in a way that does not seem too situation-specific, and in a way that seems driven by plausible economic considerations.

To illustrate with a counter example, one can think of rules that are not too complex to describe, possibly supported by some ω , and yet appear somewhat

⁹ $\mathcal{P}_{r_k} = \{\omega \mid v_\omega(r_k) = \max_{r \in R} v_\omega(r)\}$

¹⁰That the agent is unable to exploit the information that r_0 is best is precisely what makes this notion of information nonstandard. " r_0 is best" is information but it is not data that can be used and exploited to further condition behavior and improve performance.

crazy. Suppose the signals were the interval $[0,1]$, and there were two alternatives, a and b, to choose from. The rule that chooses a if the second digit in the signal is odd and chooses b if it is even appears somewhat crazy.

The estimation problem revisited.

We return to the estimation problem. We first define a more general class of preferences that allows for asymmetric costs of overshooting or undershooting the report. We next suggest plausible rules and apply our approach.

Preferences. We consider preferences characterized by a utility function $u(a, s, \beta)$, where $\beta \geq 0$ parametrizes the additional loss from overshooting the target. Specifically, we assume that $u(a, s, \beta)$ depends only on the difference $d = a - s$. We let

$$\begin{aligned} u(a, s, \beta) &= L_\beta(a - s) \text{ where} \\ L_\beta(d) &= -d^2 - \beta \max(d, 0). \end{aligned}$$

The quadratic case thus corresponds to the case where $\beta = 0$. Note that the state that fully describes preferences now consists of the pair (s, β) .

Uncertainty. Uncertainty is modelled as a joint distribution over state and observations. For simplicity we assume there is no uncertainty about the preference parameter β . Thus, an environment is described by a pair (ω, β) where $\omega \in \Delta(S \times Z)$ is defined as before.

Plausible rules. One rule that the agent could follow would be to *take his estimate at face value*:

$$r_0(z) \equiv z.$$

We shall refer to it as the *naive* rule, as this would be the optimal rule if the agent was making no mistakes in his estimates. Other rules might have him exercise caution, and pick an action that distorts, upward or downward, his initial estimate. We define rule r_γ as:

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

We fix γ_0 and analyze below the case (case 1) where the agent only compares three rules:

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

We shall also discuss the case (case 2) where the agent compares many rules:

$$R_2 = \{r_\gamma\}_{\gamma \in \mathcal{R}}.$$

Analysis. When the agent uses rule r_γ , each realization ε induces a difference $d = x - \gamma - s = \varepsilon - \gamma$. Denoting by f the distribution over errors, the expected performance associated with rule r_γ can be written as:

$$v_{\omega, \beta}(r_\gamma) = E_\varepsilon L_\beta(\varepsilon - \gamma) = -[\sigma_\varepsilon^2 + \gamma^2 + \beta \int_{\varepsilon > \gamma} (\varepsilon - \gamma) f(\varepsilon) d\varepsilon] \quad (1)$$


which implies:

$$v_{\omega,\beta}(r_\gamma) - v_{\omega,\beta}(r_0) = -\gamma^2 + \beta \left(\int_0^\gamma \varepsilon f(\varepsilon) d\varepsilon + \gamma \Pr(\varepsilon > \gamma) \right).$$

Proposition 1: *Consider case 1. If $\beta < 2\gamma_0$, then performance is maximum for the naive rule r_0 . When $\beta > 2\gamma_0$, then the cautious rule r_{γ_0} may become optimal if the noise has sufficiently high variance.*

Proof: $\int_0^\gamma \varepsilon f(\varepsilon) d\varepsilon \leq \gamma \Pr(\varepsilon \in (0, \gamma))$, so the difference $v_\omega(r_\gamma) - v_\omega(r_0)$ is below $-\gamma^2 + \gamma\beta/2$, which implies the first statement. Now if the noise has large enough variance, $\Pr(\varepsilon > \gamma)$ is close to $1/2$, which implies the second statement. \square

The following figure plots the respective regions where r_0 or r_{γ_0} is optimal when γ_0 is set equal to 1. Qualitatively, this figure illustrates the combinations of asymmetric loss functions (β) and noisiness of estimates (σ_ε) that call for cautious behavior.



Adding more rules would allow the agent to more finely adjust the cautiousness level γ to the environment he is facing, $(\sigma_\varepsilon, \beta)$. As the number of rules considered increases however, this fine tuning implicitly assumes more precise information about the environment $(\sigma_\varepsilon, \beta)$ that the agent is facing. At the limit where all rules r_γ are considered (case 2), the optimal rule is obtained by differentiating (1). Optimal shading γ^* solves:

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

implying that shading being stronger when ε is more dispersed or when the cost parameter β is larger.

Discussion.

Tracking relevant aspects of ω . A central assumption of our approach is that the agent does not know precisely the specific distribution ω he is facing.

A1 however is an assumption that the agent behaves as if he understood some aspects of ω that are relevant to his decision problem.

This is done in a crude way in case 1, but it nevertheless permits us to partition the parameter space $(\beta, \sigma_\varepsilon)$ and understand under which circumstances the agent exerts (or does not exert) caution. Case 2 gives a more complete picture, as the magnitude of the caution he exerts is fit to the dispersion of his errors (but this implicitly requires a fine knowledge of this dispersion).

What's relevant? There is an (implicit) relationship between the set of rules considered and the aspects of ω that are relevant. For example, the rules r_γ do not attempt to adjust the cautiousness level as a function of the estimate z . As a consequence, the distribution over s is irrelevant; only the distribution over errors is relevant.

Other classes of rules. Depending on the data or signals that the agent receives, other classes of rules may be seen as plausible. We provide two examples based on Example 2 (in which the agent receives two signals x and y about the true state s and the realized mean s_0 respectively).¹¹

Example 3 *The agent may consider a rule that regresses his estimate x to the (perceived) mean y , that is, for any $z = (x, y)$, use rules of the form:*

$$r_\rho(z) = (1 - \rho)x + \rho y,$$

with $\rho \in [0, 1]$.¹² *That set of rules may be preferable to R_2 , but only to the extent that y is not too noisy and that β is not too large.*

Example 4 *Observing $x - y$, the agent may decide whether x is surprisingly high, and behave normally or cautiously depending on whether he perceives that x is normal or high. Formally*

$$\begin{aligned} r_{\gamma,h}(x, y) &= x \text{ if } x < y + h \\ &= x - \gamma \text{ if } x > y + h. \end{aligned}$$

The interpretation is that under rule $r_{\gamma,h}$, the signal x either seems normal (and it is then taken at face value), or it seems high, in which case the agent is cautious.

Consider the case (case 3) where

$$R_3 = \{r_0, r_{\gamma_0}, r_{\gamma_0, h_0}\}.$$

We can find which rule is best as a function of β , σ_ε , σ_{ε_0} and σ_η . It is easy to see that as σ_{ε_0} or σ_η become large, r_{γ_0, h_0} corresponds to a random choice between r_0 and r_{γ_0} , and it is therefore dominated. If neither σ_{ε_0} nor σ_η is too high, then for intermediate values of β , rule r_{h_0, γ_0} becomes optimal.

