

An Evolutionary Perspective on Updating Risk and Ambiguity Preferences*

Philipp Sadowski[†] Todd Sarver[‡]

Current Draft:
June 2, 2021

Abstract

Using an approach predicated on evolution and adaptation, we provide foundations for a model of choice under uncertainty based on *adaptive preferences*. We argue that our approach can be applied in most contexts involving ambiguity, and we show that adaptive preferences nest variants of several established models of ambiguity and risk preferences as special cases. Our model provides a tight connection between ambiguity attitudes and violations of expected utility, and it generates novel predictions about the role of random choice in hedging against ambiguity. We also find that updating of adaptive preferences in response to new information respects dynamic consistency even at the cost of violating consequentialism, addressing a prominent tension in the ambiguity and non-expected-utility literature.

KEYWORDS: Evolution of preferences, ambiguity, updating, dynamic consistency, random choice, phenotypic flexibility

JEL CLASSIFICATION: D81, D83, D84

*[First Version: December 8, 2019](#). This research was supported by the National Science Foundation through Grant SES-1851664. We thank Jiayun Dong, Andrew Ellis, José Heleno Faro, Yoram Halevy, Cosmin Ilut, Frank Johannes, Bart Lipman, George Mailath, Massimo Marinacci, Sujoy Mukerji, Wolfgang Pesendorfer, Arthur Robson, Larry Samuelson, and numerous seminar participants for helpful comments and discussions.

[†]Duke University, Department of Economics, 213 Social Sciences/Box 90097, Durham, NC 27708. Email: p.sadowski@duke.edu.

[‡]Duke University, Department of Economics, 213 Social Sciences/Box 90097, Durham, NC 27708. Email: todd.sarver@duke.edu.

1 Introduction

There are several open questions that are central in the literature on ambiguity aversion. First, it is well known that updating ambiguous beliefs generally leads to violations of either dynamic consistency or consequentialism, which has raised the concern by some that ambiguity aversion may be a “mistake.” And if it is not, then which of these two intuitively appealing properties should be violated? The tension between consequentialism and dynamic consistency in models of ambiguity aversion, which also arises in non-expected-utility models of choice under objective risk, and the subsequent disagreement over which property to give priority are impediments to applying these models in dynamic contexts such as macroeconomics and finance where information plays a central role. Second, a plethora of models of ambiguity aversion and non-expected-utility for risk have been proposed that differ in subtle ways in the behavior they predict. What criteria should be used to select among them? Third, is there a connection between ambiguity aversion (Ellsberg-type behavior) and violations of expected utility in the context of risk (Allais-type behavior)? Fourth, when should individuals randomize over available actions in order to hedge against ambiguity?

This paper provides an evolutionary perspective on these issues, based on the notion that natural selection not only can influence physical traits, but can also shape choice behavior. Using this approach, we develop a foundation for a non-expected-utility and ambiguity-averse model of choice and study updating of this model in response to information. A key finding is that evolutionarily optimal choice must be dynamically consistent, even at the expense of consequentialism, answering the first question posed above.¹ Importantly, the evolutionary approach will provide a novel rationale for violations of consequentialism, showing that such violations should be neither surprising nor concerning.

Systematic violations of expected utility are common, but at least when risk is objective, they appear to be at odds with evolutionary optimality. A central contribution of this paper is to expand the scope of the evolutionary approach by allowing individuals to simultaneously make multiple decisions, some of which are observable and others which are hidden from the modeler. When chosen optimally, as evolution will require, the presence of such hidden actions will generate preferences that appear to violate expected utility from the perspective of the analyst. We show that the resulting class of evolutionarily optimal preferences, which we call *adaptive preferences*, includes rank-dependent expected utility in the context of risk, and variants of the smooth model, variational preferences, and multiple prior preferences in the contexts of both risk and ambiguity. Our result on dynamically consistent updating applies to this rich class of preferences. Importantly, while ambiguity-averse preferences are typically assumed to reduce to expected utility when facing objective risk, our model excludes this benchmark version of many of the ambiguity models it nests and instead closely links different uncertainty attitudes to violations of expected utility. Thus, our evolutionary

¹As we discuss in detail later, consequentialism refers to the requirement that ex post choice not be influenced by outcomes that could have been obtained on some unrealized event.

approach can help to address both the question of model selection and the potential link between Allais and Ellsberg behavior.

Finally, turning to the question of randomization in the face of ambiguity, it is well known that randomization between available options, either in the form of exogenous randomizations offered to individuals or random choice by the individuals themselves, can serve to hedge against ambiguity. Indeed, ambiguity aversion is often defined as a preference for (probabilistic) mixtures of acts. Taking this insight a step further, adaptive preferences generate the novel prediction that individuals may strictly prefer self-randomization over observable actions (i.e., random choice) to exogenous randomization (e.g., mixtures of acts).

The starting point of our analysis is an observation which dates back to a seminal paper by Robson (1996): Evolutionary optimality generates a preference for idiosyncratic uncertainty over common uncertainty, and ambiguity is closely associated with common uncertainty in many instances. Hence, natural selection favors ambiguity aversion. The intuition for why evolution can generate aversion to common uncertainty is actually quite simple. To illustrate, suppose there are two actions between which all individuals must choose in every period. For both actions, individual growth (meaning net expected number of offspring) will be either 2 or 4, each with probability $\frac{1}{2}$. The only difference is that one action bears common uncertainty, where realized per-period growth is perfectly correlated across individuals, while the other bears idiosyncratic uncertainty, where realized growth is independent across individuals. By the law of large numbers, the per-period growth of a (large subpopulation with a common) genotype who consistently chooses the idiosyncratic uncertainty will be approximately $\frac{1}{2}(2+4) = 3$. In contrast, a genotype who chooses the common uncertainty will grow by either 2 or 4, each in approximately half of the periods. Heuristically, this leads to a long-run average growth over two periods of $2 \times 4 = 8$, which is less than $3 \times 3 = 9$. This example illustrates the detrimental effect of correlation on growth: The genotype who chooses the idiosyncratic uncertainty will have a higher long-run growth rate, which implies it will almost surely dominate in the long run (Lemma 1).² We discuss and justify the close connection between ambiguity and correlated uncertainty in detail in Section 1.1.

The main innovations of our paper are the incorporation of adaptation via hidden actions, random choice, and updating following the arrival of information. Importantly, these innovations are not independent of each other: The introduction of adaptation greatly increases the scope of the evolutionary model and allows it to nest versions of a number of prominent models of ambiguity aversion and non-expected utility for risk, which in turn allows our observation concerning dynamic consistency to be applied to a much wider class of models. The presence of hidden actions also generates new predictions about the role of random choice in hedging against ambiguity. We briefly highlight the intuition behind each of these contributions and model predictions below.

²The existence and exact form of this aversion to common uncertainty depend on both the frequency of reproduction (Robatto and Szentes (2017)) and timing of reproduction within the life cycle of organisms (Robson and Samuelson (2019)). We discuss these considerations further in Section 7.1.

The incorporation of hidden actions is motivated not only by economic settings—where data sets often capture only a subset of the many decisions being made by individuals—but also by biological settings—where hidden actions might take the form of unobservable aspects of physical adaptation of organisms. In an economic context, data sets could contain information such as occupation choice, investments, or even vaccination decisions, while omitting information about other complementary decisions such as housing choice, other investments or insurance, or social distancing measures, respectively. In a biological context, hidden actions could take the form of rapid and reversible physical adaptation, known as *phenotypic flexibility*, which has recently gained increased attention in evolutionary biology.³

It is well known that hidden actions can lead to revealed preferences over *observed* choices that violated expected utility, even if the individual’s actual joint preferences over *all* choices satisfy expected utility.⁴ In particular, since different hidden actions may be optimal for different observable actions, individuals may be averse to probabilistic mixtures over observable outcomes (see Sarver (2018)). However, they are not averse to self-randomization since it still enables coordination between observed and unobserved actions. In the context of our model, this implies that self-randomization is a better hedge against ambiguity or common uncertainty than exogenous randomization.

Turning to the updating of ambiguity-averse preferences and non-expected-utility preferences for risk, it is well known that there is a tension between consequentialism and dynamic consistency: Except in very special circumstances, models of ambiguity aversion must violate at least one of these properties (Ghirardato (2002), Hanany and Klibanoff (2007)). As such, there is disagreement in the literature as to how ambiguity-averse preferences should respond to new information: Hanany and Klibanoff (2007, 2009) and Hansen and Sargent (2008) proposed maintaining dynamic consistency but dropping consequentialism; Siniscalchi (2009, 2011) instead suggested keeping consequentialism while abandoning dynamic consistency; Epstein and Schneider (2003) showed that both properties can be maintained for the multiple priors model if one imposes a strong joint restriction (“rectangularity”) on the class of information structures and beliefs; Al-Najjar and Weinstein (2009) took the more extreme position that the conflict between consequentialism and dynamic consistency is so problematic that Ellsberg-type behavior should be recognized as irrational. In an earlier literature on non-expected-utility models of choice under risk, the incompatibility of these two properties was discussed by Hammond (1988, 1989) and Machina (1989). As noted, the lack of consensus in the literature hinders the application of these preferences in dynamic contexts where information plays a central role, such as in macroeconomics and finance.

We leverage the evolutionary perspective that we develop to provide clear guidance on this issue. Since evolutionary optimality of both *ex ante* and *ex post* preferences requires

³We discuss the potential relevance of our results for evolutionary biology in more detail in Section 7.2.

⁴Prior studies of the impact of physical commitments on risk preferences include Grossman and Laroque (1990), Gabaix and Laibson (2001), and Chetty and Szeidl (2007, 2016). Unobservable commitments in particular are explored in Kreps and Porteus (1979), Machina (1984), and Ergin and Sarver (2015).

maximization of the long-run growth rate of the genotype, it follows directly that preferences must be dynamically consistent. Of course, maintaining dynamic consistency necessitates that consequentialism may be violated. Understanding why evolution may dictate these violations is more subtle.

