Learning the Krepsian State: Exploration Through Consumption

Evan Piermont*, Norio Takeoka† and Roee Teper‡

February 13, 2015

PRELIMINARY DRAFT; COMMENTS ARE WELCOME

Abstract

In this paper, we axiomatize an infinite-horizon subjective state space model with responsive learning. In contrast to many subjective state space models, the resolution of uncertainty regarding the true state is an endogenous process that depends on the decision maker's past consumption. In addition, there need not be full resolution of uncertainty between periods. When a decision maker consumes outcomes, she learns her relative preference between them; after each consumption history, the decision maker's state space is a refinement of the previous state space. We identify the (set of possible) conditional preferences induced by each series of consumption.

1 Introduction

Uncertainty regarding the state of the world, or one's preferences, underlies almost every economic environment. In such settings, the agent might have the opportunity to actively explore in order to learn about this uncertainty. When the agent takes an action, she observes its consequence and gains better understanding of the underlying uncertainty. In such domains, it is intuitive that the agent's understanding (or lack thereof) is a function of her experience, and subsequently, of her pervious choices.

The canonical models of introspective uncertainty, which identify the set of states that the decision maker believes she may learn, ignore entirely the process by which the information about the state is acquired [Kreps, 1979, Dekel et al., 2001]. Implicitly, these models are static and assume that all uncertainty is realized and in a way that is independent of the decision made by the agent. More recently, Takeoka [2007] and Dillenberger et al. [2014] introduce (static) models of gradual learning, and Krishna and Sadowski [2014] provide the behavioral foundation for a dynamic version of Kreps' model that allow for tastes shocks. However, in these models, as in Kreps, learning is not *responsive*: learning unfurls identically irrespective of the choices made by the decision maker.

*University of Pittsburgh, Department of Economics, ehp5@pitt.edu
†Yokohama National University, Department of Economics, takeoka@ynu.ac.jp
‡University of Pittsburgh, Department of Economics, rteper@pitt.edu

In many circumstances, such as a DM acknowledging the possibility of taste shocks, this is perfectly reasonable. It is likely the DM's choice of restaurant in the morning will not affect her ability to discern her preference for fish and steak in the evening. However, if we consider the scenario of a DM learning her preference over as-of-yet unconsumed alternatives, unresponsive learning is inadequate. For instance, a novice researcher who is deciding on a project, but is uncertain about her particular talents. She obtains no information without experience; there is no "exogenous" signal that will indicate her skill set. However, if she explores different research directions it will become apparent in which area she is more talented. The idea of responsive learning is in line with models of exploitation and exploration such as Bandit problems [Robbins, 1952, Bergemann and Välimäki, 2006], search [Koulayev, 2010], rational inattention [Sims, 2003], business cycles [Bachman and Moscariniz, 2012], etc, in which economic agents have the opportunity to consume information, but in order to do so they must take some action, where different actions yield different information structures.

This paper puts forth a behavioral foundation for a model of *subjective and responsive learning*. That is, the decision maker's subjective information structure changes (only) in response to her previous choices. In particular, our model considers a learning process characterized in part by two novel constraints. First, the DM's uncertainty regarding her preference over a and b is resolved once a and b are consumed. Second, if the DM is initially uncertain about her preferences over a and b then this uncertainty is not resolved unless she consumes outcomes which are correlated with her preference over a and b. Learning takes place *only* as the result of consumption.

We consider a dynamic constrained choice environment that extends the Kreps framework. Each period the DM jointly chooses a consumption outcome and a constraint for the following period. Interpreted as such, there is a clear connection with standard dynamic programing—our main result makes this connection clear. We provide the behavioral characterization of a recursive utility function, over this domain, exhibiting subjective responsive learning. Each period the DM, given her current information structure and the menu she faces, jointly maximizes her (current period) consumption utility and (discounted) continuation value. The DM takes into account that her choice of consumption today may teach her about her preferences, and thus alter her information structure tomorrow. This representation captures the notion of responsive learning by allowing the continuation utility function and information structure to explicitly depend on previous choices.

It is worth noting that our identification, just as in Kreps [1979] and the other papers mentioned above, is of *anticipated* learning. Although the model is dynamic in the sense that the decision maker contemplates intertemporal tradeoffs, preferences are from the ex-ante point of view. The DM's preferences, when choosing constrained choice environments, reflect how she expects to learn and how she expects her preferences to change once learning has taken place.[1]

The domain is requisite for the identification of responsive subjective learning. We are able to identify from behavior, not only (1) the states the DM believes possible, but also (2) how the DM expects to learn conditional on a given path of consumption, and (3) how the decision maker anticipates her preferences, over both future consumption and subsequent information, will change after learning any one of the pieces

[1]It is entirely possible, although outside the focus of this model, that the DM may learn something that she did not initially consider, or after observing an expected piece of information, she interprets it in an unexpected way. While we plan on introducing *actual* evolution of preferences to a later version of this paper, our current domain does not allow for the observation of preference at different time periods.

of information she believes possible. The identification of (2) and (3) are novel. The main complication in accommodating responsive learning is identifying (3).

What makes this identification inherently difficult is that we are interested in states that are not only subjective but also conditional. To understand the DM's conditional preferences for information, it is first necessary to be able to condition on the relevant (subjective) state. This is not straight forward, as the information structure is entirely subjective. A major contribution of this paper is the construction of strategic plans: the ability of the modeler to identify conditional preferences on subjective events using only observable choices. The following section contains an example that clarifies these notions. A more elaborate example (and the formal analysis) for how to identify (3) is discussed in Section 3.2.

1.1 A Clarifying Example

A DM is considering her preference over phones and laptops. There are four alternatives: two laptops: Mac (m) and PC (p), and two phones: iPhone (i) and Android (a). The DM has preference over dynamic constrained choice problems. A menu is denoted by curly brackets $\{\cdot\}$ and we identify each single item menu with the associated object, $\{m\} \cong m$, when it is not confusing to do so. Dynamic constrained problems are represented by nested menus: $\{m\{a,i\}, pa\}$ is the tree in which the DM can choose between m and p the first period and, if she chooses m then she can choose between a and i the second period, but if she chooses p she receives a for sure the second period.

Although it is unobserved to the modeler, lets assume that the DM considers the following subjective state space:

$$S = \begin{cases} m \sim_1 i \succ_1 p \sim_1 a \\ p \sim_2 a \succ_2 m \sim_2 i \end{cases}$$

Now, lets assume DM exhibits the following preferences:

$$m \sim p \sim \{m, p\} \tag{i}$$

$$mp\{m, p\} \succ mpm \sim mpp \tag{ii}$$

$$mp\{mi, pa\} \sim mp\{m\{a, i\}, p\{a, i\}\} \succ mp\{ma, pi\} \tag{iii}$$

In the standard Krepsian framework, we would interpret (i) as the restriction that in *all* states the DM is indifferent between a Mac and a PC. However, recall that our model imposes that uncertainty is not resolved but through the act of consumption. Thus, (i) has the additional interpretation that the DM is unsure of her preference, but does not anticipate learning her preference between now and when she must choose out of the menu. We can use (ii) to differentiate between these cases.

We know that if there is any initial uncertainty regarding m and p it must be resolved after consuming m and p: if there is ever going to be a benefit to flexibility it must be present after mp. Ex-ante, the DM anticipates this, hence (ii) is sufficient to identify the existence of state 1 where $m \succ_1 p$, and state 2 where $p \succ_2 m$. The preferences in (ii) do not, however, tell us anything about the DM's preference regarding a and

i in such states. To fully characterize the state space we appeal to (iii) and the notion of strategic planning.

Since the right hand side of "$mp\{mi, pa\} \sim mp\{m\{a, i\}, p\{a, i\}\}$" is more flexible than the left (therefore weakly preferred), (iii) is dictating that this particular flexibility will not be utilized. However, since $mp\{mi, pa\} \succ mp\{ma, pi\}$ it cannot be that a and i are always indifferent. Hence, (iii) is really stating that the assignment $m \mapsto i$ and $p \mapsto i$ is *correct* in the sense that whenever $m \succ p$ it is also true that $i \succ a$ and vice versa. In other words, choosing m from $\{m, p\}$ indicates that i will be chosen from $\{a, i\}$ and so the additional flexibility from making two choices is redundant.[2]

This mapping corresponds to the (unobservable) state space. We call any such optimal assignment, which reduces all flexibility over dynamic constraints to flexibility over consumed outcomes, a *strategic plan*. Just as in this example, strategic plans allow us to identify conditional preferences. The existence of a strategic plan indicates that everything the DM has learned can be described as information about previously consumed outcomes. At any point, the DM's information structure is a function only of her previous actions and initial beliefs. This concepts is central in the behavioral foundation of responsive learning.

1.2 Related Literature

This paper considers learning when the domain of uncertainty (the state space) is subjective. The literature on subjective states began with Kreps [1979]. In the Kreps model, the set of future preference profiles the DM considers possible is identified by examining her preferences over *menus* of consumption prizes. This framework has since been extended by Dekel et al. [2001] (DLR) to menus of lotteries, where the unique set of cardinal utility functions can be identified. While these models are interpreted with a dynamic component, there is only a single period of consumption.

Recently there have been papers that embed the DLR setup in dynamic settings. Krishna and Sadowski [2014] provide an infinite horizon model[3] where each period the DM's utility is drawn from a distribution, depending on the current state. This induces the DM to have a preference for flexibility in each period. The model is not a model of learning, and in particular, it is not responsive; The state changes each period, and in addition, the information and the period-by-period resolution of uncertainty is unrelated to the choices made by the decision maker. Their representation has a recursive structure, lending itself to examine the intertemporal tradeoff between future flexibility and current period consumption. Our model allows for similar contemplation, with the added intertemporal consideration regarding the tradeoff between future information structures and current period consumption, as is standard in models of exploration and exploitation.

Higashi et al. [2014] consider a dynamic extension of DLR with specifying the set of subjective states to that of discount factors. They axiomatize a recursive representation where subjective states evolve over time according to history of past consumption. In their model, however, correlations across subjective states (or, conditional preferences of conditional preferences, etc) are excluded, which is in clear contrast with the methodological point presented here.

[2]For the sake of exposition, this example is too simple to reveal the necessity of the strategic planning construction as the same state space could be identified from preferences over lotteries a la Dekel et al. [2001]. We bring this example to familiarize the reader with the process of learning and main identification tools considered in this paper. Section 3.3 contains a more elaborate example where strategic planning is strictly necessary for identification.

[3]Krishna and Sadowski [2014] resort the structure developed in Gul and Pesendorfer [2004]. The choice domain in our model is closely related.

In a model of subjective learning, Dillenberger et al. [2014] examine a DM who has preferences over menu-time pairs. This allows for the identification, not only of the set potential preference profiles considered possible, but also of the way that the DM expects to learn over time. As in our model, at each period the DM in Dillenberger et al. [2014] consider a state space that is a refinement of the pervious periods' state spaces. However, in contrast to our model, the path of learning is not responsive, that is, does not depend on the choices of the DM. In addition, their model is a static one.

Technically speaking, the concept of strategic planning is related to the similarly named notion in Dillenberger et al. [2014]; the benefit to flexibility can be undone by an appropriate assignment of outcomes to "states." In both models, planning axioms provide the partitional structure to information. There are two key differences between the two planning axioms. First is that in our model the states are not observable, and so, the assignment procedure must infer the state from observable actions and not from primitives. Even if we assumed that states are objective, interim information itself is subjective, and conditioning on it still requires its identification from observable actions. Second, as discussed above, strategic planning in our model is needed not only to identify conditional preferences over outcomes but also conditional preferences over information. We rely on it also to identify how the DM expects to continue to learn: her (ex-ante) anticipated path of learning given some interim information structure.

The definition of strategic plan, as formally presented in Section 3.2, allows us to condition on the event in which the DM learns a particular piece of information even though the information structure is a subjective component of the DMs preference. While our particular identification strategy and its motivation are novel, the exercise has been independently investigated by Krishna and Sadowski in a supplementary appendix to Krishna and Sadowski [2014]. There methodology relies crucially on the Markovian transition of states and cannot accommodate conditioning on subjective states in the presence of responsive learning.

2 Framework and Notation

Let X be a be a finite set, with some metric d_X. The set X represents everything that the decision maker could consume and is fully observable. Each period the decision maker consumes one element of X; let lowercase letters near the beginning of the alphabet (i.e., a, b, c) denote elements of X. For any metric space, Y, let $\mathcal{K}(Y)$ denote the set of all compact subsets of Y endowed with the hausdorff-metric.

At the beginning of each period, the decision maker faces a menu from which he must jointly choose that period's consumption and the constraint regarding the following periods choice. We can construct this in a manner similar to Gul and Pesendorfer [2004]. A one period choice problem is collection of alternatives from which the DM must choose –an element of $\mathcal{K}(X)$. The set of all one period choice problems, Z_1, corresponds to the set of all compact subsets of X (i.e., $\mathcal{K}(X)$).

An N period choice problem is a collection of pairs consisting of an alternative to consume and an $N-1$ period choice problem. As such, the set of all N period choice problems, Z_N, corresponds to $\mathcal{K}(X \times Z_{N-1})$. Note that compactness is defined in the product topology over $X \times Z_{N-1}$ (where the topology of Z_{N-1} is itself the product topology defined similarly). We can continue this iteration so as to define $\hat{Z} = \times_{n=1}^{\infty} Z_n$. We will restrict ourselves to the set of *consistent* elements of $\hat{Z}$, where $z \in \hat{Z}$ is consistent if for all n, $proj_{Z_{n-1}} z_n = z_{n-1}$. Let $Z \subset \hat{Z}$ be the restriction of $\hat{Z}$ to consistent elements. Call Z the set of Infinite Horizon Choice Problems (IHCPs). The primitive of the model is a preference relation, $\succcurlyeq$, over elements of

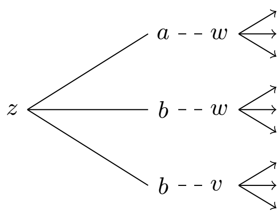

Figure 1: Representation of an IHCP z as a tree of constraints. Note that in the first period, two pairs offer w as the next period IHCP and two pairs offer b as the consumption prize.