¹¹Recall that $x = s + \varepsilon$, $y = s_0 + \varepsilon_0$ and $s = s_0 + \eta$.

¹²For example, if $R = \{r_\rho\}_\rho$ and $\beta = 0$, then the optimal rule is $\rho^* = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_{\varepsilon'}^2 + \sigma_\theta^2}$.

5 Application: cautious heuristics

A pair (z, s) corresponds to a particular situation that one faces, with z being the agent's perception of the situation. Individuals partition or categorize or classify the situations they face in possibly different ways, and given a particular classification, D corresponds to a particular class of situations, while ω corresponds to the distribution over situations in D . A rule r can then be interpreted as a *heuristic* that applies to all situations in D .

With this interpretation in mind, our approach assumes that an agent has various heuristics available, and that he is able to determine which heuristic works best for him on average over the situations in D that he faces.

We apply this idea to a simple class of choice problems with two alternatives, one of which is easier to evaluate than the other. We also extend our analysis to the case in which each alternative has two consequences, one of which is easier to evaluate than the other. We endow agents with two heuristics, a naive heuristic, and a cautious one, and derive circumstances under which the cautious heuristic is better.

This application enables us to draw a link between several much studied phenomena: ambiguity aversion, status quo bias, procrastination, temptation, and the winner's curse. All these phenomena can be understood through a single lens: *cautious behavior as a response to noisy evaluations*.

5.1 A basic choice problem.

We consider a class of choice problems with two alternatives, one of which, labelled a , can be easily evaluated by the agent, while the other, labelled b , is more difficult to evaluate. To each alternative $k \in \{a, b\}$ we associate a state s_k with the understanding that taking alternative k yields utility $u(s_k)$ to the agent.¹³ If the agent knew the state $s = (s_a, s_b)$, the optimal decision would consist of choosing the alternative $k \in \{a, b\}$ for which s_k is largest.

The agent forms a possibly noisy estimate of s_a and s_b , denoted x_a and x_b respectively.¹⁴ We assume first that $x_a = s_a$ and

$$x_b = s_b + \varepsilon.$$

That is, the agent learns s_a perfectly but gets a noisy signal of s_b . We let $z = (x_a, x_b)$ and denote by ω the joint distribution over (z, s) .

Plausible rules.

One naive rule would be that the agent takes his estimates at face value, and chooses the alternative k for which x_k is largest. For any $z = (x_a, x_b)$,

$$r_0(z) \equiv \arg \max_k x_k.$$

¹³If k is lottery, s_k would thus correspond to the certainty equivalent associated with k .

¹⁴An alternative assumption would be to assume that he forms an estimate of the difference $d = s_a - s_b$, as it may sometimes be easier to evaluate differences than each alternative separately.

A more cautious rule would be to first distort his estimate x_b by γ :

$$\begin{aligned} r_\gamma(x) &\equiv a \text{ if } x_a > x_b - \gamma \\ &\equiv b \text{ otherwise.} \end{aligned}$$

The performance of rule r_γ is given by the expected utility he obtains under ω when he follows r_γ . We denote it $v(r_\gamma)$:

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

We examine the case where the set of plausible rules is

$$R_1 = \{r_0, r_{\gamma_0}\}$$

for some $\gamma_0 > 0$, assumed to be small.

Analysis.

Observe that the rules r_0 and r_{γ_0} only differ under the event where $x_a \in (x_b - \gamma_0, x_b)$, with the agent undertaking b under rule r_0 and undertaking a under rule r_{γ_0} . Thus rule r_{γ_0} is best when the difference

$$\Delta \equiv E[u(x_b - \varepsilon) - u(x_a) \mid x_a \in (x_b - \gamma_0, x_b)]$$

is negative. We illustrate various circumstances under which r_{γ_0} is better.

Let $\theta \equiv s_a - s_b$. To fix ideas, we assume that ε and θ are drawn from independent normal distributions (respectively $\mathcal{N}(0, \sigma_\varepsilon^2)$ and $\mathcal{N}(\theta_0, \sigma^2)$), independently of s_a . We also assume that preferences are characterized by $u_\lambda(x) = -e^{-\lambda x}$ with $\lambda > 0$. The limit case where $\lambda \searrow 0$ corresponds to risk neutrality, and we shall refer to it as $\lambda = 0$. Then Δ can be rewritten as:

$$\Delta = E[e^{\lambda\theta} - 1 \mid -\theta + \varepsilon \in (0, \gamma_0)].$$

We have:

Proposition 2: *Let $\bar{\lambda} = \lambda + \frac{2\theta_0}{\sigma^2}$. If $\bar{\lambda} > 0$, then for σ_ε^2 large enough (i.e. $\sigma_\varepsilon^2 \geq 2\gamma_0/\bar{\lambda}$), performance is maximum for the cautious rule r_{γ_0} . And for σ_ε^2 small enough, performance is maximum for the naive rule.*

There are two cases worth emphasizing.

case 1: $\theta_0 = 0$ and $\lambda > 0$. On average alternatives a and b are equally attractive. Because the agent is risk averse, high enough noise in evaluations of alternative b leads to caution.

case 2: $\theta_0 > 0$ and $\lambda = 0$. The agent is risk neutral, but on average a is a better alternative. For large enough noise, caution is called for too. Intuitively, when σ_ε becomes extremely large, both rules generate a random choice (essentially driven by the error term). The cautious choice a being better on average, r_γ is a better rule.

Proposition 2 illustrates two different motives for caution. One motive is a combination of risk aversion and noisy evaluations. The other motive is that which ever rule one uses, the agent is subject to a *selection bias*: b is favored in events where ε is positive. This is not a problem when on average b is a better alternative. But it becomes a problem otherwise, and caution becomes desirable because it permits to overcome the selection bias.

Proof: Define $\Delta(x) = E[e^{\lambda\theta} - 1 \mid -\theta + \varepsilon = x]$ and let $\rho = \frac{\sigma^2}{\sigma^2 + \sigma_\varepsilon^2}$. Conditional on $-\theta + \varepsilon = x$, θ is normally distributed, with mean $\mu(x) = (1 - \rho)\theta_0 - \rho x$ and variance $(1 - \rho)\sigma^2$. This implies $\Delta(x) = \exp(\lambda\mu(x) + \frac{\lambda^2}{2}(1 - \rho)\sigma^2) - 1$. Denote by x^* the solution to $\Delta(x) = 0$. We have $x^* = \frac{\bar{\lambda}\sigma_\varepsilon^2}{2}$,¹⁵ and for any $x \in (0, x^*)$, $\Delta(x) > 0$. So if $x^* > \lambda_0$, that is if $\sigma_\varepsilon^2 \geq 2\gamma_0/\bar{\lambda}$, then $\Delta \geq \Delta(\lambda_0) > 0$. \square

5.2 Extension

We extend our basic model to deal with situations in which each alternative generates two types of consequences, one easy to evaluate, the other more difficult. We index the consequences by $t = 0, 1$ (t for time), having in mind, for example, that consequences in the future ($t = 1$) are more difficult to evaluate than consequences in the present. We let

$$\begin{aligned} s_a &= s_a^0 + s_a^1 \\ s_b &= s_b^0 + s_b^1 \end{aligned}$$

where all s_k^t are i.i.d. We focus on the case where $\lambda = 0$ (risk neutrality) so the agent's welfare coincides with s_k .