In the context of information and updating, consequentialism means that individuals only consider the outcomes that acts can generate following the actual signal realization, and not what their outcomes would have been following other possible signal realizations. For a single individual acting in isolation, consequentialism seems like a normatively appealing property. However, when a genotype consists of many individuals acting simultaneously, different individuals within this subpopulation may be receiving different signals at the same time. Since correlation in growth rates between members of the genotype plays an important role in its evolutionary success, as already highlighted above, it is in fact quite natural that consequentialism could be violated: For one individual with a given signal realization, considering the outcomes that would be obtained following other signal realizations is not paying undue attention to “what could have been,” but rather giving appropriate consideration to “what others in the population are currently experiencing.”

1.1 Ambiguity as Common Uncertainty

In many examples and applications of ambiguity, the unknown probabilities concern common factors that affect all individuals in the population. For example, in one of the earliest applications of ambiguity to economics, Dow and Werlang (1992) and Epstein and Wang (1994) examined the implications of ambiguity about asset returns. Returns to financial assets are obviously common to all individuals who invest in them. Similarly, in applications to macroeconomics, ambiguity typically concerns aggregate variables, such as factor productivity (Ilut and Schneider (2014), Bianchi, Ilut, and Schneider (2018)). Other examples of uncertainty about aggregate variables that can affect individual outcomes and where probabilities are poorly understood could include the timing of new technological breakthroughs, natural disasters such as earthquakes or tsunamis, or climate change and its implications.

One reason common uncertainty in the examples mentioned so far may be subject to greater ambiguity than idiosyncratic uncertainty is that idiosyncratic random variables can be studied using cross-sectional data, whereas aggregate variables by definition cannot. Greater abundance of data may lead to a better understanding. Nonetheless, there could be common uncertainty for which the probabilities are well understood by individuals, and our results would be equally relevant in those settings.

In addition to ambiguity taking the form of common uncertainty about aggregate variables, there is also a fundamental and systematic link between common uncertainty and any instance of ambiguity involving *model uncertainty*—ambiguity about the true data generating process. Even if the risks faced by each individual are well understood and idiosyncratic

conditional on some common underlying model parameter, if that parameter is unknown and ambiguous then all individuals share in the resulting common uncertainty.⁵ For a simple illustration, consider a medical treatment. If the efficacy (success rate) of the treatment for a population with a given set of characteristics is known, then whether it is successful for one individual is independent of whether it succeeds for another. However, if the treatment has undergone limited testing, then its success rate may be unknown and would itself be a source of common uncertainty for all individuals. In fact, most instances of ambiguity can be cast as common uncertainty about idiosyncratic probabilities.

Of course, we should be careful to point out that the correlation mechanism at play in this paper may not be the only driver of ambiguity aversion. We would not go so far as to claim that every instance of ambiguity corresponds to common uncertainty; nor would we suggest that every instance of common uncertainty involves ambiguous beliefs. Nonetheless, the main thrust of the preceding discussion is that there are indeed many situations in which ambiguity is tightly linked to common uncertainty, and our results speak specifically to these instances of ambiguity. In other cases where ambiguity is not connected to common uncertainty, we remain agnostic about whether ambiguity aversion is driven by heuristics developed by genotypes from the case of common uncertainty or whether some other source of ambiguity aversion is at play.

1.2 Outline

The remainder of the paper is structured as follows. Section 2 formally sets up our model. Section 3 establishes that adaptive preferences are evolutionarily optimal *ex ante* and then analyzes evolutionarily optimal *ex post* preferences following the arrival of a signal. We find that preferences will be dynamically consistent, yet may violate consequentialism. We also show that adaptive preferences may induce random choice.

While the treatment in Section 3 deals with the general case of random choice, the special cases we consider in Sections 4 and 5 allow us to restrict attention to deterministic action plans without loss of generality. These sections explore several applications of our evolutionary model and demonstrate its connection with other established models of ambiguity and risk preferences. In Section 4, we apply our evolutionarily optimal updating rule to the smooth model of ambiguity aversion, which overlaps with the special case of our model with no hidden actions. This special case was previously studied by Robson (1996) for preferences without signals. We discuss how our model predictions align with several recent experimental studies of ambiguity and updating, and we demonstrate how our theoretical results might help to guide experimental design for testing properties of preferences such as dynamic consistency and consequentialism.

⁵This interpretation is closely connected to the macroeconomic literature on robustness to model uncertainty (Hansen and Sargent (2001, 2008)), and is discussed in the evolutionary context in Robson (1996).

In Section 5, we explore an alternative special case involving hidden actions but with no common uncertainty, and we show that the evolutionarily optimal preferences in this case correspond to the optimal risk attitude preferences studied by Sarver (2018). In particular, we show that our model nests rank-dependent expected utility (RDU) as one special case, and we describe the evolutionarily optimal updating of RDU preferences. We find that evolution can generate a version of RDU preferences that is both dynamically consistent and consequentialist, which is perhaps surprising given the tension highlighted above. We discuss how the timing of the selection of the hidden action is critical for determining whether or not RDU preferences will be consequentialist.

Section 6 analyzes other special cases of our model when hidden actions and common uncertainty are simultaneously at play. The main result of that section is a dual formulation of our representation that greatly simplifies the analysis of random choice and the comparison to existing models. We use this result to show that versions of several prominent representations, including variational preferences, multiple priors expected utility, and rank-dependent utility, can be embedded in our general model. Importantly, these special cases provide a link between Ellsberg- and Allais-type behaviors.

In Section 7, we discuss some of the simplifying assumptions that are commonly made in economic applications of the evolutionary approach and the robustness of our results to relaxing them. We also describe the biological evidence of phenotypic flexibility, which provides an alternative interpretation and motivation for the hidden actions in our model. This connection suggests that our model may have relevance not just in economic contexts, but also in the framework of evolutionary biology. Finally, in the Online Appendix, we explore some extensions and variations of our main modeling assumptions, and we provide any proofs omitted from the main paper.

2 Evolutionary Setting

The basic idea behind the evolutionary approach is that a large population of individuals is initially made up of subpopulations with different genotypes, where a genotype specifies the physical traits as well as the programmed behavior (choices) of an organism. These choices lead to a possibly uncertain outcome, and this outcome together with the physical traits of the organism determine its evolutionary fitness, that is, its number of offspring. The offspring inherit the parent's genotype and will face a choice of their own, and so on. In this way, the number of individuals who share a particular genotype may shrink or grow over time, relative to the whole population. A genotype is evolutionarily optimal among those initially present if the relative size of its subpopulation does not vanish over time.

2.1 Uncertainty and Information

Common components of uncertainty are modeled via a state space Ω . The realization of $\omega \in \Omega$ is common to all individuals in the population. In addition, given ω , idiosyncratic uncertainty is captured via a state space S , where each individual in the population receives an independent draw of the state $s \in S$. The entire payoff-relevant state space is then $\Omega \times S$. We model information by allowing each individual to receive a private signal σ from a space of signals Σ that is informative about (ω, s) .⁶ The combined space of signals and states is thus $\Omega \times S \times \Sigma$. We assume that Ω and S are Polish spaces, that is, complete and separable metrizable spaces. We assume that Σ is finite and endowed with the discrete topology. We endow the spaces Ω , S , and Σ with their Borel σ -algebras \mathcal{B}_Ω , \mathcal{B}_S , and \mathcal{B}_Σ , respectively, and we endow the product of these spaces with the product σ -algebra $\mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma$.

Given any measurable space (Y, \mathcal{Y}) , let $\Delta(Y)$ denote the set of countably additive probability measures on Y , and let $\Delta^s(Y)$ denote the set of all simple probability measures on Y (i.e., measures with finite support). The state is drawn each period according to a measure $\mu \in \Delta(\Omega \times S \times \Sigma)$. The marginal distribution of μ on Ω assigns probability $\mu(E)$ to any measurable event $E \in \mathcal{B}_\Omega$. As noted, there is a common draw of the ω dimension of the state for all individuals in the population according to this marginal distribution. However, conditional on ω , both the s dimension of the state and the signal σ are drawn independently for each individual according to the conditional probability distribution $\mu(s, \sigma | \omega)$ on $S \times \Sigma$.⁷ Finally, the informational content of a signal $\sigma \in \Sigma$ is described by conditioning the distribution μ on σ . This information structure is quite general and includes, among other things, the partitional structures that are often used in the literature on ambiguity and updating.

2.2 Consumption and Fitness

Let Z denote a nonempty set of outcomes. Both the ω and s dimensions of the state space are potentially relevant for the outcome of an action, but the role of the signal σ is purely informational. Formally, let \mathcal{F} denote the set of simple acts, that is, the set of all measurable and finite-valued functions $f : \Omega \times S \rightarrow Z$. An *evolutionary fitness function* $\psi : Z \rightarrow \mathbb{R}$ specifies the (net expected) individual reproductive growth associated with each outcome.⁸ Given an act $f \in \mathcal{F}$, the individual growth in state (ω, s) is then $\psi(f(\omega, s))$. For example,

⁶Since S describes idiosyncratic risk, it is natural to consider private signals. In Section S3 of the Online Appendix, we briefly discuss how behavior differs between common and private signals when both are informative only about the common component Ω .

⁷More precisely, since S may be an infinite set, the conditional probability distribution given ω assigns probability $\mu(E|\omega)$ to an event $E \in \mathcal{B}_S \otimes \mathcal{B}_\Sigma$. Note that since $S \times \Sigma$ is a Polish space, the existence of a regular conditional probability distribution is ensured by Proposition 10.2.8 in Dudley (2002).

⁸Realized net individual growth, which includes both survival and offspring, must be an integer, but since reproduction may be uncertain given the outcome $z \in Z$, expected individual growth may take non-integer values. As the main results of Section 3 show, evolutionary fitness of a genotype with a large population depends only on the expected reproductive growth $\psi(z)$ its individuals attain from each outcome z .




Figure 1: Within-period timeline: after-signal adaptation

for a population of individuals, aggregate fitness of zero indicates extinction, fitness of one indicates that the birth rate is equal to the death rate and hence there is no change in the size of the population, and fitness of 1.5 indicates a 50% growth in the population. Aggregate fitness can obviously never be negative. Whether or not individual fitness functions take negative values is not important for our results on the evolutionary optimality of adaptive preferences and on the dynamic consistency of optimal updating. However, in order to derive exact dual characterizations of some special cases of our model, it will be technically useful to allow some outcomes to generate negative individual fitness, which could be interpreted as an externality that eliminates other individuals.

