

Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Cowles Foundation Discussion Papers

Cowles Foundation

1-1-1999

Preference for Information and Dynamic Consistency

Simon Grant

Atsushi Kajii

Ben Polak

Follow this and additional works at: <https://elischolar.library.yale.edu/cowles-discussion-paper-series>



Part of the [Economics Commons](#)

Recommended Citation

Grant, Simon; Kajii, Atsushi; and Polak, Ben, "Preference for Information and Dynamic Consistency" (1999). *Cowles Foundation Discussion Papers*. 1456.

<https://elischolar.library.yale.edu/cowles-discussion-paper-series/1456>

This Discussion Paper is brought to you for free and open access by the Cowles Foundation at EliScholar – A Digital Platform for Scholarly Publishing at Yale. It has been accepted for inclusion in Cowles Foundation Discussion Papers by an authorized administrator of EliScholar – A Digital Platform for Scholarly Publishing at Yale. For more information, please contact elischolar@yale.edu.

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 1208

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

PREFERENCE FOR INFORMATION AND DYNAMIC CONSISTENCY

Simon Grant, Atsushi Kajii and Ben Polak

January 1999

Preference for Information and Dynamic Consistency.*

Abstract

We provide necessary and sufficient conditions for a dynamically consistent agent always to prefer more informative signals (in single-agent problems). These conditions do not imply recursivity, reduction or independence. We provide a simple definition of dynamically consistent behavior, and we discuss whether an intrinsic information lover (say, an anxious person) is likely to be dynamically consistent.

JEL Classification: D80, D83

Keywords: information, non-expected utility, dynamic consistency, randomization, anxiety.

Simon Grant
Faculty of Economics and Commerce,
Australian National University

Atsushi Kajii
Institute of Policy and Planning Sciences
University of Tsukuba

Ben Polak
Economics Department
Yale University

*The analysis in this paper was previously part of a working paper called "Preference for Information". We thank Ngo Van Long, Tim van Zandt and two anonymous referees.

1 Introduction

Blackwell’s (1953) theorem states (among other things) that, in single-agent problems, if the agent is an expected utility maximizer, then she always prefers more information to less regardless of the set of actions available to her. By contrast, Wakker (1988) argues that if the agent is not an expected utility maximizer, then there must be a pair of signals, one more informative than the other, and a decision problem, such that the agent strictly prefers the less informative signal.¹ The key to Wakker’s result is that his agent is not dynamically consistent: her behavior within a decision tree is not always what she would choose *ex ante* if she could commit. Indeed, the reason Wakker’s agent sometimes chooses a less informative signal is in an attempt to self-commit. Machina (1989) argues that Wakker’s finding comes from too narrow a definition of dynamic consistency. Machina proposed a broader definition; one that does not require preferences to conform with expected utility theory, or even to be recursive.²

An open question confronted by this paper is whether an agent who is dynamically consistent in the Machina sense, will always prefer more information to less. Like Machina, we consider agents who violate the expected utility axioms, and whose preferences are not necessarily recursive. We go beyond Machina, however, in that we also consider violations of the reduction of compound lottery axiom. In an earlier paper (Grant, Kajii and Polak (forthcoming)), we considered what we called intrinsic preference for information: desiring information not only to make actions contingent on what you learn, but also (for example) to relieve anxiety. Such preferences entail violations of the reduction axiom. In that paper, however, we only studied the special case where the agent has just one action available to her.³ In that context, of course, dynamic consistency is not an issue. This paper extends the analysis to the many-action case. The question then becomes whether, given dynamic consistency, the conditions for an agent to prefer more information in the special single-action case are necessary and sufficient for an agent always to prefer more information in the general many-action case.

The answer is: “almost”. The reason for the “almost”, is that non-expected utility maximizers are sometimes not indifferent among randomizations over indifferent outcomes. If the agents can perform randomizations for themselves, this is not a problem. But if they rely on the realizations of a signal to serve as a randomization device, then we need some extra restrictions if they are always to prefer more informative signals. The necessary and sufficient restriction turns out to be (roughly) that the agent never strictly prefers to randomize. The exact restriction is implied by — but does not imply — recursivity.

¹ This work was later extended by Schlee (1990, 1991) and by Safra & Sulganik (1995).

² Machina’s notion of dynamic consistency is similar to what McClennen (1990) calls “resolute choice”.

³ While we borrow a result from the earlier paper, we do not rehearse the proof. The reader is referred to that paper for this and a fuller discussion of intrinsic preference for information in the single-action case.

Section 2 introduces a natural definition of dynamic consistency that captures Machina’s intuition. We say that an agent is dynamically consistent if her actual behavior within a decision tree conforms to that she would have chosen *ex ante* were she able to commit. We derive the latter by looking at the agent’s preferences among trees. Section 3 contains the main result: necessary and sufficient conditions such that a dynamically consistent agent always prefers more information to less. Section 4 discusses this result. For example, our result simply assumes dynamically consistent behavior. We suggest, by means of an example, that if an agent intrinsically strictly prefers more information, then this agent might fail to be dynamically consistent.

2 Dynamic Consistency.

We are interested in preference between, and behavior within, decision trees. We focus on trees in which an agent chooses an action after the resolution of a signal. Outcomes depend on the agent’s action and on the state of nature. For an example, consider a farmer who chooses her fertilizer after receiving a weather forecast. The success of her crop depends on the weather and on the fertilizer she chooses. In this section, we allow for decision trees that are slightly more general than this in that the set of available actions can depend on the realization of the signal. Perhaps a particular kind of fertilizer might not be available if bad weather is forecast. Preference for information is a statement about the agent’s preferences *among* such simple decision trees. Loosely (we will be more precise in the next section), it states that the agent prefers trees associated with more informative signals. Dynamic consistency is a statement about how the agent’s behavior *within* such a decision tree conforms to the behavior she would choose, *ex ante*, if she could commit.

To formalize these ideas, let Ω be the set of states of nature, let \mathcal{X} be the set of consequences, let \mathcal{A} be a set of actions, and let \mathcal{C} be the set of functions of the form $c : \Omega \times \mathcal{A} \rightarrow \mathcal{X}$, that assign a consequence to each state-action pair. In the following, unless otherwise stated, both Ω and \mathcal{X} are taken to be the interval $[0, 1]$.⁴ We will confine attention to trees involving finite signals and finite action sets. A typical tree, τ , then consists of:

- a finite set of signal realizations, (s_1, \dots, s_n) ;
- the posterior beliefs on the state space Ω , (P_1, \dots, P_n) , induced by these realizations;
- the finite action sets, (A_1, \dots, A_n) , each $A_i \subset \mathcal{A}$, associated with the realizations;
- the (prior) probabilities, (q_1, \dots, q_n) , of observing these realizations;
- and a consequence function c in \mathcal{C} .

Let \mathcal{T} denote the set of all such trees; that is, trees of the form: $\langle (s_i, P_i, A_i; q_i)_{i=1}^n; c \rangle$. Let $\succeq_{\mathcal{T}}$ be a complete and transitive preference relation over such trees. This preference relation describes

⁴ The analysis can readily be extended to outcome sets that are general compact metric spaces if we assume that all welfare-relevant risk can be characterized as risk over the ranks of outcomes; see Grant, Kajii & Polak (1992).

a choice, made ex ante, among trees. For example, by selecting ex ante among different crops to plant, each with a different response to the weather, our farmer can choose which consequence function she faces. By selecting ex ante among news sources, she can choose her signal. By selecting ex ante among supply contracts, she can choose the set of available fertilizers. Different trees can be associated with different priors over the relevant states of nature. For example, if we think of Ω as being the weather during the peak growing season, then selecting ex ante among different crops or different planting times can affect this distribution.