The data that the agent gets is $z = (z^0, z^1)$ where

$$z^0 = s_a^0 - s_b^0 \text{ and } z^1 = s_a^1 - s_b^1 + \varepsilon,$$

where ε is an independent normal noise term.

The agent is assumed to compare two rules defined as follows:

$$\begin{aligned} r_0(z) &= a \text{ iff } z^0 + z^1 > 0, \\ r_\gamma(z) &= a \text{ iff } z^0 + (1 - \gamma)z^1 > 0, \end{aligned}$$

for some fixed $\gamma \in (0, 1]$. The rule r_γ can be interpreted as a cautious rule. It puts a lower weight on the future, reflecting the difficulty in evaluating future consequences.

When γ is not too large, the cautious rule r_γ is an appropriate response to noisy evaluations of future consequences. Intuitively, the (risk neutral) agent cares about the difference $\Delta = s_a^0 - s_b^0 + s_a^1 - s_b^1$. When σ_ε is large, the signal z_1 provides poor information about Δ hence should be weighted less. When σ_ε is very large, r_0 essentially implements a random choice, and r_1 (which only considers signal z_0) is a better rule.

¹⁵Indeed, $x^* = \frac{1-\rho}{\rho}(\theta_0 + \lambda\sigma^2/2) = \frac{\sigma_\varepsilon^2}{2\sigma^2}(\lambda\sigma^2 + 2\theta_0) = \frac{\bar{\lambda}\sigma_\varepsilon^2}{2}$

5.3 Cautious behavior: examples.

Case 1 above illustrates that when an agent is facing a choice between two alternatives, and when that agent (correctly) perceives that one alternative (a) is easier to evaluate than the other alternative (b), he may prefer to be cautious and opt for a rule that favors the easier to evaluate alternative.

Caution is different than risk aversion. The state $s = (s_a, s_b)$ refers to certainty equivalents associated with a and b , so caution plays a role that is distinct from the more standard risk aversion. Caution arises because of the combination of estimation errors in evaluating s_b and risk aversion.

Caution also arises because for the situations considered (i.e., situations in D), the agent *correctly perceives* that one alternative is more difficult to evaluate than the other alternative, and then labels alternatives accordingly (as b and a respectively). If he did not and if, say the alternatives had equal chances of being labelled a and b , or b and a , then the naive rule would be best.¹⁶

These comments illustrate two channels by which cognitive abilities matter. They may affect the accuracy of evaluations (hence the correlation between z and s), and they may also affect the class of situations under scrutiny, the shape of the distribution ω and its asymmetry. The better rule is a response to **inaccuracies** in evaluations and to the perceived asymmetry in the situations considered.

From the perspective of experimental economics, which generally attempts to identify biases in decision making or cognitive limitations, the comments above suggest that observed behavior may also reflect (optimal) responses to the agent's (somewhat accurate) perception of his cognitive limitations.

We discuss below various common biases in light of this last observation.

Ambiguity aversion.

When will an alternative be perceived as more difficult to evaluate than another? We see two different characteristics that can lead to distinct characterizations or labelling of alternatives.

- *complexity*, if evaluating an alternative requires more abstract reasoning or more complex computations than the second alternative (a compound lottery as opposed to a simple lottery for example).

- *vagueness*, if the uncertainty is described in vague terms for one alternative and with precise probabilities over outcomes for the other.

The latter case is often referred to as Knightian uncertainty (as opposed to risk), and has been analyzed experimentally in Ellsberg (1961). The classic interpretation of Ellsberg urn experiment is that agents exhibit a preference for precise lotteries over vague or ambiguous outcomes, and this is usually referred to as ambiguity aversion. From our perspective however, this behavior could instead stem from standard risk aversion combined with the fact that in general vague outcomes are more prone to estimation errors or more difficult to evaluate

¹⁶More generally, if the perception of the relative difficulty in evaluating b is uncorrelated with the mistakes he makes in evaluating a and b , then the naive rule is best.

than precise lotteries.¹⁷ The first case is also consistent with recent experiments reporting that "ambiguity aversion" is correlated with the complexity of the lottery that the agent faces (Halevy (2007)).

Status quo bias.

Case 1 can help understand status quo bias if the status quo is the easier alternative to evaluate. The combination of risk aversion and noisy evaluation of the novel alternatives generates some cautiousness hence a "bias" towards the status quo.

Winner's curse.

We can apply the basic model to a choice between "not buying" (option a) and "buying" (option b), with the idea that the agent only gets a noisy estimate of the gain from buying. Case 2 shows that if on average "always buying" is a bad option, then noisy estimates call for cautious behavior, that is, using the rule r_γ that favors the "not buying" alternative. This cautious behavior looks like a response to the winner's curse, but it is just an appropriate response to noisy estimation of valuation when on average "always buying" is a bad option. With this interpretation, winner's curse is a selection bias, and it is stronger when competing with many others (because the option "buying always" worsens).

Procrastination and temptation

Cautiousness calls for shading the estimates associated with consequences that are more difficult to evaluate (for risk averse individuals). As in general, future consequences are more difficult to evaluate, cautiousness calls for shading estimates of future prospects. This cautious rule may thus generate behavior that looks like temptation (and procrastination) but it is just an appropriate response to noisy evaluations of future consequences.

Ex post Rationalization.

Finally, in repeated choice problems between two types of decisions, say a and b , we tend to have more precise information about the consequences of the decisions that we take than over those that we do not take. Once decisions of a particular type (say type a) have been taken, cautiousness consists of favoring the type of decisions we have become familiar with. We often refer to ex post rationalization as a bias. But from our perspective, it may just be an appropriate response to noisy evaluations of decisions that we are not familiar to.

¹⁷The latter view is also consistent with the experimental results reported in Fox and Tversky (1995). They consider the Ellsberg set up: Alternative 1 is the precise alternative where the agent gets \$100 if he draws the color of his choice from an urn composed of 50 red and 50 black balls; Alternative 2 is the vague alternative where the agent gets 100\$ if he draws the color of his choice from an urn composed of an unknown number of red and black balls. Fox and Tversky elicit the individuals' willingness to pay for each alternative. They find that these willingnesses to pay differ substantially across alternatives only if the individuals are confronted with both alternatives simultaneously. Our interpretation is that it is only when confronted with both alternatives simultaneously that the agent will perceive alternative 2 as more vague than alternative one, hence that more caution is called for.

5.4 Heuristics and biases

One interpretation of optimal rules in our model is that they are good heuristics: they work well on average across the problems that the individual faces. In this sense, avoiding vague or complex alternatives, or favoring alternatives that we understand better (i.e., the status quo), or shading estimates of future prospects are good heuristics, all derived from the same motive: caution.

From an experimenter’s perspective, some choices may look suboptimal or subject to biases. However, these choices may simply stem from the agent’s using heuristics based on clues such as complexity or vagueness. From the agent’s perspective, using such heuristics may be valuable, because the clues are correlated with the estimation errors that the agent makes in evaluating alternatives. Naturally, if he was certain that he wasn’t making estimation errors for the problem he currently faces, then caution would not be called for. The difficulty, however, lies not in his using a particular heuristic, but in his inability to condition more finely on the type of problem faced.