Individuals face the task of choosing acts in each period contingent on the observed signal $\sigma \in \Sigma$, but before learning the realization of the state (ω, s) . Each genotype determines preferences that are used for this choice, contingent on σ . In addition to the observable choice of act f , we assume that individuals might also take hidden actions, that is, actions that are unobservable to the modeler. Incomplete data of this sort is pervasive in economic analysis, as data sets often contain only a snapshot of one dimension of the full spectrum of decisions being made by individuals. We model hidden actions in a simple and tractable reduced form by allowing individuals to select a fitness function ψ from some feasible set Ψ in each period.⁹ As we discuss in Section 7.2, our use of multiple fitness functions can also be interpreted in terms of phenotypic flexibility in the context of evolutionary biology.

Throughout the main text, we assume that the selection of the fitness function $\psi \in \Psi$ takes place after (and in response to) the signal σ , but before the realization of the state. The timing of information and the choice of fitness function and act within each period are illustrated in Figure 1. Implicitly, we are assuming either that adaptation in the form of adjustments to the hidden action can be undertaken rapidly or, equivalently, that signals arrive sufficiently far in advance of the realization of the state to allow time for such adjustments. We briefly discuss the impact of changing the timing of our model so that adaptation must occur before the realization of signals in Section 5.2, and we provide a more in depth analysis in Section S1 of the Online Appendix.

⁹This reduced form derives immediately from a more explicit model of hidden actions, where individuals take a hidden action $y \in Y$ and have a single fixed fitness function $\hat{\psi}(z, y)$ for outcome/action pairs. The resulting set of fitness functions in our model would then be $\Psi = \{\hat{\psi}(\cdot, y) : y \in Y\}$.

We aim to uncover various preferences that can be nested within our evolutionary model, thereby illustrating both the structure imposed by our model on static uncertainty preferences and the scope of our updating results. We therefore impose only minimal technical restrictions on the set of fitness functions: We assume Ψ is nonempty and that $\sup_{\psi \in \Psi} \psi(z)$ is finite for every $z \in Z$. Of course, additional structure and restrictions on the set Ψ may be appropriate depending on the application, as the availability of various hidden actions and their impact on fitness will naturally depend on the choice context, and such restrictions will serve to refine the exact preferences under uncertainty generated by our model. However, we leave exploring such refinements as a topic for future research.

2.3 Growth Rates

In a given time period, the aggregate growth rate of a genotype will be determined by the common preferences each individual in its subpopulation is programmed to use when choosing (deterministically or possibly randomly) an act f and a fitness function ψ . We assume each decision problem is faced repeatedly, leading to a stochastic sequence of aggregate growth rates for each genotype. Our analysis of natural selection and evolutionary optimality will center around the comparison of long-run growth rates of different genotypes (with different programmed preferences).

Definition 1. Suppose the aggregate growth rate of a genotype is given by $(\lambda_t)_{t \in \mathbb{N}}$, where λ_t is the random variable that describes the aggregate growth rate in period t of the entire subpopulation of individuals with that genotype. We say that α is the *long-run growth rate* of the genotype if $\frac{1}{T} \sum_{t=1}^T \ln(\lambda_t) \rightarrow \alpha$ almost surely as $T \rightarrow \infty$.

For an arbitrary sequence $(\lambda_t)_{t \in \mathbb{N}}$ of random variables, the long-run growth rate may not exist, since the series above may not converge. However, we will see in the next section that in our model, the long-run growth rate exists for any act f and fitness function ψ .

To establish that the long-run growth rate is the appropriate statistic for comparison in our evolutionary model, the next lemma demonstrates how it relates to long-run dominance of a particular genotype over others. First, note that throughout the paper, we follow the standard convention of assuming that the number of agents of each genotype is (infinitely) large, which we formally model by treating the set of individuals of each genotype i as a continuum with measure $N^i(t)$ at time period t .¹⁰ Thus, if the sequence of aggregate growth rates of genotype i is $(\lambda_t^i)_{t \in \mathbb{N}}$ and the initial measure of this genotype is $N^i(0)$, then the measure of its subpopulation at time $T \in \mathbb{N}$ is

$$N^i(T) = N^i(0) \prod_{t=1}^T \lambda_t^i.$$

¹⁰Using results from the theory of branching processes, it can be shown that our results involving continuum populations are the correct limiting approximations for large but finite populations.

Lemma 1. Consider two genotypes $i = A, B$, where each genotype i has a sequence of stochastic aggregate growth rates $(\lambda_t^i)_{t \in \mathbb{N}}$ that converges to a long-run growth rate α^i . If $\alpha^A > \alpha^B$, then regardless of the initial measures $N^A(0) > 0$ and $N^B(0) > 0$ of their respective subpopulations at time $t = 0$, we have $N^A(t)/N^B(t) \rightarrow \infty$ almost surely as $t \rightarrow \infty$.

Note that Lemma 1 does *not* imply that a higher long-run growth rate yields higher expected population size as t grows large, as indeed it is possible to have the expected value of $N^B(t)$ exceed that of $N^A(t)$ for all t . Nonetheless, the lemma implies that the event where $N^B(t)$ exceeds $N^A(t)$ vanishes (has probability zero) in the limit as $t \rightarrow \infty$.

Evolutionary theory aims to explain which genotypes can be observed in the long run. Lemma 1 clarifies why maximizing long-run growth, rather than the expected population size, is evolutionarily optimal. If in the present moment organisms have already been evolving for t periods, then the relative population sizes of different genotypes that we observe today is a snapshot of the evolutionary process in period t . Assuming this process has been underway for some time (t is large), the probability is very high that the dominant genotype observed today is precisely the one with the highest long-run growth rate.

3 Evolutionarily Optimal Choice

Since the signal arrives prior to the choice of act and fitness function (see Figure 1), the individual can take it into account when selecting both. To analyze dynamic choice in general—and dynamic consistency in particular—it is necessary to compare ex post behavior after the arrival of information with the ex ante plan that would be formed if the individual committed to signal-contingent choices prior to the realization of the signal. We therefore begin our analysis by deriving the long-run growth rates associated with (possibly random) ex ante *plans* of action and fitness function selection. Since evolutionary optimality requires maximizing long-run growth, the optimal value function over random action plans follows immediately. In Section 3.1, we discuss the role that self-randomization plays in these formulas. In Section 3.2, we consider evolutionarily optimal ex post behavior and establish that choice is dynamically consistent.

Definition 2. A *random plan* is a function $\pi \in \mathcal{R}(\mathcal{F}, \Psi) \equiv (\Delta^s(\mathcal{F} \times \Psi))^\Sigma$ from the space of signals to the set of simple probability measures over the space of acts and feasible fitness functions. The probability π assigns to (f, ψ) following signal σ is denoted by $\pi_\sigma(f, \psi)$.

A random plan π specifies a path through a decision tree, where the randomization π_σ is selected following the signal $\sigma \in \Sigma$. We denote the special case of a deterministic plan that selects the pair (f_σ, ψ_σ) with certainty following σ by $(f_\sigma, \psi_\sigma)_{\sigma \in \Sigma}$. For a given signal and state realization (ω, s, σ) , such a deterministic plan achieves a fitness of $\psi_\sigma(f_\sigma(\omega, s))$ and,

more generally, a random plan π achieves an expected fitness of

$$\mathbb{E}_{\pi_\sigma}[\psi(f(\omega, s))] = \int_{\mathcal{F} \times \Psi} \psi(f(\omega, s)) d\pi_\sigma(f, \psi).$$

We adopt the convention that the domain of the natural logarithm includes nonpositive numbers and its range is the extended reals by setting $\ln(x) = -\infty$ for all $x \leq 0$.

Theorem 1 (Ex Ante Long-Run Growth). *Suppose Ψ and μ are fixed, and consider a genotype with an (infinitely) large subpopulation of individuals. The long-run growth rate of the genotype from choosing the random plan $\pi \in \mathcal{R}(\mathcal{F}, \Psi)$ in every period is*

$$\Lambda(\pi) = \int_{\Omega} \ln \left(\int_{S \times \Sigma} \mathbb{E}_{\pi_\sigma}[\psi(f(\omega, s))] d\mu(s, \sigma|\omega) \right) d\mu(\omega). \quad (1)$$

The concavity of the logarithm implies that Λ is more adversely affected by common uncertainty about ω than by idiosyncratic uncertainty about s . Also, since Λ expresses the long-run average growth rate in log terms, $\Lambda(\pi) = -\infty$ corresponds to extinction and $\Lambda(\pi) = 0$ corresponds to constant population size. At the heart of the proof of Theorem 1 is the same logic that is behind the seminal result of Robson (1996), who considered the special case of no signals ($\Sigma = \{\sigma\}$), no adaptation ($\Psi = \{\psi\}$), and no self-randomization.

Proof. Recall that, conditional on ω , both the s dimension of the state and the signal σ are independently distributed for each individual in the population. Self-randomization is also idiosyncratic. Therefore, by the law of large numbers, conditional on the realized ω_t at time t , the aggregate growth rate of a large population of individuals choosing a particular random plan π is approximately

$$\lambda_t(\omega_t) = \int_{S \times \Sigma} \mathbb{E}_{\pi_\sigma}[\psi(f(\omega_t, s))] d\mu(s, \sigma|\omega_t).$$

Since we consider infinite subpopulations in our model, we can treat this approximation as exact.¹¹ Taking the product over a sequence of realized common components $\omega_1, \dots, \omega_T$ and raising to the power $1/T$ gives the realized annualized growth rate over this sequence of periods:

$$\prod_{t=1}^T \left(\int_{S \times \Sigma} \mathbb{E}_{\pi_\sigma}[\psi(f(\omega_t, s))] d\mu(s, \sigma|\omega_t) \right)^{1/T}.$$

¹¹Note that an approximate (limiting) version of Theorem 1 also holds for finite populations, provided the initial population size is sufficiently large. Using the theory of branching processes (Athreya and Ney (1972, Chapter 5)), it can be shown that the average growth rate of a finite population converges to $\Lambda(\pi)$ conditional on non-extinction. Moreover, it can be shown that when $\Lambda(\pi) > 0$, the probability of extinction converges to zero as the initial population becomes large.

