LEARNING ABOUT PREFERENCE

by

Evan Piermont

Bachelor of Arts, Lake Forest College, 2011

Master of Arts, University of Pittsburgh, 2013

Submitted to the Graduate Faculty of

the Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2017

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Evan Piermont

It was defended on

March 31, 2017

and approved by

Luca Rigotti, Department of Economics, University of Pittsburgh

Roee Teper, Department of Economics, University of Pittsburgh

Teddy Seidenfeld, Department of Philosophy, Carnegie Mellon University

Stephanie Wang, Department of Economics, University of Pittsburg

Dissertation Advisors: Luca Rigotti, Department of Economics, University of Pittsburgh,

Roee Teper, Department of Economics, University of Pittsburgh

LEARNING ABOUT PREFERENCE

Evan Piermont, PhD

University of Pittsburgh, 2017

This dissertation includes three essays, each investigating the formation and evolution of preference, with particular attention paid to epistemological concerns, dynamics, and learning. Loosely speaking, each chapter provides a different approach to modeling how a Decision Maker (DM) might learn about her own preference. The work as a whole adheres to the classical notions of revealed preference and utility maximization. The first chapter considers DM who learns by observing the environment in which decisions must be made. The second chapter, written jointly with Roee Teper, considers a DM who learns by observing the consequences of her prior choices. Finally, the third chapter considers a DM who learns *novel* information of which she previously had no conception.

TABLE OF CONTENTS

1.0 INTRODUCTION . . . 1
2.0 CONTEXT DEPENDENT BELIEFS . . . 3
2.1 INTRODUCTION . . . 3
2.1.1 Organization . . . 8
2.2 GENERAL MODEL . . . 8
2.2.1 Structure And Primitives . . . 8
2.2.2 Axioms . . . 9
2.2.3 Menu Induced Belief Representation . . . 14
2.2.4 State Independence . . . 15
2.3 BAYESIAN FRAMES . . . 16
2.4 FROM EQUILIBRIUM TO MBR . . . 20
2.5 LITERATURE REVIEW . . . 21
3.0 PLANS OF ACTION . . . 24
3.1 INTRODUCTION . . . 24
3.2 THE DECISION THEORETIC FRAMEWORK . . . 31
3.2.1 Choice Objects . . . 31
3.2.2 The Axioms . . . 34
3.2.3 A Representation Result and Belief Elicitation . . . 37
3.3 THE STATISTICAL FRAMEWORK . . . 38
3.4 STRONG EXCHANGEABILITY AND CONTEMPORANEOUS CORRELATIONS . . . 41
3.4.1 AA-SYM as a Behavioral Restriction . . . 43

3.5 FURTHER DISCUSSION 44
3.5.1 Related Literature 44
3.5.2 Subjective Learning with Endogenous and Exogenous Information 45
3.5.3 A Comment on Bayesianism in Environments of Experimentation 46
4.0 INTRODUCTION UNAWARENESS AND OBSERVABLE CHOICE 48
4.1 INTRODUCTION 48
4.1.1 Observability and Unawareness 53
4.1.2 A Simple Model of Unawareness 55
4.1.3 Organization 56
4.2 LOGICAL FOUNDATIONS: PREFERENTIAL LOGIC 57
4.2.1 A Syntactic Language: $\mathcal{L}(\mathcal{X})$ 57
4.2.2 Semantics 58
4.2.3 Axioms 60
4.2.4 Implicit Preferences 61
4.3 AWARENESS STRUCTURES 64
4.3.1 A Semantic Language: $\mathcal{L}(X)$ 66
4.3.2 The Semantics of Awareness 67
4.3.3 Explicit Preference 68
4.3.4 The Structure of Awareness 69
4.4 CONTINGENT PLANNING, OR, FINALLY, THE BEHAVIORAL IMPLICATION OF AWARENESS OF UNAWARENESS 71
4.4.1 Implicitly Known Contingent Plans 71
4.4.2 Explicitly Known Contingent Plans 76
4.5 A PREFERENCE FOR FLEXIBILITY 81
4.6 UNAWARENESS AND CONTRACTS 85
4.6.1 The Principal's Problem 87
4.6.2 Incompleteness Requires Introspection 88
4.7 LITERATURE REVIEW 89
APPENDIX A. APPENDIX OF CHAPTER 2 95
A.1 LEMMAS 95

A.2 PROOFS 98
APPENDIX B. APPENDIX FOR CHAPTER 3 104
B.1 ON THE CONSTRUCTION OF PLANS. 104
B.1.1 Construction Proofs 105
B.2 LEMMAS 108
B.3 PROOF OF MAIN THEOREMS 109
APPENDIX C. APPENDIX FOR CHAPTER 4 117
C.1 SUPPORTING RESULTS 117
C.2 SOUNDNESS AND COMPLETENESS RESULTS 120
C.3 PROOFS OMITED FROM THE MAIN BODY OF TEXT 122
BIBLIOGRAPHY 129

1.0 INTRODUCTION

I am writing this, the introduction, in transit from TLV to PIT via ATL. I have never written a dissertation, nor will I likely ever again, and so, I am not entirely sure what an introduction to one comprises. There are three chapters, pairwise unrelated except at a very general level where they are all decision theory and all contemplate how a decision maker might behave when she expects her ignorance to diminish. In other words, when she expects to learn.

I have imposed an strict artificial time limit: whatever introduction is written upon my arrival at PIT is what stands as the final cut.[1] In the modern fashion, the work presented here was not written with the intention of being a dissertation, per say—rather as three stand alone works, notationally self contained and intended for publication in academic journals. The first chapter was my *second year paper*, echoes of which, residing in bygone digital folders, were written near half a decade ago. Nonetheless, all three parts are still evolving, mutating in response to a chorus of referees and seminar participants. A dissertation as a snapshot rather than a capstone.

While in Israel, I presented the second and third chapters of this dissertation, at the University of Haifa and Tel Aviv University, respectively. The chapters are ordered, loosely speaking, by their level of abstraction. Each adheres to the paradigm of revealed preference and utility theory. I ascribe to such principals almost religiously—not in the sense that they contain some deep insight, but because to deny them is to deny any foundation to human behavior.

Definition. *A decision maker satisfies the* ***revealed preference paradigm*** *if at each instance, she chooses the action that she believes will afford her the most pleasing outcome.*

[1] So any unchecked pretentious flourish and malapropistic usage errors can be chocked up to jet-lag and a lack of internet.

Many people outside neoclassical decision theory point to the cacophony of bad decision taken more or less constantly as a damning counterexample to utility theory. However, such a superficial criticism misses, or misunderstands, the inclusion of "she believes" from the definition above. Humans are computationally constrained, myopic, altruistic, addicted, indecisive, and psychologically biased in sundry, ad hoc ways. We make the best possible choices subject to this immense set of constraints. It is, in part, the decision theorist's job[2] to understand, when given a set of choices, how to identify the constraints which effected these choices.

The work contained in this dissertation is devoted to understanding one constraint in particular: that we are currently ignorant of the (Platonically) optimal choice, but acknowledge the possibility that we will shortly become less ignorant. We make choices today anticipating the arrival of better information tomorrow—information that is often a function of our present decisions. The following three chapters provide different prospectives on (i) how such anticipation shapes behavior, and (ii) how, by looking at choices, one might be able to identify what the decision maker expects she might learn.

The trip to Israel was for the purpose of visiting Roee, his family, and his alma mater, and had been discussed intermittently for years. The idea that such a trip would follow the acquisition of a job: after the myriad cover letters, the pacing in hotel hallways, the *JMP*, the *spiel*, the dinners with only one in formal attire. In contrast to popular notions of Ph.D. candidacy, it was these events, and not the preparation of this document, that were both the principal source of apprehension and marker of end of my tenure as a student. It ended well; I got a job; I got to go to Israel and present the second and third chapters of my dissertation.

I would like to thank: my friends and family for putting up with my insufferable chatter; my fellow CORE members; my always enlightening professors (heck, even the only occasionally enlightening ones); my letter writers; my committee members; Sum for putting up with my insufferable chatter and everything else; and, of course, my wonderful advisors, Luca and Roee. Y'all have made the last six years an absolute delight.

[2] I think?

2.0 CONTEXT DEPENDENT BELIEFS

In which Katya discovers the quality of a restaurant—not to mention of her date—by reflecting on its offerings.

2.1 INTRODUCTION

Both intuition and psychological evidence insist that a decision maker's (DM's) preference over alternatives is affected by the environment in which the decision is made (Kahneman and Tversky, 1984; Simonson and Tversky, 1992; Sen, 1993). While there are many external factors that potentially exert influence on the decision making process, this paper examines a model in which the set of alternatives that is *currently* available acts as a frame –a process often differentiated from general framing effects under the moniker *context dependence*. I identify the behavioral conditions for context dependent beliefs, when the DM's subjective assessment of the likelihood of events depends on the available alternatives (the menu) from which he must choose, and consider additional restrictions that correspond to particular subjective information structures.

Context dependence is often associated with notions of bounded rationality or psychological heuristics (Tversky and Simonson, 1993). This paper, however, interprets menu-induced framing as rational, exploring how and when such behavior exists within the subjective expected utility paradigm. If the DM believes the menu itself contains information regarding payoff relevant uncertainty, conditioning his preference on such information is a rational action. Specifically, the model assumes the payoff associated with each alternative is ex-ante uncertain. The DM's utility from consumption depends not only on the chosen outcome, but also on which *state of the world* is realized. The DM, before consumption, is uncertain about the state of the world, but holds a belief (a probability distribution) over the state space; in a given decision problem, the DM maximizes his expected utility according to his belief. When the DM interprets the current selection of alternatives as a signal about the state of the world, his preferences will change across different decision problems in response to his updated beliefs.

Before expounding the finer points of the model, it is worth considering two examples to better illustrate why menu-dependent preferences are indeed necessary to explain many decision making scenarios.

Example 2.1.A (Luce's diner). *On a first date, Katya finds herself in a restaurant at which she has previously never eaten, and which offers chicken (c) or steak (s). She states her strict preference for chicken ($c \succ s$). However, upon seeing the restaurant also serves frog legs (f), she now states her strict preference for steak ($s \succ c \succ f$).*

While Katya's preference reversal in the face of a (seemingly) irrelevant alternative cannot be accommodated by the standard theory (as it violates the weak axiom of revealed preference (WARP)), it has a simple, intuitive explanation. She prefers steak when the food is well prepared, but considers chicken more resilient to the inept chef. In the typical restaurant, she believes it is unlikely the food will be well cooked, and hence, has a preference for chicken. However, in the presence of an exotic dish, she deems it is more likely the restaurant employs an expert chef and so, reverses her preference.

Example 2.2.A (Sen's date). *After dinner, Katya's date, Mitya, asks whether she would like to end the date and go home (h) or go next door and get a drink (d). Thinking the date a success, Katya strictly prefers getting a drink ($d \succ h$). However, before she can respond, Mitya offers a third option: the acquisition and consumption of crystal methamphetamine (m). Katya now strictly prefers going home ($h \succ d \succ m$).*

Here, again, Katya's rather intuitive behavior cannot be explained by standard theory. She understands the offer of methamphetamine as a signal regarding Mitya's character. So, while she would prefer to continue the date as long as it is likely Mitya is reputable, his proposition is sufficient to sway her beliefs away from such a outcome.

These vignettes exemplify two main behaviors of the model. First, it is only the DM's perception of uncertainty that is changing; ex-post tastes are fixed. In other words, if the DM knew with certainty which state of the world was to be realized, he would exhibit a constant preference across menus. Second, the uncertainty is *local*. The realization regarding the quality of the food in one restaurant is not informative about the quality in a different restaurant; that a previous date was incorruptible is not evidence that a future date will be.[3]

The first part of this paper axiomatizes a particular type of context-dependence which adheres to these two restrictions. As in Anscombe and Aumann (1963), I examine a DM who ranks *acts* (i.e., functions) from a state space, S, into lotteries over consumption, $\Delta(X)$.[4] Naturally, given the motivation, not all of X will always be available. The DM's entertains a family of preferences over acts, indexed by the subset of X that is currently available. Therefore, for each $A \subseteq X$, we see the decision maker's preference, $\succcurlyeq_A$ over $\{f : S \to \Delta(A)\}$. Then, a *menu-induced belief representation* (MBR) is a single utility index, $u : S \times X \to \mathbb{R}$, and a menu-indexed family of beliefs $\{\mu_A\}_{A \subseteq X} \subseteq \Delta(S)$ such that

$$U_A(f) = \mathbb{E}_{\mu_A}\left(\mathbb{E}_{f(s)}(u(s,x))\right) \tag{MBR}$$

represents $\succcurlyeq_A$, where $\mathbb{E}_\pi(\varphi)$ denotes the expectation of the random variable φ with respect to the distribution π. Fixing the menu, the DM acts as a subjective expected utility maximizer. The utility index, u, is the same across menus. This is the consequence of the main axiom, *menu consistency*. Menu consistency dictates, conditional on the realization of a particular state, the DM's preference for alternatives is fixed across menus. Therefore, the context effect is entirely characterized by the change in the DM's beliefs regarding the state space. This places clear limits on the type of context effects that can be accommodated by a MBR. Since any change in preference is the consequence of shifting beliefs, context dependence cannot reverse preference over outcomes for which the resolution of the state is payoff

[3]Of course, one could tell a different story where there is a dynamic component by which the DM learns about the likelihood of states from experience. This is well outside of the current model.

[4]For a set Y, $\Delta(Y)$ is the set of distributions thereover.

irrelevant (note, because the tastes are state-dependent, constant acts are not necessarily certain outcomes).

Although the general model imposes *a continuity condition*[5] –if two menus differ only slightly, then so do their associated beliefs– it does not otherwise specify any restriction relating menus with their associated beliefs. The second part of this paper, therefore, explores how menus might correspond to the beliefs they induce. In particular, what restrictions correspond to the DM who, acting in Bayesian manner, interprets each menu as a collection of signals regarding the relative likelihood of each state. And, what further restrictions allow us to identify the structure of these signals. Following the *anything goes* result of Shmaya and Yariv (2015), any MBR can be rationalized by some prior and set of signals. Thus, without imposing any additional structure, it is impossible to rule out the possibility that the DM is acting in a Bayesian fashion.

I consider two more restrictive signal structures and their corresponding behavioral restrictions. In the first, a *partitional signal structure*, the DM entertains a partition of the state space and a prior belief over the likelihood of each state. Each menu induces a belief whose support coincides with some cell of the partition, and any two menus which induce beliefs with the same support carry the same informational content. In other words, the DM believes each menu can only occur in a particular subset of the state space, but, given a state and the menus possible in that state, the realized menu is chosen uniformly. This signal structure is of particular interest, as it could be seen as arising from endogenously from a separating equilibrium in a game between buyers and sellers (see Section 2.4).

In the second signal structure I consider, an *elemental signal structure*, the DM takes the elements of the menus as signals rather than the menus themselves. Specifically, he assumes that in each state, s, the inclusion or exclusion of an element x is decided according to the toss of a (potentially biased) (s, x)-coin. Therefore, the collection of included elements (the menu) is the result of a series of conditionally-independent coin tosses.

Revisiting the examples, we can see that MBR preferences and these various information structures can rationalize the choice patterns.

[5]This is a vacuous assumption when X is a discrete space.

Example 2.1.B (Luce's diner, revisited). *Katya's has MBR preferences and the following utility index:*

$$u(h,c)=8 \quad u(h,s)=16 \quad u(h,f)=8$$
$$u(m,c)=8 \quad u(m,s)=6 \quad u(m,f)=4$$

where $S=\{h,m\}$ indicates high and medium quality food, respectively. She initially believes it equally likely the food is high quality as it is mediocre: $\mu(h)=\frac{1}{2}$ and $\mu(m)=\frac{1}{2}$. She also believes only (and all) high quality restaurants offer frog legs.

So the menu $\{c,s\}$ indicates with certainty the food is mediocre, so $U_{\{c,s\}}(c)=8>6=U_{\{c,s\}}(s)$, while the menu $\{c,s,f\}$ indicates with certainty the food is good, so $U_{\{c,s,f\}}(s)=16>8=U_{\{c,s,f\}}(c)$.

So a MBR and partitional signal structure can rationalize Katya's choices over entrees. Likewise, we can see that an elemental information structure could explain her preferences regarding Mitya.

Example 2.2.B (Sen's date, revisited). *Katya's has MBR preferences and the following utility index:*

$$u(r,h)=1 \quad u(r,d)=5 \quad u(r,m)=-10$$
$$u(d,h)=1 \quad u(d,d)=-5 \quad u(d,m)=-10$$

where $S=\{r,d\}$ indicates reputable and depraved characters, respectively. She initially believes $\mu(r)=\frac{9}{10}$ and $\mu(b)=\frac{1}{10}$. She also believes that, while all dates will offer going home and getting a drink, depraved characters offer meth with probability $\frac{1}{10}$, with reputable characters with only probability $\frac{1}{100}$.

After updating upon seeing the menu $\{h,d\}$, she holds the beliefs $\mu(r)=\frac{891}{981}$ and $\mu(d)=\frac{90}{981}$; her preference is given by $U_{\{h,d\}}(d)=\frac{5(801)}{981}>1=U_{\{h,d\}}(h)$. After the menu $\{h,d,m\}$, she holds the beliefs $\mu(r)=\frac{9}{19}$ and $\mu(d)=\frac{10}{19}$; her preference is given by $U_{\{h,d,m\}}(h)=1>\frac{-5}{19}=U_{\{h,d,m\}}(d)$.

2.1.1 Organization

This paper is organized as follows. The general model is presented in Section 2.2, with the representation theorem for the main result contained in Section 2.2.3. Section 2.2.1 discusses the shortcomings of a variant model with state-independent utilities. Section 2.3 explores the additional restrictions necessary to capture particular signal structures. Section 2.4 informally explores how menu-dependent beliefs could arise naturally in a strategic environment. Finally, a survey of relevant literature is found in Section 2.5. All proofs are contained in the appendix.

2.2 GENERAL MODEL

2.2.1 Structure And Primitives

Let X be a compact and metrizable topological space, representing the grand set of consumption alternatives, and with typical elements x, y, z. Define $x^\star$ and $x_\star$ to be two distinguished elements of X, referred to as universal alternatives, and set $\bigstar = \{x^\star, x_\star\}$. Let $\mathcal{P}(X)$ denote the set of compact subsets of X; endow $\mathcal{P}(X)$ with the hausdorff metric (thus, a compact metric space). Let $\mathcal{K}(X)$ denote the subset of $\mathcal{P}(X)$ whose elements contain $\bigstar$.[6] Typical elements are A, B, C. Elements of $\mathcal{K}(X)$ are called menus, with the interpretation that they represent the set of *currently available* consumption alternatives.

For any topological space Y, let $\Delta(Y)$ denote the set of all probability measures on $(Y, \mathcal{B}(Y))$, where $\mathcal{B}(Y)$ denotes the Borel σ-algebra on Y, endowed with the topology of weak convergence. If $\mu \in \Delta(Y)$, and φ is a continuous and bounded function $\varphi : Y \to \mathbb{R}$, then

$$\mathbb{E}_\mu(\varphi(y)) = \begin{cases} \displaystyle\int_Y \varphi(y)\ d\mu(y) & \text{whenever } Y \text{ is infinite, and,} \\ \displaystyle\sum_Y \varphi(y)\mu(y) & \text{whenever } Y \text{ is finite,} \end{cases}$$

denote the expectation of φ with respect to μ.

[6] I will interpret $\bigstar$ as a set of outside options, which explains their universal availability.

Notice, for all $A \in \mathcal{K}(X)$, A is compact, (hence separable), so $\Delta(A)$ is metrized by the Lèvy–Prokhorov metric. In the standard abuse of notation, identify $x \in X$ with the degenerate distribution on x. Typical elements of $\Delta(X)$ are denoted π, ρ, τ.

Let S denote a finite state space. Endow $\Delta(X)^S$ with the product topology. The objects of choice will be menu-induced acts: for each $A \in \mathcal{K}(X)$ define $\mathcal{F}_A = \Delta(A)^S \cong \{f : S \to \Delta(X) | f(s)[A] = 1\}$. An act is a commitment to a particular consumption conditional on the realization of the type space, and so, $\mathcal{F}_A$ corresponds to the acts available given the menu A (which put probability 1 on an outcome that is available from A). For each act, $f(s)$ is the distribution over X obtained for realization s. Again, abusing notation, identify each $\pi \in \Delta(X)$ with the degenerate act such that $\pi(s) = \pi$ for all s.

For any $f, g \in \mathcal{F}_X$, and event $E \subseteq S$, let $f_{-E}g$ be the act that coincides with f everywhere except on E, where it coincides with g. Further, for some $\alpha \in (0, 1)$ let $\alpha f + (1 - \alpha)g$ be the point-wise mixture of f and g (i.e., $(\alpha f + (1 - \alpha)g)(s) = \alpha f(s) + (1 - \alpha)g(s)$ for each $s \in S$). It is immediate that if $f \in \mathcal{F}_A$ and $g \in \mathcal{F}_B$ then $f_{-E}g$ and $\alpha f + (1-\alpha)g$ both belong to $\mathcal{F}_{A \cup B}$ (in particular, note the case when $A = B$).

The primitive of the model is the family of preference relations $\{\succcurlyeq_A \subset \mathcal{F}_A \times \mathcal{F}_A\}_{A \in \mathcal{K}(X)}$. That is, the DM's preference over the acts which are available given each possible menu. With regards to notation, it is assumed whenever '$f \succcurlyeq_A g$' is written both f and g belong to $\mathcal{F}_A$. For any relation $\succcurlyeq$, let $\succ$ and $\sim$ denote the asymmetric and symmetric components, respectively.

2.2.2 Axioms

The goal of the most general representation is to provide the basic framework in which a DM might condition his beliefs regarding the state space on the menu at hand. That is, the DM treats the set of currently available consumption alternatives as a signal regarding the likelihood of different states. Given the menu, the DM acts as a subjective expected utility maximizer, with respect to his menu-induced beliefs. Clearly, each menu-dependent preference should satisfy the expected utility axioms:

[AX.2.1: EU] *For each $A \in \mathcal{K}(X)$, $\succcurlyeq_A$ satisfies the expected utility axioms, namely,*

1. ***Weak Order***. *$\succcurlyeq_A$ is a non-trivial weak order.*
2. ***Independence***. *For all $f, g, h \in \mathcal{F}_A$ and $\alpha \in (0,1)$, $f \succcurlyeq_A g \iff \alpha f + (1-\alpha)h \succcurlyeq_A \alpha g + (1-\alpha)h$.*
3. ***Continuity***. *For all $f \in F_A$, the sets $\{g \in \mathcal{F}_A | g \succcurlyeq_A f\}$ and $\{g \in \mathcal{F}_A | f \succcurlyeq_A g\}$ are closed in $\mathcal{F}_A$.*

The following well known result (so well known in fact, that it is included only to fix notation for expositional purposes) shows EU provides the expected utility structure for each menu-dependent preference.

Proposition 2.1 (Expected Utility)**.** *$\{\succcurlyeq_A \subset \mathcal{F}_A \times \mathcal{F}_A\}_{A \in \mathcal{K}(X)}$ satisfies EU if and only if for each $A \in \mathcal{K}(A)$ there exists some continuous and bounded $w : S \times X \to \mathbb{R}$ such that*

$$U_A^{VNM}(f) = \sum_s \left(\mathbb{E}_{f(s)}(w_A(s,x)) \right),$$

represents $\succcurlyeq_A$. Moreover, if $w_A(s,x)$ and $\hat{w}_A(s,x)$ both represent $\succcurlyeq_A$, then $w_A(s,x) = a_A \hat{w}_A(s,x) + b_A(s)$ where $a_A \in \mathbb{R}_{++}$ and $b_A(s) \in \mathbb{R}$ for all $s \in S$.

Recall that $f(s)$ and $g(s)$ are given, objective probability measures. The index $w(\cdot)$ can be decomposed into tastes (the utility of consuming an object given the state) and beliefs (the subjective likelihood of each state). Indeed, choose some probability distribution $\mu \in \Delta(S)$ such that $\mu(s) > 0$ and let $u_A(s,x) = \frac{w_A(s,x)}{\mu(s)}$; it is clear that

$$\mathbb{E}_\mu \left(\mathbb{E}_{f(s)}(u_A(s,x)) \right)$$

represents $\succcurlyeq_A$. Of course, this creates the classic problem of multiple rationalizing beliefs: if we consider some other $\nu \in \Delta(S)$ such that $\nu(s) > 0$, then ν and $u'_A(s,x) = \frac{w_A(s,x)}{\nu(s)}$ also represent the same preference. We cannot identify the DM's tastes for ex-post outcomes or his beliefs; the two are jointly determined.

The motivation for expanding our data to include the family of menu-induced preferences is to understand how the menu can alter the beliefs of the DM. In light of this, it becomes obvious further structure is needed to separate the effect on the perception of uncertainty (i.e., menu induced changes in belief) from other internal changes in preference (i.e., a change in tastes).

The first novel axiom, *menu-consistency*, is the first step towards such a disentanglement, and, captures the main behavior behind menu-induced beliefs. It states that the DM's tastes for outcomes do not depend on the menu at hand. This implies, any difference in preferences across menus must be the result of a change in perception of the underlying uncertainty.

Of course, the DM only cares about the assignment to state s if he believes there is a possibility s will be realized. Therefore, menu-consistency should only hold after realizations assigned positive probability according to the DM's subjective assessment. To make such ideas precise, I first need to consider null events.

Definition. *An event, $E \subset S$, is **null for menu A** (hereafter, null-A) if for all $f, g \in \mathcal{F}_A$,*

$$f_{-E}g \sim_A f.$$

Let N_A denote the set of states that are null-A, and N denote the set of everywhere null states: $N = \bigcap_{A \in \mathcal{K}(A)} N_A$.[7]

Null events, in general, have two indistinguishable interpretations. First, that the DM is indifferent between all available options conditional on the realization of the null event, E; second, that the DM places zero probability on E occurring. However, assuming the DM's tastes are consistent across different menus (the assumption that will be formalized shortly), it is possible to differentiate these two interpretations of null events. If a state, s, is null-A, but there exists a different menu, B, for which the DM displays a strict preference over elements of A (given realization of s), it must be that s was assigned zero probability when facing A. This is because the DM cannot be indifferent to all elements of A (contingent on s) since he displays strict preference in the menu B. This is formalized by evidently-null events, first considered in Karni *et al.* (1983).

Definition. *An event, $E \subset S$, is **evidently null for menu A** (hereafter, e-null-A) if E is null-A and for all $s \in E$ there exists some menu B such that*

$$(f_{-s}g) \succ_B f$$

for some f, g in $\mathcal{F}_{A \cap B}$. Let E_A denote the union of all e-null-A events.[8]

[7]Notice, the set of null-A events form a lattice with respect to set inclusion, with N_A the maximal element.

[8]Notice, the set of e-null-A events form a sub-lattice of the lattice of null-A events, with E_A the corresponding maximal element.

With this definition in mind we can now define menu consistency.

[AX.2.2: MC] *For all* $A, B \in \mathcal{K}(X)$ *and* $s \in S$ *with* $s \notin E_A \cup E_B$*, and all* $f \in \mathcal{F}_A$*,* $g \in \mathcal{F}_B$*,* $h \in \mathcal{F}_{A \cap B}$*, and such that* $f(s) = g(s)$*,*

$$f_{-s}h \succcurlyeq_A f \iff g_{-s}h \succcurlyeq_B g.$$

If $\{\succcurlyeq_A\}_{A \in \mathcal{K}(X)}$ is menu-consistent, the DM's tastes for outcomes are identical across menus. To see this, let $\pi = f(s) = g(s)$ and $\rho = h(s)$. Then MC states that the DM's preference between ρ and π, in state s, does not depend on the context in which the decision is made (i.e., does not depend on the menu from which the acts were constructed). Behaviorally, this indicates that any context effect does not alter the DM's preferences *conditional* on the realization of a particular state. In other words, if the DM knew the true state, there would be no context effect. It is this restriction that differentiates this model from a more general interpretation of context effects as psychological biases without foundation in rational behavior. The change in behavior across menus is *not* the result of a change in the state-dependent preference for outcomes (objects about which the DM is ostensibly certain) but of a change in his perception of the between-state-tradeoffs (the domain of uncertainty).

By the very nature of the problem at hand, we are losing structure in comparison to the standard model and so the axioms are weaker in comparison. As such, MC does not characterize a *new* behavior that is the result of context dependent beliefs, but rather places limits on how much structure is lost. What structure is retained by MC guarantees we can find a family of representation for $\{\succcurlyeq_A\}_{A \in \mathcal{K}(X)}$ that shares a common utility index. In other words, the primitive is represented by a single utility index and a family of menu-induced beliefs. It is important to note that this does not rule out preference reversals, even over constant acts. Each menu carries with it a perception of uncertainty, and can therefore change the DM's preferences for acts. However, given menu-consistency, any preference reversal is due entirely to the change in beliefs, and not because of changes in ex-post tastes. Setting $f = g$ in the definition, consistency guarantees that the ordering between $f_{-s}\rho$ and $f_{-s}\pi$ hold regardless of the ambient menu.

Under the definition of a *frame* as (seemingly irrelevant) information which alters the DM's perception of uncertainty, then EU and MC exactly capture the behavior where the

DM uses the menu as a frame. Unfortunately, from a practical vantage, this is insufficient, as the problem of non-uniqueness of beliefs persists. When tastes and beliefs cannot be separated, we cannot identify the avenue by which context effects alter the DM's choice process.

To overcome the issue of multiple rationalizing beliefs, Anscombe and Aumann (1963) restricted preferences to be state independent (i.e., in every non-null state, the ranking over distributions is the same). State dependency is a very restrictive assumption; it interprets states as abstract probabilistic events that have no meaning outside of their use as betting devices. Beyond this philosophical issue, state-dependence is a necessary requirement to capture the full gamut of context effects (this necessity is made precise in Remark 2.4). For these reasons, this model weakens state-independence to apply only over $\bigstar$. UV plays the same roles as state independence (equivalently, monotonicity). By ensuring, over the relatively small domain $\bigstar$, that preferences in each state coincide, beliefs can be uniquely recovered from choice data. Under the interpretation of universal elements as outside options, it is natural that the ranking of these elements does not change across different menus.

[AX.2.3: UV] *For all $A \in \mathcal{K}(X)$ and $s \in S$ with $s \notin N_A$*

$$f_{-s}x^{\star} \succ_A f_{-s}x_{\star}.$$

for all $f \in \mathcal{F}_A$.

It is also of interest (when X is infinite) to understand when the context effect acts in a continuous manner.

[AX.2.4: CC] *If $\{A_n\}_{n\in\mathbb{N}}$ converges to A in $\mathcal{K}(X)$, then for all $\alpha \in (0,1)$,*

$$\big\{\{g \in \mathcal{F}_{\star} \mid g \succcurlyeq_{A_n} \alpha x^{\star} + (1-\alpha)x_{\star}\}\big\}_{n\in N}, \text{ and}$$
$$\big\{\{g \in \mathcal{F}_{\star} \mid \alpha x^{\star} + (1-\alpha)x_{\star} \succcurlyeq_{A_n} g\}\big\}_{n\in N}$$

converge to $\{g \in \mathcal{F}_{\star} | g \succcurlyeq_A \alpha x^{\star} + (1-\alpha)x_{\star}\}$ and $\{g \in \mathcal{F}_{\star} | \alpha x^{\star} + (1-\alpha)x_{\star} \succcurlyeq_A g\}$, respectively, in $\mathcal{P}(X)$.

In other words, the menu-induced contour sets, when restricted to universal acts, converge whenever the menus converge. Because CC applies only to universal acts, it restricts only that beliefs converge (and has nothing to say about the change in tastes across menus). Convergence, of both the menu and the contour sets, is with respect to Hausdorff metric in the respective ambient spaces. So, CC states that as menus become close, the relative weights placed on each state must also converge. Notice, while the sufficiency of CC is clear, the necessity relies on the fact that the utility derived from the set of universal acts is bounded.

2.2.3 Menu Induced Belief Representation

Theorem 2.2 (Menu Induced Belief Representation)**.** *(a)* $\{\succcurlyeq_A\}_{A\in\mathcal{K}(X)}$ *satisfies* EU*,* MC*,* UV*, and* CC *if and only if there exists a state-dependent utility index,* $u : S \times X \to \mathbb{R}$*, such that* $u(\cdot, x^\star) \equiv 1$ *and* $u(\cdot, x_\star) \equiv 0$*, and such that the projections* $u|_A$ *are bounded and continuous for all* $A \in \mathcal{K}(X)$*, and a continuous function,*[9] $\{\mu_A \in \Delta(S)\}_{A\in\mathcal{K}(X)}$*, such that for all* $A \in \mathcal{K}(X)$*,*

$$U_A(f) = \mathbb{E}_{\mu_A}\left(\mathbb{E}_{f(s)}(u(s,x))\right) \tag{MBR}$$

represents $\succcurlyeq_A$*, and* $\mu_A(s) = 0$ *if and only if* $s \in E_A \cup N$*.*

(b) Moreover, the family of beliefs $\{\mu_A \in \Delta(S)\}_{A\in\mathcal{K}(X)}$ *is unique and the utility index,* $u(\cdot)$*, is unique up to null states.*

Proof. In appendix A.2.

The proof is quite straightforward. First, EU provides a linear representation for each $\succcurlyeq_A$. By UV these representations can be decomposed uniquely into a tastes (over A) and beliefs, where the utility index is normalized as in the statement of the theorem. Then, these utility indexes can be stitched together to provide a single u over the whole of X. Finally, MC ensures that this common index will jointly represent each $\succcurlyeq_A$ and CC that beliefs will change continuously.

[9] i.e., $\mu_{(\cdot)} : \mathcal{K}(X) \to \Delta(S)$ is continuous with respect to $\Delta(S)$ when endowed with the topology of weak convergence.

Because the utility index is fixed across decision problems, the shifting of probabilities is the only avenue for preferences to change. Thus, if an act f is preferred to g on a state-by-state basis, then it is preferred to g in every menu (this is, of course, precisely the content of MC). It is through the menu-dependent beliefs that this structure allows for framing effects, were by the DM's preferences change in the face of new alternatives. It follows that the types of preference reversals that are allowable is limited.

2.2.4 State Independence

In light of axiom UV, it may seem parsimonious to quit worrying about the distinguished elements, $x^\star$ and $x_\star$, and require state independence outright. This can, in fact, be accomplished by strengthening MC.

[AX.2.2*: SMC] *For all* $A, B \in \mathcal{K}(X)$ *and* $s \in S$ *with* $s \notin E_A$ *and* $s' E_B$*, and all* $f \in \mathcal{F}_A$*,* $g \in \mathcal{F}_B$*,* $h \in \mathcal{F}_{A\cap B}$*, and such that* $f(s) = g(s')$*,*

$$f_{-s}h \succcurlyeq_A f \iff g_{-s'}h \succcurlyeq_B g.$$

SMC states that tastes are consistent, not only across menus (if $A \neq B$) but also across states (if $s \neq s'$). As such, it implies the canonical form of state independence for each $\succcurlyeq_A$. When MC is replaced by SMC in Theorem 2.2, the resulting representation coincides except the utility index, $u : X \to \mathbb{R}$ is *state-independent*:

$$U_A^{SI}(f) = \mathbb{E}_{\mu_A}\left(\mathbb{E}_{f(s)}(u(x))\right), \tag{SI-MBR}$$

represents $\succcurlyeq_A$.[10] The existence of the family of beliefs, their uniqueness, and the uniqueness of the utility index are all the same as in Theorem 2.2. While this approach is only a small deviation from the general representation, it implies that there is no uncertainty regarding the preference of constant acts. As discussed before, in order for context effects to have observable content, it must be that the underlying uncertainty is payoff relevant. Together, these facts imply that SMC prohibits the DM from changing his preference over constant acts between different menus.

[10]Notice, UV is somewhat redundant in the presence of SMC. In fact, if we are willing to entertain a bit of notational juggling, we can forego it entirely.

Remark 2.3. *Let $\{\succcurlyeq_A\}_{A\in\mathcal{K}(X)}$ be represented by* (SI-MBR)*. Then for all $A, B \in \mathcal{K}(X)$, and $\pi, \rho \in \Delta(A \cap B)$, $\pi \succcurlyeq_A \rho \iff \pi \succcurlyeq_B \rho$.*

Remark 2.3 can be seen by noting that $U_A^{SI}(\pi) = \mathbb{E}_\pi(u(x))$, which does not depend on A.

2.3 BAYESIAN FRAMES

The general representation, (MBR), allows for context effects by which the DM's preferences over acts depends on the menu of currently available outcomes. It does not, however, offer any insight into the connection between the menu at hand and the effect it exerts on decision making. This section provides an exploration into the behavioral implications of particular context effects.

It is of interest to identify the restrictions on behavior that ensure the DM is acting rationally with respect to some *information structure* that gives rise to the family of menu-induced beliefs. Consider the interpretation that the DM entertains a prior belief over the state space, $\mu \in \Delta(S)$, and observes, along with the menu, some signals, drawn from a (finite) signal space, Θ. The DM also entertains a likelihood function that specifies the likelihood of a given signal, contingent on the true state, $l : \Theta \times S \to \mathbb{R}_+$. Under this interpretation, we say the information structure (μ, l, Θ) *generates* $\{\mu_A | A \in \mathcal{K}(X)\}$, if the DM's menu-induced beliefs are the posteriors generated by observing the signals. To keep things notationally clean, through this subsection, I assume that X is finite and $N = \emptyset$. I always assume the prior, μ, has full support. These assumptions ensure the updating procedures are binding everywhere, as it alleviates the concern regarding 0 probability events.

Of course, for the posteriors to be indexed by menus there must be a connection between the signals and the menu. At the most general level, the two coincide: $\Theta = \mathcal{K}(X)$.

Definition. *An information structure based on menus, $(\mu, l, \mathcal{K}(X))$, generates $\{\mu_A | A \in \mathcal{K}(X)\}$, if*

$$\mu_A(s) = \frac{\mu(s)l(A,s)}{\mathbb{E}_\mu(l(A,s'))} \tag{2.1}$$

and $\sum_{\mathcal{K}(X)} l(A,s) = 1$, $\sum_s l(A,s) > 0$ *for all* $A \in \mathcal{K}(X)$ *and* $s \in S$.[11]

Unfortunately, the requirement that the DM entertains some generating $(\mu, l, \mathcal{K}(X))$, provides no testable implications. In other words, *every* MBR can be described by some prior and likelihood function over menu realizations. The ability to choose both the signals and the prior provides enough degrees of freedom such that bayesianism can never be ruled out.

Proposition 2.4. *Let* $\{\succcurlyeq_A\}_{A\in\mathcal{K}(X)}$ *be represented by some* (MBR)*, with beliefs* $\{\mu_A | A \in \mathcal{K}(X)\}$*. Then there exists some* $(\mu, l, \mathcal{K}(X))$ *that generates* $\{\mu_A | A \in \mathcal{K}(X)\}$ *as in* (2.1)*.*

Proof. In appendix A.2.

This result is a corollary of Lemma 1 in Shmaya and Yariv (2015). Setting $\Theta = \mathcal{K}(X)$ assumes no relation between the signals associated with different menus, and it is this generality that renders behavior wholly unconstrained. However, under more specific assumptions regarding the structure of the signals, there are falsifiable restrictions on observable preference. Thus, while we can never rule out the possibility that the DM is acting in a Bayesian manner with respect to *some* signal space, we can rule out particular models of information.

First, consider a partitional information structure. Under this model of information, each menu obtains only on a cell of a partition of the state space, and any menu that obtains on the event E contains the same informational content. That is, we set $\Theta = \mathcal{K}(A)/\dot{\sim}$ for some equivalence relation $\dot{\sim}$ (with the equivalence class containing A, denoted by $[A]$).

Definition. *A partitional information structure,* $(\mu, l, \mathcal{K}(X)/\dot{\sim})$*, generates* $\{\mu_A | A \in \mathcal{K}(X)\}$ *if*

$$\mu_A(s) = \frac{\mu(s) l([A], s)}{\mathbb{E}_\mu(l([A], s'))} \tag{2.2}$$

and $l([A], s) \in \{0, 1\}$ *for all* $([A], s) \in \mathcal{K}(X)/\dot{\sim} \times S$*, and* $\sum_{s\in S}(l([A], s) \cdot l([B], s)) > 0$ *implies* $[A] = [B]$*.*

[11] The first requirement is equivalent to $\sum_{\mathcal{K}(X)} l(A,s) > 0$, and is included in the current form only for its interpretational content. Under this normalization, we can think of $l(A,s)$ as the probability of seeing menu A in state s. The second requirement states that all menus are ex-ante possible. This ensures that 2.1 is always well defined.

The requirement that the likelihoods are binary is indicative of the fact that within equivalence classes all menus have the same informational content. Hence, any two menus that have the same support, must induce the same beliefs (and therefore, we can normalize the likelihood functions to 1). The requirement that $l([A], s) \cdot l([B], s) = 0$ for any $[A] \neq [B]$ ensures that the supports of the distributions are disjoint, and hence, form a partition of S.

[AX.2.5: PS] *For all $A, B \in \mathcal{K}(X)$, such that $N_A^c \cap N_B^c \neq \emptyset$, and $f, g \in \mathcal{F}_{A\cap B}$, we have*

$$f \succsim_A g \iff f \succsim_B g.$$

PS dictates that any two menus sharing a non-null state must induce the same preference over acts. Because the general representation fixes tastes across different menus, PS implies that if the two menus induce beliefs with a overlapping supports, those beliefs must coincide completely. It is clear that this captures the behavior generated by a partitional signal structure.