This view of heuristics as rules that are good on average across problems is consistent with Tversky and Kahneman (1974)’s view.¹⁸ What we have proposed in this section is a model in which various heuristics compete.

6 Discussion and Related literature

Most decision and game theory applications assume that agents know (or behave as if they knew) the joint distribution ω over states and perceptions/signals. The theory of choice (i.e., the subjective approach) does not make that assumption, nor does the random stimuli model of Thurstone (1927). Hansen and Sargent (2001) and subsequent work model agents who have correct beliefs but that do not have full confidence in them. We discuss these models below, in light of our approach. We end with further comments.

6.1 Theory of choice.

Choices reveal not only a person’s relative preferences over outcomes but also his perception of the uncertainty he faces. The theory of choice aims at separating perceptions of uncertainty from preferences. It postulates axioms where that distinction can be made unambiguously. For an individual whose choices obey these axioms (Savage axioms), the uncertainty can be represented as beliefs that take the form of probability distributions over events: the individual acts as though he was attaching precise probabilities to uncertain events, and maximizing expected utility.

The theory makes no claim that agents would necessarily obey the axioms. However it generally argues that a *rational* agent would *want* to satisfy the

¹⁸Tversky and Kahneman describes heuristics that people employ in assessing probabilities or predicting values (for example, that of judging frequency by availability, that is, the ease with which similar occurrences are brought to one’s mind). They argue that these heuristics are useful, but that they lead to estimation errors for some types of problems.

axioms, presumably out of a concern for *coherence* across decisions. The theory then provides a way to achieve such coherence: If a decision maker wanted to satisfy the axioms, then one way to do that would be to attach (personal) probabilities to events and maximize (subjective) expected utility.

Many authors have discussed critically the theory and proposed weaker axioms. We wish to express two concerns that apply to the original theory and its various amendments.

6.1.1 On the connection with welfare

Would anyone want to follow an advice that does not enhance welfare?

In advising an agent to form beliefs and to maximize subjective expected utility (SEU), the theory is not concerned with the actual (objective) uncertainty that the agent faces, nor with the actual process that would generate the perceptions that the agent has of the uncertainty that he faces, nor with the *actual* welfare that the agent obtains if he follows the advice.

From our perspective, subjective expected utility maximization is a rule that ought to be evaluated and compared with other rules. It takes as input a belief $\beta \in \Delta(S)$, and possibly the agent's perception of his own preferences \tilde{u} , and computes, for each perception $z = (\beta, \tilde{u})$, the "optimal" decision:

$$r^{SEU}(z) = \arg \max_a E_\beta \tilde{u}(a, s).$$

Our view is that *the perception z is a signal, and r^{SEU} is just one particular rule, one that takes beliefs and perceived preferences at face value and mechanically maximizes subjective expected utility.* Defining the process that generates s and z as a joint distribution $\omega \in \Delta(S \times Z)$, one can evaluate the actual performance of that rule by computing:

$$v_\omega(r^{SEU}) = E_\omega u(r^{SEU}(z), s).$$

There is of course a special case where r^{SEU} is a good rule: when perceptions are correct, that is, when $\tilde{u} = u$ and when β coincides with the true posterior $\omega(\cdot | z)$. This case corresponds to the standard approach discussed in Section 1.

As soon as perceptions are not correct however, other rules may perform better: Rules that treat beliefs with caution, or rules that rather than using finely described probability distributions over states use cruder perceptions (possibly because these cruder perceptions may be less noisy). Thus if our aim is in understanding the behavior of agents who attempt to use rules that perform well, there is no reason to restrict attention to agents who form beliefs taking the form of probability distributions; even if we restrict attention to such agents, there is no reason to expect them to take these beliefs at face value.¹⁹

¹⁹The same critique applies to the various weakenings of the Savage axioms that have been proposed, including those that lead to representations of uncertainty in terms of probability sets. Choices are then represented as functions of these probability sets. Once viewed as advice or a decision making process that agents might consider following, these belief sets ought to be viewed as signals, and the agent should aim to use rules that best utilize these signals.

6.1.2 On the coherence requirement.

The theory of choice defines an underlying state space and asks for coherence across *all* hypothetical decisions (or comparisons) between alternatives (or acts) inducing state contingent outcomes. But would anyone insist on being coherent in this way?

To illustrate with a concrete example, think of the decision to buy (or not) an oil field at some price p_0 . The actual quantity is $s \in S$, and is not known. Buying at p_0 (f_{p_0}) and not buying (g) are two possible "acts". The set of acts also includes all call options, such as the one where an agent obtains the oil field at some predetermined price p if the quantity s falls in a given subset $E \subset S$ (call it ($f_{p,E}$)). The axioms ask for coherence across all decisions involving these call options.

There are three concerns about this structuring of the problem.

(i) The agent may not care much about coherence across decisions that he never has to make.

(ii) The coherence requirement implicitly assumes a *fixed* perception of uncertainty across all decisions. In reality, the perception of uncertainty is likely to depend on the choice proposed, as one's mind will focus or pay more attention to those states that appear more relevant given the choice proposed.²⁰

(iii) From a normative perspective, we see no legitimacy in asking for coherence unless the perception of uncertainty is correct. If perceptions are noisy, asking for coherence merely asks for correlated mistakes.

We are back to the question we asked earlier: would any one want to be coherent across decisions if this does not enhance welfare, or worse, if coherence is actually conducive to more correlation in the mistakes that one makes.

6.2 Robust decision theory

Hansen and Sargent (2001) model situations in which the agent does not have full confidence in his belief. Formulated within our framework, the agent receives data in the form of a belief $\beta \in \Delta(S)$, that is $z \equiv \beta \in Z \equiv \Delta(S)$. For an agent who would take his belief at face value, an optimal rule consists of maximizing (subjective) expected utility, that is, following r^{SEU} where

$$r^{SEU}(z) \equiv \arg \max_a E_z u(a, s).$$

For an agent who does not entirely trust his belief, a plausible rule is the one suggested by Hansen and Sargent:

$$r^\theta(z) \equiv \arg \max_a \min_{z'} E_{z'} u(a, s) + \theta d(z', z),$$

where θ is some positive parameter and $d(\cdot, \cdot)$ is a distance between beliefs. For any $\theta < \infty$, rule r^θ can be interpreted as a caution rule, with $\frac{1}{\theta}$ measuring the

²⁰ Agents presumably care about the dual problem: the perceptions of the uncertainty vary a great deal across the problems they face, and their task is not so much to achieve coherence in the face of *unchanging* uncertainty, but to determine what signals to process and how to process them, in a way that works reasonably well across the decisions they face.

degree of caution. At the limit where θ gets arbitrarily large, r^θ coincides with r^{SEU} .