Taking the logarithm of this expression and then the limit as $T \rightarrow \infty$, we have

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \ln \left(\int_{S \times \Sigma} \mathbb{E}_{\pi_\sigma} [\psi(f(\omega_t, s))] d\mu(s, \sigma | \omega_t) \right) \\ \rightarrow \int_{\Omega} \ln \left(\int_{S \times \Sigma} \mathbb{E}_{\pi_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma | \omega) \right) d\mu(\omega) \quad \text{a.s.}, \end{aligned}$$

by the law of large numbers. ■

The long-run growth rate of the population is optimized if individuals choose π to maximize Equation (1). However, since only the random choice of act is observed while the choice of fitness function corresponds to some unobservable action, it will be useful to decompose π_σ into its (observable) marginal distribution over acts and (unobservable) conditional distribution over fitness functions given the act.

Definition 3. An *action plan* is a function $\rho \in \mathcal{R}(\mathcal{F}) \equiv (\Delta^s(\mathcal{F}))^\Sigma$ from the space of signals to the set of simple probability measures over acts, where $\rho_\sigma(f)$ is the probability assigned to f following signal σ . An *adaptation plan* is a function $\tau \in \mathcal{R}(\Psi | \mathcal{F}) \equiv (\Delta^s(\Psi))^{\Sigma \times \mathcal{F}}$ from the space of signals and acts to the set of simple probability measures over the feasible fitness functions, where $\tau_\sigma(\psi | f)$ is the probability assigned to fitness function ψ following signal σ and the observable choice of act f .

The choice of random plan π can equivalently be expressed as the choice of ρ and τ . Formally, let $\tau_\sigma \otimes \rho_\sigma$ denote the measure with marginal distribution ρ_σ on \mathcal{F} and conditional distribution $\tau_\sigma(\cdot | f)$ on Ψ . Then, the expectation of $\psi(f(\omega, s))$ with respect to this measure is

$$\mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] = \int_{\mathcal{F}} \int_{\Psi} \psi(f(\omega, s)) d\tau_\sigma(\psi | f) d\rho_\sigma(f).$$

Given an action plan ρ and adaptation plan τ , the corresponding joint plan over both actions and adaptation is $\pi = \tau \otimes \rho \equiv (\tau_\sigma \otimes \rho_\sigma)_{\sigma \in \Sigma}$. Therefore, the highest possible long-run growth rate associated with an action plan ρ (and subsequent optimal choice of adaptation plan) is

$$\begin{aligned} V(\rho) &= \sup_{\tau \in \mathcal{R}(\Psi | \mathcal{F})} \Lambda(\tau \otimes \rho) \\ &= \sup_{\tau \in \mathcal{R}(\Psi | \mathcal{F})} \int_{\Omega} \ln \left(\int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma | \omega) \right) d\mu(\omega). \end{aligned} \tag{2}$$

Robson (1996) considered the special case with a single fitness function ψ , and without signals or random choice, in which case the long-run growth rate associated with the deterministic choice of act f reduces to

$$V(f) = \int_{\Omega} \ln \left(\int_S \psi(f(\omega, s)) d\mu(s | \omega) \right) d\mu(\omega). \tag{3}$$

In an application to status and relative consumption effects, Nöldeke and Samuelson (2005) considered a formula similar to Equation (3) that also incorporates signals, but without adaptation ($\Psi = \{\psi\}$) and without random choice.¹² The survey by Robson and Samuelson (2011) summarizes these and other recent developments in the literature on evolution of preferences.

By Lemma 1, the evolutionarily optimal genotype is the one that maximizes the long-run growth rate; that is, it selects among ex ante action plans to maximize Equation (2). We refer to the preferences over action plans represented by this function V as *adaptive preferences*. As is usual in random choice contexts, we do not directly observe these preferences, only the implied random choice rule. Formally, a *decision problem* $A = (A_\sigma)_{\sigma \in \Sigma}$ specifies a nonempty and finite set of available acts A_σ following each signal σ . The resulting set of feasible action plans is

$$\mathcal{R}(A) \equiv \{\rho \in \mathcal{R}(\mathcal{F}) : \text{supp}(\rho_\sigma) \subset A_\sigma, \forall \sigma \in \Sigma\}.$$

Corollary 1 (Ex Ante Choice). *Suppose Ψ and μ are fixed. Then, for every infinitely repeated decision problem A , the genotype that chooses an action plan in $\text{argmax}_{\rho \in \mathcal{R}(A)} V(\rho)$ achieves a weakly higher long-run growth rate than all others.*

The adaptive preferences represented by Equation (2) specify the optimal response to correlated and uncorrelated uncertainty, but do not concern ambiguity per se. However, as laid out in Section 1.1, in many examples and applications of ambiguity, the unknown probability concerns a common factor that affects all individuals in the population. Thus, the evolutionary mechanism described in Theorem 1 may capture one important source of ambiguity aversion. In particular, the Robson (1996) representation in Equation (3), which applies to the case without information, is a special case of the issue-preference model studied by Nau (2006) and Ergin and Gul (2009), and it is a special case of the smooth model of Klibanoff, Marinacci, and Mukerji (2005) when restricted to acts f that depend only on s . We discuss this special case and its extension to signals and updating in detail in Section 4.

3.1 The Role of Random Choice

In the literature on evolutionary biology, several studies have highlighted the potential benefits of randomization in behavior or in the assignment of physical characteristics to organisms (e.g., Cooper and Kaplan (1982), Bergstrom (2014)). In the context of ambiguity, the discussion of random choice and its role as a hedging device dates back to Raiffa (1961), and it has been explored axiomatically more recently by Saito (2015) and Ke and Zhang (2020).¹³ One

¹²Given the specific assumptions of their model, self-randomization is not necessary for evolutionary optimality. We discuss other specific instances where deterministic plans are optimal in Sections 4 and 5. See also Section S2 of the Online Appendix for a related discussion of signal response in lieu of self-randomization.

¹³For a related discussion of random choice induced by quasiconcave non-expected-utility preferences for risk, see Machina (1985) and Cerreia-Vioglio, Dillenberger, Ortoreva, and Riella (2018).

significant prediction of our model that has not received attention in this prior literature is the possibility of strict preferences to self-randomize rather than use some exogenous source of randomization such as mixtures of acts.¹⁴

The following example illustrates the role of random choice in our model and how the presence of hidden actions causes self-randomization to become a better hedge against ambiguity than exogenous sources of randomization. The key conceptual point in this example is that self-randomization over actions by individuals allows them to coordinate between observed actions (choice of act) and unobserved actions (adaptation). In contrast, exogenous randomization that takes place after the choice of hidden action hinders such coordination. Moreover, in this case, whether the exogenous randomization over acts takes place before or after the realization of the state is not important for payoffs, so we can conveniently model exogenous sources of randomization using probabilistic mixtures of acts (in the sense of Anscombe–Aumann).

Example 1. *The farmers in a community have to choose between planting one of two crops, f and g . There are two common states of the world, a rainy state and a dry state, $\Omega = \{r, d\}$, and $\mu(r) = \mu(d) = 0.5$. Crop f produces high yield \bar{f} in state r and low yield \underline{f} in d . Crop g instead produces high yield \bar{g} in d and low yield \underline{g} in r . Suppose that there are two harvesting technologies, and each farmer makes an unobserved (hidden) investment in harvesting equipment at the time of choosing a crop and before the state resolves. Denote the two resulting fitness functions by ψ_1 and ψ_2 . The first technology is suited to crop f : $\psi_1(\bar{f}) = 2$ and $\psi_1(z) = 1$ for $z \neq \bar{f}$. The second technology is suited to crop g : $\psi_2(\bar{g}) = 2$ and $\psi_2(z) = 1$ for $z \neq \bar{g}$. The individual reproductive fitness associated with each crop and technology combination in the two states $\omega = r, d$ is summarized in Table 1.*

	ψ_1	ψ_2
f	2, 1	1, 1
g	1, 1	1, 2

Table 1: Individual growth in states (r, d)

The signal space Σ and the idiosyncratic state space S play no role in this example and can be dropped from the long-run growth formula in Equation (2). We now consider deterministic action plans, action plans involving self-randomization, and exogenous randomization over acts, and we compare the long-run growth rates associated with each.

¹⁴Indeed, the aforementioned axiomatic models invoke the opposite preference. Empirically, the question of whether self-randomization or mixtures serve as a better hedge against ambiguity has received little attention in the experimental literature, but there is some evidence that subjects have limited or even negative willingness to pay for mixtures of acts (see Dominiak and Schmedler (2011) and Agranov and Ortoleva (2017)), indicating a potential preference for self-randomization.

- **Deterministic choice:** All farmers planting the same crop f or g exposes the population to common uncertainty and leads to the following long-run growth rates:

$$V(f) = \Lambda(f, \psi_1) = \int_{\Omega} \ln(\psi_1(f(\omega))) d\mu(\omega) = \frac{1}{2} \ln(2) + \frac{1}{2} \ln(1) \approx 0.3466$$

$$V(g) = \Lambda(g, \psi_2) = \int_{\Omega} \ln(\psi_2(g(\omega))) d\mu(\omega) = \frac{1}{2} \ln(1) + \frac{1}{2} \ln(2) \approx 0.3466.$$

- **Self-randomization:** To analyze random choice, we first determine the optimal joint plan π over crops and technology, and then deduce from it the optimal (observable) action plan ρ over crops. Let $\pi = \frac{1}{2}\delta_{(f, \psi_1)} + \frac{1}{2}\delta_{(g, \psi_2)}$ be the plan that randomizes uniformly over (f, ψ_1) and (g, ψ_2) , pairing each crop with the appropriate harvest technology.¹⁵ From a quick examination of Table 1, we see that this equal weight randomization eliminates common uncertainty and gives an average individual fitness of 1.5 in each state. The resulting action plan $\rho = \frac{1}{2}\delta_f + \frac{1}{2}\delta_g$ is the marginal distribution of π , and

$$V\left(\frac{1}{2}\delta_f + \frac{1}{2}\delta_g\right) = \Lambda\left(\frac{1}{2}\delta_{(f, \psi_1)} + \frac{1}{2}\delta_{(g, \psi_2)}\right) = \ln(1.5) \approx 0.4055.$$