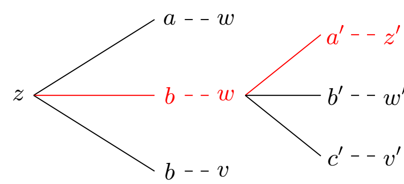

Figure 2: The red path corresponds the the feasible sequence $\sigma = ba'$, and $z_\sigma = z'$. The IHCP $z_{-\sigma}z''$ would be identical to the above except with z' replaced with z''.

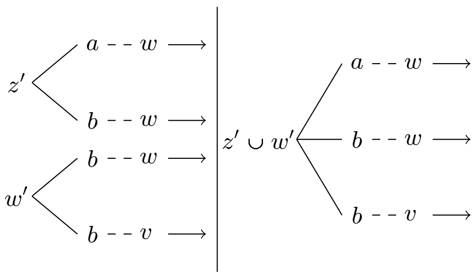

Figure 3: Representation of IHCPs z' and w', and their union $z' \cup w'$. Note, (b, w) is available both in z' and w' and appears only one in $z' \cup w'$, and $z' \cup w' = z$ as shown in Figure 1.

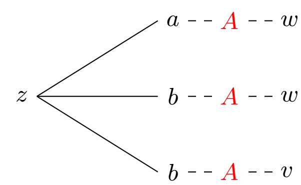

Figure 4: The IHCP $z_{-n}Az_n$ for $n = 1$.

Z. One can show, again following Gul and Pesendorfer [2004], that there is a natural identification between Z and $\mathcal{K}(X \times Z)$.

2.1 Notation

By the identification between Z and $\mathcal{K}(X \times Z)$, we can describe each z as a collection in $X \times Z$: $z = \{a_i, z(a_i)\}_{i \in M(z)}$ for some index set $M(z)$.[4] We call $M(z)$ the menu induced by z. Each element, $i \in M(z)$, is a pair: a consumption today, $a_i \in X$, and a continuation problem tomorrow, $z_i \in Z$. Since z_i is also an IHCP, we can write it too as a menu of pairs of consumptions and continuation problems: $z_i = \{(b_j, z_j)\}_{j \in M(z_i)}$. It is helpful to think of an IHCP as an infinite tree, where each node is an IHCP and the branches emanating from a node correspond to the elements of the induced menu, see Figure 1.[5]

As a prerequisite, we need to develop a notation of "tree-traversal", in order to identify particular branches of a given IHCP. For a given IHCP, $z \in Z$, we can define the set of all feasible (decision) sequences of length n, $n(z)$. Given an IHCP, z, a decision is feasible if it is an element of the implied menu: $i \in M(z)$ with the corresponding continuation tree is z_i. A feasible decision sequence of length n, for IHCP z, is a sequence of

[4] $M(z)$ is identified with z by considering it to be a collection of "singleton for today followed by an IHCP for tomorrow". For convenience we will refer to $M(z)$ as the set of indexes of such trees.

[5] In Figures 1-4, we have suppressed the $M(z)$ indexing notation in order to make the illustrations cleaner (both visually and conceptually). The indexes of $M(z)$ correspond to the branches emanating from IHCPs. It is worth noting that the *true* description of σ in Figure 2 is actually $\sigma = b_2 a'_1$ (under the obvious index) so as to distinguish it from a different path (for example if a' was available in v).

Notation	Meaning	Notes
a, b, c	Single period consumption prizes	
z, w, v	IHCPs	
σ, ρ, π	Streams of consumption	Finite unless otherwise noted, also represents feasible sequences when subscripted as in z_σ or $z_{-\sigma}$
W, V	Collections of IHCPs	Identified with their unions (and so also IHCPs)
Az	A finite choice problem followed by an IHCP	
$\boldsymbol{\sigma}, \boldsymbol{\rho}$	Degenerate IHCPs	Specify a predetermined (infinite) sequence of consumption

n choices, where the m^{th} choice is a feasible choice out of the continuation tree implied by the $(m-1)^{th}$ choice and the first choice feasible out of $M(z)$. In other words, each decision sequence corresponds to a particular path up the tree. We will denote such sequences with lower case greek letters (σ, ρ). Note that we will also use the same notation for *unindexed* streams of consumption. We do not anticipate this to be an issue as the two are not used in the same context, and it eases the notation burden; σ generically represents a stream of consumption (indexed or not).

Using this notation we can define "tree-pruning" operations. For some z, let σ be a finite, feasible decision sequence. Then we can define z_σ as the continuation tree after choosing σ. From this, we can define $z_{-\sigma}w$ as the IHCP that agrees with z everywhere but after the feasible sequence of decision σ, where z_σ is replaced by w. See Figure 2. Similarly, for an IHCP z, a finite feasible decision sequence σ and a finite tree A, let $z_{-\sigma}Az_\sigma$ be the IHCP in which, if the path σ was chosen then the continutaion probalem is Az_σ. That is, $z_{-\sigma}$ was concatenated with Az_σ. Finally, we denote by $z_{-n}Az_n$ the IHCP resulting from the concatenation of $z_{-\sigma}$ with Az_σ for *every* feasible sequence of length n (that is, $\sigma \in n(z)$), see Figure 4.

Lastly, we need to define notation for "tree-amalgamating" operations. For any two IHCPs, $z, w \in Z$, we can define their union, $(z \cup w)$ as the IHCP such that $M(z \cup w) = M(z) \cup M(w)$. See Figure 3. We will refer to (finite) collections of IHCPs with capital letters near the end of the alphabet (W, V). In an abuse of notation, we will identify each collection of IHCPs with the IHCP that is the union of its elements: the collection W can be treated as the IHCP: $W = \bigcup_{i \in I} w_i$. With this identification, any IHCP, z, that does not have a degenerate first period choice, is also a collection of smaller IHCPs. As such, the distinction between collections and individual IHCPs is not strict (i.e., for every W there is an equivalent w) but the notation will be helpful later.

For any finite choice problem, $A \in Z_n$, we can define Az as the IHCP that assigns z as a continuation tree to every feasible decision sequence of length n (contained in A). In particular, az is the IHCP which begins with consumption of $a \in X$ followed by the IHCP z. Likewise, for some finite sequence of consumption, $\sigma \in \bigtimes_{n=1}^{N} X$, σz is the IHCP which begins with consumption σ followed by z.

The use of bold lettering signifies degenerate IHCPs – predetermined sequence of consumption. So for each $\sigma \in \bigtimes_{n=1}^{\infty} X$, we identify $\boldsymbol{\sigma}$ as the IHCP such that there is only one feasible decision sequence of any given length, where the unique sequence coincides with the first n terms of σ.

2.2 Histories

At any given node in an IHCP we can define the set of consumption prize that must have been consumed to reach that node. Formally, define the history at a node as the support[6] (in X) of the (unique) decision sequence that leads to that node. In other words, the history does not remember the index but only remembers the consumption (i.e., ignores the menu from which the consumption was chosen, the order of consumption, and the number of times a prize has been consumed). We assume as an axiom 0 that the preference for continuation trees depends *only* on the history. As such, we can define the history dependent preferences.

Definition. *For each h let $\succcurlyeq_h$ denote the projection of preferences after history h defined by*

$$z \succcurlyeq_h w \iff hz \succcurlyeq hw$$

The DM's preference is entirely from the ex-ante view point. Her preference $\succcurlyeq_h$ is how she *expects* to behave after consuming h.

2.3 Discussion

Although our axioms are explicitly regarding preferences from an ex-ante point of view, the framework is flexible enough so that we can infer the decision maker's preference after hypothetical events. This is, in fact, the crux of this exercise; we are to understand how the decision maker *expects* to learn after a given consumption.

In particular, there are two types of hypothetical preferences in which we are interested: history induced ($\succcurlyeq_h$, defined above) and information induced ($\succcurlyeq_{h|t}$, not yet formally defined). The former refers to the decision maker's (anticipated) preference conditional on a path of consumption (the history). The later, more subtly, refers to the decision maker's preferences conditional on learning a particular bit of information (being in a particular cell of the partition induced by the history). Neither of these preferences are directly observable, but they are indirectly identifiable from the preferences. To see the relation, consider the following conceptual derivation of representation.

Let S capture the uncertainty the DM faces regarding her preference. For practical purposes we think of S as being the possible preference orderings that the DM believes are possible. Each hypothetical history induces a partition of this state space: when the DM consumes a new element she anticipates learning how to identify different sets of states, thereby refining her partition. While from her ex-ante point of view the decision maker can only identify the pieces of information she *could* learn (the partition over the state space), she anticipates that when making a decision, she will have learned a particular piece of information (the cell of the partition). Her preference of flexibility is the result of this anticipated learning.

Consider a DM who is ranking contingent IHCPs to be consumed after consuming the history h. She is aware that when she chooses a path of consumption from an IHCP she will have the information structure induced by h: the partition $\Sigma(h)$. This imparts an ex-ante preference for flexibility. Hence her value for

[6]The support of a (finite) decision problem, and in particualr a feasible sequence, is all the elements of X that can be consumed in that problem.

IHCPs after history h, V_h, takes the form,[7]

$$V_h(z) = \sum_{t \in \Sigma(h)} \max_{i \in M(z)} \sum_{t' \in \Sigma(h \cup a_i | t)} \big(u_{t'}(a_i) + \delta V_{h \cup a_i | t'}(z_i)\big). \tag{4}$$

The DM is able to make distinct choices contingent on the resolution of uncertainty embodied by the information provided by the consumption and h and a. As the DM consumes, both the information structure and the conditioning events evolve. The information structure is refined, $\Sigma(h) \to \Sigma(h \cup a)$, as a direct result of the addition of new elements to the history. In turn, the conditioning events become more informative, $S \to t'$, as the DM utilizes this available information in each period's consumption choice.

Note that the continuation value does not take the same form. This is because, when choosing from z she utilized her available information, so subsequent choices will take into account this information. Hence, any future choice must be conditioned on being in the true cell of the partition of $\Sigma(h)$. To make $V_{h \cup a | t'}$ well defined we need to consider the DM's valuation when she conditions on a particular piece of information t from the information $\Sigma(h)$ she expects to learn having consumed h. The preference, $\succcurlyeq_{h|t}$, with information structure $\Sigma(h)$ and conditioning on $t \in \Sigma(h)$, can be represented recursively by the following *Additive Learning Through Consumption* functional form:

$$V_{h|t}(z) = \max_{i \in M(z)} \sum_{t' \in \Sigma(h \cup a_i | t)} \big(u_{t'}(a_i) + \delta V_{h \cup a_i | t'}(z_i)\big). \tag{A-LTC}$$

And thus,

$$V_h(z) = \sum_{t \in \Sigma(h)} V_{h|t}(z).$$

In addition, the information dependent utility indices have the following aggregation property: if $h \subset h'$ then $\Sigma(h')$ is a refinement of $\Sigma(h)$ and $u_t = \sum_{t' \in \Sigma(h'|t)} u_{t'}$, for every $t \in \Sigma(h)$. This implies that (4) can be written as[8]

$$V_h(z) = \sum_{t \in \Sigma(h)} \max_{i \in M(z)} \big(u_t(a_i) + \delta V_{h \cup a_i | t}(z_i)\big). \tag{5}$$

The focus of this draft is on a more general non additively separable case, both across pieces of information and the consumption-learning utilities, which is presented formally below. While the axiomatization presented here corresponds to the more general case, a similar axiomatic structure can provide the foundation for the additive representation. To do so, one needs to consider a richer environment as in Gul and Pesendorfer [2004] and include an independence axiom. This will be included in the next draft of this paper.

[7] When consumption experience increases then the expected information structure becomes finer, that is, $\Sigma(h \cup a)$ is a finer partition than $\Sigma(h)$. Moreover, $\Sigma(h \cup a|t)$ stands for the partition of the cell t generated by the consumption of a.

[8] Similar to equation A-LTC one can define history dependent preferences, conditional on a piece of information coarser than that generated by the history (that is, not utilizing all possible information). For history $h \subseteq h'$ and $t \in \Sigma(h)$ let $V_{h'|t}(z) = \sum_{t' \in \Sigma(h'|t)} \max_{i \in M(z)} \sum_{\theta \in \Sigma(h' \cup a_i | t')} \big(u_\theta(a_i) + \delta V_{h' \cup a_i | \theta}(z_i)\big)$.

The non-additive representation takes the following form:

$$V_h(z) = \phi_h\big(\max_{i\in M(z)} F[u_t(a_i), V_{h\cup a_i|t}(z_i)]_{t\in\Sigma(h)}\big). \tag{6}$$

where ϕ_h (that aggregates information dependent utilities) and F (that aggregates consumption and exploration utilities) are strictly increasing aggregators. The preferences following history h and conditional on the information $t \in \Sigma(h)$ are represented by the *Learning Through Consumption* functional:

$$V_{h|t}(z) = \max_{i\in M(z)} F[u_t(a_i), V_{h\cup a_i|t}(z_i)], \tag{LTC}$$

where the aggregation property of information dependent utilities takes the form $u_t = \phi_{h'|t}(u_{t'})_{t'\in\Sigma(h'|t)}$ for every $t \in \Sigma(h)$ and $h' \supset h$.