The rules r^θ are plausible, yet it is not a priori clear that agents would want to use them. As for subjective expected utility, our view is that a rule will be used if it performs well, and to evaluate the performance of a rule, we need to make some assumption on the joint distribution $\omega \in \Delta(S \times Z)$ over states and data. Once this is done, one may indeed define:

$$v_\omega(r^\theta) = E_\omega u(r^\theta(z), s),$$

enabling us to make comparisons between rules.

One may distinguish two cases:

- *correct beliefs*: z coincides with $\omega(\cdot | z)$. This is the case considered in Barillas and Sargent (2009). Beliefs are correct, and using a caution rule r^θ with $\theta < \infty$ is clearly not called for.

- *noisy beliefs*: z does not coincide with $\omega(\cdot | z)$. Then, *among rules* r^θ , one cautious rule r^{θ_0} may turn out to be optimal.

The last statement emphasizes ‘*among rules* r^θ ’ because in the absence of a restriction on the set of rules considered, the optimal rule would be:

$$r^*(z) \equiv r^\infty(\omega(\cdot | z))$$

In other words, an agent may find attractive to follow r^{θ_0} rather than r^{SEU} . But this is the case only because his beliefs are noisy *and* because there is an implicit restriction on the set of rules that the agent compares. As in our basic model, cautiousness can be seen as an appropriate response to noisy perceptions when the set of rules considered is constrained.

There is no reason however to expect the set of rules that agents consider to take the form $\{r^\theta\}_\theta$. This happens to be a restriction imposed by the modeler, and other restrictions such as the ones we considered (for example rules in which beliefs do not necessarily take the form of a probability distribution) may be equally illuminating.

6.3 Random stimuli model

Thurstone (1927) introduced a random stimuli model that can be described as follows. Think of an individual trying to assess which of two colors $k = 1, 2$ is brighter. Call R_k the (true) brightness level and call S_k the random stimulus received by the agent, assuming that

$$S_k = R_k + \sigma_k \varepsilon_k.$$

In assessing whether $R_1 > R_2$, the individual is assumed to compare the stimuli and report that $R_1 > R_2$ when $S_1 > S_2$. The noisy perception may thus lead to mistakes in evaluating whether R_1 lies above or below R_2 .

Going back to our choice problem between two alternatives $k = 1, 2$, and interpreting s_k as some physical property of alternative k that the agent cares

about, our model may be viewed as a random stimuli model in which the individual has a noisy perception of s_k :

$$z_k = s_k + \sigma_k \varepsilon_k.$$

The true welfare (or utility) associated with alternative k is $u_k = u(s_k)$, but the agent observes a noisy signal of that utility, say $\hat{u}_k = u(z_k)$. One rule consists of maximizing perceived utility (thus taking at face value the signals \hat{u}_k).

However, as **with** maximization of subjective utility, there is no reason that maximizing perceived utility is a good rule for welfare. Using our terminology, maximization of perceived utility corresponds to the naive rule r_0 , but other rules such as the cautious rule r_γ might improve welfare.

Intuitively, returning to the color example, there are two types of mistake one can make: saying that $R_1 > R_2$ when $R_2 > R_1$, and saying that $R_2 > R_1$ when $R_1 > R_2$. If only one type of mistake is costly, then clearly a cautious rule will perform better. Another example, in line with the winner's curse type of example, is as follows: if option 1 is better on average, and if perceptions regarding option 2 are noisier, then one ought to be cautious when $z_2 > z_1$.

It is interesting to note that all the random utility models following Block and Marschak (1960) and McFadden (1973), or the applications of that model to games (i.e., the quantal response models following McKelvey and Palfrey (1995)) have followed the same path, namely, assuming that agents maximize \hat{u}_k . For the econometrician who interprets \hat{u}_k as the true utility to the agent, the assumption that the agent maximizes \hat{u}_k is clearly legitimate. But if one truly interprets \hat{u}_k as a perceived utility,²¹ then the joint distribution over true and perceived utility matters, and the rule that maximizes perceived utility has no reason to be welfare maximizing, unless specific assumptions are made relative to that joint distribution.

7 Perspectives

Our approach brings up new perspectives on modelling imperfect information. We first make precise why our main assumption (A1) falls outside any standard notion of imperfect information. We then outline implications of our approach for decision problems, games and mechanism design.

7.1 Limited information, skill and awareness

Our main assumption (A1) has various interpretations, either as imperfect or limited information, or as limited skill or awareness.

Limited information. The agent does not know with precision the process ω that generates data, but being able to identify that " r_{k_0} is best in R " is equivalent to him being *partially informed of* ω . This notion of partial information is

²¹In both Block and Marschak (1960) and McKelvey and Palfrey (1995) for example, imperfections in perceptions motivate stochastic choices.

quite different from the standard one: to model imperfect information on ω , a standard model would presume that the process ω is drawn from a distribution over processes (call ϕ that distribution), and " r_{k_0} is best in R " would then be viewed as a *signal* about the *actual* process ω . Based on that signal and the presumption that the agent knows ϕ , the agent would then derive the optimal decision rule r .²²

This does not mean that the information " r_{k_0} is best in R " cannot be viewed as a signal under our approach. It can, but then it should be included in the description of the data that the agent obtains. Data would then consist of a pair $z^e \equiv (z, k_0)$, where k_0 stands for the signal " r_{k_0} is best in R ", and a rule would then map z^e to decisions. The modeler's task would then be to define the subset of rules R^e that the agent can plausibly compare, and determine what additional insights one would derive from that more elaborate data generating process.²³

Limited skill or limited awareness. Our agent is aware that his estimate z is not correct, and a rule defines a particular way to exploit the estimate. Because he only considers a restricted set of rules, he has a limited ability to exploit that estimate (compared to an omniscient agent who would know the process ω that generates z). This limitation may be interpreted as a limit on skill: an agent who uses a more limited set of rules $R' \subset R$ has fewer abilities to exploit z than one using R . This limitation may also be interpreted as stemming from limited awareness: an agent may not consider rules because he is not aware that they might be useful.

To illustrate further, assume than an agent receives signals from two different sources, according to

$$z_1 = s + \alpha_1 \varepsilon_1 \text{ and } z_2 = s + \alpha_2 \varepsilon_2.$$

Define $r_{k,\lambda}(z) = z_k - \lambda$ as the rule that uses signal z_k and shades it by λ . Also define $\hat{r}_{k,\lambda,a}(z)$ as the rule that uses signal z_k , and shades it by λ only if $|z_1 - z_2| > a$ (and by 0 otherwise). An agent who considers:

$$R = \{r_{1,0}, r_{2,0}\},$$

is able to distinguish which source is most reliable. One who considers

$$R = \{r_{1,0}, \hat{r}_{1,\lambda_0,a_0}\},$$

cannot distinguish which of the two sources is more reliable, but he may adjust to events where α_1 or α_2 are likely to be high, and exert caution in these case.

²²Rational inattention models (see Sims (2003)) adopt that perspective. These models constitute an attempt to model agents who are imperfectly informed about the data generating process, but they keep the standard assumption that the actual data generating process is drawn from a distribution that the agent knows.

²³In comparison, the standard approach presumes that the agent can find the optimal rule across all possible rules r^e . In the standard model, modelling partial or limited information on ω thus requires that the agent distinguishes between rules belonging to a set larger than if he knew ω with precision.