- **Exogenous randomization (mixtures):** As noted above, we can represent any exogenous randomization between the two crops that is not carried out well in advance as the probabilistic mixture (in the sense of Anscombe-Aumann) of the two acts, $h = \frac{1}{2}f + \frac{1}{2}g$. In this case, the farmer must choose (perhaps randomly) which harvesting equipment to acquire prior to learning the realized crop. This choice of adaptation plan τ can be represented as choosing the probability $\alpha \in [0, 1]$ of selecting ψ_1 . Averaging the individual fitness from f and g in Table 1, we see that the state-contingent average fitness is $(1.5, 1)$ and $(1, 1.5)$ for the fitness functions ψ_1 and ψ_2 , respectively. Therefore,

$$\Lambda(\alpha\delta_{(h, \psi_1)} + (1 - \alpha)\delta_{(h, \psi_2)}) = \frac{1}{2} \ln(1 + 0.5\alpha) + \frac{1}{2} \ln(1.5 - 0.5\alpha).$$

This long-run growth rate is maximized by taking $\alpha = 0.5$, and hence

$$V\left(\frac{1}{2}f + \frac{1}{2}g\right) = \Lambda\left(\frac{1}{2}\delta_{(h, \psi_1)} + \frac{1}{2}\delta_{(h, \psi_2)}\right) = \ln(1.25) \approx 0.2231.$$

This example shows that when there are hidden actions, self-randomization serves as a better hedge against common uncertainty or ambiguity (we will explore the ambiguity interpretation of the representation in more detail in Section 4) than mixtures of acts. Indeed, in this example, mixtures of acts perform even worse than the original acts since the hedging benefit of the mixture is outweighed by the loss of fitness associated with sometimes mis-coordinating the hidden action (harvesting equipment) with the realized act (crop).

¹⁵As is standard, we use δ_x to denote the Dirac probability measure that assigns probability one to x .

3.2 Updating and Dynamic Consistency

We begin this section with definitions of consequentialism and dynamic consistency. Since it is impossible to directly observe preferences over random action plans,¹⁶ we focus on the choice of action plan given a decision problem $A = (A_\sigma)_{\sigma \in \Sigma}$. As a natural extension of the standard choice correspondence used for deterministic choice, let $\mathcal{C}(A)$ denote the set of all random action plans ρ that the individual is willing to choose ex ante. Similarly, let $\mathcal{C}(A_\sigma|\sigma, \rho)$ denote the set of all random actions $\hat{\rho}_\sigma \in \Delta^s(A_\sigma)$ that the individual is willing to choose after observing the signal σ and given the ex ante plan ρ .

Definition 4. Random choices satisfy *consequentialism* if ex post choices do not depend on ex ante plans: $\mathcal{C}(A_\sigma|\sigma, \rho) = \mathcal{C}(A_\sigma|\sigma, \hat{\rho})$ for all A and all $\rho, \hat{\rho} \in \mathcal{R}(A)$ and $\sigma \in S$.

Definition 5. Random choices satisfy *dynamic consistency* if ex ante plans are carried out ex post: For any A and $\rho, \hat{\rho} \in \mathcal{R}(A)$ such that $\rho_{\sigma'} = \hat{\rho}_{\sigma'}$ for all $\sigma' \neq \sigma$,

1. $\rho \in \mathcal{C}(A)$ and $\hat{\rho} \notin \mathcal{C}(A)$ together imply $\rho_\sigma \in \mathcal{C}(A_\sigma|\sigma, \rho)$ and $\hat{\rho}_\sigma \notin \mathcal{C}(A_\sigma|\sigma, \rho)$.
2. $\rho \in \mathcal{C}(A)$ implies $\rho_\sigma \in \mathcal{C}(A_\sigma|\sigma, \rho)$ whenever $\mu(\sigma) > 0$.

These conditions extend the standard definitions used in the special case of partitional information structures and deterministic choice (e.g., Machina and Schmeidler (1992), Epstein and Le Breton (1993), or Hanany and Klibanoff (2007)) to our more general framework. Specifically, in the case of deterministic choice, the random action plan ρ reduces to a deterministic plan $(f_\sigma)_{\sigma \in \Sigma}$. Partitional learning corresponds to the special case where Σ is a partition of S , so each signal σ is a subset of S and, conditional on the signal σ , the measure μ assigns probability zero to states outside of the event σ . In this case, a deterministic action plan can be reduced to an act by defining $f(s) = f_\sigma(s)$ for $s \in \sigma \in \Sigma$. Finally, when $f, g \in A$, $f \in \mathcal{C}(A)$ and $g \notin \mathcal{C}(A)$ means that $f \succsim g$, and similarly for ex post preferences $\succsim_{\sigma, f}$. Our version of dynamic consistency then implies the standard definition: If $f(s) = g(s)$ for all $s \notin \sigma$,

$$f \succ g \implies f \succ_{\sigma, f} g \quad \text{and} \quad f \succsim g \implies f \succsim_{\sigma, f} g \text{ whenever } \mu(\sigma) > 0.$$

As noted above, ambiguity-aversion and violations of expected-utility in general imply that consequentialism and dynamic consistency cannot be simultaneously satisfied, as we will

¹⁶There are several reasons for this: First, we already observed in the last section that exogenous randomizations offered to an individual are not treated the same as self-randomization, so the individual's ranking of a pair of exogenous randomizations over acts may not reflect her ranking of self-randomizations with the same distributions. Second, whatever options are made available to an individual, be they acts or lotteries (mixtures) over acts, the individual is always able to randomize between them, making it impossible to directly infer the ranking of the options provided. That is, given the option set $\{\rho, \hat{\rho}\}$, the individual may instead prefer to self-randomize to obtain another distribution over acts $\alpha\rho + (1 - \alpha)\hat{\rho}$. With the exception of Saito (2015), the decision theory literature on randomization and ambiguity largely ignores this fundamental issue with the observability of preferences over random choices.

illustrate in Section 4. Fortunately, the evolutionary approach gives clear guidance about which property to favor: The evolutionarily optimal ex ante plans are precisely those that maximize the long-run growth rate of the genotype. The evolutionarily optimal ex post plans are those that achieve the same objective. Thus, dynamic consistency is necessarily satisfied, as the following results demonstrate.

Theorem 2 (Ex Post Long-Run Growth). *Suppose Ψ and μ are fixed, and suppose the genotype forms an action plan $\rho \in \mathcal{R}(\mathcal{F})$ ex ante, which it follows after every signal $\sigma' \neq \sigma$, but it deviates from this plan after signal σ by instead implementing the ex post random action $\hat{\rho}_\sigma \in \Delta^s(\mathcal{F})$. Then, its long-run growth rate is*

$$V(\hat{\rho}_\sigma | \sigma, \rho) = \sup_{\tau \in \mathcal{R}(\Psi | \mathcal{F})} \int_{\Omega} \ln \left(\mu(\sigma | \omega) \int_S \mathbb{E}_{\tau_{\sigma} \otimes \hat{\rho}_\sigma} [\psi(f(\omega, s))] d\mu(s | \sigma, \omega) + \int_{S \times \Sigma \setminus \{\sigma\}} \mathbb{E}_{\tau_{\sigma'} \otimes \rho_{\sigma'}} [\psi(f(\omega, s))] d\mu(s, \sigma' | \omega) \right) d\mu(\omega). \quad (4)$$

Theorem 2 follows directly from Theorem 1. The growth rate formula in Equation (4) simply evaluates $\hat{\rho}_\sigma$ (following σ) in conjunction with the ex ante plan ρ (following other signals) according to Equation (2). It follows that choices are dynamically consistent.

Corollary 2 (Dynamic Consistency). *Given a decision problem $A = (A_\sigma)_{\sigma \in \Sigma}$, the long-run growth rate is optimized if individuals maximize Equation (2) ex ante and Equation (4) ex post, so $\mathcal{C}(A) = \operatorname{argmax}_{\rho \in \mathcal{R}(A)} V(\rho)$ and $\mathcal{C}(A_\sigma | \sigma, \rho) = \operatorname{argmax}_{\hat{\rho}_\sigma \in \Delta^s(A_\sigma)} V(\hat{\rho}_\sigma | \sigma, \rho)$. Thus, evolutionarily optimal random choice is dynamically consistent.*

Given the tension between dynamic consistency and consequentialism, one implication of Corollary 2 is that choice may violate consequentialism. In Section 4, we discuss such a violation in the context of the Ellsberg example. The evolutionary approach provides a natural interpretation for why consequentialism may be violated. For expositional clarity, consider deterministic plans. Consequentialism states that preferences between acts f and g following a signal σ do not depend on what act would have been obtained following other signals σ' . This property could therefore be interpreted as preferences not depending on “what might have been.” In our model, the genotype consists of a large subpopulation of individuals. Even if one individual receives the signal σ , other members of this subpopulation are simultaneously receiving different signals. From the individual perspective, choice after updating can be thought of as the best response to other individuals who are all playing the Pareto optimal equilibrium of the game that has long-run population growth as the payoff. In other words, individuals may violate consequentialism because they care about the outcomes others of their genotype are experiencing; in particular, they care about the correlation between their own fitness and the fitness of others with the same genotype.

Hanany and Klibanoff (2007, 2009) similarly studied dynamically consistent (and hence non-consequentialist) conditional preferences. In particular, they showed that for a variety

of models of ambiguity aversion, such conditional preferences between acts f and g can be represented using updated beliefs within an otherwise unchanged value function. Crucially, since conditional preferences may violate consequentialism, their updating rule for beliefs is typically not Bayesian and depends nontrivially on the original choice set and the ex ante plan. Therefore, their approach necessarily conflates beliefs and tastes, since the updated beliefs depend not just on the information structure, but on the decision problem itself. In contrast, updated beliefs in Equation (4) are derived using standard Bayesian updating and hence are independent of the decision problem. The violation of consequentialism in this expression comes instead from an externality—in the sense that each individual is programmed to care about correlation with other individuals—which requires the ex ante plan to be a part of the ex post value function. In the context of our evolutionary model, this strikes us as the most natural formulation of the conditional growth rate, as it emphasizes the underlying reason for the dependence of the optimal ex post choices on the ex ante plan.

