

Discussion Paper No. 1249

**SUBJECTIVE REPRESENTATION
OF COMPLEXITY**

By

Nabil I. Al-Najjar*

Ramon Casadesus-Masanell*

and

Emre Ozdenoren*

November 1998

Revised: January 1999

<http://www.kellogg.nwu.edu/research/math>

(*) Department of Managerial Economics and Decision Sciences, J.L. Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL, 60208. E-mail addresses: <al-najjar@nwu.edu>, <rcasades@kellogg.nwu.edu>, <eozenor@skew2.kellogg.nwu.edu>.

We are grateful to Peter Klibanoff, Tim van Zandt and Chris Sanchirico for very helpful comments.

SUBJECTIVE REPRESENTATION OF COMPLEXITY

by

NABIL I. AL-NAJJAR*

RAMON CASADESUS-MASANELL*

and

EMRE OZDENOREN*

November 1998

Revised: January 1999

(*) Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL, 60208. E-mail addresses: <al-najjar@nwu.edu>, <rcasades@kellogg.nwu.edu>, <eozenor@skew2.kellogg.nwu.edu>.

We are grateful to Peter Klibanoff, Tim van Zandt, and Chris Sanchirico for very helpful comments.

Abstract:

We study how individuals cope with the complexity of their environment by developing subjective models, or representations, to guide their predictions and decisions. Formally, an individual who believes his environment is deterministic, but too complex to permit tractable deterministic representation, builds a probabilistic model embodying perceived regularities of that environment. In this model, the individual's inability to think through all possible instances of the problem is represented by an uncertainty about random states. The resulting behavior is fully rational in the traditional sense, yet consistent with an agent who believes his environment is too complex to warrant precise planing, forgoes finely detailed contingent rules in favor of vaguer plans, and expresses a preference for flexibility. We consider applications to time-inconsistent preferences, delegation, and two-player simultaneous games.

1. INTRODUCTION

One remarkable aspect of human behavior is individuals' ability to cope with complexity. The role of complexity in everyday life is almost self-evident: people engage in complex social and economic interactions where it is pointless to examine each of the vast set of possible contingencies. Yet somehow people manage to formulate coherent, goal-seeking plans of action. How individuals confront situations they recognize as too complex to permit a full analysis remains a key problem in cognitive sciences, and it is of obvious importance in understanding economic and strategic behavior.

This paper studies how individuals cope with complexity by developing models, or representations, to guide their predictions and decisions. We consider a rational agent facing a 'problem', a loose collection of situations or instances he views as sharing a common structure. Examples of problems are: taking an appropriate action on behalf of a superior, adjudicating a legal case, categorizing an object, diagnosing a disease, ... etc. The agent never gets to directly confront a problem at this level of abstraction; rather, he faces a series of randomly drawn, specific instances of that problem (*e.g.*, a specific legal case, specific patient). Each instance is therefore unique, with special idiosyncratic aspects not shared by others.

We consider behavior displaying: (1) coherent assessment of the average (expected) performance of plans of action; (2) *ex post*, once told which instance he is actually facing, it is obvious what the optimal action should be; nevertheless, (3) it is not possible to describe a full *ex ante* rule specifying what the optimal action should be in every possible contingency. Far from odd or pathological, this seemingly contradictory behavior is pervasive in decision making, arising in planning in intertemporal settings, devising rules for defining classes of objects, categorization and pattern recognition (see Section 2).

We propose to explain such behavior in terms of an agent who believes his environment is deterministic, but may be too complex to permit tractable deterministic representation. To see the intuition, an agent facing a problem may initially be tempted to exhaustively enumerate all data, *i.e.*, an instance-by-instance enumeration of his utility, hence optimal choice. The vastness of the

set of possible instances quickly makes such crude representation unfeasible, so the agent is likely to seek simpler representations to exploit any useful patterns or regularities he perceives in the data. For example, the agent may believe that one particular feature completely determines his utility. Although the number of possible instances may be unfathomably large, the agent would have no difficulty providing a simple representation in terms of that single relevant feature.

Complexity (as perceived by the agent and reflected in his choice behavior) arises not because of the number of possible instances per se, but because of the number of *independent* pieces of information that must be assimilated and integrated in the decision process. But an agent able to formulate coherent plans must perceive enough regularity in his environment to evaluate the average performance of these plans. This suggests an agent who thinks in terms of generalizations embodying useful perceived patterns, yet recognizes the limits of these patterns, that they miss instance-specific variations accounted for only by essentially enumerating all data. The agent believes these variations have no further useful structure; in short, to him they ‘look random’.

Formally, our main result formulates behavioral assumptions that identify an agent’s subjective representation of his deterministic (and, in principle, knowable) environment in terms of random states reflecting his inability to think through all instances. In other words, the agent (behaves as if he) thinks of his cognitive *limitations* due to the complexity of his environment as cognitive *uncertainty* about a suitably chosen state space.

The behavior we describe is that of a fully rational agent in the traditional sense. He optimizes given a coherent model of the environment; he understands and takes advantage of all the implications of this model, unhampered by cognitive limitations that prevent him from carrying out reasoning or drawing inferences that we, as modelers, can perform. In particular, the agent is able to break up problems into simpler sub-problems and condition his plans of action accordingly without incurring any cost for doing so.

Yet behavior displays many features traditionally considered the hallmark of ‘bounded rationality’. For instance, there is a discontinuity between the ex ante and the ex post perspectives: the agent

may think the environment is too complex to warrant precise planing, yet once the particular instance he is facing is known, what to do becomes ‘obvious’. He may then forgo finely detailed planning for future contingencies, choosing to rely instead on coarser and vaguer plans of actions, which he fills in and completes as events unfold. Such behavior would display a preference for flexibility, suggested by Kreps (1992) as one of the key implications of an agent’s inability to foresee future contingencies—see the discussion below.

Our model generates this sort of behavior because it may be interpreted as capturing the limiting properties of environments with increasing complexity but decreasing computational/cognitive cost of gathering and processing information (Section 2). The limiting behavior (which is what we formally model) is that of a fully rational agent who confronts a vastly more complex environment. This agent optimizes, but he does so against a subjective representation that captures perceived patterns of his environment, rather than detailed, instance-by-instance examination.

Since agents are rational, their behavior displays the internal consistency of traditional models of rational behavior. Closed equilibrium models displaying interesting cognitive limitations are possible, without resorting to exogenously imposed limits on agents’ cognitive abilities, hard-wiring repetitive tasks, and introspection. This is a considerable advantage in multi-agent settings because it allows predictions that do not hinge on requiring agents to be dumb, ignore potentially useful information, or use ad hoc rules that the modeler or a rival could easily exploit. Another advantage of this approach is that many of the theoretical tools and concepts about individual decision making and game theory extend in a natural way. In Section 6, we illustrate this by considering how our model can be used in examples of time-inconsistent preferences, delegation, and two-player simultaneous games.

This paper relates to several strands of literature. In particular, it is related to the recent literature emerging from Kreps’ interpretation of preference for flexibility (Kreps (1979, 1992)) in terms of an agent’s inability to foresee future contingencies. Important works in this vein include Nehring (1996).

and Dekel, Lipman, and Rustichini (1997).¹ We discuss the link with this literature in Section 4.4-5.

The idea that a probabilistic model may be an effective way to represent deterministic but complex phenomenon is, of course, not new. For example, Lipman's (1995) representation of an agent's perception of the complexity of the logical implications of his assumptions about the world may be interpreted in this manner. While similar in flavor, our setups and results obviously differ. The point that plans of action should correspond to procedures that can be effectively carried out (formally, they must be algorithmic) was made by several authors, including McAfee (1984) Binmore (1987), Gilboa and Schmeidler (1994), and Anderlini and Felli (1994). See Section 4.5 for more discussion.

Complexity, of course, is a central issue in cognitive sciences.² Our formal model of individuals facing instances of a problem, each identified with a potentially large collection of features, shares many similarities with learning models in the pattern recognition literature (*e.g.*, Devroye, Györfi, and Lugosi (1996)). Our focus, however, is quite different. Unlike these works, we do not attempt to develop procedural models of decision making (*i.e.*, answering questions like: 'how do people think?'), or focus on the convergence properties of various classes of algorithms. Rather, our approach is to take an agent's behavior in an economic environment as given, and use it to infer how the agent sees his world, what he considers complex and simple, and how this impacts on his decisions.

The paper is organized as follows. Section 2 provides a simple example containing the main ideas of the paper. Section 3 describes the formal setup of our model, while Section 4 contains the definition of subjective representation of complexity and proves the main representation theorem. Section 4 also discusses preference for flexibility and the relationship to Kreps' model. Section 5 considers two important extensions of the basic model. Section 6 discusses applications and Section 7 concludes. All proofs are contained in the appendix.

¹ See Dekel *et al.* (1998) for a survey of this literature.

² See, for example, Medin and Ross (1992) and Holland *et al.* (1989).

2. MOTIVATION AND EXAMPLES

This section is independent from the rest of the paper, and may be skipped without loss of continuity; it provides informal motivation of the main points developed in this paper.

An agent faces *instances* of a problem, each defined in terms of the values of objectively given features $\{1, \dots, I\}$. We focus here on a special type of problem which involves devising a rule to categorize each instance by placing it into one of a predetermined set of categories $B = \{b_1, \dots, b_K\}$.

Categorization is a basic and pervasive aspect of decision making. Mundane tasks in pattern recognition (*e.g.*, Is a given object a chair or a table?) are examples of categorization. More complex examples may be found in legal, managerial, and problem solving contexts. In a sense, categorization appears in virtually every decision process because decisions tend to be made based on concepts and categories, rather than raw data.³

What makes a concept or category complex? Consider the problem of devising a rule for categorizing legal cases based on legal concepts, *e.g.*, perjury, breach, or obscenity. Such decision would have to be based on the cases' objectively given features (details of evidence, parties' actions and statements, and so on). Imagine an agent able to classify each specific instance, but has a hard time exhibiting the rule followed in his classification. A well-known example is Justice Stewart's famous statement confessing that he is unable to come up with a good definition of obscenity, "But I know it when I see it".⁴ This behavior is not odd or pathological: people often have no difficulty performing mundane tasks, like telling whether an object is a chair or a table, yet have considerable difficulty devising a corresponding definition. More relevant to economic applications are contractual settings in which it is obvious to the parties what should be done *ex post* when facing a specific instance, yet they have hard time codifying such understanding explicitly *ex ante*.

Before proceeding, it is useful to note the key features of such behavior. First, there is a dis-

³ This is the prevailing view in cognitive sciences; see Medin and Ross (1992), and Holland *et al.* (1989).

⁴ "I shall not today attempt further to define the kinds of material I understand to be embraced within that short-hand description; and perhaps I could never succeed in intelligibly doing so. But I know it when I see it." *Jacobellis vs. Ohio*, 378 U.S. 184, 197 (1964) (Stewart, J., concurring).

continuity between the ex post problem, when the agent is facing one specific instance, and the ex ante problem where he has to provide a recipe covering all conceivable instances. We interpret the behavior “I cannot define it, but I know it when I see it” as evidence that the agent considers the ex ante problem an order of magnitude more complex than the ex post problem. Second, a categorization rule predicts the instance’s category based on objectively given features, not whim, intuition, or subjective criteria. In applications like writing a contract or formulating instructions to subordinates, it is natural to require rules to constitute *procedures* that can, at least in principle, be codified, communicated to others, or used in evaluating performance. Third, an agent may well be able to evaluate the average performance of a rule or definition, even though he knows it is impossible to make an instance-by-instance evaluation of such rule. People have little difficulty judging one definition of an object, say a ‘chair’, to be better than another, without going over all conceivable objects to check the fitness of these definitions. In the same vein, an individual may be able to judge one incomplete contract as better or worse than another, without evaluating the full implications of each contract in every conceivable contingency.