3 Axioms

3.1 Ex-ante Preferences

We assume that the following axioms hold for all h. To achieve the basic utility structure we have require that $\succcurlyeq$ is a continuous weak order:

A1. (Weak Order). *The binary relation $\succcurlyeq$ is a continuous weak order.*

The first substantive axiom governs the learning process. We are assuming that the preference does not change unless the decision maker learns something (or, more accurately, anticipates learning). If the decision maker has already consumed a then consuming a again will not teach her anything new, and her preference will not change. Hence, we postulate that consuming items that are already in the history will not change her preferences.

A2. (H-Stationarity). *For any $z, w, v \in Z$, ρ feasible for v, and $n > 0$, if A is a finite decision problem with supp$(A) \subseteq h \cup \bigcap_{\sigma\in n(z)\cup n(w)} \sigma$, then:*

i. $z \succcurlyeq_h w \iff z_{-n}Az_n \succcurlyeq_h w_{-n}Aw_n$; and

ii. $v_{-\rho}A\{z \cup w\} \sim_h v_{-\rho}\{Az \cup Aw\}$.

H-stationarity has two components. First, (i) dictates that the decision maker does not change her preference when consuming prizes she has already consumed. The motivation for this is straightforward, as the only mechanism of learning is consumption of new prizes. The formal restriction states that if consuming the decision problem A is unavoidable (i.e., it will be reached on *every* path of consumption after n periods) and if all of the consumption prizes in A will already have been consumed (i.e., *supp*$(A) \subseteq h \cup \bigcap_{\sigma\in n(z)\cup n(w)} \sigma$), then the inclusion of A will not influence the DM's preferences between z and w. A special case of (i) states that if *supp*$(A) \subseteq h$ then $z \succcurlyeq_h w \iff Az \succcurlyeq_h Aw$, which is reminiscent of the canonical stationarity axiom of Koopmans [1960] (applied to previously consumed prizes). We require the more general formulation to enforce that preferences remain stationary (with respect to previously consumed prizes) even as they change (with respect to novel prizes).

Second, (ii) dictates that the decision maker is indifferent to making a choice before or after the consumption of previously consumed prizes. In general, the decision maker would prefer to delay making decisions, as it would allow her to condition on more information. However, if in the interim the decision maker learns no new information (because she consumes items from already in the history), then she is just as well off making the decision now.

Without learning the union of two IHCPs would never be preferred to both; the decision maker knows already which stream she would choose and would (weakly) prefer the menu that contained this stream. But, if the decision maker anticipates learning then flexibility is beneficial as it allows her to capitalize of the information she learns. Flexibility allows the DM to make different decisions conditional on the realization of the learning process (which piece of information she learns).

A3. (Preference for Flexibility). *For all $v, z, w \in Z$ and for all feasible decision sequences given v:*

$$v_{-\sigma}(z \cup w) \succcurlyeq_h v_{-\sigma} z$$

Note that this axiom looks slightly different than its conical form. The difference being that we require that flexibility is not disadvantageous, not only at time zero, but also on any branch of any IHCP. Because we have relaxed stationarity it is not enough to impose that $z \cup w \succcurlyeq_h z$ since this could be true for initial choices but be violated for continuation problems on a particular branch of some IHCP.

Next we need some notion that flexibility is consistent. That is to say, if the decision maker's preferences at history h are such that she anticipates no benefit to flexibility at history h', it must be also that her preference admits no benefit to the same flexibility at h' from the point of view of ex-ante preference.

A4. (Consistent Flexibility). *For every $h' \supseteq h$, and set of sequences contained in the history, $A \subseteq \times_{i=1}^{N} h$, and functions, $f, g : A \to Z$ then*

$$\bigcup_{\sigma \in A} \sigma\pi f(\sigma) \sim_h \bigcup_{\sigma \in A} \sigma\pi\{f(\sigma) \cup g(\sigma)\}$$
$$\iff$$
$$\bigcup_{\sigma \in A} \sigma f(\sigma) \sim_{h'} \bigcup_{\sigma \in A} \sigma\{f(\sigma) \cup g(\sigma)\}$$

for all π such that $h \cup \pi = h'$.

The DM must choose out of A today, which will determine the IHCP she will receive after consuming π which will create the history h'. If she sees no benefit to the additional flexibility provided by the assignment g, then it is because under all possible realization of uncertainty brought about by her consumption of h', the IHCP assigned by f is preferred to that by g. But if this is the case, then it must there is no benefit to flexibility according to the preference $\succcurlyeq_{h'}$. Further, since A is contained in the history, the DM consumption of h' will not affect her preference out of A (ignoring the continuation IHCP).

3.2 Strategic Planning

The above four axioms above provide the framework for a dynamic choice from menus. The preference for flexibility, after a given history, is being driven by the decision maker's subjective belief that she will better know her preference by the time a choice problem is chosen from the menu. That is, the decision maker entertains a subjective state space (a list of possible preference orderings over outcomes) for each history. Unlike the canonical models of choice from menus, there is not an immediate revelation of the state –after the consumption of a new prize the decision maker can distinguish between more states, but there may still be residual uncertainty. As such, we wish to impose the axioms above on the state dependent preferences as well. This ensures that after some learning has taken place, the DM anticipates continuing to learn in the same manner, albeit with more information about her true preference.

However, because the states are not observable, we need a way to construct choices that reveal the state dependent preferences. As discussed above, when the decision maker anticipates acquiring additional information regarding the state space, she will value future flexibility –allowing her future self to capitalize on the new information. However, if the state is fully revealed by previous choices, then flexibility no longer has value.

To understand this point recall the example from the introduction. The DM learns about her preference regarding phones by consuming laptops. As such, her preference after mp places a strict benefit to flexibility over phones:

$$\{a, i\} \succ_{mp} a \sim_{mp} i.$$

However, the ability to choose out of the menu $\{a, i\}$ can be rendered redundant when the phones are correctly paired with the laptops:

$$\{mi, pa\} \sim_{mp} \{m\{a, i\}, p\{a, i\}\}$$

This is because a choice out of $\{m, p\}$ fully characterizes the state; it utilized *all* of the information the decision maker has so far learned. So, by providing the appropriate state dependent optimal choice out of $\{a, i\}$, the DM gains no benefit from the ability to make two distinct choices.

Of course we are working backwards in the example: we already know the state space and are constructing the choices. However, if there always exists a collection of choices (such as $\{m, p\}$) that fully utilize the anticipated information structure, and a mapping from flexible menus to these choices (such as $a \mapsto p$ and $i \mapsto m$) then the converse is possible. We can start with the mappings and back out the state dependent utilities. Our next axiom ensures these mappings (called strategic plans) exist.

We introduce the following bit of notation:

Definition. (i) *For every h, such that $|h| \geqslant 2$, fix an arbitrary ordering, o, over all 2 (unique) element subsets of h. Let $\Sigma^o(h)$ denote the set of all sequences of length $\binom{|h|}{2}$ such that each element in $\Sigma^o(h)$ corresponds to the choice function for a strict preference ordering over h. Moreover, if $|h| = 0, 1$ then let $\Sigma^o(h) = \varnothing$.* (ii) *Let $\Sigma(h) = \bigcup \Sigma^o(h)/\sim_o$ to be the set of equivalence classes generated by the preference on h (that is, ignoring the ordering over subsets).*

To make clear the definition of $\Sigma(h)$, consider the following example. Let $h = \{a, b, c\}$, and take the

	$\{a,b\}$	$\{a,c\}$	$\{b,c\}$	σ
$a \succ b \succ c$	a	a	b	aab
$a \succ c \succ b$	a	a	c	aac
$b \succ a \succ c$	b	a	b	bab
$b \succ c \succ a$	b	c	b	bcb
$c \succ a \succ b$	a	c	c	acc
$c \succ b \succ a$	b	c	c	bcc

Figure 5: Example of $\Sigma(h)$

following ordering over all subsets with two unique elements: $o = \{a,b\},\{a,c\},\{b,c\}$. Then, $\Sigma^o(h) = \{aab, aac, bab, bcb, acc, bcc\}$. Each of the 6 sequences in $\Sigma^o(h)$ corresponds to one of the six possible strict preferences over $\{a,b,c\}$ as shown in Figure 5. If we change the ordering to $o' = \{a,c\},\{a,b\},\{b,c\}$, we permute each element of $\Sigma^o(h)$ accordingly, switching the first two elements of each sequence. But the mapping from preferences to sequences remains intact; each permuted sequence corresponds to the same strict preference.

Hence, the elements of $\Sigma(h)$ represent the maximizing elements for any possible strict preference over h (under some common ordering). For the remainder of this section, we deal with $\Sigma(h)$ corresponding to strict preference. In Appendix B, we show that this assumption can be weakened. A similar, although notationally cumbersome, exercise can be done starting with weak preferences over h.

The astute reader may have noticed that $\Sigma(h)$ was the notation in previous sections for the partition of the state space. It will turn out that $\Sigma(h)$ will describe the DMs information structure. It partitions the subjective state space into equivalence classes of states that respect the same preference over h. With this in mind it becomes clear that subsequent partitions are refinements of previous ones.

If the decision maker, after consuming h, learns her preferences over h, then there is an element of $\Sigma(h)$ that maximizes that preference. We assume (via axiom 5) that there is an injective mapping between the subjective states that a decision maker considers contingent on history h, and the elements of $\Sigma(h)$. In other words, we assume that everything that can be learned from the history h is identifiable from the decision makers preference over $\Sigma(h)$ after history h. This is a reasonable assumption: if the decision maker consumes a and b, she cannot learn about her preference between c and d unless her preference over c and d is somehow related to her preference over a and b. The mechanism we are interested in is experiential consumption.

Therefore, to reveal the decision makers preference contingent in being in a state we introduce the notion of strategic plans. A strategic plan is an allocation of choice problems to each element of $\Sigma(h)$ such that flexibility is no longer beneficial (beyond choosing from $\Sigma(h)$, that is).

Definition. *Given a history, h, and $W \subseteq Z$, a* ***strategic plan*** *is a function $p^W : \Sigma(h) \to W$ such that*

$$\bigcup_{\sigma \in \Sigma(h)} \sigma(p^W(\sigma)) \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma W$$

for all $\Sigma(h)$.

A strategic plan for W assigns to each sequence in $\Sigma(h)$ a subset of the choice problem (i.e., a element in the collection W) in such a way that there in no additional benefit from flexibility. Since each element of $\Sigma(h)$ corresponds to a strict preference, then $z = p^W(\sigma)$ implies that z is the maximizing element of W in the state that corresponds to σ.

A5. (Strategic Planning). *For all h and all $W = \{z_i\}_{i=1}^N \subset Z$ there exists a strategic plan over W with respect to $\succcurlyeq_h$.*

Strategic planing states that, for any finite collection of choice problems, the benefit to flexibility can be completely muted by appropriately allocating the choice problems to sequences in $\Sigma(h)$. While the intuition behind Strategic Planning is clear from the examples above, it is somewhat cumbersome to define and so raises the question as to the precise behavior it captures. The economic content of the axiom is precisely the following: the existence of a strategic plan assumes that each state is completely identifiable from its preference over h, and that the learning process is deterministic. That is, the axiom rules out (i) any learning (objective or subjective) which is not the direct result of consumption, and (ii) learning models that allow for residual uncertainty over outcomes that have been consumed.

Scenarios such as bandit problems, where the DM receives stochastic signals rather than deterministic ones, do not satisfy Strategic Planning. With such a signal structure it would not be possible to reduce all flexibility as the DM continues to learn from repeated consumption of the same alternative. It also rules out scenarios where the DM receives signals from sources other than her own consumption (i.e., non-responsive learning) as in Krishna and Sadowski [2014]. With non-responsive learning there is no corresponding sequence of choices that reveals the information the DM learned and so the construction in simply not possible.

It is also worth noting that although the axiom is existential, it is still falsifiable. Since strategic planning concerns only finite collections of IHCPs, there is a finite number of restrictions that need to be tested to verify the existence for a given collection.

Because states are mutually exclusive, we impose that the decision maker cannot hedge across them. In terms of the observables this implies that the maximizing elements of a strategic plan depend only on the state and not on the menus offered in other states of the world. The following axiom, closely resembling the Weak Axiom of Revealed Preference from standard choice theory, creates such a restriction.

A6. (WARP). *For all $W, V \subseteq Z$, if there exist strategic plans p^W and p^V, such that $p^W(\rho), p^V(\rho) \in V \cap W$ for some $\rho \in \Sigma(h)$ then*

$$\rho p^V(\rho) \cup \bigcup_{\sigma \in \Sigma(h) \setminus \rho} \sigma(p^W(\sigma)) \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma W$$

WARP states that if z is the maximal element in W conditional on choosing σ, and z is available in V but some other alternative w was chosen, then w must be as good as z. This is required to hold only in strategic plans, since in a generic IHCP hedging concerns may cause the decision maker to prefer choose elements based on multi-state contingencies.