4 Applications: Ambiguity Aversion

In this section, we focus on the special case of a single fitness function ($\Psi = \{\psi\}$), which allows the supremum over Ψ to be dropped from the representation in Equation (2). For expositional ease, we also focus on deterministic choice. While in general self-randomization over available acts may serve as a hedging device even absent hidden actions, it will be easy to see that in the examples of this section deterministic choice is optimal. We begin with a simple example without signals and then incorporate signals in Section 4.1.

Example 2 (Ellsberg—no signals). *Consider an Ellsberg urn with one black ball and two balls that could each be either red or yellow. Each individual independently draws one ball from the urn, which we model using the state space $S = \{b, r, y\}$ for independent risk. The individual may be offered the following bets on colors of the ball drawn:*

	b	r	y
B	1	0	0
R	0	1	0
BY	1	0	1
RY	0	1	1

In this table, B denotes the act that pays \$1 if the ball drawn is black and \$0 otherwise, BY indicates the act that pays \$1 if the ball is either black or yellow, and so on. The typical preference pattern documented by Ellsberg (1961) is $B \succ R$ and $BY \prec RY$, in violation of Savage’s sure-thing principle.

To understand such preferences within the evolutionary model described above, note that although the draw of the ball is independent across individuals, the composition of the urn

itself may be common for all individuals. In this case, we can model the possible urn compositions using the set $\Omega = \{\omega_1, \omega_2, \omega_3\}$, where $\omega_1 = (b, r, r)$, $\omega_2 = (b, r, y)$, and $\omega_3 = (b, y, y)$. Even if individuals form subjective probability assessments on the possible urn compositions, this correlated uncertainty is treated differently than uncorrelated uncertainty. For ease of illustration, suppose that μ assigns equal weight to each urn composition and that there is a single fitness function ψ that takes values $\psi(0) = 0$ and $\psi(1) = 1$. Since when acts only depend on s the long run growth rate for deterministic choice in Equation (3) is a special case of the smooth model (Klibanoff, Marinacci, and Mukerji (2005)) with a concave transformation function, these evolutionarily optimal preferences exhibit Ellsberg behavior:

$$V(B) = \ln \left[\frac{1}{3} \right] > \frac{1}{3} \ln \left[\frac{2}{3} \right] + \frac{1}{3} \ln \left[\frac{1}{3} \right] + \frac{1}{3} \ln[0] = V(R),$$

and

$$V(BY) = \frac{1}{3} \ln \left[\frac{1}{3} \right] + \frac{1}{3} \ln \left[\frac{2}{3} \right] + \frac{1}{3} \ln[1] < \ln \left[\frac{2}{3} \right] = V(RY).$$

Simple calculations show that randomizations $\alpha\delta_B + (1 - \alpha)\delta_R$ for $0 < \alpha < 1$ yield long-run growth rates strictly between those of B and R , and likewise for the second choice scenario. Thus, deterministic choice is indeed optimal and we have $\mathcal{C}(\{B, R\}) = \{B\}$ and $\mathcal{C}(\{BY, RY\}) = \{RY\}$.¹⁷

In Example 2, the crucial assumption for generating ambiguity aversion is that the composition of the urn is common across all individuals. In contrast, if a different urn is composed for each individual and if there is no correlation in how these urns are constructed, then correlation aversion alone would not produce ambiguity aversion—a different mechanism would be required to generate Ellsberg behavior. This example is therefore useful for illustrating both the scope and the limitations of the evolutionary model: Adaptive preferences generate ambiguity aversion anytime there is uncertainty about the model itself or some other factor that is common to all individuals in the population, which we contend is the case in the vast majority of examples and applications of ambiguity.¹⁸ As noted earlier, in cases involving idiosyncratic ambiguity, we do not take a stand on whether ambiguity aversion is driven by heuristics developed by the genotypes from the case of common uncertainty or if it arises from some other source.¹⁹

¹⁷It is important to keep in mind that for other decision problems random choice may be optimal. For instance, suppose instead that individuals are able to bet on any single color of their choosing: black, red, or yellow. An independent 50–50 randomization between betting on red or yellow will yield a 1/3 probability of winning for each of the three possible urn compositions; more importantly, the outcomes for each individual taking this randomization will be independently distributed. Thus, evolutionarily optimal preferences will be indifferent between betting on black and making this randomized bet on red and yellow.

¹⁸Halevy and Feltkamp (2005) suggested another mechanism by which correlation can generate ambiguity aversion: Risk aversion alone implies that an individual who makes repeated bets on an urn would rather draw from a risky than an ambiguous urn. In our evolutionary context, instead, the maximization of long-run growth generates an aversion to correlation in the contemporaneous draws of different individuals.

¹⁹If one is not convinced that the Ellsberg urn is a perfect fit for our model, the objects in the example

In line with the interpretation of ambiguity as model uncertainty, we favor a statistical interpretation of the smooth model used in this section, where each $\omega \in \Omega$ is a candidate for the true model (the law of nature governing the distribution of $s \in S$) and the marginal distribution of μ on Ω is a prior over the candidate models.²⁰ For simplicity, we treat μ as constant over time. In that case, evolutionary optimality requires that individuals’ preferences (eventually) assign the correct weights, so that μ becomes objective—it accurately reflects the data generating process.²¹ However, our evolutionary approach can easily be extended to allow μ_t to change with time t . For a simple example, suppose there is an index set K and a set of possible distributions $\mu^k \in \Delta(\Omega \times S)$, where $k \in K$ is redrawn periodically after finitely many periods. Then, in each period t , it is again evolutionarily optimal for individuals to maximize the growth rate in Equation (2), this time using their “best guess” of the distribution μ_t given all information available at time t . This information evolves as follows: One $\omega \in \Omega$ is commonly drawn each period, so that in between draws of k the marginal of μ^k on Ω is gradually revealed. At the same time, with a large number of individuals who each independently draw a state $s \in S$ each period, the conditional $\mu^k(\cdot|\omega)$ on S can be fully revealed in one period. In other words, in this situation ambiguity will only linger in the case of common uncertainty, in line with the discussion in Section 1.1.

4.1 Understanding Violations of Consequentialism

We now consider signals and updating in the special case of a single fitness function, $\Psi = \{\psi\}$. As above, we will focus on an example where deterministic choice is optimal, and we therefore restrict attention to deterministic action plans $(f_\sigma)_{\sigma \in \Sigma}$. In this case, the long-run growth rate in Equation (2) becomes

$$V((f_\sigma)_{\sigma \in \Sigma}) = \int_{\Omega} \ln \left(\int_{S \times \Sigma} \psi(f_\sigma(\omega, s)) d\mu(s, \sigma|\omega) \right) d\mu(\omega).$$

can be recast in terms of other examples discussed in the introduction. For instance, the acts B, R, Y could represent different medical treatments for a condition and the idiosyncratic states b, r, y could represent the events in which each treatment is successful for an individual, with B being a better understood treatment than R and with the efficacy of the combined treatment RY being better understood than that of BY .

²⁰See Klibanoff, Marinacci, and Mukerji (2005) or Marinacci (2015) for a discussion of this interpretation. An alternative interpretation is that the marginal of μ on Ω is a preference parameter that captures subjective plausibility of different first-order probabilistic beliefs $\mu(\cdot|\omega)$ on S .

²¹Halevy (2007) found that individuals who are ambiguity averse when betting on Ellsberg acts often also fail to reduce compound lotteries, in the sense that their preferences satisfy the pattern in the example even when the distribution over possible urns is objective. This seminal and often replicated finding has been hard to reconcile with the notion of ambiguity aversion. Our evolutionary approach can explain the equivalence between compound risk and ambiguity when the urn is common. In the original experiment a different urn was randomly chosen for each individual, and we remain agnostic about whether subjects rely on a heuristic from the case with common uncertainty, or whether a different mechanism is responsible for the failure to reduce purely idiosyncratic compound lotteries.

The long-run growth rate in Equation (4) from deviating from the plan $(f_\sigma)_{\sigma \in \Sigma}$ by instead selecting g following the signal $\bar{\sigma}$ becomes

$$V(g|\bar{\sigma}, (f_\sigma)_{\sigma \in \Sigma}) = \int_{\Omega} \ln \left(\mu(\bar{\sigma}|\omega) \int_S \psi(g(\omega, s)) d\mu(s|\omega, \bar{\sigma}) + \int_{S \times \Sigma \setminus \{\bar{\sigma}\}} \psi(f_\sigma(\omega, s)) d\mu(s, \sigma|\omega) \right) d\mu(\omega).$$

In particular, choice is dynamically consistent (this is a special case of Corollary 2).

We now revisit the special case of the smooth model from Example 2 (where acts f depend only on s) to illustrate the tension between consequentialism and dynamic consistency.

Example 3 (Ellsberg—with signals). *In the context of the urn in Example 2, suppose that individuals each receive a private signal that tells them whether the ball drawn for them is yellow (y) or not yellow ($\neg y$).²² As is standard in models of partitional learning, preferences over signal-contingent action plans for this information structure are entirely pinned down by preferences over acts. For example, since B and R both pay zero in state $s = y$, the action plan $R\neg yB$ that selects act R following the signal $\neg y$ and selects B following the signal y gives the same outcome in every state/signal combination (that occurs with positive probability) as the act R . Similarly, the action plan $R\neg yY$ gives the same outcome in every non-null state/signal combination as the act RY , and so on. Thus, the Ellsberg preferences over acts described above imply the following preferences over action plans:²³*

$$B\neg yB \succ R\neg yB \quad \text{and} \quad B\neg yY \prec R\neg yY.$$

Therefore, dynamic consistency requires that

$$B \succ_{\neg y, B\neg yB} R \quad \text{and} \quad B \prec_{\neg y, R\neg yY} R.$$

However, this pattern is incompatible with consequentialism, which would require that preferences between B and R following the signal $\neg y$ be independent of the ex ante action plan.

Note that the tension illustrated in this example depends neither on a particular choice of updating rule nor on the specific model of ambiguity aversion: Ellsberg behavior of the form $B\neg yB \succ R\neg yB$ and $B\neg yY \prec R\neg yY$ together with this specific information structure cannot satisfy both dynamic consistency and consequentialism.²⁴ Our model and results

²²Formally, for each $\omega \in \Omega$, we have $\mu(y|s, \omega) = 1$ if $s = y$ and $\mu(\neg y|s, \omega) = 1$ if $s = b, r$.