Finally, an agent who considers:

$$R = \{r_{1,0}, r_{1,\lambda_0}\}$$

is focuses only on the first source. This can thought of a way to model an agent who is only skilled at exploiting the first source, or an agent who is unaware that the second source is relevant to assessing s .

To illustrate with a famous example, consider a situation in which more or less able detectives must decide, based on a series of indices, whether they ought to continue investigations thinking that (a) intrusion on the premise has occurred, (b) no intrusion has occurred, or (c) it is uncertain whether (1) or (2) is true or not. The underlying state $s \in \{I, N\}$ indicates whether intrusion has occurred or not, the data $z = (z^b, \dots)$ consists of a sequence of indices, among which $z^b \in \{0, 1\}$ indicates whether dog barking has been reported or not. A rule maps a sequence of indices to a decision $d \in \{a, b, c\}$. Detectives do not know the joint distribution ω over state and data, but they are able to compare a limited set rules. Differences in abilities (or awareness) across detectives can be modelled as differences in the set of rules that detectives consider. A detective who considers all possible mappings from data to decisions will figure that reports of dog barking indicates intrusion,²⁴ and he will also possibly deduce – depending on the probability that barking is reported when it occurs, that no report indicates no intrusion. A detective who considers fewer rules may miss the latter connection.

7.2 Perceptions and representations of uncertainty.

The traditional approach models uncertainty as a belief taking the form of a probability distribution over some underlying state space. That belief is either viewed as an endowment/perception (the agent comes to decision problems equipped with a belief) or as a representation (the agent acts as if he had that belief), but in applications, that distinction is immaterial, and the former view prevails.

Our perspective is that an agent does not necessarily come to decision problems with beliefs taking the form of a probability distribution over some underlying state space. They come to decision problems with potentially rich data, consisting of whatever the agent perceives as being relevant to the decision problem at hand, and this potentially rich data gets aggregated into a "perception". A "perception" may consist of a simple point estimate, but it may also include some rough idea of the confidence that the agent has in his estimate. In principle, a perception could also take the form of a probability distribution over states, or a set of probability distributions over states. In complex problems, we may model the way data gets aggregated into a crude perception using "mental states"²⁵ and sometimes not, as in this paper. In the latter case we directly endow the agent with a perception.

²⁴This is to the extent that there are no fake reports or that one does not confuse real dog barking with dreams of dog barking.

²⁵See, e.g., Compte and Postlewaite (2010, 2011).

The agent faces uncertainty in the sense that his "perception" (z) may not be an accurate description of the state (s) that the agent currently faces. Rather, a "perception" is a *signal* (imperfectly) correlated with the actual "state", and that signal may be used by the agent to make appropriate decisions. Our analysis assumes a joint distribution ω over "perceptions" and "states".²⁶

For any given perception z_0 , one could follow Savage's path and construct beliefs. If the agent were confronted with hypothetical choices (decisions over Savage acts), and if the agent's choices were to satisfy Savage's axioms, one could *represent* the choices made utilizing a belief taking the form of a probability distribution over states. That belief would be the *implicit representation of uncertainty* that would account for the agent's reported choices over Savage acts. We objected earlier, however, that choices may satisfy Savage axioms, and yet bear no relationship whatsoever with the true state s , nor with the perception z_0 at which these hypothetical choices were made.

In contrast, we assume that the agent is concerned with performance/welfare given the decisions he actually has to make, and that he uses perceptions as a guide to decision making.

Ideally, the agent would want to compare all possible ways to use perceptions, that is, all possible rules that map perceptions into decisions, and rank them. If he could do that, one could represent the choices made at any given perception through a belief taking the form of a probability distribution over states. That belief would be an implicit representation of the uncertainty that the agent faces, and for each perception z_0 , it would coincide with the Bayesian posterior $\omega(. | z_0)$.²⁷

Our approach limits the agent's ability to compare rules. One can define an implicit representation of uncertainty associated with the choices (over rules) made. However, since choices are more succinct (there are few rules), that representation does not take the form of a distribution over states. With two rules r_0 and r_1 , there are distributions $\omega \in \mathcal{P}_0$ for which r_0 turns out to be a better rule, and distributions $\omega \in \mathcal{P}_1$ for which r_1 turns out to be a better rule. Being able to identify which rule is better among r_0 and r_1 can be viewed as an implicit ability to identify whether the decision environment ω belongs to \mathcal{P}_0 or \mathcal{P}_1 . The (implicit) representation of uncertainty is not a detailed probability distribution over states, but rough properties of ω that make r_0 or r_1 the better alternative.

In Proposition 1 for example, $R = \{r_0, r_1\}$, with r_1 being a cautious rule, and we may refer to \mathcal{P}_0 as a set of "standard environments" and \mathcal{P}_1 as a set of "risky" environments. A1 assumes that the agent can identify which rule is best in R . This is equivalent to assuming that the agent can identify whether the environment is "standard" or "risky". The kind of environment the agent

²⁶This view of perceptions as signals applies *even if* perceptions take the form of probability distributions over states (or "beliefs" in standard terminology). We would then view "beliefs" as signals that agents need to employ appropriately to make decisions.

²⁷Again, that belief would not be a perception, it would simply account for the choices made.

faces (standard or risky) IS the way he (implicitly) assesses the uncertainty he faces.

The assessment of uncertainty is vague (as most of our uncertainty assessments are). Also, there is no presumption that this assessment is made prior to making the decision. Deciding that caution should be exerted is equivalent to assessing that the environment is risky.

7.3 Games

As for decision problems, the traditional approach to games presumes that agents come to strategic interactions endowed with a belief taking the form of a probability distribution over some underlying state space. A possible "state" corresponds as before to whatever is relevant to the game being played. In games however, the state space is more difficult to define because it potentially includes a description of the other player's beliefs: starting from a parameter space S , the relevant state space consists of the universal beliefs space generated by S (Mertens and Zamir (1985)).

Working on universal beliefs space being quite impractical, the standard practice follows Harsanyi (1967). This consists of summarizing a player's information as a type $t_i \in T_i$, with the understanding that a type t_i also determines his belief $\beta_i^{t_i} \in \Delta(T_{-i})$ over other player's type. Though Harsanyi starts with the subjective view, thus with beliefs possibly disconnected from the true "state", he argues for restricting attention to mutually consistent beliefs, that is, derived from a joint distribution $\omega \in \Delta(T)$ over type profiles and Bayes rule: $\beta_i^{t_i} = p(\cdot | t_i)$, and for having the type profile $t = (t_1, \dots, t_n)$ fully determining preferences over strategy profiles.

This model can be viewed as one of dispersed information: the true parameter space is T , and each agent sees a piece of the realized parameter t . But t need not coincide with the true parameter space, say S . Type profiles t and states s are then drawn from an objective joint distribution $\omega \in \Delta(T \times S)$. Players use ω to form beliefs according to Bayes rule: $\beta_i^{t_i} = \omega(\cdot | t_i) \in \Delta(T_{-i} \times S)$, and to determine, for each realized type profile t , *indirect* preferences over strategy profiles.