3.3 An Example Showing the Necessity of Strategic Planning

The following example shows the necessity of the strategic planning construction. Let the decision problem be as in Section 1.1, but assume the (unobservable) state space is now:

$$S = \begin{cases} m \sim_1 i \succ_1 p \sim_1 a \\ m \sim_2 a \succ_2 p \sim_2 i \\ p \sim_3 a \succ_3 m \sim_3 i \end{cases}$$

We can observe the following preferences:

$$m \sim p \sim \{m, p\} \tag{iv}$$

$$mp\{m, p\} \succ mpm \sim mpp \tag{v}$$

$$\underbrace{ai\{a, i\}}_{v} \succ \underbrace{aia}_{w} \succ \underbrace{aii}_{z} \tag{vi}$$

$$mp\{mv, pw\} \sim mp\big\{m\{v, w, z\}, p\{v, w, z\}\big\} \succ mp\big\{m\{w, z\}, pw\big\} \tag{vii}$$

Preferences (iv) and (v) are the same as (i) and (ii) in the example in Section 1.1 and play the same role. The difference is that here, and unlike in the simpler example, learning the preference over m and p does not fully reveal the state of the world. Even after learning $m \succ p$, there is residual uncertainty. It is not enough to simply identify the conditional preference regarding a and i (conditional on $m \succ p$ or $p \succ m$). We also need to identify how the DM expects to continue learning (about a and i) after future consumption in each state. That is, we need to identify the DM's future preference for flexibility over all decision problems containing a and i.

Preference (ii) is sufficient to identify (the existence of) a group of states I in which $m \succ_i p$ for $i \in I$ and a group of states, J where $p \succ_j m$ for $j \in J$. Thus, identifying the DMs preference of a and i after learning that she prefers m to p (or p to m) entails identifying her aggregated preference in I (or J). Hence, we need to condition on being I or J. To do this, we construct a strategic plan –a tree such each branch is optimal in either I or J but not both. Clearly, if the DM knows the true state is in I then she prefers m to p. So consider, $mp\{m, p\}$; by the time the DM must choose out of $\{m, p\}$ she will know her preference and be able to choose optimally.

Thus, the optimal assignment of continuation trees to the branches of $mp\{m, p\}$ will coincide with the DMs preference in I and J. Hence, (vii) tells us that, conditional on $m \succ p$, $v \sim_i \{v, w, z\} \succ_i \{w, z\}$, echoing what we already inferred from (vi). Learning that she preferred the Mac to the PC did not teach her which phone she prefers; she still benefits from flexibility after consuming the phones. This implies that there are at least two states in which $m \succ p$.

On the other hand, if we condition on $p \succ m$, then $w \sim_j \{v, w, z\}$, or $aia \sim_j ai\{a, i\}$. Learning that she preferred the PC was indicative of her preference for the android phone. Any subsequent flexibility is redundant as no residual uncertainty remains. Of course, this identifies the single state where $p \succ m$.

	$\{a,b\} = a$	$\{a,c\}$	$\{b,c\}$	σ
$a \succ b \succ c$	a	a	b	ab
$a \succ c \succ b$	a	a	c	ac
~~$b \succ a \succ c$~~	~~b~~	~~a~~	~~b~~	~~ab~~
~~$b \succ c \succ a$~~	~~b~~	~~c~~	~~b~~	~~cb~~
$c \succ a \succ b$	a	c	c	cc
~~$c \succ b \succ a$~~	~~b~~	~~c~~	~~c~~	~~cc~~

Figure 6: Example of $\Sigma(h'|\sigma)$

3.4 Interim Preferences

Using the language of strategic planning we can define the interim (i.e., state-dependent) preferences. That is, the anticipated preference of the decision maker, conditional on having *learned* that her preference over h is the preference corresponding to the sequence $\sigma \in \Sigma(h)$ (we will refer to σ as the state of the world at h).

Definition. *For each h and $\sigma \in \Sigma(h)$, let $\succcurlyeq_{h|\sigma}$ denote the projection of preferences after history h and conditioned on the sequence σ defined by $z \succcurlyeq_{h|\sigma} w$ if and only if there exists a strategic plan $p : \Sigma(h) \to \{z, w\}$ such that $p(\sigma) = z$.*

The binary relation $\succcurlyeq_{h|\sigma}$ represents the preferences of the decision maker, conditional on being in the state of the world where σ is the maximizing element of $\Sigma(h)$. Although we do not have the language to directly observe the DM's hypothetical preferences, we can infer them by constructing $\succcurlyeq_{h|\sigma}$, which is fully characterized by $\succcurlyeq$.

As the decision maker continues to consume, her interim preference will change. As such, we can consider preferences of the form $\succcurlyeq_{h'|\sigma}$ where $\sigma \in \Sigma(h)$ and $h' \supset h$. These preferences are defined in the same way as $\succcurlyeq_h$ was defined from $\succcurlyeq$.

Definition. *For each h, $\sigma \in \Sigma(h)$ and $h' \supset h$ let $\succcurlyeq_{h'|\sigma}$ denote the projection of preferences after history h' conditional on σ defined by*

$$z \succcurlyeq_{h'|\sigma} w \iff h'z \succcurlyeq_{h|\sigma} h'w$$

The role of WARP becomes immediately clear: $\succcurlyeq_{h|\sigma}$ is a well defined preference relation.

Since we will want also to define strategic plans from the point of view of interim preferences (to examine the DM's anticipated future learning conditional on past learning), it will be helpful to define $\Sigma(h'|\sigma)$ for $h' \supset h$ and $\sigma \in \Sigma(h)$. $\Sigma(h'|\sigma)$ is the subset of $\Sigma(h')$ that is consistent with σ, restricted to the two element subsets not contained by h. That is, defined by weak orders that respect the weak ordering that generated σ and leaving out the redundant pairs. For example, if $h = \{a, b\}$ and $h' = \{a, b, c\}$ and $\sigma = a$ then $\Sigma(h'|\sigma) = \{ab, ac, cc\}$, as shown in Figure 6.

Strategic plans with respect to interim preferences will be defined on this domain. Since Lemma 3 guarantees that in any conditional preference the generating weak order is respected, we can define strategic plans from the point of view of $\succcurlyeq_{h'|\sigma}$ over $\Sigma(h'|\sigma)$ rather than over $\Sigma(h')$. This has the notational benefit being able to write future strategic plans in terms of previous ones.

Our first result, which is incremental in constructing the LTC representation provides the recessive structure of the preferences. The conditional preference, $\succcurlyeq_{h'|\sigma}$, inherits all of the structure imposed on $\succcurlyeq_h$; it satisfies *all* of the previous axioms.

Theorem 1. *Assume $\succcurlyeq$ is a Continuous Weak Order that satisfies H-stationarity, Flexibility, Consistent Flexibility, Strategic Planning, and WARP. Then for each h, $\sigma \in \Sigma(h)$, and $h \subseteq h'$, $\succcurlyeq_{h'|\sigma}$ satisfies Continuous Weak Order, H-stationarity, Flexibility, Consistent Flexibility, Strategic Planning, and WARP.*

This result allows us to obtain the branching, recursive representation. Even as the preferences change, their structure –and so, the functional form of the representation– remains intact. Additionally, Theorem 1 indicates that there is nothing special about the ex-ante preference, $\succcurlyeq$, other than being associated with a particular state space. If preferences were elicited at a later date, when some items had been consumed, the resulting axiomatic structure and representation would be identical save for a smaller (conditional) state space.

Our last axiom imposes consistency on the way the DM's preferences over outcomes (i.e., over X) can change as she learns. In particular we impose that the value of a singleton is constant from an ex-ante point of view. This implies that the value a decision maker places on an outcome a is the aggregated value she places on a at each state which she considers possible. We want this to hold, not only at period zero, but also for each interim preference, when looking forward. We also impose that there is a best and a worst outcome that are ex-ante identified. This assumption is simply a convenient way to calibrate states to one another.

A7. (Singleton Recursivity). *There exist universal best and worst element, $\bar{a}$ and $\underline{a}$, such that $A = \times^{\infty}\{\bar{a}, \underline{a}\}$ is order-dense in Z. Moreover,*

$$\boldsymbol{\rho} \succcurlyeq_{h|\sigma} \boldsymbol{\tau} \iff \boldsymbol{\rho} \succcurlyeq_{h'|\sigma} \boldsymbol{\tau}$$
$$\boldsymbol{\nu} \succcurlyeq \boldsymbol{\pi} \iff \boldsymbol{\nu} \succcurlyeq_{h'|\sigma} \boldsymbol{\pi}$$

for all $\sigma \in \Sigma(h'')$ such that $h'' \subseteq h \cap h'$, $\boldsymbol{\rho}, \boldsymbol{\tau} \in \times^{\infty} X$, and $\boldsymbol{\pi}, \boldsymbol{\nu}, \in A$.

Even though a decision maker expects to learn, a pre-determined stream of consumption offers no way to capitalize on the new information. Hence, the DM's value of a stream of singleton prizes is the aggregation of her value in each state she still considers possible. Similarly, since the DM's value over streams concerning $\bar{a}$ and $\underline{a}$ are ex-ante known, her preference regarding such streams does not change, even when conditioned on particular information.

4 The Characterization Result

Theorem 2 (Learning Through Consumption Representation)**.** *The preference $\succcurlyeq$ satisfies Weak Order, H-stationarity, Flexibility, Consistent Flexibility, Strategic Planning, WARP and Singleton Recursivity, if and only if there exists a state space S and a partition thereof, $\Sigma(h)$, corresponding to each history h, a family of (information-dependent) utility functions $u_\sigma : X \to \mathbb{R}$ for every $\sigma \in \Sigma(h)$, a strictly increasing aggregator*

$F : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, *and a family of strictly increasing aggregators* $\phi_{h'|\sigma} : \mathbb{R}^{|\Sigma(h'|\sigma)|} \to \mathbb{R}$ *for every* $h' \supseteq h$, *such that*

$$V_{h'|\sigma}(z) = \phi_{h'|\sigma}\big((V_{h'|\tau}(z))_{\tau\in\Sigma(h'|\sigma)}\big) = \phi_{h'|\sigma}\big(\max_{i\in M(z)} F(u_\tau(a_i), V_{h'\cup a_i|\tau}(z_i))_{\tau\in\Sigma(h'|\sigma)}\big) \tag{8}$$

represents $\succcurlyeq_{h'|\sigma}$. *Moreover, (i)* $\Sigma(h')$ *is a refinement of* $\Sigma(h)$, *(ii)* $u_\sigma = \phi_{h'|\sigma}(u_\tau)_{\tau\in\Sigma(h'|\sigma)}$ *and (iii)* $u_\sigma(a) \geqslant u_\sigma(b) \iff u_\tau(a) \geqslant u_\tau(b)$ *for any* $\tau \in \Sigma(h'|\sigma)$ *and* $a, b \in h$.

5 The Dominance Relation Behind Strategic Planning

In this section we provide equivalent statements for Strategic Planning, which postulates the existence of a strategic plan given any history. Given a history h and a finite IHCP W, we are going to conisder all possibe assignments from optional consumptions (that is, $\Sigma(h)$) into W, *including* non-optimal assignments. Conceptually, we will define a notion of dominance over IHCPs that will allow us to find the maximal assignment with respect to this notion. This assignment will be a strategic plan.

5.1 The Dominance Relation

Definition. *Given any collection of IHCPs,* W, *and a history,* h, *an* ***assignment*** *is a function from* $\Sigma(h)$ *to* W. *Let* As(W,h) *be the set of all such assignments.*

Definition. *Set* $z, w \in Z$. *Then we say that* z ***σ-dominates*** w, *denoted* $z \succcurlyeq_\sigma w$, *if and only if,*

$$\sigma z \cup \bigcup_{\rho\in\Sigma(h)\setminus\sigma} \rho\big(z \cup w\big) \succcurlyeq_h \sigma w \cup \bigcup_{\rho\in\Sigma(h)\setminus\sigma} \rho\big(z \cup w\big).$$

We say that z σ-dominates w when the assignment is preferred in the absence of any cross state hedging. Hedging concerns are removed by ensuring that z and w can be chosen in any other "state."

Definition. *Set* $p \in$ As(W,h). *Then we say that* p *is in the* ***envelope of W*** *if* $p(\sigma)$ σ*-dominates* w *for all* $w \in W$ *and* $\sigma \in \Sigma(h)$.

An assignment is in the envelope of W if each assignment is σ-dominance maximal element of the collection; it is the maximal assignment according to dominance. Without further restrictions, the envelope of a collection need not exist. The following two restrictions will imply that (i) the envelope of any *finite* W exists, and (ii) an assignment in the envelope is a strategic plan for W. Moreover, since σ-dominance does not depend on the assignment, these restrictions will also guarantee that WARP is satisfied.

For these claims to hold, we need to first ensure that σ-dominance is a weak order. Since the completeness of $\succcurlyeq_\sigma$ is inherited by the completeness of $\succcurlyeq$, we need only to impose transitivity.

A8. (σ-Transativity). *For all* $\sigma \in \Sigma(h)$, $\succcurlyeq_\sigma$, *is transitive.*

While σ-Transativity is written in terms of the auxiliary relation $\succcurlyeq_\sigma$, it is straight forward to write as a restriction on our primitive $\succcurlyeq$. Axiom σ-Transativity, has the following immediate consequence:

Lemma 1. *If* $\succcurlyeq$ *is a weak order that satisfies* σ*-Transativity, then* $\succcurlyeq_\sigma$ *is a weak order over* Z.

The next axiom is a modularity condition, ensurring the maximal element according to the dominance relation (i.e., maximal for all σ) is a strategic plan.

A9. (MODULARITY). *For all finite W and $p \in As(W,h)$, p is in the envelope of W if and only if*

$$\bigcup_{\rho \in \Sigma(h)} \rho p(\rho) \sim_h \bigcup_{\rho \in \Sigma(h)} \rho W.$$

Modularity characterizes preferences for flexibility entirely in terms of dominance: flexibly is of no benefit unless it can improves the dominance of its assignments. This is essentially the content behind Strategic Planning (i.e., that flexibility can be muted) and WARP (that plans are evaluated without hedging concerns). The next result states that in Theorem 2, our main representation result, Strategic Planning and WARP could be replaced with σ-Transativity and Modularity.