²³We are making two implicit assumptions in this argument (both of which are implied by our model). The first is that individuals are indifferent between action plans that differ only on μ -measure zero events, such as $R\neg yB$ and R . The second is that preferences between any two acts (trivial action plans) f and g does not change after the introduction of signals.

²⁴Note, in particular, that Ellsberg preferences with this information structure are therefore incompatible with the Epstein and Schneider (2003) model of multiple priors expected utility with rectangular priors. A similar example can be found in Hanany and Klibanoff (2007).

imply that individuals with these ex ante preferences will exhibit the conditional ex post preferences listed above. Thus, individuals will be dynamically consistent but will violate consequentialism.

As noted previously, consequentialism is violated because evolutionarily optimal preferences include an “externality” that incorporates the growth rate of other individuals in the population who are simultaneously receiving different signals. For example, the conditional preference between B and R following signal $\neg y$ given the ex ante action plan $R\neg yY$ is based on the following comparison:

$$\begin{aligned}
V(R|\neg y, R\neg yY) &= \sum_{\omega \in \Omega} \mu(\omega) \ln \left(\underbrace{\mu(\neg y|\omega)}_{\substack{\text{fraction} \\ \text{getting} \\ \text{signal } \neg y}} \underbrace{\mu(r|\omega, \neg y)}_{\substack{\text{average fitness} \\ \text{from } R \text{ after} \\ \text{signal } \neg y}} + \underbrace{\mu(y|\omega)}_{\substack{\text{fraction} \\ \text{getting} \\ \text{signal } y}} \underbrace{1}_{\substack{\text{fitness} \\ \text{from } Y \\ \text{after } y}} \right) \\
&> \sum_{\omega \in \Omega} \mu(\omega) \ln \left(\mu(\neg y|\omega) \mu(b|\omega, \neg y) + \mu(y|\omega) 1 \right) = V(B|\neg y, R\neg yY).
\end{aligned}$$

Notice the complementarity between r and y : The probability of seeing signal $\neg y$ and then state r is higher for urn compositions ω where the probability of signal y (and hence state y) is lower, as the first two columns of Table 2 illustrate.

	$\mu(\neg y \omega)\mu(r \omega, \neg y)$	$\mu(y \omega)$	$\mu(\neg y \omega)\mu(b \omega, \neg y)$
$\omega_1 = (b, r, r)$	$1 \cdot \frac{2}{3} = \frac{2}{3}$	0	$1 \cdot \frac{1}{3} = \frac{1}{3}$
$\omega_2 = (b, r, y)$	$\frac{2}{3} \cdot \frac{1}{2} = \frac{1}{3}$	$\frac{1}{3}$	$\frac{2}{3} \cdot \frac{1}{2} = \frac{1}{3}$
$\omega_3 = (b, y, y)$	$\frac{1}{3} \cdot 0 = 0$	$\frac{2}{3}$	$\frac{1}{3} \cdot 1 = \frac{1}{3}$

Table 2: Calculating the fitness from choosing R or B following signal $\neg y$

Choosing R following signal $\neg y$ thus achieves higher expected individual growth in precisely those instances when there are fewer individuals who contribute to aggregate growth by receiving signal y and then choosing Y . In contrast, choosing B does not hedge against this aggregate growth-rate risk, because the probability of state b is independent of the urn composition, as shown in the last column of the table. When the ex ante plan is instead $B\neg yB$, the hedging motive for the choice of R following $\neg y$ disappears, as now the growth rate following signal y is zero. In this case, we have the opposite conditional preference:

$$\begin{aligned}
V(B|\neg y, B\neg yB) &= \sum_{\omega \in \Omega} \mu(\omega) \ln \left(\mu(\neg y|\omega)\mu(b|\omega, \neg y) \right) \\
&> \sum_{\omega \in \Omega} \mu(\omega) \ln \left(\mu(\neg y|\omega)\mu(r|\omega, \neg y) \right) = V(R|\neg y, B\neg yB).
\end{aligned}$$

4.2 Lessons from and for the Lab

Our evolutionary approach to ambiguity aversion applies directly to situations where ambiguity can be identified with common uncertainty, that is, where the same uncertainty is faced by a large subpopulation. In contrast, the number of subjects in lab experiments is small. It may be that ambiguity aversion in the lab is due to common uncertainty about the motives of experimenters or is a heuristic based on the adaptive model, in which cases our insights would apply. Alternatively, other sources of ambiguity aversion might be at play, in which case our insights may not apply. Therefore, the validity of our model’s prediction of dynamic consistency in the laboratory strikes us as an empirical question. We now compare the predictions of our model to existing evidence, and in the process explain how the model can provide guidance on how to successfully test for dynamic consistency.

4.2.1 Testing Dynamic Consistency

When investigating dynamic consistency between ex ante signal-contingent plans and ex post choice after a particular signal realization, it is often implicitly assumed that the information structure that gives rise to that signal realization is irrelevant for ex post choice. For instance, Dominiak, Duersch, and Lefort (2012) considered a decision situation similar to Example 3. To efficiently collect ex post preferences contingent on the signal $\neg y$, their experimental design redraws the ball from the urn until the ball comes up black or red, so that no individual ever learns y .²⁵ While this protocol may be adequate in the context of some theoretical models, our model of adaptive preferences predicts that dynamic consistency and consequentialism are at odds with each other only when there are complementaries in payoffs between individuals who receive different signals. If no individuals receive the counterfactual signal (in this case y), then our model will not predict any violations of consequentialism.

To illustrate, recall from Table 2 that it is optimal to follow the plan $R\neg yY$ after learning $\neg y$ only because this choice hedges against fluctuations in population growth from individuals who learn y . However, if subjects never learn y , then this hedging motive disappears and adaptive preferences will instead favor B over R following $\neg y$, in line with the experimental findings in Dominiak, Duersch, and Lefort (2012):

$$\begin{aligned} V(B|\text{draw until } \neg y, (f_\sigma)_{\sigma \in \Sigma}) &= \sum_{\omega \in \Omega} \mu(\omega) \ln \left(\mu(b|\omega, \neg y) \right) \\ &> \sum_{\omega \in \Omega} \mu(\omega) \ln \left(\mu(r|\omega, \neg y) \right) = V(R|\text{draw until } \neg y, (f_\sigma)_{\sigma \in \Sigma}), \end{aligned}$$

for any ex ante action plan $(f_\sigma)_{\sigma \in \Sigma}$. Note that our model predicts this ranking both ex ante and ex post, so preferences are both dynamically consistent and consequentialist given this

²⁵Dominiak, Duersch, and Lefort (2012) used different colors for the balls in their experiment. We have translated to the colors used in our Example 3 for ease of exposition.

decision setting. This demonstrates that in order to test our model’s prediction of dynamic consistency, the entire event tree must be implemented.

A recent experiment by Bleichrodt et al. (2020) that did implement the entire game tree found that consequentialism was satisfied slightly more often than dynamic consistency (73.2% versus 66.2% of subjects). Interestingly, although not the focus of their experiment, some of their experimental evidence seems to suggest that violations of dynamic consistency may be connected with cognitive constraints and narrow bracketing. In particular, the primary violation of dynamic consistency observed in their experiment also constitutes a violation of monotonicity (stochastic dominance), and in many cases of consequentialism, if subjects properly integrate payoffs across the different questions in the experiment.²⁶

	Red	Blue	Yellow
Odd	33	M	$67 - M$
Even	33	M	$67 - M$

Table 3: Composition of Ambiguous Bag (Table 1 from Bleichrodt et al. (2020))

Specifically, Bleichrodt et al. (2020) offered subjects bets on cards drawn from an ambiguous bag, where half of the 200 cards carried an even number and the other half an odd number. For each parity, 33 cards were red, and 67 cards were either blue or yellow (see Table 3). For each subject, a card was drawn from this ambiguous bag at random. Subjects were asked two consecutive questions that were equally likely to be the one to determine their final reward. The first question asked subjects to bet on the color of a card drawn for them, contingent on its parity. Note that betting on yellow for odd parity and blue for even parity (or vice versa) is a dominant choice for a weakly ambiguity averse subject: it hedges perfectly against ambiguity and achieves a higher expected payoff than betting on red, since $67/2 > 33$. After answering the first question, subjects were told the parity of their card and asked again to bet on its color. Roughly half (49%) of the violations of dynamic consistency observed in this experiment involved switching from the unconditionally optimal bets to betting on red after learning the parity of the card. However, this choice pattern also constitutes a violation of monotonicity and potentially consequentialism. When considering the combination of the two questions, subjects should realize that if, for instance, the revealed parity is odd and they previously bet on yellow for that case in the first question, then betting on blue (rather than red) in the second question hedges perfectly against ambiguity and yields the highest possible expected payoff given the randomized payment scheme.

These experimental findings draw to light an interesting parallel and potential future research question: Narrow bracketing constitutes a failure to integrate payoffs from differ-

²⁶In a similar vein, Kuzmics (2017) showed that if subjects integrate payoffs across questions in an experiment and can self-randomize, then any behavior that respects monotonicity can be rationalized by expected utility for some subjective prior on the state space. Hence, any behavior inconsistent with expected utility when considering the experiment as a whole suggests possible narrow bracketing of questions.

ent aspects of one’s overall choice situation when making decisions. Consequentialism by definition involves not integrating other unrealized branches of the decision tree when making choices. In the experiment by Bleichrodt et al. (2020), there seems to be a connection between the two, since making the purportedly consequentialist choice of red after learning the parity of the card could only be viewed as optimal if subjects narrowly bracket the payments from different questions in the experiment. This begs the empirical question of whether these two behavioral patterns—narrow bracketing and consequentialism (at the expense of dynamic consistency)—correlate across individuals, and if so, whether both should be considered mistakes as our model would suggest.

4.2.2 Ambiguous Signals

While many models of updating with ambiguity consider an ambiguous prior with unambiguous signals, it is also possible that the information content of the signals themselves is ambiguous. For instance, there may be no prior ambiguity at all until the result of a poorly understood test becomes available. In a recent laboratory experiment, Epstein and Halevy (2020) examined the response of subjects to signals that have ambiguous precision, and they documented violations of the martingale property of beliefs.