Theorem 2.5. *Let $\{\succsim_A\}_{A\in\mathcal{K}(X)}$ be represented by some (MBR), with beliefs $\{\mu_A | A \in \mathcal{K}(X)\}$. Then, there exists some $(\mu, l, \mathcal{K}(X)/\dot{\sim})$ that generates $\{\mu_A | A \in \mathcal{K}(X)\}$ as in (2.2) if and only if $\{\succsim_A\}_{A\in\mathcal{K}(X)}$ satisfies PS.*

Proof. In appendix A.2.

An even more restrictive signal structure would be $\Theta = X$ with signals that are conditionally independent. Under this interpretation, the element x is made available (in state s) according to the flip of an (s, x)-coin. Conditional on the state, the coins are independent of one another, although their bias can vary with both the element and the state.

Definition. *An information structure based on elements, (μ, l, X), with conditionally independent signals generates $\{\mu_A | A \in \mathcal{K}(X)\}$ if*

$$\mu_A(s) = \frac{\mu(s) \prod_{x\in A} l(x, s) \prod_{y\notin A}(1 - l(y, s))}{\mathbb{E}_\mu \left(\prod_{x\in A} l(x, s') \prod_{y\notin A}(1 - l(y, s'))\right)}, \tag{2.3}$$

and $l(x, s) \in (0, 1)$ for all $(x, s) \in X \setminus \bigstar \times S$, and $l(x^\star, s) = l(x_\star, s) = 1$ for all $s \in S$.

The requirement that likelihoods lie in the interior of $(0,1)$ is tantamount to assuming there are no null states, and ensures that (2.3) is well defined for all menus and states. To include null states in such a set up adds little intuition and greatly increases the level of attention to technical detail that needs to be paid. The requirement regarding $x^\star$ and $x_\star$, stems from the fact that they are necessarily realized in every state, and hence, uninformative.

A menu, A, acts as the frame induced by the relative probabilities of inclusion on each x-coin with $x \in A$ and exclusion for each y-coin with $y \notin A$. The fact that signals are independent, indicates that the inclusion or exclusion of a particular element carries the same informational content regardless of the composition of the menu. Of course, even though the informational value is the same, the *effect* of this information on beliefs is relative to the information provided by the other elements included (or excluded) from the menu. This behavior is captured by the following axiom.

[AX.2.6: IID] *For all $x \in X$, and $A, B \in \mathcal{K}(X)$, such that $x \notin A \cup B$, and states $s, s' \notin N_A \cup N_B$, if for some distributions $\pi^A, \rho^A, \in \Delta(A)$ and $\pi^B, \rho^B, \in \Delta(B)$: $(x_\star)_{-s}\pi^A \sim_A (x_\star)_{-s'}\rho^A$ and $(x_\star)_{-s}\pi^B \sim_B (x_\star)_{-s'}\rho^B$, then for all $\alpha \in (0,1)$,*

$$(x_\star)_{-s}\pi^A \succcurlyeq_{A\cup x} (x_\star)_{-s'}(\alpha\rho^A + (1-\alpha)x_\star) \iff (x_\star)_{-s}\pi^B \succcurlyeq_{B\cup x} (x_\star)_{-s'}(\alpha\rho^B + (1-\alpha)x_\star). \tag{2.4}$$

IID states that the proportional change in belief, in response to the inclusion of an element x, is the same across all menus. Without x, obtaining π^A in state s and ρ^A in state s' (and $x_\star$ everywhere else) are equally appealing, given menu A. When x is included, the beliefs change (i.e., the observation of the x-coin switches from exclude to include), and therefore, so do preferences. IID states that the same proportional change in preferences must occur, regardless of the initial menu. So if the change in preferences is such that, π^A in state s is now indifferent to $\alpha\rho^A + (1-\alpha)x_\star$ in state s' (and $x_\star$ everywhere else) given A, then the same α proportional shift preserves indifference when moving from B to $B \cup x$. This behavior, along with the general representation, exactly captures the updating procedure given by (2.3).

Theorem 2.6. *Let $\{\succsim_A\}_{A\in\mathcal{K}(X)}$ be represented by some* (MBR)*, with beliefs $\{\mu_A|A \in \mathcal{K}(X)\}$, all of which have full support. Then, there exists some (μ, l, X) that generates $\{\mu_A|A \in \mathcal{K}(X)\}$ as in* (2.3) *if and only if $\{\succsim_A\}_{A\in\mathcal{K}(X)}$ satisfies* IID*.*

Proof. In appendix A.2.

2.4 FROM EQUILIBRIUM TO MBR

This section briefly (and informally) describes how a MBR could arise as the natural consequence of a game between buyers and sellers (i.e., the observed behavior of the buyer in equilibrium satisfies the MBR axioms). Consider the environment where, first, a seller constructs a menu of goods to offer the buyer at posted prices, and then, the buyer decides whether or not to buy any of the offered goods. In other words, the sellers act as *stores*, who can curate their selections. Sellers are privately endowed with a type (read: the seller's quality or ability), and this type governs both the cost of *stocking* a particular good, and also, the utility a buyer derives from its consumption. In this environment, under standard single-crossing conditions, different types of sellers might differentiate themselves in equilibrium by offering different menus of goods. Hence, in such an equilibrium, the sellers beliefs regarding the type of seller, and hence the value of the offered goods, is dependent on the offered menu. Specifically, if the equilibrium is in pure strategies, this induces a MBR with partitional signal structure.

Example 2.1.C (Luce's diner, one last time). *There are three types of restaurants, high (h), medium (m), and low (l) quality. Each can offer any selection of chicken (c), steak (s), or frog legs (f). The cost for a particular restaurant to keep an item on the menu (train the chef, provide a wine pairing, keep fresh ingredients, etc), is given by the following matrix:*

$$c_h(c) = 1 \qquad c_h(s) = 2 \qquad c_h(c) = 6$$
$$c_m(c) = 1 \qquad c_m(s) = 4 \qquad c_m(c) = 10$$
$$c_l(c) = 2 \qquad c_l(s) = 6 \qquad c_l(c) = 10$$

A patron, given that the quality of the food is known, has preferences (in dollar terms) according to

$$u(h,c) = 8 \qquad u(h,s) = 16 \qquad u(h,f) = 8$$
$$u(m,c) = 8 \qquad u(m,s) = 6 \qquad u(m,f) = 4$$
$$u(l,c) = 4 \qquad u(l,s) = 1 \qquad u(l,f) = 0$$

Each type of restaurant can select any subset of the main courses (along with posted prices) to offer potential diners. Given the observed menu and the subsequently updated beliefs, a diner will select the course that maximizes her utility (her expected utility from consumption less the posted price). All diners can take an outside option with utility 0.

Assume, initially, the diner has a uniform prior over the different types of restaurants. Then the following is a Bayes-Nash equilibrium. The high type offers $\{\langle c{:}8\rangle, \langle s{:}16\rangle, \langle f{:}8\rangle\}$*, the medium type offers* $\{\langle c{:}8\rangle, \langle s{:}6\rangle\}$ *and the low type offers* $\{\langle c{:}4\rangle\}$*. As this is a separating equilibrium, after observing any of these menus, the diner places probability 1 on the corresponding type, and chooses* s*,* c*,* c*, respectively. When seeing any other menu, she places probability 1 on low, and takes the outside option.*

Notice that in this example, both the utilities for outcomes and the beliefs after the observation $\{c, s, f\}$ and $\{c, s\}$ map exactly to Katya's tastes and beliefs given the same observations. As such, the behavior of buyers in such a separating equilibrium would correspond exactly to the MBR with the partitional information structure described in Example 1.B.

2.5 LITERATURE REVIEW

Tversky and Kahneman (1981) developed the notion of framing –the idea that a decisions are influenced by their surrounding context. Framing has a large literature, both in the theoretical, experimental, and psychological settings Kahneman and Tversky (1984); Rubinstein and Salant (2008); Tversky and Shafir (1992). A particular type of framing concerns the consideration of menu, or currently available alternatives, referred to in the literature as context dependence. In contrast to this model, context dependence is often associated with particular psychological heuristics such as a basing choice on the difference between the attributes of outcomes or reluctance to choose extreme outcomes Simonson and Tversky (1992).

That a menu may contain information relevant to the DM's choice over the objects it contains was first articulated by Luce and Raiffa (1957) and expounded upon by Sen (1993, 1997). Sen describes the notion of the *epistemic value* of a menu with more tact than I could hope to achieve: "What is offered for choice can give us information about the underlying situation, and can thus influence our preference over the alternatives, as we see them. For example, the chooser may learn something about the person offering the choice on the basis of what he or she is offering." It is by paraphrasing/formalizing the vignettes in Luce and Raiffa (1957) and Sen (1997) that I constructed the examples that run throughout this paper.

There have been several decision theory papers which deal with characterizing framing effects that stem from informational sources. Ahn and Ergin (2010) considers a DM whose beliefs, and hence preferences, depend on the description of the state space. There a depiction of the state space is a partition of it, and preferences are defined over all acts measurable with respect to the partition. The interpretation is that different descriptions of the state space might alter the DM of contingencies which he would otherwise be unaware. Bourgeois-Gironde and Giraud (2009) construct a model of "rational" framing in the domain of Bolker–Jeffrey decision theory. They take as motivation, and provide an axiomatic foundation for, the observation of Sher and McKenzie (2006) that (seemingly) logically equivalent statements might in fact contain different information because the choice to use one description over another might itself impart information. As such, Bourgeois-Gironde and Giraud (2009) consider a set of frames and allow two different, but logically equivalent, statements that belong to different frames to induce different beliefs of the DM.

The epistemic aspect of decision problems has been studied by Kochov (2010) in a model that shares many philosophical motivations with this one. Kochov's model defines a decision problem as a collection of menus, and imposes the canonical axioms (i.e., Dekel *et al.* (2001)) on a preference relation over each decision problem to back out a problem-specific subjective state space. The primary mechanism by which epistemic content alters the decision makers preference in Kochov's model is by changing the composition of the subjective state-space (i.e., the difference in preference is mitigated through a change in tastes, rather than beliefs). The interpretation of menus revealing different unforeseen contingencies is problematic from the modelers point of view: it is impossible to observe a decision maker who is both aware and unaware of a particular contingency. This paper, on the other hand, explains the same behavior by confining the context effect to be a local one.

The appeal to a DM who holds multiple beliefs has been explored in the literature on state dependent preferences. Karni *et al.* (1983) propose a DM who ranks alternatives after some hypothetical event with a known probability. Like this paper, there is an imposed consistency in ex-post tastes across different decision problems. Also related, is Karni and Safra (2014), which takes a somewhat converse approach. There, the decision maker has beliefs regarding his state dependent preferences, or *states of mind* which induce a preference over menus, rather than the menu inducing the belief about the state space. As such, it is the DM's beliefs regarding a subjective state space (his *state of mind*) that is invariant across decision problems.

It is also worth noting that models of endogenous reference dependence can be interpreted as context dependence. In these models the decision problem is associated with a reference level of utility by which the DM evaluates each outcome Koszegi and Rabin (2006); Ok *et al.* (2015). As such, adding outcomes that will effect the reference point will thereby change the DM's preferences. These models can be thought of as a specific case of epistemic concerns; the reference point is information about some underlying state variable. A decision problem associated with reference point, r, is an indicator that the state-of-the-world is s_r.

Finally, this paper is related to the decision theoretic literature on identifying the conditions under which a decision maker is Bayesian updating with respect to subjective (and hence unseen) signals, for example, Lehrer and Teper (2015). In particular, the general model can be seen as a special case of the subjective signal structure discussed in Shmaya and Yariv (2015).

3.0 PLANS OF ACTIONS[12]

In which Ivan, in trying to make an investment, learns, simultaneously, that (i) AAPL *and* GOOG *move more or less in tandem, and, having only enough capital for a single share, that (ii) such information is of no value.*

3.1 INTRODUCTION

Exploration models capture a common trade off between an immediate payoff and new information, which can potentially impact future decisions and payoffs.[13] In such models an agent has to choose, every period, one project out of several in which to invest. By observing the outcome of an investment, the agent learns both about the chosen project and, in case the outcomes across different projects are correlated, about other projects as well. Each decision is predicated on the tradeoff between the immediate value of the investment and the future value of the information obtained by observing the outcome. Therefore, the agent's optimal investment strategy is a function of the history of observed outcomes, the projects that will be feasible in the future, and her beliefs over the true *joint* distribution of the outcomes of each project. While it is the correlation between projects that allows the agent to extrapolate her observations to future outcomes of the different projects, we show in this paper that *contemporaneous correlations* (i.e., the likelihood of an outcome of

[12]ROEE TEPER IS A COAUTHOR OF THIS CHAPTER.

[13]Exploration models were introduced by Robbins (1952) and have been extensively studied in the statistics literature (as bandit problems), and widely incorporated in economic models (as search problems, stopping problems, research and development, experimentation, portfolio design, etc). See Berry and Fristedt (1985) for an overview of classic results within the statistics literature. For a survey of economic applications see Bergemann and Välimäki (2008).

project a in a period given the outcome of project b in the *same* period) carry no economic content in such exploration problems. In other words, when solving an exploration problem, contemporaneous correlations can be ignored without changing the set of optimal strategies.

Consider, for example, a scenario in which Ivan has to invest each period in one of two projects. The outcome of each project depends on the state of the economy, which could be either high or low with equal probabilities, independently across periods. In case the state is high, both projects yield high outcomes, and if the state is low, low outcomes. In particular, the projects are (fully) correlated. Note, each project yields high and low outcomes with equal probabilities. Alternatively, consider another scenario in which the projects are independent (that is, the outcomes do not depend on the state of the economy) and both yield each outcome with equal probabilities. Since Ivan needs to commit to exactly one project in each period, his investment strategy will depend on the marginal probabilities of each project. Thus, Ivan's investment strategy will be the same in these two scenarios. Putting differently, only the process of marginal distributions affect the optimal investment strategy; we need not worry about the effects of contemporaneous correlations on the optimal strategy.

This example is stylized; the specified projects are fully correlated and there is no learning (that is, the agent knows exactly what is the underlying distribution governing the state of the economy). We show, however, that the example can be generalized to any exploration problem where the outcome generating distribution (the object which the agent is attempting to learn) is stationary.[14] In particular, any such problem can be represented as another exploration problem in which the Bayesian model dictates that projects are not correlated, without affecting the agent's preferences over strategies, and hence the optimal strategy.

While this result allows an agent to simplify her decision problem, it also has a down-side from the modeler's vantage. We show there is an inherent limitation on the type of beliefs that can be identified by observing an agent's preferences in such decision making environments; the general stochastic process governing beliefs can only be partially identified. Given an agent's preference over investment strategies in the two scenarios discussed

[14] We believe similar issues will arise for Markovian transitions. This seems to require different tools and we expect the characterizations to be different than in the stationary case.

above, there is no behavior that would identify which of the two scenarios the agent has in mind. Fortunately, even under the partial identification, we show there are meaningful behavioral restrictions allowing the modeler to test whether the agent is acting according to *some* Bayesian model.

Because in exploration environments the agent can chose only one project in each period, her preferences over the different strategies depend only on the *margins* of her beliefs. And vice versa, the agent can only reveal—through choice or preference over investment strategies—her history dependent beliefs over each project separately. To show this, we provide a decision theoretic model, where we introduce a new dynamic and recursive framework capturing the exploration-exploitation tradeoffs faced by a decision maker (henceforth, DM). Our primitive is a preference relation over the different strategies that can be implemented by an agent facing a bandit problem. In this framework we first provide the behavioral (axiomatic) restrictions of subjective discounted expected utility maximization. The principal observation arising from our result is that the representation pins down only the processes of marginal beliefs of the different projects separately.

To better understand the economic relevance of this identification, we proceed with an analysis of a statistical framework. Here we consider stochastic processes that are determined by observing the outcome of a single project in each period, where potentially different projects are chosen across periods. We refer to these processes as *observable*, in light of the fact that they are precisely the output of our decision theoretic result. If, to the contrary, we had been able to observe the process over the joint realizations of all experiments, then the classic *exchangeability* property (or symmetry, as referred to in the decision theoretic terminology) would characterize Bayesianism. Given the limits of what can been observed, we cannot resort directly to such classical results. We provide instead a necessary and sufficient condition, *Across-Action Symmetry* (AA-SYM), for the observable processes to be consistent with an exchangeable process over the collection of all projects. It is not surprising, we obtain only a partial identification; a consistent exchangeable process, when it exists, need not be unique. We show however, that whenever our condition is met, there exists an exchangeable process, consistent with the observables, wherein there is no (contemporaneous) correlation across the different actions. Moreover, such a consistent exchangeable process *is* unique.

Finally, we show AA-SYM can be written in terms of the decision theoretic primitive. Combining these results, we conclude that the DM's subjective joint distribution is not fully identified. Put differently, contemporaneous correlations across actions do not affect preferences and optimal strategies in bandit problems. Nonetheless, capitalizing on the sufficiency of AA-SYM for the marginals to be consistent with an underlying exchangeable model, we obtain an axiomatization for classic Bayesianism in *any* bandit problems.

In this framework a DM is tasked with ranking sequential and contingent choice objects: the action taken by the agent at any stage depends on the outcomes of previous actions. Formally, our primitive is a preference over *plans of action* (*PoAs*). Each action, a, is associated with a set of consumption prizes the action might yield, S_a. Then, a PoA is recursively defined as a *lottery over pairs* (a, f), where a is an action and f is a mapping that specifies the continuation PoA for each possible outcome in S_a. Theorem 3.2 shows that the construction of PoAs is well defined. So, a PoA specifies an action to be taken each period that can depend on the outcome of all previously taken actions.

The actions in our model is in direct analogy to the arms of bandit problem (or actions in a repeated game). PoAs correspond to the set of all (possibly mixed) strategies in these environments. Note, however, the DM's perception of which outcome in S_a will result form taking action a is not specified. This is subjective and should be identified from the DM's preferences over PoAs. As discussed above, the main question is to what extent these beliefs can be identified and what are the economic implications of belief identification in this framework?

Theorem 3.3 axiomatizes preferences over PoAs of a DM who at each history entertains a belief regarding the outcome of future actions. That is, at each history h and for every action a, the DM entertains a belief $\mu_{h,a}$ over the possible outcomes S_a; $\mu_{h,a}(x)$ is the DM's subjective probability that action a will yield outcome x, contingent on having observed the history h. Given this family of beliefs, the DM acts as a subjective discounted expected utility maximizer, valuing a PoA p, after observing h, according to a *Subjective Expected Experimentation* (*SEE*) representation:

$$U_h(p) = \mathbb{E}_p \left[\mathbb{E}_{\mu_{h,a}} \left[u(x) + \delta U_{h'}(f(x)) \right] \right], \tag{SEE}$$

where h' is the updated history (following h) when action a is taken and x is realized. All the parameters of the model –the consumption utility over outcomes, u, the discount factor, δ, and the history dependent subjective beliefs, $\{\mu_{h,a}\}_{h\in\mathcal{H},a\in\mathcal{A}}$– are identified uniquely.

Our setup requires a formulation of a novel axiom termed *proportionality* (PRP): at any given history, the manner in which the DM evaluates continuation problems must be proportional to the manner in which she evaluates the consumption utility. Indeed, in order to ensure that the DM is acting consistently with a family of beliefs it must be that she assesses the value of each action according to the expectation of the consumption utility and discounted continuation utility it induces. Furthermore, it is necessary that the probabilistic weight she places on a given consumption utilities is the same as the weight she places on the corresponding continuation value.

Theorem 3.3 shows that PRP, along with (some of the) standard behavioral conditions for discounted expected utility, is necessary and sufficient for an SEE representation. While the axiomatization does not point to the optimal strategy in general strategic experimentation problems, which is known to be a hard problem to solve when actions are correlated, it provides (like most axiomatization theorems) a unifying guidance as to what might or might not be ruled out.

The identification result accompanying the representation concerns the marginal beliefs, $\{\mu_{h,a}\}_{h\in\mathcal{H},a\in\mathcal{A}}$, and not a stochastic process over all actions, as is the starting point in the standard approach to bandit problems. To be clear, this is not a limitation of the current setup any more than of bandit problems in general: observing a single pull of an arm each period simply does not provide sufficient data to identify the joint distribution. As such, we would like to know when the family of identified beliefs is consistent with an underlying exchangeable process. Further, given consistency, what are the limits of identification regarding this exchangeable model. To answer these questions we turn to the statistical model.

The Statistical Model. As above, there is a set of actions, $\mathcal{A}$, each element of which, a, is associated with the outcome space S_a. We are considering a family of processes over the outcomes of the different actions—where each period one and only one action is observed.

Let $\mathbf{T} = (T_1, T_2, ...)$, where $T_i \in \{S_a\}_{a \in \mathcal{A}}$ for every i. Let $\mathcal{T}$ denote the set of all such sequences. For any $\mathbf{T}$ in $\mathcal{T}$, let $\zeta_{\mathbf{T}}$ be a distribution over $\mathbf{T} = (T_1, T_2, ...)$. We refer to these distributions as our observables; and, denoting $S = \prod_{a \in \mathcal{A}} S_a$, we assume a distribution, ζ, over $S^{\mathbb{N}}$ is not observable.

We are motivated by the decision theoretic identification of $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$. While a process ζ over $S^{\mathbb{N}}$ specifies each $\mu_{h,a}$ (as the ζ-probability that following history h, action a will yield outcome x), it conveys strictly more information. For example, the ζ-probability that action a yields outcome x at the same time that action b yields outcome y has no counterpart in the identified family of marginals. The family of processes $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$, on the other hand, contains exactly the same information as $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$. In both models, we do not have direct access to the probability of joint realizations of different actions. We only have access to the marginal distributions. Therefore, the exercise at hand concerns a direct translation of the decision theoretic observables into the statistical language.

In this framework, we introduce a condition referred to as *Across-Action Symmetry (AA-SYM)* and Theorem 3.5 shows that it is necessary and sufficient for the observables to be consistent with an exchangeable process over the joint realizations of all actions in every period (that is, $S^{\mathbb{N}}$). Informally, AA-SYM states that the probability of obtaining outcome x when taking action a followed by outcome y when taking action b, is the same as the probability of obtaining outcome y when taking action b followed by outcome x when taking action a. This is reminiscent of the symmetry (exchangeability) property, but note, in each period the outcome space may change as different actions can be taken.

The inherent observability constraint in this framework bears a cost; the exchangeable process with which our observables are consistent is typically not unique. Bearing in mind this generic non-uniqueness, we introduce what we term *strongly exchangeable* processes –a subclass of the widely studied exchangeable processes. We elaborate. Assume there is an underlying distribution governing the joint realization of actions that is *inter-temporally* i.i.d. This distribution is not known exactly, but there exists a prior probability over what it might be. The prior is updated every period upon the observation of the realization of actions. Due to de Finetti (1931); Hewitt and Savage (1955), these classical Bayesian updating processes are referred to as exchangeable. In a strongly exchangeable process, where the periodic state-space takes a product structure, the set of possible underlying distributions are such that outcomes across actions are independent. Thus, a strongly exchangeable process is one in which, conditional on the distributional parameter, outcomes are both inter-temporally and *contemporaneously* independent.

Despite strong exchangeability having more structure than classic exchangeability, it imparts no additional restrictions in our statistical model. Theorem 3.7 shows that a family of observables satisfying AA-SYM is consistent with a strongly exchangeable process, and this process is unique. We conclude, strong exchangeability is the full characterization of Bayesianism in our statistical framework and the lack of contemporaneous correlations carry no constraints beyond AA-SYM.

Finally, returning to our decision theoretic model, we show that AA-SYM can be represented as an axiom over the primitives. Proposition 3.8 states that the additional axiom is both necessary and sufficient for the beliefs of an SEE representation to follow an exchangeable processes. This, of course, implies that two decision makers whose beliefs (as exchangeable processes) induce the same family of observable processes, will have the same preferences over strategies in *any* Bandit problem. In other words, Proposition 3.8 implies that contemporaneous correlations across actions do not impose any additional restrictions beyond classic Bayesianism when analyzing bandit problems, and no behavior can identify them in such an environment.

Organization. The paper is broadly broken into the two halves outlined above; Section 3.2 contains the decision theoretic framework, and Sections 3.3 and 3.4 the statistical one. Within Section 3.2, we first formally introduce the environment and the construction of plans of action (Sections 3.2.1). Next, we provide the axioms and representation result for an SEE structure (Sections 3.2.2 and 3.2.3). Section 3.3 introduces the observable processes that represent SEE belief structures. Here, we provide a statistical condition on observable processes, AA-SYM, so that the SEE belief structure is consistent with an exchangeable process. Section 3.4 introduces the notion of strong exchangeability and presents our (non) uniqueness result. The translation of AA-SYM back into decision theoretic terms is presented in Section 3.4.1. Section 3.5.1 discusses the related literature. An informal discussion regarding how a decision theoretic model would incorporate exogenous information appears in Section 3.5.2. Lastly, Section 3.5.3 discusses the point of disagreement among Bayesians in environments of experimentation. All the proofs are in the Appendices.

3.2 THE DECISION THEORETIC FRAMEWORK

3.2.1 Choice Objects

The purpose of the current section is to construct the different choice objects, termed *plans of action* (*PoAs*). The primitive of our model, as presented in the subsequent section, is a preference relation over all PoAs.

Let X be a finite set of outcomes, endowed with a metric d_X. Outcomes are consumption prizes. For any metric space, M, let $\mathcal{K}(M)$ denote the set non-empty compact subsets of M, endowed with the Hausdorff metric. Likewise, for any metric space M, denote $\Delta^{\mathcal{B}}(M)$ as the set of Borel probability distributions over M, endowed with the weak*-topology, and $\Delta(M)$ the subset of distributions with denumerable support.

Let $\mathcal{A}$ be a compact and metrizable set of actions. Each action, a, is associated with a set of outcomes, $S_a \in \mathcal{K}(X)$, which is called the support of the action. We assume the map $a \mapsto S_a$ is continuous and surjective. For any metric space M, let $\mathcal{A} \otimes M = \{(a, f) | a \in \mathcal{A}, f\colon S_a \to M\} = \{(a, \{(x_i, m_i)\}_{i \in I}) \in \mathcal{A} \times \mathcal{K}(X \times M) | \bigcup_{i \in I}\{x_i\} = S_a \text{ and } x_i \neq x_j, \forall i \neq j \in I)\}$, endowed with the subspace topology inherited from the product topology. By the continuity of $a \mapsto S_a$ we know that the relevant subspace is closed and hence the topology on $\mathcal{A} \otimes M$ is compact whenever M is. We can think of f as the assignment into M for each outcome in the support of action a. For any $f : X \to M$ we will abuse notation and write (a, f) rather than $(a, f|_{S_a})$.

With these definitions we can define PoAs. A PoA is a tree of actions, such that each period the DM receives a lottery (with denumerable support) over actions conditional on the outcomes for each of the previous actions.

We begin by constructing slightly more general objects. Set $R_0 = \Delta^{\mathcal{B}}(\mathcal{A})$; a 0-period plan is a lottery over actions. Given an action, an element of its support is realized and the plan is over. Then a 1-period plan, r_1, is a lottery over actions, and continuation mappings into 0-period plans: $r_1 \in \Delta^{\mathcal{B}}(\mathcal{A} \otimes P_0)$. Given the realization of an action-continuation pair, (a, f), in the support of r_1, and the realized element of the support, $x \in S_a$, the DM receives a 0-period plan, as given by $f(x)$. Continuing in this fashion, we can define recursively,

$$R_n = \Delta^{\mathcal{B}}(\mathcal{A} \otimes R_{n-1}).$$

Define $R^* = \prod_{n \geq 0} R_n$. R^* is the set of all PoAs (including inconsistent plans and plans whose support is arbitrary). We first restrict ourselves to the set of *consistent* elements of R^*: those elements such that, the $(n-1)$-period plan implied by the n-period plan is the same as the $(n-1)$-period plan. To see why this is important, consider an element, r_n, of R_n. The plan r_n specifies an action to be taken in period 0 and, conditional on the outcome, the plan r_{n-1} which itself specifies the action to be taken in the next period and the continuation plan r_{n-2} for the next, etc. If we stop this process at any period $m < n$, ignoring whatever continuation plans are assigned, the output is an m period plan. Hence, each n period plan specifies a (unique) m period plan for each $m < n$. Moreover, an element $r \in R^*$ specifies an n period plan for each $n \in \mathbb{N}$. Intuitively, we would like to view each r as an infinite plan, by considering the sequence of arbitrarily large, and expanding, finite plans. Consistency is the requirement that makes this work, that for $r_n = \text{proj}_n r$, the first $m < n$ periods specify exactly $r_m = \text{proj}_m r$. Let R denote the set of all consistent plans.[15]

Proposition 3.1. *There exists a homeomorphism,* $\lambda : R \to \Delta^{\mathcal{B}}(\mathcal{A} \otimes R))$ *such that*

$$\text{marg}_{\mathcal{A} \times \mathcal{K}(X \times R_{n-1})}(\lambda(r)) = \text{proj}_n r. \tag{3.1}$$

Proof. In appendix B.1.

[15]Precisely specifying the $m < n$ period plan implied by r_n requires a more cumbersome notation than we wish to introduce in the text; for the formal definitions see Appendix B.1.

Next we want to consider plans whose support is denumerable. It is easy enough to set $P_0 = \Delta(\mathcal{A}) \subset R_0$, and define recursively $P_n = \Delta(\mathcal{A} \otimes P_{n-1}) \subset R_n$. Of course, there is a potential pitfall still lurking: for a given $\prod_{n\geq 0} P_n$, although each p_n is a denumerable lottery, the associated element, $\lambda(p)$ might live in $\Delta^B(\mathcal{A} \otimes P)$ rather than $\Delta(\mathcal{A} \otimes P)$. Indeed, we need also to restrict our attention to the set of plans that have countable support not just for each finite level, but also "in the limit," and whose implied continuation plans are also well behaved in such a manner. Fortunately, this can be done:[16]

Theorem 3.2. *There exists maximal set $P \subset R$ such that for each $p \in P$, $\text{proj}_n p \in P_n$, and λ is a homeomorphism between P and $\Delta(\mathcal{A} \otimes P)$.*

Proof. In appendix B.1.

The set P is our primitive.[17] As a final notational comment, we would like to consider a further specification of *objective* plans, denoted by $\Sigma \subset P$. Σ denotes the set of plans which contain no subjective uncertainty; in every period, every possible action yields some outcome with certainty. Recall, for each $x \in X$ there is an associated action, a_x such that $S_{a_x} = \{x\}$. Associate this set of actions with X. Then $\Sigma_0 = \Delta(X)$ and, recursively, $\Sigma_n = \Delta(X \times \Sigma_{n-1})$. Finally $\Sigma = P \bigcap \prod_{n\geq 0} \Sigma_n$. That is, these plans specify only actions with deterministic outcomes at every stage. It is straightforward to show λ takes Σ to $\Delta(X \times \Sigma)$.

PoAs are infinite trees; each node, therefore, is itself the root of a new PoA—a distribution over action-continuation pairs. Each action-continuation, (a, f), in the support of a node contains branches to new nodes (PoAs). The branches emanating from an action coincide with the outcomes in the support of that action, $x \in S_a$. The node that follows x is the PoA specified by $f(x)$. Each node, therefore, is reached after a unique history: the history specifies

[16] One can also consider measurable lotteries (instead of lotteries with countable support). In fact, the construction of the homeomorphism in Appendix B.1 considers measurable lotteries. In the paper we focus on discrete support for notational cleanliness (see footnote 18) and tractability (to avoid measurability issues in proofs). We justify our focus by noting that $\Delta(\mathcal{A} \otimes P)$ is dense in $\Delta^{\mathcal{B}}(\mathcal{A} \otimes R)$ and so, given continuity (Axiom vNM), preferences over the more general objects are recoverable.

[17] One might consider an alternative framework of "adapted processes" of Anscombe-Aumann acts (see, for example, Epstein and Schneider (2003)), modified to our multi-action environment. In such a setup there is a distinction between exogenous states and outcomes (of the different actions). However, in a classical exploration problem, an outcome of an action is *simultaneously* an object from which the agent derives utility and from which the agent learns regarding the uncertainty underlying the (different) action(s). Similar results to those presented here would be obtained had we adopted the framework of adapted processes, but it seems conceptually appropriate to resort to plans of actions.

the realization of the distribution of each pervious node, and outcome of the action realized. Thus, for a given PoA, p, each history of length n is an element of $\prod_{t=1}^{n} P \times [\mathcal{A} \otimes P] \times X$ such that $p^1 = p$ and

$$(a^t, f^t) \in supp(p^t)$$

$$x^t \in S_{a^t}$$

$$p^{t+1} = f^t(x^t)$$

Define the set of all histories of length n for p as $\mathcal{H}(p, n)$ and the set of all finite histories as $\mathcal{H}(p)$. Let $\mathcal{H}(n) = \bigcup_{p \in P} \mathcal{H}(p, n)$ and, $\mathcal{H} = \bigcup_{n \in \mathbb{N}} \mathcal{H}(n)$. For each $h \in \mathcal{H}(p, n)$, h corresponds to the node (PoA) defined by $f^n(x^n)$. Lastly, for any $p, q \in P$ and $h \in \mathcal{H}(p)$ define $p_{-h}q$ as the (unique!) element of P that coincides with p everywhere except after h in which case $f^n(x^n)$ is replaced by q. Note that the n period plan implied p and $p_{-h}q$ are the same. For any $p, q \in P$ and $n \in \mathbb{N}$, let $p_{-n}q \equiv \bigcup_{h \in \mathcal{H}(p,n)} p_{-h}q$.

For any $h = (p^1, a^1, f^1, x^1 \ldots p^n, a^n, f^n, x^n)$ and $\hat{h} = (\hat{p}^1, \hat{a}^1, \hat{f}^1, \hat{x}^1 \ldots \hat{p}^n, \hat{a}^n, \hat{f}^n, \hat{x}^n)$ both in $\mathcal{H}(n)$, we say that h and h' are $\mathcal{A}$-*equivalent*, denoted by $h \overset{\mathcal{A}}{\sim} h'$ if $a^i = \hat{a}^i$ and $x^i = \hat{x}^i$ for $i \leq n$. That is, two histories of length n are $\mathcal{A}$-equivalent, whenever they correspond to the same sequence of action-realization pairs, ignoring the objective randomization stage of each period and the continuation assignment to outcomes that did not occur. It will turn out, we are only interested in the $\mathcal{A}$-equivalence classes of histories. Technically, this is the consequence of the linearity of preference and indifference to the resolution of uncertainty (as shown in Lemma 7); conceptually, this is because all uncertainty in the model regards the realization of actions, and so, observing objective lotteries has no informational benefit.

3.2.2 The Axioms

The primitive in our model is a preference relation $\succcurlyeq \subseteq P \times P$ over all PoAs. When specific PoA and history are fixed, the preferences induce history dependent preferences as follows: for any $p \in P$, and $h \in \mathcal{H}(p)$ define $\succcurlyeq_h \subseteq P \times P$ by

$$q \succcurlyeq_h r \iff p_{-h}q \succcurlyeq p_{-h}r.$$

The following axioms will be employed over all history induced preferences.[18] A history is *null* if $\succcurlyeq_h$ is a trivial relation. This first four axioms are variants on the standard fare for discounted expected utility. They guarantee the expected utility structure, non-triviality, stationarity and separability (regarding objects over which learning cannot take place), respectively.

[AX.3.1: vNM] *The binary relation, $\succcurlyeq_h$ satisfies the expected utility axioms. That is: weak order, continuity (defined over the relevant topology, see Appendix B.1) and independence.*

We require a stronger non-triviality condition that is standard, because of the subjective nature of the dynamic problem. We need to ensure the DM believes *some* outcome will obtain. Therefore, not all histories following a given action can be null.

[AX.3.2: NT] *For any non-null h, and any (a, f), not all $h' \in h \times \mathcal{H}((a, f), n)$ are null.*

Of course, the nature of the problem at hand precludes stationarity and separability in full generality. Since the objective is to let the DM's beliefs depend on prior outcomes explicitly, her preferences will as well. However, the DM's beliefs do not influence her assessment of objective plans (i.e., elements of Σ), and so it is over this domain that stationarity and separability are retained. This means, the DM's preferences *in utility terms* are stationary and separable, but we still allow the conversion between actions and utils to depend on her beliefs which change responsively.

[AX.3.3: SST] *For all non-null $h \in \mathcal{H}$, and $\sigma, \sigma' \in \Sigma$,*

$$\sigma \succcurlyeq \sigma' \iff \sigma \succcurlyeq_h \sigma'.$$

[AX.3.4: SEP] *For all $x, x' \in X, \rho, \rho' \in \Sigma$ and $h \in \mathcal{H}$,*

$$\big(\frac{1}{2}(x, \rho) + \frac{1}{2}(x', \rho')\big) \sim_h \big(\frac{1}{2}(x, \rho') + \frac{1}{2}(x', \rho)\big).$$

[18] It is via the use of this construction that our appeal to denumerably supported lotteries provides tractability. If we were to employ lotteries with uncountable support, then histories would, in general, be zero probability events; under the expected utility hypothesis, $\succcurlyeq_h$ would be null for all $h \in \mathcal{H}$. This could be remedied by appealing to histories as *events* in $\mathcal{H}$, measurable with respect to the filtration induced by previous resolutions of lottery-action-outcome tuples. We believe that this imposes a unnecessary notational burden.

Because of the two-stage nature of the resolution of uncertainty each period (first, the resolution of lottery over $\mathcal{A} \otimes P$, and then the resolution of the action over X), we need an additional separability constraint. From the point of view of period n, and when considering the continuation problem beginning in period $n+1$, the DM should not care if uncertainty is resolved in period n (when the action-continuation pair is realized), or in period $n+1$. That is, we also assume the DM is indifferent to the timing of objective lotteries *given a fixed action.*

[AX.3.5: IT] *For all* $a \in \mathcal{A}$, $h \in \mathcal{H}$, $\alpha \in (0,1)$, *and* $(a,f),(a,g) \in \hat{P}$,

$$\alpha(a,f) + (1-\alpha)(a,g) \sim_h (a, \alpha f + (1-\alpha)g),$$

where mixtures of f *and* g *are taken point-wise.*

Thus far the axioms introduced are somewhat standard. However, in our particular framework these assumptions do not guarantee that the value of the action is in any way related with its realization of consumption alternatives. This is because, unlike other environments, the set of outcomes, X, plays a dual role in exploration models: representing both the space of outcomes and the state space regarding future actions.

The realization of an outcome x delivers utility according to both of these roles, and, to ensure consistency between them requires two steps. First, construct a subjective distribution over each action by treating X as a state space. This will be done by looking at the ranking of continuation mappings for each action (i.e., (a,f) compared to (a,g)). Interpreting X as the periodic state space, these continuation mappings are analogous to "acts" in the standard subjective expected utility paradigm–and so, standard techniques allow for the identification of such a subjective belief. Second, we need to ensure that the value assigned to arbitrary PoAs is the expectation according to these beliefs. Towards this, the following notation is introduced.

Definition. *For any function* $f : X \to \mathcal{P}$, *define* $p.f \in P$ *as* $p.f[(a,g)] = p[\{(b,h) | b = a\}]$ *if* $g = f$, *and* $p.f[(a,g)] = 0$ *if* $g \neq f$.

Take note, because we are dealing with distributions of denumerable support, we have no measurability concerns. The plan of action $p.f$ has the same distribution over actions in the first period, but the continuation plan is unambiguously assigned by f, as shown in Figures **??** and **??**. If the original plan is in $\mathcal{A} \otimes P$, then the dot operation is simply a switch of the continuation mapping: $(a, g).f = (a, f)$. This operation is introduced because it allows us to isolate the subjective distribution of the first period's action.

Definition. *$p, q \in P$ are **h-proportional** if for all $f, g : X \to \Sigma$.*

$$p.f \succcurlyeq_h p.g \iff q.f \succcurlyeq_h q.g$$

Since the images of f and g are in Σ, there is no informational effect from observing the outcome of p. Hence, f and g can be thought of as objective assignments into continuation utilities. The ranking '$p.f \succcurlyeq p.g$' is really a ranking over f and g as functions from $X \to \mathbb{R}$. Thus, h-proportionality states that the DM's subjective uncertainty regarding X is the same when faced with p or with q.[19]

[AX.3.6: PRP] *For all $p, q \in P$, and $f : X \to \Sigma$ if p and q are h-proportional then $p.f \sim_h q.f$.*

The outcomes of an action represent not only the uncertainty regarding continuation, but also the utility outcome for the current period. So, when p and q are h-proportional, and thus induce the same uncertainty regarding X, the DM's uncertainty about her current period utility is the same across the plans. Therefore, if we replace the continuation problems with objectively equivalent plans, the DM should be indifferent between p and q.

3.2.3 A Representation Result and Belief Elicitation

The following is our general axiomatization result. It states that the properties above characterize a DM who, when facing a PoA, calculates the subjective expected utility according to a collection of history dependent beliefs over action-outcome pairs, and among different PoAs contemplates the benefits of consumption versus learning.