This paper develops a model that explains this and other related behavior in terms of the agent’s subjective representation of the perceived complexity of his environment. The underlying intuition is easily explained in the following special setting. Assume, without loss of generality, that each feature is binary (*i.e.*, takes values of either 0 or 1), so the space of instances, X , consists of 2^J individual instances. The agent’s belief that ‘he knows it when he sees it’ is a function $b : X \rightarrow B$, where $b(x)$ gives the ‘true’ category of x . Data is generated by random draws of instances according to a distribution λ on X .⁵ We require λ to have a diffuse enough support (*e.g.*, λ is uniform on X) so ex ante each instance has vanishing probability of being encountered.

The agent’s problem is to devise a ‘tractable’ rule $f : X \rightarrow B$ to match the true category $b(x)$ as closely as possible. How difficult is this task? If b can be pinned down using a small number of features, then it can be easily replicated by a rule f , regardless of the size of the space of instances X .

Complexity arises when details matter: If all instance-specific information is potentially relevant, it

⁵ This rules out active experimentation. Examples are a doctor who has no choice about which patient to treat, or a judge who has no choice about the cases assigned to him.

may be impossible to compress this information into a simple sufficient statistic. There is a standard way to make this formal, the Kolmogorov complexity measure, defined as the length of the shortest program that can reproduce b .⁶ Roughly, b is complex if describing it requires a nearly exhaustive enumeration of its values. A basic fact is that complexity is the norm: most functions are complex in this sense.

How would an agent cope with hard problems? Exhaustive enumeration is clearly pointless: the number of possible instances, 2^I , increases exponentially with I , so no conceivable amount of resources can capture a complex b , even for moderate values of I . If the agent resorts to rules whose complexity is bounded by T , then the impossibility of exhaustive enumeration requires T to be substantially less than 2^I . This is just saying that our agent faces the familiar ‘curse of dimensionality’: the problem has too many potentially relevant dimensions relative to available resources and data.

Our agent gives the problem his ‘best shot’, identifying the significant patterns or regularities he perceives. But patterns necessarily lump together otherwise distinct instances, leaving some residual variability unaccounted for. The agent, who recognizes these limitations, copes with this residual variability by developing a subjective, probabilistic model in which states represent potentially important variations too complex to warrant detailed examination. The agent uses this subjective model as his guide for predictions and decisions. In particular, he uses it to compute the expected payoff of rules without having to examine what these rules imply on an instance-by-instance basis. However, his recognition of imperfections in his model implies behavior that may display preference for flexibility as well as the discontinuity between the ex ante and ex post problems mentioned above.

Our formal model in fact considers the limiting case of the above informal story. Specifically, we take the set of features to be countable and impose no exogenous bounds on the procedures used by the agent. This limiting model offers not only a clean and tractable statement of these ideas, but also a way to capture agents fully rational in the traditional sense as the limit of nearly rational agents in a vastly more complex environment.

⁶ See Cover and Thomas (1991) chapter 7 for an introduction, and Theorems 7.2.4 and 7.5.1 for results showing that most functions are complex.

3. THE MODEL

A *problem* Q consists of the primitives $(X, \lambda, B, V, \mathcal{A})$, where X is the set of *instances*; λ is the probabilistic process generating the data; B is a finite set of *actions*; V is a finite set of *ex post* utilities; and \mathcal{A} is an algebra of *conditioning events*. We formally describe each one of these elements below.

3.1. Instances

An *instance* is a collection of ‘objectively’ given data about the problem. In light of the example of Section 2, we want to allow for the possibility that the agent believes that the number of relevant features may be unbounded. That is, we do not want to rule out an agent belief that the problem he is facing is subject to a ‘curse of dimensionality’: for any set of features the agent considers, there may be other relevant dimensions he hasn’t considered.

Formally, then, there is a countable number of features, with the i th feature \tilde{x}_i taking values in the set X_i . For notational simplicity, we assume that each X_i is binary (that is, $\tilde{x}_i \in X_i$ takes one of two possible values). The set of instances of the problem is therefore the infinite product $X = \prod_i X_i$: with each instance corresponding to a sequence (x_1, x_2, \dots) of values of the features.⁷

3.2. Actions and ex post Utilities

There is a finite set of actions B , and a finite set, V , of utility functions $v : B \rightarrow \mathbb{R}$, which we later interpret as ex post rankings of actions.

An example is the class of *categorization problems*, where an agent has to place an object into one of K possible categories discussed in Section 2. Here the set of actions is $B = \{b_1, \dots, b_K\}$, where b_k

⁷ Sometimes it may be useful to think of X as the set of binary expansions (sequences of 0’s and 1’s) of numbers in the interval $[0,1]$, though this may be misleading. We attach no particular meaning to the ordering of features, so the ‘location’ of an instance x on the interval is of no relevance.

represents ‘place the instance in category k ’. The set of ex post utilities is $V = \{v_1, \dots, v_K\}$, where v_k is a K -vector, with value 1 at the k th entry, and 0 elsewhere. We interpret v_k as the utility if the instance is truly of category k , in which case a match (categorizing it correctly) yields a payoff of 1, while a mismatch yields 0.

3.3. Breaking up Problems into Sub-problems

In tackling a given problem, it is natural to think that the agent breaks it up into smaller, simpler sub-problems—presumably based on common properties of the underlying instances. We shall think of this as a finite sequential *procedure*, consisting of an algebra \mathcal{A} of sub-problems, and a *tree* T , representing the sequence followed in breaking up the problem into these sub-problems.

Before describing \mathcal{A} and T formally, we emphasize that a procedure represents an *effective* way of breaking-up the problem: it must consist of steps that can be, at least in principle, codified, written down, communicated, and instructed to subordinates. This rules out ‘procedures’ based on subjective criteria that cannot be made explicit, whim, gut-feeling, or oracles.

The algebra of *conditioning events*, \mathcal{A} , consists of all subsets $A \subset X$ of instances for which there is an ‘effective’ method, formally an algorithm α (*i.e.*, a Turing machine, or a computer program), to determine membership in A .⁸ For a conditioning event $A \in \mathcal{A}$, the corresponding *sub-problem* Q_A is the original unconditional problem Q , except that the agent may condition his actions on the knowledge that he will face instances drawn from A . We will often abuse the definition and not distinguish between A and Q_A .

Next, the sequence followed in breaking up the problem is modeled as a finite tree with nodes $T = \{t_0, t_1, \dots, t_N\}$, a precedence relation \rightarrow , and root t_0 . Informally, we consider an agent who, starting with Q_A , partitions A into a finite number of conditioning sub-events, takes each one of these

⁸ That is, the characteristic function of A , χ_A , is algorithmic: there is an algorithm α such that $\chi_A(x) = \alpha(x)$ for every x . See Gilboa and Schmeidler (1994) for more detailed formal definitions of what an algorithm formally means in this context.

and further partition it, and so on. More formally, let $S(t)$ denote the set of immediate successors of node t , and let Z denote the set of terminal nodes. Then a *decomposition* of a sub-problem Q_A by a tree T (more precisely, (T, \rightarrow, t_0)) is a function $\gamma : T \rightarrow \mathcal{A}$ such that (writing $A(n)$ for $\gamma(t_n)$): (1) $Q_{A(0)} = Q_A$ (the tree has root A); and (2) for every $t \in T$, $\{A(t') : t' \in S(t)\}$ is a partition of $A(t)$. Of interest is the *terminal partition* generated by a tree, defined as $\{A(t) : t \in Z\}$.

3.4. Behavior, Rules and Options

We now describe the setting in which the agent confronts the problem. Let $\bar{\mathcal{A}}$ be the σ -algebra generated by \mathcal{A} (Appendix A.1 provides characterizations of $\bar{\mathcal{A}}$). We assume that λ is non-atomic and has full support on X .⁹

For reasons that will become transparent later, we model the agent's decisions as a sequential problem:

Ex ante stage: *The agent believes he will face an instance drawn at random from X according to the probability distribution λ .*

Interim stage: *A specific instance x is drawn; the agent knows x but does not yet know his utility.*

Ex post stage: *The agent finds out his utility $v \in V$.*

Continuation: *Having observed the pair (x, v) , the agent faces the ex ante problem again.*

The 'continuation' stage introduces a dynamic aspect useful in the representation theorem and in discussing dynamic choice. For the moment, however, it may be ignored.

⁹ The support assumption seems innocuous, since we can re-define X to be the support of λ . Non-atomicity is more substantial; it formalizes the idea that each instance is unique, with zero probability of ever being encountered. Our main results are preserved if we allow atoms—although their statements would be more cumbersome. What is essential is that λ has a non-trivial non-atomic component.

At the ex ante stage, the agent must devise a plan to deal with all conceivable instances, which we model as *options*. Let \mathcal{C} be the set of all non-empty subsets of B . An option is any function $g : X \rightarrow \mathcal{C}$ measurable with respect to \mathcal{A} . The interpretation of an option g is: “you may wait until an instance is drawn, then choose any action b you like subject to the constraint $b \in g(x)$ ”. It is important to note that the option must be exercised at the interim stage, *i.e.*, after x is drawn, but before the utility itself is known.¹⁰

An important subset of options are those that involve no flexibility: a *rule* is any function $f : X \rightarrow B$ measurable with respect to \mathcal{A} . Unlike options, a rule is an explicit, rigid plan prescribing the precise action to take in each possible instance. Let \mathcal{F} denote the set of rules, and \mathcal{G} the set of options. Obviously, $\mathcal{F} \subset \mathcal{G}$; we refer to $\mathcal{G} - \mathcal{F}$ as the set of non-trivial options. When clear from the context, we use b and C to denote the constant rule $f(x) = b$ and option $g(x) = C$ respectively.

Note that options can be used to represent varying degrees of flexibility (*e.g.*, an option allowing full flexibility ($g(x) = B$) for some instances, but much less flexibility elsewhere). Property rights, discretionary powers, and laws, involve varying degrees of discretion, ranging from completely vague, option-like plans, to strict rule-like restrictions.

¹⁰ Later (Section 4) we consider other types of options that differ by the time at which they may be exercised.

4. UTILITY AND CHOICE

4.1. Behavioral Assumptions

We describe the agent's behavior in terms of assumptions on his von Neumann-Morgenstern utility in the hypothetical decision problem described earlier. To focus on complexity considerations, we take as primitive the agent's utility which, in the context of our experiment, consists of an ex ante utility U_0 reflecting his choices in the ex ante stage, and a continuation utility U_1 for the continuation stage.¹¹

The ex ante utility has the form $U_0 : \mathcal{G} \times \mathcal{A} \rightarrow \mathbb{R}$, where $U_0(g, Q_A)$ represents his payoff from an option g given a sub-problem Q_A . Let L be the set of probability distributions on $V = \{v_1, \dots, v_M\}$. Our first assumption says that, when faced with a sub-problem Q_A , the agent thinks of it as a lottery over ex post utilities:

A.1: Lottery equivalence: For every $A \in \mathcal{A}$ there is a lottery $l = (l_1, \dots, l_M) \in L$ such that, for every $C \in \mathcal{C}$,

$$U_0(C, Q_A) = l_1 \max_{b \in C} v_1(b) + \dots + l_M \max_{b \in C} v_M(b).$$

The sense in which this equivalence holds is important. First, it assumes that the agent has von Neumann-Morgenstern utility over lotteries. Second, the max is taken with respect to the ex post utilities. This implicitly says that once the agent sees the realized instance x , he knows what to do. We relax this in Section 5.1.