Within each given tree, the agent chooses an action at each decision node; that is, after each signal realization. More generally, the agent might randomize over the available actions but, for the moment, consider just non-random behavior. Let b be a *behavior function* that, for each tree τ in \mathcal{T} , assigns an action from A_i to each signal realization s_i of τ . Let b_τ be the corresponding behavior in the specific tree τ , and $b_\tau(s_i)$ be the corresponding action taken in the tree τ after seeing the specific signal realization s_i . That is, $b_\tau(s_i)$ is an element of A_i . The behavior $b(\tau)$ is *not* to be thought of as a ‘plan of action’ (or ‘strategy’) formed ex ante, but rather as a description of the agent’s actual behavior within the tree τ . These do not necessarily coincide. That is, what the agent actually does within the tree τ is not necessarily the same thing as she would choose were she able to commit to those choices ex ante.

Given a decision tree, each behavior function b induces a two-stage lottery over outcomes. In our example, suppose that there are two weather states, wet or dry. Suppose there are two possible forecasts of the weather, s_1 and s_2 , which occur with probabilities q_1 and q_2 respectively. The posterior probability of wet given s_1 is p_1 , and given s_2 is p_2 . There are two types of fertilizer available at each state, a_1 and a_2 . If a_1 is chosen then the crop yield is x if wet and x' if dry. Conversely, if a_2 is chosen, then the yield is x' if wet and x if dry. Let us refer to this tree as $\hat{\tau}$. Suppose that the agent’s behavior in this tree, $b_{\hat{\tau}}$, is to choose a_1 at s_1 and a_2 at s_2 . Then the induced two-stage lottery assigns first-stage probability q_1 to the second-stage lottery $[x, p_1; x', (1 - p_1)]$, and assigns first-stage probability q_2 to the second-stage lottery $[x, (1 - p_2); x', p_2]$.

More formally, let $\mathcal{L}(\mathcal{X})$ be the set of probability measures on \mathcal{X} . Let $M(P, a, c)$ be the probability measure in $\mathcal{L}(\mathcal{X})$ where, for any subset E of \mathcal{X} , $M(P, a, c)(E) = P(\{\omega \in \Omega : c(\omega, a) \in E\})$. That is, given a probability distribution P over the set of states, $M(P, a, c)(E)$ is the probability that an outcome in E occurs if the action a is chosen. Let $\mathcal{L}_0(\mathcal{L}(\mathcal{X}))$ be the set of probability measures on $\mathcal{L}(\mathcal{X})$ with finite support. Let $\psi(b_\tau; \tau)$ be the two-stage lottery induced by the behavior function b in the tree τ . This is the element of $\mathcal{L}_0(\mathcal{L}(\mathcal{X}))$ that assigns to each (second-stage) lottery F in $\mathcal{L}(\mathcal{X})$ the sum of all the q_i in the set $\{q_i : M(P_i, b_\tau(s_i), c) = F\}$.⁵

A special class of trees are those where the agent is ‘committed to an action ahead of time’.

⁵ Notice that, by the definition of a two-stage lottery, if $[(M(P_1, b_\tau(s_1), c), q_1; \dots; (M(P_n, b_\tau(s_n), c), q_n)]$ and $[(M(P_1, b'_\tau(s_1), c), q_1; \dots; (M(P_n, b'_\tau(s_n), c), q_n)]$ both assign the same total first-stage probability to each second-stage lottery in $\mathcal{L}(\mathcal{X})$ then they are identified as the same element of $\mathcal{L}_0(\mathcal{L}(\mathcal{X}))$.

Formally, let the set of *commitment trees* be the subset of \mathcal{T} for which each action set A_i is a singleton. In our farmer story, an example of a commitment tree would be one in which the only fertilizer available after seeing forecast s_1 is a_1 , and the only fertilizer available after seeing s_2 is a_2 . In this tree, the farmer is committed beforehand to the same fertilizers that her behavior $b_{\hat{\tau}}$ would have chosen within the less restrictive tree $\hat{\tau}$ described above. That is, the only possible behavior in this commitment tree induces the same two-stage lottery over outcomes as that induced by $b_{\hat{\tau}}$ in the tree $\hat{\tau}$.

We are going to use the agent’s preferences among commitment trees to help define dynamic consistency. Let \succeq_2 denote the restriction of $\succeq_{\mathcal{T}}$ to the set of commitment trees. Since the agent’s behavior within each commitment tree is degenerate, each such tree maps directly to a two-stage lottery in $\mathcal{L}_0(\mathcal{L}(\mathcal{X}))$. Therefore, we will interpret \succeq_2 as representing the agent’s preferences over such two-stage lotteries.⁶

We have now defined an agent’s preference relation over two-stage lotteries and her behavior within all trees in \mathcal{T} . Informally, it is natural to say that an agent is dynamically consistent if her behavior within each tree induces a two-stage lottery that is best among those lotteries that were achievable in that tree, where ‘best’ means according to her ex ante preference relation over such lotteries. In our farmer story, in the game $\hat{\tau}$ described above, there are four possible (non-random) behaviors. We can denote these by (a_1, a_1) , (a_1, a_2) , (a_2, a_1) and (a_2, a_2) . If each of these behaviors induces a different two-stage lottery then there are also four achievable two-stage lotteries: $\psi((a_1, a_1); \hat{\tau})$, $\psi((a_1, a_2); \hat{\tau})$, $\psi((a_2, a_1); \hat{\tau})$ and $\psi((a_2, a_2); \hat{\tau})$. More generally, let $B(\tau)$ be the set of two-stage lotteries that can be induced by some non-random behavior b_{τ} in the tree τ .

Definition 1 *A preference relation between commitment trees \succeq_2 (i.e., a preference relation over two-stage lotteries), and a (non-random) behavior within trees b , are dynamically consistent if, for all trees τ in \mathcal{T} and all achievable two-stage lotteries X in $B(\tau)$, $\psi(b_{\tau}, \tau) \succeq_2 X$.*

The above notion of dynamic consistency links ex ante preferences and actual (interim) behavior. This definition is closely related to those of Johnsen & Donaldson (1985), McClennen (1990), Machina (1989), and Karni & Schmeidler (1991). Like theirs, our definition of dynamic consistency does not restrict preferences to be expected utility, or to be recursive, or to respect reduction of compound lotteries.

If an agent can predict her own behavior in any tree (that is, she knows her behavior function b), then, given an ex ante choice between two decision trees, she should choose the one in which

⁶ This interpretation implicitly assumes away ‘framing effects’; that is, if two commitment trees map to the same two-stage lottery then the agent is indifferent between them. Formally, if two commitment trees τ and τ' are such that, the (total) probability assigned by $\psi(b_{\tau}; \tau)$ to each (second-stage) lottery F in $\mathcal{L}(\mathcal{X})$ equals that assigned by $\psi(b_{\tau'}; \tau')$ to F , then $\tau \sim_{\mathcal{T}} \tau'$. (Since these are commitment trees, the behaviors within the trees, b_{τ} and $b_{\tau'}$, are uniquely determined.)

her behavior will induce the two-stage lottery that is better according to her ex ante preferences. This idea is usually called “sophisticated choice” to distinguish it from “myopic choice” where the agent fails to predict correctly.⁷ The following formalizes this standard idea.