[19]To see this, note that the relation R on $\mathbb{R}^X \times \mathbb{R}^X$ defined by fRg if and only if $p.f \succcurlyeq p.g$ is a preference relation over acts that satisfies the Anscombe and Aumann (1963) axioms, and therefore encodes the DM's subjective likelihood of each $E \subset X$. From a functional standpoint, h-proportionality states the subjective distribution over X induced by p is the same as that induced by q.

Theorem 3.3 (Subjective Expected Experimentation Representation). *$\succsim_h$ satisfies* vNM*,* NT*,* SST*,* SEP*,* IT *and* PRP *if and only if there exists a utility index $u : X \to \mathbb{R}$, a discount factor $\delta \in (0, 1)$, and a family of beliefs $\{\mu_{h,a} \in \Delta(S_A)\}_{h\in\mathcal{H},a\in\mathcal{A}}$ such that*

$$U_h(p) = \mathbb{E}_p\left[\mathbb{E}_{\mu_{h,a}}\left[u(x) + \delta U_{h'(a,x)}(f(x))\right]\right], \quad \text{(SEE)}$$

jointly represents $\{\succsim_h\}_{h\in\mathcal{H}}$, where $h'(a, x) = (h, p, (a, f), x)$. Moreover, u is cardinally unique, δ is unique, the family of beliefs is unique, and $\mu_{h,a} = \mu_{h',a}$ whenever $h \overset{\mathcal{A}}{\sim} h'$.

Proof. In Appendix B.3.

The theorem states that we can (uniquely) elicit the beliefs, following every history, over the outcomes of each action separately. We will henceforth refer to such beliefs as an *SEE belief structure*. The axioms do not impose any restrictions on the dynamics of such beliefs. More importantly, the theorem shows that, when ranking the different strategies in a bandit problem, the decision maker does not reveal her beliefs over the *joint* realizations of the different actions.

3.3 THE STATISTICAL FRAMEWORK

In order for a modeler to understand the DM's updating process, and whether it follows Bayes rule, we need to construct her beliefs regarding not only each action individually but also her beliefs regarding the correlation between actions. As we will see, in the generic case we have insufficient data to uniquely identify a (subjective) joint distribution. We will still, however, be able to identify a representative with unique properties.

Consider the family $\mathcal{T}$ of all sequences of individual experiments (alternatively, individual actions), where different experiments can be taken in the different periods. Let $\mathbf{T} = (T_1, T_2, ...)$ where $T_i \in \{S_a : a \in \mathcal{A}\}$ for every $i \geq 1$; so, each T_i corresponds to taking an action, say a, and expecting one of its outcomes, S_a. (Like before S_a corresponds to the set of possible outcomes.) $\mathcal{T}$ is then the collection of all such $\mathbf{T}$'s. For each

$\mathbf{T} = (T_1, T_2, ...)$ let $\zeta_{\mathbf{T}} \in \Delta^{\mathcal{B}}(\mathbf{T})$ be a process over $\mathbf{T}$; a distribution over all possible outcomes when taking action T_1, followed by T_2, followed by T_3, etc. For a given history of outcomes $h \in (T_1, T_2, ..., T_n)$, we denote $h \in \mathbf{T}$ if $\mathbf{T} = (T_1, T_2, ..., T_n, T_{n+1}, ...)$. Lastly, for a sequence of experiments $\mathbf{T} = (T_1, ..., T_n, T_{n+1}, ...)$ and a permutation $\pi : n \to n$, denote $\pi\mathbf{T} = (T_{\pi(1)}, ..., T_{\pi(n)}, T_{n+1}, ...)$.

We consider a family of processes, $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$, indexed by the possible sequences of experiments, $\mathcal{T}$. For a given family, we require that for sequences of experiments $\mathbf{T}, \mathbf{T}' \in \mathcal{T}$, if there is some history, $h \in \mathbf{T} \cap \mathbf{T}'$, then $\zeta_{\mathbf{T}}(h) = \zeta_{\mathbf{T}'}(h)$. This condition imposes that the probability of outcomes today do not depend on which experiments are to be conducted in the future. The set of all families of processes that meet this condition is in bijection to the set of all SEE belief structures,[20] which is exactly the output of Theorem 3.3. Thus, in case there is no confusion, we will refer to a family of processes $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ as an SEE belief structure as well.

Let $\mathcal{S}_{\mathcal{A}} \equiv \prod_{a\in\mathcal{A}} S_a$, and $\mathcal{S} \equiv \prod_{n\geq 0} \mathcal{S}_{\mathcal{A}}$. $\mathcal{S}$ represents the grand state-space; a state, s, determines the realization of each action in each period –an entity which is unobservable to the modeler.

Definition. *An SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ is **consistent** with $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$ if $\zeta|_{\mathbf{T}} = \zeta_{\mathbf{T}}$ for every $\mathbf{T} \in \mathcal{T}$.*

That is, $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ is consistent with some process ζ over $\mathcal{S}$ if for every sequence of experiments $\mathbf{T}$, the marginal of ζ to $\mathbf{T}$ coincides with $\zeta_{\mathbf{T}}$. In such a case the processes ζ, which we cannot observe, explains all our data.

[20]Indeed, fix a family of history dependent beliefs $\{\mu_{h,a}\}_{h\in\mathcal{H},a\in\mathcal{A}}$ and consider a sequence $\mathbf{T} = (S_{a_1}, S_{a_2}, ...)$. Let $h = (x_1, x_2, ..., x_n) \in \mathbf{T}$ and define

$$\zeta_{\mathbf{T}}(h) = \mu_{\emptyset,a_1}(x_1) \cdot \mu_{(a_1,x_1),a_2}(x_2) \cdots \mu_{(a_1,x_1,...,a_{n-1},x_{n-1}),a_n}(x_n).$$

Then $\zeta_{\mathbf{T}}$ is defined as the unique (continuous) processes over $\mathbf{T}$ that is consistent with $\zeta_{\mathbf{T}}(h)$ for every history $h \in \mathbf{T}$. This procedure can be inverted: Fix, $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ and some history $h = (a_1, x_1) \dots (a_n, x_n)$. Let $\mathbf{T}$ be any sequence such that $T_i = S_{a_i}$ for $i \leq n$ and $T_{n+1} = S_a$. Then define

$$\mu_{h,a}(x) = \zeta_{\mathbf{T}}(T_{n+1} = x | T_1 = x_1 \dots T_n = x_n).$$

From here it is clear why we need to impose the condition that the probability of an event is not affected by future experiments: this condition arises naturally in the conditions of the SEE representation (see, Theorem 3.3). Recall that $\mu_{h,a} = \mu_{h',a}$ whenever $h \overset{\mathcal{A}}{\sim} h'$. Further, it is easy to check the above maps are continuous so that the bijection is in fact a homeomorphism.

Because it forms the basis subjective Bayesianism and for the statistical literature on bandit problems, we will pay particular attention to the class of *exchangeable* processes.

Definition. *Let Ω be a probability space and $\hat{\Omega} = \prod_{n\geq 1}\Omega$. The process $\zeta \in \Delta^{\mathcal{B}}(\hat{\Omega})$ is* ***exchangeable****, if there exists probability measure θ over $\Delta^{\mathcal{B}}(\mathcal{S}_{\mathcal{A}})$, such that*

$$\zeta(E) = \int_{\Delta^{\mathcal{B}}(\Omega)} \hat{D}(E)d\theta(D), \tag{3.2}$$

where for any $D \in \Delta^{\mathcal{B}}(\Omega)$, $\hat{D}$ is the corresponding product measure over $\hat{\Omega}$.

Remark 3.4. *If ζ is exchangeable, then θ is unique.*

Exchangeable processes were first characterized by de Finetti (1931, 1937) and later extended by Hewitt and Savage (1955). Their fundamental result states that a process $\zeta \in \Delta^{\mathcal{B}}(\hat{\Omega})$ is exchangeable if and only if for any finite permutation $\pi : \mathbb{N} \to \mathbb{N}$ and event $E = \prod_{n\in\mathbb{N}} E_n$, we have

$$\zeta(E) = \zeta(\prod_{n\in\mathbb{N}} E_{\pi(n)}). \tag{3.3}$$

Exchangeable processes are of clear statistical importance, in particular within the subjectivist paradigm (see, for example Schervish (2012)). From the economic vantage, a DM who understands there to be an exchangeable process governing the outcome of actions would be considered Bayesian.[21] This is because, given the representation in Eq. 3.2, the DM (acts as if she) entertains a second order distribution, which she updates following every observation.

We would like to understand under what circumstances an SEE belief structure is a result of Bayesian updating. If we could infer from preferences the beliefs over the joint realizations of all actions, that is $\prod_{a\in\mathcal{A}} S_a$, then our questions would boil down to verifying whether this process satisfies exchangeability. However, we can only infer the beliefs over each action separately, and thus, our task remains. We need to find a condition on the family of $\zeta_{\mathbf{T}}$'s that determines whether it follows Bayes rule.

[21] It is possible to consider more general Bayesian models than exchangeable processes. At least for the case of independent actions, for example, it is not hard to adapt a local consistency axiom as in Lehrer and Teper (2015) that will imply that beliefs follow a general martingale process.

Definition. *An SEE belief structure* $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ *is* **Across-Arm Symmetric (AA-SYM)** *if*

$$\zeta_{\mathbf{T}}(h) = \zeta_{\pi\mathbf{T}}(\pi h)$$

for every $\mathbf{T} \in \mathcal{T}$, $h \in \mathbf{T}$ *and a permutation* $\pi : n \to n$.

Intuitively, AA-SYM requires that if we consider a different order of experiments, then the probability of outcomes (in the appropriate order) does not change.

The next theorem states that across-arm symmetry is a necessary and sufficient condition for an SEE belief structure to be consistent with Bayesian updating of some belief over the *joint* realizations of all actions.

Theorem 3.5. *An SEE belief structure* $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ *satisfies AA-SYM if and only if it is consistent with an exchangeable process* $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$.

Theorem 3.5 is stated without proof. Necessity is trivial and sufficiency will be a straightforward application of Theorem 3.7.

3.4 STRONG EXCHANGEABILITY AND CONTEMPORANEOUS CORRELATIONS

Unfortunately, AA-SYM is not sufficient to obtain a unique exchangeable process consistent with an SEE belief structure. This lack of identification stems directly from the inability to observe the DM's beliefs regarding *contemporaneous* correlations. Consider two coins, a and b, which can both take values in $\{H, T\}$. Both coins are flipped each period. Consider the following two governing processes, which are i.i.d. across time periods. (1) the coins are perfectly correlated (with equal probability on HH and TT), or (2) the coins are identical and independent (and both have equal probability on H and T). Notice, the two cases induce the same marginal distributions over each coin *individually.* Thus, if the modeler has access only to the DM's marginal beliefs, the two processes are indistinguishable.

In this section we introduce a strengthening of exchangeability, which we aptly call *strongly exchangeable*, under which stochastic independence is preserved both intertemporally (as in vanilla exchangeability) and *contemporaneously*.[22]

Definition. *A process* $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$ *is* ***strongly exchangeable*** *if there exists a probability measure* θ *over* $\Delta^{IN} \equiv \prod_{a\in\mathcal{A}} \Delta(S_a)$*, such that*

$$\zeta(E) = \int_{\Delta^{IN}} \hat{D}(E) d\theta(D),$$

where for any $D \in \Delta^{IN}$*,* $\hat{D}$ *is the corresponding product measure over* $\mathcal{S}$*.*

Under a strongly exchangeable process the outcomes of actions that occur at the same time are independently resolved. Of course, this does not impose that there is no informational cross contamination between actions. Information regarding the distribution of action a is informative about the underlying parameter governing the exchangeable process, and therefore, also about the distribution of action b. Since exchangeable processes were first characterized as being invariant to permutations, for the sake of completeness we provide a similar characterization of strongly exchangeable processes.

Theorem 3.6. *The process* $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$ *is strongly exchangeable if and only if for any set of finite permutations* $\{\pi_a : \mathbb{N} \to \mathbb{N}\}_{a\in\mathcal{A}}$ *and event* $E = \prod_{n\in\mathbb{N}} \prod_{a\in\mathcal{A}} E_{n,a}$*, we have*

$$\zeta(E) = \zeta(\prod_{n\in\mathbb{N}} \prod_{a\in\mathcal{A}} E_{\pi_a(n),a}). \tag{3.4}$$

Proof. In Appendix B.3.

Following the intuition above, it should come as no surprise that under AA-SYM strong exchangeability can never be ruled out. In other words, there is no SEE belief structure, therefore no preferences over PoAs, that distinguishes exchangeability from strong exchangeability. Strongly exchangeable processes are ones where each dimension can be permuted independently. If $\pi_a = \pi_b$ for all $a, b \in \mathcal{A}$, the condition is exactly exchangeability. Strongly exchangeable process are especially relevant with respect to the current focus because they act as representative members to the equivalence classes of exchangeable processes consistent with the same SEE belief structure.

[22] We feel reasonably certain that strong exchangeability must have been studied previously in the statistics literature. However, we have found no references.

Theorem 3.7. *An SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ satisfies AA-SYM if and only if it is consistent with a strongly exchangeable process. Furthermore, such a strongly exchangeable process is unique.*

Proof. In Appendix B.3.

3.4.1 AA-SYM as a Behavioral Restriction

In this section we introduce the axiomatic counterpart of AA-SYM, and so we can identify Bayesianism in exploration environments directly from preferences over the strategies.

Definition. *Let π be an n-permutation and $p, q \in P$. We say that q is π-permutation of p if for all $h \in \mathcal{H}(p,n)$, $h' \in \mathcal{H}(q,n)$, $\text{proj}_{\mathcal{A}^n} h = \pi\big(\text{proj}_{\mathcal{A}^n} h'\big)$.*

If p admits any π-permutations it must be that the first n actions are assigned unambiguously (i.e., it does not depend on the realization of prior actions nor the objective randomization).

[AX.3.7: AA-SYM] *Let π be an n-permutation and $p, p' \in P$ with p' a π-permutation of p. Then, for all $a \in \mathcal{A}$, $\tau, \sigma, \sigma' \in \Sigma$, and $h \in \mathcal{H}(p,n)$, $h' \in \mathcal{H}(p',n)$, if h is a permutation of h' then*

$$p_{-n}\tau \succcurlyeq (p_{-n}\sigma)_{-h}\sigma' \iff p_{-n}\tau \succcurlyeq (p_{-n}\sigma)_{-h'}\sigma'.$$

After n periods the plan $p_{-n}\tau$ provides τ with certainty, while the plan $(p_{-n}\sigma)_{-h}\sigma'$ provides σ *unless* the history h occurs. Hence, the DM's preference between the plans depends on their ex-ante subjective assessment of how likely h is to occur. Similarly to the logic behind h-proportionality, AA-SYM states that the DM's assesses h to be exactly as probable as h'. In other words, the DM's likelihood of outcome realizations is invariant to the order in which the actions are taken. The intuition behind the next result is correspondingly straightforward.

Proposition 3.8 (Correlated Arms, Exchangeable Process)**.** *Let $\succcurlyeq$ admit an SEE representation with the associated observable processes $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$. Then, the following are equivalent:*

1. *$\succcurlyeq_h$ satisfies AA-SYM;*

2. $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ *satisfies AA-SYM;*
3. $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ *is consistent with an exchangeable process; and*
4. $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$ *is consistent with a (unique) strongly exchangeable process.*

Proof. The proof that condition 1 is equivalent to condition 2 is provided Appendix B.3. Conditions 2, 3, and 4 are equivalent due to Theorem 3.7.

The proposition implies that strong-exchangeability carries no additional restrictions, beyond those of exchangeability, on agents' preferences over the different strategies in bandit problems, and in particular on their optimal strategies.

3.5 FURTHER DISCUSSION

3.5.1 Related Literature

Within decision theory, the literature on learning broadly considers how a DM incorporates new information, generally via notions of Bayesianism and Exchangeability, and often in the domain of uncertainty: see Epstein and Le Breton (1993); Epstein and Seo (2010); Klibanoff *et al.* (2013); Lehrer and Teper (2015). Recently, there has been an interest in subjective learning, or, the identification of the set of possible "signals" that the DM believes she might observe. At it's most simple, this is the elicitation of the set of potential tastes (often referred to as subjective states) the decision maker anticipates, accomplished by examining the DM's preference over *menus* of choice objects: see Kreps (1979); Dekel *et al.* (2001). By also incorporating consumption goods that contract on an objective state space, the modeler can interpret the DM's preference for flexibility as directly stemming from her anticipation of acquiring information regarding the likelihood of states, as in Dillenberger *et al.* (2014); Krishna and Sadowski (2014).

There is also a small but highly relevant literature working on the identification of responsive learning. Hyogo (2007) considers a two-period model, with an objective state space, in which the DM ranks action-menu pairs. The action is taken in the first period and provides information regarding the likelihood of states, after the revelation of which, the DM choose a state-contingent act from the menu. The identification of interest is the DM's subjective interpretation of actions as signals. Similarly, Cooke (2016) entertains a similar model without the need for an objective state-space, and in which the consumption of a single object in the first period plays the role of a fully informative action. Cooke, therefore, identifies both the state-space and the corresponding signal structure. Piermont *et al.* (2016) consider a recursive and infinite horizon version of Kreps' model, where the DM deterministically learns about her preference regarding objects she has previously consumed. Dillenberger *et al.* (2015) consider a different infinite horizon model where the DM makes separate choices in each period regarding her information structure and current period consumption. It is worth pointing out, all of these models, unlike the this paper, capitalize on the "preference for flexibility" paradigm to characterize learning. We are able to identify subjective learning without appealing to the menu structure because of the purely responsive aspect of our model. In other words, flexibility is "built in" to our setup, as a different action can be taken after every possible realization of the signal (action).

3.5.2 Subjective Learning with Endogenous and Exogenous Information

As witnessed the literature covered above, there seems to be a divide in the literature regarding subjective learning. In one camp, are models that elicit the DM's perception of exogenous flows of information (as a canonical example, take Dillenberger *et al.* (2014)), and in the other are models that assume information is acquired only via actions taken by the DM (where this paper lies). Realistically, neither of these information structures capture the full gamut of information transmission in economic environments.

Consider the following example within the setup of the current paper. A firm is choosing between two projects (actions), a and b. Assume that each project has a high-type and a low type. The firm believes (after observing h) the probability that each project is the high-type is $\mu_{h,a}$ and $\mu_{h,b}$, respectively. By experimenting between a and b the firm's beliefs and preferences will evolve.

But, what happens if the firm anticipates the release of a comprehensive report regarding project a just before period 1? This report will declare project a high quality with probability

$\alpha^h > \frac{1}{2}$ if the projects true type is high and with probability $\alpha^l < \frac{1}{2}$ if it is low. Hence, the report is an informative signal. Now, if the firms belief after observing h in period 0 is given by $[\mu_{h,a}, \mu_{h,b}]$ then, according to Bayes rule, the firms belief regarding project a being the high-type, at the beginning of period 1 will be $\mu^+_{h,a} = \frac{\alpha^h \cdot \mu_{h,a}}{\alpha^h \cdot \mu_{h,a} + \alpha^l (1-\mu_{h,a})}$, if the report is positive, and $\mu^-_{h,a} = \frac{(1-\alpha^h)\cdot \mu_{h,a}}{(1-\alpha^h)\cdot \mu_{h,a} + (1-\alpha^l)\cdot(1-\mu_{h,a})}$ if the report is negative.

Unfortunately, however, the ex-ante elicitation of preferences in our domain cannot capture the anticipation of information. The firm is ranking PoAs according to its aggregated belief from the ex-ante perspective, and thus, so as to maximize its expected belief:

$$\big(\alpha^h \mu_{h,a} + \alpha^l(1-\mu_{h,a})\big)\mu^+_{h,a} + \big((1-\alpha^h)\mu_{h,a} + (1-\alpha^l)(1-\mu_{h,a})\big)\mu^-_{h,a} = \mu_{h,a}.$$

Because of the Bayesian structure, the DM's beliefs must form a martingale, so her expectation of her anticipated beliefs are exactly her ex-ante beliefs. This fact, coupled with the linearity of expected utility, imply that the DM's ex-ante preference over PoAs is unaffected by her anticipation of exogenous information arrival.

All hope is not lost, however, of fully characterizing the DM's subjective information structure. The approach of Dillenberger *et al.* (2014) is orthogonal to our's, leading us to conjecture that the two models can co-exist and impart a clean separation between exogenous and endogenous information flows. Going back to the example, imagine there are two PoAs, p and q such that p is preferred to q under beliefs μ^+_h, and q to p under μ^-_h. The DM would therefore strictly desire flexibility after period 0, even after she is able to condition her decision on h. Of course, because the report is released after period 0, irrespective of the action taken by the DM, for any 0-period history h', there must exist some other PoAs, p' and q', for which flexibility is strictly beneficial (after h').

3.5.3 A Comment on Bayesianism in Environments of Experimentation

The results in Section 3.4 have two related implications to Bayesianism in general models of experimentation. First, it is well known that the beliefs of two Bayesians observing the same sequence of signals will converge in the limit. Our results imply that in a setup of experimentation, different Bayesians obtaining the same information, might still hold different views of the world in the limit. Their beliefs over the uncertainty underlying each action will be identical, but they can hold different beliefs over the joint distribution.

The second point has to do with the possible equivalence with non-Bayesian DMs. Theorem 3.7 states that AA-SYM is necessary and sufficient for an SEE belief system to be consistent with some exchangeable process. As discussed in the Introduction, AA-SYM projected to stochastic processes is weaker than the standard symmetry axiom applied in the literature, because the standard assumption requires that histories fully specify the evolution of the state, while in our setup, histories can only specify cylinders. Because AA-SYM is a weaker assumption, de Finetti's theorem implies that processes satisfying such an assumption need not be exchangeable and have a Bayesian representation as in Eq. (3.2).

Consider the following example of a stochastic process. In every period two coins are flipped. In odd periods the coins are perfectly correlated (with equal probability on HH and TT), and in even periods the coins are identical and independent (and both have equal probability on H and T). The associated observable processes satisfy AA-SYM, but the process itself is clearly not exchangeable. Nevertheless, Theorem 3.7 guarantees that there is a (unique) strongly exchangeable process that is consistent with the SEE belief structure. In this case it is easy to see that that process would be the one in which every period we toss two coins that are identical and independent (and both have equal probability on H and T).

4.0 INTRODUCTION UNAWARENESS AND OBSERVABLE CHOICE

In which Alyosha, cognizant of his own limitations, and, not wanting to appear desultory, resolves to take things as they come.

4.1 INTRODUCTION

There is a marked difference between being *unaware* of one's preferences and not *knowing* (also referred to as being *uncertain* about) one's preferences. While certainty and uncertainty (about some piece of information, φ) together constitute awareness of φ, unawareness describes total ignorance –a complete lack of perception of, or ability to reason directly about φ.[23] This paper explores the behavioral (i.e., observable) identification between unawareness and uncertainty and contemplates the type of data required to make such a separation. Due to the consideration of observability, the primary interest is in the decision maker's preference (hypothetically embodied by choice data), and how patterns in preference change in response to the structure of awareness. I will show that unawareness produces distinct patterns, and so, attempting to model unawareness with uncertainty, regardless of how complex, will fail. As an example of when such issues may arise and how they might alter predictions, I consider a simple contracting environment. When unawareness is taken into account, players can have an incentive to conceal mutually beneficial information, leading to the optimality of incomplete contracts.

To highlight the distinction between uncertainty and unawareness, consider Alyosha, who will buy a new smartphone in six months. He will have three options at the time of purchase: x, y, and z, but, of course, might not know *now* which phone he would most like

[23] This was first noted by Modica and Rustichini (1994).

to purchase six months from now. This uncertainty could arise because he does not know the technical specifications of the phones, their price, etc., and his true preference depends on the realization of these variables. Contrast this to the case where Alyosha has never heard of phone z. Here, he is unaware of z, and so naturally, of any preferences there regarding. The key aspect, if Alyosha is unaware of a piece of information (the existence of phone z), he is unable to make any choice based directly on this information.

More subtle, but just as fundamental, is our acknowledgement of our own unawareness. Indeed, most people would readily admit to the possibility that they cannot conceive of all future technologies or trends, or exhaustively list the set of possible occurrences for the upcoming week. This recognition of unawareness is important because it suggests the things a decision maker (DM) is unaware of may play an indirect role in her decision making, even if they cannot be directly acted upon. Central to the analysis, then, is the DM who is (1) unaware, (2) aware she is unaware, and (3) unaware about what she is unaware. A DM in such an epistemic state is referred to as *introspectively unaware*. By contrast, a DM who does not satisfy condition (2) would be referred to as *naively unaware*. In the presence of introspective unawareness, Alyosha might envision a world in which he prefers *something* other than x and y. Of course, he cannot know this *something* is z, since that would require he is aware of it.

Notice, under either uncertainty or introspective unawareness, the DM has a natural inclination to delay making her choice (i.e., if she cannot start using the phone for six months, she might as well wait until then to choose). However, the motivation for delay is different. Under uncertainty, she would like to wait so as to make a decision based on the realization of the relevant variables (the technical specs, price, etc.). Under (introspective) unawareness, she would like to wait in case she becomes aware of something better than whatever she would have chosen today. Notice also, if the DM is unaware she is unaware, she has no reason to delay; she does not consider the possibility she becomes aware of new information.

Now, Alyosha is going to Father Zosima to purchase the phone on his behalf, and has to instruct him *today* about which phone to purchase in six months. If Alyosha is either uncertain or introspectively unaware of his preference, it will not be optimal for him to specify any single phone. In the case of uncertainty, however, he could leave detailed instructions for Zosima to carry out the optimal choice: in the event the technical specs are (s_x, s_y, s_z), and the prices are $(\$x, \$y, \$z)$, purchase phone x, in the event ... etc. A commitment to consume (in the future) a particular alternative given the state of affairs is referred to as a *contingent plan*. Since Alyosha's optimal decision depends on the realization of some variables,[24] it is enough for her to specify a contingent plan that depends on these variables. Contrast this to the case in which Alyosha introspectively unaware. No plan, at least no articulable one,[25] can carry out his optimal decision (this vignette is captured formally in Example 4.4.2). This is because he would need to alert Zosima to events that are described by information he is unaware of—to include such information in a contingent plan would require he is aware of it.

The main result of this paper shows the behavioral condition for introspective unawareness is a strict preference for delaying choices rather than committing to a contingent plan, even when every possible articulable plan is offered. In particular:

★1 *When the decision maker is fully aware, she is always willing to commit to some articulable contingent plan.* Intuitively, the DM's language is rich enough that she can contract on the resolution of any relevant uncertainty, effectively imitating whatever her dynamic behavior would have been.

★2 *Without full awareness, the DM might find every articulable contingent plan unacceptable.* A strict preference for delay (relative to all articulable contingent plans) is possible only if the DM is not fully aware –this behavior is an indication of unawareness. Intuitively, if DM believes she may become aware of new alternatives, she understands her future self may have options better than any she could currently articulate in a contingent plan.

★3 *If the DM does finds every articulable contingent plan unacceptable, she must be introspectively unaware.* Intuitively, the DM must be aware enough to come to the conclusion that waiting might afford new possibilities. When she is ignorant of her own unawareness, she cannot consider this possibility.

[24] Variables of which she is aware, of course.

[25] A contingent plan –a function from events to outcomes– is *articulable* if the DM is aware of its constituent parts, its domain and image.

So a preference for delay that cannot be appeased by the appeal to contingent planning is the behavioral indication –in an exact sense– of introspective unawareness. The intuition is exactly as in the above example: the DM's language is not rich enough to specify the optimal contingent plan (unawareness), but is rich enough that she knows this fact (awareness of unawareness).

Directly incorporating unawareness into a decision theoretic model introduces subtleties that need to be dealt with judiciously. First, one must take care to ensure the process of eliciting preferences from a DM does not affect her preferences. While asking a DM to rank risky prospects should not affect her risk preference, asking her to contemplate objects of which she was formerly unaware certainty would affect her awareness (for a longer discussion on this topic, see Section 4.1.1). Second, the *type* of unawareness considered (i.e., naive or introspective, object-based or state-based, etc.) must be rich enough to produce observable patterns, even when keeping in mind the pervious concern. Finally, Modica and Rustichini (1994); Dekel *et al.* (1998) show that within the context of state space models, simply assuming that the DM is unaware of certain states (while retaining desirable properties of knowledge) is insufficient: the DM will either be fully aware or fully unaware.

To overcome these obstacles, I develop a logical framework that directly incorporates the DM's preference, as well as her knowledge and awareness thereof. This ensures that my notion of awareness is well founded and rigorous, and allows me to directly verify that the elicitation method does not require the DM to contemplate objects she herself could not have articulated. Finally, this construction indicates that contingent plans are precisely the type of data required to observe unawareness. Giving the DM either less flexibility (for example, deterministic choice of a single element) or more flexibility (for example, choice over incomplete contingent plans) fails to separate unawareness and uncertainty. Although, in the end, choices will be observable, this more general framework will provide the tools to analyze the *epistemic conditions* (i.e., knowledge and awareness) that generate the observable patterns in choice data.

I begin with an epistemic modal logic, based on set of logical statements that include a formal description of the DM's preference, and adapted from Fagin and Halpern (1988), Halpern and Rêgo (2009) and Board and Chung (2011). Each *state of the world* is defined by a set of statements which are true: statements that include how the DM ranks objects (for example, "x is preferred to y"), what the DM *implicitly knows* (for example "the DM implicitly knows 'x is preferred to y' ") and what the DM *is aware of* (for example "the DM is aware of the statement 'x is preferred to y' "). The intersection of implicit knowledge and awareness is *explicit knowledge*. Implicit knowledge can be thought of as idealized knowledge –what the DM would know if she was fully aware and logically omnipotent. In contrast, explicit knowledge can be thought of as working knowledge, subject to cognitive limitations. Then, on top of this logic, I build a decision theory. This allows me to speak of a preference as true (irrespective of the DM's epistemic state), implicitly known (if the DM's preference is the logical consequence of things she knows), and explicitly known (if the DM implicitly knows her preference and is aware of it). If the DM is fully aware, her implicitly and explicitly known preferences coincide. The characterization of unawareness arises from the contrast between the structure of these different preferences.

Failing to account for the effect of unawareness can distort predictions. Specifically, ex-ante solution concepts in dynamic environments.[26] However, including the formal machinery needed to deal appropriately with these concerns can impose a large cost in simple models, and so, should not be done needlessly. This paper, therefore, provides a test for the presence for introspective unawareness, so that economists might better understand in which context unawareness is present and where it can be safely ignored.

To exemplify how introspective unawareness can alter behavior in economic settings, I examine a highly stylized strategic environment in which a (fully aware) principal offers a take-it-or-leave-it contract to a (introspectively unaware) agent. I show, when constrained to offer complete contracts, the principal might have an incentive to conceal Pareto improving information. Intuitively, this is because the introspectively unaware agent is unable to properly anticipate the value of the actions she is unaware of but might become aware of (i.e., her value to not committing). By unveiling new actions by including them in his contract, the principal alters the agent's epistemic state –her perceived value of delay– potentially increasing her aversion to commitment. When unconstrained, the principal can overcome this by

[26]For more concrete examples, the divergence from standard predictions in applications to game theory Feinberg (2012), contract theory Filiz-Ozbay (2012), information economics Heifetz *et al.* (2013), and mechanism design Auster (2013).

leaving the contract incomplete. I show, because this behavior is motivated by the agent's fuzzy perceived value to delay, it arises only in the presence of introspective unawareness.

In addition, the epistemic preference framework relates to the decision theoretic literature on subjective state spaces, which began with Kreps (1979). I show that the Krepsian paradigm, epitomized by DM who ranks *menus* of objects, and whose ranking respects set inclusion (so that larger menus are preferred to their subsets), can be faithfully captured by a special case of the models studied in this paper. Such behavior arises without the need to appeal to unawareness. Intuitively, the flexibility that arises by appealing to menus is, just as with contingent plans, bounded by the DM's articulation. Hence, introspective unawareness produces the same unwillingness to commit in this domain.

In particular, this characterization is of interest in relation to models of subjective learning. To identify what the DM believes she might learn, axiomatizations generally include the requirement that any dynamic choice behavior is indifferent to some contingent plan –in essence, assuming the existence of acceptable, and articulable, plans.[27] As such, the results of this paper mandate that a theory of subjective learning under unawareness cannot be built on the same machinery. Put differently, current models of subjective learning necessarily reduce all learning to resolution of uncertainty rather than from the arrival unanticipated information.

4.1.1 Observability and Unawareness

Most decision theoretic models begin with the declaration of some set, X, over which the DM's preferences are described by a binary relation, $\succcurlyeq \subseteq X \times X$. Under the revealed preference interpretation, the modeler precludes from the outset the possibility of alternatives of which the decision maker is unaware. Indeed, if the modeler were to ask a person on the street, or a subject in the lab, to choose between phone x and phone y, it is unreasonable to believe, at the time of her answer, she is unaware of either x or y. The very act of asking

[27] The literature on learning has been principally interested in the case where the DM entertains a subjective state space, and identification regards the set of events in this state space the DM believes she might learn. Dillenberger *et al.* (2014, 2015); Piermont *et al.* (2016) have considered constructions that exactly correspond to contingent plans in this paper. While Ergin and Sarver (2010) and Riella (2013) do not directly construct such plans, the interpretation of both papers concerns a DM who constructs a contingent plan *after* observing the menu they receive.

forces the DM's hand. When moving to the realm of contingent plans –functions from events to consumption– this problem not only persists, but is compounded. Now, the modeler is precluding unawareness both of the consumption space and the set of contingencies. Finally, when trying to identify introspection, the problem becomes even more precarious. Any question regarding even the *existence* of unforeseen objects has the potential to change the DM's epistemic state.

The decision theoretic literature has posited many different behavioral markers for unawareness: incomplete preferences, preference for flexibility Kreps (1979), reverse Bayesianism Karni and Vierø (2016), unmeasurable states Kochov (2015); Minardi and Savochkin (2015); Grant and Quiggin (2014), etc.[28] However, these papers, formulated from a revealed preference approach, must make restrictive (ex-ante) assumptions on the DM's epistemic state. For example, Kochov (2015) and Minardi and Savochkin (2015) assume the modeler has a strictly more complete view of the world (is aware of more) than the agents within the model. This means it is impossible to detect unawareness in agents who are more aware than the modeler, severely narrowing the contexts in which these models can be applied. On the other hand, Karni and Vierø (2016) and Grant and Quiggin (2014) allow the DM's unawareness to be largely unrestricted. They do, however, assume that the DM is introspectively unaware, requiring her to rank objects explicitly containing *surprising* outcomes.

This paper asks the precedent question as to how such epistemic states might be identified. Such issues are dealt with by relaxing what is meant by the revealed preference approach. Instead of providing the decision maker with a set and asking her to indicate her preferences thereover, the modeler asks the DM to provide both the set of alternatives and the list of preference restrictions between them. That is to say, the DM provides a set of statements such as "x is preferred to y," where x and y are object of which she herself conceives. While this is clearly inadequate for some practical purposes, the complexity of the DM's task is no greater than in the standard model, and it does not beg the question of awareness.

A crucial aspect to the identification of unawareness contained in this paper is that

[28]For a detailed account of these papers' and how they relate to this one, see the literature review in Section 4.7.

it never requires the DM to contemplate objects she herself could not have conceived. It suffices for the modeler to consider the DM's preference over the set of objects she herself reported. The main contribution of this paper is, therefore, the assertion of a framework that characterizes introspective unawareness from choices regarding only information of which the DM is aware.

4.1.2 A (Far Too) Simple Model of Decision Making Under Unawareness.

Outside of the limits of real cognition, let X represent the objective set of all possible alternatives in all possible worlds for all possible decision makers. Assume the modeler asks the DM to report her personal set of preference restrictions (a set of ordered pairs of objects the DM conceived of). Then, the reported preference restrictions are a subset of $X \times X$. If the DM is unaware of some of the elements of X then she cannot include them in her report.

So, let $\succcurlyeq \subseteq X \times X$ be the decision maker's personal set of preference restrictions.[29] The highest standard of rationality would dictate that $\succcurlyeq$ is a transitive, complete, and reflexive relation. If we allow for unawareness, however, what properties should $\succcurlyeq$ possess? Transitivity has no reason to be discarded; if the decision maker's true preference is transitive and she is aware $x \succcurlyeq y$ and $y \succcurlyeq z$, then she has all the constituent parts to deduce that $x \succcurlyeq z$.

[AX.4.1: Trv] *For all $x, y, z \in X$, $x \succcurlyeq y$ and $y \succcurlyeq z$ then $x \succcurlyeq z$.*

Completeness, however, is clearly too strong; so too reflexivity. If the DM is unaware of x, she cannot report $x \succcurlyeq x$, even if it is true. However, if the DM is aware that $x \succcurlyeq y$ and that her preferences are complete, she can deduce that $x \succcurlyeq x$ and $y \succcurlyeq y$. Likewise if she is aware of both $x \succcurlyeq x$ and $y \succcurlyeq y$, she can deduce that there must be some preference between x and y. This is the idea behind local completeness:

[AX.4.2: LCmp] *For all $x, y \in X$, '$x \succcurlyeq y$ or $y \succcurlyeq x$' if and only if '$x \succcurlyeq x$ and $y \succcurlyeq y$'.*

Trv and LCmp provide enough structure to provide a basic, but mostly pointless, result. There exists some $\mathcal{A} \subseteq X$ of objects the DM is *aware of*, and a utility function $U : \mathcal{A} \to \mathbb{R}$

[29]For any relation $\succcurlyeq$, let $\sim$ an $\succ$ denote the symmetric and antisymmetric components.

that represents $\succcurlyeq$.[30] If $x \notin \mathcal{A}$, then $U(x)$ is undefined and there is no $y \in X$ such that $x \succcurlyeq y$ or $y \succcurlyeq x$: the DM is unaware of x. In addition, LCmp can be weakened to local reflexivity, to allow the decision maker to have incomplete preferences, even over objects she is aware of:

[AX.4.3: LRfx] *For all $x, y \in X$, if $x \succcurlyeq y$ then '$x \succcurlyeq x$ and $y \succcurlyeq y$'.*

Trv and LRfx provide the analog to discarding completeness (while retaining reflexivity) in the standard model Ok (2002). There exists some $\mathcal{A} \subseteq X$ and a set of utility functions $U_k : \mathcal{A} \to \mathbb{R}$ such that $x \succcurlyeq y$ if and only if $U_k(x) \geq U_k(y)$ for all k.

These results show, in static contexts, unawareness has only trivial behavioral implications. Is there any economic content to considering X rather than $\mathcal{A}$? The behavior is *locally* identical –the patters of observable choice will be the same but for the inclusion or exclusion of particular elements. As alluded to in the pervious section, to see the full effects of unawareness, dynamic environments are needed. Of course, even there, if the DM is unaware she is unaware then she ignores the consequences of her limited understanding entirely; she acts as in the static case outlined above. In light of this, my focus is the introspectively unaware decision maker in a dynamic setting.

4.1.3 Organization

The structure of the paper is as follows. Section 4.2 introduces the logical underpinnings of the decision theory and expounds on the choice patterns based on implicitly known preference. Section 4.3 formally introduces awareness structures and explicit knowledge. The main results are contained in Section 4.4, which introduces contingent plans and the notions of implicit and explicit acceptability. Section 4.5 show the connection to subjective state space models and a preference for flexibility. Section 4.6 explores a simple strategic contracting game. A survey of the relevant literature can be found in Section 4.7. Additional results and proofs omitted from the text are contained in the appendix.

[30] Later, it is assumed that the set of consumable objects is denumerable. Here, however, if X is uncountable, the relevant continuity assumptions must also be made.

4.2 LOGICAL FOUNDATIONS: PREFERENTIAL LOGIC

This section outlines the formal construction of the logic used in this paper. First Section 4.2.1 provides the syntax for well defined formulae. That is, a purely mechanical account of which strings of characters will be *well defined*. Then Section 4.2.2 endows well defined formulae with meaning by providing a semantic interpretation. This interpretation is the standard *possible worlds semantics* adapted to consider preferential statements. Finally, section 4.2.3, considers an axiomatization (a method of deriving new true statements from old ones) corresponding to the semantic models.

4.2.1 A Syntactic Language: $\mathcal{L}(\mathcal{X})$

Preferential choices will be described directly by an epistemic logic. To this end, for each $n \geq 1$, define a countable set of n place predicates denoted by $\alpha, \beta, \gamma, \ldots$. Assume the existence of a countably infinite set of variables denoted by $\mathcal{X} = a, b, c \ldots$. Then, any n place predicate followed by n variables is a well formed *atomic* formula. That is, if α is a 2 place predicate, then αab is a well formed atomic formula, with the interpretation that a and b stand in the α relation to one another. For example, if α is "greater than", then αab states that a is greater than b. Also assume the existence of a distinguished predicate, $\succcurlyeq$, representing weak preference (where $(a \succcurlyeq b)$ is used rather than $(\succcurlyeq ab)$). Take note that variables are placeholders, and, until endowed with an interpretation, do not refer to any specific object.

Define the set of well formed formulae recursively: for any well formed formulae, φ and ψ, $\neg\varphi$, $\varphi \wedge \psi$ and $K_t\varphi$ are also well formed (for $t = 0 \ldots T$). K_t represents implicit knowledge at time t, the interpretation of which will be standard: the DM implicitly knows φ (at time t) if φ is true in every state of affairs she considers possible (at time t).[31] In addition, this language allows for universal quantification, $\forall$. So for any well formed φ, $\forall a\varphi$ is well formed, where a any individual variable. The resulting language is $\mathcal{L}(\mathcal{X})$.