Under lottery equivalence, the agent knows he will be facing one of finitely many possible utilities $v \in V$, and has beliefs about their probabilities, assessed for the entire set of instances A . It is natural to allow the agent to devise more elaborate plans based on more refined sorting of instances. We

¹¹ A more complete treatment would derive the agent's von Neumann-Morgenstern utility from more primitive assumptions about his preferences.

model this by assuming the agent uses trees to refine Q_A (see Section 3.3). Our next assumption puts some structure on the way the agent evaluates the outcome of such procedures:

A.2: *Reduction of Compound Sub-Problems:* For any $A \in \mathcal{A}$ with $\lambda(A) > 0$, any tree T with root A and terminal partition $\{A_1, \dots, A_N\}$, and any option g ,

$$U_0(g, Q_A) = \sum_{i=1}^N \lambda(A_i | A) U_0(g, A_i)$$

This says that the agent thinks that if A is broken into $\{A_1, A_2\}$, say, he views Q_A as equivalent to a lottery over the sub-problems Q_{A_1} and Q_{A_2} with weights given by λ . Stated differently, the agent views the uncertain prospect (g, Q_A) as a two-stage lottery in which a sub-problem Q_{A_i} is drawn first according to λ , then faces the restriction of g to whichever sub-problem has been drawn.

Our final assumption concerns the continuation utility which, like the ex ante utility, evaluates option/sub-problem pairs, but taking into account the outcome of one round of facing the problem. This outcome takes the form of an instance/utility pair (x, v) , so the continuation utility takes the form $U_1 : \mathcal{G} \times \mathcal{A} \times X \times V \rightarrow R$, where $U_1(\cdot, \cdot; x, v)$ is interpreted as the agent's cardinal valuation of option/sub-problems given that he observed an instance x and utility v .

We shall focus on the benchmark case of a long-run, steady-state where the agent believes he has learned everything there is to be learned from repeated encounters with this problem. In Section 5.2 we discuss how to relax this assumption.

A.3: For every outcome (x, v) , the continuation utility $U_1(\cdot, \cdot; x, v) : \mathcal{G} \times \mathcal{A} \rightarrow R$ coincides with U_0 .

The assumption says that observing the outcome of the decision problem does not cause the agent to alter his evaluation. It may be possible to replace A.3 by a more primitive restriction on behavior, namely that the agent is not willing to pay anything for an extra piece of data. Later we discuss how to extend the model to accommodate dynamic adjustments of preferences as a result of learning useful correlations in the data.

4.2. Subjective Representation of the Environment

How are we to think of behavior satisfying these assumptions? One such behavior is that of an agent who strictly prefers to decide after knowing the instance x , even though x was not unforeseen (it was part of his description of states), and he understands the way x can help him predict his ex post utility v . We shall interpret this as evidence of the agent believing that there is too much variability in his utility to make it worthwhile to give a detailed contingent account of how it varies with each instance. Instead, he subjectively represents his environment as if his von Neumann-Morgenstern utilities fluctuates randomly. This randomness represents cognitive uncertainty about the outcome of a deterministic, but complex process.

We first give a general definition of random utility:

DEFINITION: A *random utility* \mathcal{U} is a probability measure λ on $(X, \tilde{\mathcal{A}})$, a probability space (Ω, Σ, P) , and random objects $\{\tilde{v}(x) : x \in X\}$ on (Ω, Σ) .

Anticipating our representation result, we focus on random utilities satisfying:

B.1: $\tilde{v}(x) : \Omega \rightarrow V$, i.e., $\tilde{v}(x)$ is a V -valued random vector;

B.2: For each m , $l_m(x) = P\{\tilde{v}(x) = v_m\}$ is a measurable function in x .

B.3: For any finite set of instances $\{x_1, \dots, x_J\} \subset X$, the random vectors $\{\tilde{v}(x_1), \dots, \tilde{v}(x_J)\}$ are independent.

Assumption B.1 says that the agent knows that he will never have state-contingent utilities outside the finite set V . Assumption B.2 roughly says that, although the instance-by-instance description of his utility may be complicated, the *average* performance of an action is well-behaved. B.3 says that knowing the state-contingent utility at an instance x conveys no additional information about his utility at another instance x' . Violations of B.3 would reflect the agent's belief that knowledge gathered at one instance x may be helpful in dealing with new instances never encountered before. See Section 5.2 for discussion.

Given a random utility \mathcal{U} satisfying B.1-3, we *define* the ex ante expected utility it generates by

$$\lambda(A)U_0(g, Q_A) = E \int_A \max_{b \in g(x)} \tilde{v}(b, x, \omega) d\lambda(x). \quad (1)$$

and set $U_1 = U_0$ for all outcomes (x, v) . The ‘max’ captures the idea that $g(x)$ represents the option to pick the action ex post, subject to the constraint that the action picked belongs to the constraint set $g(x)$.¹² The integral is the Pettis integral, a well-known integral first used in economics by Uhlig (1996)—the Appendix provides background and references. For our purpose, this integral reduces to a very simple form:

THEOREM 1. *Under assumptions B.1-3,*

i) *For every C and A , the Pettis integral is a degenerate random variable, and*

$$\lambda(A)U_0(C, Q_A) = \int_A \left[\sum_m l_m(x) \max_{b \in C} v_m(b) \right] d\lambda(x). \quad (2)$$

ii) *For any option $g \in \mathcal{G}$,*

$$\lambda(A)U_0(g, Q_A) = \int_A \left[\sum_m l_m(x) \max_{b \in g(x)} v_m(b) \right] d\lambda(x). \quad (3)$$

Local Problems and Cognitive Discontinuity: One interpretation of this result is that the agent computes his utility by thinking of Q_A as an aggregation of small *local problems*. Formally, define let $l(x) = (l_1(x), \dots, l_M(x))$, let Q_x be the *local problem* in which the agent chooses actions facing a lottery $l(x)$, and let $U_0(C, Q_x)$ be the von Neumann-Morgenstern utility relative to $l(x)$.¹³ Local problems are perfectly standard ‘choice under uncertainty’ decision problems that involve no complexity considerations. Then (2) is equivalent to:

$$\lambda(A)U_0(C, Q_A) = \int_A U_0(C, Q_x) d\lambda(x). \quad (2')$$

¹² In the special case of a constant action b , this reduces to $E \int_A \tilde{v}(b, x, \omega) d\lambda(x)$.

¹³ That is: $l_1(x) \max_{b \in C} v_1(b) + \dots + l_M(x) \max_{b \in C} v_M(b)$.

This says that the agent computes his expected utility in the potentially complex sub-problem Q_A by aggregating his utility at local problems which involve no complexity considerations.

It is essential to note that Q_x is not the instance x , but the ‘average problem’ in vanishingly small neighborhoods around x . Interesting behavior in our model arises when Q_x does not degenerate to a lottery assigning unit mass to the ex post utility at x . In this case, the agent faces a cognitive discontinuity: ex ante he must think in terms of local problems Q_x around an instance x , which are qualitatively different from the ex post problem when he knows which instance x he is facing for sure.

4.3. Representation Theorem

THEOREM 2. *A random utility \mathcal{U} satisfying B.1-3 generates expected utilities U_0 and U_1 satisfying A.1-3.*

We now turn to the converse of the last theorem. Suppose we observe an agent making choices and ranking options in a manner consistent with A.1-3. The next theorem shows that we can think of this behavior ‘as if’ the agent formulated a probabilistic model of his environment, and ranked options based on their expected utility calculated via the Pettis integral:

THEOREM 3. *If U_0 and U_1 satisfy A.1-3, then there exists a random utility \mathcal{U} satisfying B.1-3 such that U_0 is its expected utility (i.e., U_0 satisfies (1) relative to \mathcal{U}).*

Sketch of the proof: While the proof is involved, a basic sketch is instructive and easy to present. Starting with a utility U_0 , imagine an agent who uses trees to successively refine Q_A into simpler sub-problems. Each sub-problem has an equivalent lottery, but these lotteries need not be unique. What is uniquely defined, however, is the vector of utilities, with entry for the utility of each option $C \in \mathcal{C}$. Our reduction of compound sub-problems implies that in dividing a problem into sub-problems, the utility vector of the original problem is the λ -average of the vectors of utilities of the sub-problems. This implies that the vector of utilities is a martingale on X with respect to an

appropriately chosen filtration that generates $\bar{\mathcal{A}}$ in the limit. By the martingale convergence theorem, a limiting ‘local’ utility $U_\infty(x)$ exists. Any such utility has at least one equivalent lottery: in fact a measurable selection $l(x)$ of such lotteries exists, picking one lottery at each x . This selection is used to construct the random utility using standard techniques (in particular, Kolmogorov extension theorem). The independence assumption A.3 restricts our construction by requiring that B.3 must hold.

4.4. Preference for Flexibility

So far we haven’t discussed when the agent ‘thinks’ his environment is complex. As an example, suppose the agent believes that only two utilities are possible, v and v' , and that the utility at x is v if and only if the first feature has value 1. This is a very simple world the agent believes he can sort out based on a single feature. Complexity arises when the agent believes ex post utilities (hence optimal choices) are ‘mixed-up’, so he cannot disentangle sets of instances over which various actions should be chosen.

What are the behavioral implications of complexity? Motivated by Kreps (1992), we focus on preference for flexibility as evidence of the agent’s belief about how complex his environment is.¹⁴ To formalize this, we say that an option g' *refines* g if $g'(x) \subset g(x)$ for every instance x .

DEFINITION: A random utility \mathcal{U} satisfying B.1-3 displays no *preference for flexibility* over Q_A if every $g \in \mathcal{G}$ can be refined by a rule f such that $U_0(f, Q_A) = U_0(g, Q_A)$.

Our goal is to characterize such preference for flexibility in terms of the agent’s random utility. To do this, we need the following:

¹⁴ Our agent, as in Kreps, always weakly prefers more flexibility. Indeed, maximum flexibility (*i.e.*, $g(x) = B$ for all x) is the *unconstrained* optimal plan. In specific applications, however, other considerations generate constraints reflecting benefits to committing to narrower sets of actions (see Section 6). Thus, the more interesting question is whether flexibility is strictly valued by the agent.

DEFINITION: Two utilities $v, v' \in V$ are *ordinally equivalent* if they induce the same ranking over actions. That is, for every b, b' : $v(b) > v(b') \iff v'(b) > v'(b')$. A lottery l has an *ordinally equivalent support* if every v, v' in the support of l are ordinally equivalent.

THEOREM 4. Suppose that \mathcal{U} is a random utility satisfying B.1-3, and such that each $v \in V$ induces a strict ordering on the set of actions. Then \mathcal{U} displays no preference for flexibility if and only if, for almost every instance x , $l(x)$ has ordinally equivalent support.

We now compare our notion of preference for flexibility with that in Kreps (1992) using his example:

Kreps' Luggage: An agent is told he will travel at a time and to a destination picked at random. Let B denote the set of all bundles of luggage he can carry with him. The agent comes up with a list S of objectively describable features of his (yet unknown) time and destination of travel. For instance, S may include the month in which travel occurs, the destination's latitude, altitude, average annual temperature, and so on. The agent is given the possibility of conditioning his bundle of luggage on the list of states he can come up with. Thus, the objects of choice are all functions $f : S \rightarrow B$. Suppose now we also allow him to choose contingent opportunity sets $g : S \rightarrow \mathcal{C}$, where \mathcal{C} is the set of all non-empty subsets of B , with the interpretation that in state s , the agent is free to pick ex post any bundle he likes from the opportunity set $g(s)$. The agent recognizes that the list of contingencies S he thought about may be incomplete, that he may have missed some relevant contingencies $s' \notin S$. Kreps' idea is that the agent's recognition that he did not foresee everything may be inferred from this agent's preference for flexibility.