As already pointed out, two views are possible. Either ω is known and each player can use his own type t_i to form a belief and define his preferences. Or types are viewed as data that players use to find out optimal play: a strategy consists of a rule r_i that maps a type t_i to an action, and one assumes that he can find out the optimal rule r_i in the set of all rules. If he could do so, he would behave *as if* he knew ω , and the belief $\beta_i^{t_i} = \omega(\cdot | t_i)$ could then be viewed as a way to represent the uncertainty that the agent faces.

Our perspective is that it is not reasonable to assume that ω is known, and that with rich type spaces, it is not reasonable either to assume that determining the optimal rule within the set of all rules can be achieved. Our approach amounts to reducing the number of rules that the agent can compare. Types are interpreted as data that the agent can use to condition behavior. But he

cannot compare all possible ways to use data, and thus does not end up behaving as if he knew ω with great precision.²⁸

7.4 Mechanism design

Our simple decision problem (illustration 1) can be examined through the lens of mechanism design, as a game between the agent and a social planner whose preference coincides with that of the agent. The usual assumption is that both the agent and the social planner know the joint distribution ω over signals and states. Under that assumption, delegated and centralized mechanisms are essentially equivalent, and they result in implementation of the rule r^* that maximizes expected utility across all possible rules.²⁹

Our perspective is that neither the agent nor the social planner knows ω with precision: the set of rules R that an agent can compare implicitly determines what the agent knows about ω , or his skill in interpreting and using the signals he gets, or his cognitive ability. Most of our models and intuitions rely on people having unbounded cognitive ability, or unbounded ability to exploit the signals they get. In practice however, organizations must deal with people's limited cognitive abilities, and the design of organizations (who talks to whom, who decides) may actually be a response to these limitations.

Illustration 1: Assume that there are two types of agent, 1 and 2 (with observable characteristics that enable the social planner to distinguish them), and that they both observe the same signal z . Also assume that type 1 uses R_1 while type 2 uses R_2 . The social planner considers delegating the decision to either agent 1 or to agent 2. If $R_2 \subset R_1$, then the social planner (and the agents) will be better off if the decision is delegated to agent 1. If $R_2 \not\subset R_1$, with no inclusion, then it is not a priori obvious to whom the decision ought to be delegated. It could be that for $\omega \in \mathcal{P}_2$, R_2 generates higher welfare, while R_1 generates higher welfare when $\omega \in \mathcal{P}_1$. Being able to identify whether the decision ought to be delegated to agent 1 or 2 amounts to being able to determine whether $\omega \in \mathcal{P}_1$ or $\omega \in \mathcal{P}_2$.

Issues regarding how to exploit signals arise more forcefully with more agents, or when signals arise from different sources. In team problems in which agents

²⁸One could insist that the only data that agents use are beliefs. This would amount to defining each possible realized belief β_i as a possible type $t_i : t_i \equiv \beta_i$, and then assume as above a joint distribution ω over types and states. An agent who could determine the optimal way to use his data t_i would in the end behave as if he knew ω , that is as if he had belief $\beta_i^{t_i} = \omega(\cdot | t_i)$, and maximize expected utility with respect to that belief $\beta_i^{t_i}$ (rather than with respect to β_i).

²⁹Under delegation: if the decision is delegated to the agent, he will find it optimal to choose $r^*(z)$ when he observes z .

Under communication without commitment: If the decision process is that the agent reports \hat{z} and then the social planner decides, then there is an equilibrium in which the agent reports the truth $\hat{z} = z$ and the social planner chooses $r^*(\hat{z})$. (Alternatively, one could assume that the agent reports both \hat{z} and $\hat{\omega}$.)

Under communication with commitment: if the social planner commits to taking $r(\hat{z})$ when the agent reports z , then, to the extent that r is invertible, it is optimal for the agent to report $\hat{z} = r^{-1}(r^*(z))$.

have identical preferences but possibly divergent opinions, standard theory sees no difficulty in aggregating those opinions. Since the joint distribution over states and signals/opinions is assumed to be known, standard models sweep that issue away, focusing on incentive issues.³⁰

Illustration 2: Consider a team problem in which agents $i = 1, 2$ have identical preferences but each agent i has a distinct perception z_i of the appropriate decision s . The distribution ω now refers to the joint distribution over (z_1, z_2, s) . With the usual assumption that players and the social planner know ω , there is no difficulty in exploiting the signals z_1 and z_2 and implementing the rule $r^*(\cdot, \cdot)$ that maximizes expected utility across all possible rules.³¹ In particular, if the decision is delegated to 1, performance is improved if agent 2 first communicates to agent 1: If we wish to compare mechanism D1 where the decision is fully delegated to agent 1 (he then exploits signal z_1), to mechanism C1 where player 1 decides after player 2 reports his signals (player 1 then exploits (z_1, \hat{z}_2)), the standard view is that C1 cannot do worse than D1, because player 1 can always ignore player 2's report.

From our perspective however, one cannot presume that ω is known perfectly, so to each mechanism, one needs to associate rules that agent 1 compares, and these rules implicitly define what player 1 knows about ω . In particular, if player 1 is limited in the total number of rules that he can compare, a trade-off may appear, as he may end up exploiting z_1 with more skill if he only focuses on signal z_1 , while the returns from trying to exploit simultaneously z_1 and z_2 may be small, because there are so many ways z_1 and z_2 could be combined. As a result, for a possibly large range of distributions ω , D1 may be optimal.

7.5 Wilson's critique.

We often refer to Wilson (1987) to acknowledge our unease with common knowledge assumptions in game theoretic interactions, in particular when they involve "one agent's probability assessment about another's preferences or information". That uneasiness should not be restricted to games; it ought to arise even in decision problems, when one assumes that an agent knows (or acts as if he knew) the joint distribution over states and signals.

The uneasiness stems from two reasons: (i) lack of realism of the assumption, and more importantly, (ii) lack of realism of the consequences of the assumption, with behavior adjusting very finely to details of the environment. The drastic

³⁰There are historical reasons for this. Mechanism design originates from the work of Hurwicz, and comparisons between market and centralized mechanisms. When private information refers to information about own preferences only, then, given a social objective, aggregation of information is not an issue: only transmission of information matters.

³¹For example, if player 1 reports \hat{z}_1 to player 2, who then decides, then there exists an equilibrium where player 1 reports $\hat{z}_1 = z_1$ and player 2 follows r^* . If the central planner commits to following $r^*(\hat{z}_1, \hat{z}_2)$, there exists an equilibrium in which both players report their true signals. Note however that if the central planner commits to some rule $r(\hat{z}_1, \hat{z}_2) \neq r^*(\hat{z}_1, \hat{z}_2)$, then in general there cannot be reporting strategies $\hat{z}_i = \phi_i(z_i)$ that implement r^* : in other words, it is essential that the central planner knows the appropriate way to aggregate the signals z_i .

simplification that we propose – looking at few alternative rules – partitions the set of environments ω into few subsets, and small changes in the environment ω "rarely" induce changes in the rule used (and consequently, changes in behavior).