Taking the standard shorthand, $\varphi \vee \psi$ is short for $\neg(\neg\varphi \wedge \neg\psi)$, $\varphi \implies \psi$ is short for

[31] The reason this type of knowledge is qualified as implicit, is in contrast to explicit knowledge that requires the DM also be aware of φ.

$\neg\varphi \vee \psi$, and $\exists a\varphi$ is short for $\neg\forall a\neg\varphi$. In addition, let P_t denote $\neg K_t\neg$, with the intended interpretation of $P_t\varphi$ as the DM considers φ possible; she does not know it is not the case. Lastly, to expedite notation, for any φ let $\nabla_t\varphi$ denote $((\varphi \implies K_t\varphi) \wedge (\neg\varphi \implies K_t\neg\varphi))$ with the intended interpretation, the DM knows the truth value of φ at time t.

Per usual, an occurrence of a variable a is *free* in a formula φ if a is not under the scope of a quantifier, and is *bound* otherwise. A formula with no free occurrences is called a *sentence*. Two formulae φ and ψ are called *bound alphabetic variants* of one another if φ and ψ differ only because where φ has well formed sub-formulae of the form $\forall a\zeta$ where ψ has $\forall b\eta$ and ζ has free occurrences of a in exactly the same places as η has free occurrences of b.

If φ is a formula, then $\varphi[a/b]$ denotes the formula created by replacing all (and possibly no) free occurrences of a with b. Because this can change the interpretation of the formula in unintended ways, (in particular, if there was a free a in φ that corresponds to a bound b in $\varphi[a/b]$) I introduce the following notation: $\varphi[[a/b]]$ denotes the formula created first by taking a bound alphabetic variant of φ with no bound occurrences of b, and then replacing every free a with b.

4.2.2 Semantics

I will work with a fixed domain model. This means the variables (and hence quantification) range over the same domain in every possible world. A word should be said on this, as there is considerable philosophical debate regarding constant domains. On one hand, it simplifies matters considerably to assume the same objects *hypothetically* exist in each possible world. One the other, the very intention that possible worlds be distinct means they might be defined by different objects. Here, I take the view that different possible worlds are defined by the different relation between objects (that may or may not exists, but always hypothetically exist). This view coincides with the main emphasis of looking at different worlds primarily as embodiments of different preferences. The DM conceives of possible worlds in which she entertains different preferences (where these preferences may be the result of different knowledge or awareness) – this is perfectly possible without the introduction of varying domains.

So, to begin, let X denote a domain of the individual variables, the class of all possible values a variable might take. Elements of X are referred to using $x, y, z \ldots$. Notice, this is same notation as used for the domain of consumption alternatives in previous sections. This is intentional, as I will interpret X as being a domain of consumable objects. Let an *assignment* be a function from the set of individual variables into X. If μ and μ' are both assignments such that differ only in the object assigned to a then they are referred to as a-variants, and related by $\mu \sim_a \mu'$.

Then, for a given language, $\mathcal{L}(\mathcal{X})$, each DM is characterized by the tuple

$$M = \langle S, X, \mathcal{V}, \{R_t\}_{t \leq T}, \{\succcurlyeq_s\}_{s \in S} \rangle.$$

M is referred to as a model (or, a model of decision making). $S = \{s, s', \ldots\}$ is a set of states of the world. $\{R_t\}_{t \leq T}$ is a time indexed family of accessibility relations on S; as is standard, the interpretation of $R_t(s) = \{s' | sRs'\}$ is the states the DM considers possible when the true state is s. Truth values will be assigned by $\mathcal{V}$, a function that assigns to each n place predicate and state of the world s, a class of n-tuples from X. The intended interpretation is, for some $\langle X, \mathcal{V} \rangle$, if $(x_1 \ldots x_n) \in \mathcal{V}(\alpha, s)$, with $x_1 \ldots x_n \in X$, then $\alpha x_1 \ldots x_n$ is true in that model in state s.

Let $\mathcal{M}(\mathcal{X})$ be the class of all models based on $\mathcal{L}(\mathcal{X})$, and for a given domain, X, let $\mathcal{M}^X$ denote the subclass of models based on X. Then a DM, M, is represented semantically via the operator $\models$, recursively, as

$$
\begin{array}{ll}
(M, s) \models_\mu \alpha a_1 \ldots a_n & \text{iff } (\mu(a_1) \ldots \mu(a_n)) \in \mathcal{V}(\alpha, s), \text{ for atomic formula (except } \succcurlyeq), \\
(M, s) \models_\mu (a \succcurlyeq b) & \text{iff } \mu(a) \succcurlyeq_s \mu(b), \\
(M, s) \models_\mu \neg \varphi & \text{iff not } (M, s) \models_\mu \varphi, \\
(M, s) \models_\mu (\varphi \wedge \psi) & \text{iff } (M, s) \models_\mu \varphi \text{ and } (M, s) \models_\mu \psi, \\
(M, s) \models_\mu K_t \varphi & \text{iff for all } s' \in R_t(s), (M, s') \models_\mu \varphi, \\
(M, s) \models_\mu \forall a \varphi & \text{iff for } \mu' \sim_a \mu', (M, s) \models_{\mu'} \varphi.
\end{array}
$$

A formula φ is *satisfiable* if there exists a M, and a state thereof, s, and an interpretation μ, such that $(M, s) \models_\mu \varphi$. If a $(M, s) \models_\mu \varphi$ for every assignment μ, write $(M, s) \models \varphi$. Given a DM, M, φ is *valid* in M, denoted as $M \models \varphi$, if $(M, s) \models \varphi$ for all s. Likewise, for some class of DMs, $\mathcal{N}$, φ is *valid* in $\mathcal{N}$, denoted as $\mathcal{N} \models \varphi$, if $N \models \varphi$ for all $N \in \mathcal{N}$. Finally, φ is *valid* (i.e., without qualification) if $M \models \varphi$, for all models M.

4.2.3 Axioms

Consider the following axiom schemata (and inference rules) regarding the language $\mathcal{L}(\mathcal{X})$:

[AX.4.4: Prop] *All substitution instances of valid formulae in propositional logic.*

[AX.4.5: K] $(K_t\varphi \wedge K_t(\varphi \implies \psi)) \implies K_t\psi$.

[AX.4.6: 1∀] $\forall a\varphi \implies \varphi[[a/b]]$

[AX.4.7: Barcan] $K_t\forall a\varphi \implies \forall aK_t\varphi$

[AX.4.8: MP] *From* φ *and* $(\varphi \implies \psi)$ *infer* ψ.

[AX.4.9: GenK] *From* φ *infer* $K_t\varphi$.

[AX.4.10: Gen∀] *From* $\varphi \implies \psi$ *infer* $\varphi \implies \forall a\psi$*, provided* a *is not free in* φ.

Denote $\boldsymbol{K_T} = (\text{Prop} \cup \text{K} \cup \text{MP} \cup \text{GenK})$ and $\forall\boldsymbol{K_T} = \boldsymbol{K_T} \cup (1\forall \cup \text{Barcan} \cup \text{Gen}\forall)$. It is well known, $\forall\boldsymbol{K_T}$ is a sound and complete axiomatization of the first order language $\mathcal{L}(\mathcal{X})$ with respect to the $\mathcal{M}(\mathcal{X})$ (for example see Chapters 13 and 14 of Hughes and Cresswell (1996)).[32] That is, the above axiom system exactly captures the semantic structure imposed above.

Further axioms can impose structure on the DMs knowledge, and hence, semantically, on the family of accessibility relations. Consider the following:

[AX.4.11: D] $K_t\varphi \implies P_t\varphi$.

[AX.4.12: T] $K_t\varphi \implies \varphi$.

[AX.4.13: 4] $K_t\varphi \implies K_tK_t\varphi$.

[32] Given an axiom system $\boldsymbol{AX}$, and a language $\mathcal{L}$, we say that the formula $\varphi \in \mathcal{L}$ is a *theorem* of $\boldsymbol{AX}$ is it is an axiom of $\boldsymbol{AX}$ or derivable from previous theorems using rules of inference contained in $\boldsymbol{AX}$. Further, $\boldsymbol{AX}$ is said to be *sound*, for the language $\mathcal{L}$ with respect to a class of structures $\mathcal{N}$ if every theorem of $\boldsymbol{AX}$ is valid in $\mathcal{N}$. Conversely, $\boldsymbol{AX}$ is said to be *complete*, for the language $\mathcal{L}$ with respect to a class of structures $\mathcal{N}$ if every valid formula in $\mathcal{N}$ is a theorem of $\boldsymbol{AX}$.

[AX.4.14: 5] $P_t\varphi \implies K_t P_t \varphi$.

It is well known, in the presence of $\forall \boldsymbol{K_T}$, T, 4, and 5 correspond to the class of models where $\{R_t\}_{t\leq T}$ is a family whose members are reflexive, transitive, and Euclidean,[33] respectively (see Fagin *et al.* (2003) for the propositional case and Hughes and Cresswell (1996) for a first order treatment). Of note is the system $\boldsymbol{S5} = (\forall \boldsymbol{K_T} \cup \mathrm{T} \cup 4 \cup 5)$, corresponding to the class of models where $\{R_t\}_{t\leq T}$ is a family whose members are equivalence relations, and therefore, partition the state space.

4.2.4 Implicit Preferences

Just as we can axiomatize the structure of knowledge, so to can we provide the structure to preferences. Consider the following basic axioms:

[AX.4.15: CMP] $\forall a \forall b \big(\neg(a \succcurlyeq b) \implies (b \succcurlyeq a)\big)$.

[AX.4.16: TRV] $\forall a \forall b \forall c \big((a \succcurlyeq b) \wedge (b \succcurlyeq c) \implies (a \succcurlyeq c)\big)$.

Denote $\boldsymbol{P} = (\textsc{Cmp} \cup \textsc{Trv})$. In a similar spirit to the above results, CMP, TRV correspond to models where $\{\succcurlyeq_s\}_{s\in S}$ is a family whose members are complete and transitive, respectively. Formally, denote by the superscripts, r, t, e, the classes of DMs such that $\{R_t\}_{t\leq T}$ are reflexive, transitive, and Euclidean, and by cmp, trv the class of DMs such that $\{\succcurlyeq_s\}_{s\in S}$ that are complete, and transitive, respectively. For example, $\mathcal{M}^{r,e,cmp,trv}(\mathcal{X})$ is the class of DM for which R_t is reflexive and Euclidean for all t, and $\succcurlyeq_s$ is complete and transitive for all $s \in S$.

Proposition 4.1. *Let $\mathcal{C}$ be a possibly empty subset of $\{\mathrm{T}, 4, 5, \textsc{Cmp}, \textsc{Trv}\}$, and let C be the corresponding subset of $\{r, e, t, cmp, trv\}$. Then $\forall \boldsymbol{K_T} \cup \mathcal{C}$ is a sound and complete axiomatization of $\mathcal{L}(\mathcal{X})$ with respect to $\mathcal{M}^C(\mathcal{X})$.*

Proof. In appendix C.2.

Preferential axioms play the role of traditionally decision theoretic restrictions (i.e., completeness, transitivity, etc); any (satisfiable) theory including these restrictions will have a

[33] Recall, a relation is *Euclidian* if xRy and xRz imply yRz.

model of decision making adhering to the corresponding decision theoretic framework. The importance, therefore, of this framework is that it provides us a language to make a clean distinction between non-modal preference and its epistemic counterpart and to analyze the interplay there between. Specifically, the distinction between some elementary (read, true) preference and the preference the decision maker *knows* or is *aware of*.

The discrepancy between the DM's "true" preferences and her implicitly known preferences (and later, in the presence of unawareness, her explicitly known preferences), can be made formal. To do this, define, in addition to $\succcurlyeq_s$, implicitly known preference, $\succcurlyeq_{K_t,s}$.

Definition. *Let M be a model of decision making. Define the implicit preference relation, as*

$$\succcurlyeq_{K_t,s} = \bigcap_{s' \in R_t(s)} \succcurlyeq_{s'} \tag{4.1}$$

The interpretation is as follows: If the true state is s, then the DM is endowed with the preference $\succcurlyeq_s$. However, she may not know this is her preference. In particular, if $(M, s) \models_\mu (a \succcurlyeq b) \wedge \neg K_t(a \succcurlyeq b)$, then in state s at time t, the DM prefers $\mu(a)$ to $\mu(b)$, but does not (implicitly) know her preference. For the DM to make a report $\mu(a) \succcurlyeq \mu(y)$, or take an action on the basis of said preference, she must, as a prerequisite, know she prefers $\mu(a)$ to $\mu(b)$.

Remark 4.2. *$x \succcurlyeq_{K_t,s} y$ if and only if $(M, s) \models_\mu K_t(a \succcurlyeq b)$ holds for any μ such that $\mu(a) = x$ and $\mu(b) = y$.*

If M models a theory that includes preferential axioms, then the DM knows these restrictions. For example, if $M \in \mathcal{M}^{cmp,trv}(\mathcal{X})$, then the DM knows her preference are complete and transitive (since every instance of CMP and TRV hold at every state, as implied by GENK). This observation implies that while the DM may not know her true preference, $(M, s) \models_\mu \neg K_t(a \succcurlyeq b) \wedge \neg K_t(b \succcurlyeq a)$, she nonetheless knows the structure of her true preference, $(M, s) \models_\mu K_t((a \succcurlyeq b) \vee (b \succcurlyeq a))$. As a result, structural ignorance, where the DM does not know the structure of her preferences, cannot be captured here. This behavior can be captured once awareness is introduced.

When the decision maker's true preferences are a weak order, her implicit preferences are a preorder (i.e., reflexive and transitive). While the implicit preferences inherit reflexivity and transitivity, the same can not be said about completeness. Even if the DM knows her preferences are complete, she might not know what her preference is. The framework takes a complete preference relation and returns an incomplete preference relation by nature of introducing uncertainty about the true state of affairs. Indeed, consider the following example.

Example 4.2. *There is a single time period. There are two states of the world,* $S = \{s, s'\}$ *and two elements that can be consumed,* $X = \{x, y\}$*. The preference relations in each state are given by* $\succcurlyeq_s = \{(x,x),(x,y),(y,y)\}$ *and* $\succcurlyeq_{s'} = \{(x,x),(y,x),(y,y)\}$ *and the accessibility relation is the trivial* $R = S^2$*. So,* $\succcurlyeq_{K,s} = \succcurlyeq_{K,s'} = \{(x,x),(y,y)\}$*, which are not complete. So,* $M \models_\mu K((a \succcurlyeq b) \vee (a \succcurlyeq b)) \wedge \neg K(a \succcurlyeq b) \wedge \neg K(a \succcurlyeq b)$*, for any* μ *such that* $\mu(a) = x$ *and* $\mu(b) = y$*. Notice, the DM also satisfies negative introspection (i.e., 5), so she knows she does not know her preference:* $M \models_\mu K\big(\neg K(a \succcurlyeq b) \wedge \neg K(a \succcurlyeq b)\big)$.

It is well known, in finite domains, incomplete preferences can be represented by family of utility functions Ok (2002); Evren and Ok (2011). Multi-utility representations identify a set of utility functions, $\mathcal{U}$, such that x is preferred to y if and only if $u(x) \geq u(y)$ for all $u \in \mathcal{U}$. By nature of producing an incomplete preference relation, this framework also allows such a representation. From the conceptual analogy between the representation and the set $\mathcal{U}$ and the implicitly known preferences and the set $R(s)$, it is clear this framework instills the natural interpretation to the set of utility functionals.

Remark 4.3. *Let* M *be a model of* $\forall \boldsymbol{K_T} \cup \boldsymbol{P}$*. For each* $s \in S$*, let,*

$$\mathcal{U}_{s,t} = \{u_{s'} : X \to \mathbb{R} \mid u_{s'} \text{ represents } \succcurlyeq_{s'} \text{ and } s' \in R_t(s)\}. \tag{4.2}$$

Then, for every $s \in S$ *and* $t \leq T$*,*

1. $\succcurlyeq_{K_t,s}$ *is reflexive and transitive, and*
2. $x \succcurlyeq_{K_t,s} y$ *if and only if* $u(x) \geq u(y)$ *for all* $u \in \mathcal{U}_{s,t}$.

Proof. Fix $s \in S$, and $t \leq T$. By Proposition 4.1, $(x, x) \in \succsim_s$ for all $s \in S$ and $x \in X$. So $(x, x) \in \bigcap_{s' \in R_t(s)} \succsim_{s'} = \succsim_{K_t,s}$. $\succsim_{K_t,s}$ is reflexive. Now let, $(x, y), (y, z) \in \succsim_{K_t,s}$. So, for each $s' \in R_t(s)$, $(x, y), (y, z) \in \succsim_{s'}$. By Proposition 4.1, $\succsim_{s'}$ is transitive, so $(x, z) \in \succsim_{s'}$. Hence, $(x, z) \in \bigcap_{s' \in R_t(s)} \succsim_{s'} = \succsim_{K_t,s}$. $\succsim_{K_t,s}$ is transitive.

Now, to establish the representation claim: $u(x) \geq u(y)$ for all $u \in \mathcal{U}_{s,t}$ if and only if $u_{s'}(x) \geq u_{s'}(y)$ for all $s' \in R(s)$, where $u_{s'}$ represents $\succsim_{s'}$. By the definition of a representation, this is if and only if $x \succsim_{s'} y$ for all $s' \in R_t(s)$, the definition of implicit preference.

Moreover, any incomplete preference (over a finite domain) can be generated by preferential logical model. This can be seen by constructing, for each $u \in \mathcal{U}$ some s_u such that $\succsim_s$ is the order induced by u. Then any model such that $S \supseteq \{s_u | u \in \mathcal{U}\}$ and $R(s) = \{s_u | u \in \mathcal{U}\}$ will suffice. As such, simple models of incomplete preferences can be seen as models of epistemic uncertainty: a DM who is aware of all of her options but is unsure about the true state of affairs. Her incompleteness is derived form a lack of knowledge, even though her preferences exist and she knows her preferences exist. The set of preference relations she considers in her multi utility model are her true preferences in each world she considers possible.

4.3 AWARENESS STRUCTURES

To directly incorporate unawareness, add two time indexed modal operators, A_t and E_t for $t = 1 \dots T$, to the logic. Starting with the same set of atomic propositions as above, refer to the resulting language as $\mathcal{L}^A(\mathcal{X})$. The interpretation of the modal operators is as in Fagin and Halpern (1988). K_t is implicit knowledge at time t; an agent knows φ, denoted $K_t\varphi$, if φ is true in all state she considers possible at time t. A_t is awareness at time t; $A_t\varphi$ is interpreted as the DM is aware of φ at time t. Lastly, E_t is explicit knowledge —the conjunction of K_t and A_t. The DM explicitly knows φ if she implicitly knows it and is aware of it: $K_t\varphi \wedge A_t\varphi$. Set P_t^E as short hand for explicit possibility, or $\neg E_t \neg$.[34]

[34]Explicit possibility states the DM does not explicitly know the negation of a formula –hence she would not act in response to its negation. Note that while explicit knowledge is a stricter requirement that implicit

The choice of semantics for awareness in a first order model is non-trivial. One choice is, for example, to assume awareness is comprised of some subset of the domain over which variables range –implying if the decision maker is aware of all the objects specified in a formula, she is aware of the formula itself (i.e., "Object Based Unawareness" investigated by Board and Chung (2011); Board *et al.* (2011)). But this is very limiting, as it prohibits the DM from from being aware of an object but unaware of some of its attributes –a necessary feature to model many economic environments. For example, a DM is choosing between a Mac and a PC, and currently does not know her preference. However, after learning that her job requires her to use a particular piece of software, which runs only on Windows, and which she had been unaware of, she prefers the PC.

Another option, as in the logic of Halpern and Rêgo (2009), is that awareness is comprised of some subset of the language (a set of formulae or sentences). This is closer to ideal but still not exactly right. The reason being, I wish to interpret awareness semantically (i.e., by assigning meaning to the variables in a formula), rather than purely syntactically (i.e., leaving variables as variables). To see why, consider

$$A_0 \forall a (a \succcurlyeq a), \tag{4.3}$$

$$A_0 \forall a (a \succcurlyeq b), \tag{4.4}$$

where both a and b are variables. The first formula states the DM is aware $\succcurlyeq$ is a reflexive relation. The assignment of variables plays no role in the interpretation of the first formula, since the only variable, a, is under the scope of a quantifier. The second formula, however, has a variable b that is not under the scope of any quantifier. To provide this formula with semantic meaning, first b must be given a meaning via a semantic interpretation. The intention is that the DM is aware of information, statements that have specific interpretations, and when b is left as a variable, (4.4) does not have a specific interpretation.

My approach strikes a balance between these two options. I formally construct the awareness structure such that it ranges over formulae (unlike Board and Chung (2011); Board *et al.* (2011)), but *after* they have been assigned a semantic meaning (unlike Halpern

knowledge, the opposite relation is true for possibly. For this reason, Remark 4.10 and Theorem 4.12 require that the DM explicitly knows the explicitly possibility in question.

and Rêgo (2009)). This construction uses a second language intimately related to $\mathcal{L}^A(\mathcal{X})$ –the language is exactly $\mathcal{L}^A(\mathcal{X})$ but with free variables replaced with logical constants.

4.3.1 A Semantic Language: $\mathcal{L}(X)$

The following is done with the intention of representing the DM's awareness over semantic statements, formulae that have interpretations –whose variables are assigned to elements of X. Of course, in order to allow for introspection, the DM must be able to be aware of formulae that contain bound variables. To do this, I define an auxiliary language $\mathcal{L}(X)$. The construction of $\mathcal{L}(X)$ is as follows. First, any n place predicate followed by n *members of* X is a well formed *atomic* formula. Then for well formed formulae, $\bar{\varphi}$ and $\bar{\psi}$, $\neg\bar{\varphi}$, $\bar{\varphi} \wedge \bar{\psi}$, $\Box\bar{\varphi}$ are also well formed. Lastly for any formula $\bar{\varphi}$, $\forall a \bar{\varphi}[x/a]$ is well formed, where $\bar{\varphi}[x/a]$ is the formula $\bar{\varphi}$ with all (and possibly no) occurrences of x replaced with a, for some $x \in X$ and $a \in \mathcal{X}$. Notationally, $\bar{\varphi}$ (with a bar) will always denote a formula in $\mathcal{L}(X)$, whereas φ (without a bar) will denote a formula in $\mathcal{L}^A(\mathcal{X})$.

The formulae of $\mathcal{L}(X)$ correspond precisely the formulae of $\mathcal{L}^A(\mathcal{X})$ but with semantic interpretation attached to the free, and only free, variables. To see this, consider *reduced assignments*, which are purely syntactic transformations of formulae. A reduced assignment $\bar{\mu}$, which corresponds to the assignment μ, is a function that takes as an input a well formed formula $\varphi \in \mathcal{L}(\mathcal{X})$ and returns the string of characters created by taking φ and replacing every free variable, a with $\mu(a)$. Lemma 8 states the output of a reduced assignment is always in $\mathcal{L}(X)$, and every formula in $\mathcal{L}(X)$ is the output of some reduced assignment acting on some formula of $\mathcal{L}^A(\mathcal{X})$.

This connection allows us to describe awareness as a set, $\mathcal{A} \subseteq \mathcal{L}(X)$, and hence naturally endowed with a semantic interpretation. Under a given assignment μ, $\mathcal{A}$ represents the DM's awareness of the formulae $\{\varphi \in \mathcal{L}(X) | \bar{\mu}(\varphi) \in \mathcal{A}\}$. Lemma 8, ensures this conception is not limiting: for any $\Gamma \subseteq \mathcal{L}(X)$ there exist some $\mathcal{A}(\Gamma, \mu) = \{\bar{\varphi} \in \mathcal{L}(X) | \bar{\varphi} = \bar{\mu}(\varphi), \varphi \in \Gamma\} \subseteq \mathcal{L}(X)$ such that DM is aware of exactly Γ, given the semantic interpretation μ.

4.3.2 The Semantics of Awareness

A DM with an awareness structure is described as $M = \langle S, X, \mathcal{V}, \{R_t\}_{t \leq T}, \{\succcurlyeq_s\}_{s \in S}, \{\mathcal{A}_t\}_{t \leq T} \rangle$, where previously introduce components are as they were before and $\mathcal{A}_t : S \to 2^{\mathcal{L}(X)}$ indicates the set of propositions of which the DM is aware in state s at time t. The interpretation is captured by the following semantics, which are added to the definition of $\models$, in Section 4.2.2:

$$(M, s) \models_\mu A_t\varphi \quad \text{iff } \bar{\mu}(\varphi) \in \mathcal{A}_t(s),$$
$$(M, s) \models_\mu E_t\varphi \quad \text{iff } (M, s) \models_\mu K_t\varphi \text{ and } (M, s) \models_\mu A_t\varphi.$$

At times, it will be helpful to speak of the truth of a formula in $\mathcal{L}(X)$, provided $M \in \mathcal{M}^X$. This will be represented by

$$(M, s) \models_X \bar{\varphi} \quad \text{iff } (M, s) \models_\mu \varphi, \text{ for some } (\mu, \varphi) \text{ such that } \bar{\mu}(\varphi) = \varphi.$$

Lemma's 8 and 9 guarantee $\models_X$ is well defined. Indeed, by Lemma 8, there always exists a (μ, φ) that stands in the correct relation to $\bar{\varphi}$ and by Lemma 9 the choice of such (μ, φ) does not impact the truth value.

Notice, if φ is a sentence in $\mathcal{L}^A(\mathcal{X})$, then $(M, s) \models_\mu \varphi$ implies $(M, s) \models \varphi$ and $(M, s) \models_X \varphi$ (where $\varphi \in \mathcal{L}(X)$ since it has no free variables). This is well known in general first order modal logic, but it worth pointing out, the same holds in the presence of awareness and explicit knowledge modalities. To see this note the only case not covered by previous arguments is if $\varphi = A_t\psi$. Since φ is a sentence, so too must be ψ. But then, $\bar{\mu}(\psi) = \psi$ for all reduced assignments $\bar{\mu}$, and so, the truth of $A_t\psi$ does not depend on the interpretation.

Fagin and Halpern (1988) showed this semantic structure of unawareness is captured axiomatically by

[AX.4.17: A0] $E_t\varphi \iff (K\varphi \wedge A_t\varphi)$.

Axiom A0 simply states that explicit knowledge is the conjunction of implicit knowledge and unawareness. Even with first order models and semantic awareness, this is in the only additional axiom needed for general awareness structures.

Proposition 4.4. $\forall \boldsymbol{K_T} \cup$ A0 *is a sound and complete axiomatization of* $\mathcal{L}^A(\mathcal{X})$ *with respect to* $\mathcal{M}(\mathcal{X})$.

Proof. In appendix C.2.

Take note, Lemma's 8 and 9 imply that $\forall \boldsymbol{K_T} \cup$ A0 is also a complete and sound axiomatization of $\mathcal{L}(X)$ with respect to $\mathcal{M}^X$.

4.3.3 Explicit Preference

Just as implicit knowledge gave rise to the DM's implicitly known preference, explicit knowledge will likewise define her explicitly known preferences. As a prerequisite, define the preferences of which the DM is aware.

Definition. *Let* M *be a model of decision making. Define*

$$\succcurlyeq_{A_t,s} = \{(x,y)|(x \succcurlyeq y) \in \mathcal{A}_s(t)\} \tag{4.5}$$

Just as explicit knowledge is the intersection of knowledge and awareness, explicit preference is the intersection of implicit preference and awareness.

Definition. *Let* M *be a model of decision making. Define the explicit preference relation, as*

$$\succcurlyeq_{E_t,s} = \succcurlyeq_{A_t,s} \bigcap \succcurlyeq_{K_t,s} \tag{4.6}$$

Remark 4.5. $x \succcurlyeq_{E_t,s} y$ *if and only if* $(M,s) \models_\mu E_t(a \succcurlyeq b)$ *holds for any* μ *such that* $\mu(a) = x$ *and* $\mu(b) = y$.

Just as the DM could have a true preference and not implicitly know it, if $(M,s) \models_X K_t(x \succcurlyeq y) \wedge \neg A_t(a \succcurlyeq y)$, then at time t, in state s, the DM implicitly knows her preference but is unaware of it, and hence does not explicitly know it. As argued earlier implicit knowledge is necessary to act upon a preference. Similarly, when awareness is taken into account, the DM must *explicitly* know her preference in order to act in accordance. The major contribution of this paper, and what is meant as the behavioral implications of unawareness, is the contrast between $\succcurlyeq_{K_t,s}$ and $\succcurlyeq_{E_t,s}$.

Without further restrictions the sets $\succcurlyeq_{E_t,s}$ can any arbitrary subset of $\succcurlyeq_{K_t,s}$. However, given the intent to describe a DM's preferences under unawareness, there are several natural assumptions on the structure of awareness. At this stage, I wish to examine the behavioral implications of a DM who is unaware of some outcomes and contingencies, but acts as rationally as possible with respect to those of which she is aware. In Section 4.1.2, without the underlying epistemic framework, LCmp stated the DM was aware of some $Y \subseteq X$, over which she had complete preferences. A similar restriction can be made in this environment. For example, if the DM is aware of a preference ranking x weakly to y, $A_t(x \succcurlyeq y)$, it seems only natural she is also aware of the preference ranking y weakly to x. Likewise, if the DM is aware of rankings over x and y, and over z and w, she ought to be aware of rankings over x and z.

[AX.4.18: A1] $\forall a \forall b (A_t(a \succcurlyeq b) \iff (A_t(a \succcurlyeq a) \wedge A_t(b \succcurlyeq b)))$.

This restriction does not say anything about the DM's actual or implicitly known preferences –only about her awareness. Nonetheless, there is an intuitive connection between A1 and LCmp: there exists some $Y \subseteq X$, such that the DM is aware of all preference relations among Y, and none outside of Y. This connection can be seen by the following result:

Proposition 4.6. *Let M be a model of $\forall \boldsymbol{K} \cup \boldsymbol{P} \cup$ A0 $\cup$ A1, then $\succcurlyeq_{E_t,s}$ satisfies* LRfx *and* Trv*.*

Proof. Let $(x, y) \in \succcurlyeq_{E_t,s}$. Then $(x, y) \in \succcurlyeq_{A_t,s}$: $(M, s) \models_X A_t(x \succcurlyeq y)$. So by A1, $(M, s) \models_X A_t(x \succcurlyeq x)$, or, $(x, x) \in \succcurlyeq_{A_t,s}$. By Remark 4.3, $(x, x) \in \succcurlyeq_{K_t,s}$. So, $(x, x) \in \succcurlyeq_{A_t,s} \bigcap \succcurlyeq_{K_t,s} = \succcurlyeq_{E_t,s}$. The proof for transitivity similar.

4.3.4 The Structure of Awareness

By a similar motivation, if the DM is aware of the statement φ implies x is preferred to y, $A_t(\varphi \implies (x \succcurlyeq y))$, then reasonably she should be aware of the statement φ and the preference $(x \succcurlyeq y)$. This is captured by the notion that $\mathcal{A}$ is closed under "subformulae".

[AX.4.19: A↓] $$\begin{cases} \big((A_t\neg\varphi) \vee (A_t\Box\varphi) \vee (A_t(\varphi \wedge \psi)) \vee (A_t(\psi \wedge \varphi))\big) \implies A_t\varphi \\ A_t\forall a\varphi \implies \neg\forall b \neg A_t\varphi[[a/b]] \end{cases}$$

A↓ specifies the DM is able to extrapolate downwards: if the DM is aware of φ, her understanding is full enough to contemplate on the constituent parts of φ. A↓ is not exactly closure under subformulae, since this would entail $A_t \forall a\varphi \implies \forall b A_t \varphi[[a/b]]$: the DM is aware of every instance embodied by $\forall a\varphi$. This is clearly too strong, it implies the DM who is aware her preferences are reflexive, is aware of every object of consumption. To accommodate this subtly, the quantified case is relaxed so that the DM only need be aware of some instance embodied by $\forall a\varphi$. There are two notable consequences of A↓. First, the DM's explicit knowledge is closed under implication: $(E_t\varphi \wedge E_t(\varphi \implies \psi)) \implies E_t\psi$. Second, the DM is unaware of what she is unaware of.

The converse of A↓ can also be dictated.

[AX.4.20: A↑] $(A_t\varphi \wedge A_t\psi) \implies \big((A_t\neg\varphi) \wedge (A_t\Box\varphi) \wedge (A_t(\varphi\wedge\psi)) \wedge (A_t(\psi\wedge\varphi)) \wedge A_t\forall a\varphi[[b/a]]\big)$

A↑ specifies the DM is able to extrapolate upwards: if the DM is aware of all of the constituent parts of φ, her understanding is full enough to be aware of φ itself. While A↓ does not seem to be a particularly restrictive notion, A↑ is slightly more controversial as it implies that the DM who is aware of anything is aware of formula of arbitrary complexity.

As in the propositional case, if both A↓ and A↑ are jointly satisfied then $\mathcal{A}_t$ is generated by a set of primitive propositions (i.e., generated according to the construction rules of $\mathcal{L}(X)$). The unawareness of DM who satisfies A↓ and A↑ does not arise from cognitive limitations (she is aware of all of the logical entailments of her own awareness). Instead, her unawareness regards statements completely disjoint from her current world view, things she has *never heard of*.

Lastly, consider the axiom that dictates the DM knows what she is aware of,

[AX.4.21: KA] $\nabla_t(A_t\varphi)$.

It is well known that KA distinguishes models in which $\mathcal{A}_t(s) = \mathcal{A}_t(s')$ for all $s' \in R_t(s)$. The interpretation here is that the DM can always (implicitly) reflect on her on awareness, and therefore distinguish between states in which her awareness differs. Of course, this delineation does not necessarily occur at the explicit knowledge level.

Finally, let $\boldsymbol{A}$ denote the axiom system A0 ∪ A↓ ∪ A↑ and $\boldsymbol{A}^\star$ as $\boldsymbol{A}$ ∪ A1 ∪ KA.

4.4 CONTINGENT PLANNING, OR, FINALLY, THE BEHAVIORAL IMPLICATION OF AWARENESS OF UNAWARENESS

With the requisite foundational matters taken care of, I turn to delineating the behavioral effects of incorporating awareness structures. Because I would like to talk about maximality of contingent plans, assume "utilities" are bounded from above by some outcome. This is captured axiomatically by

[AX.4.22: BND] $\exists a \forall b (a \succcurlyeq b)$.

and, when the DM's preferences is to be bounded when considering the restriction to her awareness set,

[AX.4.23: ABND] $\exists a \forall b (A_t(a \succcurlyeq a) \wedge (A_t(b \succcurlyeq b) \implies (a \succcurlyeq b)))$.

These axioms ensure the decision maker, when fully informed about her preference, would have a optimal consumption choice –so any lack of maximality is not arising from a lack of closure. Notice, also, in any model of full awareness, BND and ABND are materially equivalent.

4.4.1 Implicitly Known Contingent Plans

It is in this set up that I show awareness of unawareness produces a natural incompleteness in rankings over contingent plans, in particular, even when states are fully contractable. A mapping $(\bar{\varphi} \mapsto x)_t$, where $\bar{\varphi} \in \mathcal{L}(X)$ and $x \in X$, is the commitment to consume x in period t if $\bar{\varphi}$ is true. This is a *partial contingent plan*, since it does not specify what happens in states where $\bar{\varphi}$ is not true. Regardless, the following dictates when a rational DM would refuse accept committing to $(\bar{\varphi} \mapsto x)_t$.

Definition. *A partial contingent plan, $(\bar{\varphi} \mapsto x)_t$, is **implicitly unacceptable** to a decision maker, M, in state s, if*

$$(M, s) \models_\mu P_0\big(\varphi \wedge \exists a K_t(a \succ b)\big), \tag{4.7}$$

whenever $\bar{\mu}(\varphi) = \bar{\varphi}$ and $\bar{\mu}(b) = x$.

Parsing the above definition, a DM is unwilling to commit to $(\bar{\varphi} \mapsto x)_t$ if at time 0 the decision maker considers it possible $\bar{\varphi}$ is true and by time t, she will know the existence of an object y, such that y is preferred to x. There are three relevant outcomes at time t: (1) $\bar{\varphi}$ is not true, so her commitment does not bind, (2) $\bar{\varphi}$ is true but she does not know any object she prefers to x, or, (3) $\bar{\varphi}$ is true and she knows an object she prefers to x. So, the DM is unwilling to commit to $(\bar{\varphi} \mapsto x)_t$ if she believes it is possible (3) might occur – by waiting until time t and then making a decision she could be strictly better off. On the other hand, say a contingent plan is acceptable if it is not unacceptable.

Definition. *A partial contingent plan, $(\bar{\varphi} \mapsto x)_t$, is **implicitly acceptable** to a decision maker, M, in state s, if*

$$(M, s) \models_\mu K_0\big(\varphi \implies \forall a P_0(b \succcurlyeq a)\big) \tag{4.8}$$

whenever $\bar{\mu}(\varphi) = \bar{\varphi}$ and $\bar{\mu}(b) = x$.

Parsing this formula, a DM is willing to commit if she knows at time 0, whenever $\bar{\varphi}$ is true she will consider it possible at time t that x is preferred to any other element. In other words, she knows she will not know φ is true and an object y is preferred to x. If the DM knows $\neg\varphi$ then $(\varphi \implies \forall a P_0(b \succcurlyeq a))$ is always true and the DM finds the partial contingent plan acceptable.

Stringing together these partial contingent plans can produce a complete contingent plan. Of course, to be well defined the formulae ought to partition the state space.

Definition. *A finite set of formulae, $\Lambda \subset \mathcal{L}(X)$, is a **contractable set**, or simply contractable, for some class of models, $\mathcal{N} \subseteq \mathcal{M}^X$, if*

1. *$\bigvee_{\bar{\varphi} \in \Lambda} \bar{\varphi}$ is valid in $\mathcal{N}$, and,*
2. *for any distinct $\bar{\varphi}, \bar{\psi} \in \Lambda$, $\bar{\varphi} \wedge \bar{\psi}$ is unsatisfiable in $\mathcal{N}$.*

Sometimes a contractable set will be referred to as (Λ, Γ, μ), under the acknowledgement that Λ is contractable, $\Gamma \subset \mathcal{L}(\mathcal{X})$ and $\bar{\mu}$ defines a bijection between Λ and Γ. An obvious example of the basis for a contingent plan is $\{\bar{\varphi}, \neg\bar{\varphi}\} \subset \mathcal{L}(X)$. Using Λ as a basis, a *period t contingent plan*, c_t, is a mapping from $c_t : \Lambda \to X$. We can define implicit acceptability and unacceptability of complete contingent plans by similar conditions as above.

Definition. *A contingent plan,* $c_t : \Lambda \to X$*, based on* (Λ, Γ, μ)*, is **implicitly acceptable** to a decision maker,* M*, in state* s*, if*

$$(M,s) \models_\mu K_0 \bigwedge_{\varphi \in \Gamma} \Big(\varphi \implies \forall a P_t\big(c_t(\varphi) \succcurlyeq a\big)\Big) \tag{4.9}$$

whenever $\bar{\mu}(\varphi) = \bar{\varphi}$ *and* $\bar{\mu}(c(\varphi)) = c(\bar{\varphi})$*, and provided there is no free occurrence of* a *in* φ*, for all* $\varphi \in \Lambda$*. It is **implicitly unacceptable**, if, under the same conditions,*

$$(M,s) \models_\mu P_0 \bigvee_{\varphi \in \Gamma} \Big(\varphi \wedge \exists a K_t\big(a \succ c_t(\varphi)\big)\Big). \tag{4.10}$$

Intuitively, a contingent plan is acceptable if it provides outcomes that are no worse than what could have been selected by the DM had she waited until time t and then made a decision. It is unacceptable if this is not the case. Lemma 10 shows, in any model of $\forall \boldsymbol{K_T} \cup$ Cmp, in a given state, every contingent plan is either acceptable or unacceptable (logically, (4.9) and (4.10) are material negations is such models).

Of course, for a contingent plan to be deemed acceptable, the set of formulae on which it is based must be rich enough so that it could mimic any decision making process the DM could have implemented without a contingent plan (i.e., by waiting until t and making a single decision). If this condition is not met by a set Λ, then there may not exist any acceptable contingent plan. Before formalizing this idea, consider the following example.