Theorem 3. *If the preference $\succcurlyeq$ satisfies Weak Order, H-stationarity, Flexibility, and Consistent Flexibility, then the following are equivalent:*

1. *$\succcurlyeq$ satisfies Strategic Planning and WARP;*
2. *$\succcurlyeq$ satisfies σ-Transativity and Modularity.*

5.2 The Dominance Relation and Kreps' Independence

To conclude this discussion, we bring a second alternative to the axiomatization of learning through consumption. The axiom below is a reformulation of Kreps' independnece axiom in the context of σ-domination.

A10. (DOMINANCE INDEPENDENCE). *For all $f, g, k \in$ As(Z,h) such that $f(\sigma) \succcurlyeq_\sigma g(\sigma)$ for all $\sigma \in \Sigma(h)$,*

$$\bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup k(\sigma)\} \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup g(\sigma) \cup k(\sigma)\}.$$

The following propostion states that in the main characterization result, Strategic Planning can be replaced with σ-Transativity and Dominance Independence.

Proposition 1. *If the preference $\succcurlyeq$ satisfies Weak Order, H-stationarity, Flexibility, Consistent Flexibility, and σ-Transativity, then the following are equivalent:*

1. *$\succcurlyeq$ satisfies Modularity.*
2. *$\succcurlyeq$ satisfies WARP and Dominance Independence.*

A Proofs

A.1 Preliminary Results

Lemma 2. *The decision maker has a weak preference for delaying choices. That is, for all $z, w, v \in Z$, σ feasible with respect to v and finite decision problems A,*

$$v_{-\sigma}A\{z \cup w\} \succcurlyeq_h v_{-\sigma}\{Az \cup Aw\}.$$

Proof. Proof of Lemma 2. This is an immediate consequence of the decision makers preference for flexibility. That is, we know that the left hand side is identically equal to $v_{-\sigma}\big\{A\{z \cup w\} \cup A\{z \cup w\}\big\}$, which by flexibility and transitivity

$$v_{-\sigma}\big\{A\{z \cup w\} \cup A\{z \cup w\}\big\} \succcurlyeq_h v_{-\sigma}\big\{A\{z\} \cup A\{z \cup w\}\big\} \succcurlyeq_h v_{-\sigma}\big\{A\{z\} \cup A\{w\}\big\},$$

which is identically equal to the right hand side. ◻

Recall that for all $a \in X$, $\mathbf{a}$ denotes the sequence $(a, a, ...)$. For some $\sigma \in \Sigma(h)$ define $\succcurlyeq_\sigma$ as the unique preference order over h that generates σ as the sequence of choices out of the two element subsets of h that generated $\Sigma(h)$.

Lemma 3. *Let $a, b \in h$ and let σ be generated by a weak order such that $a \succcurlyeq_\sigma b$ then $\boldsymbol{a} \succcurlyeq_{h|\sigma} \boldsymbol{b}$.*

Proof. Proof of Lemma 3. We will show this for $h = \{a, b\}$. The same argument extends to a general history, but the notation becomes more complicated. Assume that $\boldsymbol{b} \succcurlyeq_{h|a} \boldsymbol{a}$. This implies that $\boldsymbol{b} \succcurlyeq_{h|a} a\boldsymbol{b} \succcurlyeq_{h|a} \boldsymbol{a}$ and $\boldsymbol{b} \succcurlyeq_{h|a} b\boldsymbol{a}$ by singleton recursivity. Case 1: $\boldsymbol{b} \succcurlyeq_{h|b} \boldsymbol{a}$. By Strategic Planning and WARP,

$$\{a\boldsymbol{b}, b\boldsymbol{b}\} \sim_h \bigcup_{\sigma=a,b} \sigma\{\boldsymbol{a}, a\boldsymbol{b}, b\boldsymbol{a}, \boldsymbol{b}\}.$$

By H-stationarity

$$\boldsymbol{b} \sim_h \{\boldsymbol{a}, a\boldsymbol{b}, b\boldsymbol{a}, \boldsymbol{b}\},$$

and by flexibility we get

$$\{\boldsymbol{a}, \boldsymbol{b}\} \sim_h \{\boldsymbol{a}, a\boldsymbol{b}, b\boldsymbol{a}, \boldsymbol{b}\},$$

which is a strategic plan that respects the correct ordering.

Case 2: $\boldsymbol{a} \succcurlyeq_{h|b} \boldsymbol{b}$. By appealing again to singleton recursivity we have

$$\{a\boldsymbol{b}, b\boldsymbol{a}\} \sim_h \bigcup_{\sigma=a,b} \sigma\{\boldsymbol{a}, a\boldsymbol{b}, b\boldsymbol{a}, \boldsymbol{b}\}.$$

So flexibility dictates

$$\{a\{\boldsymbol{a}, \boldsymbol{b}\}, b\{\boldsymbol{a}, \boldsymbol{b}\}\} \sim_h \bigcup_{\sigma=a,b} \sigma\{\boldsymbol{a}, a\boldsymbol{b}, b\boldsymbol{a}, \boldsymbol{b}\}.$$

Finally, by H-stationarity we have

$$\{\boldsymbol{a}, \boldsymbol{b}\} \sim_h \{\boldsymbol{a}, a\boldsymbol{b}, b\boldsymbol{a}, \boldsymbol{b}\},$$

which is the same strategic plan.

◻

Lemma 4. *For all $\sigma \in \Sigma(h)$, and $W \subset Z$, there exists a strategic plan, $p^W : \Sigma(h) \to W$, such that $p^W(\sigma) = z$ if any only if*

$$\sigma z \ \cup \bigcup_{\rho \in \Sigma(h) \setminus \sigma} \rho W \sim_h \bigcup_{\rho \in \Sigma(h)} \rho W.$$

Proof. Proof of Lemma 4 First, assume such a strategic plan exists. So we have

$$\begin{aligned}\bigcup_{\rho \in \Sigma(h)} \rho W &\succcurlyeq_h \sigma z \ \cup \bigcup_{\rho \in \Sigma(h) \setminus \sigma} \rho W \\ &\succcurlyeq_h \sigma z \ \cup \bigcup_{\rho \in \Sigma(h) \setminus \sigma} \rho(p^W(\rho)) \\ &\sim_h \bigcup_{\rho \in \Sigma(h)} \rho W,\end{aligned}$$

where both weak preferences are consequences of Flexibility and the indifference relation by Strategic Planning and our assumption.

Now assume that

$$\sigma z \ \cup \bigcup_{\rho \in \Sigma(h) \setminus \sigma} \rho W \sim_h \bigcup_{\rho \in \Sigma(h)} \rho W. \tag{1}$$

By Strategic Planning we know there exists some $p^W : \Sigma(h) \to W$ such that

$$\bigcup_{\rho \in \Sigma(h)} \rho p^W(\rho) \sim_h \bigcup_{\rho \in \Sigma(h)} W. \tag{2}$$

Consider the collection $V = \{W, w_i\}_{w_i \in W}$ (In an effort to be clear, $V = \{\{\bigcup_{w \in W} w\}, w_1, w_2, \dots w_n\}$. It is immediate that $\bigcup_{w \in W} w = \bigcup_{v \in V} v$ and so, utilizing this identity, we can rewrite (1) and (2):

$$\sigma z \ \cup \bigcup_{\rho \in \Sigma(h) \setminus \sigma} \rho W \sim_h \bigcup_{\rho \in \Sigma(h)} \rho V, \tag{3}$$

$$\bigcup_{\rho \in \Sigma(h)} \rho p^W(\rho) \sim_h \bigcup_{\rho \in \Sigma(h)} \rho V. \tag{4}$$

So, by (3) and (4) we satisfy the conditions for WARP and so:

$$\sigma z \ \cup \bigcup_{\rho \in \Sigma(h) \setminus \sigma} \rho p^W(\rho) \sim_h \bigcup_{\rho \in \Sigma(h)} \rho V,$$

or, since $\bigcup_{w \in W} w = \bigcup_{v \in V} v$,

$$\sigma z \ \cup \bigcup_{\rho \in \Sigma(h) \setminus \sigma} \rho p^W(\rho) \sim_h \bigcup_{\rho \in \Sigma(h)} \rho W.$$

□

Lemma 5. *Let $h \subset h'$. The collection of functions, $\big(p_\sigma : \Sigma(h'|\sigma) \to W\big)_{\sigma \in \Sigma(h)}$ constitute a collection of strategic plans for the corresponding preference relations $\big(\succcurlyeq_{h'|\sigma} \big)_{\sigma \in \Sigma(h)}$ if and only if the map*

$$p : \Sigma(h') \longrightarrow W$$
$$p : \sigma\tau \mapsto p_\sigma(\tau)$$

is a strategic plan for $\succcurlyeq_{h'}$. Where $\sigma \in \Sigma(h)$ and $\tau \in \Sigma(h'|\sigma)$.

Proof. Proof of Lemma 5.

Let $\big(p_\sigma : \Sigma(h'|\sigma) \to W\big)_{\sigma \in \Sigma(h)}$ constitute a collection of strategic plans for the corresponding preference relations $\big(\succcurlyeq_{h'|\sigma} \big)_{\sigma \in \Sigma(h)}$. Then we know that, for each $\sigma \in \Sigma(h)$:

$$\bigcup_{\tau \in \Sigma(h'|\sigma)} \tau p_\sigma(\tau) \sim_{h'|\sigma} \bigcup_{\tau \in \Sigma(h'|\sigma)} \tau W.$$

Using the definition of $\sim_{h'|\sigma}$ this is equivalent to

$$(h' \backslash h) \bigcup_{\tau \in \Sigma(h'|\sigma)} \tau p_\sigma(\tau) \sim_{h|\sigma} (h' \backslash h) \bigcup_{\tau \in \Sigma(h'|\sigma)} \tau W. \tag{1}$$

Hence, we know that there must exist two strategic plans (with respect the $\succcurlyeq_h$) such that each side of (1) gets assigned to σ in one of the two plans. Then, utilizing WARP to combine these strategic plans, we know

$$\bigcup_{\sigma \in \Sigma(h)} \sigma(h' \backslash h) \bigcup_{\tau \in \Sigma(h'|\sigma)} \tau p_\sigma(\tau) \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma(h' \backslash h) \bigcup_{\tau \in \Sigma(h'|\sigma)} \tau W. \tag{2}$$

Since by Flexibility, the right-hand side of (2) is indifferent to

$$\bigcup_{\sigma \in \Sigma(h)} \sigma(h' \backslash h) \bigcup_{\tau \in \Sigma(h'|\sigma)} \{\tau W, \tau p_\sigma(\tau)\},$$

Consistent Flexibility implies

$$\bigcup_{\sigma \in \Sigma(h)} \sigma \bigcup_{\tau \in \Sigma(h'|\sigma)} \tau p_\sigma(\tau) \sim_{h'} \bigcup_{\sigma \in \Sigma(h)} \sigma \bigcup_{\tau \in \Sigma(h'|\sigma)} \{\tau W, \tau p_\sigma(\tau)\}.$$

Again by Flexibility,

$$\bigcup_{\sigma \in \Sigma(h)} \sigma \bigcup_{\tau \in \Sigma(h'|\sigma)} \tau p_\sigma(\tau) \sim_{h'} \bigcup_{\sigma \in \Sigma(h)} \sigma \bigcup_{\tau \in \Sigma(h'|\sigma)} \tau W,$$

which is equivalent to

$$\bigcup_{\sigma \in \Sigma(h)} \bigcup_{\tau \in \Sigma(h'|\sigma)} \sigma\tau p_\sigma(\tau) \sim_{h'} \bigcup_{\sigma \in \Sigma(h)} \bigcup_{\tau \in \Sigma(h'|\sigma)} \sigma\tau W.$$

And so, the specified map, p, is indeed a strategic plan for $\succcurlyeq_{h'}$. Each step is a bi-directional implication so this proves the if and only if statement. □

A.2 Proof of Theorem 1

Fix some h, $h' \supset h$ and some $\sigma \in \Sigma(h)$.

(Weak Order). That $\succcurlyeq_{h'|\sigma}$ is complete is immediate from the definition. Now assume that $z \succcurlyeq_{h'|\sigma} w$ and $w \succcurlyeq_{h'|\sigma} v$. Let $\hat{V} = \{h'z, h'w\}$ and $\bar{V} = \{h'w, h'v\}$. By the definition of $\succcurlyeq_{h'|\sigma}$ there exist strategic plans, $p^{\hat{V}} : \Sigma(h) \to \hat{V}$, and $p^{\bar{V}} : \Sigma(h) \to \bar{V}$, such that $p^{\hat{V}}(\sigma) = h'z$ and $p^{\bar{V}}(\sigma) = h'w$.

Define $V = \{h'z, h'w, h'v\}$ and let $p^V : \Sigma(h) \to V$ be the plan ensured by Strategic Planning. We claim that there exits some strategic plan, q^V, thereover, such that $q^V(\sigma) = h'z$. There are three cases: Case 1: $p^V(\sigma) = h'z$. Then the claim holds. Case 2: $p^V(\sigma) = h'w$, then since $p^{\hat{V}}(\sigma) = h'z$ we satisfy the conditions of WARP:

$$\sigma h'z \ \cup \bigcup_{\rho\in\Sigma(h)\setminus\sigma} \rho p^V(\rho) \sim_h \bigcup_{\rho\in\Sigma(h)} \rho\{h'z \cup h'w \cup h'v\}, \tag{1}$$

which defines the desired plan. Case 3: $\hat{p}(\sigma) = h'v$, we can iterate the process used in case 2 to prove the claim.