Definition 2 *An agent is sophisticated if her preference relation between trees $\succeq_{\mathcal{T}}$ and her (non-random) behavior within trees b are such that, for all pairs of trees τ and τ' in \mathcal{T} , $\tau \succeq_{\mathcal{T}} \tau'$ if and only if $\psi(b_{\tau}, \tau) \succeq_2 \psi(b_{\tau'}, \tau')$.*

Combining these ideas, an agent who is both sophisticated and dynamically consistent, given an ex ante choice between two decision trees, should first look at the set of achievable two-stage lotteries in each tree. She should then pick the tree that permits the best two-stage lottery according to her ex ante preferences. Since she is dynamically consistent, she knows that her behavior within the tree will indeed induce this lottery. Again, the following formalizes this standard idea.

Observation *If an agent whose preference relation between trees is $\succeq_{\mathcal{T}}$, and whose (non-random) behavior within trees is b is both dynamically consistent and sophisticated then, for all pairs of trees τ and τ' in \mathcal{T} , $\tau \succeq_{\mathcal{T}} \tau'$ if and only if there is a achievable two-stage lottery X in $B(\tau)$ such that $X \succeq_2 X'$ for all achievable two-stage lotteries X' in $B(\tau')$.*

Random behavior. So far, we have defined dynamic consistency for non-random behavior, and we have determined what this implies for sophisticated ex ante choices among trees. We now want to extend these ideas to allow for randomization. A complication that arises once we allow the agent to randomize is that the precise two-stage lottery that is induced by such behavior depends on the time at which the randomization takes place. In our farmer story, an example of random behavior in the tree $\hat{\tau}$ could be for the farmer, at the weather forecast s_1 , to choose fertilizer a_1 with probability α and fertilizer a_2 with probability $(1 - \alpha)$; and, at the weather forecast s_2 , to choose fertilizer a_2 with probability 1. If the farmer observes the randomization device at the same time as she observes the forecast then this random behavior induces the two-stage lottery that assigns first-stage probability αq_1 to the second-stage lottery $[x, p_1; x', (1 - p_1)]$; assigns first-stage probability $(1 - \alpha)q_1$ to the second-stage lottery $[x, (1 - p_1); x', p_1]$; and assigns first-stage probability q_2 to the second-stage lottery $[x, (1 - p_2); x', p_2]$. It is as if there were three weather reports that the farmer could use to condition her fertilizer choice: a report s_{11} that occurred with probability αq_1 ; an identical report s_{12} that occurred with probability $(1 - \alpha)q_1$, and the report s_2 .

On the other hand, suppose the farmer only observes the realization of the randomization device when the true state is revealed. We can think of this as the farmer choosing a random

⁷ See, for example, O'Donoghue & Rabin (1996).

action at s_1 . In this case the random behavior induces the two-stage lottery that assigns: first-stage probability q_1 to the second-stage lottery $[x, \alpha p_1 + (1-\alpha)(1-p_1); x', \alpha(1-p_1) + (1-\alpha)p_1]$; and assigns first-stage probability q_2 to the second-stage lottery $[x, (1-p_2); x', p_2]$. These are different two-stage lotteries. As we shall see in Section 3, once we move away from standard preferences (in particular, once we drop the reduction of compound lotteries axiom), this distinction can matter.

More generally, we can think of two types of randomized behavior within our set of decision trees \mathcal{T} . With *early-resolution randomization*, the realization of the ‘randomizing device’ is contemporaneous with the resolution of the signal: the agent can be thought of as selecting a non-random action at the second stage, with that choice depending jointly on the realizations of the randomizing device and of the signal. With *late-resolution randomization*, the resolution of the randomizing device is contemporaneous with the final resolution of the state: that is, the agent mixes at the second stage. If the agent has access to randomization devices, it seems natural to allow her both early- and late-resolution random behaviors. Let \tilde{b} be a random behavior function. It describes how the agent randomizes (both early and late) within each tree in \mathcal{T} . Let \tilde{b}_τ describe the early and late randomizations of the agent in the specific tree τ . Let $\psi(\tilde{b}_\tau, \tau)$ be the two-stage lottery induced by the random behavior \tilde{b}_τ in the tree τ . And let $\tilde{B}(\tau)$ be the set of two-stage lotteries that can be induced by some random behavior \tilde{b}_τ in the tree τ .

Clearly a commitment tree cannot commit an agent to randomize. Nevertheless, both dynamic consistency and of sophistication allowing for randomization can be defined exactly analogously to their definitions when randomization is excluded. A preference relation between commitment trees \succeq_2 (i.e., a preference relation over two-stage lotteries) and a random behavior within trees \tilde{b} , are dynamically consistent if, for all trees τ in \mathcal{T} and all achievable two-stage lotteries X in $\tilde{B}(\tau)$, $\psi(\tilde{b}_\tau, \tau) \succeq_2 X$. An agent whose preference relation between trees is $\succeq_{\mathcal{T}}$, and whose random behavior within trees is \tilde{b} is sophisticated if, for all pairs of trees τ and τ' in \mathcal{T} , $\tau \succeq_{\mathcal{T}} \tau'$ if and only if $\psi(\tilde{b}_\tau, \tau) \succeq_2 \psi(\tilde{b}_{\tau'}, \tau')$. If an agent whose preference relation between trees is $\succeq_{\mathcal{T}}$, and whose random behavior within trees is \tilde{b} is both dynamically consistent and sophisticated, then, for all pairs of trees τ and τ' in \mathcal{T} , $\tau \succeq_{\mathcal{T}} \tau'$ if and only if there is a achievable two-stage lottery X in $\tilde{B}(\tau)$ such that $X \succeq_2 X'$ for all achievable two-stage lotteries X' in $\tilde{B}(\tau')$.

3 Preference for Information

In this section, we provide necessary and sufficient conditions for a sophisticated dynamically consistent agent always to prefer more information to less. The conditions will be restrictions on the agent’s ex ante preferences and, in particular, on her preferences over commitment trees (that is, over two-stage lotteries).

The following generic definitions are useful. Let \mathcal{Z} be a convex set in a linear space, and let $\mathcal{L}_0(\mathcal{Z})$ denote the set of probability measures on \mathcal{Z} with finite support. It is conventional to write

$[z_1, p_1; \dots; z_n, p_n]$ to denote the measure (or lottery) in $\mathcal{L}_0(\mathcal{Z})$ that assigns, to each z in \mathcal{Z} , the probability $\sum_{i: z_i=z} p_i$. Writing lotteries in this way does *not* imply that there are n elements in the support of this lottery, since the z_i need not be distinct. For example, $[z_1, p_1; z_1, (p - p_1); z_2, (1 - p)]$ is identical to $[z_1, p; z_2, (1 - p)]$.

Definition 3 *An elementary linear bifurcation of a measure $[(z_j, p_j)_{j=1}^m]$ in $\mathcal{L}_0(\mathcal{Z})$ is a measure of the form $[z_1, p_1; \dots; z_{k-1}, p_{k-1}; z'_k, \beta p_k; z''_k, (1 - \beta)p_k; z_{k+1}, p_{k+1}; \dots; z_m, p_m]$ in $\mathcal{L}_0(\mathcal{Z})$, where β is in $[0, 1]$, and $z_k = \beta z'_k + (1 - \beta)z''_k$.*

These two lotteries are identical except that the ‘ k th outcome’, z_k , of the first lottery is the conditional mean of the two outcomes z'_k and z''_k in the second lottery. That is, the probability weight on the k th outcome is distributed between two outcomes in mean-preserving way. It does not follow from this definition that the support of the original lottery is a subset of the support of the elementary linear bifurcation: the outcomes z'_k and z''_k could both coincide with outcomes in the support the original lottery. This will be important when we consider preference for randomization below.