Example 4.3. *Let* $T = 1$*. There are two states of the world* $S = \{s, s'\}$*. The domain consists of 2 distinct objects,* $X = \{x, y\}$*. The preference relations, as induced by* $\mathcal{V}(\succcurlyeq)$*, in each state are given by* $\succcurlyeq_s = \{(x,x),(x,y),(y,y)\}$ *and* $\succcurlyeq_{s'} = \{(x,x),(y,x),(y,y)\}$*. The accessibility relations are* $R_0 = S^2$*, and* $R_1 = \{(s,s),(s',s')\}$*; the DM is initially uncertain about the state but will be able to distinguish in period 1. It is easily checked that this is a model of* $\boldsymbol{S5} \cup \boldsymbol{P} \cup \forall \cup$ Bnd $\cup$ F $\cup\, S_1$*.*

For any $\bar{\varphi} \in \mathcal{L}(\Delta)$ *let* $\Lambda = \bar{\psi} = (\bar{\varphi} \vee \neg\bar{\varphi})$*. Let* μ *be any assignment such that* $\bar{\mu}(\psi) = \bar{\psi}$ *and* $\bar{\mu}(b) = x$*. Since,* ψ *is a tautology, and there are no two distinct elements of* Λ*,* Λ *is* *contractable. There are only two contingent plans that can be based on* Λ, $\bar{\psi} \mapsto x$, *and* $\bar{\psi} \mapsto y$. *Now, notice*

$$(M, s') \models_{\mu'} (a \succ b),$$

for the a-variant of μ *where* $\mu'(a) = y$. *Since,* $R_t(s') = \{s'\}$,

$$(M, s') \models_{\mu'} K_t(a \succ b).$$

So, by definition, (and the fact that φ *is tautological),*

$$(M, s') \models_{\mu} \varphi \wedge \exists a K_t(a \succ b).$$

Now and finally, since $s' \in R_0(s'')$ *for all* $s'' \in S$,

$$(M, s') \models_{\mu} P_0\big(\varphi \wedge \exists a K_t(a \succ b)\big).$$

But since μ *was arbitrary,* $\bar{\psi} \mapsto x$ *is unacceptable in any state* $s'' \in S$. *Of course the same argument goes to show* $\bar{\psi} \mapsto y$ *is likewise unacceptable, and hence, there is no acceptable contingent plan.*

Example 4.4.1 raises the question as to the existence of restrictions on a contractable set Λ that will guarantee the existence of an acceptable contingent plan. In Example 4.4.1, no contingent plan allows the DM to consume different objects in different states, which of course, she would want to do. If, however, she waited until period 1, she would be able to make an informed decision and choose the optimal consumption. In order to ensure the existence of acceptable contingent plans, I will consider contractable sets that are rich enough that any decision making process the decision maker could implement without a contingent plan can reproduced with one.

Definition. *Fix a model* M. *Let* $s^{K_t} = \{\bar{\varphi} \in \mathcal{L}(X) | (M, s) \models_X K_t\bar{\varphi}\}$. *A contractable set* Λ *is* ***t-separable*** *if whenever* $s^{K_t} \neq s'^{K_t}$ *then* $(M, s) \models_X \bar{\varphi}$ *implies* $(M, s') \models_X \neg\bar{\varphi}$, *for any* $\varphi \in \Lambda$.

Theorem 4.7. *Let M be a model $\forall \boldsymbol{K_T} \cup \boldsymbol{P} \cup \mathrm{T} \cup \mathrm{BND}$. For any t-separable and contractable set Λ, there exists a contingent plan, c_t, based on Λ that is acceptable at every s.*

Proof. In appendix C.3.

Theorem 4.7 states that so long as the contractable set is rich enough to allow the DM to make decisions with at least as much distinction as if she had waited until time t, then she will be willing to commit to something. The intuition is simple: there exists a contingent plan that can implement her optimal choice behavior in time t, so the optimal contingent plan must be weakly better. But since the optimal contingent plan weakly out preforms waiting, she is willing to commit to it.

Without further qualification, Theorem 4.7 does not help to establish the existence of an optimal contingent plan since it requires the existence of a t-separable contingent plan, itself not guaranteed to exist. The next result mollifies this concern, at least in finite models.

Proposition 4.8. *Let M be a model of $\forall \boldsymbol{K_T} \cup \boldsymbol{P} \cup \mathrm{T} \cup \mathrm{BND}$ with a finite state space, S. Then there exists a t-separable contractable set containing $|S|$ formulae.*

Proof. In appendix C.3.

Putting these two results together, provides a behavioral characterization for full awareness regarding the acceptability of contingent plans.

Corollary 4.9. *Let M be a model of $\forall \boldsymbol{K_T} \cup \boldsymbol{P} \cup \mathrm{T} \cup \mathrm{BND}$ with a finite state space, S. Then for all t, there exists a contractable set $\Lambda(t)$ and a contingent plan there based, c_t, that is acceptable in every $s \in S$.*

While t-separability is sufficient to ensure existence, it is not necessary. For example, consider the case where preferences are constant across states, and so, fully known. Then, even if Λ is a single formula, an acceptable contingent plan exists. This indicates that perhaps the notion of t-separability can be weakened to provide a necessary condition. Indeed, this can be done,[35] but the resulting restriction is more convoluted than insightful. Further, in light of Proposition 4.8, this additional complexity is superfluous. In particular, as it is

[35]Specifically, if two states, s and s', can be separated via a contingent plan whenever the set of maximal elements of $\succcurlyeq_{K_t,s}$ and $\succcurlyeq_{K_t,s'}$ have an empty intersection.

my focus to show explicit acceptability is harder to satisfy than its implicit counterpart, it suffices to provide any weak condition that provides existence of implicitly acceptable contingent plans in a large class of models.

4.4.2 Explicitly Known Contingent Plans

The takeaway from the previous section is that if the decision maker is acting according to her implicitly known preference, then she is always willing to commit to *something*. There exists an acceptable contingent plan based on any t-separable set of formulae, itself guaranteed to exist in finite models. The intuition is clear: if a contingent plan can be written in a sufficiently flexible way (i.e., such that it will allow the decision maker to use all available information) there is no reason not to commit. Of course, this line of reasoning relies on the dictate the DM knows what it is she might learn. In other words, the contingent plan allows the DM to specify consumption in the event she learns a particular piece of information, and so it is requisite she know (at time the contingent plan is written) every piece of information she might learn. This is markedly impossible in the event she is unaware of things, and aware of her unawareness!

This is, finally, the behavioral implication of unawareness: the unwillingness to commit to *any* contingent plan, even under circumstances that make knowledge very well behaved. The acceptance of a contingent plan is given by the following, in analogy to (4.9) and (4.10):

Definition. *A contingent plan, $c_t : \Lambda \to X$, based on (Λ, Γ, μ), is **explicitly acceptable** to a decision maker, M, in state s, if*

$$(M, s) \models_\mu E_0 \bigwedge_{\varphi \in \Gamma} \Big(\varphi \implies \forall a P_t^E\big(c_t(\varphi) \succcurlyeq a\big)\Big) \tag{4.11}$$

*whenever $\bar{\mu}(\varphi) = \bar{\varphi}$ and $\bar{\mu}(c(\varphi)) = c(\bar{\varphi})$, and provided there is no free occurrence of a in φ, for all $\varphi \in \Lambda$. It is **explicitly unacceptable**, if, under the same conditions,*

$$(M, s) \models_\mu P_0^E \bigvee_{\varphi \in \Gamma} \Big(\varphi \wedge \exists a E_t\big(a \succ c_t(\varphi)\big)\Big). \tag{4.12}$$

With implicit knowledge, DM a contact is either acceptable or unacceptable, as shown in Lemma 10. With explicit knowledge, the relation is weaker. If a contingent plan is acceptable it is not unacceptable, and if a contingent plan is unacceptable it is not acceptable (each implies the negation of the other), but it is possible neither condition holds.[36] As such, it is not enough to show a contingent plan is not acceptable, instead, it is required to show (4.12) holds.

There exists one additional concern when relaxing full awareness. It may be the DM is unaware of some aspect of the contingent plan itself, and therefore could not make reasonable choices regarding it. For instance, if the formulae on which the plan was based are not in the DM's awareness set. Therefore, it is necessary to impose that contingent plans are articulable to the DM. That is to say, the DM is aware of the constituent parts of the contingent plan (the contingencies, Λ, and the outcomes, $\mathrm{Im}(c)$).

Definition. *Given a DM, M, a contingent plan $c : \Lambda \to X$ is* ***articulable*** *(in state s) if*

1. $\Lambda \subseteq \mathcal{A}_0(s)$.
2. $(x \succcurlyeq y) \in \mathcal{A}_0(s)$ *for any* $x, y \in \mathrm{Im}(c)$.

If a contingent plan is articulable, then DM is able to conceive of it at the time when she would have to commit. If a DM was asked to report her set of potential contingent plans, these are the contingent plans which she would be able to articulate; the contingent plans that are constructible given the language of which she is currently aware. Under full awareness, every contingent plan is articulable. As stated in the introduction, the behavioral characterization relies only on rankings over objects the DM can articulate, and so I will restrict my attention to articulable contingent plans.

The following example shows even under very well behaved knowledge, the existence of unawareness can render *every* articulable contingent plan unacceptable. Of course, this trivially holds when $A_0 = \emptyset$ (since there is no explicit knowledge *and* no articulable contingent plans). Hence, to make a meaningful claim, I am obliged to show something more strict:

[36] In particular, this arises when a contingent plan is implicitly acceptable (therefore not explicitly unacceptable), but the DM is unaware of the necessary statements for (4.11) to hold. This misalignment is a consequence of the fact that the awareness modality, A, does not respect material equivalence (i.e., $\varphi \iff \psi$ is a theorem does not imply $A_t\varphi \iff \mathcal{A}_t\psi$) is a theorem). Of course, this property is potentially desirable as it is a very natural relaxation of logical omniscience.

even when articulable contingent plans exist, they may all be unacceptable and the DM will explicitly know this!

Remark 4.10. *There exist models of* $\boldsymbol{S5} \cup \boldsymbol{PA}^\star \cup$ Bnd*, that admit articulable contingent plans and such that the DM explicitly knows every articulable contingent plan is explicitly unacceptable.*

Proof. Example 4.4.2.

Example 4.4. *Let* S *be a single state,* s*, assessable from itself in every period. Let* $X = \{x, y, z\}$*, with* $\succcurlyeq_s$ *defined by* $z \succ_s y \succ_s x$*.* $\mathcal{A}_0$ *be the closure (under the construction rules for* $\mathcal{L}(X)$*) of* $\{\{x, y\}^2\}$*,* $\mathcal{A}_t$ *be the closure of* $\{\{x, y, z\}^2\}$*. Let* $\bar{\varphi}$ *be any statement of the form* $(\bar{\psi} \vee \neg\bar{\psi})$ *for some* $\bar{\psi} \in \mathcal{A}_0$*. This defines a model of* $\boldsymbol{S5} \cup \boldsymbol{P} \cup \forall \cup \boldsymbol{A}^\star \cup$ Bnd $\cup$ F $\cup$ S_t *(note:* F *and* S_t *are defined in Section 4.5).*

Note, it is without loss of generality to consider contingent plans of the form $\bar{\varphi} \mapsto x$ *or* $\bar{\varphi} \mapsto y$*; since there is only one state there will always be a unique element of the contractable set that is satisfied in the model. Fix,* (Λ, Γ, μ) *such that* $\bar{\mu}(\varphi) = \bar{\varphi}$ *and* a *does not occur free in* φ*. And denote by* a, b, c *variables such that* $\bar{\mu}(a) = x$*,* $\bar{\mu}(b) = y$*, and* $\bar{\mu}(c) = z$*. We will show* $\bar{\varphi} \mapsto y$ *is unacceptable (the argument for* $\bar{\varphi} \mapsto x$ *is the same).*

By assumption $(z \succ y) \in A_t(s)$*:*

$$(M, s) \models_\mu A_t(c \succ b)). \tag{4.13}$$

Moreover, since both $\bar{\varphi}$ *and* $(x \succ y)$ *are in* $A_0 \subset A_t$*,* $P_0^E(\bar{\varphi} \wedge \exists a E_t(a \succ y))$ *is also in* $A_0 \subset A_t$*:*

$$(M, s) \models_\mu A_0 P_0^E(\varphi \wedge \exists a E_t(a \succ b)). \tag{4.14}$$

By assumption $(M, s) \models_\mu (c \succ b)$*. Since* $R_t(s) = \{s\}$*, this implies,* $(M, s) \models_\mu K_t(c \succ b)$*. Combining this with (4.13), we have* $(M, s) \models_\mu E_t(c \succ b)$*. This means, for any* $\mu' \sim_a \mu$ *such that* $\mu'(a) = \mu(c)$*,* $(M, s) \models_{\mu'} E_t(a \succ b)$*. By definition,*

$$(M, s) \models_\mu \exists a E_t(a \succ b).$$

Applying the fact $R_t(s) = \{s\}$ twice, (and, φ is true at s) implies

$$(M, s) \models_\mu K_0 P_0^E\big(\varphi \wedge \exists a E_t(a \succ b)\big). \tag{4.15}$$

Combining (4.14) and (4.15),

$$(M, s) \models_\mu E_0 P_0^E\big(\varphi \wedge \exists a E_t(a \succ b)\big).$$

as desired.

Example 4.4.2 shows once awareness is introduced, there is no longer a guarantee of acceptability. The introduction of unawareness (and subsequent focus on explicit rather than implicit acceptability) has fundamentally changed the behavior of the DM –creating a preference for delay that cannot be assuaged by allowing the DM to make conditional decisions. And, by Proposition 4.8, this behavior, unlike incompleteness or a preference for flexibility, cannot be explained in a framework with full awareness, no matter how much uncertainty exists. It is a behavioral trait that indicates the presence of unawareness.

Recall, unacceptability refers here only to articulable contingent plans. In the example, all of the conditions for Proposition 4.8 are met; there exists an implicitly acceptable contingent plan, namely $\bar{\varphi} \mapsto z$. However, this contact is inarticulable because the DM is unaware of z. If the modeler were to ask the DM about $\bar{\varphi} \mapsto z$, she would, after some reflection, be willing to accept it. Of course, in doing so, the modeler would have *made* her aware of z, thereby changing the structure of the very entity whose identification is of interest. It is this subtlety that motivates the departure from the standard revealed preference framework.

As informally argued in the introduction, an unwillingness to commit to any articulable contingent plan is the result of a language that is not rich enough to specify the optimal contingent plan (unawareness), but is rich enough to articulate this fact (awareness of unawareness). The following results formalize this intuition.

Theorem 4.11. *Let M be a model of $\forall \boldsymbol{K_T} \cup \boldsymbol{P} \cup \boldsymbol{A}^\star \cup$ $\textsc{Bnd}$, such that $(x \succcurlyeq x), (y \succcurlyeq y) \in \mathcal{A}_0(s)$ for some $x, y \in X$. Then if c, based on (Λ, Γ, μ) is an articulable and implicitly acceptable contingent plan in state s, it is explicitly acceptable in state s.*

Proof. In appendix C.3.

Theorem 4.11 states that if the DM can articulate an optimal contingent plan, she will find it explicitly acceptable. Since the contingent plan is implicitly acceptable, the consumption alternatives it specifies are optimal in each state. Moreover, since the DM is aware of these alternatives, she is also aware of the statements professing their optimality (by her ability to extrapolate from explicitly known statements, i.e., A↑). Putting these two facts together delivers the result. This result shows the given formulation of awareness (to wit, under $\boldsymbol{A}^{\star}$), places clear limits on how unaware a DM can be. Because of the structure of explicit knowledge, the DM cannot be explicitly uncertain if a consumption alternative, x, is optimal: if $\forall a(x \succcurlyeq a) \in \mathcal{A}$ and x is indeed optimal (at all s), the DM explicitly knows this. More generally, unawareness, captured as in this paper, does not allow the DM to question statements she implicitly knows. This limitation is discussed in more detail in Halpern and Rêgo (2013).

Theorem 4.11 places an upper-bound on the DM's awareness such that she finds all articulable contingent plans explicitly unacceptable; Theorem 4.12 places the corresponding lower-bound.

Theorem 4.12. *Let M be a model of $\forall \boldsymbol{K_T} \cup \boldsymbol{P} \cup \mathrm{T} \cup \boldsymbol{A}^{\star} \cup \mathrm{ABnd}$ with a finite state space. Then if M admits articulable contingent plans in state s, and the DM explicitly knows every articulable contingent plan is unacceptable, the DM explicitly knows it is possible she is unaware. Specifically,*

$$(M, s) \models E_0 P_0^E \left(\exists a \left(\neg A_0(a \succcurlyeq a) \wedge A_t(a \succcurlyeq a)\right)\right).$$

Proof. In appendix C.3.

A DM cannot be so unaware she is not even aware waiting will afford her a more complete world view. That is, the DM must be introspectively unaware. The intuition of this result is straightforward. The DM, acting on explicit knowledge, must explicitly know all contingent plans are unacceptable and this requires she is aware she will have more choices if she does not commit.

4.5 A PREFERENCE FOR FLEXIBILITY

One interpretation of Kreps (1979) is the anticipation of learning induces a *preference for flexibility.* That is, the DM's preference over menus (i.e, subsets of X), respects set inclusion: if $m' \subseteq m \subseteq X$ then m is preferred to m'. A DM who expects to learn her true preference, but is currently uncertain, will prefer the flexibility to make choices contingent on the information she learns. The Krepsian model has a clear connection to the notion of contingent planning (a menu is a restriction on which contingent plans are feasible) as well as more generally to the epistemic framework where the anticipation of learning can be defined precisely. In this section, I will show that the Krepsian framework can be faithfully reproduced as a special case of the general model outlined above. In particular, this special case is one of full awareness; as such, the unforeseen contingencies interpretation is not strictly needed, and a preference for flexibility in not alone the behavioral indication of unawareness.

To reproduce the anticipation of learning, two axioms are needed. First,

[AX.4.24: F] $K_t\varphi \implies K_{t'}\varphi$ *for all* $t' \geq t$.

The restriction F captures learning by ensuring the decision maker knows (weakly) more at later time periods.

Proposition 4.13. $\forall \boldsymbol{K_T} \cup$ F *is a sound and complete axiomatization of* $\mathcal{L}^A(\mathcal{X})$ *with respect to the class of models such that* $R_{t'} \subseteq R_t$ *for all* $t' \geq t$ *(denoted* $\mathcal{M}^f$*). Moreover, in the* $\boldsymbol{S}$**5** *framework,* $\{R_t\}_{t\in T}$ *is a filtration of* S.

Proof. In appendix C.2.

Second, that all uncertainty will be realized by time t,

[AX.4.25: S_t] $\nabla_t(x \succcurlyeq y)$.

Axiom S_t dictates the DM implicitly knows her preference at time t. From a modeling perspective S_t corresponds to the class of models such that $\succcurlyeq_s = \succcurlyeq_{s'}$ if $s' \in R_t(s)$.

Remark 4.14. *Let* M *be a model of* $\forall \boldsymbol{K_T} \cup \boldsymbol{P} \cup S_t$*. The implicit preferences in time* t*,* $\{\succcurlyeq_{K_t,s}\}_{s\in S}$*, are complete. Moreover, if* M *satisfies* A0 $\cup$ A1*, then* $\succcurlyeq_{E_t,s}$ *satisfies* LCmp*.*

Proof. By Proposition 4.1, $\succcurlyeq_s$ is a preference relation. By S_t, $\succcurlyeq_{K_t,s} = \succcurlyeq_s$, and so trivially inherits completeness. Now, let $x \succcurlyeq_{E_t,s} y$ or $y \succcurlyeq_{E_t,s} x$. By A0, this implies $(M,s) \models_X A_t(x \succcurlyeq y) \vee A_t(x \succcurlyeq y)$, and so, by A1, $(M,s) \models_X A_t(x \succcurlyeq x) \wedge A_t(y \succcurlyeq y)$. Moreover, since $\succcurlyeq_{K_t,s}$ is reflexive, $(M,s) \models_X K_t(x \succcurlyeq z) \wedge K_t(y \succcurlyeq u)$. Finally, this implies, $(M,s) \models_X E_t(x \succcurlyeq x) \wedge E_t(y \succcurlyeq y)$, as desired. Conversely, let $x \succcurlyeq_{E_t,s} x$ and $y \succcurlyeq_{E_t,s} y$. Then by A1, $(M,s) \models_X A_t(x \succcurlyeq y)$, and, switching the roles of x and y, $(M,s) \models_X A_t(y \succcurlyeq x)$. The completeness of $\succcurlyeq_{K_t,s}$ delivers the result.

To see how axiom F$\cup S_t$ can generate a preference for flexibility consider the simple two period case; let $T = 1$. It is immediate that if $\succcurlyeq_{K_0,s}$ is complete for any s, it follows $\succcurlyeq_{K_1,s}$ is likewise complete. The interesting case is when $\succcurlyeq_{K_0,s}$ is incomplete. So let x, y be elements for which the DM does not have implicitly known preference at time 0 in state s. So:

$$\begin{aligned}
(M,s') &\models_X (x \succcurlyeq y),\\
(M,s'') &\models_X (y \succcurlyeq x),\\
M &\models_X K_1(x \succcurlyeq y) \vee K_1(y \succcurlyeq x),\\
M &\models_X K_0(K_1(x \succcurlyeq y) \vee K_1(y \succcurlyeq x)),
\end{aligned}$$

where $s', s'' \in R_0(s)$. The first two lines rely on Cmp, and the third from S_1 (and Cmp). The fourth line follows directly. This last line shows if the DM is going to learn her true preference by time 1, then at time 0, she knows she will learn her true preference (this is *not* the case once awareness is introduced). As such, she would like the option of choosing either x or y, contingent on what she learns. A preference for flexibility is not the product of learning alone, but requires also the DM acknowledge the possibility she will learn.[37] Now, consider the problem of a DM choosing a menu in period 0 to be the choice set in period t. If the DM knows in period 0 she will know in period t one element of m is preferred to every element in m' then she prefers m to m'. This behavior is captured by the definition of dominance.

[37]This sentiment is echoes the fact that unawareness alone in not sufficient to perturb behavior, and introspective unawareness is required. A preference for flexibility could be seen as the behavioral marker of introspective uncertainty.

Definition. *A menu, m, **s-dominates** a menu, m', (at period t), if and only if*

$$(M, s) \models_X K_0 \bigvee_{z \in m} K_t \bigwedge_{z' \in m'} (z, z'). \tag{4.16}$$

*Further, m **strictly** s-dominates m', if it dominates m' and m' does not dominate m.*

That is, m dominates m' if the DM knows no matter what state of affairs is the true state, she will choose out of the menu m rather than m'.[38] Of course, she does not need to know which element is the maximal one. For example:

Example 4.5. *Let $S = \{s, s'\}$ and $X = x, y, z$. Then, $x \succ_s y \succ_s z$ and $z \succ_{s'} y \succ_{s'} x$. $R_0 = S^2$ and $R_1 = \{(s, s), (s', s')\}$. M is a model of $\forall \boldsymbol{K_1} \cup \boldsymbol{P} \cup \mathrm{F} \cup S_1$. Notice, $\{x, z\}$ strictly s-dominates $\{y\}$. Indeed, $(M, s) \models K_1(x \succcurlyeq y)$ and $(M \succcurlyeq s') \models K_1(z \succcurlyeq y)$. So, $M \models K_0(K_1(x \succcurlyeq y) \vee K_1(z \succcurlyeq y))$. So, the DM knows, in the true state of affairs, either x or z is preferred to y, but does not know which preference is her true preference.*

The dominance relation is generally incomplete. So, to connect this statement to previous work on preference over menus, where the ranking is usually a weak ordering, s-dominance must be extended to a complete and transitive relation.

Definition. *A preference $\geq$ over menus (i.e., a subset of $2^X \times 2^X$) is **$\boldsymbol{FS_1}$-generated** if $\geq$ is complete and transitive and there exists some some model, M, of $\forall \boldsymbol{K_1} \cup \boldsymbol{P} \cup \mathrm{F} \cup S_1$, and some state s thereof, such that if m s-dominates m' then $m \geq m'$, and if m strictly s-dominates m' then $m > m'$.*

The following result shows that beginning with the s-dominance relation, generated by some epistemic model, and extending it to a weak order captures exactly the "preference for flexibility" described in Kreps (1979). In fact, the converse it also true: every Krepsian decision maker can be formulated as the extension of an s-dominance relation with respect to some epistemic model.

[38]Under the restriction S_t, the DM's knowledge (regarding her preference) in period t is exactly her preference, and so inclusion of K_t in (4.16) is superfluous. It is included for conceptual reasons; the DM must know her preference in order to act on it, so if she remains ignorant in period t, there is no reason to desire non-singleton menus. This becomes important to the analysis when S_t is relaxed.

Theorem 4.15. *The preference $\geq$ satisfies Kreps (1979)' axioms if and only if it is FS_1-generated. Moreover, for any FS_1-generated $\geq$, there exists a strictly increasing aggregator $\Gamma : \mathbb{R}^{|R_0(s)|} \to \mathbb{R}$ such that*

$$m \geq m' \iff \Gamma(\{\max_{x \in m} u_{s'}(x)\}_{u_{s'} \in \mathcal{U}_{s,0}}) \geq \Gamma(\{\max_{x \in m'} u_{s'}(x)\}_{u_{s'} \in \mathcal{U}_{s,0}}), \tag{4.17}$$

where $\mathcal{U}_{s,0}$ is as in (4.2).

Proof. In appendix C.3.

Notice that the dominance relation, projected onto singleton menus, produces a reflexive and transitive relation on X. In fact, in such models, $x \succcurlyeq_{K_{s,0}} y$ if and only if $\{x\}$ s-dominates $\{y\}$. Hence a preference for flexibility can be seen as a natural extension of multi-utility models.

Finally, I investigate the connection between implicit acceptability and the dominance relation over menus defined in Section 4.5. Intuitively, when the menu is thought of as the image of a contingent plan, then the contingent plan specifies how the DM will choose out of the menu. Under this interpretation, Theorem 4.16 the provides the connection between having a well defined preference over menus and being willing to accept a contingent plan.

Theorem 4.16. *Let M be a model of $\forall \boldsymbol{K_T} \cup \boldsymbol{P} \cup S_t \cup$ T $\cup$ F $\cup$ Bnd. A finite menu $m \subseteq X$ is not strictly s-dominated (with regards to time t consumption) if and only if it is the image of an acceptable (in state s) contingent plan, $c_t : \Lambda \to X$.*

Proof. In appendix C.3.

If a menu is undominated (according to the definition given by (4.16)), one must be able to construct an acceptable contingent plan from it, and if a contingent plan is acceptable it must induce an undominated menu. This formally establishes the behavior being captured by a preference for flexibility is exactly what facilitates a contingent plan as being acceptable (or not) in models of full awareness. In fact, this result can prove Corollary 4.9 for the case with finite X: since $m = X$ must be undominated it must contain the image of an acceptable contingent plan, and hence, such an object must exist. Conversely, this result suggests a well defined preference (i.e., weak order) over menus is only sensible if acceptable contingent plans exist. It is this last observation (as well as Theorem 4.15) that serves as motivation for moving away from a theory predicated on a preference over menus.

4.6 UNAWARENESS AND CONTRACTS

This section contains a simple example to show how the framework presented above could be used in applications. Assume there are two players: a *principal* (player p), who is offering a take-it-or-leave-it contract to an *agent* (player a). The model takes place in an interactive awareness structure

$$M = \langle S, X, \mathcal{V}, \{R_{t,i}\}_{t\leq T, i=1,2}, \{\succcurlyeq_{s,i}\}_{s\in S, i=1,2}, \{\mathcal{A}_{t,i}\}_{t\leq T, i=1,2}\rangle,$$

in which players' knowledge and awareness are defined over atomic statements and both their own and their opponents knowledge and awareness.[39] It is in this framework that I will show the principal has an incentive to conceal mutually beneficial information. The intuition being that, although certain novel actions are Pareto improving in every ex-post scenario, the agent will react to the discovery of novel information by becoming more sensitive to her own unawareness, hence increasing her aversion to commitment. In other words, the display of surprising outcomes indicates to the agent that the novel outcomes are more valuable than she previously thought; the added value to waiting (and taking an outside option) is greater than the value added by the novel outcome itself. Further, I will show that this incentive can naturally lead to the optimality of incomplete contracts.

If M is any interactive model of decision making under unawareness, define $Y_{t,s,i} \equiv \{x | (x \succcurlyeq_i x) \in \mathcal{A}_{s,i}(t)\} \subseteq X$, the set of outcomes player i is aware of at time t is state s. Further, to make matters simple and tractable, assume each player has a state dependent (expected) utility index from $u_{s,i} : X \to \mathbb{R}$ that represents $\succcurlyeq_{s,i}$, and $\mu_t \in \Delta(S)$, such that

$$\sum_{s\in S} \mu_{t,s,i}(s) u_{s,i}(x)$$

[39] It is clear that the same axiomatization will suffice, simply by adding additional indexes to the modalities.

is a completion of $\succcurlyeq_{K_{t,s,i}}$. From Remark 4.3, we know $\mu_{t,s,i}$ must put positive probability on s' if and only if $s' \in R_{t,i}(s)$. Lastly, let $\bar{u}_{s,i} \equiv u_{s,i}|_{Y_{t,s,i}}$. Proposition 4.6 guarantees that $\sum_{s \in S} \mu_{t,s,i}(s)\bar{u}_{s,i}(x)$ is a completion of $\succcurlyeq_{E_{t,s,i}}$ on $Y_{t,s,i}$.

The timing is as follows. In period 0, the principal offers the agent a continent plan to be executed in period 1, which I assume is simply a function from S to X. If the offer is rejected the the agent can take an outside offer, some action in $Q \subseteq X$. If the contract is accepted the principal's and agent's evaluations of contracts are based on their explicitly known preferences:

$$\sum_{s' \in S} \mu_{0,s,i}(s')\bar{u}_{s,i}(c(s'))$$

If the contract is rejected, p gets a utility of 0, and a gets a utility according to her outside option. Of course, to make our problem well defined, we also have quantify the agents perceived value of the outside option. In the case of full awareness (or, naive unawareness, as in Auster (2013)), the agent would have a well defined expected utility over the outside option. This is not the case with introspective unawareness, as the DM is aware of the possibility that waiting will afford novel actions. So consider a mapping $\delta_{t,s,i} : 2^{\mathcal{L}(X)} \to \mathbb{R}$ with the restriction that

$$\delta_{t,s,i}(\mathcal{A}_{t,i}(s)) \geq \sum_{s \in S} \max_{x \in Q \cap Y_{0,s,a}} \mu_{t,s,i}(s)\bar{u}_{s,i}(x), \tag{4.18}$$

which, by the characterization in section 4.4, holds with equality if the DM is fully aware or naively unaware. δ captures the DM's attitude towards unawareness, her perceived value to the objects that she is currently unaware of (also, implicitly, the likelihood of discovering these novel actions).[40]

I focus on the case where actions are verifiable and the principal is fully aware. As such, the principal's problem is simply to offer the acceptable contract that maximizes his payoff. That is, maximize his payoff subject to a participation constraint on behalf of the agent.

[40] δ is in principle elicitable via the method described in previous sections –the additional utility the DM would require to accept some contingent plan–although certainty not for counterfactual worlds. Nonetheless, behavioral and experimental data could easily capture the "shape" of δ. Also, under the interpretation of Q as a menu of actions to be chosen from in period 1, the lower bound for δ is the value such a menu (assuming awareness stays the same) as given by (4.17).

Moreover, we will assume the agent explicitly knows the principal is fully aware. However, and unlike prior application of awareness, the agent is introspectively unaware.

4.6.1 The Principal's Problem

Let $S = \{s_1, s_2\}$ and $X = \{x, y, z, w\}$, $T = \{0, 1\}$. Assume that neither agent knows the state in period 0, but both will know the state in period 1 (this is a model of $\boldsymbol{S5} \cup S_t \cup \mathrm{F} \cup \mathrm{B}{\small ND}$). Let $Q = X$. Assume that the belief in period 0, for both players and both states, is $(\frac{1}{2}, \frac{1}{2})$, and that utilities are given by,

$u_{s,p}$

	x	y	z	w
s_1	3	1	3	0
s_2	1	3	4	6

$u_{s,a}$

	x	y	z	w
s_1	3	1	3	6
s_2	1	3	4	0

Since we will only consider a limited number of awareness sets, let, for some $D \subseteq X$, $\mathcal{A}(D)$ denote the awareness structure in which the agent is aware only (and all) statements that contain objects in D. In period 0, let $\delta_0(\mathcal{A}(D)) = 3$ if $z, w \notin D$ and $\delta_0(\mathcal{A}(D)) = 5$ otherwise. Since the game ends after period 1, δ_1 is given by (4.18), holding with equality. Consider the case where the agent's initial awareness is $\mathcal{A}(\{x, y\})$ and, if left unperturbed by the principal's offer, it remains her awareness in period 1. That is, without any outside influence, the agent does not become aware of any new objects. However, if the principal offers the contract c, then the agent's awareness set in period 1 is $\mathcal{A}(\{x, y\} \cup Im(C))$.

What is the principal's optimal strategy, given that he is constrained to offer complete contracts? Notice that if he offers the contract $c = (x, y)$ (read: x in state 1, y in state 2), it is accepted and the expected utility (for both players) is 3. It is accepted because the agent is indifferent between accepting and rejecting, given that her perception of the value of the outside option is also 3. It is easy to verify that this is the best the principal can do. To see this, note that if the principal offers a contract containing either z or w, he must provide the agent with a utility of at least 5 (the agent's new value participation constraint). Clearly, this is only possible with the contract $c = (w, z)$. But this gives the principal an expected utility of 2, worse than $c = (x, y)$.

Nonetheless, the contract $c = (x, z)$ makes both players strictly better off, providing a utility of 3.5. Hence, when the principal is constrained to offer a compete contract, he willingly conceals a Pareto improving contract. The intuition is simple: expanding the agent's awareness makes her more aware of her own awareness, and hence she displays a larger aversion to commitment. This second effect outweighs the first, so the principal chooses not to disclose the information.

Now, consider the case where the principal can offer an incomplete contract. Such a contract does not provide any alternative for a particular state, upon the realization of which the players renegotiate. Now the principle can offer the contract $c = (x, \cdot)$ (read: x in state 1, re-negotiate in state 2). This is acceptable to the agent, since $\delta_{0,s_1} = 3 = u_{s_1,a}(x)$. In period 1, if state s_2 is realized, the principal offers the new contract $c = z$. This is again acceptable since $\delta_{1,s_2} = 4 = u_{s_2,a}(z)$. Therefore by appealing to incomplete contracts, the principal can implement his unconstrained optimal contract.

4.6.2 Incompleteness Requires Introspection

The above example, while highly stylized, is indicative of a general phenomena. Although the effect of unawareness can be quantified via δ, and delay can be calculated, unawareness introduces behavior that intrinsically different than uncertainty. Unlike in the more standard framework, the value of delay (i.e., the outside option) changes with the agents epistemic state, and therefore is itself a function of the contract being offered. As such, there may exist feasible contracts which are initially individually rational, but cease to be so when offered. It is this effect, driven by introspective unawareness, that can make incompleteness strictly beneficial.

Remark 4.17. *Assume the principal is fully aware. Further assume that contracts take the form $S \to X \setminus Q$ (that is, there are a distinct set of outside options). If there is an incomplete contract c that is strictly better for the principal than any complete contract, then the agent is introspectively unaware.*

It is worth briefly addressing the relation between this environment and previous work connecting awareness with incomplete contracts. There is a large body of literature on incomplete contracts arising from the indescribably of states, leading to the well known discussion of Maskin and Tirole (1999). They show show, so long as players understand the utility consequences of states, indescribability should not matter. Here, of course, the unawareness is in the domain of actions rather than states, and so directly and intrinsically regards the understanding of utility. In other words, the utility of the outside option is defined in an articulable manner. Incompleteness, even in the single agent case (i.e., over contingent plans) is welfare improving.

More specifically, when the stipulation that $Q \cap Im(c) = \emptyset$ is dropped, naive awareness can also induce the optimality of incomplete contracts. There, the principal may withhold information strategically, as the novel outcomes may be of direct value as an outside option. This is similar in spirit to the arguments put forth in Filiz-Ozbay (2012) and Auster (2013), where the agents are naively unaware. However, leaving in the stipulation, we see that the discovery of novel outcomes can effect the value of the outside option *indirectly*.

One obviously missing component from the above example, is how the agent's perception of unawareness reacts to the offer of an incomplete contract. It is reasonable to assume the agent believes when such a contract is offered, it must be due to to strategic concerns relating to options outside if her current awareness. Hence the offer of an incomplete contract is itself reason to change her perception of the value of delay. This effect cannot be captured at all by naive unawareness, and highlights the importance of creating a richer epistemic framework. However, because this behavior is complicated, and the agent's shifting perception is likely subject to equilibrium effects, I leave any formal analysis to future work.

4.7 LITERATURE REVIEW

This paper is within the context of two distinct, albeit related, literatures: that on epistemic logic and unawareness, and that on unawareness and unforeseen contingencies in decision theory. Unawareness was first formalized within modal logic by Fagin and Halpern (1988), who introduced the modal operator for awareness, A, and explicit knowledge, E. This was extended later by Halpern and Rêgo (2009) to include first order statements that allow for introspective unawareness, and extended further by Halpern and Rêgo (2013), to allow the agent to be uncertain about whether she has full awareness or not.

The structure of quantification and awareness in Halpern and Rêgo (2009) is significantly different than the one presented here. They use a logic where quantification is over formulae. Because the foremost concern is over the alternatives that can be consumed, (and over which preferences can be defined), I use a first order modal logic where the variables range only over some fixed domain and are present in formulae only under the scope of predicates. Moreover, I am interested in the awareness of particular objects and so explicitly consider awareness of formula that contain free variables; awareness is a subset of the auxiliary semantic language $\mathcal{L}(X)$.

Relatedly, Board and Chung (2011) and Board *et al.* (2011) propose an alternate structure of awareness that is, like the one presented here, based on objects and predicates. They, however, assume awareness is fully characterized by a set of objects. Then the DM is aware of a statement like "phone x is preferred to phone y" if she is aware of x and y. This does not allow, as I do, for the DM to be aware of an object but not some of its attributes; for example, a DM cannot be aware of x and y but unaware of statements like "phone x has a higher pixel density than phone y." For many practical applications, it seems necessary to disconnect the DM's awareness of attributes of objects from her awareness of objects themselves. Note, this issue cannot be resolved by a simple relabeling of alternatives: for example, phone x' is similar to phone x in all respects save its pixel density. This resolution does not work because phone x has all of its properties regardless of the DM's awareness –a commitment to consume x is a commitment to the relevant properties.

The logic presented in this paper, is a happy medium between that of Halpern and Rêgo (2009) and of Board and Chung (2011). It allows for general awareness structures, ranging over formulae, but *after* the formulae have been given semantic interpretations. Therefore, we can speak of the DM's awareness of particular objects, but not necessarily all of the attributes, or relations between, these objects.

In economics, *state space models* –the semantic structure that include states, and define knowledge and unawareness as operators thereon, as in this paper– have been of particular interest. Modica and Rustichini (1994) and Dekel *et al.* (1998) both provide beautiful, albeit negative, results in this domain. They show, under mild conditions, unawareness must be in some sense trivial; the DM is either fully aware or fully unaware. While Modica and Rustichini (1994) consider a specific awareness modality, Dekel *et al.* (1998) show, under reasonable axioms, state-space models do not allow any non-trivial unawareness operator. As stated, this would be a very damning result for this paper, as it would imply either $\succcurlyeq_E = \succcurlyeq_K$ or $\succcurlyeq_E = \emptyset$, either way, not making for an interesting decision theory. The resolution comes from disentangling explicit and implicit knowledge. Considering these forms of knowledge separately reasonably avoids ever simultaneously satisfying the necessary axioms for their negative result. A far more succinct and intuitive discussion than I could hope to achieve is found in Section 4 of Halpern and Rêgo (2013), and so, I refer the reader there.

Beyond the separation of implicit and explicit knowledge, there have been other approaches to the formalization of unawareness that circumvent the problems outlined in the pervious paragraph. Modica and Rustichini (1999) propose propose models in which the DM is aware only of a subset of formulae (similar in spirit to the awareness sets proposed here, albeit necessarily generated by primitive propositions), and entertains a subjective state space (a coarsening of the objective state space) in which the DM cannot distinguish between any two states that differ only by the truth of a proposition of which she is unaware. Heifetz *et al.* (2006) and Heifetz *et al.* (2008) consider a lattice of state spaces that are ordered according to their expressiveness. In this way, unawareness is captured by events that are not expressible from different spaces –events that are not contained in the event nor the negation of the DM's knowledge. Li (2009) also provides a model with multiple state spaces, where the DM entertains a subjective state space (similar to the above papers, the set of possible state spaces forms a lattice). This allows the DM to be unaware of events in finer state spaces, while having non-trivial knowledge in coarser state spaces.

The decision theoretic take on unawareness is primarily based on a revealed preference framework, and so, unlike its logical counterpart does not dictate the structure of awareness but rather tries to identify it from observable behavior. The first account of this approach (and which predates the literature by a sizable margin) is Kreps (1979). Kreps considers a DM who ranks menus of alternatives, and whose preferences respect set inclusion. The motivation being larger menus provide the DM with the flexibility to make choices after *unforeseen contingencies*. This interpretation, while not strictly ruled out by the model, is certainly not its most obvious interpretation, especially in light of the titular representation theorem. That Krepsian behavior can always be rationalized in a model without appealing to unawareness is shown formally in Theorem 4.15; a longer discussion in relation to this paper is found in Section 4.5.

More recently there has been a growing interest in modeling the unaware DM. Kochov (2015) posits a behavioral definition of unforeseen contingencies. He considers the DM's ranking over streams of acts (function from the state space to consumption). An event, E, is considered foreseen if all bets on E do not distort otherwise prefect hedges. That is to say, an event is unforeseen if the DM cannot "properly forecast the outcomes of an action" contingent on the event. Kochov shows the events a DM is aware of form a coarsening of the modeler's state space. In a similar vein, Minardi and Savochkin (2015) also contemplate a DM who has a coarser view of the world than the modeler. This coarse perception manifests itself via imperfect updating; the DM cannot "correctly" map the true event onto an event in her subjective state space. The events that are inarticulable in the subjective language of the DM can be interpreted as unforeseen. However, in these works, the objects of which the DM is supposedly unaware are encoded objectively into the alternatives she ranks. Because of this, I argue they are behavioral models of *misinterpretation* rather than unawareness.