Finally, Strategic Planning ensures that there exists a plan p^W over $W = \{h'z, h'v\}$. If $p^W(\sigma) = h'z$ then it cannot be that $h'v \succ_{h|\sigma} h'z$ and we are done. If $p^W(\sigma) = h'v$, then the fact that $p^V(\sigma) = h'z$ and WARP provide a new strategic plan $q^W(\sigma) = h'z$, so again it cannot be that $h'v \succ_{h|\sigma} h'z$.

(Continuity). Let $\{z_n\}_{n\in\mathbb{N}}$ be a convergent sequence in Z, with limit point z, such that $z_n \succcurlyeq_{h'|\sigma} w$ for all n and some $w \in Z$. If it is the case that $z \succcurlyeq_{h'|\sigma} z_n$ for some $n \in \mathbb{N}$, then by the transitivity of $\succcurlyeq_{h'|\sigma}$ we have that $z \succcurlyeq_{h'|\sigma} w$. So assume that for all n, $z_n \succcurlyeq_{h'|\sigma} z$. By Lemma 4 (and WARP) this implies that:

$$\sigma h'z_n \ \cup \bigcup_{\rho\in\Sigma(h)\setminus\sigma} \rho\{h'z_n \cup h'z \cup h'w\} \sim_h \bigcup_{\rho\in\Sigma(h)} \rho\{h'z_n \cup h'z \cup h'w\}.$$

Taking the limit, the continuity of $\succcurlyeq_h$ provides:

$$\sigma h'z \ \cup \bigcup_{\rho\in\Sigma(h)\setminus\sigma} \rho\{h'z \cup h'w\} \sim_h \bigcup_{\rho\in\Sigma(h)} \rho\{h'z \cup h'w\},$$

implying that $z \succcurlyeq_{h'|\sigma} w$. Hence the contour sets of of $\succcurlyeq_{h'|\sigma}$ are closed.

(H-Stationarity). Let $z \succcurlyeq_{h'|\sigma} w$, and A be a finite decision problem which is contained in $h' \cup \bigcap_{\sigma\in n(z)\cup n(w)} \sigma$ for some n. So, there exists a strategic plan, $p^V : \Sigma(h) \to V$, with $V = \{h'z, h'w\}$, such that

$$\hat{z} = \sigma h'z \cup \bigcup_{\rho\in\Sigma(h)\setminus\sigma} \rho\{h'z, h'w\} \sim_h \bigcup_{\rho\in\Sigma(h)} \rho\{h'z, h'w\} = \hat{w}.$$

Let $m = |\sigma| + |h'|$. Consider $H = h \cup \bigcap_{\sigma\in[n+m](\hat{z})\cup[n+m](\hat{w})} \sigma$. For any $a \in supp(A)$ we claim $a \in H$.

There are three cases: (1) $a \in h$, in which case we are done since $h \subset H$. (2) $a \in h'\setminus h$ in which case it is in the first m periods of $\hat{z}$ and $\hat{w}$ (by construction of the above plans you *have* to consume h'). (3)

$a \in \bigcap_{\sigma \in n(z) \cup n(w)} \sigma$ in which case it is simply pushed back m periods (by construction of the above plans you *have* to consume either the first n periods of z or the first n periods of w). Hence, by H-stationarity we can interject A after $n+m$ periods. Letting $W = \{h' z_{-n} A z_n, h' w_{-n} A w_n\}$ this gives

$$\sigma h' z_{-n} A z_n \cup \bigcup_{\rho \in \Sigma(h) \backslash \sigma} \rho W \sim_h \bigcup_{\rho \in \Sigma(h)} \rho W,$$

so, by Lemma 4,

$$z_{-n} A z_n \succcurlyeq_{h'|\sigma} w_{-n} A w_n.$$

The converse holds from the bi-directional implication of each step.

Call $V = \{h' v_{-\pi} A\{z \cup w\}, h' v_{-\sigma}\{Az \cup Aw\}$. There exists a strategic plan $p^V : \Sigma(h) \in V$. Assume that $p^V(\sigma) = h' v_{-\pi} A\{z \cup w\}$. So

$$\sigma h' v_{-\pi} A\{z \cup w\} \cup \bigcup_{\rho \in \Sigma(h) \backslash \sigma} \rho p^V(\rho) \sim_h \bigcup_{\rho \in \Sigma(h)} \rho V,$$

but by the (ii) of H-stationarity we have

$$\sigma h' v_{-\pi}\{Az \cup Aw\} \cup \bigcup_{\rho \in \Sigma(h) \backslash \sigma} \rho p^V(\rho) \sim_h \bigcup_{\rho \in \Sigma(h)} \rho V,$$

and so, $v_{-\pi} A\{z \cup w\} \succcurlyeq_{h'|\sigma} v_{-\sigma}\{Az \cup Aw\}$. Of course, if we assume $p^V(\sigma) = h' v_{-\pi}\{Az \cup Aw\}$ the same argument holds. We have $v_{-\pi} A\{z \cup w\} \sim_{h'|\sigma} v_{-\sigma}\{Az \cup Aw\}$ as desired.

(Preference for Flexibility). Assume, to the contrary, that there exists some $v, z, w \in Z$, and π feasible for v such that $v_{-\pi} z \succ_{h'|\sigma} v_{-\pi}(z \cup w)$. Let $V = \{h' v_{-\pi} z, h' v_{-\pi}(z \cup w)\}$. Lemma 4 therefore implies

$$\sigma h' v_{-\pi} z \cup \bigcup_{\rho \in \Sigma(h) \backslash \sigma} \rho V \succ_h \sigma h' v_{-\pi}(z \cup w) \cup \bigcup_{\rho \in \Sigma(h) \backslash \sigma} \rho V,$$

an immediate contradiction to the DMs preference for flexibility.

(Consistent Flexibility). Assume that, for some $A \subseteq \bigtimes_{i=1}^{N} h'$, finite π such that $h' \cup \pi = h''$ and $f, g : A \to Z$ we have

$$\bigcup_{\tau \in A} \tau \pi f(\tau) \sim_{h'|\sigma} \bigcup_{\tau \in A} \tau \pi \{f(\tau) \cup g(\tau)\}. \tag{2}$$

Define, $V = \left\{h' \bigcup_{\tau \in A} \tau \pi f(\tau), h' \bigcup_{\tau \in A} \tau \pi \{f(\tau) \cup g(\tau)\}\right\}$. Then, by the definition of $\succcurlyeq_{h'|\sigma}$, WARP, and lemma

4 we know that there exists some strategic plan $p^V : \Sigma(h) \to V$ such that

$$\sigma h' \bigcup_{\tau \in A} \tau\pi f(\tau) \cup \bigcup_{\rho \in \Sigma(h)\setminus\sigma} \rho p^V(\rho) \sim_h \bigcup_{\rho\in\Sigma(h)} \rho V, \tag{3}$$

$$\sigma h' \bigcup_{\tau \in A} \tau\pi \{f(\tau) \cup g(\tau)\} \cup \bigcup_{\rho \in \Sigma(h)\setminus\sigma} \rho p^V(\rho) \sim_h \bigcup_{\rho\in\Sigma(h)} \rho V. \tag{4}$$

Define the mapping $\hat{f} : \Sigma(h)A \to Z$ as

$$\rho\tau \mapsto \begin{cases} f(\tau) & \text{if } \rho = \sigma \\ f(\tau) & \text{if } \rho \neq \sigma \text{ and } p^V(\rho) = \bigcup_{\tau\in A} \tau\pi f(\tau) \\ \{f(\tau) \cup g(\tau)\} & \text{if } \rho \neq \sigma \text{ and } p^V(\rho) = \bigcup_{\tau\in A} \tau\pi \{f(\tau) \cup g(\tau)\}, \end{cases}$$

and the mapping $\hat{g} : \Sigma(h)A \to Z$ as

$$\rho\tau \mapsto \begin{cases} \{f(\tau) \cup g(\tau)\} & \text{if } \rho = \sigma \\ f(\tau) & \text{if } \rho \neq \sigma \text{ and } p^V(\rho) = \bigcup_{\tau\in A} \tau\pi f(\tau) \\ \{f(\tau) \cup g(\tau)\} & \text{if } \rho \neq \sigma \text{ and } p^V(\rho) = \bigcup_{\tau\in A} \tau\pi \{f(\tau) \cup g(\tau)\}. \end{cases}$$

Then, we can re-write the (3) and (4) as

$$\bigcup_{\rho\in\Sigma(h)} \rho h' \bigcup_{\tau\in A} \tau\pi \hat{f}(\rho\tau) \sim_h \bigcup_{\rho\in\Sigma(h)} \rho h' \bigcup_{\tau\in A} \tau\pi \{\hat{f}(\rho\tau) \cup \hat{g}(\rho\tau)\}, \tag{5}$$

$$\bigcup_{\rho\in\Sigma(h)} \rho h' \bigcup_{\tau\in A} \tau\pi \hat{g}(\rho\tau) \sim_h \bigcup_{\rho\in\Sigma(h)} \rho h' \bigcup_{\tau\in A} \tau\pi \{\hat{f}(\rho\tau) \cup \hat{g}(\rho\tau)\}, \tag{6}$$

so by Consistent Flexibility (expanding h to h' by eliminating h'[9])

$$\bigcup_{\rho\in\Sigma(h)} \rho \bigcup_{\tau\in A} \tau\pi \hat{f}(\rho\tau) \sim_{h'} \bigcup_{\rho\in\Sigma(h)} \rho \bigcup_{\tau\in A} \tau\pi \{\hat{f}(\rho\tau) \cup \hat{g}(\rho\tau)\},$$

$$\bigcup_{\rho\in\Sigma(h)} \rho \bigcup_{\tau\in A} \tau\pi \hat{g}(\rho\tau) \sim_{h'} \bigcup_{\rho\in\Sigma(h)} \rho \bigcup_{\tau\in A} \tau\pi \{\hat{f}(\rho\tau) \cup \hat{g}(\rho\tau)\},$$

again by Consistent Flexibility (expanding h' to h'' by eliminating π)

$$\bigcup_{\rho\tau\in\Sigma(h)A} \rho\tau \hat{f}(\rho\tau) \sim_{h''} \bigcup_{\rho\tau\in\Sigma(h)A} \rho\tau \{\hat{f}(\rho\tau) \cup \hat{g}(\rho\tau)\}, \tag{7}$$

$$\bigcup_{\rho\tau\in\Sigma(h)A} \rho\tau \hat{g}(\rho\tau) \sim_{h''} \bigcup_{\rho\tau\in\Sigma(h)A} \rho\tau \{\hat{f}(\rho\tau) \cup \hat{g}(\rho\tau)\}. \tag{8}$$

[9]This is permissible by defining $\dot{f} : \rho \mapsto \bigcup_{\tau\in A} \tau\pi \hat{f}(\rho\tau)$, $\dot{g}$ likewise, and $\dot{h} : \rho \mapsto \bigcup_{\tau\in A} \tau\pi\{\hat{f}(\rho\tau) \cup \hat{g}(\rho\tau)\}$. Then using WARP, we can show that the antecedent for CF is satisfied (via $\dot{f} \cup \dot{h}$ for (5) and $\dot{g} \cup \dot{h}$ for (6).

Let $\hat{V} = \{\bigcup_{\tau\in A}\tau f(\tau), \bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\}\}$ and Let $\hat{p}^V : \Sigma(h)\backslash\sigma \to \hat{V}$ as

$$\rho \mapsto \begin{cases} \bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\} & \text{if } p^V(\rho) = \bigcup_{\tau\in A}\tau\pi\{f(\tau)\cup g(\tau)\} \\ \bigcup_{\tau\in A}\tau f(\tau) & \text{if } p^V(\rho) = \bigcup_{\tau\in A}\tau\pi f(\tau). \end{cases}$$

Then, (7) and (8) are rewritten as

$$\sigma\bigcup_{\tau\in A}\tau f(\tau) \cup \bigcup_{\rho\in\Sigma(h)\backslash\sigma}\rho\hat{p}^V(\rho) \sim_{h''} \bigcup_{\rho\in\Sigma(h)}\rho\hat{V},$$
$$\sigma\bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\} \cup \bigcup_{\rho\in\Sigma(h)\backslash\sigma}\rho\hat{p}^V(\rho) \sim_{h''} \bigcup_{\rho\in\Sigma(h)}\rho\hat{V}.$$

Again by Consistent Flexibility,

$$\sigma h''\bigcup_{\tau\in A}\tau f(\tau) \cup \bigcup_{\rho\in\Sigma(h)\backslash\sigma}\rho h''\hat{p}^V(\rho) \sim_{h} \bigcup_{\rho\in\Sigma(h)}\rho h''\hat{V},$$
$$\sigma h''\bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\} \cup \bigcup_{\rho\in\Sigma(h)\backslash\sigma}\rho h''\hat{p}^V(\rho) \sim_{h} \bigcup_{\rho\in\Sigma(h)}\rho h''\hat{V}.$$

By appealing to Lemma 2 we have

$$\sigma h''\bigcup_{\tau\in A}\tau f(\tau) \cup \bigcup_{\rho\in\Sigma(h)\backslash\sigma}\rho h''\hat{p}^V(\rho) \succcurlyeq_{h} \bigcup_{\rho\in\Sigma(h)}\rho\{h''\bigcup_{\tau\in A}\tau f(\tau), h''\bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\}\}, \tag{9}$$
$$\sigma h''\bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\} \cup \bigcup_{\rho\in\Sigma(h)\backslash\sigma}\rho h''\hat{p}^V(\rho) \succcurlyeq_{h} \bigcup_{\rho\in\Sigma(h)}\rho\{h''\bigcup_{\tau\in A}\tau f(\tau), h''\bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\}\}. \tag{10}$$

The opposite weak preference preference to (9) and (10) are guaranteed by the DMs preference for flexibility. Hence (9) and (10) hold with indifference and thus define a pair of strategic plans indicating the conditional preference:

$$h''\bigcup_{\tau\in A}\tau f(\tau) \sim_{h|\sigma} h''\bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\},$$

or

$$\bigcup_{\tau\in A}\tau f(\tau) \sim_{h''|\sigma} \bigcup_{\tau\in A}\tau\{f(\tau)\cup g(\tau)\},$$

as desired. The converse is also true by the same argument.