Following Wilson's critique, the literature on robust mechanism design and on robust equilibrium analysis (in particular the work by Bergemann and Morris (2012)) attempts to weaken the common knowledge assumption that we usually make when analyzing games of incomplete information. The route follows Harsanyi's suggestion, and it consists of working with a richer type space. Specifically, in the face of possibly imperfect knowledge of the joint distribution ω over states and signals, the traditional route consists of adding one extra layer of uncertainty: one perturbs the model by assuming some randomness over the actual distribution ω and by adding signals correlated with the realized ω . One thus obtains an enriched information structure, characterized by a new joint distribution, say $\bar{\omega}$, over states and signals, and it is this enriched or perturbed game that is analyzed, under the presumption that this more sophisticated signal structure would be commonly known, with the hope that behavior would not be too sensitive to these perturbations.³²

As an internal test to the standard approach (which builds on a strong common knowledge assumption), the robustness test above is clearly a legitimate one, as we would not want equilibrium predictions to depend finely on knowledge that seems out of reach (not only to an outside observer, but even to players themselves).³³ The route we propose is not an internal test to the standard approach, but an alternative way to deal with player's lack of detailed knowledge about the distribution ω they happen to face.

7.6 Concluding remarks

Heterogenous beliefs. The traditional view is that differences in beliefs may only stem from differences in information. One implication of our approach is that differences in assessment of uncertainty may stem from differences in the rules that agents contemplate. Consider for example two individuals, $i = 1, 2$, each using $R_i = \{r_0, r_{\gamma_i}\}$ with $\gamma_1 \neq \gamma_2$. These sets of rules induce different partitions of the set of environments: $(\mathcal{P}_0^i, \mathcal{P}_{\gamma_i}^i)$, hence despite facing the same environment ω , they may end up with different assessments of the uncertainty they are facing: a standard situation for one, a risky situation for the other.

Ad hoc assumptions. The standard model is also ad hoc in the following sense. There is a joint probability distribution ω on states and signals for which no uncertainty is permitted. In the spirit of Example 2, one might want to relax that assumption, and add some uncertainty in the way ω is generated.

³²Note that strictly speaking, enriching the type space does not weaken the common knowledge assumption. It replaces it with another that might be viewed as even less plausible.

³³Bergemann and Morris (2011, page 4) argue: "the absence of the observability of the belief environment constitutes a separate reason to be skeptical towards an analysis which relies on very specific and detailed assumptions about the belief environment".

We share that view. However we view that it applies not only to outsiders observing the interaction, *but also to players themselves*.

This is done at the cost of enriching the set of signals, but the problematic assumption remains. One may go one step further in the hierarchy of beliefs about distributions, but eventually, we have to close the model and model as certainty what seems most doubtful.

This is a highly unrealistic assumption, but we do not object to it because we believe it helps us shape our intuition. We hope that as we move to more challenging environments than the one described here, the value of restricting choice to few decision rules will be judged by the same criterion.

References.

- Barillas, F. and T. J. Sargent, "Doubts or Variability?" *Journal of Economic Theory*, Volume 144, Issue 6, November 2009, pp. 2388-2418.
- Bergemann, D and S. Morris (2012), *Robust Mechanism Design*. World Scientific Publishing, forthcoming January 2012.
- Bergemann, D and S. Morris (2011), Robust Predictions in Games with Incomplete Information, Discussion paper n°1821, Cowles Foundation
- Block, H. D. and J. Marschak, (1960), "Random orderings and stochastic theories of responses", in I. Olkin et al., (Eds.), *Contributions to probability and statistics: Essays in honor of Harold Hotelling*. Stanford University Press, Stanford, pp. 97-132.
- Compte, Olivier and Andrew Postlewaite (2010), "Plausible Cooperation," mimeo.
- Compte, Olivier and Andrew Postlewaite (2011), "Belief Formation," mimeo.
- Crawford, Vincent P. and Joel Sobel (1982), "Strategic Information Transmission," *Econometrica*, vol. 50(6), pages 1431-51.
- Ellsberg, Daniel (1961), "Risk, Ambiguity and the Savage Axioms", *Quarterly Journal of Economics*, 75: 643-669.
- Fox, Craig R. and Amos Tversky (1995), "Ambiguity Aversion and Comparative Ignorance." *The Quarterly Journal of Economics*, Vol. 110, No. 3.
- Geanakoplos, J. (1989) Game Theory Without Partitions, and Applications to Speculation and Consensus, Discussion paper n. 914, Cowles Foundation
- Gilboa, Itzhak, Andrew Postlewaite and David Schmeidler (2009), "Is It Always Rational to Satisfy Savage's Axioms?," *Economics and Philosophy*, 25: 285-296.
- Gilboa, Itzhak, Andrew Postlewaite and David Schmeidler (2008), "Probability and Uncertainty in Economic Modeling", *Journal of Economic Perspectives*, 22(3):173-188.
- Gilboa, Itzhak, and David Schmeidler (1989), "Maxmin Expected Utility with a Non-Unique Prior", *Journal of Mathematical Economics*, 18, 141-153.
- Gilboa, Itzhak, and David Schmeidler (1995), "Case-Based Decision Theory," *Quarterly Journal of Economics*, 110(3): 605-639.
- Halevy, Yoram (2007), "Ellsberg Revisited: an Experimental Study," *Econometrica*, 75 (2) 503-536.
- Hansen, L. and T. Sargent (2001), "Robust Control and Model Uncertainty," *American Economic Review* 91(2): 60-66.

- Harsanyi, J. (1967), Games with Incomplete Information Played by ‘Bayesian’ Players, I-III,” *Management Science*, 4(3): 159-182.
- Jehiel, P. (2005), “Analogy-Based Expectation Equilibrium,” *Journal of Economic Theory* 123(2): 81-104.
- McFadden, Daniel (1973), "Conditional Logit Analysis of Discrete Choice Behavior", in *Frontiers in Econometrics*, Academic Press, New York, pp. 105-142.
- McKelvey, R. and T. Palfrey (1995), “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior* 10(1): 6-38.
- Mertens, J-F and S. Zamir (1985), “Formulation of Bayesian analysis for games with incomplete information,” *International Journal of Game Theory* 14(1): 1-29.
- Osborne, Martin J. and Ariel Rubinstein, Ariel (1998), "Games with Procedurally Rational Players," *American Economic Review*, vol. 88(4), pages 834-47.
- Savage, Leonard J. (1954), *The Foundations of Statistics*. New York: John Wiley and Sons. (Second edition in 1972, Dover.)
- Schmeidler, David (1989), “Subjective Probability and Expected Utility without Additivity”, *Econometrica*, 57: 571-587.
- Sims, C. (2003), “Implications of Rational Inattention,” *Journal of Monetary Economics* 50(3): 665-690.
- Thurstone, L. L. (1927), "Psychophysical Analysis", *The American Journal of Psychology* Vol. 38, No. 3, pp. 368-389
- Tversky, A. and D. Kahneman (1974), "Judgment under uncertainty: Heuristics and biases", *Science*, vol. 185, pp1124-31
- Wilson, R. (1987), "Game-Theoretic Analyses of Trading Processes," in: *Advances in Economic Theory: Fifth World Congress*, Truman Bewley (ed.); Chapter 2, pp. 33-70. Cambridge UK: Cambridge University Press.