Karni and Vierø (2016); Grant and Quiggin (2014) are more explicit about modeling unawareness, and, along with their companion papers, are (to my knowledge) the only decision theoretic paper that deals with unawareness of consumption alternatives, rather than contingencies. They examine a DM who evaluates acts which may specify an alternative explicitly demarcated as "something the DM is unaware of," and who can be interpreted as possessing probabilistic belief regarding the likelihood of discovering such an outcome. They observe the DM's preferences over acts, both before and after the discovery of a new alternative. Of particular interest to the authors, is the process by which the DM updates her beliefs; in particular they provide the axiomatic characterization of *reverse Bayesianism* first developed by the same authors in Karni and Vierø (2013). The assertion of the existence of a novel consumption alternative without dictating what that alternative is has clear parallels with this paper: the DM in their paper is necessarily introspectively unaware. However, their framework takes as given the epistemic state of the decision maker –naive in Karni and Vierø (2013) and introspective in Karni and Vierø (2016).

Grant and Quiggin (2012) develop a general model to deal with unawareness in games, founded on a modal logic which incorporates unawareness in a similar way to Modica and Rustichini (1994). They show that while this model is rich enough to provide non-trivial unawareness, it fails to allow for introspective unawareness, even when first order quantification is permitted. (This limitation arises because of the desired interplay between the structure of knowledge and the structure of awareness as facilitated by the game theoretic environment.) By relaxing the connection with the modal underpinnings, they then consider possible heuristics that a player might exhibit when she inductively reasons that she is introspectively unaware. In a companion paper, Grant and Quiggin (2014) provide a (decision theoretic) axiomatization of such heuristics.

Several recent papers examine the value of information (and expanded awareness) when agents have bounded perception. Quiggin (2015) defines the value of awareness, in analogy to the value of information, and shows the two measures are perfectly negatively correlated. Galanis (2015) examines the value of information under unawareness and shows, in contrast to standard results, the value of information can be negative. Extending this line of thought to the multi agent case (with risk sharing) Galanis (2016), shows, under unawareness, public information might be treated asymmetrically, allowing some agents to prosper at the expense of others.

There are also (very few) economic papers that directly investigate the connection between observable choice and the underlying logical structure. Morris (1996) works somewhat in the reverse direction of the current paper, providing a characterization of different logical axioms (for example, K, T, 4, etc) in terms of preferences over bets on the state of the world. Schipper (2014) extends this methodology to include unawareness structures as described in Heifetz *et al.* (2006). In a similar set up, Schipper (2013) constructs an expected utility framework to elicit (or reveal) a DM's belief regarding the probability of events (when she might be aware of some events). Schipper concludes, the behavioral indication of unawareness of event E is that the DM treats both E *and* its complement as null. This is a very nice result, and the intuition maps nicely to the idea of unawareness. To achieve such a representation, the objects of choice are maps from the “true” state-space to outcomes, and as such, raise the previously discussed issues of observability and the limits of an unaware modeler.

APPENDIX A

APPENDIX OF CHAPTER 2

A.1 LEMMAS

Lemma 1. *If* $\{\succsim_A\}_{A\in\mathcal{K}(X)}$ *satisfies* UV*, then for all* $A \in \mathcal{K}(X)$*,* $N_A = E_A \cup N$*.*

Proof. Fix some $A \in \mathcal{K}(X)$. By definition both E_A and N are subsets of N_A, so, $E_A \cup N \subseteq N_A$. Towards the opposite inclusion, let $s \in N_A$. We will show that if $s \notin E_A$ then $s \in N$. So assume further, that $s \notin E_A$. Since s is null-A, $x^\star \sim_A x_\star$. Since s is not e-null-A, for every $B \in \mathcal{K}(X)$, $x^\star \sim_B x_\star$. By the contrapositive of UV we have $s \in N_B$. Since this holds for all B, $s \in N$.

Lemma 2. *Let* $\{\mu_{A_n} \in \Delta(S)\}_{n\in\mathbb{N}}$ *converge weakly to* $\mu_A \in \Delta(S)$*. Endow* $[0,1]^S$ *with the* $||\cdot||_1$ *norm. Fix some* $\beta \in (0,1)$*. Then, the Kuratowski limit of* $\left\{\hat{g} \in [0,1]^S | \, \mathbb{E}_{\mu_{A_n}}(\hat{g}) \geq \beta\right\}$ *is* $\left\{\hat{g} \in [0,1]^S | \, \mathbb{E}_{\mu_A}(\hat{g}) \geq \beta\right\}$*.*[41]

Proof. Define $\dot{C}(A,\beta) = \left\{\hat{g} \in [0,1]^S | \, \mathbb{E}_{\mu_A}(\hat{g}) \geq \beta\right\}$. First, we show, $\hat{g} \in \dot{C}(A,\beta) \implies \hat{g} \in \underset{n\to\infty}{LI}\, \dot{C}(A_n,\beta)$. Indeed, assume $\hat{g} \in \dot{C}(A,\beta)$, and consider $B_\epsilon(\hat{g})$. Let $in(\hat{g}) = \{s \in S | \hat{g}(s) < 1\}$. Let $M = \sum_{s\in in(\hat{g})} \mu_A(s)$. There are two cases: where $M = 0$ and where $M > 0$.

[41] Recall, for a metric space $\mathcal{Y}$ the Kuratowski limit exists if and only if the *Kuratowski limit superior* and *Kuratowski limit inferior* coincide. These are defined, respectively, as follows, (for some $\{Y_n \subseteq \mathcal{Y}\}_{n\in\mathbb{N}}$), as

$$\underset{n\to\infty}{LS}\, Y_n = \{y \in \mathcal{Y} | \forall B_\epsilon(y), B_\epsilon(y) \cap Y_n \neq \emptyset, \text{ for infinitely many } n\},$$
$$\underset{n\to\infty}{LI}\, Y_n = \{y \in \mathcal{Y} | \forall B_\epsilon(y), B_\epsilon(y) \cap Y_n \neq \emptyset, \text{ for sufficiently large } n\}$$

where $B_\epsilon(y)$ is the ball or radius ϵ around y. Note, these always exists and necessarily $\underset{n\to\infty}{LI}\, Y_n \subseteq \underset{n\to\infty}{LS}\, Y_n$.

First, let $M = 0$. Since S is a discrete space, and so the indicator functions for each state are continuous, weak convergence implies strong convergence. Thus, $\lim_n \mu_{A_n}(in(\hat{g})) \to \mu_A(in(\hat{g})) = 0$: in particular, for sufficiently large n, $\mu_{A_n}(in(\hat{g})) < 1 - \beta$, and so $\mathbb{E}_{\mu_{A_n}}(\hat{g}) = \mathbb{E}_{\mu_{A_n}}(\hat{g}|in(\hat{g})) + \mathbb{E}_{\mu_{A_n}}(\hat{g}|in((\hat{g}))^c) \geq \beta$.

So, let $M > 0$ and fix some $\epsilon > 0$. Let $\epsilon' = \frac{\min\{\epsilon, \{1-\hat{g}(s)\}_{s \in in(\hat{g})}\}}{|S|}$. Let, $\hat{f} \in B_\epsilon(\hat{g})$ be given by

$$\hat{f}(s) = \begin{cases} \hat{g}(s) & \text{if } s \notin in(\hat{g}) \\ \hat{g}(s) + \epsilon' & \text{if } s \in in(\hat{g}) \end{cases}$$

Let $m \in \mathbb{N}$ be such that for all $n > m$, $|\mu_A(s) - \mu_{A_n}(s)| < \frac{M\epsilon'}{|S|}$. Then,

$$\begin{aligned} \mathbb{E}_{\mu_{A_n}}(\hat{f}) &> \mathbb{E}_{\mu_A}(\hat{f}) - \sum_s \frac{M\epsilon'}{|S|} \hat{f}(s) \\ &> \sum_{s \notin in(\hat{g})} \mu_A(s)\hat{g}(s) + \sum_{s \in in(\hat{g})} \mu_A(s)\hat{g}(s) + \sum_{s \in in(\hat{g})} \mu_A(s)\epsilon' - M\epsilon' \\ &= \beta. \end{aligned}$$

So, $\hat{f} \in \dot{C}(A_n, \beta)$ for all $n > m$.

Now, we show that $\hat{g} \in \underset{n\to\infty}{LS}\, \dot{C}(A_n, \beta) \implies \hat{g} \in \dot{C}(A, \beta)$. Indeed, assume $\hat{g} \in \underset{n\to\infty}{LS}\, \dot{C}(A_n, \beta)$. Then, let n_k denote a subsequence for which $B_{\frac{1}{k}}(\hat{g}) \cap \dot{C}(A_{n_k}, \beta) \neq \emptyset$, with $\hat{g}^k$ a point in the intersection. Then $\lim_{n\to\infty}(\mu_{A_{n_k}}(s)\hat{g}^k(s)) = \mu_A(s)\hat{g}(s)$ for all $s \in S$. So, $\mathbb{E}_{\mu_{A_{n_k}}}(\hat{g})$ converges to $\mathbb{E}_{\mu_A}(\hat{g})$, and so, by the preservation of inequalities, $\hat{g} \in \dot{C}(A, \beta)$.

Therefore, $\dot{C}(A, \beta) \subseteq \underset{n\to\infty}{LI}\, \dot{C}(A_n, \beta) \subseteq \underset{n\to\infty}{LS}\, \dot{C}(A_n, \beta) \subseteq \dot{C}(A, \beta)$; the Kuratowski limit exists and is equal to $\dot{C}(A, \beta)$.

Definition. *For a menu $A \in \mathcal{K}(X)$, define the* equalizer *of A, $e_A : N_A \times N_A \to R_{++}$ as*

$$e_A(s, s') \mapsto \begin{cases} \frac{1}{\alpha} \text{ such that } (x_\star)_{-s}(\alpha x^\star + (1-\alpha)x_\star) \sim_A (x_\star)_{-s'}x^\star & \text{if } (x_\star)_{-s}x^\star \succcurlyeq_A (x_\star)_{-s'}x^\star \\ \alpha \text{ such that } (x_\star)_{-s'}(\alpha x^\star + (1-\alpha)x_\star) \sim_A (x_\star)_{-s}x^\star & \text{if } (x_\star)_{-s'}x^\star \succ_A (x_\star)_{-s}x^\star \end{cases}$$

That e_A is well defined follows from the following observation.

Lemma 3. *Let $\{\succcurlyeq_A\}_{A\in\mathcal{K}(X)}$ be represented by some (MBR), with beliefs $\{\mu_A | A \in \mathcal{K}(X)\}$, all of which have full support. Then, for all $A \in \mathcal{K}(X)$, $e_A(s, s') = \frac{\mu_A(s)}{\mu_A(s')}$.*

Proof. If $(x_\star)_{-s}x^\star \succcurlyeq_A (x_\star)_{-s'}x^\star$, then for some $\alpha \in (0, 1)$, we have $(x_\star)_{-s}\big(\alpha x^\star + (1-\alpha)x_\star\big) \sim_A$

$(x_\star)_{-s'}x^\star$. Using (MBR), we have that $U_A\Big((x_\star)_{-s'}x^\star\Big) = \mu_A(s')$, and $U_A\Big((x_\star)_{-s}\big(\alpha x^\star + (1-\alpha)x_\star\big)\Big) = \alpha\mu_A(s)$. Setting $e_A(s,s') = \frac{1}{\alpha}$, delivers ther result. The other case is similar.

Lemma 4. *Let $\{\succcurlyeq_A\}_{A\in\mathcal{K}(X)}$ be represented by some (MBR) with beliefs $\{\mu_A | A \in \mathcal{K}(X)\}$, all of which have full support. Then $\{\succcurlyeq_A\}_{A\in\mathcal{K}(X)}$ satisfies IID if and only if, for all $x \in X$ and $A, B \in \mathcal{K}(X)$ with $x \notin A \cup B$, and states $s, s' \in S$, we have*

$$\frac{e_A(s,s')}{e_{A\cup x}(s,s')} = \frac{e_B(s,s')}{e_{B\cup x}(s,s')}. \tag{A.1}$$

Proof. **Necessity.** Assume that (A.1) holds, with $x \in X$, $A, B \in \mathcal{K}(X)$, and $s, s' \in S$ satisfying the relevant constraints. Denote by A' and B', $A \cup x$ and $B \cup x$, respectively. Towards a contradiction, assume that there exists some $\pi^A, \rho^A, \in \Delta(A)$, $\pi^B, \rho^B, \in \Delta(A)$, and $\alpha = (0,1)$ be such that,

$$(x_\star)_{-s}\pi^A \sim_A (x_\star)_{-s'}\rho^A, \qquad \text{implying} \qquad \frac{\mu_A(s)}{\mu_A(s')} = \frac{(\rho^A \cdot u)}{(\pi^A \cdot u)}, \tag{A.2}$$

$$(x_\star)_{-s}\pi^B \sim_B (x_\star)_{-s'}\rho^B, \qquad \text{implying} \qquad \frac{\mu_B(s)}{\mu_B(s')} = \frac{(\rho^B \cdot u)}{(\pi^B \cdot u)}, \tag{A.3}$$

$$(x_\star)_{-s}\pi^A \succcurlyeq_{A'} (x_\star)_{-s'}(\alpha\rho^A + (1-\alpha)x_\star), \qquad \text{implying} \qquad \frac{\mu_{A'}(s)}{\mu_{A'}(s')} \geq \alpha\frac{(\rho^A \cdot u)}{(\pi^A \cdot u)}, \tag{A.4}$$

$$(x_\star)_{-s}\pi^B \prec_{B'} (x_\star)_{-s'}(\alpha\rho^B + (1-\alpha)x_\star), \qquad \text{implying} \qquad \frac{\mu_{B'}(s)}{\mu_{A'}(s')} < \alpha\frac{(\rho^B \cdot u)}{(\pi^B \cdot u)} \tag{A.5}$$

Dividing the implications of (A.2) by (A.4) and (A.3) by (A.5), and applying Lemma 3, we get a direct contradiction to (A.1).

Sufficiency. Assume IID holds. Let $x \in X$, $A, B \in \mathcal{K}(X)$, and $s, s' \in S$ satisfy the relevant constraints for IID. Let $M = \max\{\frac{\mu_A(s)}{\mu_A(s')}, \frac{\mu_B(s)}{\mu_B(s')}, 1\}$. Finally, for any $\beta \in [0, M]$, and $s \in S$, let $f^*(s,\beta) = (x_\star)_{-s}(\frac{\beta}{M}x^\star + (1-\frac{\beta}{M})x_\star)$. Using (MBR), we have

$$U_A\big(f^*(s,1)\big) = U_A\big(f^*(s', \frac{\mu_A(s)}{\mu_A(s')})\big) = \frac{\mu_A(s)}{M}, \text{ and} \tag{A.6}$$

$$U_B\big(f^*(s,1)\big) = U_B\big(f^*(s', \frac{\mu_B(s)}{\mu_B(s')})\big) = \frac{\mu_B(s)}{M}. \tag{A.7}$$

Let $\alpha = \frac{\mu_A(s')\mu_{A'}(s)}{\mu_A(s)\mu_{A'}(s')}$. Case: $\alpha \leq 1$. Applying (MBR) again delivers,

$$U_{A'}\big(f^*(s,1)\big) = U_{A'}\big(f^*(s', \alpha\frac{\mu_A(s)}{\mu_A(s')})\big) = \frac{\mu_{A'}(s)}{M}.$$

By (A.6) and (A.7), we can apply IID, so,

$$U_{B'}\big(f^*(s,1)\big) = U_{B'}\big(f^*(s', \alpha\frac{\mu_B(s)}{\mu_B(s')})\big). \tag{A.8}$$

Expanding (A.8) according to (MBR):

$$\mu_{B'}(s) = \mu_{B'}(s')\frac{\mu_A(s')\mu_{A'}(s)}{\mu_A(s)\mu_{A'}(s')}\frac{\mu_B(s)}{\mu_B(s')},$$

which by Lemma 3, is equivalent to (A.1). In the case where $\alpha > 1$, consider $f^*(s, \frac{1}{\alpha})$ and $f^*(s', \frac{\mu_A(s)}{\mu_A(s')})$, and proceed in a similar manner.

A.2 PROOFS

Proof of Theorem 2.2. **Part (a), necessity.** The necessity of EU, MC, UV are obvious from the inspection of the representing functionals. CC follows from the continuity of $\mu_{(\cdot)}$. Fix some $\{A_n\}_{n\in\mathbb{N}}$ with limit point A. Notice that $\mathcal{F}_\star$ is homeomorphic to $[0,1]^S$ (endowed with the $||\cdot||_1$ norm), via the identification of $f = \big((\alpha_1 x^\star + (1-\alpha_1)x_\star), \ldots, (\alpha_{|S|}x^\star + (1-\alpha_{|S|})x_\star)\big)$ with $\hat{f} = (\alpha_{s_1}, \ldots, \alpha_{s_{|S|}})$. Fix some $\beta \in (0,1)$. Then, by the representation, we have, for every $B \in \mathcal{K}(X)$,

$$\dot{C}(B,\beta) = \{g \in \mathcal{F}_\star | g \succcurlyeq_B \beta x^\star + (1-\beta)x_\star\} \cong \left\{g \in \mathcal{F}_\star | \mathop{\mathbb{E}}_{\mu_B}(\hat{g}) \geq \beta\right\}. \tag{A.9}$$

As $\mu_\cdot$ is continuous, we can apply Lemma 2: the Kuratowski limit of $\dot{C}(A_n,\beta)$ is $\dot{C}(A,\beta)$. Since for all A_n, $\dot{C}(A_n,\beta)$ is convex, they are connected. By the equivalence of Kuratowski convergence and convergence in Hausdorff metric for sequences of connected sets (Salinetti and Wets (1979), Corollary 3A), we are done.

Part (a), sufficiency. It is a direct application of the expected utility theorem that EU delivers for each A the existence of some continuous and bounded $w : S \times X \to \mathbb{R}$ such that

$$U_A^{VNM}(f) = \sum_s \left(\mathop{\mathbb{E}}_{f(s)}(w_A(s,x))\right),$$

represents $\succcurlyeq_A$. Moreover, if $w_A(s,x)$ and $\hat{w}_A(s,x)$ both represent $\succcurlyeq_A$, then $w_A(s,x) = a_A\hat{w}_A(s,x) + b_A(s)$ where $a_A \in \mathbb{R}_{++}$ and $b_A(s) \in \mathbb{R}$ for all $s \in S$.[42]

[42]For a reference using the same framework, see "NM Theorem" of Karni *et al.* (1983).

By exploiting the degrees of freedom from the scalars $b_A(s)$, we can set $w_A(s, x_\star) = 0$, for all A and all $s \in S$. The resulting functionals are unique up to linear transformations. Note, this implies that for all $s \in N_A$, $w_A(s, \cdot)$ is identically 0.

For each $A \in \mathcal{K}(X)$, let $u_A(s, x) : N_A^c \times A \to \mathbb{R}$ be the mapping

$$u_A : (s, x) \mapsto \frac{w_A(s, x)}{w_A(s, x^\star)},$$

and $\mu_A \in \Delta((E_A \cup N)^c)$ as the distribution defined by

$$\mu_A(s) = \frac{w_A(s, x^\star)}{\sum_s w_A(s, x^\star)}.$$

Notice, μ_A is well defined and has full support, since by the non-triviality of $\succcurlyeq_A$, $N_A \neq S$, and for each $s \in (E_A \cup N)^c$, $s \in N_A^c$ (Lemma 1), and so by UV, $w_A(s, x^\star) > w_A(s, x_\star) = 0$. Define,

$$U_A^{MD}(f) = \mathop{\mathbb{E}}_{\mu_A}\left(\mathop{\mathbb{E}}_{f(s)}(u_A(s, x))\right). \tag{A.10}$$

Following standard algebraic manipulations, we can see $\mu_A(s)u_A(s, x) = \frac{1}{\sum_s w_A(s,x^\star)} w_A(s, x)$, and therefore U_A^{MD} represents $\succcurlyeq_A$.

Let $D = \{(s, x) \in S \times X \mid \exists A \in \mathcal{K}(X), x \in A, s \notin N_A\}$. For each $(s, x) \in D$, let $A_{s,x}$ be any menu such that $x \in A_{s,x}$ and $s \notin N_{A_{s,x}}$. Define the mapping $u : D \to \mathbb{R}$ as,

$$u : (s, x) \mapsto u_{A_{s,x}}(s, x).$$

and extend u to $S \times X$, by defining $u(s, x) = 0$ for all $(x, s) \in D^c$.

We now claim, for any $A \in \mathcal{K}(X)$, $s \notin N_A$ and $x \in A$, we have $u(s, x) = u_A(s, x)$. Indeed, for every such $A, B \in \mathcal{K}(X)$ and $s \notin N_A \cup N_B$. Let $\succcurlyeq_{A|B|s} \subseteq (\Delta(A \cap B))^2$ be defined by:

$$\pi \succcurlyeq_{A|B|s} \rho \iff \mathop{\mathbb{E}}_{\pi}(u_A(s, x)) \geq \mathop{\mathbb{E}}_{\rho}(u_A(s, x)).$$

Since $\succcurlyeq_{A|B|s}$ is represented by a linear utility function, it satisfies EU, and so, by the expected utility theorem, $u_A(s, \cdot)$ is the unique utility index, up to affine transformations.

Fix some A and $s \notin N_A$, and $x \in A$. By (A.10), $\mathbb{E}_\pi(u_A(s, x)) \geq \mathbb{E}_\rho(u_A(s, x))$ holds if and only if, for all $f \in F_A$, $f_{-s}\pi \succcurlyeq_A f_{-s}\rho$. Applying MC, we immediately have $g_{-s}\pi \succcurlyeq_{A_{s,x}} g_{-s}\rho$ for any $g \in F_{A_{s,x}}$ (here we use the fact that $s \notin N_{A_{s,x}}$). From (A.10) again, $\succcurlyeq_{A|A_{s,x}|s} = \succcurlyeq_{A_{s,x}|A|s}$. So $u_A(s, \cdot)$ is an affine transformation of $u_{A_{s,x}}(s, \cdot)$. Moreover, both are twice normalized: $u_A(s, x^\star) = u_{A_{s,x}}(s, x^\star) = 1$ and $u_A(s, x_\star) = u_{A_{s,x}}(s, x_\star) = 0$. Hence they must coincide on $A \cap A_{s,x}$. Finally, since $x \in A \cap A_{s,x}$, we have $u_A(s, x) = u_{A_{s,x}}(s, x) = u(s, x)$. Clearly, since

$u_A = u|_A$ and u_A is continuous and bounded, $u|_A$ is continuous and bounded.

Because it eases exposition, we well prove that $\mu_{(\cdot)} : \mathcal{K}(X) \to \Delta(S)$ is continuous after we have shown that it is unique.

Part (b). Uniqueness results are standard. It is clear from the argument above that $u(\cdot,\cdot)$ is unique (given the normalization on $\bigstar$), as it must represent $\succcurlyeq_{A_{s,x}|A_{s,x}|s}$. With regards to beliefs, assume to the contrary, for some $A \in \mathcal{K}(X)$, μ and ν both represent (in conjunction with u, as in (MBR)) $\succcurlyeq_A$. Then there must be some s, s', such that $\mu(s) < \nu(s)$ and $\mu(s') > \nu(s')$. Assume (*with* loss of generality, but the other case follows from the reflected argument) that $\mu(s) \leq \mu(s')$. Set π as the probability distribution given by,

$$\pi(x) = \begin{cases} \frac{\mu(s)}{\mu(s')} & \text{if } x = x^\star, \\ 1 - \frac{\mu(s)}{\mu(s')} & \text{if } x = x_\star, \\ 0 & \text{otherwise.} \end{cases}$$

Given that (μ, u) represents $\succcurlyeq_A$, it follows from (MBR) that $(x_\star)_{-s'}\pi \sim_A (x_\star)_{-s}x^\star$. But, since (ν, u) also represents $\succcurlyeq_A$: $(x_\star)_{-s'}\pi \prec_A (x_\star)_{-s}x^\star$, a clear contradiction.

Part (a), sufficiency continued. Assume, by way of contradiction, that $\mu_{(\cdot)} : \mathcal{K}(X) \to \Delta(S)$ was not continuous. Then there exists some $\{A_n\}_{n\in\mathbb{N}}$ converging to A, such that $\{\mu_{A_n}\}_{n\in\mathbb{N}}$ does not converge to A. This implies that there exists some bounded and continuous function, $\hat{f} : S \to \mathbb{R}$ such that $\mathbb{E}_{\mu_{A_n}}(\hat{f})$, does not converge to $\mathbb{E}_{\mu_A}(\hat{f})$, say it is strictly above (if it is below, or does not exist, the arguments are essentially the same). Since f is bounded, it is without loss of generality that we consider $\hat{f} : S \to [0,1]$. Let $\beta = \mathbb{E}_{\mu_A}(\hat{f})$. Then for some $\epsilon > 0$ and all $m \in \mathbb{N}$, there is some $n > m$, such that $\mathbb{E}_{\mu_{A_n}}(\hat{f}) > \beta + \epsilon$. Hence, there is some subsequence, A_{n_k} such that $\hat{f} \in \dot{C}(A_{n_k}, \beta + \frac{\epsilon}{2})$, and $\hat{f} \notin \dot{C}(A, \beta + \frac{\epsilon}{2})$ (notice, $\dot{C}(\cdot,\cdot)$ is defined in (A.9)). But, $\dot{C}(A, \beta + \frac{\epsilon}{2})$ is closed by EU. Hence, $\dot{C}(A_{n_k}, \beta + \frac{\epsilon}{2})$ is bounded away from $\dot{C}(A, \beta + \frac{\epsilon}{2})$, completing the proof as it is tantamount to a contradiction of CC.

Proof of Proposition 2.4. This follows directly from Lemma 1 of Shmaya and Yariv (2015) which states (in the language of this paper) that such a given a prior μ and a set of posteriors $\{\mu_A\}_{A\in\mathcal{K}(X)}$ one can find a generating signal structure, that transforms μ into $\{\mu_A\}_{A\in\mathcal{K}(X)}$, so long as the prior beliefs lie in the relative interior of the convex hull of the set of posteriors, i.e., $\mu \in \text{ri}(\text{Conv}(\{\mu_A\}_{A \in \mathcal{K}(X)}))$. Given the additional flexibility in choosing the prior, and the fact the relative interior of a non-empty convex set is non-empty, we can always find such a μ.

Proof of Theorem 2.5. **Necessity.** Assume there was some $(\mu, l, \mathcal{K}(X)/\dot{\sim})$, satisfying the relevant assumptions, that generates $\{\mu_A | A \in \mathcal{K}(X)\}$. Let $A, B \in \mathcal{K}(X)$ such that $N_A^c \cap N_B^c \neq \emptyset$. Let $s \in N_A^c \cap N_B^c$. Thus, it must be that $l([A], s) \neq 0$ and $l([B], s) \neq 0$, implying that $[A] = [B]$. But then, $l([A], \cdot) \equiv l([B], \cdot)$. Therefore, $\mu_A = \mu_B$. Clearly, these induce the same preferences over the intersection of A and B, and PS is satisfied.

Sufficiency. Assume PS holds. By assumption $N = \emptyset$, so, for each $s \in S$, choose some menu $A(s)$ such that $s \notin N_A$. Define $\mu \in \Delta(S)$ as

$$\mu(s) = \frac{\mu_{A(s)}(s)}{\sum_{s' \in S} \mu_{A(s')}(s')}.$$

Define $\dot{\sim}$ over $\mathcal{K}(X)^2$, by $A \dot{\sim} B$ if and only if $N_A = N_B$. It is obvious that $\dot{\sim}$ is a equivalence relation, and so, let $\mathcal{K}(X)/\dot{\sim}$ be the resulting quotient set. For each $[A] \in \mathcal{K}(X)/\dot{\sim}$ and $s \in S$, set

$$l([A], s) = \begin{cases} 1 \text{ if } s \notin N_A, \\ 0 \text{ if } s \in N_A. \end{cases}$$

Now, notice that if $l([A], s) = l([B], s) = 1$, then $N_A^c \cap N_B^c \neq \emptyset$. So by PS the projections of $\succcurlyeq_A$ and $\succcurlyeq_B$ onto $\mathcal{F}_\star$ must coincide. Therefore, it must be that $\mu_A = \mu_B$, implying that $[A] = [B]$.

Finally, pick some A in $\mathcal{K}(X)$. Now notice that for all $s \in N_A^c$, $s \in N_{A(s)}^c$ and so, by the above line of reasoning $\mu_A = \mu_{A(s)}$. Clearly, since $l([A], s)|_{N_A^c} \equiv 1$, (2.2) holds.

Proof of Theorem 2.6. **Necessity.** Assume there exists some (μ, l, X), with $l(x, s) \in (0, 1)$ for all $(x, s) \in X \times S$, that generates $\{\mu_A | A \in \mathcal{K}(X)\}$. For some A that does not contain x and $s, s' \in S$, we have $e_A(s, s') = \frac{\mu_A(s)}{\mu_A(s')}$, and $e_{A \cup x}(s, s') = \frac{\mu_{A \cup x}(s)}{\mu_{A \cup x}(s')}$. Using (2.3), we have

$$\mu_{A \cup x}(s) = \frac{\mu_A(s) \frac{l(x,s)}{1 - l(x,s)}}{\mathbb{E}_{\mu_A}\left(\frac{l(x,s')}{1 - l(x,s')}\right)}$$

for all $s \in S$. So,

$$\begin{aligned} e_{A\cup x}(s,s') &= \frac{\mu_A(s)\frac{l(x,s)}{1-l(x,s)}}{\mu_A(s')\frac{l(x,s')}{1-l(x,s')}}, \\ &= \frac{\frac{l(x,s)}{1-l(x,s)}}{\frac{l(x,s')}{1-l(x,s')}} e_A(s,s'). \end{aligned}$$

Hence the ratio of equalizers does not depend on the menu. By Lemma 4, IID holds.

Sufficiency. Assume IID holds. Let $\alpha(x,s) = \frac{\mu_{\{\bigstar\cup x\}}(s)}{\mu_{\bigstar}(s)}$, set

$$l(x,s) = \frac{\alpha(x,s)}{1+\alpha(x,s)}, \tag{A.11}$$

for all $(x,s) \in X \setminus \bigstar \times S$ and $l(x^\star,\cdot) \equiv l(x_\star,\cdot) \equiv 1$. Let $\gamma(s) = \prod_{x\in X\setminus\bigstar}(1-l(x,s))$. Define $\mu \in \Delta(S)$ by,

$$\mu(s) = \frac{\frac{\mu_\bigstar(s)}{\gamma(s)}}{\mathbb{E}_{\mu_\bigstar}\left(\frac{1}{\gamma(s')}\right)}$$

By construction, $\mu_\bigstar$ is generated according to (2.3).

We will now show that as defined, (μ, l, X) generates the remainder of $\{\mu_A | A \in \mathcal{K}(X)\}$. We proceed by induction on the cardinality of A.

Define $\nu_x \in \Delta(S)$

$$\nu_x(s) = \frac{\mu_\bigstar(s)\frac{l(x,s)}{1-l(x,s)}}{\mathbb{E}_{\mu_\bigstar}\left(\frac{l(x,s')}{1-l(x,s')}\right)}$$

Now, using the algebraic identity $\frac{\frac{\alpha}{1+\alpha}}{1-\frac{\alpha}{1+\alpha}} = \alpha$, we have $\frac{l(x,s')}{1-l(x,s')} = \alpha(x,s) = \frac{\mu_{\{\bigstar\cup x\}}(s)}{\mu_\bigstar(s)}$. Therefore $\nu_x(s) = \mu_{\{\bigstar\cup x\}}(s)$. This completes the base case (for $|A| = 3$).

Now assume that (μ, l, X) generates $\{\mu_A | A \in \mathcal{K}(X), |A| \leq n\}$. Fix any A with n elements, and $x \notin A$. Let A' denote $A \cup x$. Set,

$$\nu_{A'} = \frac{\mu_A(s)\frac{l(x,s)}{1-l(x,s)}}{\mathbb{E}_{\mu_A}\left(\frac{l(x,s')}{1-l(x,s')}\right)}$$

Towards a contradiction, assume that $u_{A'} \neq \nu_{A'}$. Therefore, there must exist some s such that $u_{A'}(s) > \nu_{A'}(s)$, and s' such that $u_{A'}(s') < \nu_{A'}(s')$. Therefore we have:

$$\frac{e_A(s,s')}{e_{A\cup x}(s,s')} = \frac{\frac{\mu_A(s)}{\mu_A(s')}}{\frac{\mu_{A'}(s)}{\mu_{A'}(s')}} < \frac{\frac{\mu_A(s)}{\mu_A(s')}}{\frac{\nu_{A'}(s)}{\nu_{A'}(s')}}$$

$$= \frac{\frac{l(x,s')}{1-l(x,s')}}{\frac{l(x,s)}{1-l(x,s)}} = \frac{e_{\bigstar}(s,s')}{e_{\bigstar \cup x}(s,s')}$$

Which, by Lemma 4 is a contradiction to IID. Therefore, the inductive step holds, and (μ, l, X) generates $\{\mu_A | A \in \mathcal{K}(X)\}$.

APPENDIX B

APPENDIX FOR CHAPTER 3

B.1 ON THE CONSTRUCTION OF PLANS.

Generalized Plans. We will begin by constructing a more general notion of Plans (reminiscent of IHCPs, first constructed in Gul and Pesendorfer (2004), and then refine our notion to capture only the elements of interest. This methodology serves two purposes. First, the more general approach allows us to use standard techniques for the construction of infinite horizon choice objects. Second, generalized plans may be of direct interest in future work, when, for example, denumerable support is not desirable. To begin, let $Q_0 = \Delta^{\mathcal{B}}(\mathcal{A})$ and, for define recursively for each $n \geq 1$

$$Q_n = \Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q_{n-1})).$$

Finally, define $Q^* = \prod_{n \geq 0} Q_n$. Q^* is the set of generalized plans.

Consistency. For the first step, we will follow closely GP of consistent IHCPs, but with enough difference that it makes sense to define things outright. Formally, let $G_1 : \mathcal{A} \times \mathcal{K}(X \times Q_0) \to \mathcal{A}$ as the mapping $(a, \{x, q_0\}) \mapsto a$. Let $F_1 : Q_1 \to Q_0$ as the mapping $F_1 : q_1 \mapsto \big(E \mapsto q_1(G_1^{-1}(E))\big)$, for any $E \in \mathcal{B}(\mathcal{A})$. Therefore, for any $E \in \mathcal{B}(\mathcal{A})$, $F_1(p_1)(E)$ is the probability of event E in period 0 as implied by p_1; $F_1(p_1)$ is the distribution over period 0 actions implied by p_1. From here we can recursively define $G_n : \mathcal{A} \times \mathcal{K}(X \times Q_n) \to$

$\mathcal{A} \times \mathcal{K}(X \times Q_{n-1})$ as:

$$G_n : (a, \{x, q_{n-1}\}) \mapsto (a, \{x, F_{n-1}(q_0)\}))$$

and $F_n : Q_n \to Q_{n-1}$ as:

$$F_n : q_n \mapsto \big(E \mapsto q_n(G_n^{-1}(E))\big)$$

for any E in $\Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q_{n-1}))$. A consistent generalized plan is one such that

$$F_n(q_n) = q_{n-1},$$

for all n. Let Q denote all such generalized plans.

B.1.1 Construction Proofs

Lemma 5. *There exists a homeomorphism, $\lambda : Q \to \Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q))$ such that*

$$\text{marg}_{\mathcal{A}\times\mathcal{K}(X\times Q_{n-1})}(\lambda(q)) = q_n. \tag{B.1}$$

Proof. [Step 1: Extension Theorem.] Let $C_n = \big\{(q_0, \ldots q_n) \in \prod_{k=0}^{n} Q_k | q_k = F_{k+1}(q_{k+1}), \forall k = 1 \ldots n-1\big\}$, and $T_n = \mathcal{K}(X \times C_n)$ for $n \geq 0$. Let $T^* = \prod_{n=0}^{\infty} T_n$ and $T = \big\{t \in T^* | (\text{proj}_{T_n} t_{n+1} = t_n\big\}$. Let $Y_0 = \Delta^{\mathcal{B}}(\mathcal{A})$ and for $n \geq 1$ let $Y_n = \Delta^{\mathcal{B}}(\mathcal{A} \times T_0 \times \ldots \times T_n))$. We say the the sequence of probability measures $\{\nu_n \in Y_n\}_{n\geq 0}$ is consistent if $\text{marg}_{A\ldots T_{n-1}} \nu_{n+1} = \nu_n$ for all $n \geq 0$. Let Y^c denote the set of all consistent sequences. Then we know by Brandenburger and Dekel (1993), for every $\{\nu_n\} \in Y^c$ there exists a unique $\nu \in \Delta^{\mathcal{B}}(\mathcal{A} \times T^*)$ such that $\text{marg}_{\mathcal{A}} \nu = \nu_0$ and $\text{marg}_{A\ldots T_n} \nu = \nu_n$. Moreover, the map $\psi : Y^c \to \Delta^{\mathcal{B}}(\mathcal{A} \times T^*)$:

$$\psi : \{\nu_n\} \mapsto \nu$$

is a homeomorphism.

[Step 2: Extending Backwards.] Let $D_n = \big\{(t_0, \ldots t_n) \in \times_{n=0}^{n} T_n | t_k = \text{proj}_{T_n}(t_{k+1}), \forall k = 1 \ldots n-1\big\}$. Let $Y^d = \big\{\{\nu_n\} \in Y^c | \nu_n(\mathcal{A} \times D_n) = 1, \forall n \geq 0\big\}$. For each $q \in Q$, there exists a unique $\{\nu_n\} \in Y^d$, such that $\nu_0 = q_0$ and $\text{marg}_{\mathcal{A}\times\mathcal{K}(X\times Q_{n-1})}(\text{marg}_{\mathcal{A}\times T_{n-1}}(\nu_n)) = q_n$ for all $n \geq 1$. Indeed, let m_0, m_1 be the identify function on $\mathcal{A}$ and $\mathcal{A} \times \mathcal{K}(X \times Q_0)$, respectively.

Then for each $n \geq 2$ let $m_n : \mathcal{A} \times D_{n-1} \to \mathcal{A} \times \mathcal{K}(X \times Q_{n-1})$ as follows:

$$m_{n+1} : (a, \{x^0, q_0^0\}, \{x^1, q_0^1, q_1^1\} \dots \{x^n, q_0^n \dots q_n^n\}^n) \mapsto (a, \{x^n, q_n^n\}).$$

Note: for $n \geq 0$, each m_n is a Borel isomorphism. Indeed, continuity of m_n is obvious, and measurability follows immediately from the fact that canonical projections are measurable in the product σ-algebra. It is clear that m_n is surjective, and —since (given F_k for $k \in 1 \dots n$) q_n uniquely determines $q_0 \dots q_{n-1}$, which, (given the projection mappings) uniquely determines $T_0 \dots T_{n-1}$— m_i is also injective. As for, m_n^{-1}, continuity follows from the continuity of F_k for $k \in 1 \dots n$ and the projection mappings. Lastly, measurability of m_n^{-1} comes from the fact that a continuous injective Borel function is a Borel isomorphism (see Kechris (2012) corollary 15.2).

So, let $\psi : Q \to Y^d$ as the map

$$\varphi : q \mapsto \{E_n \mapsto q_n(m_n(E_n))\}_{n \geq 0},$$

for any $E_n \in \mathcal{B}(A \times T_0 \times \dots \times T_n)$. The continuity of φ and φ^{-1} follow from the fact that they are constructed from the pushforward measures of m_n^{-1} and m_n, respectively, which are themselves continuous (or, explicitly, see GP lemma 4).

Finally, let $\Gamma_n = \mathcal{A} \times D_n \times_{k=n+1}^{\infty} T_k$. Let $\nu = \psi(\{\nu_n\})$ for some $\{\nu_n\}$ in Y^d. Then $\nu(\Gamma_n) = \nu(\mathcal{A} \times D_n) = 1$. So, $\nu(\mathcal{A} \times T) = \nu(\cap_{n \geq 0} \Gamma_n) = \lim \nu(\Gamma_n) = 1$. Also, note, if $\nu(\mathcal{A} \times T) = 1$, then $\nu(\Gamma_n) = 1$ for all $n \geq 0$. So, $\nu \in Y^d$ if and only if $\nu(\mathcal{A} \times T) = 1$, i.e., if, $\psi(Y^d) = \big\{\nu \in \Delta^{\mathcal{B}}(\mathcal{A} \times T^*) | \nu(\mathcal{A} \times T) = 1\big\}$.