To draw a close comparison with Kreps' example, consider a given finite set of features I , and let $S = 2^I$. What we called rules and options correspond to Kreps' contingent actions and opportunity sets, respectively.¹⁵ Kreps' agents may value flexibility because of complexity, but they may also value it because they have an intrinsic taste for flexibility, suffer from cognitive limitations, or face exogenous restrictions on the set of contingencies they can condition on. In Kreps' general setting, it

¹⁵ In Dekel *et al.* (1987) these correspond to actions and menus.

may be impossible to separate these effects (Kreps (1992, p. 268)).

On the other hand, one would expect the source of an agent's preference for flexibility to be important when introspection, learning, dynamic choice, and multi-agent considerations are introduced. These considerations raise the issue of how the agent goes about improving his understanding of the world, how he incorporates evidence that contingencies he hadn't thought about before (*i.e.*, $s' \notin S$) may be relevant.

In this paper we focused on environments with considerably more structure (by starting with the agent's von Neumann-Morgenstern utilities, and by including a more detailed description of the states and an explicit description of the multistage decision process). Our treatment is therefore less general than Kreps. On the other hand, the added structure allows us to more narrowly identify why the agent values flexibility: since he can refine his rules by breaking up problems into sub-problems, our framework takes into account the agent's attempts at discovering useful regularities or patterns in the underlying problem that would make flexibility less valuable, or even redundant. Theorem 4 pins down the value of flexibility in terms of the agent's inability to figure out an effective way of sorting out his ex post utilities due to the complexity of the environment.

4.5. Further Remarks on the Interpretation of the Model

- i) *Interpretation of Expected Utility as a Law of Large Numbers:* Equation (2) can be stated as:

$$U_0(C, Q_A) = \sum_m l_m(A) \max_{b \in C} v_m(b). \quad (2'')$$

where $\lambda(A)l_m(A) = \int_A l_m(x)d\lambda$ is the average probability of v_m . This says that the agent behaves by computing the maximum utility for each v_m , then weighs these using the average probability of $l_m(A)$. This may be *interpreted* as a law of large numbers, in the sense that in computing his average utility, the agent uses the average of expectations $l_m(A)$, effectively ‘assuming’ that idiosyncratic variations wash out.¹⁶ Here, we do not *assume* that the Pettis integral is the correct way to aggregate idiosyncratic variations in utilities. Rather, we *prove* that observed behavior satisfying our behavioral assumptions can be thought of as if the agent believed that this integral provided the correct aggregation. Thus, an interesting by-product of Theorems 1 and 3 is a decision theoretic interpretation of the Pettis integral.

- ii) *How Natural is the restriction to \mathcal{A} ?* Our analysis is based on restricting agents to plans for which there is an effective way of finding out what they imply at each possible instance. Formally, we require any rule or option to be measurable with respect to the algebra \mathcal{A} of algorithmic subsets. Is this too strong? If we think of plans as procedures to be implemented or communicated to others, codified in a set of rules or laws, then there must be a finite list of instructions that can reproduce the implications of such plans (*e.g.*, what subset of actions is allowed) at each instance. This is precisely the definition of \mathcal{A} . This is also the sense in which behavior of an agent who examined all improvements possible within \mathcal{G} is rational: no one, including the modeler, is able to suggest effective ways of improving on the agent’s decisions. In economics, game theory, and contracting, the idea that agents are limited to procedures that can be carried out effectively (and so must be algorithmic in this sense) may be found, for example, in McAfee (1984), Binnmore (1987), Gilboa and Schmeidler (1994), and Anderlini and Felli (1994).

¹⁶ The interpretation of the Pettis integral as a law of large numbers was proposed by Uhlig (1996) in his study of economies with a continuum of agents.

- iii) *Introspection:* Uncertainty about states in our representation is subjective: it is the agent's way of thinking about what his utility would be at instances he hadn't considered. By introspection, the agent may decide to take up a particular instance x and figure out what his utility at that instance (assuming no interim uncertainty-see Section 5.1). It is pointless, however, for the agent to try to think through every possible instance, because there are too many of them, each with zero probability of being actually encountered. The agent's probabilistic model of his world may be thought of as reflecting his answer to the question: To what extent does the knowledge gained from introspection about a single instance generalize to other instances? When the problem is simple, for a typical instance x , introspection pins down not just the utility at x , but also the utility at neighboring instances. Otherwise, such introspection gives only a rough idea about what happens around x , because regardless of how close the agent gets to x , there are always details left out that interact in a complex way in determining his utility. (This is the essence of the discontinuity between the local problem Q_x and the corresponding instance x). In the representation, $l(x)$ may be interpreted as containing all that is generalizable from introspection about x .
- iv) *Subjective derivation of utility:* In this paper subjectivity is limited to the new subjective states the agent creates to represent the perceived complexity of his environment. Aside from this, we have taken as given his assessment of the probability distribution λ , his von Neumann-Morgenstern utilities over lotteries, and his utility over options and sub-problems. A worthwhile extension of this model would be to derive these from more primitive assumptions about preferences. In the context of Kreps' framework, Nehring (1998) derives a probability over the foreseen states from assumptions about the agent's behavior, while Dekel *et al.* (1997) derive a unique subjective state space and a representation where additivity is meaningful. The methods developed by these authors may provide a basis for extending our model to derive subjective λ and ex post utilities from the agent's preferences.

5. EXTENSIONS: INTERIM UNCERTAINTY AND CORRELATION

The analysis of last section presented the starkest example of behavior consistent with our model: complete discontinuity between the agent's perspective before and after knowing the instance x . Here we extend our basic model of behavior by allowing for interim uncertainty and correlation.

5.1. *Interim Uncertainty*

Consider the problem of diagnosing a medical condition. We may think of such problem as consisting of: the *ex ante* problem of formulating procedures that practitioners can apply in performing such diagnosis, and the *interim* problem of deciding what to do given a specific instance of the problem.¹⁷ A rigid diagnosis procedure will be deficient because practitioners may want to amend it to take into account the specific nature of each instance. However, even facing an instance, one may still be unsure of the correct diagnosis. This is an example in which uncertainty is resolved in two stages: first, cognitive uncertainty due to complexity is resolved once a specific instance is encountered, but a residual interim uncertainty remains; second, interim uncertainty is resolved ex post (*e.g.*, the patient lives or dies).

Our model so far focused exclusively on the behavior of an agent for whom all uncertainty is resolved once he knows which instance x he is facing. An example of such behavior is an agent unable to give a precise definition of an object (say, a chair) but believes he can correctly categorize such object when he sees it. This is an example of sharp cognitive uncertainty where ex ante utility is random, but ex post utility is deterministic once x is known.

Our model can be modified to allow for a distinction between cognitive uncertainty due to complexity, and residual uncertainty. For simplicity, we will not attempt to start from primitive behavior

¹⁷ An instance in this case is a particular patient with a given set of symptoms, family history, genetic composition, and so on.

assumptions as we did earlier, but work directly from the representation. We also make other simplifying assumptions that can be easily dispensed with.

Let D be the set of all degenerate distributions δ_v , $v \in V$ and let L^0 be a finite set of distributions that includes all degenerate distributions.¹⁸

To allow for interim uncertainty, we amend our definition of random utility by representing it as a collection $\{\tilde{d}(x) : x \in X\}$ of functions on (Ω, Σ) such that:

B.1': $\tilde{d}(x) : \Omega \rightarrow L^0$, *i.e.*, $\tilde{d}(x)$ is a D -valued random vector;

B.2': For each $l \in L^0$, $d_l(x) = P\{\tilde{d}(x) = l\}$ is a measurable function in x .

B.1' allows for the possibility that all the agent knows when he sees an instance x is a distribution $l(x)$ on ex post utilities. Our earlier definition is the special case that precludes interim uncertainty by requiring that $l(x) \in D$, *i.e.*, a degenerate lottery on ex post utilities. Condition B.2' is an obvious modification of B.2.

With these assumptions, we can show (using essentially the same argument as Theorem 1) that the agent evaluates options by :

$$\lambda(A)U_0(g, Q_A) = \int_A \sum_{l \in L^0} d_l(x) \max_{b \in g(x)} E_l v(b) d\lambda.$$

What are the potentially observable implications of the distinction between the two types of uncertainty? To address this imagine that we enlarge the set of options available by introducing new, *ex post options*, that may be exercised only at the ex post stage, after v is known. These differ from the *interim options* we have been using so far, which must be exercised in the interim stage, *i.e.*, after x is known but before v is known. An individual who believes he knows what to do once he sees the instance x attaches no additional value to ex post options above and over the value of the interim

¹⁸ That is, δ_v is the element of L that puts unit mass on the ex post utility v , and $D \subset L^0 \subset L$. The finiteness of L^0 is used here to economize on inessential technical complications. There does not seem to be any special difficulty in extending the analysis by taking $L^0 = L$.

options. On the other hand, an individual who believes there is substantial interim uncertainty (*e.g.*, the doctor in the diagnosis problem) attaches value to the ex post option. Thus, one can in principle detect some of the behavioral implications of the two types of uncertainty by offering different types of options.

A second, related, question we can address with the introduction of these new options is the difference between uncertainty due to complexity and the usual uncertainty where an agent faces a lottery on a set of payoff functions. As we have seen earlier, complexity is what makes interim options valuable. Uncertainty about payoffs, on the other hand, makes ex post options valuable. Thus, an environment with traditional uncertainty about payoffs and no complexity is one where an agent may be willing to pay a considerable premium for an ex post options, but nothing for an interim option.¹⁹

5.2. Correlation and Learning

Our independence assumption B.3 is a useful benchmark describing a steady-state reached after an agent who has accumulated a large number of observations, either directly or by observing other agents dealing with a similar problem. Such agent believes he had extracted all useful patterns or regularities in his environment, so observing a new instance adds nothing, has no effect on how he assess the probability at other instances. Interesting learning possibilities appear if we weaken B.3 to allow for correlation across instances. While a more general treatment is possible, we focus instead on a simple example that illustrates the basic points.

Modify Kreps' luggage example by introducing a travel agent who makes travel plans for a new

¹⁹ To see how this is reflected in the representation, suppose there are just two ex post utilities v and v' . Let \bar{d} be the distribution that puts equal weight on δ_v and $\delta_{v'}$. Consider first an agent with random utility \mathcal{U} where the interim lottery puts unit mass on \bar{d} . That is, regardless of which x the agent encounters, he believes the two utilities obtain with probability $\frac{1}{2}$ each. Knowledge of the instance has no predictive value in this case. Contrast this with another agent with random utility \mathcal{U}' such that the interim lottery $d(x)$ puts probability $\frac{1}{2}$ over each of δ_v and $\delta_{v'}$. This agent believes knowledge of the instance x is a perfect predictor of the outcome (*i.e.*, he knows it when he sees it), but that the environment is too complex for him to think through all possible instances.

client. The travel agent, who evaluates luggage bundles and flexibility to maximize the client's expected utility, continues to face the complexity caused by her inability to think of what is needed in every possible contingency. However, not knowing the new client very well, the problem is compounded by her uncertainty about the client's preferences (his taste, travel habits, medical needs, and so on). More formally, suppose the travel agent believes the client's random utility is generated by either P with probability μ , or P' with probability $\mu' = 1 - \mu$. We assume P and P' satisfy the independence assumption B.3 and display preference for flexibility.