We use the standard Blackwell (1953) definition of more information. A signal is defined by its (finite) list of possible realizations S , and its likelihood function $\lambda : S \times \Omega \rightarrow [0, 1]$, where for any ω in Ω , $\sum_{s \in S} \lambda(s|\omega) = 1$. Then:

Definition 4 *The signal (S, λ) is more informative than the signal (S', λ') with respect to the prior belief π on Ω if there exists a function $\alpha : S' \times S \rightarrow [0, 1]$ such that $\sum_s \alpha(s', s) = 1$ for all s in S , and $\lambda'(s'|\omega) = \sum_s \alpha(s', s)\lambda(s|\omega)$ for all ω in Ω such that $\pi(\omega) > 0$ and all s' in S' .⁸*

For example, in our farmer story, suppose that there are two possible weather-forecast services, each of which gives two possible forecasts, (s_1, s_2) and (s'_1, s'_2) . As before, suppose that the weather is either wet or dry. If the weather is going to be wet, the probability of the first service reporting s_1 is λ_1 , and the probability of the first service reporting s_2 is λ_2 . The second forecast service, however, operates by listening to the first service forecast, copying it down quickly, and then re-broadcasting it in its own words. This copying can lead to errors. Let α_1 be the probability that s_1 gets wrongly re-reported as s'_2 , and let α_2 be the probability that s_2 gets wrongly re-reported as s'_1 . Intuitively, the imperfectly copied weather forecast service is less informative than the original, and this is confirmed by Blackwell’s definition. For example, if it is wet, the probability of the second service reporting s'_1 is $(1 - \alpha_1)\lambda_1 + \alpha_2\lambda_2$.

To apply the apparatus of the previous section, we want to associate signals with decision trees in \mathcal{T} . In our farmer story, suppose that the farmer’s prior probability of it being wet is π . Suppose

⁸ More generally, we say that one signal is more informative than another if it is more informative with respect to all priors.

that he has the same two fertilizers (with the same consequences for his crop) as in the tree $\hat{\tau}$. The first weather-forecasting service then induces a tree like $\hat{\tau}$ where $q_1 := \pi\lambda_1 + (1 - \pi)(1 - \lambda_2)$, $p_1 := \pi\lambda_1/q_1$ and where q_2 and p_2 are defined similarly.

More generally, for each signal (S, λ) , given a prior belief π on Ω , we can compute the unconditional probability of observing each realization. We can ignore signal realizations that occur with zero probability, so let $(s_1, \dots, s_N) \subseteq S$ denote the list of signal realizations that occur with positive probability. For each such s_i , let q_i in $(0, 1]$ denote this probability, and let P_i denote the posterior belief on Ω induced by that realization. Thus, given a prior π , a consequence function c in \mathcal{C} , and a fixed finite action set $A \in \mathcal{A}$, the signal (S, λ) induces the decision tree $\langle (s_i, P_i, A; q_i)_{i=1}^n; c \rangle$.

These trees induced in this way are special in that the action set, A , is the same for each realization. We shall call such trees *Blackwell decision trees*. Informally, we say that an agent always prefers more information if he always prefers the Blackwell decision trees induced by more informative signals, regardless of the (fixed) set of actions he has available. That is, information loving is a statement about the agent's ex ante preferences among decision trees, $\succeq_{\mathcal{T}}$. More formally,

Definition 5 *An agent is an information lover if: for all priors π on Ω , all consequence functions c in \mathcal{C} , and all finite action sets $A \in \mathcal{A}$, if the signal (S, λ) is more informative than the signal (S', λ') with respect to π , then the Blackwell tree induced by (S, λ) is weakly preferred to that induced by (S', λ') according to the agent's (ex ante) preference relation $\succeq_{\mathcal{T}}$.*

Blackwell (1953) showed that, regardless of their prior beliefs, available action sets and consequence functions, all atemporal expected utility maximizers are information lovers. The term 'atemporal' here means that preferences respect the reduction of compound lotteries axiom (hereafter, 'reduction'). Expected utility implies (among other things) that preferences respect the independence axiom. We will relax both reduction and independence. A loose intuition for Blackwell's result is that more information weakly enlarges the set of reduced lotteries achievable by the agent. Given reduction, the agent only cares about reduced lotteries. Thus, given a larger choice of such lotteries, the agent must be weakly better off. The reason this intuition is loose is that, strictly speaking, the achievable set is only enlarged if we allow the agent to randomize. Given independence, however, the agent will never strictly prefer to randomize: any utility level that can be attained by randomization can also be attained by some non-random choice. Once we drop reduction and independence, both steps in the above argument fail: the set of available multi-stage lotteries is *not* increased by more information; and the agent may strictly prefer randomization. We discuss each in turn.

The first major maintained hypothesis for Blackwell, then, was that preferences over two-stage

lotteries, \succeq_2 , satisfy the reduction axiom. This axiom is very restrictive. It is equivalent to saying that, if an agent has only one action available to her, she would be indifferent to information. In our story, however, uncertainty about the weather may cause the farmer to suffer anxiety. If so, she may want to get a weather forecast even if there is no actions she can take contingent on what she learns (say, she has only one fertilizer available). This case provides extra reason for the farmer to be an information lover. On the other hand, the farmer might enjoy the hope that ignorance permits. If so, she might strictly prefer not to get the forecast unless she can make her actions contingent on what she learns. In this case, the farmer is not an information lover, since the definition requires her (weakly) to prefer more informative signals regardless of her available action set.

In an earlier paper (Grant, Kajii & Polak, forthcoming), we focussed on such intrinsic preferences for and against information. That is, we explicitly relaxed the reduction axiom. In that paper, we restricted attention to the special case where the agent had only a single action available. Clearly, a *necessary* condition for an agent to be an information lover in general is that she (weakly) prefers more information to less in that special case. The earlier paper shows this is equivalent to the following property of preferences over two-stage lotteries. That is, if we drop reduction but we want information loving, we are going to need this property.

Definition 6 *An agent's preference relation over two-stage lotteries, \succeq_2 , satisfies single-action information loving if and only if, for any pair of two-stage lotteries X and Y in $\mathcal{L}_0(\mathcal{L}(\mathcal{X}))$ such that Y is an elementary linear bifurcation of X , $Y \succeq_2 X$.*

In this paper, we are looking for conditions that give information loving in the more general many-action case. It is natural to ask whether the above necessary property is also sufficient in the general case. It turns out that the answer depends on the agent's ability to randomize and her preferences concerning randomization.

The second major maintained hypothesis for Blackwell was that preferences satisfied the independence axiom, and hence the agent would never strictly prefer to randomize. To see why randomization may be an issue consider the following variant on the famous indivisible good allocation problem of Diamond (1967) and Machina (1989). Suppose that a mother has three children and one indivisible good. For reasons of procedural fairness, she may strictly prefer to randomize over which child should get the good. If the mother can randomize for herself then there is no problem. But suppose that the only 'randomization device' available to her is a signal about whether or not it will rain tomorrow. Suppose that the signal has three possible realizations, each of which occurs with probability $1/3$, inducing the posteriors $1/4$, $1/2$ and $3/4$, respectively. Then, the mother can use the signal to achieve her desired fair randomization. Now suppose that the signal is replaced with a second signal which has only two possible realizations, each of

which occurs with probability 1/2, inducing the posteriors 1/4 and 3/4, respectively. It is easy to check that the second signal is more informative by the standard Blackwell definition. Now, however, the mother cannot use the signal to achieve her desired fair randomization between the three children. Under our assumption that she cannot randomize for herself, it is possible for the mother to prefer the less informative signal.