[Step 3: Extending Forwards.] Let τ denote the map from $\mathcal{A} \times \mathcal{K}(X \times Q) \to \mathcal{A} \times T$ as

$$\tau : \big(a, \{x, q\}\big) \mapsto \big(a, (\{x, q_0\}, \{x, q_0, q_1\}, \dots)\big)$$

That τ it is a bijection follows from the consistency conditions on Q, T, and C_n for $n \geq 1$. Now takes some measurable set $E \subseteq T$. Then $\tau^{-1}(E) = \bigcap_{n \geq 0} \big\{\{x, q_0, \dots q_n \times_{k=n_1^{\infty}} Q_k\} \in K(X \times Q^*)\big\}$, the countable intersection of measurable sets, and hence measurable. That τ and τ^{-1} are continuous is immediate. Therefore, by the same argument as in [Step 2], τ is a Borel isomorphism and $\kappa : \Delta^{\mathcal{B}}\big(\mathcal{A} \times T\big) \to \Delta^{B}\big(\mathcal{A} \times \mathcal{K}(X \times Q)\big)$,

$$\kappa : \nu \mapsto \big(E \mapsto \nu(\tau(E))\big)$$

for all E in $\Delta^{\mathcal{B}}\big(\mathcal{A} \times \mathcal{K}(X \times Q)\big)$. $\text{marg}_{\mathcal{A}}(\kappa(\nu)) = \text{marg}_{\mathcal{A}}(\nu)$ and $\text{marg}_{\mathcal{A} \times \mathcal{K}(X \times Q_{n-1})}(\kappa(\nu)) =$

$\mathrm{marg}_{\mathcal{A}\times\mathcal{K}(X\times Q_{n-1})}(\mathrm{marg}_{\mathcal{A}\times T_{n-1}}(\nu))$ for all $n \geq 1$.

Behold, $\lambda = \kappa \circ \psi \circ \varphi$ is the desired homeomorphism.

Definition. *Let $R_0 = Q_0$ and $R_1 = \{r_1 \in Q_1 | r_1(\mathcal{A} \otimes R_0) = 1\}$. Then, recursively let $R_n = \{r_n \in Q_n | r_n(\mathcal{A} \otimes R_{n-1}) = 1\}$. Set $R = \prod_{n=0}^{\infty} R_n$.*

Proof of Proposition 3.1. We show that λ is a homeomorphism between R and $\Delta^{\mathcal{B}}(\mathcal{A} \otimes R)$. Identify $\Delta^{\mathcal{B}}(\mathcal{A} \otimes R)$ with $\left\{\nu \in \Delta^{\mathcal{B}}(\mathcal{A} \times K(X \times Q)) | \nu(\mathcal{A} \otimes R) = 1\right\}$. Let $r \in R$. For each $n \geq 0$ let $\Gamma_n^r = \{(a, \{x, q\}) \in \mathcal{A} \otimes Q | q_k \in R_k,\ k = 0 \dots n\}$. Then $\lambda(r)(\Gamma_n^r) = \mathrm{marg}_{A\times K(X\times Q_n)}(\lambda(r))(\mathcal{A} \otimes R_n) = r_{n+1}(\mathcal{A} \otimes R_n) = 1$ for all $n \geq 1$. So $\lambda(r)(\mathcal{A} \otimes R) = \lambda(r)(\cap_{n\geq 0}\Gamma_n^r) = \lim \lambda(r)(\Gamma_n^r) = 1$. Now, fix $q \in Q$ with $\lambda(q)(\mathcal{A}\otimes R) = 1$, then $q_n(\mathcal{A}\otimes R_{n-1}) = \mathrm{marg}_{A\times K(X\times Q_{n-1})}(\lambda(q))(\mathcal{A}\otimes R_{n-1}) = \lambda(r)(\Gamma_n^r) \geq \lambda(r)(\mathcal{A}\otimes R) = 1$ for all $n \geq 0$ and so $q \in R$.

Definition. *For a metric space, M, let $\Delta(M) \subseteq \Delta^{\mathcal{B}}(M)$ denote the set of all distributions with countable support. I.e., for all $\nu \in \Delta(M)$, there exists a countable set S_ν such that $m \notin S_\nu \implies \nu(m) = 0$, and $\sum_{m\in S_\nu} \nu(m) = 1$.*

Definition. *Set $W : \mathcal{P}(R) \to \mathcal{P}(R)$ as the function:*

$$W : E \mapsto \{r' \in R | r' \in Im(f) \text{ for some } (a, f) \in \mathrm{supp}(\lambda(r)), r \in E\}$$

Definition. *Let $P_0 = \Delta(\mathcal{A})$ and $P_1 = \{p_1 \in R_1 | p_1 \in \Delta(\mathcal{A}\otimes P_0)\}$. Then, recursively let $P_n = \{p_n \in R_n | p_n \in \Delta(\mathcal{A} \otimes P_{n-1})\}$. Set $P = \left\{p \in \prod_{n=0}^{\infty} P_n | \prod_{n=0}^{\infty} \lambda(W^n(r)) \subset \prod_{n=0}^{\infty} \Delta(\mathcal{A} \otimes R)\right\}$.*

Proof of Theorem 3.2. We show that λ is a homeomorphism between P and $\Delta(\mathcal{A}\otimes P)$. First note, by construction, for all $r \in R$, $\lambda(r) \in \Delta^B(A \otimes W(r))$. Let $p \in P$; by the conditions on P, $\lambda(p) \in \Delta(\mathcal{A} \otimes R)$. Therefore, it suffices to show that for any $p \in P$, and $r \in W(p)$, $r \in P$. So fix some $r \in W(p)$. It follows from an analogous argument to Corollary 3.1 that $r \in \prod_{n=0}^{\infty} P_n$. Finally, note that $W^{n-1}(r) \subseteq W^n(r)$, for all $n \geq 2$.

B.2 LEMMAS

Lemma 6. *If $\succsim_h$ satisfies* vNM *and* IT*, then $\succsim_h$ satisfies the* sure thing principal*:*

[AX. B.1: STP] *For all $a \in \mathcal{A}$ and $f, f', g, g' : X \to P$, such that, for all $x \in X$, either (i) $f(x) = f'(x)$ and $g(x) = g'(x)$ or (ii) $f(x) = g(x)$ and $f'(x) = g'(x)$. Then,*

$$(a, f) \succsim_h (a, g) \iff (a, f') \succsim_h (a, g').$$

Proof. Assume this was not true and, without loss of generality, that $(a, f) \succsim_h (a, g)$ but $(a, g') \succ_h (a, f')$. Now notice, when mixtures are taken point-wise, $\frac{1}{2}f + \frac{1}{2}g' = \frac{1}{2}g + \frac{1}{2}f'$. Therefore,

$$\begin{aligned}\big(\frac{1}{2}(a, f) + \frac{1}{2}(a, g')\big) &\succ_h \big(\frac{1}{2}(a, g) + \frac{1}{2}(a, f')\big) \\ &\sim_h (a, \frac{1}{2}g + \frac{1}{2}f') = (a, \frac{1}{2}f + \frac{1}{2}g') \\ &\sim_h \big(\frac{1}{2}(a, f) + \frac{1}{2}(a, g')\big),\end{aligned}$$

where the first line follows from vNM, and the indifference conditions from IT. This is a contradiction.

Lemma 7. *If $\succsim_h$ satisfies* vNM *and* IT *for all $h \in \mathcal{H}$, then, if $h \overset{\mathcal{A}}{\sim} h'$ then $\succsim_h = \succsim_{h'}$.*

Proof. We will show the claim on induction by the length of the history. So let $h, h' \in \mathcal{H}(1)$ such that $h \overset{\mathcal{A}}{\sim} h'$. Therefore, $h = (p, (a, f), x)$ and $h' = (p', (a, g), x)$. Notice, by definition we have, $p = \alpha(a, f) + (1 - \alpha)r$ and $p' = \alpha'(a, g) + (1 - \alpha')r'$, for some $\alpha, \alpha' \in (0, 1]$ and $r, r' \in P$.

Let $q, q' \in P$; we want to show that $q \succsim_h q' \iff q \succsim_{h'} q'$. So let $q \succsim_h q'$, or by definition, $p_{-h}q \succsim p_{-h}q'$, which by the above observation is equivalent to

$$\alpha(a, f)_{-((a,f),(a,f),x)}q + (1 - \alpha)r \succsim \alpha(a, f)_{-((a,f),(a,f),x)}q + (1 - \alpha)r.$$

By independence (i.e., vNM) this is if and only if $(a, f)_{-((a,f),(a,f),x)}q \succsim (a, f)_{-((a,f),(a,f),x)}q'$, which by STP is if and only if $(a, g)_{-((a,g),(a,g),x)}q \succsim (a, g)_{-((a,g),(a,g),x)}q'$. Using independence again, this is if and only if $p'_{-h'}q \succsim p'_{-h'}q'$. This completes the base case.

So assume the claim holds for all histories of length n. So let $h, h' \in \mathcal{H}(n + 1)$ such that $h \overset{\mathcal{A}}{\sim} h'$. Therefore, $h = (h_n, p, (a, f), x)$ and $h' = (h'_n, p', (a, g), x)$, for some $h_n, h'_n \in H(n)$

such that $h_n \overset{\mathcal{A}}{\sim} h'_n$. By the inductive hypothesis $\succcurlyeq_{h_n} = \succcurlyeq_{h'_n}$.

Let $q, q' \in P$, and $q \succcurlyeq_h q'$, or by definition, $p_{-(p,(a,f),x)}q \succcurlyeq_{h_n} p_{-(p,(a,f),x)}q'$. By independence and the sure thing principle this is if and only if $(a,g)_{-((a,g),(a,g),x)}q \succcurlyeq_{h_n} (a,g)_{-((a,g),(a,g),x)}q'$, which by independence again (and the equivalence of $\succcurlyeq_{h_n}$ and $\succcurlyeq_{h'_n}$), is if and only if $p'_{-(p',(a,g),x)}q \succcurlyeq_{h'_n} p'_{-(p',(a,g),x)}q'$.

B.3 PROOF OF MAIN THEOREMS

Proof of Theorem 3.3. [Step 0: Value Function.] Since $\succcurlyeq_h$ satisfies vNM, there exists a $v_h : \mathcal{A} \otimes P \to \mathbb{R}$ such that

$$U_h(p) = \mathop{\mathbb{E}}_{p}\left[v_h(a,f)\right] \tag{B.2}$$

represents $\succcurlyeq_h$, with v_h unique un to affine translations.

[Step 1: Recursive structure.] To obtain the skeleton of the representation, lets consider $\hat{\succcurlyeq}$, the restriction of $\succcurlyeq$ to Σ (i.e., using the natural association between streams of lotteries and degenerate trees). The relation $\hat{\succcurlyeq}$ satisfies vNM (it is continuous by the closure of Σ in P). Hence there is a linear and continuous representation: i.e., an index $\hat{u} : X \times \Sigma \to \mathbb{R}$ such that:

$$\hat{U}(\sigma) = \mathop{\mathbb{E}}_{\sigma}\left[\hat{u}(x,\rho)\right] \tag{B.3}$$

unique upto affine translations.

Following Gul and Pesendorfer (2004), (henceforth GP), fix some $(x', \rho') \in \Sigma$. From SEP we have $\hat{U}(\frac{1}{2}(x,\rho) + \frac{1}{2}(x',\rho')) = \hat{U}(\frac{1}{2}(x,\rho') + \frac{1}{2}(x',\rho))$, and hence, $\hat{u}(x,\rho) = \hat{u}(x,\rho') + \hat{u}(x',\rho) - \hat{u}(x',\rho')$. Then setting $u(x) = \hat{u}(x,\rho') - \hat{u}(x',\rho')$ and $W(\rho) = \hat{u}(x',\rho)$, we have,

$$\hat{U}(\sigma) = \mathop{\mathbb{E}}_{\sigma}\left[u(x) + W(\rho)\right] \tag{B.4}$$

Now, consider $p' = (x', \rho)$. Notice that p' has unique 1-period history: $h = (p', p', x')$. By NT, h cannot be null. So, by SST, $\hat{\succcurlyeq}_h = \hat{\succcurlyeq}$. This implies, of course that $W = \delta\hat{U} + \beta$ for some $\delta > 0$ and $\beta \in \mathbb{R}$. Following Step 3 of Lemma 9 in GP exactly, we see that $\delta < 1$ and without loss of generality we can set $\beta = 0$:

$$\hat{U}(\sigma) = \mathop{\mathbb{E}}_{\sigma} \left[u(x) + \delta \hat{U}(\rho) \right] \tag{B.5}$$

Both representing $\hat{\succcurlyeq}$ and being unique up to affine translations, we can normalize each U_h to coincide with $\hat{U}$ over Σ.

[Step 2: The existence of subjective probabilities.] For each $a \in \mathcal{A}$ consider

$$\mathcal{F}(a) = a \otimes \Sigma$$

i.e., the elements of $\hat{P}$ that begin with action a and from period 2 onwards are in Σ. Associate $\mathcal{F}(a)$ with the set of "acts": $f : S_a \to \Sigma$, in the natural way. For any acts f, g let $f_{-x}g$ denote the act that coincides with f for all $x' \in S_a$, $x' \neq x$, and coincides with g after x. For each $h \in \mathcal{H}$, and acts $f, g \in \mathcal{F}(a)$, say $f \mathrel{\dot{\succcurlyeq}}_{h,a} g$ if and only if $(a, f) \succcurlyeq_h (a, g)$.

It is immediate that $\dot{\succcurlyeq}_{h,a}$ is a continuous weak order (where, as before, continuity follows from the closure of $\mathcal{F}$ in P). Further, $\dot{\succcurlyeq}_{h,a}$ satisfies independence. Indeed: fix $f, g, h \in \mathcal{F}(a)$ with $f \mathrel{\dot{\succcurlyeq}}_{h,a} g$. Then

$$\begin{aligned} f \mathrel{\dot{\succcurlyeq}}_{h,a} g &\implies (a, f) \succcurlyeq_h (a, g) \\ &\implies \alpha(a, f) + (1-\alpha)(a, h) \succcurlyeq_h \alpha(a, g) + (1-\alpha)(a, h) \\ &\implies (a, \alpha f + (1-\alpha)h) \succcurlyeq_h (a, \alpha g + (1-\alpha)h) \\ &\implies \alpha f + (1-\alpha)h \mathrel{\dot{\succcurlyeq}}_{h,a} \alpha g + (1-\alpha)h, \end{aligned}$$

where the third line uses IT. Lastly, $\dot{\succcurlyeq}_{h,a}$ satisfies monotonicity, a direct consequence of SST and STP. Hence, we have state-independence which gives us the full set of Anscombe and Aumann (1963) axioms for an SEU representation of $\dot{\succcurlyeq}_{h,a}$ with state space S_a. That is, a belief $\mu_{h,a} \in \Delta(S_a)$ and a utility index from $\Sigma \to \mathbb{R}$ (which is of course, $\hat{U}$, and so will be denoted as such), such that

$$\hat{V}_{h,a}(f) = \mathop{\mathbb{E}}_{\mu_{h,a}} \left[\hat{U}(f(x)) \right] \tag{B.6}$$

represents $\dot{\succcurlyeq}_{h,a}$.

[Step 3: Proportional Actions.] Now, fix some $h \in \mathcal{H}$ and consider an arbitrary $(a, f) \in \mathcal{A} \otimes P$. Let $\rho \in \Sigma$ be such that $\text{marg}_X \rho = \mu_{h,a}$. We claim, (a, f) and ρ are h-proportional.

Fix some $g, g' : X \to \Sigma$. From (B.6), we know

$$(a, g) \succcurlyeq_h (a, g') \iff \mathop{\mathbb{E}}_{\mu_{h,a}} \left[\hat{U}(g(x))\right] \geq \mathop{\mathbb{E}}_{\mu_{h,a}} \left[\hat{U}(g'(x))\right] \tag{B.7}$$

From (B.5) we have

$$\begin{aligned} \hat{U}(\rho.g) &= \mathop{\mathbb{E}}_{\rho} \left[u(x) + \delta \hat{U}(g(x))\right] \\ &= \mathop{\mathbb{E}}_{\mathrm{marg}_X \rho} \left[u(x) + \delta \hat{U}(g(x))\right] \\ &= \mathop{\mathbb{E}}_{\mu_{h,a}} \left[u(x)\right] + \delta \mathop{\mathbb{E}}_{\mu_{h,a}} \left[\hat{U}(g(x))\right] \end{aligned}$$

In corresponding fashion we obtain the analogous representation for $\hat{U}(\rho.g')$, and hence

$$\rho.g \succcurlyeq_h \rho.g' \iff \mathop{\mathbb{E}}_{\mu_{h,a}} \left[\hat{U}(g(x))\right] \geq \mathop{\mathbb{E}}_{\mu_{h,a}} \left[\hat{U}(g'(x))\right] \tag{B.8}$$

Combining the implications of (B.7) and (B.8), we see that (a, f) and ρ are h-proportional.

[Step 4: Proportional Plans.] We now claim that for any $h \in \mathcal{H}$ and $p \in P$ there exists some $\sigma \in \Sigma$ such that $p \sim_h \sigma$. Fix some $p \in P$, and for each $n \in \mathbb{N}$ define p^n to be any PoA that agrees with p on the first n periods, then provides elements of Σ unambiguously. Note that $p_n \to p$ point-wise and hence converges in the product topology. Therefore, the claim reduces to finding a convergent sequence $\{\sigma_n\}_{n \in \mathbb{N}} \subset \Sigma$ such that $\sigma^n \sim_h p^n$, as continuity ensures the limits are indifferent.

We will prove the subsidiary claim by induction. Consider p^1, for each $(a, f) \in supp[p^1]$, note, by assumption, $f : X \to \Sigma$. Let $\tau^{1,(a,f)} \in \Sigma$ be such that $\mathrm{marg}_X \tau^{1,(a,f)} = \mu_{h,a}$. By [Step 3], (a, f) and $\tau^{1,(a,f)}$ are h-proportional. And thus, $\tau^{1,(a,f)}.f \sim_h (a, f).f = (a, f)$, by PRP. Let $\sigma^1 \in \Sigma$ be such that $\sigma^1[E] = p^1[\{(a, f) | \tau^{1,(a,f)}.f \in E\}]$. Therefore,

$$\begin{aligned} U_h(p^1) &= \mathop{\mathbb{E}}_{p^1} \left[v_h(a, f)\right] \\ &= \mathop{\mathbb{E}}_{p^1} \left[\hat{U}(\tau^{1,(a,f)}.f)\right] \\ &= \mathop{\mathbb{E}}_{\sigma^1} \left[\hat{U}(\rho)\right] \\ &= \hat{U}(\sigma^1) \end{aligned}$$

where the third line comes from the change of variables formula for pushforward measures. This completes the base case.

Now, assume the claim hold for all h and $m \leq n - 1$ for some $n \in \mathbb{N}$. Consider p^n. Note that for all h' of the form $h(x) = (h, p^n, (a, f), x)$, the implied continuation problem $p^n(h')$ satisfies the inductive hypothesis. Therefore, there exists a $\sigma^{n-1,h'} \sim_{h'} p(h')$ for all such h'.

Let $\star$ denote the mapping: $(a, f) \mapsto (a, f)^\star = (a, x \mapsto \sigma^{n-1,h(a,x)})$, where $h(a, x) = (h, p^n, (a, f), x)$. By construction, for each (a, f) in $supp(p^n)$, and $x \in S_a$ we have $(a, f) \sim_h (a, f_{-x}\sigma^{n-1,h(a,x)})$ (using the notation from [Step 2]). Employing STP we have $(a, f) \sim_h (a, f)^\star$ (i.e., enumerating the outcomes in S_a and changing f one entry at a time, where STP ensures that each iteration is indifferent to the last).

Let $\hat{p}^n \in P$ be such that $\hat{p}^n[E] = p^n[\{(a, f)|(a, f)^\star \in E\}]$. So,

$$\begin{aligned} U_h(p^n) &= \mathop{\mathbb{E}}_{p^n}\big[v_h(a, f)\big] \\ &= \mathop{\mathbb{E}}_{p^n}\big[v_h((a, f)^\star)\big] \\ &= \mathop{\mathbb{E}}_{\hat{p}^n}\big[v_h(b, g)\big] \\ &= U_h(\hat{p}^n) \end{aligned}$$

Applying the base case to $\hat{p}^n$ concludes the inductive step. Notice also, the convergence of $\{\sigma^n\}_{n\in\mathbb{N}}$ is easily verified, following the fact that the marginals on p_n are fixed for any σ^m with $m \geq n$.

[Step 5: Representation.] Consider any $(a, f) \in \mathcal{A} \otimes P$. We claim that there exists an $(a, f') \in \mathcal{F}(a)$ such that $(a, f) \sim_h (a, f')$. Indeed, by [Step 4], for any $x \in S_a$, there exists some $\rho(a, x)$ such that $\rho(a, x) \sim_{h(a,x)} f(x)$, where $h(a, x) = (h, (a, f), (a, f), x)$. Define $f' \in \mathcal{F}(a)$ as $x \mapsto \rho(a, x)$. It follows from STP that $(a, f) \sim_h (a, f')$.

We know by [Step 3] that there exists a $\rho \in \Sigma$, h-proportional to (a, f), with $\text{marg}_X\rho = \mu_{h,a}$. Hence $(a, g) = (a, f).g \sim_h \rho.g$ for all $g : X \to \Sigma$. We have,

$$\begin{aligned} v_h(a, g) &= \hat{U}(\rho.g) \\ &= \mathop{\mathbb{E}}_{\mu_{h,a}}\big[u(x) + \delta\hat{U}(g(x))\big], \end{aligned}$$

and so, for (a, f'):

$$v_h(a, f') = \mathop{\mathbb{E}}_{\mu_{h,a}}\big[u(x) + \delta\hat{U}(\rho(a, x))\big].$$

By the indifference condition $\rho(a, x) \sim_{h(a,x)} f(x)$,

$$v_h(a, f) = \mathop{\mathbb{E}}_{\mu_{h,a}}\big[u(x) + \delta U_{h(a,x)}(f(x))\big]. \tag{B.9}$$

Notice, $h(a,x) \overset{\mathcal{A}}{\sim} h'(a,x) = (h, p, (a, f), x)$, so by Lemma 7, $\succcurlyeq_{h(a,x)} = \succcurlyeq_{h'(a,x)}$. Applying this fact, and plugging (B.9) into (B.2) provides

$$U_h(p) = \mathop{\mathbb{E}}_p \left[\mathop{\mathbb{E}}_{\mu_{h,a}} \left[u(x) + \delta U_{h'(a,x)}(f(x)) \right] \right] \tag{B.10}$$

as desired.

Proof of Theorem 3.6. First we show, if a strongly exchangeable process ζ over $\mathcal{S}$ is induced by an i.i.d distribution D over $\mathcal{S}_{\mathcal{A}}$, then it must be that the marginals of D (on $\{S_a\}_{a\in\mathcal{A}}$) are independent, that is $D \in \Delta^{IN}$. Indeed, consider two non-empty, disjoint collection of actions, $\hat{\mathcal{A}}, \hat{\mathcal{A}}' \subset \mathcal{A}$. Let $E, F \in S_{\hat{\mathcal{A}}}$, $E', F' \in S_{\hat{\mathcal{A}}'}$, be measurable events. Identify E^n with the cylinder it E generates in $\mathcal{S}$ when in the n^{th} coordinate: $E^n = \{s \in \mathcal{S} | s_{n,\mathcal{B}} \in E\}$. Since ζ is strongly exchangeable we have that

$$\zeta\left(E^n \cap E'^n \cap F^{n+1} \cap F'^{n+1}\right) = \zeta\left(E^n \cap F'^n \cap F^{n+1} \cap E'^{n+1}\right). \tag{2Sym}$$

We will refer to the latter weaker property as *two symmetry.* Now, since ζ is i.i.d generated by D, we have that (abusing notation by identifying E with the cylinder it generates in $S_{\mathcal{A}}$)

$$D(E \cap E') \cdot D(F \cap F') = D(E \cap F') \cdot D(F \cap E').$$

Substituting via the rule of conditional probability:

$$D(E|E') \cdot D(E') \cdot D(F|F') \cdot D(F') = D(E|F') \cdot D(F') \cdot D(F|E') \cdot D(E').$$

This implies that

$$\frac{D(E|E')}{D(E|F')} = \frac{D(F|E')}{D(F|F')}.$$

Since this is true for all events, we have that $D(E|E') = D(E|F')$ for every $E \in S_{\hat{\mathcal{A}}}$ and $E', F' \in S_{\hat{\mathcal{A}}'}$, implying $\hat{\mathcal{A}}$ and $\hat{\mathcal{A}}'$ are independent.

We now move to show that strong exchangeability is sufficient for the representation specified in the statement of the result. Since strong exchangeability implies exchangeability, we can apply de Finetti's theorem and represent the process ζ by

$$\zeta(\cdot) = \int_{\Delta(S_{\mathcal{A}})} \hat{D}(\cdot) d\psi(D).$$

We need to show that ψ's support lies in Δ^{IN}.

For $s \in \mathcal{S}$ and $t \in \mathbb{N}$ let s_t be the projection of s into the first t periods. Now, let $\zeta(\cdot|s_t) : S_{\mathcal{A}} \to [0,1]$ be the *one period ahead predictive probability*, given that the history of realizations in the first t periods is s_t. Since ζ is exchangeable, $\zeta(\cdot|s_t)$ converges (as $t \to \infty$) with ζ probability 1. Moreover, the set of all limits is the support of ψ. Denote the limit for a particular s by D_s. Of course, the exchangeability of ζ also guarantees that $\zeta(\cdot,\cdot|s_t) : S_{\mathcal{A}} \times S_{\mathcal{A}} \to [0,1]$, that is the *two period ahead predictive probability*, converges to $D_s \times D_s$. Furthermore, ζ is strongly exchangeable; the limit itself satisfies (2Sym), and the arguments above imply that $D_s \in \Delta^{IN}$ with ζ probability 1.

Proof of Theorem 3.7. Fix an SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T}\in\mathcal{T}}$. We first construct a pre-measure $\hat{\zeta}$ on the semi-algebra of cylinder sets. Fix any ordering over $\mathcal{A}$. Set $\hat{\zeta}(\emptyset) = 0$ and $\hat{\zeta}(\mathcal{S}) = 1$. Let $E \neq \mathcal{S}$ be an arbitrary cylinder, i.e., $E = \prod_{n\in\mathbb{N}}\prod_{a\in\mathcal{A}} E_{n,a}$, such that for only finitely many (n,a), is $E_{n,a} \neq S_a$. Clearly, there are a finite number of $a \in \mathcal{A}$ such that $E_{k,a} \neq S_a$ for any k. By the ordering on $\mathcal{A}$ denote these $a_1 \dots a_n$. For each a_i let m_i denote the number of components such that $E_{k,a_i} \neq S_{s_i}$, and for $j = 1 \dots m_i$, let $k_{i,j}$ denote the j^{th} such component. Finally, for each a_i, let π_{a_i} denote any permutation such that $\pi_{a_i}(k_{i,j}) = j + \sum_{i'<i} m_{i'}$. Consider $\hat{E} = \prod_{n\in\mathbb{N}}\prod_{a\in\mathcal{A}} E_{\pi_a(n),a}$, where $\pi_a = \pi_{a_i}$ if $a \in a_1 \dots a_n$ and the identity otherwise. That is, for $n \in 1 \dots m_1$, $\hat{E}_{n,a} = S_a$ for all a except a_1, for $n \in m_1 \dots m_1 + m_2$, $\hat{E}_{n,a} = S_a$ for all a except a_2, etc. Let $\mathbf{T}(E)$ denote the sequence such that $T_n = S_i$ for $\sum_{i'<i} m_{i'} < n \leq \sum_{i'\leq i} m_{i'}$. Again that is, for $n \in 1 \dots m_1$, $T_n = S_{a_1}$, for $n \in m_1 \dots m_1 + m_2$, $T_2 = S_{a_2}$, etc.

For the remainder of this proof, for any cylinder E, $\hat{E}$ denotes the corresponding cylinder generated by the above process, in which at most a single action is restricted in each period. Let $\mathbf{T}(E)$ denote any observable process which observes the sequence of restricted actions. Finally, for any cylinder, E, which is restricted in most one action each period, and any $\mathbf{T}$ which observes each restricted set, identify E the relevant event in $\mathbf{T}$. So, Set $\hat{\zeta}(E) = \zeta_{\mathbf{T}(E)}(\hat{E})$.

To apply the Carathéodory extension theorem for semi-algebras, we need to show that for any sequence of disjoint cylinders $\{E^k\}_{k\in\mathbb{N}}$ such that $E = \bigcup_{k\in\mathbb{N}} E^k$ is a cylinder, $\hat{\zeta}(E) = \sum_{k\in\mathbb{N}} \hat{\zeta}(E^k)$. Towards this, assume that E, E' are disjoint cylinders such that $E \cup E'$ is a cylinder. Then it must be that there exists a unique (n, a) such that $E_{n,a} \cap E'_{n,a} = \emptyset$ and for all other (m, b), $E_{m,b} = E'_{m,b}$. Indeed, if this was not the case, then there exists some (m, b) and some x such that (WLOG) $x \in E_{m,b} \setminus E'_{m,b}$. But then, for all $s \in E \cup E'$, $s_{m,b} = x \implies s_{n,a} \in E_{n,a} \neq (E \cup E')_{n,a}$ a contradiction to $E \cup E'$ being a cylinder. But this implies $\hat{E}$ and $\hat{E}'$ induce the same sequence of restricted coordinates, differing on the restriction of single coordinate, and therefore, $\mathbf{T}(E) = \mathbf{T}(E')$. This implies that $\hat{E} \cup \hat{E}' \subseteq \mathbf{T}(E)$. Since $\zeta_{\mathbf{T}(E)}$ is finitely additive, so therefore $\hat{\zeta}(E \cup E') = \zeta_{\mathbf{T}(E)}(\hat{E} \cup \hat{E}') = \zeta_{\mathbf{T}(E)}(\hat{E}) + \zeta_{\mathbf{T}(E)}(\hat{E}') = \hat{\zeta}(E) + \hat{\zeta}(E')$.

Since $\hat{\zeta}$ is finitely additive over cylinder sets, countable additivity follows if we show that for all decreasing sequences of cylinders $\{E^k\}_{k \in \mathbb{N}}$, such that $\inf_k \hat{\zeta}(E^k) = \epsilon > 0$, we have $\bigcap_{k \in \mathbb{N}} E^k \neq \emptyset$. But this follows immediately from the finiteness of S_a. Since $E^{k+1} \subseteq E^k$, it must be that $E^k_{n,a} \subseteq E^k_{n,a}$. But each $E^k_{n,a}$ is finite, hence compact, and nonempty, because $\zeta(E^k) \geq \epsilon$. Therefore $\bigcap_{k \in \mathbb{N}} E^k_{n,a} \neq \emptyset$. The result follows by noting that the intersection of cylinder sets is the cylinder generated by the intersection of the respective generating sets. Let ζ denote the unique extension of $\hat{\zeta}$ to the σ-algebra on $\mathcal{S}$.

That ζ is consistent with $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ is immediate. We need to show that ζ is strongly exchangeable. Let E be a cylinder. Let $\bar{\pi}_a$ denote a finite permutation for each $a \in \mathcal{A}$. Let $F = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\bar{\pi}_a(n),a}$. Let π_{a_i} denote the permutation given by the construction of $\hat{F}$. Then $\hat{F} = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\pi_a(\bar{\pi}_a(n)),a}$. In particular, this implies there exists some permutation π^* such that $\hat{F} = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\pi^*(n),a}$. By AA-SYM, $\zeta_{\mathbf{T}(\hat{E})}(\hat{E}) = \zeta_{\pi^* \mathbf{T}(\hat{E})}(\pi^* \hat{E}) = \zeta_{\mathbf{T}(\hat{F})}(\hat{F})$. Therefore, $\zeta(E) = \zeta(F)$ and so, by Theorem 3.6, ζ is strongly exchangeable.

Finally, the similar logic show that ζ is unique. Towards a contradiction, assume there was some distinct, strongly exchangeable ζ', also consistent with $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$. Then, since the cylinder sets form a π-system, there must be some cylinder such that $\zeta(E) \neq \zeta'(E)$. But, by strong exchangeability, $\zeta(\hat{E}) = \zeta(E)$ and $\zeta'(\hat{E}) = \zeta'(E)$, so $\zeta(\hat{E}) \neq \zeta'(\hat{E})$ –a contradiction to their joint consistency with $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$.

Proof of Theorem 3.8. Let $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in A}$ be an SEE structure for $\succcurlyeq$ that satisfies AA-SYM. Let $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ be the associated family of observable processes. Fix $\mathbf{T}$ and some n period history $h \in \mathbf{T}$. Let, $(a_1, x_1) \ldots (a_n, x_n)$, where for each $i \leq n$ let a_i is such that $T_i = S_{a_i}$ and x_i is the i^{th} component of h. This represents an $\mathcal{A}$-equivalence class of decision theoretic histories. In out standard abuse of notation, let h also denote this class of histories. Following this abuse, when it is not confusing to do so, let πh denote both the permuted statistical history and the $\mathcal{A}$-equivalence class represented by $(a_{\pi(1)}, x_{\pi(1)}) \dots (a_{\pi(n)}, x_{\pi(n)})$.

Fix some n-permutation π. Let p denote the PoA that assigns a_i in the i^{th} period with certainty. Let p' be the π-permutation of p. We have

$$\alpha = \zeta_{\mathbf{T}}(h) = \mu_{\emptyset,a_1}(x_1) \cdot \mu_{(a_1,x_1),a_2}(x_2) \cdots \mu_{(a_1,x_1,\dots,a_{n-1},x_{n-1}),a_n}(x_n.$$

Let $\sigma, \sigma' \in \Sigma$ be such that $U_h(\sigma) = 1$ and $U_h(\sigma') = 0$. Then, by (SEE) we have

$$p_{-n}(\alpha\sigma + (1-\alpha)\sigma') \sim (p_{-n}\sigma')_{-h}\sigma$$

so, by AA-SYM, we have,

$$p'_{-n}(\alpha\sigma + (1-\alpha)\sigma') \sim (p'_{-n}\sigma')_{-h'}\sigma$$

which implies, again by (SEE),

$$\alpha = \mu_{\emptyset,a_{\pi(1)}}(x_{\pi(1)}) \cdot \mu_{(a_{\pi(1)},x_{\pi(1)}),a_{\pi(2)}}(x_{\pi(2)}) \cdots \mu_{(a_{\pi(1)},x_{\pi(1)},\dots,a_{\pi(n-1)},x_{\pi(n-1)}),a_{\pi(n)}}(x_{\pi(n)}) = \zeta_{\pi\mathbf{T}}(\pi h).$$

Hence, $\zeta_{\mathbf{T}}(h) = \zeta_{\pi\mathbf{T}}(\pi h)$ as desired.

APPENDIX C

APPENDIX FOR CHAPTER 4

C.1 SUPPORTING RESULTS

Lemma 8. *$\bar{\varphi} \in \mathcal{L}(X)$ if and only if there exists some reduced assignment $\bar{\mu}$ and formula in $\varphi \in \mathcal{L}(\mathcal{X})$ such that $\bar{\varphi} = \bar{\mu}(\varphi)$.*

Proof. The proof is by induction on the construction of $\bar{\varphi}$. First enumerate the elements of $\mathcal{X}$. Assume $\bar{\varphi}$ is atomic in $\mathcal{L}(X)$. Then $\bar{\varphi} = \alpha x_1 \dots x_n$, for some (not necessarily distinct) $x_1 \dots x_n \in X$. So let φ be the formula $\alpha = a_1 \dots a_n$, with $a_1 \dots a_n$ (note, the a_i's are distinct elements); and μ, any assignment that extends the mapping $\mu : a_i \mapsto x_i$ for $i = 1 \dots n$. It is immediate that $\bar{\varphi} = \bar{\mu}(\varphi)$. In the other direction, if for some assignment μ, $\bar{\varphi} = \bar{\mu}(\varphi) = \alpha\mu(a_1) \dots \mu(a_n)$, then $\bar{\varphi}$ is an atomic $\mathcal{L}(X)$.

So assume the theorem holds for arbitrary $\bar{\varphi}$ and $\bar{\psi}$, with corresponding formulae φ, ψ, and assignments μ and ν. Then $\neg\bar{\varphi} = \neg\bar{\mu}(\varphi) = \bar{\mu}(\neg\varphi)$ and $\Box\bar{\varphi} = \Box\bar{\mu}(\varphi) = \bar{\mu}(\Box\varphi)$, by nature of the fact that μ acts only on variables. Proving both directions.

Next, let $k(\bar{\varphi})$ be the number of occurrences of elements of D in $\bar{\varphi}$ and $k(\bar{\psi})$ the number of occurrences of elements of D in $\bar{\psi}$. Then, let τ be any assignment that extends the mapping

$$\tau : \begin{cases} a_i \mapsto \mu(a_i) \text{ for } i = 1 \dots k(\bar{\varphi}) \\ a_i \mapsto \nu(a_i) \text{ for } i = k(\bar{\varphi}) + 1 \dots k(\bar{\varphi}) + k(\bar{\psi}). \end{cases}$$

Then, $\bar{\varphi} \wedge \bar{\psi} = \bar{\tau}(\varphi \wedge \tau[a_i/a_{i+k(\bar{\varphi})}])$. In the other direction, assume there exists some assignment μ and $\eta \in \mathcal{L}(\mathcal{X})$ such that $\bar{\mu}(\eta) = \bar{\varphi} \wedge \bar{\psi}$. Since, μ acts only on variables and one character at a time, it is immediate that η must be for the form $\varphi \wedge \psi$ where ψ and ψ are such that $\bar{\mu}(\varphi) = \bar{\varphi}$ and $\bar{\mu}(\psi) = \bar{\psi}$ and so, by the inductive hypothesis, $\bar{\varphi}, \bar{\psi} \in \mathcal{L}(X)$ and therefore so is $\bar{\varphi} \wedge \bar{\psi}$.

Finally, let $C = \{c \in \mathcal{X} | \mu(c) = x, c \text{ occurs in } \varphi\}$ (if x appears in $\bar{\varphi}$, then by the inductive hypothesis, C is non empty). Let $\zeta \in \mathcal{L}(\mathcal{X})$ be the formula that coincides with φ except all (and possibly no) free occurrence of a are replaced with free occurrences of $b \in \mathcal{X}$, where b does not occur (free or bound) in φ, and all (and possibly no) free occurrence of any $c \in C$ are replaced with free occurrences of a (notice there are no free occurrences of any $c \in C$ in ζ (except possible if $a \in C$)). Then let τ be an assignment that coincides with μ everywhere but for a and b, where $\tau(a) = x$ and $\tau(b) = \mu(a)$. Then,

$$\begin{aligned}
\bar{\tau}(\forall a \zeta) &= \forall a \bar{\tau}(\zeta)[\tau(a)/a] \\
&= \forall a \bar{\mu}(\varphi)[\tau(a)/a] \\
&= \forall a \bar{\varphi}[\tau(a)/a] \\
&= \forall a \bar{\varphi}[x/a],
\end{aligned}$$

where the first equality is definitional, since τ acts only on free variables and there are no free $c \in C$ in ζ, the second follows from the construction of ζ and τ, which ensures $\bar{\mu}(\varphi) = \bar{\tau}(\zeta)$, the third from the inductive hypothesis, and the fourth since $\tau(a) = x$.

In the other direction, assume there exists some assignment μ and $\eta \in \mathcal{L}(\mathcal{X})$ such that $\bar{\mu}(\eta) = \forall a \bar{\varphi}[x/a]$. Since, μ acts only on variables and one character at a time, it is immediate that η must be for the form $\forall a \varphi$, were $\bar{\mu}(\varphi) = (\bar{\varphi}[x/a])[a/\mu(a)]$. So let $\tau \sim_a \mu$ and $\tau(a) = x$. Then, $\bar{\tau}(\varphi) = (\bar{\varphi}[x/a])[a/\tau(a)] = (\bar{\varphi}[x/a])[a/x] = \bar{\varphi}$. So $\bar{\varphi} \in \mathcal{L}(X)$, by the inductive hypothesis, and so, $\forall a \bar{\varphi}[x/a] \in \mathcal{L}(X)$.

Lemma 9. *Fix some $\bar{\varphi} \in \mathcal{L}(X)$. Let (φ, μ) and (φ', μ') be such that $\bar{\mu}(\varphi) = \bar{\mu}'(\varphi') = \bar{\varphi}$. Then $(M, s) \models_\mu \varphi$ if and only if $(M, s) \models_{\mu'} \varphi'$*

Proof. The proof is by induction. First, assume $\bar{\varphi}$ is an atomic proposition, i.e., $\bar{\varphi} = \alpha x_1 \ldots x_n$. Then $\varphi = \alpha a_1 \ldots a_n$ and $\varphi' = \alpha b_1 \ldots b_n$, where $\mu(a_i) = \mu'(b_i) = x_i$. Then $(M, s) \models_\mu \varphi$ if any only if $(\mu(a_1) \ldots \mu(a_n), s) = (x_1 \ldots x_n, s) = (\mu'(b_1) \ldots \mu'(ba_n), s) \in \mathcal{V}(\alpha)$, if and only $(M, s) \models_{\mu'} \varphi'$.

So assume the result holds for all formulae of order n. Since reduced assignments preserve the structure of $\neg, \wedge, \Box$, we need only consider the case of $\forall a\bar{\varphi}[x/a]$. So let $\bar{\mu}(\forall a\varphi) = \bar{\mu}'(\forall a\varphi') = \forall a\bar{\varphi}[x/a]$, and $(M,s) \models_\mu \forall a\varphi$. Then for any $\mu^a \sim_a \mu$, $(M,s) \models_{\mu^a} \varphi$. Let $\mu'^a \sim \mu'$ be such that $\mu^a(a) = \mu'^a(a)$. Then notice $\bar{\mu}^a(\varphi) = \bar{\mu}'^a(\varphi')$ (this follows from the fact that when $\mu^a(a) = \mu'^a(a) = x$, it must be that $\bar{\mu}^a(\varphi) = \bar{\mu}'^a(\varphi') = \bar{\varphi}$). So by the inductive hypothesis, $(M,s) \models_{\mu'^a} \varphi'$, for all a-variants of μ'.