The travel agent's valuation of flexibility at the first stage takes into account both the problem's complexity and uncertainty about the client's type. Although first stage preferences satisfy Kreps' axioms, these preferences cannot discriminate between the two uncertainties, a point that was emphasized by Kreps (1992).

On the other hand, the two uncertainties are conceptually distinct and lead to different patterns of dynamic choice. Imagine, for instance, if the problem is repeated over time, then one would expect the travel agent's evaluation of rules and options to change to reflect learning the type of the client. In particular, with accumulation of information about the client's preferences, preference for flexibility due to uncertainty about whether the client is type P or P' may diminish or disappear, while that due to the complexity of P or P' will not.

In summary, an agent may distinguish between (what he views as) learnable patterns, or the *complexity type*, of his environment, on the one hand, and inherent complexity due to irreducible variability, on the other. This distinction, which may be mute in static settings, has potentially important implications for dynamic choice, *i.e.*, how preference for flexibility changes over time with the arrival of new information.

6. APPLICATIONS TO MULTI-AGENT PROBLEMS

We consider applications to two-player games, where the game played may depend on the instance x . We make standard game-theoretic assumptions that players have common knowledge of the structure of their environment. In our context, these include the space of instances X , the distribution λ , and each other's random utility.

6.1. Flexibility with Time-Inconsistent Preferences

Here we consider the classic Strotz-Pollak problem of time-inconsistent preferences.²⁰ Although this problem involves a single agent, it is formally equivalent to one in which an agent interacts with his future 'self'. Time-inconsistency and preference for flexibility have opposite effects, since time-inconsistency tends to push the agent in the direction of narrowing his future options. When both effects are present, the agent's preference may violate Kreps' monotonicity axiom (which requires that the agent weakly prefers larger subsets of actions).

Consider an agent who, in a first period (time 0), has a utility \mathcal{U} , while in the second period (period 1), his utility may change to \mathcal{U}' . The agent's preferences is time-inconsistent if $\mathcal{U} \neq \mathcal{U}'$.²¹ For example, consider an agent who at time 0 must decide whether to accept a credit card solicitation. He anticipates that flexibility (in the form of having available cards with substantial combined credit limits) is desirable because he cannot think in advance of every possible contingency that might arise. On the other hand, this agent also anticipates that his preferences at time 1 might be such that he does not resist the temptation of excessive spending. Time-inconsistency generates a cost for flexibility, so it is no longer true that the option $g(x) = B$ for all x is always optimal.

²⁰ Strotz (1955-56) and Pollak (1968). See Asheim (1997) for a recent analysis and more extensive references.

²¹ The characterization in Theorem 5 below is valid in general, although its implications are useful primarily when \mathcal{U} and \mathcal{U}' induce ordinally different rankings on a set of positive measure of instances. The reason is that, by Theorem 4, preference for flexibility is ordinal in nature. We state the results without additional restrictions for simplicity of exposition.

Our goal is to characterize the agent's planning at time 0 and show how it depends on his environment. To fix notation, suppose that \mathcal{U} and \mathcal{U}' are defined on the same finite set of utilities V (with generic elements denoted v_m and u_n , respectively).²² Let l and l' denote the distributions on V determined by \mathcal{U} and \mathcal{U}' . We maintain the independence assumption (B.3), so the agent at time zero believes the utilities are independent across instances. However, we now allow for the random utilities to be correlated at each instance. At a given x , let $p(x)$ be the joint distribution on $V \times V$, which we require to be measurable in x and to have marginals $l, l' : X \rightarrow L$. Thus, p measures the extent of the discrepancy between the agent's current and future preferences.

Assume that each utility induces a strict ordering on actions, and define $b(C, u) = \operatorname{argmax}_{b \in C} u(b)$ to be the agent's optimal action at $u \in V$ and $C \in \mathcal{C}$. In a given instance x , and state ω , the agent at time 1 chooses $b = b(C, \tilde{u}(x, \omega))$, so the utility of time 0's agent is the random variable $\tilde{v}(b, x, \omega)$, and his expected utility is given by the Pettis integral:

$$\begin{aligned} U_0(g) &= E \int_X \tilde{v}(b(g(x), \tilde{u}(x, \omega)), x, \omega) d\lambda \\ &= \int_X \sum_{m,n} p(v_m, u_n) v_m(b(g(x), u_n)) d\lambda. \end{aligned}$$

The second equality, which may be proven analogously to Theorem 1, simply asserts that one can evaluate g by computing its value at each local problem Q_x , then aggregate the resulting payoffs.

Our goal is to prove the existence of an optimal plan, and to characterize such plan in terms of the primitives. We call g^+ *globally optimal* if $U_0(g^+) \geq U_0(g)$ for every $g \in \mathcal{G}$. Finding such an optimal plan involves searching through the function space \mathcal{G} , a process which may be both difficult (since \mathcal{G} is infinite dimensional) and not very insightful in terms of clarifying what an optimal rule looks like (*e.g.*, how it varies with the primitives of the environment).

A more useful way to proceed is to consider what kind of plans would solve the problem locally, in the sense of being optimal at any given local problem Q_x . Recall that Q_x is the artificial problem

²² Assuming a common set V is without loss of generality because we can always take this set to be the union of the support the two players' utilities.

obtained as the limit of a sequence of sub-problems Q_{A_k} that decrease to x , which may be interpreted as representing how the problem looks like near x (which, in the interesting case, may be quite different from the problem had x been known). Since each local problem Q_x may be identified with a joint distribution $p(x)$, finding the optimal plan is the straightforward problem of finding C which maximizes $E_{p(x)}v_m(b(C, u_n))$. An option g is *locally optimal* if, for almost every x :

$$g(x) \in \operatorname{argmax}_{C \in \mathcal{C}} E_{p(x)}v_m(b(C, u_n))$$

THEOREM 5.

- i) *A plan is globally optimal if and only if it is locally optimal.*
- ii) *A locally optimal plan exists.*

Part (i) ensures that to verify optimality, it is enough to verify optimality at local problems that involve no complexity considerations. Part (ii), on the other hand, ensures that one can thread together locally computed optima into a well-defined option.

Theorem 5 provides an insight into the agent's planning by reducing his problem to that of examining the trade-off between commitment and flexibility at local problems. The theorem helps identify how this trade-off derives from the complexity of the environment and the expected discrepancy between his current and future utility.

6.2. Delegation Games

A principal decides on how much discretionary control over an action to delegate to an agent. We model this in the form of an option $g \in \mathcal{G}$, with the interpretation that the agent is given discretion to choose $b \in g(x)$ for each x . For example, the agent may be a manager of a firm on behalf of an owner who delegates to him the running of day to day operation of the firm. Another example is a legislature delegating to a judge control over the adjudication of legal cases, subject to the constraints and guidelines reflected in the relevant laws.

Different choices of g may be used to transfer full, partial, or no discretion to the agent. The principal's problem is to determine the optimal level of discretion, reflecting the trade-off between flexibility and precision.

To formally model this trade-off, we assume that the principal and agent have utilities \mathcal{U} and \mathcal{U}' . The difference between these utilities reflects (un-modeled) incentive problems or a different understanding about what should be done in each possible instance. If the principal and agent had the same utility functions and the environment is complex, then it is strictly optimal for the principal to choose $g(x) = B$ for all x . On the other hand, if the two utilities do not coincide, then increasing discretion (*i.e.*, making g less strict and less rule-like) increases the manager's freedom to adapt his action to the particular circumstances he faces, but also increases the chance that he chooses an action not in the interest of the principal.

Theorem 5 applies the present context without modification. Its interpretation in the delegation context is as follows. When the marginal l is degenerate (or has an ordinally equivalent support), the principal is able to be precise about what action he wants the agent to take. Consequently, no discretion is given: a singleton $g(x)$ dictating the principal's optimal action is weakly optimal. The more interesting case is that where the support of l contains ordinally non-equivalent utilities. We may interpret this as indicating the principal's uncertainty about the right action due to the complexity of his environment. In this case, partial delegation will typically result, to balance the incentive to capture some of the gains from flexibility against the probability that the agent takes a suboptimal action.

As we see from these applications, Theorem 5 is an instance of a more generally applicable optimality principle for solving optimization problems in our setup. The basic idea is that our setup allows to solve the seemingly difficult (infinite dimensional) global problem by threading together (finite dimensional) local problems.

6.3. Simon's Model of Authority in Employment Relationships

The delegation model above may be adapted to provide a formulation of Simon's (1951) seminal model on authority in employment relationships. Briefly, Simon considers that the distinguishing feature of employment contracts is the employee's acceptance, in exchange for a wage, of his employer's subsequent discretionary authority to tell him what to do. This creates a trade-off between flexibility and abuse of authority: since the task the agent is asked to perform depends on an unknown state, it is efficient to give the employer some flexibility in directing the agent as contingencies unfold. On the other hand, the employer may be tempted to abuse his authority (*e.g.*, by directing the employee to taking onerous tasks). An employment contract gives the employer a limited discretion so as to balance these forces.

One issue left unanswered by Simon is: what prevents the parties from agreeing to a detailed enough contract that eliminates the need for flexibility (hence authority and its abuse). One answer may be that the complexity of the environment makes it difficult to adapt the contract's provisions to every possible contingency. To model this, we modify the delegation example by imagining that a contract g is proposed which the agent either accepts or rejects. If accepted, then g endows the employer with the authority to direct the employee to take action b conditional on the realization of x , provided that action falls within the employer's authority, in the sense of $b \in g(x)$.

An employment contract is complete if it belongs to \mathcal{F} , *i.e.*, a rule that prescribes a single action at every possible instance. Typically, this is not optimal because the employer's discretion may be valuable in a complex environment. On the other hand, a completely vague contract (*e.g.*, $g(x) = B$) may not be optimal either because it leaves the employer greater flexibility to abuse his authority over the employee. As in Simon, the optimal contract g^* is drawn so as to balance these effects. The characterization in Theorem 5 above carries through without changes. In particular, we can solve for g^* by looking at local solutions for the local problems Q_x .

Testing for complexity: The predictive value of this model obviously depends on whether we can say something about the form of the optimal contract g^* (how it should look like and the form of its

incompleteness) as a function of objective, potentially observable aspects related to the complexity of the environment. Contracting parties often have access to massive data (generated both internally within an organization, and through observing similar interactions elsewhere) that can be used to evaluate the relative complexity of different work environments, tasks, and job assignments. For example, it may be clear from past experience that the problem of describing the tasks of a company's CEO is more complex than that of a security guard or a assembly line worker in one of this company's plants. In our model, complexity of a task is (at least in principle) measurable in terms of the variability of the action required as a function of a given objective description of instances. Testing for complexity in this case boils down to tests of randomness, *i.e.*, that there is no remaining structure or pattern one can uncover in the data. Thus, our formulation of Simon's model provides a criterion to distinguish between routine aspects that may be hard-wired into a formal contract, from novel, non-routine aspects that warrant discretion.

6.4. Nash equilibrium in a two stage game

Here we consider a simple setting in which we define and prove the existence of a Nash equilibrium in a two player game. The purpose is to illustrate how our framework may be used to formulate and analyze simultaneous move games (define strategies, equilibria, and prove existence).