(Strategic Planning). We need only show that there exists a strategic plan for any 2 element subset of Z: induction and the arguments used above can extend to any finite set. So take some $z, w \in Z$. We want to show that there exists some $\hat{p}^{\{z,w\}} : \Sigma(h'|\sigma) \to \{z, w\}$ such that

$$\bigcup_{\tau\in\Sigma(h'|\sigma)}\tau\hat{p}^{\{z,w\}}(\tau) \sim_{h'|\sigma} \bigcup_{\tau\in\Sigma(h'|\sigma)}\tau\{z\cup w\}.$$

From strategic planning (as applied to $\succcurlyeq_{h'}$), we know that there exists some $p^{\{z,w\}} : \Sigma(h') \to \{z, w\}$ such that

$$\bigcup_{\pi \in \Sigma(h')} \pi p^{\{z,w\}}(\pi) \sim_{h'} \bigcup_{\pi \in \Sigma(h')} \pi \{z \cup w\}.$$

Lemma 5 completes the claim.

(WARP). For some finite $W, V \subset Z$ let $\hat{p}^W : \Sigma(h'|\sigma) \to W$ be a strategic plan (over $\succcurlyeq_{h'|\sigma}$). Assume that $\hat{p}^W(\rho) \in W \cap V$. So, by lemma 5, there exists a strategic plan, $p^W : \Sigma(h') \to W$ such that $p^W(\sigma\rho) = \hat{p}^W(\rho)$. Let $\hat{p}^V : \Sigma(h'|\sigma) \to V$ be some other strategic plan such that $\hat{p}^V(\rho) \in W \cap V$. Again we have an extension such that $p^V(\sigma\rho) = \hat{p}^V(\rho)$.

So by WARP the function

$$\bar{p}^W(\pi) = \begin{cases} p^V(\pi) & \text{if } \pi = \sigma\rho \\ p^W(\pi) & \text{if } \pi \neq \sigma\rho \end{cases}$$

is also a strategic plan for W, according to $\succcurlyeq_{h'}$. Applying lemma 5 again provides the result.

A.3 Proof of Theorem 2

Step-1: Recursive Structure on Consumption Streams Let us begin by considering $\hat{\succcurlyeq}$, the restriction of $\succcurlyeq$ to degenerate IHCPs that assign pre-determined streams of consumption. Since the set of such IHCPs is closed in in Z, continuity is inherited by $\hat{\succcurlyeq}$. Let $\hat{V} : \times_{i=1}^{\infty} X \to \mathbb{R}$ be the numerical representation of $\hat{\succcurlyeq}$. Further, note that Singleton Recursivity is equivalent on this domain to fully stationary and separable preferences a la Koopmans [1960] (Section 16, line (11)). Thus, there exists a function $u : X \to \mathbb{R}$ and a strictly increasing function, continuous in its second argument $F : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ such that

$$\hat{V}(\boldsymbol{\sigma}) = F[u(\sigma_1), \hat{V}(\{\boldsymbol{\sigma}\}_{\boldsymbol{n=2}}^{\infty})]. \tag{1}$$

Step-2: From $\boldsymbol{\times^{\infty} X}$ to $\boldsymbol{Z}$. Let $\bar{V}$ be a continuous representation of $\succcurlyeq$. Since $\bar{V}$ also represents $\times^{\infty} X$, there exists some continuous, strictly increasing, $\bar{\psi} : \bar{V}[\times^{\infty} X] \to \mathbb{R}$ such that $\bar{\psi}(\bar{V}(\boldsymbol{\sigma})) = \hat{V}(\sigma)$. Note that $\bar{V}[\times^{\infty} X]$ is compact by its closure in Z and the continuity of $\bar{V}$. So, according to Husseinov (2010) there exists a continuous, strictly increasing $\psi : \mathbb{R} \to \mathbb{R}$ such that $\bar{\psi}(x) = \psi(x)$ for $x \in \bar{V}[\times^{\infty} X]$. Define $V = \psi \circ \bar{V}$. V is a continuous representation of $\succcurlyeq$ that coincides with (1) over elements of $\times^{\infty} X$.

For each $x \in \mathbb{R}$ we can define the function $F_x : \mathbb{R} \to \mathbb{R}$ by $F_x(y) = F[x, y]$. For any x, F_x is strictly monotone. Using these couriering functions, we can define $V_a : Z \to \mathbb{R}$ for each $a \in X$ as

$$V_a : z \mapsto F_{u(a)}^{-1}(V(az)).$$

The functional V_a represents $\succcurlyeq_a$ by construction. Indeed,

$$\begin{aligned} z \succcurlyeq_a w &\iff az \succcurlyeq aw \\ &\iff V(az) \geqslant V(aw) \\ &\iff F^{-1}_{u(a)}(V(az)) \geqslant F^{-1}_{u(a)}(V(aw)). \end{aligned}$$

Note that for consumption streams, $V_a(\boldsymbol{\sigma}) = V(\boldsymbol{\sigma})$ by (1). Now given any V_h, we can define $V_{h\cup a}$ using the same map,

$$V_{h\cup a} : z \mapsto F^{-1}_{u(a)}(V_h(az)),$$

to inductively define all hypothetical preferences. Note, of course, that we have:

$$V_h(az) = F[u(a), V_{h\cup a}(z)]. \tag{2}$$

Notice that if $a \in h$ then $V_{h\cup a} = V_h$. To see this assume that for some z, $V_h(z) > V_{h\cup a}(z)$. Now by the density of streams of universal outcomes, there exists some $\boldsymbol{\sigma}$ such that $V_h(z) > V_h(\boldsymbol{\sigma}) = V_{h\cup a}(\boldsymbol{\sigma}) > V_{h\cup a}(z)$. But this is a direct contradiction of H-stationarity.

Step-3: The existence of a Krepsian State Space. Theorem 1 states that there for each $\sigma \in \Sigma(h)$ the corresponding preference relation $\succcurlyeq_{h|\sigma}$ is a continuous weak order, and therefore representable by a continuous value function $V_{h|\sigma} : Z \longrightarrow \mathbb{R}$. We claim that these states form a Krepsian state space.

Lemma 6. *Let $z = (a_i, z_i)_{i\in M(z)}$ be a choice problem and $\sigma \in \Sigma(h)$. There exists some $i \in M(z)$ such that $(a_i, z_i) \sim_{h|\sigma} z$.*

Proof. We first claim that the lemma holds if z is a finite set. By identifying z as a collection of singleton menus $W = \{\{(a_i, z_i)\} \mid i \in M(z)\}$, the Strategic Planing axiom ensures that there exist $(a^\sigma_i, z^\sigma_i) \in z$, $\sigma \in \Sigma(h)$ such that

$$\bigcup_{\sigma\in\Sigma(h)} \sigma\{(a^\sigma_i, z^\sigma_i)\} \sim_h \bigcup_{\sigma\in\Sigma(h)} \sigma z.$$

By Lemma 4,

$$\sigma\{(a^\sigma_i, z^\sigma_i)\} \cup \bigcup_{\rho\in\Sigma(h),\rho\neq\sigma} \rho z \sim_h \bigcup_{\sigma\in\Sigma(h)} \sigma z,$$

which implies that there exists a strategic plan for $V = \{z, \{(a^\sigma_i, z^\sigma_i)\}\}$ such that $\{(a^\sigma_i, z^\sigma_i)\}$ is assigned to σ. By definition of $\succcurlyeq_{h|\sigma}$, $(a^\sigma_i, z^\sigma_i) \succcurlyeq_{h|\sigma} z$. On the other hand, by Preference for flexibility, $z \succcurlyeq_{h|\sigma} (a^\sigma_i, z^\sigma_i)$.Hence, we have $(a^\sigma_i, z^\sigma_i) \sim_{h|\sigma} z$, as desired.

Take any $z \in Z$. Since Z is metrizable with the Hausdorff metric, there exists a sequence $z^n \in Z$ such that $z^n \to z$ and z^n is a finite subset of z. By the above claim, there exists $(a^n, w^n) \in z^n$ such that $z^n \sim_h (a^n, w^n)$. Since $X \times Z$ is compact, without loss of generality, assume that (a^n, w^n) converges to some point (a, w). By continuity, $z \sim_h (a, w)$ with $(a, w) \in z$. This completes the proof. ☐

By the consequence of Lemma 6,

$$V_{h|\sigma}(z) = \max_{i \in M(z)} V_{h|\sigma}(a_i z_i). \tag{1}$$

For any $z \in Z$ define

$$f_h(z) = \{w \in Z | z \succcurlyeq_{h|\sigma} w, \forall \sigma \in \Sigma(h)\}.$$

The set $f_h(z)$ returns the intersections of the lower contour sets of each state induced preference. We now claim that $w \in f_h(z)$ if and only if $z \sim_h z \cup w$. This is the characterization of the dominance relation in Kreps [1979]:

$$\begin{aligned} w \in f_h(z) & \iff \\ \bigcup_{\sigma \in \Sigma(h)} \sigma z \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma\{z \cup w\} & \iff \\ z \sim_h \{z \cup w\}, & \end{aligned}$$

where the first implication follows from the definition of $\succcurlyeq_{h|\sigma}$ and WARP, and the second implication from H-stationarity.

From here we can follow Kreps' proof. Now define $\phi_h : \mathbb{R}^{\Sigma(h)} \to \mathbb{R}$ as any strictly increasing function that extends the map $\big(V_{h|\sigma}(z)\big)_{\sigma \in \Sigma(h)} \mapsto V_h(z)$. This is well defined since $\big(V_{h|\sigma}(z)\big)_{\sigma \in \Sigma(h)} = \big(V_{h|\sigma}(w)\big)_{\sigma \in \Sigma(h)}$ implies that $w \in f_h(z)$ and $z \in f_h(z)$ which in turn implies that $z \sim_h z \cup w \sim_h w$. Likewise, if $\big(V_{h|\sigma}(z)\big)_{\sigma \in \Sigma(h)} \geqslant \big(V_{h|\sigma}(w)\big)_{\sigma \in \Sigma(h)}$ with some strict equality then $w \in f_h(z)$ but $z \notin f_h(w)$ implying $z \sim_h z \cup w \succ_h w$. Plugging (1) into the aggregator we have

$$V_h(z) = \phi_h\big(\max_{i \in M(z)} V_{h|\sigma}(a_i, z_i)_{\sigma \in \Sigma(h)}\big).$$

Step-4: States within States. Now Theorem 1 tells us that $\succcurlyeq_{h|\sigma}$ obeys the same set of axioms as $\succcurlyeq$. Therefore, if the above claim holds, and $V_{h|\sigma}$ represents some $\succcurlyeq_{h|\sigma}$ then we can repeat the previous steps (1-3) to obtain the state dependent version of equation (2), which implies

$$(a, z) \succcurlyeq_{h|\sigma} (b, w) \iff F_\sigma[u_\sigma(a), V_{h \cup a|\sigma}(z)] \geqslant F_\sigma\left[u_\sigma(b), V_{h \cup b|\sigma}(w)\right],$$

From Singleton Recursivity we know that it is WLOG that $F_\sigma = F$ for all σ. Assume that this was not the case, so that for some σ, $\dot{V} : z \mapsto \max_{i \in M(z)} F(u_\sigma(a_i), \dot{V}(z_i))$ does not represent $\succcurlyeq_{h|\sigma}$. So there must be some z and w such that $z \succcurlyeq_{h|\sigma} w$ but $\dot{V}(w) > \dot{V}(z)$. By H-stationarity we have that $\sigma_1 z \succcurlyeq_{h|\sigma} \sigma_1 w$. Also, by the strict increasingness of F we have that $\dot{V}(\sigma_1 w) > \dot{V}(\sigma_1 z)$. Therefore, by the continuity of $\succcurlyeq_{h|\sigma}$ and by Singleton-Recursivity we can find sequences of universal outcome, $\boldsymbol{\pi}$ and $\boldsymbol{\tau}$ such that $\boldsymbol{\pi} \succcurlyeq_{h|\sigma} z \succcurlyeq_{h|\sigma} w \succcurlyeq_{h|\sigma} \boldsymbol{\tau}$ and $\dot{V}(\sigma_1 \boldsymbol{\tau}) > \dot{V}(\sigma_1 \boldsymbol{\pi})$, and hence $\dot{V}(\boldsymbol{\tau}) > \dot{V}(\boldsymbol{\pi})$.

But this contradicts the invariance of preferences regarding universal streams. Indeed, we choose some common normalization, over $\bar{a}$ and $\underline{a}$, then it is clear that $\dot{V} \equiv V$ on such a domain. But $\boldsymbol{\pi} \succcurlyeq_{h|\sigma} \boldsymbol{\tau} \iff$

$\boldsymbol{\pi} \succcurlyeq \boldsymbol{\tau}$, which is a contradiction.