The problem here is that a signal can be less informative but still have more signal realizations. Conditioning on signal realizations generates two-stage lotteries similar to those generated by early-resolution randomization. Thus, there is no advantage to having extra realizations if the agent can perform early resolution randomization for herself. Alternatively there is no advantage if the agent never strictly prefers such randomization. The following property rules out such preference, without imposing independence.

Definition 7 *An agent's preference relation over two-stage lotteries, \succeq_2 , satisfies conditional quasi-convexity if for all pairs of two-stage lotteries of the form $X = [(F_i, q_i)_{i=1}^n]$ and $Y = [(F_1, q_1; \dots; F_{j-1}, q_{j-1}; F_j^l, q_j; F_{j+1}, q_{j+1}; \dots; F_n, q_n)]$ in $\mathcal{L}_0(\mathcal{L}(\mathcal{X}))$, with $q_j > 0$: if $X \succeq_2 Y$ then, for all α in $(0, 1)$, $X \succeq_2 [(F_1, q_1; \dots; F_{j-1}, q_{j-1}; F_j, \alpha q_j; F_j^l, (1 - \alpha)q_j; F_{j+1}, q_{j+1}; \dots; F_n, q_n)]$.*

The following proposition provides necessary and sufficient conditions for sophisticated dynamically consistent agents always to prefer more information. We already know that single-action information loving is a necessary condition. If the agent can perform early-resolution randomizations then it is also sufficient. If the agent cannot perform early-resolution randomizations other than by conditioning on signal realizations, then conditional quasi-convexity is also necessary. Formally,

Proposition 1 *Suppose that an agent is sophisticated and dynamically consistent and that her preferences over two-stage lotteries are continuous.*

- (i) *If she can perform early-resolution randomizations then: she is an information lover if and only if her preferences over two-stage lotteries satisfy single-action information loving.*
- (ii) *If she cannot perform early-resolution randomizations other than by conditioning on signal realizations then: she is an information lover if and only if her preferences over two-stage lotteries satisfy both single-action information loving and conditional quasi-convexity.*

Proof. See appendix.

4 Discussion.

Recursivity and Conditional Quasi-Convexity. It is often assumed that preferences over two-stage lotteries are recursive. Formally, a preference relation over two-stage lotteries, \succeq_2 ,

satisfies recursivity if for all pairs of two-stage lotteries of the form $X = [(F_i, q_i)_{i=1}^N]$ and $Y = [F_1, q_1; \dots; F_{j-1}, q_{j-1}; F'_j, q_j; F_{j+1}, q_{j+1}; \dots; F_N, q_N]$ in $\mathcal{L}_0(\mathcal{L}(\mathcal{X}))$, with $q_j > 0$: $X \succeq_2 Y$ if and only if $[F_j, 1] \succeq_2 [F'_j, 1]$.⁹ For example, a property like this is assumed by Kreps & Porteus (1978), Johnsen & Donaldson (1985), LaValle & Wapman (1985), Chew & Epstein (1989), Segal (1990), Epstein (1992), and Sarin & Wakker (1994, 1996). On the other hand, Machina (1989) argues against assuming this degree of consequentialism. Following Machina, our definition of dynamic consistency does not require preferences to be recursive. Moreover, neither part of Proposition 1 requires recursivity. More specifically, conditional quasi-convexity is strictly weaker than recursivity. Thus, given single-action information loving, recursivity is sufficient but not necessary for an agent to be an information lover (even if the agent cannot perform randomizations).

Even if we assumed recursivity, the conditions for information loving shown in Proposition 1 do not imply expected utility. This contrasts with the conclusions of Wakker (1988) and Safra & Sulganik (1995). The reason is that, unlike these authors, we do not assume the reduction axiom. Reduction would imply single-action information neutrality, whereas we only require (weak) single-action information loving.

Although weaker than recursivity, conditional quasi-convexity is still restrictive in that it rules out strict preference for randomization. A referee points out, however, that there is anyway a tension between strict single-action preference for information and strict preference for randomization. The agent who likes information because she is anxious, or simply dislikes living with risk, might be expected strictly to prefer *not* to randomize. It does not follow, however, that an agent who exhibits strict single-action information loving will never strictly prefer to randomize. By definition, single-action information loving describes preference between trees within which there is no possibility of randomization. Thus, we can think of single-action information loving as isolating intrinsic attitudes towards information from other considerations such as procedural fairness. In the absence of the opportunity to randomize, the mother in the Diamond example might prefer more information to less. Perhaps, she is an anxious person. But, once randomizations are available, her concerns for procedural fairness outweigh her concern for anxiety. Thus, while there is a tension between anxiety and randomization in the sense that there is a trade off, there is no contradiction.

Even if we relax conditional quasi-convexity, given single-action information loving, a sophisticated dynamically consistent agent will still ‘almost always’ prefer more informative signals, even if she cannot perform randomizations other than by conditioning on signal realizations. Loosely speaking, whenever a more informative signal offers fewer opportunities to randomize, we can always slightly perturb the signal to increase the number of possible realizations.

⁹ Once again, the $(F_i)_{i=1}^N$, and F'_j are not necessarily distinct.

Preference for information and Dynamic Inconsistency. So far, we have shown conditions under which a sophisticated dynamically consistent agent is an information lover. These conditions included at least weak preference for information in the absence of any choices within the tree: that is, single-action information loving. Suppose, however, that this preference were strict, perhaps because the agent suffers from anxiety. Is it still reasonable to assume that the agent will be dynamically consistent? Consider the following example.

Consider the dilemma of a prisoner, let's call her Jo, who is anxiously awaiting the outcome of her trial for murder. If she is found guilty she will spend the next ten years in jail. The verdict has been delivered, but Judge Oti has sealed it in an envelope overnight. Jo would much rather that the envelope be opened immediately; she has intrinsic preference for information. Meanwhile, the DA is considering offering Jo a plea bargain involving only four years in jail. Suppose that if the DA were to offer Jo a choice between opening the envelope now or taking the plea, Jo would be more or less indifferent. Perhaps, by a small margin, she would choose to rip open the envelope. Suppose instead that the DA offered either the plea or simply waiting for the envelope to be opened in the morning. By transitivity, Jo would want to take the plea. But, now suppose that the plea must be declared in front of Judge Oti in the morning, and Jo cannot commit ahead of time. Come the morning, the choice between the plea or the envelope will look much the same as the choice looked the night before between the plea and immediately opening the envelope, not waiting overnight. Therefore, Jo may refuse the plea come morning. But, anticipating the inconsistency of her future self, Jo is destined for an uneasy night.

In this story, Jo's intrinsic strict preference for information led her to be dynamically inconsistent. Jo is sophisticated in that she realizes her behavior the next morning will differ from the behavior she would choose if she could commit herself tonight. Thus, she would prefer to commit. This suggests that there is a tension between dynamic consistency and intrinsic preference for information. In Kreps & Porteus's (1978) study of intrinsic preference for information, however, agents are dynamically consistent (indeed, they have recursive preferences). The reason is that, in Kreps & Porteus, intrinsic preference for information can be thought of as concern about the time that separates the resolution of uncertainty from the impact of the event itself. For Jo, this is the period between the verdict being announced and either the sentence being carried out or her going free. If the agent is concerned with this time period then dynamic consistency seems a natural assumption. But introspection and our story suggest that intrinsic preference for information should instead be thought of as concern about the time that separates the resolution of uncertainty from the decision maker herself; i.e., from the present. That is, if agents' desire to know is a desire to alleviate anxiety then they are likely to be concerned about the time period over which anxiety will otherwise be endured. In the story, Jo wants to know now because she foresees that the night between now and the verdict tomorrow will otherwise be unpleasant.