To economize on notation, we keep the setting as close as possible to that of Section 5.1. For notational simplicity, we assume there are two players $i, j = 1, 2$. The set of instances X is as before. At each x the two players have payoff functions v_m^i and v_n^j respectively. The payoffs are generated according to the random utilities \mathcal{U}^i and \mathcal{U}^j , with joint distribution $p(x)$ on the players' payoff functions. Thus, at each instance x , a simultaneous move game with payoff vector (v_m^1, v_n^2) is realized with probability $p(v_m^1, v_n^2; x)$.²³

In this setup, x represents an instance of the environment facing the two players; the situation

²³ As before, we assume that p is independent across the instances, although at a given x the utilities may be correlated.

they find themselves when they actually get to choose their actions. The dependence of payoffs on x reflects the possibility that the environment may influence the interaction between players in potentially complex ways. To make the model tractable, we assume that the players agree on the probability λ of various contingencies and have common knowledge of each other's random utility.

We consider Nash equilibrium of the normal form of this game, with the restriction that contingent strategies must be algorithmic. Aside from this restriction, everything else is completely standard. To prove existence, we need to expand our definition of rules to include randomizations. Formally, a *mixed rule* for player i , $i = 1, 2$, is a function $\phi^i : X \rightarrow \Delta^{B^i}$ that is $\bar{\mathcal{A}}$ -measurable. That is, a mixed rule assigns probability distributions on actions, rather than pure actions as in the case of regular (pure) rules. Given a profile of mixed rules, (ϕ^1, ϕ^2) , we define the expected payoff of player i as

$$U_0^i(\phi^1, \phi^2) = \int_X \sum_{m,n} p(c_m^i, c_n^j; x) c_m^i(\phi^1(x), \phi^2(x)) d\lambda$$

DEFINITION: A profile $(\bar{\phi}^1, \bar{\phi}^2)$ is an equilibrium (in rules) if for every i, j , $\bar{\phi}^i$ is a best response to $\bar{\phi}^j$.

THEOREM 6. *An equilibrium exists.*

The idea of the proof is similar to that of Theorem 5 above. Specifically, we imagine the game as consisting of playing local games Q_x , each of which is a normal form game with random payoffs determined by the joint distribution $p(x)$. Since each such game is finite, it has at least one equilibrium. The proof shows that it is possible to select a mixed rule from the equilibrium correspondence.

Finally, call a player's strategy ϕ^i *sequentially consistent* if, with probability one on x , the action $\phi^i(x)$ is optimal when the instance is known to be x .²⁴ We may use the analysis of flexibility (Section 4.4) to characterize when an equilibrium in rules is sequentially consistent. By Theorem 4, in a complex problem, players are likely to revise their strategies ex post as contingencies unfold (*i.e.*, they 'cross a bridge when they get to it'). These agents will not work out a fully detailed contingent plan of action that anticipates every contingency, but rather deal with contingencies as they arise.

²⁴ In this discussion we assume that $\phi^i(x)$ is pure for simplicity.

7. SUMMARY AND CONCLUSIONS

While its importance and pervasiveness is beyond dispute, complexity remains an illusive concept to formally model. Rather than prescribing computational or thought procedures for agents to follow in coping with their complex environment, this paper attempts to elicit their assessment of the complexity or simplicity of the environment, the world as they see it. We explain an agent's behavior as the result of a coherent model representing the agent's attempt to find order and regularities in his world, while at the same recognizing that his model cannot explain everything.

Three features of our approach are worth emphasizing. First, we show the possibility that fully rational agents confronting complex situations may display behavior often associated with bounded rationality. Second, despite the complexity of the environment, the agent's model of it may be remarkable simple. This, in turn, means that we can develop tractable models of behavior, as illustrated by Theorems 5 and 6 where complex multi-agent settings may be analyzed using essentially standard tools and concepts.

The third, and perhaps most important feature of our model is its *closure*: as in traditional economic and strategic models, agents optimize given an understanding of their environment which is as good as that of the modeler. There are no arbitrage opportunities or money-pumps through which an agent may be systematically exploited. This requirement, well entrenched in many areas of economics, imposes considerable discipline on the model's predictions.

There obviously remain important questions for future work. One is to explore the implication of the model in terms of generating similarity- and analogy-driven behavior, along the lines described in Gilboa and Schmeidler (1995). Another pressing issue is incorporating learning considerations: our focus has been on a steady-state where the agent believes he learned all there is to be learned from the environment. How the agent gets to use data to formulate, refine and extend his model is an obvious next step.

APPENDIX

A.1. Topological and Measurable Properties of $\bar{\mathcal{A}}$

Our analysis of \mathcal{A} and $\bar{\mathcal{A}}$ builds on a result due to Gilboa and Schmeidler (1994, Proposition 3.1, p. 377). Call a set $A \subset X$ *finitely defined* if there is a finite set of features I sufficient for determining membership in A . That is, the characteristic function of A , χ_A , has the property that $\chi_A(x) = \chi_A(x')$ for every x and x' that agree on I (i.e., for every feature $i \in I$, $x_i = x'_i$). For example, the set $\{x : x_1 = 0\}$ is finitely defined, while the singleton set $\{x\}$ is not. Clearly, every finitely defined set is algorithmic. What is more surprising is the converse:

PROPOSITION A.1: *(Gilboa-Schmeidler (1994)) A subset $A \subset X$ belongs to \mathcal{A} if and only if it is finitely defined.*

The reason why this is not entirely obvious is that algorithms allow for unbounded computation: for instance, the machine could scan a subset of features and, depending on some internally generated output, may decide to examine more features, and so on. Nevertheless, the proposition says that an algorithm will never look beyond a predetermined finite set of features. Critical for this result is the assumption that the algorithm always halts. See Gilboa and Schmeidler (1994) for further elaboration on this and related points.

Aside from its intrinsic interest, Proposition A.1 makes it quite easy to derive the topological and measurable structures of \mathcal{A} and $\bar{\mathcal{A}}$:

PROPOSITION A.2:

- i) \mathcal{A} is a base for the product topology, τ , on X ;
- ii) τ is a complete, separable metrizable topology;
- iii) $\bar{\mathcal{A}}$ coincides with the Borel σ -algebra generated by τ .

Proof: (i) Let τ' denote the topology generated by taking \mathcal{A} as a base (that is the collection of sets obtained by taking arbitrary unions of sets in \mathcal{A}). We need to show that $\tau = \tau'$. Call a set *simple* if it is of the form $\{x : x_i = \alpha\}$, for some feature i , and $\alpha \in X_i$. Note that simple sets are inverse images of projections, and the product topology is the coarsest topology that makes all projections continuous. Thus, τ -open sets are those sets which can be obtained as arbitrary unions of finite intersections of simple sets. Obviously, every simple set is in \mathcal{A} , and so is any finite union of simple sets. Consequently, $\tau \subset \tau'$. In the other direction, it is enough to show that $\mathcal{A} \subset \tau$. By Proposition A.1, any $A \in \mathcal{A}$ is finitely defined in terms of a finite set of features I , say. Then A can be generated by taking finite intersections and unions of simple sets corresponding to features in I . Since every simple set is in τ , A must also be in τ .

(ii) τ is the product topology of a countable collection of compact spaces.²⁵ Thus, τ itself is compact (Royden (1968), Theorem 19, p. 166), metrizable (p. 151), and complete and separable (Propositions 13-15, pp. 163-4).

(iii) Since \mathcal{A} is countable, the product topology involves only countable unions of sets in \mathcal{A} . Thus, the σ -algebra generated by \mathcal{A} , namely $\bar{\mathcal{A}}$, coincides with that generated by the product topology. But the later is just the Borel σ -algebra.

Q.E.D.

A.2. The Pettis Integral: Proof of Theorem 1

Here we provide the definition of the Pettis integral as it applies to our setup, and prove Theorem 1. The interested reader may consult Diestel and Uhl (1977) for more detailed account.

The Pettis integral is a method to integrate random variables (like our random utilities $\tilde{r}(x) : \Omega \rightarrow V$ or functions defined on them) by viewing them as elements of a linear space. Here we pursue the

²⁵ In fact discrete spaces, as each X_i is assumed to be binary (*i.e.*, each feature can take only one of two possible values).

most straightforward way of doing this. Let L^2 to be the linear (Hilbert) space of all random variables on (Ω, Σ, P) with finite mean and variance.²⁶ The inner product of two points $f, f' \in L^2$ is defined in the usual way: $(f | f') = \int_{\Omega} f(\omega) f'(\omega) dP = \text{cov}(f, f') + EfEf'$. We let $\mathbf{1}$ denote the (equivalence class of) random variables that take the value 1 with probability 1, and note that $(\mathbf{1} | f) = Ef$.

Consider a function $x \mapsto f(x)$, which maps X into L^2 (specifically, we later consider the function $f(x) = \max_{b \in C} \tilde{v}(b, x, \omega)$). The Pettis integral is a way to integrate such mapping by averaging the random variables $f(x)$ as points in L^2 . While this implies that the Pettis integral itself is a point in L^2 , we will see that under our assumptions, the Pettis integral is a degenerate random variable (*i.e.*, constant almost everywhere), so there is an obvious way to identify it with a real number.

We can now provide a formal definition: $\int_A f(x) d\lambda$ is the Pettis integral of $x \mapsto f(x)$ over $A \in \mathcal{A}$ if for any $z \in L^2$,

$$\left(z \mid \int_A f(x) d\lambda \right) = \int_A (z | f(x)) d\lambda, \quad (4)$$

where the integral on the RHS is the ordinary Lebesgue integral.

General results on the existence and characterization of the integral $\int_A f(x) d\lambda$ are available. However, these results are unnecessary for our purpose since our special structure (especially assumption B.3) enables us to display and characterize the integral directly.²⁷

Proof of Theorem 1: (i) Fix A and C . To economize on notation, define $f(x) = \max_{b \in C} \tilde{v}(b, x, \omega)$ and let $\bar{f}(x) = Ef(x)$. Note that any finite collection of $f(x)$'s is independent, being functions of independent random variables.

²⁶ More precisely, L^2 consists of equivalence classes of functions because the L^2 norm cannot distinguish between two random variables that differ only on a set of measure zero. Here, we will abuse notation and use the same symbols to denote the random variable and its equivalence class, as the difference plays no role in what follows.

²⁷ This proof derives the value of integral from equation (4) using only elementary properties of means and covariances. A much shorter proof is possible, although the argument would depend on Hilbert space techniques which may be unfamiliar and less transparent.

If the Pettis integral of $f(x)$, $\int_A f(x)d\lambda \in L^2$, existed then from (4), this integral would have to satisfy, for every $z \in L^2$, the equation:

$$\text{cov} \left(z, \int_A f d\lambda \right) + Ez \ E \int_A f d\lambda = \int_A \text{cov}(z, f(x)) d\lambda + Ez \int_A \bar{f}(x) d\lambda.$$

In particular, a necessary condition for the Pettis integrability of $f(x)$ is that $\bar{f}(x)$ is measurable in x , a condition satisfied in our case since, from assumption B.2, $\bar{f}(x) = E \max_{b \in C} \tilde{v}(b, x, \omega) = \sum_m l_m(x) \max_{b \in C} v_m(b)$, which is measurable in x .