This provides,

$$V_h(z) = \phi_h\big((V_{h|\sigma}(z))_{\sigma\in\Sigma(h)}\big) = \phi_h\big(\max_{i\in M(z)} F[u_\sigma(a_i), V_{h\cup a_i|\sigma}(z_i)]_{\sigma\in\Sigma(h)}\big).$$

Moreover, we can repeat this entire exercise starting with the conditional preferences, retaining any normalizations. Note we use the definition of a conditional strategic plan (i.e., over sequences in $\Sigma(h'|\sigma)$). Finally,

$$V_{h'|\sigma}(z) = \phi_{h'|\sigma}\big((V_{h'|\tau}(z))_{\tau\in\Sigma(h'|\sigma)}\big) = \phi_{h'|\sigma}\big(\max_{i\in M(z)} F[u_\tau(a_i), V_{h'\cup a_i|\tau}(z_i)]_{\tau\in\Sigma(h'|\sigma)}\big).$$

This is the representation we are after. It is clear from the construction of the information structure (i.e., $\Sigma(h)$), that if $h \subset h'$ then $\Sigma(h')$ is a refinement of $\Sigma(h)$, hence (i). (ii), i.e., that $u_\sigma = \phi_{h'|\sigma}(u_\tau)_{\tau\in\Sigma(h'|\sigma)}$ for $\sigma \in \Sigma(h)$, is an immediate consequence of Singleton-Recursivity, most easily seen with the normaliziation $u_\sigma(\underline{a}) = 0$ and $F(x,0) = x$. Lastly, (iii) follows from Lemma 3 and the definition of $\Sigma(h'|\sigma)$.

A.4 Proof of Theorem 3

(2 ⇒ Strategic Planning) Fix h and finite W. By Lemma 1 the σ-dominance relation induces a weak ordering, $\succcurlyeq_\sigma$, over Z for each each $\sigma \in \Sigma(h)$. So, by the finiteness of W there is a maximal element (in W) for each σ: call this w_σ. Define the assignment: $\bar{p} : \sigma \mapsto w_\sigma$.

We claim that $\bar{p}$ is a strategic plan. By construction $\bar{p}$ is in the envelope of W. Therefore, by applying Axiom M we have

$$\bigcup_{\rho\in\Sigma(h)} \rho p(\rho) \sim_h \bigcup_{\rho\in\Sigma(h)} \rho W.$$

(2 ⇒ WARP) Let $W, V \subseteq Z$, and strategic plans p^W and p^V, such that $p^W(\sigma), p^V(\sigma) \in V \cap W$ for some $\sigma \in \Sigma(h)$. Define $w = p^W(\sigma)$ and $v = p^V(\sigma)$. Since

$$\bigcup_{\rho\in\Sigma(h)} \rho p^W(\rho) \sim_h \bigcup_{\rho\in\Sigma(h)} \rho W,$$

Axiom M implies that p^W is in the envelope of W. Likewise, p^V in the envelope of V. So, by definition $v \succcurlyeq_\sigma z$ for all $z \in V$, hence $v \succcurlyeq_\sigma w$. Lastly, this implies that $\hat{p}^W$, defined by

$$\hat{p}^W(\pi) = \begin{cases} v \text{ if } \pi = \sigma \\ p^W(\pi) \text{ if } \pi \neq \sigma, \end{cases}$$

is in the envelope of W. Applying Axiom M again:

$$\sigma p^V(\rho) \cup \bigcup_{\rho\in\Sigma(h)\backslash\sigma} \rho(p^W(\sigma)) \sim_h \bigcup_{\rho\in\Sigma(h)} \rho W,$$

so WARP holds.

(1 ⇒ σ-Transativity). Claim: $z \succcurlyeq_\sigma w$ if and only if $z \succcurlyeq_{h|\sigma} w$. Indeed, letting $p \in As(\{z,w\},h)$ be a strategic plan we have:

$$z \succcurlyeq_{h|\sigma} w \iff \sigma z \cup \bigcup_{\rho \in \Sigma(h)\setminus\sigma} \rho p(\rho) \succcurlyeq_h \sigma w \cup \bigcup_{\rho \in \Sigma(h)\setminus\sigma} \rho p(\rho) \tag{2}$$

$$\iff \sigma z \cup \bigcup_{\rho \in \Sigma(h)\setminus\sigma} \rho\big(z \cup w\big) \succcurlyeq_h \sigma w \cup \bigcup_{\rho \in \Sigma(h)\setminus\sigma} \rho\big(z \cup w\big) \tag{3}$$

$$\iff z \succcurlyeq_\sigma w, \tag{4}$$

where the equivalence in (2) invokes WARP, and between (2) and (3) invokes Lemma 4.

Given the claim, the transitivity of $\succcurlyeq_{h|\sigma}$ (by Theorem 1 of the previous file) directly implies the transitivity of $\succcurlyeq_\sigma$.

(1 ⇒ Modularity). Fix $p \in As(W,h)$ with p in the envelope of W. By Strategic planning there exists strategic plan p^W. By the definition of the envelope of W, $p(\sigma) \succcurlyeq_\sigma w$ for all $w \in W$, which by the claim implies $p(\sigma) \succcurlyeq_{h|\sigma} w$ for all $w \in W$, in particular, $p(\sigma) \succcurlyeq_{h|\sigma} p^W(\sigma)$. By WARP we have that p is a SP: the indifference condition of Axiom M holds. The opposite direction is near identical and thus omitted.

A.5 Proof of Proposition 1

(WARP and Dominance ⇒ Modularity) Suppose that p is in the envelope of W. For all $w \in W$ and $\sigma \in \Sigma(h)$, $p(\sigma) \succcurlyeq_\sigma w$. Since W is finite, denote $W = \{w^1, \cdots, w^m\}$. Let $f(\sigma) = p(\sigma)$ and $g^i(\sigma) = w^i$ (constant function) $i = 1, \cdots, m$. Since f σ-dominates g^i state-by state, by taking $k = f$, Dominance implies

$$\bigcup_{\sigma \in \Sigma(h)} \sigma f(\sigma) \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup g^1(\sigma)\}.$$

By taking $k = g^1$, Dominance implies

$$\bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup g^1(\sigma)\} \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup g^1(\sigma) \cup g^2(\sigma)\}.$$

By repeating this argument finite times with setting $k(\sigma) = \cup_{i=1}^n g^i(\sigma)$, we have

$$\bigcup_{\sigma \in \Sigma(h)} \sigma f(\sigma) \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma W.$$

On the other hand, assume that there exists $p \in As(W,h)$ such that

$$\bigcup_{\sigma \in \Sigma(h)} \sigma p(\sigma) \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma W.$$

By WARP and Lemma 4,[10] for all $\sigma \in \Sigma(h)$ and $w \in W$,

$$\sigma p(\sigma) \cup \bigcup_{\rho \in \Sigma(h) \setminus \sigma} \rho\{p(\sigma), w\} \sim_h \bigcup_{\rho \in \Sigma(h)} \rho\{p(\sigma), w\}$$

So by definition, $p(\sigma) \succcurlyeq_\sigma w$. Hence, p is in the envelope of W.

(Modularity $\Rightarrow$ Dominance) Suppose that $\succcurlyeq$ satisfies Modularity. Take any $f, g, k \in$ *As(Z,h)* such that $f \succcurlyeq_\sigma g$ for each $\sigma \in \Sigma(h)$. Let $W = \{f(\sigma), g(\sigma), k(\sigma)\}_{\sigma \in \Sigma(h)}$. By Lemma 1, there exists $p \in$ *As(W,h)* such that p is in the envelope of W. In particular, since $f(\sigma) \succcurlyeq_\sigma g(\sigma)$, we can assume that p is a selection from $\{f(\sigma), k(\sigma)\}$ for all σ. By Modularity,

$$\bigcup_{\sigma \in \Sigma(h)} \sigma p(\sigma) \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma W.$$

By Flexibility,

$$\bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup k(\sigma)\} \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma p(\sigma) \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma W \succcurlyeq_h \bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup g(\sigma) \cup k(\sigma)\}.$$

On the other hand, it follows from Flexibility that

$$\bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup g(\sigma) \cup k(\sigma)\} \succcurlyeq_h \bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup k(\sigma)\}.$$

Therefore,

$$\bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup k(\sigma)\} \sim_h \bigcup_{\sigma \in \Sigma(h)} \sigma\{f(\sigma) \cup g(\sigma) \cup k(\sigma)\}.$$

B Strategic Plans over Weak Orderings

The analysis by which strategic planning separates the preference ordering over the history via the construction of $\Sigma(h)$ is not restricted to strict orderings. Although the actual sequences that identify each state may become very long they are ensured to exist and be finite.

We will first examine the case with $|h| = 2$ to gain intuition then show how this can be extended. For this section let m denote the universally worst outcome. If $h = \{a, b\}$, then there are 3 states: $a \succ b$, $a \sim b$, and $b \succ a$.

The complication is creating a sequence that is optimal when the DM is strictly prefers a to b but *not* optimal when the preference is weak. This of course cannot be simply a choice of a since indifference does not rule out this choice as optimal. However, by using a longer sequence and the universal worst outcome we can find the desired sequences. Consider,

[10] Note: the other direction of this proof shows that Dominance is sufficient to ensure the existence of a strategic plan for all finite W, and hence, with the addition of WARP we have satisfied the conditions to apply Lemma 4.

1:	$a \succ b$	$aa, aa, \ldots, am$
2:	$a \sim b$	$ab, ab, \ldots, ab$
3:	$b \succ a$	$bb, bb, \ldots, bm$

Where "..." denote the repetition of the previous consumption entries. If the decision maker is indifferent, is is clear that 2 is the unique optimal choice as it is the only sequence to avoid consumption of m. Moreover, for sufficiently long repetitions 1 becomes the unique choice for a DM who prefers a to b strictly, since a one time (highly discounted) consumption of m is better that the repeated consumption of b.

This intuition is extendable to the general case. To outline how, we show the case with $h = \{a, b, c\}$. There are 13 weak orderings:

1:	$a \succ b \succ c$	$aab, aab, \ldots, mab$	$aab, aab, \ldots, amb$	$aab, aab, \ldots, aam$
2:	$a \succ c \succ b$	$aac, aac, \ldots, mac$	$aac, aac, \ldots, amc$	$aac, aac, \ldots, aam$
3:	$b \succ a \succ c$	$bab, bab, \ldots, mab$	$bab, bab, \ldots, bmb$	$bab, bab, \ldots, bam$
4:	$b \succ c \succ a$	$bcb, bcb, \ldots, mcb$	$bcb, bcb, \ldots, bmb$	$bcb, bcb, \ldots, bcm$
5:	$c \succ a \succ b$	$acc, acc, \ldots, mcc$	$acc, acc, \ldots, amc$	$acc, acc, \ldots, acm$
6:	$c \succ b \succ a$	$bcc, bcc, \ldots, mcc$	$bcc, bcc, \ldots, bmc$	$bcc, bcc, \ldots, bcm$
7:	$a \succ b \sim c$	$aab, aab, \ldots, mab$	$aab, aab, \ldots, amb$	$aab, aac, \ldots, aab$
8:	$b \succ a \sim c$	$bab, bab, \ldots, mab$	$bab, bcb, \ldots, bab$	$bab, bab, \ldots, bam$
9:	$c \succ a \sim b$	$acc, bcc, \ldots, acc$	$acc, acc, \ldots, amc$	$acc, acc, \ldots, acm$
10:	$a \sim b \succ c$	$aab, bab, \ldots, aab$	$aab, aab, \ldots, amb$	$aab, aab, \ldots, aam$
11:	$a \sim c \succ b$	$aac, aac, \ldots, mac$	$aac, acc, \ldots, aac$	$aac, aac, \ldots, aam$
12:	$b \sim c \succ a$	$bcb, bcb, \ldots, mcb$	$bcb, bcb, \ldots, bmb$	$bcb, bcc, \ldots, bcb$
13:	$a \sim b \sim c$	$aab, bab, \ldots, aab$	$aac, acc, \ldots, aac$	$aac, aab, \ldots, aac$

One only needs to break up each pairwise distinct aspect of different preference relations into "blocks" that imitate the $|h| = 2$ case. I.e., the first block separates strict and weak indifference between a and b the second between a and c and finally between a and c. Since DMs preferences (over IHCPs) are stationary (since we consume only elements in the history and universal outcomes) this process separates all the weak preference relations from one another.

References

Rudiger Bachman and Giuseppe Moscariniz. Business cycles and endogenous uncertainty. *mimeo*, 2012.

Dirk Bergemann and Juuso Välimäki. Bandit problems. Technical report, Cowles Foundation for Research in Economics, Yale University, 2006.

Eddie Dekel, Barton L. Lipman, and Aldo Rustichini. Representing preferences with a unique subjective state space. *Econometrica*, 69(4), 2001.

David Dillenberger, Juan Sebastian Lleras, Philipp Sadowski, and Norio Takeoka. A theory of subjective learning. *Journal of Economic Theory*, 153, 2014.

Faruk Gul and Wolfgang Pesendorfer. Self-control and the theory of consumption. *Econometrica*, 72, 2004.

Youichiro Higashi, Kazuya Hyogo, and Norio Takeoka. Stochastic endogenous time preference. *Journal of Mathematical Economics*, 51:77–92, 2014.

Tjalling C Koopmans. Stationary ordinal utility and impatience. *Econometrica: Journal of the Econometric Society*, pages 287–309, 1960.

Sergei Koulayev. Search with dirichlet priors: estimation and implications for the search with dirichlet priors: estimation and implications for the search with dirichlet priors: estimation and implications for the price elasticity of demand. *mimeo*, 2010.

David M. Kreps. A representation theorem for 'preference for flexibility'. *Econometrica*, 47:565–576, 1979.

R Vijay Krishna and Philipp Sadowski. Dynamic preference for flexibility. *Econometrica*, 82(2):655–703, 2014.

Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58, 1952.

Christopher A Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690, 2003.

Norio Takeoka. Subjective probability over a subjective decision tree. *Journal of Economic Theory*, 136(1): 536–571, 2007.