If the period between now and resolution is what matters, then dynamic consistency seems less natural. Over time, the period between the present and resolution gets shorter while the period focussed on by Kreps & Porteus, between resolution and punishment/consumption stays the same. As resolution gets closer and the period of anxiety gets shorter, so the future burden of anxiety figures less prominently in the agent's calculation. Past anxiety is a sunk cost. Therefore, the agent will be less willing to make costly choices to relieve anxiety, such as buying insurance, paying for information, or taking precautions. Dynamic consistency excludes such behavior.

This discussion suggests that to offer a better description of how intrinsic concern for information affects behavior, we should drop dynamic consistency. For example, whereas the Kreps-Porteus model allows an agent's degree of risk aversion to depend on the time interval that separates resolution and consumption, we could allow it to depend on the time between resolution and the present.¹⁰ Like models of 'hyperbolic' discounting, this will lead to violations of dynamic consistency.¹¹ The nature of and the reasons for these violations, however, are quite distinct. We believe this would be a fruitful avenue for further research.

Appendix

We first recall the relationship between more informative signals and the distribution of posteriors that they induce. The following Lemma is well-known. The first equivalence is from Blackwell & Girshick (1954) theorem 12.2.2. The last equivalence shows that any increase in the informational content of a signal can be achieved by a sequence of elementary linear bifurcations. For a proof, see Grant, Kajji & Polak (forthcoming).

Lemma A *Suppose that $[(P_i, q_i)_{i=1}^N]$ and $[(P'_j, q'_j)_{j=1}^{N'}]$ are the distributions of posteriors on Ω induced by the signals (S, λ) and (S', λ') respectively. Suppose that both distributions of posteriors have the same prior, that is, $\sum_i q_i P_i = \sum_j q'_j P'_j = \pi$ in $\mathcal{L}_0(\mathcal{X})$. Then the following three statements are equivalent:*

- (i) *The signal (S, λ) is more informative than the signal (S', λ') with respect to the prior belief π*
- (ii) *There exist weights $\beta_{ij} \geq 0$, $i = 1, \dots, N$, $j = 1, \dots, N'$, such that $\sum_i \beta_{ij} = 1$ for all j , $\sum_j \beta_{ij} q'_j = q_i$ for all i , and $\sum_i \beta_{ij} P_i = P'_j$ for all ω in Ω such that $\pi(\omega) > 0$ and all j .*
- (iii) *There exists a sequence of distributions of posteriors $[(P_i^k, q_i^k)_{i=1}^{N^k}]_{k=1}^K$, with $[(P_i^1, q_i^1)_{i=1}^{N^1}] = [(P'_j, q'_j)_{j=1}^{N'}]$ and $[(P_i^K, q_i^K)_{i=1}^{N^K}] = [(P_i, q_i)_{i=1}^N]$, such that $[(P_i^{k+1}, q_i^{k+1})_{i=1}^{N^{k+1}}]$ is a linear bifurcation of $[(P_i^k, q_i^k)_{i=1}^{N^k}]$ for $k = 1, \dots, K - 1$.*

¹⁰ See Cook (1989) for a related idea in the psychology literature.

¹¹ See, for example, Loewenstein & Prelec (1992), Laibson (1997), and O'Donoghue & Rabin (1996).

Proof of Proposition 1: Sufficiency. Fix a prior belief π on Ω , and two signals $((s_1, \dots, s_N), \lambda)$ and $((s'_1, \dots, s'_{N'}), \lambda')$ with the former signal more informative than the latter with respect to π . Let $[(P_i, q_i)_{i=1}^N]$ and $[(P'_j, q'_j)_{j=1}^{N'}]$ denote the distributions of posteriors induced by the signals. By Lemma A(iii) there exists a sequence of distributions of posteriors $([(P_i^k, q_i^k)_{i=1}^{N^k}])_{k=1}^K$, with $[(P_i^1, q_i^1)_{i=1}^{N^1}] = [(P'_j, q'_j)_{j=1}^{N'}]$ and $[(P_i^K, q_i^K)_{i=1}^{N^K}] = [(P_i, q_i)_{i=1}^N]$, such that $[(P_i^{k+1}, q_i^{k+1})_{i=1}^{N^{k+1}}]$ is an elementary linear bifurcation of $[(P_i^k, q_i^k)_{i=1}^{N^k}]$ for $k = 1, \dots, K - 1$. That is, in each step from k to $k + 1$ of the sequence, a probability mass q_j^k on one of the posteriors P_j^k is ‘split’ into two parts, $\beta^k q_j^k$ and $(1 - \beta^k) q_j^k$, and placed on two posteriors P_{j1}^k and P_{j2}^k , for which β^k is in $[0, 1]$ and $\beta^k P_{j1}^k + (1 - \beta^k) P_{j2}^k = P_j^k$. For each k , let $\tau^k = \langle (s_i^k, P_i^k, A, q_i^k)_{i=1}^{N^k}, c \rangle$ be the Blackwell tree associated with the distributions of posteriors $([(P_i^k, q_i^k)_{i=1}^{N^k}])_{k=1}^K$, the consequence function c , and the action set A .

At the first step in the sequence, there are three cases to consider. For ease of notation but without loss of generality, assume that the probability mass on a posterior that is split in the first step is $[P'_1, q'_1]$ and write β for β^1 , so that $\beta P'_{11} + (1 - \beta) P'_{12} = P'_1$. If (a) $P'_{11} \neq P'_j$ and $P'_{12} \neq P'_j$ for all $j = 2, \dots, N'$, then $N^2 = N' + 1$. If (b), without loss of generality, $P'_{11} = P'_2$ but $P'_{12} \neq P'_j$ for all $j = 2, \dots, N'$, then $N^2 = N'$. If (c), without loss of generality, $P'_{11} = P'_2$ and $P'_{12} = P'_3$, then $N^2 = N' - 1$.

Proof of (i). Let m denote a early-resolution randomization over the actions A .¹² Let $(m'_i)_{i=1}^{N'}$ denote the agent’s early-resolution random behavior in the Blackwell tree $\tau' = \langle (s'_i, P'_i, A, q'_i)_{i=1}^{N'}, c \rangle$ induced by the signal $((s'_1, \dots, s'_{N'}), \lambda')$. For any α in $[0, 1]$, and any pair of early-resolution randomizations m and m' , the mixture $[\alpha m + (1 - \alpha) m']$ is a well-defined randomization; that is, it can form part of an early-resolution random behavior. For each of the above cases, we will construct an early-resolution random behavior as follows.

In case (a), $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ can be written as $[P'_{11}, \beta q'_1; P'_{12}, (1 - \beta) q'_1; P'_2, q'_2; \dots; P'_{N'}, q'_{N'}]$. Since $N^2 = N' + 1$, any signal that induces $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ must have at least $N' + 1$ realizations. Thus, the early-resolution random behavior with $N' + 1$ components, $(m'_1, m'_1, m'_2, \dots, m'_{N'})$ is feasible in the tree $\tau^2 = \langle (s_i^2, P_i^2, A, q_i^2)_{i=1}^{N^2}, c \rangle$.

In case (b), $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ can be written as $[P'_2, q'_2 + \beta q'_1; P'_{12}, (1 - \beta) q'_1; P'_3, q'_3; \dots; P'_{N'}, q'_{N'}]$. Since $N^2 = N'$, any signal that induces $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ must have at least N' realizations. Thus, the early-resolution random behavior with N' components, $([\frac{q'_2}{q'_2 + \beta q'_1} m'_2 + \frac{\beta q'_1}{q'_2 + \beta q'_1} m'_1], m'_1, m'_3, \dots, m'_{N'})$ is feasible in the tree $\tau^2 = \langle (s_i^2, P_i^2, A, q_i^2)_{i=1}^{N^2}, c \rangle$.