Taking $z = \mathbf{1}$, we further conclude that

$$E \int_A f d\lambda = \int_A \bar{f}(x) d\lambda \tag{5}$$

which is the equation asserted in part (i) of the theorem. Next, taking z to be any random variable with $Ez = 0$ implies

$$\text{cov} \left(z, \int_A f d\lambda \right) = \int_A \text{cov}(z, f(x)) d\lambda.$$

However, since the random variables $f(x) - \bar{f}(x)$ are independent, they must be orthogonal as points in L^2 , so $\text{cov}(z | f(x)) = 0$ except for at most countably many x 's. This implies

$$\text{cov} \left(z, \int_A f d\lambda \right) = 0 \tag{6}$$

for every mean-zero z . But then $\int_A f d\lambda$ must be degenerate.

In summary, if the Pettis integral existed, then it must be a degenerate random variable with mean $\int_A \bar{f}(x) d\lambda$ (*i.e.*, it must satisfy equations 5 and 6). Conversely, the point in L^2 corresponding to any degenerate random variable with mean $\int_A \bar{f}(x) d\lambda$ satisfies (5) and (6), and would indeed be the Pettis integral.²⁸ Such random variable exists by our earlier observation that the (Lebesgue) integral $\int_A \bar{f}(x) d\lambda$ exists.

²⁸ There are many degenerate random variables with mean $\int_A \bar{f}(x) d\lambda$, but they must all agree off sets of measure zero in Ω , so they correspond to a single point in L^2 .

Part (ii) follows from the additivity of the integral and the fact that g must be measurable: Fix $A \in \mathcal{A}$ and define $A_C = \{x : g(x) = C\} \cap A$, then

$$\begin{aligned} \lambda(A)U_0(g, Q_A) &= \sum_{C \in \mathcal{C}} E \int_{A_C} \max_{b \in g(x)} \tilde{v}(b, x, \omega) d\lambda \\ &= \sum_{C \in \mathcal{C}} \int_{A_C} \sum_m l_m(x) \max_{b \in g(x)} v_m(b) d\lambda \\ &= \int_A \sum_m l_m(x) \max_{b \in g(x)} v_m(b) d\lambda, \end{aligned} \tag{*}$$

where (*) follows by noting that $g(x) = C$ for any $x \in A_C$ and applying part (i) of the theorem.

Q.E.D.

A.3. Proof of Theorems 2 and 3

Throughout the proof, we use the notation $l(A) = \frac{1}{\lambda(A)} \int_A l(x) d\lambda \in L$.

Proof of Theorem 2: A.3 is obvious from the construction. A.1 follows by using, for every C and A , the lottery $l(A)$. To see this, note that Theorem 1 implies

$$\lambda(A)U_0(C, Q_A) = \int_A \left[\sum_m l_m(x) \max_{b \in C} v_m(b) \right] d\lambda.$$

But this is just $\lambda(A) \sum_m [\int_A l_m(x) d\lambda] \max_{b \in C} v_m(b)$. It only remains to show A.2, which follows by the additivity of the integral (using Theorem 1):

$$\begin{aligned} \lambda(A)U_0(g, Q_A) &= \int_A \left[\sum_m l_m(x) \max_{b \in g(x)} v_m(b) \right] d\lambda \\ &= \sum_{i=1}^N \int_{A_i} \left[\sum_m l_m(x) \max_{b \in g(x)} v_m(b) \right] d\lambda \\ &= \sum_{i=1}^N \lambda(A_i)U_0(g, Q_{A_i}). \end{aligned}$$

Q.E.D.

Turning to the proof of Theorem 3, we start with a preliminary result, which uses a martingale convergence argument to show that our assumptions (especially the reduction of compound sub-problems, A.2) imply the existence of limiting lotteries. First, we introduce the notation $U(C, l) = l_1 \max_{b \in C} v_1(b) + \dots + l_M \max_{b \in C} v_M(b)$.

PROPOSITION A.3: *Under assumptions A.1-3, there is a $\bar{\mathcal{A}}$ -measurable function $l : X \rightarrow L$ such that for every $A \in \bar{\mathcal{A}}$ with $\lambda(A) > 0$, and $C \in \mathcal{C}$, $U_0(C, Q_A) = U(C, l(A))$.*

Proof: Let \mathcal{A}_n be the (finite) partition of X determined by conditioning on the first n features, and let $\bar{\mathcal{A}}_n$ be the σ -algebra it generates. By Proposition A.2, $\cup_{n=1}^{\infty} \bar{\mathcal{A}}_n = \bar{\mathcal{A}}$. Let $\{C_1, C_2, \dots, C_{2^K-1}\}$ be an enumeration of \mathcal{C} . To each $v \in V$, define

$$\hat{v} = \left(\max_{b \in C_1} v(b), \dots, \max_{b \in C_{2^K-1}} v(b) \right).$$

Let $\hat{V} = \{\hat{v} : v \in V\}$, and note that \hat{V} and $co\hat{V}$ are subsets of \mathbb{R}^{2^K-1} . Define $U_n : X \rightarrow co\hat{V}$ by

$$U_n(x) = \begin{pmatrix} \sum_{i=1}^{2^n-1} U_0(C_i, Q_{A_i}) \chi_{A_i}(x) \\ \vdots \\ \sum_{i=1}^{2^n-1} U_0(C_{2^K-1}, Q_{A_i}) \chi_{A_i}(x) \end{pmatrix}.$$

where χ_{A_i} is the characteristic function of A_i . Note that U_n is \mathcal{A}_n -measurable.

We now show that U_n is a martingale with respect to the filtration $\{\bar{\mathcal{A}}_n\}$. That $E[U_n] < \infty$ is obvious. We need only show that $E[U_{n+1} | \bar{\mathcal{A}}_n] = U_n$, λ -a.s. Let A be an element of the partition \mathcal{A}_n , and let $\{B_1, B_2\}$ be the partition of A in \mathcal{A}_{n+1} (recall our assumption that each feature can take only two possible values). Then,

$$\begin{aligned} E[U_{n+1}(x) | x \in A] &= \lambda(B_1 | A) \begin{pmatrix} U_0(C_1, Q_{B_1}) \\ \vdots \\ U_0(C_{2^K-1}, Q_{B_1}) \end{pmatrix} + \lambda(B_2 | A) \begin{pmatrix} U_0(C_1, Q_{B_2}) \\ \vdots \\ U_0(C_{2^K-1}, Q_{B_2}) \end{pmatrix} \\ &= \begin{pmatrix} U_0(C_1, Q_A) \\ \vdots \\ U_0(C_{2^K-1}, Q_A) \end{pmatrix} = U_n(x) \end{aligned}$$

(The last equality follows from the reduction of compound sub-problems.) Thus, $(U_n, \bar{\mathcal{A}}_n)$ forms a martingale. By the martingale convergence theorem, there is a $\bar{\mathcal{A}}$ -measurable function, $U_\infty : X \rightarrow coV$ such that $\lim_{n \rightarrow \infty} U_n(x) = U_\infty(x)$ for λ -a.e. x (Shiryayev (1984), p. 476).

Let $\Gamma : X \rightarrow L$ be given by $\Gamma(x) = \{l \in L : U_\infty(C, x) = U(C, l), \forall C\}$, where $U_\infty(C, x)$ denotes the j th entry in the vector $U_\infty(x)$. This correspondence is non-empty valued because $U_\infty \in coV$. In fact, there is a measurable selection l , as shown by the lemma below.

To complete the proof, we show that for any measurable selection l , $A \in \bar{\mathcal{A}}$, and $C \in \mathcal{C}$, we have $U_0(C, Q_A) = U(C, l(A))$:

$$\begin{aligned}
\lambda(A)U(C, l(A)) &= \sum_{m=1}^M \int_A l_m(x) \max_{C'} v_m d\lambda \\
&= \int_A \sum_m l_m(x) \max_{C'} v_m d\lambda \\
&= \int_A U_\infty(C, x) d\lambda \\
&= \int_A \lim_{n \rightarrow \infty} U_n(C, x) d\lambda \\
&= \lim_{n \rightarrow \infty} \int_A U_n(C, x) d\lambda \tag{a} \\
&= \lim_{n \rightarrow \infty} \int_A \sum_{i=1}^{2^n-1} U_0(C, Q_{A_i}) \chi_{A_i}(x) d\lambda \\
&= \lim_{n \rightarrow \infty} \sum_{i=1}^{2^n-1} U_0(C, Q_{A_i}) \int_A \chi_{A_i}(x) d\lambda \\
&= \lim_{n \rightarrow \infty} \sum_{i=1}^{2^n-1} U_0(C, Q_{A_i}) \lambda(A \cap A_i) \\
&= \lambda(A)U_0(C, Q_A).
\end{aligned}$$

Here, (a) follows from the dominated convergence theorem and the fact that U_n converges to U_∞ for almost every x .

Q.E.D.

Lemma A.1: *The correspondence Γ admits a measurable selection.*

Proof: From Klein and Thompson (1984, Theorem 14.2.1, p. 163), it suffices to show that Γ is a measurable correspondence, *i.e.*, for every closed subset $F \subset L$, $\{x : \Gamma(x) \cap F \neq \emptyset\} \in \bar{\mathcal{A}}$ (for the definition, see Klein and Thompson, p. 153). Let $W = \{y : y = \sum_{i=1}^{2^K-1} l_j v_j, l \in F\}$ and note that this is a measurable subset of $co\hat{V}$. Since U_∞ is $\bar{\mathcal{A}}$ -measurable, we have $\{x : U_\infty(x) \in W\} \in \bar{\mathcal{A}}$. But $\{x : U_\infty(x) \in W\} = \{x : \exists l \in F \text{ such that } U_\infty(x) = \sum_{i=1}^{2^K-1} l_j v_j\} = \{x : \Gamma(x) \cap F \neq \emptyset\}$.

Q.E.D.

Proof of Theorem 3: We construct the state space with standard argument using the Kolmogorov Extension Theorem (Shiryayev (1984)). The set of states Ω will be the set of all functions $\omega : X \rightarrow V$. Viewing Ω as a product space, projection on the x coordinate, denoted $\pi_x : \Omega \rightarrow V$, is $\pi_x(\omega) = \omega(x)$. We define Σ to be the σ -algebra generated by projections; that is, Σ is the smallest σ -algebra containing all sets of the form $\{\omega : \pi_x^{-1}(v)\}$, for $x \in X$ and $v \in V$.

For each x , define $\tilde{v}_x : \Omega \rightarrow V$ by $\tilde{v}_x(\omega) = \omega(x)$ (*i.e.*, the projection of the state ω on x). Clearly, \tilde{v}_x is measurable with respect to Σ .

For every finite set of instances $\{x_1, \dots, x_S\}$, define the probability distribution $P_{\{x_1, \dots, x_S\}}$ on the finite set V^S so that the random variables $\{\tilde{v}_{x_1}, \dots, \tilde{v}_{x_S}\}$ are independently distributed with $E\tilde{v}_{x_s} = l(x_s)$. Clearly, the $P_{\{x_1, \dots, x_S\}}$'s satisfy the consistency condition in Kolmogorov's Extension Theorem (Shiryayev (1984), Theorem 4, p. 165), so there is a unique probability measure P that agrees with every finite dimensional distribution $P_{\{x_1, \dots, x_S\}}$.²⁹

We now have a random utility \mathcal{U} on a probability space (Ω, Σ, P) . By construction, \mathcal{U} satisfies B.1-3. It is also immediate that A.3 is satisfied. It only remains to show that expected utility maximization

²⁹ The consistency condition here says that for any two subsets of instances $\{y_1, \dots, y_T\} \subset \{x_1, \dots, x_S\}$, and any $(v_1, \dots, v_T) \in V^T$, $P_{\{y_1, \dots, y_T\}}(v_1, \dots, v_T) = P_{\{x_1, \dots, x_S\}}(v_1, \dots, v_T)$. This is true by our construction of the P 's as independent distributions based on the same expectation function $l(x)$.

relative to \mathcal{U} induces the utility U_0 . From Proposition A.3, we have that for every $A \in \bar{\mathcal{A}}$ and $C \in \mathcal{C}$,

$$U_0(C, Q_A) = \sum_m l(A) \max_C v_m \quad (7)$$

where l is the measurable selection derived in Proposition A.3, and used in the construction of P above. But by Theorem 1, expected utility maximization of U generates the same values for the average lotteries $l(A)$ in (7). This means that expected utility maximization induces the same preference on constant options (C, Q_A) . Assumption A.2 implies that this pins down the preference over general options (g, Q_A) .