Lemma 10. *Let M be a model of $\forall \boldsymbol{K_T} \cup \boldsymbol{P}$. Then for any contingent plan, c_t, based on the contractable set (Λ, Γ, μ), in each state, c_t is either implicitly acceptable or implicitly unacceptable.*

Proof. Assume c_t is not implicitly acceptable in state s. Then,

$$(M,s) \models_\mu \neg K_0 \bigwedge_{\varphi\in\Gamma} \big(\varphi \implies \forall a P_t(b \succcurlyeq a)\big).$$

Utilizing $\psi \equiv \neg\neg\psi$, we have

$$(M,s) \models_\mu P_0 \neg \bigwedge_{\varphi\in\Gamma} \big(\varphi \implies \forall a P_t(b \succcurlyeq a)\big).$$

Applying De Morgan's Law, and the definition of $\implies$,

$$(M,s) \models_\mu P_0 \bigvee_{\varphi\in\Gamma} \neg\big(\neg\varphi \vee \forall a P_t(b \succcurlyeq a)\big),$$

and De Morgan's once more, and $\neg\exists a \equiv \forall a \neg$,

$$(M,s) \models_\mu P_0 \bigvee_{\varphi\in\Gamma} \big(\varphi \wedge \exists a \neg P_t(b \succcurlyeq a)\big).$$

Then, since, $\neg P_t \equiv K_t \neg$,

$$(M,s) \models_\mu P_0 \bigvee_{\varphi\in\Gamma} \big(\varphi \wedge \exists a K_t \neg(b \succcurlyeq a)\big).$$

and lastly, from the material equivalence of $\neg(b \succcurlyeq a)$ and $(a \succ b)$ (under $\boldsymbol{P}$, by Proposition 4.1), and the fact that K_t respects material equivalence,

$$(M,s) \models_\mu P_0 \bigvee_{\varphi\in\Gamma} \big(\varphi \wedge \exists a K_t(a \succ b)\big),$$

as desired.

C.2 SOUNDNESS AND COMPLETENESS RESULTS

Because propositions 4.1, 4.4, and 4.13 all relate to the soundness and completeness of particular axiomatizations they are grouped together. Moreover, since Proposition 4.4 is the most general result, it is proven first. The notation there introduced is used without reintroduction in the subsequent proofs.

Proof of Proposition 4.4. The proof of soundness is standard (with perhaps the A0 as the exception, which is immediate). Towards completeness, we will construct a canonical structure. To this end, define $\mathcal{L}^+(\mathcal{X},\mathcal{Y})$ and the extension of $\mathcal{L}^A(\mathcal{X})$ that contains as atomic formulae exactly the same predicates, and a set of variables that contains $\mathcal{X}$ but, in addition, countably many variables, $\mathcal{Y}$, not in $\mathcal{X}$. A set, $\Lambda \subset \mathcal{L}^+(\mathcal{X},\mathcal{Y})$ is *admissible* if for every $\varphi \in \mathcal{L}^+(\mathcal{X},\mathcal{Y})$ for every $a \in \mathcal{X}\cup\mathcal{Y}$, there exists some $b \in \mathcal{X}\cup\mathcal{Y}$ such that $(\varphi[[a/b]] \implies \forall a\varphi) \in \Lambda$.

Now, let S^c be the set of all admissible and maximally $\forall \boldsymbol{K_T} \cup$ A0 consistent sets of formulae in $\mathcal{L}^+(\mathcal{X},\mathcal{Y})$. Note, by Theorem 14.1 of Hughes and Cresswell (1996), if $\Delta \subset \mathcal{L}(\mathcal{X})$ is consistent then there exists an $s \in S$ such that $\Delta \subseteq s$. Also, define $s^{K_t} = \{\varphi \in \mathcal{L}^+(\mathcal{X},\mathcal{Y}) | K_t\varphi \in s\}$ and $s^{A_t} = \{\varphi \in \mathcal{L}^+(\mathcal{X},\mathcal{Y}) | A_t\varphi \in s\}$.

Define the canonical model $M^c = \langle S^c, X^c, \mathcal{V}^c, \{R^c_t\}_{t\leq T}, \{\succcurlyeq^c_s\}_{s\in S^c}, \{\mathcal{A}^c_t\}_{t\leq T}\rangle$, where S^c is defined as above, $X^c = \mathcal{X}\cup\mathcal{Y}$, $\mathcal{V}^c$ is defined by $(a_1 \dots a_n, s) \in \mathcal{V}^c(\alpha)$ if and only if $\alpha a_1 \dots a_n \in s$, R^c_t is defined by $sR^c_t s'$ if and only if $s^{K_t} \subseteq s'$, $\succcurlyeq^c_s = \{(a,b) | (a \succcurlyeq b) \in s\}$, and $\mathcal{A}^c_t(s) = s^{A_t}$ for all t. Finally, define the conical assignment as the identity, $\mu^c : a \mapsto a$.

We will now show, for any $s \in S^c$ and $\varphi \in \mathcal{L}^+(\mathcal{X},\mathcal{Y})$, $(M^c, s) \models_{\mu^c} \varphi$ if and only if $\varphi \in S$. This will complete the proof, because any $\forall \boldsymbol{K_T} \cup$ A0 consistent $\varphi \in \mathcal{L}(\mathcal{X})$ is contained in some $s \in S^c$, and hence satisfiable. The proof is by induction on the construction of φ. For the base case, note, for any atomic $\alpha a_1 \dots a_n$ we have $(M^c, s) \models_{\mu^c} \alpha a_1 \dots a_n$ if any only if $(\mu^c(a_1) \dots \mu^c(a_n), s) \in \mathcal{V}^c(\alpha)$ if and only if $(a_1 \dots a_n, s) \in \mathcal{V}^c(\alpha)$ if and only if $\alpha a_1 \dots a_n \in s$, as desired. Likewise, for any $(a \succcurlyeq b)$, $(M^c, s) \models_{\mu^c} (a \succcurlyeq b)$ if any only if $\mu^c(a) \succcurlyeq^c_s \mu^c(b)$ if and only if $a \succcurlyeq^c_s b$ if and only if $(a \succcurlyeq b) \in s$.

Assume this holds for arbitrary φ and ψ. The inductive step for the cases of $\neg\varphi$, $\varphi \wedge \psi$, $K_t\varphi$ and $\forall a\varphi$ are exactly as in the proof of Theorem 14.3 of Hughes and Cresswell (1996). So assume $(M^c, s) \models_{\mu^c} A_t\varphi$. This is if and only if $\bar{\mu}^c(\varphi) \in \mathcal{A}^c_t(s)$. Since μ is the identity, so is $\bar{\mu}$, and therefore, this is if and only if $\varphi \in \mathcal{A}_t^c(s) = s^{A_t}$ which is if and only if $A_t\varphi \in s$. Lastly, let $(M^c, s) \models_{\mu^c} E_t\varphi$. This is if and only if $(M^c, s) \models_{\mu^c} K_t\varphi$ and $(M^c, s) \models_{\mu^c} A_t\varphi$, so, by the induction hypothesis, if and only if $K_t\varphi \in s$ and $A_t\varphi \in s$. By the properties of maximally consistent sets, $(K_t\varphi \wedge A_t\varphi) \in s$, which by the validity of A0 is if and only if $E_t\varphi \in s$.

Proof of Proposition 4.1. The evaluation of propositions in $X \times X$ via of the semantic structure $\{\succcurlyeq_s\}_{s\in S}$ imposes no additional restrictions (namely, no restrictions on R); we could just as well begin with $\mathcal{V}$ and derive $\{\succcurlyeq_s\}_{s\in S}$ there from. In addition, Cmp and Trv, do not involve statements regarding K_t and so impose no restrictions on K_T, and therefore $\{R_t\}_{t\leq T}$. Hence, the result when $\mathcal{C}$ is restricted to be in $\{\mathrm{T}, 4, 5\}$ is as in the well known case (see for example Hughes and Cresswell (1996)).

So let $\{R_t\}_{t\leq T}$ be arbitrary. Let $M \in \mathcal{M}^{cmp}$. Let μ be arbitrary. Assume $(M, s) \models_\mu \neg(a \succcurlyeq b)$. By definition $(\mu(a), \mu(b)) \notin \succcurlyeq_s$, and by the completeness of $\succcurlyeq_s$ this implies $(\mu(b), \mu(a)) \in \succcurlyeq_s$, so, $(M, s) \models_\mu (b \succcurlyeq a)$: since μ was arbitrary, Cmp is valid in $\mathcal{M}^{cmp}$. So $\boldsymbol{K} \cup$ Cmp is sound with respect to M^{cmp}.

To show completeness, we construct the canonical structure, $M^{c:cmp}$. To this end, let $S^{c:cmp}$ denote the set of all maximally $\forall\boldsymbol{K_T} \cup$ Cmp consistent sets of formulae in $\mathcal{L}^+(\mathcal{X}, \mathcal{Y})$. Let the rest of the canonical model be defined as in the proof of Proposition 4.4. The result follows if $\succcurlyeq_s^{c:cmp}$ is complete for all s, since then $M^{c:cmp} \in \mathcal{M}^{cmp}$, implying any $\forall\boldsymbol{K_T} \cup$ Cmp consistent formula is satisfiable in $\mathcal{M}^{cmp}$. Fix, $(a, b) \in X^c \times X^c$. Since s is maximally consistent it contains Cmp and either $(a \succcurlyeq b)$ or $\neg(a \succcurlyeq b)$. If $(a \succcurlyeq b) \in s$ then $a \succcurlyeq_s^{c:cmp} b$ and we are done. If $\neg(a \succcurlyeq b) \in s$, then, since s contains Cmp and every instance of 1$\forall$ it contains $\neg(a \succcurlyeq b) \wedge (\neg(a \succcurlyeq b) \implies (b \succcurlyeq a)) \in s$ and consequently, $(b \succcurlyeq a)$. Therefore, $b \succcurlyeq_s^{c:cmp} a$. $\succcurlyeq_s^{c:cmp}$ is complete, as desired.

Now, let $M \in \mathcal{M}^{trv}$, and μ be arbitrary. Assume $(M, s) \models_\mu (a \succcurlyeq b) \wedge (b \succcurlyeq c)$. By definition, this implies $(\mu(a), \mu(b)), ((\mu(b), \mu(c)) \in \succcurlyeq_s$; by the transitivity of $\succcurlyeq_s$ this implies $(\mu(a), \mu(c)) \in \succcurlyeq_s$, so, $(M, s) \models_\mu (a \succcurlyeq c)$. Since μ was arbitrary, $\forall\boldsymbol{K} \cup$ Trv is sound with respect to M^{trv}.

Again, we construct the canonical structure, $M^{c:trv}$, as usual. Assume $a \succcurlyeq_s^{c:trv} b$ and $b \succcurlyeq_s^{c:trv} c$. So, $(a \succcurlyeq b), (b \succcurlyeq c) \in s$. Since s is maximally $\forall\boldsymbol{K} \cup$ TRV consistent it contains TRV and every instance of 1$\forall$, therefore $((a \succcurlyeq b) \wedge (a \succcurlyeq b)) \wedge (((a \succcurlyeq b) \wedge (b \succcurlyeq c)) \implies (a \succcurlyeq c)) \in s$. This implies, $(a \succcurlyeq c) \in s$. Therefore, $\succcurlyeq_s$ is transitive, as desired.

Proof of Proposition 4.13. Let $M \in \mathcal{M}^f$, and μ be arbitrary. Let $t \geq t'$. Assume $(M, s) \models_\mu K_t\varphi$. So, $(M, s') \models \varphi$ for all $s' \in R_t(s)$. In particular, $(M, s'') \models_\mu \varphi$ for all $s'' \in R_{t'} \subseteq R_t$. By definition, $(M, s) \models_\mu K_{t'}\varphi$. F is valid in $\mathcal{M}^f$; $\forall\boldsymbol{K_T} \cup$ F is sound with respect to M.

Towards completeness, we construct the canonical structure, $M^{c:f}$, as usual. The result follows if $R_{t'}(s) \subseteq R_t(s)$ is true for all $S^{c:f}$, and $t \leq t'$. So fix some $s \in S^{c:f}$, and let s' be such that $sR_{t'}^{c:f}s'$. By definition this implies $s^{K_{t'}} \subseteq s'$. Now, let $\varphi \in s^{K_t}$: by definition $K_t\varphi \in s$. Since s contains every instance of F, $(K_t\varphi \implies K_{t'}\varphi) \in s$, and consequently, $K_{t'}\varphi \in s$. By definition $\varphi \in s^{K_{t'}}$ Since φ was arbitrary, $s^{K_t} \subseteq s^{K_{t'}} \subseteq s'$, implying $sR_t^{c:f}s'$, as desired.

C.3 PROOFS OMITED FROM THE MAIN BODY OF TEXT

Proof of Theorem 4.7. Let $\Lambda \subset \mathcal{L}(X)$ be t-separable and contractable, and (Γ, μ) be such that $\Gamma \subset \mathcal{L}(\mathcal{X})$ and $\bar{\mu}$ defines a bijection between Λ and Γ such that there is no free occurrence of a in any formula of Γ. For each state s, let $x(s)$ denote any $\succcurlyeq_s$ maximal element, guaranteed to exists by BND. Denote by $\bar{\Lambda}$ the subset of Λ that are satisfied in some state: $\bar{\Lambda} = \{\bar{\varphi} \in \Lambda | (M, s) \models_X \bar{\varphi}, s \in S\}$. For each $\bar{\varphi} \in \bar{\Lambda}$, define, $\bar{c}_t(\bar{\varphi})$ to be any element of $\{x(s) | s \in S$, such that $(M, s) \models_X \bar{\varphi}\}$. Finally, let c_t be any extension of $\bar{c}_t$ to Λ. It remains to show c_t is acceptable.

It suffices to show there no state such that $(M, s) \models_\mu \neg \bigwedge_{\varphi \in \Gamma} \big(\varphi \implies \forall a P_0(c_t(\varphi) \succcurlyeq a)\big)$. So assume to the contrary, this was true for some state s'. Applying De Morgan's, $(M, s') \models_\mu \bigvee_{\varphi \in \Gamma} \big(\varphi \wedge \exists a K_t(a \succ c_t(\varphi))\big)$. By the definition of contractable, there is a unique $\bar{\psi} \in \Lambda$ (with corresponding $\psi \in \Gamma$) such that $(M, s') \models_X \bar{\psi}$, and so, it must be that $(M, s') \models_\mu$

$\psi \wedge \exists a K_t(a \succ c_t(\psi))$. Therefore, for some $\mu' \sim_a \mu$,

$$(M, s') \models_{\mu'} K_t(a \succ c_t(\psi)), \tag{C.1}$$

Further, by T, $(M, s') \models_{\mu'} (a \succ c_t(\psi))$. i.e., $c_t(\bar{\varphi})$ is not $\succcurlyeq_s$ maximal.

By the construction of c_t, (and the fact that $\bar{\psi} \in \bar{\Lambda}$) there must be some other state, s'', such that $c_t(\varphi)$ is $\succcurlyeq_{s''}$ maximal and $(M, s'') \models_X \bar{\psi}$. But then, by (the contrapositive of) t-separability, $s'^{K_t} = s''^{K_t}$. So, by (C.1), $(M, s'') \models_{\mu'} K_t(a \succ c_t(\psi))$, a contradiction, via T, to the $\succcurlyeq_{s''}$ maximality of $c_t(\varphi)$.

Proof of Proposition 4.8. Let M be such a model. Let $\sim_{K_t}$ be the equivalence relation on S defined by $s \sim_{K_t} s'$ if $s^{K_t} = s'^{K_t}$. Let $S/\sim_{K_t}$ denote the resulting quotient space of S, with elements $[s]$. Enumerate the elements of $S/\sim_{K_t}$. The proof is by induction on the number of elements in $S/\sim_{K_t}$. If $S/\sim_{K_t}$ contains a single element, any single tautological statement provides t-separable contractable set.

Assume the result hold for n, with the corresponding set $\Lambda_n = \{\bar{\lambda}_{1,n} \dots \bar{\lambda}_{n,n}\}$. Finally, let $S/\sim_{K_t}$ contain $n+1$ elements. By definition of $S/\sim_{K_t}$, it must be that for each $[s_i]$, $i \leq n$, there exists some statement φ_i, such that (abusing notation: $[s]$ denoting any of its elements),

1. $\bar{\varphi}_i \in [s_i]^{K_t} \setminus [s_{n+1}]^{K_t}$, or,
2. $\bar{\varphi}_i \in [s_{n+1}]^{K_t} \setminus [s_i]^{K_t}$.

So, for each $i \leq n$, define $\bar{\psi}_i = \neg K_t \bar{\varphi}_i$ if (1) holds, and $\bar{\psi}_i = K_t \bar{\varphi}_i$ if (2) holds. Define,

$$\bar{\lambda}_{n+1,n+1} = \bigwedge_{i \leq n} \bar{\psi}_i,$$

$$\bar{\lambda}_{i,n+1} = \bar{\lambda}_{i,n} \wedge \neg \bar{\lambda}_{n+1,n+1},$$

for $i \leq n$. We claim $\Lambda_{n+1} = \{\bar{\lambda}_{1,n+1} \dots \bar{\lambda}_{n+1,n+1}\}$ is a t-separable contractable state. So let (Γ, μ) be any set such that $\Gamma \subset \mathcal{L}(\mathcal{X})$ and $\bar{\mu}$ defines a bijection between Λ and Γ. Let $\lambda_{i,n+1}$ be the corresponding element to $\bar{\lambda}_{i,n+1}$. It must be that $\lambda_{n+1,n+1}$ is of the form $\bigwedge_{i \leq n} \psi_i$, and from the properties of μ, we know $\bar{\mu}(\psi_i) = \bar{\psi}_i$.

Towards contractibility, let M' be any model satisfying the conditions of the theorem, and s' any state thereof. We claim $(M', s') \models_\mu \neg(\lambda_{j,n+1} \wedge \lambda_{k,n+1})$ for $j \neq k$ and $j \neq n+1$.

Indeed, if $k = n+1$ this is immediate. If $k \neq n+1$ then the fact that the conjunction of any two distinct formulae in $\Lambda(n)$ is unsatisfiable provides the claim. Further, we claim $(M', s') \models_\mu \bigwedge_{j \leq n+1} \lambda_{j,n+1}$. If $(M', s') \models_\mu \lambda_{n+1,n+1}$ we are done, if not, then the validity of the disjunction of all formulae of $\Lambda(n)$ provides the claim.

Towards t-separability, assume $\bar{\psi}_i$ is of the form $\neg K_t \bar{\varphi}_i$ (i.e., (1) holds), and let φ be the corresponding formula of $\mathcal{L}^A(\mathcal{X})$ such that $\psi_i = \neg K_t \varphi$. Then $\bar{\varphi} \notin [s_{n+1}]^{K_t}$. So there does not exist any assignment μ' and φ' such that $\bar{\mu}(\varphi') = \bar{\varphi}_i$ and $(M, [s_{n+1}]) \models_{\mu'} K_t \varphi'$. In particular, this is true for μ and φ; therefore $(M, [s_{n+1}]) \models_\mu \psi_i$. Now, assume $\bar{\psi}_i$ is of the form $K_t \bar{\varphi}_i$ (i.e., (2) holds), again with corresponding φ. Then $\bar{\varphi} \in [s_{n+1}]^{K_t}$. So there exists some μ' and φ' such that $(M, [s_{n+1}]) \models_{\mu'} K_t \varphi'$. But, $\mu'(\varphi') = \varphi = \mu(\varphi)$, so by lemma 9, $(M, [s_{n+1}]) \models_\mu \psi_i$. Hence,

$$(M, [s_{n+1}]) \models_\mu \bar{\lambda}_{n+1,n+1}$$

A similar logic applies to show, for each $i \leq n$, $(M, [s_i]) \models_\mu \neg\psi_i$, and so, by the inductive hypothesis,

$$(M, [s_i]) \models_\mu \bar{\lambda}_{i,n+1}$$

so $\Lambda(n+1)$ is t-separable.

Proof of Theorem 4.11. Let c_t be articulable and implicitly acceptable –so (4.9) holds. So for every $s' \in R_0(s)$,

$$(M, s') \models_\mu \bigwedge_{\varphi \in \Gamma} \big(\varphi \implies \forall a P_t(c_t(\varphi) \succcurlyeq a)\big)$$

Since possibility implies explicit possibility, (i.e., $P_t \varphi \implies P_t^E \varphi$ is a theorem), we have

$$(M, s') \models_\mu \bigwedge_{\varphi \in \Gamma} \big(\varphi \implies \forall a P_t^E(c_t(\varphi) \succcurlyeq a)\big)$$

which implies,

$$(M, s) \models_\mu K_0 \bigwedge_{\varphi \in \Gamma} \big(\varphi \implies \forall a P_t^E(c_t(\varphi) \succcurlyeq a)\big). \tag{C.2}$$

Now, notice, $(M, s) \models_\mu A_0(a \succcurlyeq a)$ and $(M, s) \models_\mu A_0(b \succcurlyeq b)$ (by the assumption there are two distinct elements of D of which the DM is aware), and $(M, s) \models_\mu A_0 \varphi$ for every $\varphi \in \Gamma$ since

c_t is articulable. This implies by A↑, and A1,

$$(M,s) \models_\mu A_0 \bigwedge_{\varphi\in\Gamma} \big(\varphi \implies \forall a P_t^E(c_t(\varphi) \succcurlyeq a)\big). \tag{C.3}$$

Combining (C.2) and (C.3), provides the result.

Proof of Theorem 4.12. Assume this was not true. First note, since M admits articulable contingent plans, there must be at least some $x \in X$ such that $(x \succcurlyeq x) \in \mathcal{A}_0(s)$. Let $\mu(a) = x$. By A↑ and KA this implies that, for all $s' \in R_0(s)$

$$(M,s') \models_\mu A_0 P^E\left(\exists a \left(\neg A_0(a \succcurlyeq a) \wedge A_t(a \succcurlyeq a)\right)\right),$$

Hence, by our hypothesis, and A0, we have, for some $s' \in R_0(s)$

$$(M,s') \models_\mu E_t \neg\left(\exists a \left(\neg A_0(a \succcurlyeq a) \wedge A_t(a \succcurlyeq a)\right)\right),$$

so by A0 again, we have, for all $s'' \in R_0(s')$,

$$(M,s') \models_\mu \neg\left(\exists a \left(\neg A_0(a \succcurlyeq a) \wedge A_t(a \succcurlyeq a)\right)\right),$$

Clearly, this implies

$$\{(x \succcurlyeq x) : (x \succcurlyeq x) \in \mathcal{A}_t(s'')\} \subseteq \mathcal{A}_0(s'') = \mathcal{A}_{0'} = \mathcal{A}_0(s), \tag{C.4}$$

where the last two equalities come from KA and the fact that $s' \in R_0(s')$ for all $s' \in R_0(s)$, by T.

Order the states in $R_0(s)$, $s_1 \ldots s_n$. So, for each $s_i \in R_0(s)$, let $x(s_i)$ denote any $\succcurlyeq_{s_i}$ maximal element of $\{x : (x \succcurlyeq x) \in \mathcal{A}_t(s')\}$ guaranteed to exist by ABnd. Let μ be such that $\mu(a_i) = x(s_i)$ for each $i \leq n$. By (C.4), we have $(M,s) \models_\mu A_0(a_i \succcurlyeq a_i)$ and so by A1,

$$(M,s) \models_\mu A_0(a_i \succcurlyeq a_j), \tag{C.5}$$

for any $i,j \leq n$.

Now define for each $i \leq n$, define recursively $\bar\varphi_i = \bigwedge_{j\leq n}(x(s_i) \succcurlyeq x(s_j)) \wedge \neg \bigvee_{j<i} \bar\varphi_j$. By Proposition 4.1, $\succcurlyeq_s$ is a weak order and so clearly, $\Lambda = \{\bar\varphi_i | i \leq n\}$ forms the a contractable set. Moreover, by (C.5) and A↑, $c_t : \bar\varphi_i \mapsto x(s_i)$ is articulable. But, by construction, is clear c_t is explicitly acceptable, a contradiction.

Proof of Theorem 4.15. First, fix some any model, M of $\forall \boldsymbol{K_1} \cup \boldsymbol{P} \cup \mathrm{F} \cup S_1$. Note, by the completeness of $\geq$ we have $m \geq m$ for all m, so in order for the definition of an FS_1 generated preference to be well defined, we need dominance to be reflexive. Indeed, for each $s \in S$, $\succcurlyeq_s$ is a preference relation by Proposition 4.1. So for each $s \in S$ let $\bar{m}(s)$ be the $\succcurlyeq_s$-maximal element of m. So by definition of $\succcurlyeq_s$-maximal, and S_1, we have, for all $s \in S$, $(M,s) \models_X K_1 \bigwedge_{z\in m}(\bar{m}(s) \succcurlyeq z)$. Since this holds for all $s \in S$ (in particular, for all $s' \in R_0(s)$), and $\bigcup_{s\in S} \bar{m}(s) \subseteq m$, it follows that $(M,s) \models_X K_0 \bigvee_{y\in m} K_1 \bigwedge_{z\in m}(y \succcurlyeq z)$.

Now, to show sufficiency, we will prove that the representation (4.17) holds. Since $\geq$ is complete and transitive, there exists some $V : 2^X \to \mathbb{R}$ that represents it. Define $\xi(m) \equiv \{\max_{x\in m} u_{s'}(x)\}_{u_{s'}\in\mathcal{U}_{s,0}}$. So let Γ be any strictly increasing extension of the map: $\xi(m) \mapsto V(m)$. It remains to show that Γ is well defined. Indeed, if $\xi(m) = \xi(m')$, then for all $s' \in R_0(s)$, we have $\bar{m}(s') \sim_{s'} \bar{m}'(s')$, implying (via S_1),

$$(M,s') \models_X K_1 \bigwedge_{z'\in m'} (\bar{m}(s') \succcurlyeq z') \wedge K_1 \bigwedge_{z\in m} (\bar{m}'(s') \succcurlyeq z).$$

It follows that m s-dominates m' and that m' s-dominates m, so by the requirements of an FS_1 generated preference, $V(m) = V(m')$. Now if $\xi(m) > \xi(m')$ (i.e,. component wise inequality with some strict), we have likewise have for all $s' \in R_0(s)$, $\bar{m}(s') \succcurlyeq_{s'} \bar{m}'(s')$, (with some strict preference) implying (via S_1),

$$(M,s') \models K_1 \bigwedge_{z'\in m'} (\bar{m}(s') \succcurlyeq z'),$$

and for at least one state $s'' \in R_0(s)$,

$$(M,s'') \models \neg K_1 \bigwedge_{z\in m} (\bar{m}'(s'') \succcurlyeq z').$$

It follows that m strictly s-dominates m', so by the requirements of an FS_1 generated preference, $V(m) > V(m')$, as desired.

Towards necessity, we will construct the FS_1 model that generates $\geq$. So let $\geq$ satisfy the axioms of Kreps (1979), and so, the representation therein holds, (i.e., of the form of (4.17), with an arbitrary state space, Ω). It is easy to check the following model suffices, $S \cong \Omega$, $\mathcal{V}$ can be arbitrary, $R_0 = \Omega^2$, $R_1 = \bigcup_{\omega\in\Omega}(\omega,\omega)$, and for each ω, let $\succcurlyeq_\omega$ be the order generated by u_ω in the initial representation.

Proof of Theorem 4.16. First, assume m is the image of such a contingent plan: $m = \mathrm{Im}(c_t)$, with c_t based on (Λ, Γ, μ). By way of contradiction, assume m is strictly dominated by m'. So,

$$(M,s) \models_\mu \neg K_0 \bigvee_{b \in \mu^{-1}(m)} K_t \bigwedge_{a \in \mu^{-1}(m')} (b \succcurlyeq a),$$

which implies for some $s' \in R_0(s)$, and all $b \in \mu^{-1}(m)$ we have

$$(M,s') \models_\mu \neg K_t \bigwedge_{a \in \mu^{-1}(m')} (b \succcurlyeq a),$$

which in turn implies, for some $s'' \in R_t(s')$ and some $a \in \mu^{-1}(m')$ and all $b \in \mu^{-1}(m)$, we have $(M, s'') \models_\mu (a \succ b)$. Now S_t implies

$$(M,s') \models_\mu K_t(a \succ b).$$

or, $(M,s') \models_\mu \exists a K_t(a \succ b)$. Let $\bar{\psi}$ (with $\psi = \bar{\mu}^{-1}(\bar{\psi})$) denote the unique element of Λ such that $(M,s') \models_X \bar{\psi}$. By assumption, $c_t(\bar{\psi}) \in \mu^{-1}(m)$, so, $(M,s') \models_\mu \psi \wedge \exists a K_t(a \succ (c_t(\psi))$. This clearly implies

$$(M,s') \models_\mu \bigvee_{\varphi \in \Gamma} \varphi \wedge \exists a K_t(a \succ c_t(\varphi)).$$

Lastly, since s' is assessable from s at 0, we have a contradiction to the acceptability of c.

Now, assume m is an undominated menu. For each $s \in S$, let $x(s)$ denote any $\succcurlyeq_s$ maximal element in m. Let Λ be a t-separable set. The proof proceeds as in Theorem 4.7. Denote by $\bar{\Lambda}$ the subset of Λ that are satisfied in some state: $\bar{\Lambda} = \{\bar{\varphi} \in \Lambda | (M,s) \models_X \bar{\varphi}\}$. For each $\bar{\varphi} \in \bar{\Lambda}$, define, $\bar{c}_t(\bar{\varphi})$ to be any element of $\{x(s) | s \in S$, such that $(M,s) \models_X \bar{\varphi}\}$. Finally, let c_t be any extension of $\bar{c}_t$ to Λ. It remains to show c_t is acceptable.

Assume it was not. Then there exists some $s' \in R_0(s)$ such that,

$$(M,s') \models_\mu \bigvee_{\varphi \in \Gamma} \big(\varphi \wedge \exists a K_t(a \succ c_t(\varphi))\big). \tag{C.6}$$

By the definition of contractable, there must be some $\varphi \in \Gamma$ such that $(M,s') \models_\mu \varphi$. As shown in Theorem 4.7, it must be $c_t(\bar{\varphi})$ is $\succcurlyeq_{s'}$ maximal element in m. By (C.6), there must be some $\mu' \sim_a \mu$ such that $(M,s') \models_{\mu'} (a \succ c_t(\varphi))$, and so, by S_t, $(M,s') \models_{\mu'} K_t(a \succ c_t(\varphi))$. It is immediate that $m \cup \bar{\mu}'(a)$ strictly s-dominates m, a contradiction.

N.B. this argument only shows that m *contains* the image of an acceptable contingent plan. However, we can always add a list of logical contradictions to Λ in order to exhaust the remainder of m (which is finite by assumption).

BIBLIOGRAPHY

Ahn, David and Haluk Ergin (2010). Framing Contingencies. *Econometrica* **78**(2), 655–695.

Anscombe, Frank and Robert Aumann (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics* **34**(1), 199–205.

Auster, Sarah (2013). Asymmetric awareness and moral hazard. *Games and Economic Behavior* **82**, 503–521.

Bergemann, Dirk and Juuso Välimäki (2008). Bandit Problems. In: *The New Palgrave Dictionary of Economics* (Steven N. Durlauf and Lawrence E. Blume, Eds.). second ed.

Berry, Donald A and Bert Fristedt (1985). *Bandit problems: sequential allocation of experiments (Monographs on statistics and applied probability)*. Springer.

Board, Oliver J. and Kim Sau Chung (2011). Object-Based Unawareness: Axioms. *Working paper*.

Board, Oliver J., Kim Sau Chung and Burkhard C. Schipper (2011). Two models of unawareness: Comparing the object-based and the subjective-state-space approaches. *Synthese* **179**(SUPPL. 1), 13–34.

Bourgeois-Gironde, Sacha and Raphaël Giraud (2009). Framing effects as violations of extensionality. *Theory and Decision* **67**(4), 385–404.

Brandenburger, Adam and Eddie Dekel (1993). Hierarchies of Beliefs and Common Knowledge. *Journal of Economic Theory* **59**(1), 189–198.

Cooke, Kevin (2016). Preference Discovery and Experimentation. *Working paper*.

de Finetti, Bruno (1931). Funzione caratteristica di un fenomeno aleatorio.

de Finetti, Bruno (1937). La prévision: ses lois logiques, ses sources subjectives. In: *Annales de l'institut Henri Poincaré*. Vol. 7. pp. 1–68.

Dekel, Eddie, Barton L. Lipman and Aldo Rustichini (1998). Standard State-Space Models Preclude Unawareness. *Econometrica* **66**(1), 159–173.

Dekel, Eddie, Barton L. Lipman and Aldo Rustichini (2001). Representing Preferences with a Unique Subjective State Space. *Econometrica* **69**(4), 891–934.

Dillenberger, David, Juan Sebastián Lleras, Philipp Sadowski and Norio Takeoka (2014). A theory of subjective learning. *Journal of Economic Theory* pp. 1–31.

Dillenberger, David, R. Vijay Krishna and Philipp Sadowski (2015). Dynamic Rational Inattention. *Working paper.*

Epstein, Larry G and Kyoungwon Seo (2010). Symmetry of evidence without evidence of symmetry. *Theoretical Economics* **5**(3), 313–368.

Epstein, Larry G and Martin Schneider (2003). Recursive multiple-priors. *Journal of Economic Theory* **113**(1), 1–31.

Epstein, L.G. and M. Le Breton (1993). Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory* **61**, 1–22.

Ergin, Haluk and Todd Sarver (2010). A Unique Costly Contemplation Representation. *Econometrica* **78**(4), 1285–1339.

Evren, Özgür and Efe A. Ok (2011). On the multi-utility representation of preference relations. *Journal of Mathematical Economics* **47**(4-5), 554–563.

Fagin, Ronald and Joseph Y. Halpern (1988). Belief, Awarness and Limited Reasoning. *Artificial Intelligence* **34**(1988), 491–501.

Fagin, Ronald, Joseph Y. Halpern, Yoram Moses and Moshe Y. Vardi (2003). *Reasoning About Knowledge*. Vol. 1.

Feinberg, Yossi (2012). Games with unawareness. *Working paper.*

Filiz-Ozbay, Emel (2012). Incorporating unawareness into contract theory. *Games and Economic Behavior* **76**(1), 181–194.

Galanis, Spyros (2015). The value of information under unawareness. *Journal of Economic Theory* **157**, 384–396.

Galanis, Spyros (2016). The value of information in risk-sharing environments with unawareness. *Games and Economic Behavior* **97**, 1–18.

Grant, Simon and John Quiggin (2012). Induction reasoning about unawareness. *Economic Theory* pp. 1–39.

Grant, Simon and John Quiggin (2014). A preference model for choice subject to surprise. *Theory and Decision.*

Gul, Faruk and Wolfgang Pesendorfer (2004). Self-Control and the Theory of Consumption. *Econometrica* **72**(1), 119–158.

Halpern, Joseph Y. and Leandro C. Rêgo (2009). Reasoning about knowledge of unawareness. *Games and Economics Behavior*.

Halpern, Joseph Y. and Leandro C. Rêgo (2013). Reasoning about knowledge of unawareness revisited. *Mathematical Social Sciences*.

Heifetz, Aviad, Martin Meier and Burkhard C. Schipper (2006). Interactive unawareness. *Journal of Economic Theory* **130**(1), 78–94.

Heifetz, Aviad, Martin Meier and Burkhard C. Schipper (2008). A canonical model for interactive unawareness. *Games and Economic Behavior* **62**, 304–324.

Heifetz, Aviad, Martin Meier and Burkhard C. Schipper (2013). Unawareness, beliefs, and speculative trade. *Games and Economic Behavior* **77**(1), 100–121.

Hewitt, Edwin and Leonard J. Savage (1955). Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society* **80**(2), 470.

Hughes, G. E. and M. J. Cresswell (1996). *A New Introduction to Modal Logic*. Vol. 107. Routledge.

Hyogo, Kazuya (2007). A subjective model of experimentation. *Journal of Economic Theory* **133**(1), 316–330.

Kahneman, Daniel and Amos Tversky (1984). Choices, values, and frames. *American Psychologist* **39**(4), 341.

Karni, Edi and Marie-Louise Vierø (2013). "Reverse Bayesianism": A choice-based theory of growing awareness. *American Economic Review* **103**(7), 2790–2810.

Karni, Edi and Marie-Louise Vierø (2016). Awareness of Unawareness : A Theory of Decision Making in the Face of Ignorance. *Working paper*.

Karni, Edi and Zvi Safra (2014). A Theory of Stochastic Choice under Uncertainty. *Working paper*.

Karni, Edi, David Schmeidler and Karl Vind (1983). On State Dependent Preferences and Subjective Probabilities. *Econometrica* **51**(4), 1021–1031.

Kechris, Alexander (2012). *Classical descriptive set theory*. Vol. 156. Springer Science & Business Media.

Klibanoff, Peter, Sujoy Mukerji and Kyoungwon Seo (2013). Perceived Ambiguity and Relevant Measures. *Econometrica* **82**(5), 1–46.

Kochov, Asen (2010). The Epistemic Value of a Menu and Subjective States. *Working paper*.

Kochov, Asen (2015). A Behavioral Definition of Unforeseen Contingencies. *Working paper*.

Koszegi, Botond and Matthew Rabin (2006). A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics* **122**(4), 1133–1165.

Kreps, David M (1979). A Representation Theorem for "Preference for Flexibility". *Econometrica* **47**(3), 565–577.

Krishna, R. Vijay and Philipp Sadowski (2014). Dynamic Preference for Flexibility. *Econometrica* **82**(133), 655–703.

Lehrer, Ehud and Roee Teper (2015). Who is a Bayesian?. *Working paper.*

Li, Jing (2009). Information structures with unawareness. *Journal of Economic Theory* **144**(3), 977–993.

Luce, R. Duncan and Howard Raiffa (1957). *Games and Decisions: Introduction and Critical Survey.*. Wiley New York, NY.

Maskin, Eric and Jean Tirole (1999). Unforeseen Contingencies and Incomplete Contracts. *Review of Economic Studies* **66**, 83–114.

Minardi, Stefania and Andrei Savochkin (2015). Subjective Contingencies and Limited Bayesian Updating.

Modica, Salvatore and Aldo Rustichini (1994). Awareness and partitional information structures. *Theory and Decision* **37**(1), 107–124.

Modica, Salvatore and Aldo Rustichini (1999). Unawareness and Partitional Information Structures. *Games and Economic Behavior* **27**(2), 265–298.

Morris, Stephen (1996). The Logic of Belief and Belief Change: A Decision Theoretic Approach. *Journal of Economic Theory* **69**(35), 1–23.

Ok, Efe A. (2002). Utility Representation of an Incomplete Preference Relation. *Journal of Economic Theory* **449**(November), 429–449.

Ok, Efe A, Pietro Ortoleva and Gil Riella (2015). Revealed (P)Reference Theory. *American Economic Review* **105**(1), 299–321.

Piermont, Evan, Norio Takeoka and Roee Teper (2016). *Games and Economic Behavior* **100**, 69–94.

Quiggin, John (2015). The value of information and the value of awareness. *Theory and Decision* **80**(MAY), 167–185.

Riella, Gil (2013). Preference for Flexibility and Dynamic Consistency. *Journal of Economic Theory* **148**(6), 2467–2482.

Robbins, Herbert (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* **58**(5), 527–536.

Rubinstein, Ariel and Yuval Salant (2008). {(A,f)} Choice with Frames. *The Review of Economic Studies* **75**, 1287–1296.

Salinetti, Gabriella and Roger Wets (1979). On the convergence of sequences of convex sets in finite dimensions. *SIAM Review* **21**(1), 18–33.

Schervish, Mark J (2012). *Theory of statistics*. Springer Science & Business Media.

Schipper, Burkhard C. (2013). Awareness-dependent subjective expected utility. *International Journal of Game Theory* **42**(3), 725–753.

Schipper, Burkhard C. (2014). Preference-based unawareness. *Mathematical Social Sciences* **70**, 34–41.

Sen, Amartya (1993). Internal Consistency of Choice. *Econometrica* **61**(3), 495–521.

Sen, Amartya (1997). Maximization and the Act of Choice. *Econometrica* **65**(4), 745–779.

Sher, Shlomi and Craig R M McKenzie (2006). Information leakage from logically equivalent frames. *Cognition* **101**(3), 467–494.

Shmaya, Eran and Leeat Yariv (2015). Experiments on Decisions Under Uncertainty: A Theoretical Framework. *American Economic Review*.

Simonson, Itamar and Amos Tversky (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research* **29**(3), 281–295.

Tversky, Amos and Daniel Kahneman (1981). The Framing of Decisions and the Psychology of Choice. *Science*.

Tversky, Amos and Eldar Shafir (1992). Choice Under Conflict: The Dynamics of Deferred Decision. *Psychological Science*.

Tversky, Amos and Itamar Simonson (1993). Context-dependent preferences. *Management Science* **39**(10), 1179–1189.