In case (c), $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ can be written as $[P'_2, q'_2 + \beta q'_1; P'_3, q'_3 + (1 - \beta) q'_1; P'_4, q'_4; \dots; P'_{N'}, q'_{N'}]$. Since $N^2 = N' - 1$, any signal that induces $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ must have at least $N' - 1$ realizations. Thus, the early-resolution random behavior with $N' - 1$ components, $([\frac{q'_2}{q'_2 + \beta q'_1} m'_2 + \frac{\beta q'_1}{q'_2 + \beta q'_1} m'_1],$

¹² The proof is essentially the same if we also allow for late-resolution randomizations. In this case, we would interpret m as an early-resolution randomization over late-resolution randomizations on A .

$[\frac{q'_3}{q'_3+(1-\beta)q'_1}m'_3 + \frac{(1-\beta)q'_1}{q'_3+(1-\beta)q'_1}m'_1], m'_4, \dots, m'_{N'}$) is feasible in the tree $\tau^2 = \langle (s_i^2, P_i^2, A, q_i^2)_{i=1}^{N^2}, c \rangle$.

In each case, the prescribed behavior induces a two-stage lottery that is a elementary linear bifurcation of the two-stage lottery resulting from the behavior $(m'_i)_{i=1}^{N'}$ for the tree τ' . Single-action information loving implies that this bifurcation is preferred. Sophisticated dynamic consistency then implies that $\tau^2 \succeq_{\mathcal{T}} \tau'$. And, by a similar argument, $\tau^{k+1} \succeq_{\mathcal{T}} \tau^k$ for all $k = 1, \dots, K-1$. So by transitivity of $\succeq_{\mathcal{T}}$ we have $\tau = \tau^K \succeq_{\mathcal{T}} \tau'$, where τ is the Blackwell tree induced by the more informative signal $((s_1, \dots, s_N), \lambda)$. \square

Proof of (ii). Let $(a'_i)_{i=1}^{N'}$ in $A^{N'}$ denote the agent's non-random behavior $b_{\tau'}$ in the Blackwell tree τ' and suppose that this behavior induces the two-stage lottery $[M(P'_i, a'_i, c), q'_i]_{i=1}^{N'}$. Again, we will go case by case.

In case (a), $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ can be written as $[P'_{11}, \beta q'_1; P'_{12}, (1-\beta)q'_1; P'_2, q'_2; \dots; P'_{N'}, q'_{N'}]$. Thus, the $N' + 1$ component non-random behavior $(a'_1, a'_1, a'_2, \dots, a'_{N'})$ is feasible in the tree $\tau^2 = \langle (s_i^2, P_i^2, A, q_i^2)_{i=1}^{N^2}, c \rangle$. This behavior induces a two-stage lottery that is a simple linear bifurcation of the two-stage lottery resulting from the behavior $(a'_i)_{i=1}^{N'}$ for the tree τ' . Single-action information loving implies that this bifurcation is preferred. Sophisticated dynamic consistency then implies that $\tau^2 \succeq_{\mathcal{T}} \tau'$.

In case (b), $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ can be written as $[P'_2, q'_2 + \beta q'_1; P'_{12}, (1-\beta)q'_1; P'_3, q'_3; \dots; P'_{N'}, q'_{N'}]$. Thus, both the N' component non-random behavior $(a'_1, a'_1, a'_3, \dots, a'_{N'})$ and the N' component non-random behavior $(a'_2, a'_1, a'_3, \dots, a'_{N'})$ are feasible in the tree $\tau^2 = \langle (s_i^2, P_i^2, A, q_i^2)_{i=1}^{N^2}, c \rangle$. Let X be the two-stage lottery $[M(P'_2, a'_1, c), \beta q'_1 + q_2; M(P'_{12}, a'_1, c), (1-\beta)q'_1; M(P'_3, a'_3, c), q'_3; \dots; M(P'_{N'}, a'_{N'}, c), q'_{N'}]$ induced by the former behavior and let X' be the two-stage lottery $[M(P'_2, a'_2, c), \beta q'_1 + q_2; M(P'_{12}, a'_1, c), (1-\beta)q'_1; M(P'_3, a'_3, c), q'_3; \dots; M(P'_{N'}, a'_{N'}, c), q'_{N'}]$ induced by the latter behavior. Let Y be the two-stage lottery $[M(P'_2, a'_1, c), \beta q'_1; M(P'_{12}, a'_1, c), (1-\beta)q'_1; M(P'_2, a'_2, c), q'_2; \dots; M(P'_{N'}, a'_{N'}, c), q'_{N'}]$. Since Y is an early-resolution mixture of X and X' , by conditional quasi-convexity, at least one of X or X' is (weakly) preferred to Y . Moreover, Y is a elementary linear bifurcation of the two-stage lottery resulting from the behavior $(a'_i)_{i=1}^{N'}$ for the tree τ' . Thus, transitivity, single-action information loving and sophisticated dynamic consistency imply that $\tau^2 \succeq_{\mathcal{T}} \tau'$.

In case (c), $[(P_i^2, q_i^2)_{i=1}^{N^2}]$ can be written as $[P'_2, q'_2 + \beta q'_1; P'_3, q'_3 + (1-\beta)q'_1; P'_4, q'_4; \dots; P'_{N'}, q'_{N'}]$. Thus, all four $N' - 1$ component pure strategies $\mathbf{a}^{13} := (a'_1, a'_3, \dots, a'_{N'})$, $\mathbf{a}^{23} := (a'_2, a'_3, \dots, a'_{N'})$, $\mathbf{a}^{11} := (a'_1, a'_1, \dots, a'_{N'})$ and $\mathbf{a}^{21} := (a'_2, a'_1, \dots, a'_{N'})$ are feasible in the tree $\tau^2 = \langle (s_i^2, P_i^2, A, q_i^2)_{i=1}^{N^2}, c \rangle$. Let X^{jk} be the two-stage lottery $[M(P'_2, a'_j, c), \beta q'_1 + q_2; M(P'_3, a'_k, c), (1-\beta)q'_1 + q'_3; \dots; M(P'_{N'}, a'_{N'}, c), q'_{N'}]$ induced by the behavior \mathbf{a}^{jk} , for $j = 1, 2$ and $k = 1, 3$. Let Y^{12k} be the two-stage lottery $[M(P'_2, a'_1, c), \beta q'_1; M(P'_2, a'_2, c), q'_2; M(P'_3, a'_k, c), (1-\beta)q'_1 + q'_3; \dots; M(P'_{N'}, a'_{N'}, c), q'_{N'}]$, for $k = 1, 3$. But Y^{12k} is an early-resolution mixture of X^{1k} and X^{2k} . Thus by conditional quasi-convexity, at least one of X^{1k} or X^{2k} is (weakly) preferred to Y^{12k} . Let Z be

the two-stage lottery $[M(P'_2, a'_1, c), \beta q'_1; M(P'_2, a'_2, c), q'_2; M(P'_3, a'_1, c), (1-\beta)q'_1; M(P'_3, a'_3, c), q'_3; \dots; M(P'_{N'}, a'_{N'}, c), q'_{N'}]$. But Z is an early-resolution mixture of Y^{121} and Y^{123} . Thus, by conditional quasi-convexity, at least one of Y^{121} or Y^{123} is (weakly) preferred to Z . Hence by transitivity at least one of the two-stage lotteries X^{13} , X^{23} , X^{11} or X^{21} is weakly preferred to Z . Moreover, Z is a simple linear bifurcation of the two-stage lottery resulting from the behavior $(a'_i)_{i=1}^{N'}$ for the tree τ' . Thus, transitivity, single-action information loving and sophisticated dynamic consistency imply that $\tau^2 \succeq_{\mathcal{T}} \tau'$.