Q.E.D.

A.4. Proof of Theorem 4

Suppose that $l(x)$ has ordinally equivalent support for almost every x (we drop all ‘almost every’ qualifications below as they play no role in the argument). Fix an option g and define $f(x) = \operatorname{argmax}_{b \in g(x)} U_0(g(x), Q_x)$. Obviously, f refines g . Furthermore, the fact that $l(x)$ has ordinally equivalent support implies that f is single-valued. Next we show that $f \in \mathcal{F}$ (i.e., f is measurable with respect to $\bar{\mathcal{A}}$). For $C \in \mathcal{C}$ and $V' \subset V$, define the set $A(C, V') = \{x : \operatorname{supp} l(x) \subset V' \text{ and } g(x) = C\}$; this set clearly belongs to $\bar{\mathcal{A}}$. Partition V into equivalence classes of ordinally equivalent utilities $V_s \in \{V_1, \dots, V_S\}$. By assumption, for each x , $\operatorname{supp} l(x) \subset V_s$ for one such equivalence class. The assumption of ordinally equivalent support implies that sets of the form $A(C, V_s)$, for $C \in \mathcal{C}$ and $V_s \in \{V_1, \dots, V_S\}$, partition \bar{X} , and that f is constant on each $A(C, V_s)$. Hence $f \in \mathcal{F}$.

To complete the proof, we show that $U_0(f, Q_A) = U_0(g, Q_A)$. Since l puts positive weight only on utilities with the same ranking of actions, for any C and l that is ordinally equivalent, $\sum_m l_m(x) \max_{b \in C} v_m(b) = \max_{b \in C} \sum_m l_m(x) v_m(b)$. In terms of local problems Q_x , this equality is just $U_0(g(x), Q_x) = U_0(f(x), Q_x)$. Theorem 1 (ii) implies that $U_0(g(x), Q_A) = U_0(f(x), Q_A)$.

In the other direction, suppose that \mathcal{U} displays no preference for flexibility. By way of contradiction, suppose, that there is A with positive measure such that the support of $l(x)$ is not ordinally equivalent

for every $x \in A$. Then for every such x , there are utilities v, v' in the support of $l(x)$ and actions b, b' (all depending on x) such that $v(b) > v(b')$, $v'(b) < v'(b')$. Since the sets of actions and utilities are finite, there is $A' \subset A$ with $A' \in \mathcal{A}$ and $\lambda(A') > 0$, such that this holds for fixed v, v', b, b' . Since \mathcal{U} displays no preference for flexibility, there is $f \in \mathcal{F}$ such that $U_0(f, Q) = U_0(g, Q)$. Define the sets $B_1 = A'^c$, $B_2 = A \cap \{f = b\}$, $B_3 = A \cap \{f = b'\}$. Clearly, $\{B_1, B_2, B_3\}$ is a partition of X by sets in \mathcal{A} . However, $U_0(f, Q_{B_i}) \leq U_0(g, Q_{B_i})$, $i = 1, 2, 3$, and with strict inequality for $i = 2, 3$. Contradiction.

A.5. Proof of Theorems 5 and 6

Proof of Theorem 5: To prove part (i), suppose that g^+ is a locally optimal option, and let g be any other option. From Theorem 1 (ii), we have $U_0(g) = \int_X U_0(g, Q_x) d\lambda$ and $U_0(g^+) = \int_X U_0(g^+, Q_x) d\lambda$, where Q_x is the local problem with joint distribution $p(x)$. But $U_0(g^+, Q_x) \geq U_0(g, Q_x)$ for every x , since g^+ is locally optimal. Thus, $U_0(g^+) - U_0(g) = \int_X U_0(g^+, Q_x) - U_0(g, Q_x) d\lambda \geq 0$.

To prove part (ii), define first the correspondence

$$G(x) = \operatorname{argmax}_{C \in \mathcal{C}} \sum_{m,n} p(v_m, u_n)(x) v_m(b(C, u_n))$$

of locally optimal options at a local problem Q_x . Our problem is to show that we can select an option g , such that $g(x) \in G(x)$ for all x . To prove this, we show that G is a measurable correspondence, in which case a measurable selection exists by Klein and Thompson (1984, p. 163). That is, we must show that for every closed $F \subset \mathcal{C}$,

$$\{x : G(x) \cap F \neq \emptyset\} \in \bar{\mathcal{A}}.$$

Since \mathcal{C} is finite, we may restrict attention to singleton sets $F = \{C\}$ (since $\bar{\mathcal{A}}$ is preserved under finite unions and intersections). Thus, for $C \in \mathcal{C}$, we have

$$\begin{aligned} \{x : G(x) \cap C \neq \emptyset\} &= \{x : C \in G(x)\} \\ &= \{x : E_{p(x)} v_m(b(C, u_n)) \geq E_{p(x)} v_m(b(C', u_n)), \forall C' \in \mathcal{C}\} \end{aligned}$$

$$\begin{aligned}
&= \cap_{C'' \in \mathcal{C}} \{x : E_{p(x)} v_m(b(C, u_n)) \geq E_{p(x)} v_m(b(C'', u_n))\} \\
&= \cap_{C'' \in \mathcal{C}} \{x : E_{p(x)} [v_m(b(C, u_n)) - v_m(b(C'', u_n))] \geq 0\}. \tag{**}
\end{aligned}$$

The expression $E_{p(x)} [v_m(b(C, u_n)) - v_m(b(C'', u_n))]$ depends only on the probability weights $p(v_m, u_n)(x)$. Since the joint distribution $p(x)$ is measurable with respect to $\bar{\mathcal{A}}$, every set in the finite intersection in (**) is in $\bar{\mathcal{A}}$. Thus, G is indeed a measurable correspondence.

Q.E.D.

The idea of the proof of the next theorem is analogous to that of Theorem 5. We first identify equilibrium behavior at artificial local games, where at each x players believe they are facing a lottery over payoff vectors given by $p(x)$. Complexity considerations in such games do not arise. We use a measurable selection argument to show that we can select an equilibrium at each instance, measurably relative to $(X, \bar{\mathcal{A}})$. It is then immediate that the resulting profile of locally optimal strategies constitutes a globally optimal rule.

Proof of Theorem 6: Let \bar{L} be the set of probability distributions on $V^1 \times V^2$. l^1 and l^2 are the marginals of p on V^1 and V^2 respectively. Define $\mathcal{E} : \bar{L} \rightarrow \Delta^{B^1} \times \Delta^{B^2}$ to be the correspondence where $\mathcal{E}(\bar{l})$ is the set of Nash equilibrium profiles of the local game Q_x , *i.e.*, the game in which payoffs are randomly distributed according to \bar{l} . Then \mathcal{E} is a compact valued and upper hemicontinuous correspondence. By Klein and Thompson (1984, Example 13.3.5 p. 159 and Theorem 14.2.1, p. 163), there exists a measurable selection from \mathcal{E} . If e one such selection, then since $p : X \rightarrow \bar{L}$ is measurable by assumption, $e \circ p$ is a function on X , measurable with respect to $\bar{\mathcal{A}}$, and such that $e \circ p(x)$ is an equilibrium of the local game Q_x . Note that no local deviation is profitable, since $e \circ p(x)$ is an equilibrium for each x . Integrating over X , no global deviation is optimal. Thus, the pair of rules implied by $e \circ p$ constitute an equilibrium.

Q.E.D.

REFERENCES

- Anderlini, L. and L. Felli (1994): "Incomplete Written Contracts: Undescribable States of Nature," *Quarterly Journal of Economics*, **109**, 1085-1124.
- Asheim, G. (1997): "Individual and Collective Time-Consistency," *Review of Economic Studies*, **64**, 427-443.
- Binmore, K. (1987): "Modeling Rational Players. Part I," *Economics and Philosophy*, **3**, 179-214.
- Cover, T.M. and J.A. Thomas (1991): *Elements of Information Theory*. New York: John Wiley & Sons.
- Dekel, E., B. Lipman, and A. Rustichini (1998): "Recent Developments in Modeling Unforeseen Contingencies". Forthcoming in: *European Economic Review (Proceedings of European Economic Association Meetings)*.
- Dekel, E., B. Lipman, and A. Rustichini (1997): "A Unique Subjective State Space for Unforeseen Contingencies," Northwestern University.
- Devroye, L., L. Györfi, and G. Lugosi (1996): *A Probabilistic Theory of Pattern Recognition*. Berlin: Springer.
- Diestel, J.J. and Uhl, Jr. (1977): *Vector Measures*. Providence, R.I.: AMS (Mathematical Surveys no. 15).
- Gilboa, I. and D. Schmeidler (1994): "Infinite Histories and Steady Orbits in Repeated Games," *Games and Economic Behavior*, **6**, 370-99.
- Gilboa, I. and D. Schmeidler (1995): "Case-Based Decision Theory," *Quarterly Journal of Economics*, **110**, 605-39.
- Holland, J.H., K.J. Holyoak, R.E. Nisbett, and P.R. Thagard (1989): *Induction*. MIT Press: Cambridge.
- Klein, E. and A. C. Thompson (1984): *Theory of Correspondences*. New York: Wiley-Interscience Publication.

- Kreps, D. M. (1979): "A Representation Theorem for 'preference for flexibility'," *Econometrica*, **47**, 565-77.
- Kreps, D. (1992): "Static Choice in the Presence of Unforeseen Contingencies," in *Economic Analysis of Markets and Games*, ed. by P. Dasgupta, D. Gale, O. Hart and E. Maskin. Cambridge, MA: MIT Press, 258-81.
- Lipman, B. (1995): "Decision Theory without Logical Omniscience: Toward an Axiomatic Framework for Bounded Rationality". Forthcoming in: *Review of Economic Studies*.
- Lipman, B. (1993): "Information Processing and Bounded Rationality: A Survey," *Canadian Journal of Economics*, **28**, 42-67.
- McAfee, P. (1984): "Effective Computability in Economic Decisions," Western Ontario.
- Medin, D.L. and B.H. Ross (1992): *Cognitive Psychology*. New York: Harcourt Brace Jovanovich.
- Nehring, K. (1998): "Preference for Flexibility and Freedom of Choice in a Savage Framework". Forthcoming in: *Econometrica*.
- Pollak, R.A. (1968): "Consistent Planning," *Review of Economic Studies*, **35**, 201-8.
- Royden, H. L. (1968): *Real Analysis, 2ed Edition*. New York: MacMillan Publishing Co.,Inc..
- Shiryayev, A. N. (1984): *Probability*. New-York: Springer-Verlag.
- Simon, H. (1951): "A Formal Theory of Employment Relationship," *Econometrica*. **19**, 293-305.
- Strotz, R.H. (1955-56): "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies*, **23**, 165-180.
- Uhlig, H. (1996): "A Law of Large Numbers for Large Economies," *Economic Theory*, **8**, 41-50.