By analogous reasoning, $\tau^{k+1} \succeq_{\mathcal{T}} \tau^k$ for all $k=1, \dots, K-1$, and so by sophisticated dynamic consistency and transitivity, we have $\tau \succeq_{\mathcal{T}} \tau'$. \square

Necessity. Given Lemma A, the fact that weak single action information loving is necessary follows immediately from considering Blackwell problems in which the action set is a singleton. To see that conditional quasi-convexity is necessary when early-resolution randomization is not possible, suppose that preferences are not conditionally quasi-convex. Then there exists a pair of two-stage lotteries $X = [(F_i, q_i)_{i=1}^n]$ and $Y = [F_1, q_1; \dots; F_{j-1}, q_{j-1}; F'_j, q_j; F_{j+1}, q_{j+1}; \dots; F_n, q_n]$ where $X \succeq_2 Y$ and a β such that the two-stage lottery $Z := [F_1, q_1; \dots; F_{j-1}, q_{j-1}; F_j, \alpha q_j; F'_j, (1-\alpha)q_j; F_{j+1}, q_{j+1}; \dots; F_n, q_n] \succ_2 X$. For ease of notation, we consider the case where $X = [F_1, q_1; F_2, (1-q_1)]$, $Y = [F'_1, q_1; F_2, (1-q_1)]$, and $Z = [F_1, \hat{\beta}q_1; F'_1, (1-\hat{\beta})q_1; F_2, (1-q_1)]$. The proof for the general case is analogous. By continuity, if there exists one such Z then there exist a continuum of such early-resolution mixtures preferred to both X and Y . That is, there exist positive weights $\alpha_1, \alpha_2, \alpha_3$, and α_4 (where $\sum_i \alpha_i = 1$) such that $Z_{123} = [F_1, (\alpha_1 + \alpha_2 + \alpha_3)q_1; F'_1, \alpha_4 q_1; F_2, (1-q_1)] \succ_2 Z_{12} = [F_1, (\alpha_1 + \alpha_2)q_1; F'_1, (\alpha_3 + \alpha_4)q_1; F_2, (1-q_1)] \succ_2 Z_1 = [F_1, \alpha_1 q_1; F'_1, (\alpha_2 + \alpha_3 + \alpha_4)q_1; F_2, (1-q_1)] \succ_2 X \succeq_2 Y$.

We next construct Blackwell decision trees in which behavior will induce the two-stage lotteries above. Let posteriors P_1, P'_1 and P_2 , actions a_1 and a_2 , and consequence function c be defined such that: $M(\bar{P}, a_1, c) = F_1$ and $M(\bar{P}, a_2, c) = F'_1$ for any $\bar{P} = \gamma P_1 + (1-\gamma)P'_1$ with $\gamma \in [0, 1]$; and such that $M(P_2, a_1, c) = M(P_2, a_2, c) = F_2$. Let \hat{P} be given by $(\alpha_2 + \alpha_3)\hat{P} = \alpha_2 P_1 + \alpha_3 P'_1$. Consider a four-realization signal that induces the distribution of posteriors on Ω : $[P_1, \alpha_1 q_1; \hat{P}, (\alpha_1 + \alpha_2)q_1; P'_1, \alpha_4 q_1; P_2, (1-q_1)]$. By construction (and Lemma A), this signal is more informative than the three-realization signal that induces the distribution of posteriors: $[P_1, (\alpha_1 + \alpha_2)q_1; P'_1, (\alpha_3 + \alpha_4)q_1; P_2, (1-q_1)]$. Suppose that the available action set is $\{a_1, a_2\}$ and the consequence function is c . Then, without randomization the best available two-stage lottery in the Blackwell tree induced from the first signal is Z_{123} ; while that from the second signal is Z_{12} . But since $Z_{123} \succ Z_{12}$, this contradicts preference for the more informative signal. \blacksquare

References

- Blackwell, David (1953): "Equivalent Comparison of Experiments", *Annals of Mathematics and Statistics*, 24, 265-272.
- Blackwell, David and Meyer Girshick (1954): *Theory of Games and Statistical Decisions*, New York: John Wiley & Sons.
- Chew, Soo Hong and Larry Epstein (1989): "The Structure of Preferences and Attitudes Towards the Timing of the Resolution of Uncertainty", *International Economic Review*, 30: 103-117.
- Cook, Victoria (1989): *The Effects of Temporal Uncertainty Resolution on the Overall Utility and Suspense of Risky Monetary and Survival Gambles*, Ph.D. thesis, McGill University.
- Epstein, Larry (1992): "Behavior under risk: recent developments in theory and applications", in Jean-Jacques Laffont (ed.) *Advances in Economic Theory, sixth world congress*, Vol. 2, Cambridge: Cambridge University Press, 1-63.
- Grant, Simon, Atsushi Kajii and Ben Polak (1992): "Many Good Choice Axioms: When Can Many be Treated As One?", *Journal of Economic Theory*, 56, 313-337.
- Grant, Simon, Atsushi Kajii and Ben Polak (forthcoming): "Intrinsic Preference for Information", *Journal of Economic Theory*, forthcoming.
- Johnsen, Thore and John Donaldson (1985): "The Structure of Intertemporal Preferences under Uncertainty and Time Consistent Plans," *Econometrica*, 53, 1451-1458.
- Karni, Edi and David Schmeidler (1991): "Atemporal Dynamic Consistency and Expected Utility Theory", *Journal of Economic Theory*, 54, 401-408.
- Kreps, David and Evan Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory", *Econometrica*, 46, 185-200.
- Laibson, David (1997): "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics*, 112, 443-77.
- LaValle, Irving and Kenneth Wapman (1985): "Rolling Back Decision Trees Requires the Independence Axiom!", *Management Science*, 32, 382-385.
- Loewenstein George & Drazen Prelec (1992): "Anomalies in Intertemporal Choice: Evidence and an Interpretation", *Quarterly Journal of Economics*, 107, 573-597.
- Machina, Mark (1989): "Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty", *Journal of Economic Literature*, 27, 1622-1668.
- McClellenn, Edward (1990): *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge University Press.
- O'Donoghue, Ted, and Matthew Rabin (1996): "Doing it Now or Later", Northwestern University Center for Mathematical Studies in Economics and Management Science Discussion Paper 1172.
- Safra, Zvi and Eyal Sulganik (1995): "On the Nonexistence of Blackwell's Theorem-Type Results with General Preference Relations", *Journal of Risk and Uncertainty*, 10, 187-201.
- Sarin, Rakesh and Peter Wakker (1994): "Folding Back in Decision Tree Analysis", *Management Science*, 40, 625-628.

Sarin, Rakesh and Peter Wakker (1997): “Dynamic Choice and Nonexpected Utility”, working paper, University of Tilburg.

Schlee, Edward (1990): “The Value of Information in Anticipated Utility Theory”, *Journal of Risk and Uncertainty*, 3, 83-92.

Schlee, Edward (1991): “The Value of Perfect Information in Nonlinear Utility Theory”, *Theory and Decision*, 30, 127-131.

Segal, Uzi (1990): “Two-Stage Lotteries without the Reduction Axiom”, *Econometrica*, 58, 349-377.

Wakker, Peter (1988): “Nonexpected Utility as Aversion of Information”, *Journal of Behavioral Decision Making*, 1, 169-175.