Shishkin and Ortoleva (2020) subsequently tested one striking implication of all common consequentialist approaches to updating models of ambiguity aversion, namely that ambiguous information can have negative value. This implication seems counterintuitive, and indeed their evidence casts doubt on it: Ambiguity-averse subjects appear to ignore information unless it is valuable. Reacting to information only when this adds value is precisely what maximizes *ex ante* expected utility, and is hence the dynamically consistent course of action. In other words, their experimental findings are in line with our model of adaptive preferences, and its evolutionary foundation provides a rationale for them.

5 Applications: Non-Expected Utility

The tension between dynamic consistency and consequentialism is not exclusive to environments with ambiguity, but can also arise when updating models of non-expected utility for risk. Machina (1989) prominently argued that those models should be updated in a way that is dynamically consistent, even at the cost of consequentialism. The adaptive model accommodates violations of expected utility, and since updating in the adaptive model is dynamically consistent, our results support this general position for the models that it nests as special cases. To illustrate, in this section we consider another canonical special case of our model: rank-dependent utility. Perhaps surprisingly, we show in Section 5.1 that evolution can generate a version of dynamic RDU that is both dynamically consistent and consequentialist. This result depends crucially on the timing in the model, and in Section 5.2 we

discuss how the non-consequentialist version of dynamic RDU suggested by Machina (1989) is obtained when adaptation (the hidden action) instead takes place before the signal arrives.

5.1 DC and Consequentialist Updating of RDU

To focus on risk preferences, in this section we restrict attention to the special case of our model without common uncertainty ($\Omega = \{\omega\}$), but we now permit non-degenerate after-signal adaptation. In this case, there is no strict benefit to self-randomization, so it is without loss of generality to restrict attention to deterministic action plans $(f_\sigma)_{\sigma \in \Sigma}$ and adaptation plans $(\psi_\sigma)_{\sigma \in \Sigma}$.²⁷ Therefore, Equation (2) becomes

$$\begin{aligned} V((f_\sigma)_{\sigma \in \Sigma}) &= \sup_{(\psi_\sigma)_{\sigma \in \Sigma} \in \Psi^\Sigma} \ln \left(\int_{S \times \Sigma} \psi_\sigma(f_\sigma(s)) d\mu(s, \sigma) \right) \\ &= \ln \left(\int_{\Sigma} \sup_{\psi \in \Psi} \left[\int_S \psi(f_\sigma(s)) d\mu(s|\sigma) \right] d\mu(\sigma) \right). \end{aligned} \quad (5)$$

Note that in this case the logarithm can also be dropped by taking a monotone transformation, but we will retain it for consistency in expressing growth rates in log terms and for ease of comparing the formulas in this section to later results. Although the connection is non-trivial, the following result shows that rank-dependent utility with a pessimistic probability distortion function can be expressed as a special case of our model.

Proposition 1 (Updating Rank-Dependent Utility). *Suppose $\Omega = \{\omega\}$ and $Z \subset \mathbb{R}$. Fix μ , and fix any bounded nondecreasing function $u : Z \rightarrow \mathbb{R}$ and any function $\varphi : [0, 1] \rightarrow [0, 1]$ that is nondecreasing, concave, and onto. Then, there exists a set Ψ of functions $\psi : Z \rightarrow \mathbb{R}$ such that the ex ante value function V defined by Equation (5) can be equivalently expressed as*

$$V((f_\sigma)_{\sigma \in \Sigma}) = \ln \left(\int_{\Sigma} \int_Z u(z) d(\varphi \circ F_{f_\sigma, \mu(\cdot|\sigma)})(z) d\mu(\sigma) \right)$$

and ex post adaptive preferences following a signal $\bar{\sigma}$ are represented by

$$\hat{V}(g|\bar{\sigma}, (f_\sigma)_{\sigma \in \Sigma}) = \int_Z u(z) d(\varphi \circ F_{g, \mu(\cdot|\bar{\sigma})})(z),$$

where

$$F_{g, \mu(\cdot|\bar{\sigma})}(z) = \int_S \mathbf{1}[g(s) \leq z] d\mu(s|\bar{\sigma})$$

²⁷Formally, after dropping the expectation over Ω from Equation (2), we have

$$V(\rho) = \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \ln \left(\int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(s))] d\mu(s, \sigma) \right).$$

Since the expression inside the logarithm is linear in both τ and ρ , it is maximized by a deterministic action plan and adaptation plan.

denotes the cumulative distribution function of g given μ and $\bar{\sigma}$.

For the intuition behind this result, consider first the case of no signals. A key step in the logic of Proposition 1 is a duality result that shows that for any probability distortion function φ as in the proposition, there exists a set of fitness functions Ψ such that, for any act f ,²⁸

$$\sup_{\psi \in \Psi} \int_S \psi(f(s)) d\mu(s) = \int_Z u(z) d(\varphi \circ F_{f,\mu})(z). \quad (6)$$

In the case with signals, we have an expression similar to the left side of this equation inside an expectation over signals, and hence the same duality applies. Specifically, applying the dual formula in Equation (6) to the acts f_σ and measures $\mu(\cdot|\sigma)$ in Equation (5) immediately yields the ex ante value function in Proposition 1. The formula for the ex post value function follows similarly from Theorem 2. However, we denote this value function by $\hat{V}(g|\bar{\sigma}, (f_\sigma))$ rather than $V(g|\bar{\sigma}, (f_\sigma))$ to emphasize that it differs from Equation (4) not only because it is expressed as a rank-dependent utility, but also because we drop the logarithm and the conditional fitness associated with other signals. This is possible because ex post preferences do not depend on what happens after signals $\sigma \neq \bar{\sigma}$ in the case without common uncertainty.

Since μ only captures idiosyncratic uncertainty in this section and since we identify idiosyncratic uncertainty with pure risk, the distribution of outcomes $F_{f_\sigma, \mu(\cdot|\sigma)}$ following each signal σ amounts to an unambiguous risky prospect. The rank-dependent utility representation with concave φ in Proposition 1 suggests that individuals violate expected utility when choosing over risk by overweighting the probability assigned to worse outcomes. In other words, given the appropriate set of fitness functions Ψ , adaptive preferences are equivalent to maximizing expected utility with distorted probability weights.²⁹

5.2 Importance of the Timing of Adaptation and Information

In contrast to the examples considered in Machina (1989), ex post preferences in Proposition 1 are actually independent of the plan $(f_\sigma)_{\sigma \in \Sigma}$. That is, consequentialism is not violated by this dynamically consistent version of the rank-dependent utility model with information. The evolutionary intuition behind this result is that σ realizes prior to adaptation, and in our model, idiosyncratic risk that resolves before the selection of the hidden action is evaluated

²⁸The dual formula in Equation (6) is similar to several existing results in the literature. See, for example, Wakker (1994), Chatterjee and Krishna (2011), or the Supplementary Material of Sarver (2018).

²⁹In this paper, we focus on exploring the scope of adaptive preferences, and hence of our insights on dynamically consistent updating, by identifying special cases that can be nested. Sarver (2018) considers a similar representation to Equation (5) (but without signals), which also nests RDU. He further shows that his model does not overlap with other prominent non-expected-utility models (disappointment aversion, betweenness preferences, cautious expected utility) except in the case of expected utility. These insights are easily extended to our model and help delineate the boundary of the set of preferences that it nests. As we briefly touch on in Section 6.3, a natural next step is to characterize further restrictions on preferences in terms of properties of the set Ψ .

in accordance with expected utility. This is reflected by the ex ante value function, where only the cdf $F_{f_\sigma, \mu(\cdot|\sigma)}$ over outcomes *given* σ is distorted by φ , rather than the unconditional distribution that also incorporates uncertainty about the realization of σ itself.

In Section S1 of the Online Appendix, we show that if adaptation instead takes place before information, so that the hidden action ψ has to be chosen before the realization of σ , then the RDU distortion function will be applied to all uncertainty, including the signal realization. That is, for any concave distortion function φ , we show there exists a set Ψ such that the value function for *before-signal adaptation* is

$$V((f_\sigma)_{\sigma \in \Sigma}) = \ln \left(\int_Z u(z) d(\varphi \circ F_{(f_\sigma), \mu})(z) \right)$$

where

$$F_{(f_\sigma), \mu}(z) = \int_{S \times \Sigma} \mathbf{1}[f_\sigma(s) \leq z] d\mu(s, \sigma).$$

In this case, the rank of an outcome depends on the entire plan $(f_\sigma)_{\sigma \in \Sigma}$, and by dynamic consistency the ex post value of an act g must also depend on this ex ante action plan, violating consequentialism. This is exactly the approach to modeling rank-dependent utility with information suggested by Machina (1989). Thus, the evolutionary perspective serves not only to support Machina’s well-established approach in the case of before-signal adaptation, but also generates a novel and perhaps unexpected version of dynamic RDU preferences in the case of after-signal adaptation (Proposition 1).³⁰

6 Ambiguity Aversion and Non-Expected Utility

Many models of choice under uncertainty are founded on behavioral axioms rather than evolution. We already observed in Sections 4 and 5 that our model of adaptive preferences nests as special cases rank-dependent utility in the context of risk and a version of the smooth model in the context of ambiguity. Indirectly, our approach thus provides evolutionary foundations for the *behavior* those models represent and determines how they should take into account information. In this section, we expand our analysis to special cases of our representation that simultaneously incorporate both ambiguity aversion and non-expected

³⁰Existing evidence on the updating of non-expected-utility preferences and dynamic consistency is somewhat inconclusive. Cubitt, Starmer, and Sugden (1998) used a between-subject experimental design and found violations of our notion of dynamic consistency, but some more recent studies offer more favorable evidence: Hey and Paradiso (2006) examined subjects’ valuations for different decision problems and found that a slight majority of subjects (56%) behaved in accordance with dynamic consistency (their data could not distinguish whether these subjects had expected-utility or non-expected-utility preferences). Hey and Panaccione (2011) conducted a study that distinguished between expected-utility and non-expected-utility subjects and found that among the latter, between 64% and 91% were best classified as “resolute”: They followed through on their ex ante plans, in accordance with our definition of dynamic consistency.

utility for risk.³¹

One impediment to the analysis of special cases of our general representation is that it has a logarithm between the two layers of integration. For example, our results for rank-dependent utility in the previous section assumed that there was no common uncertainty, and it is not immediately obvious how those results might be extended to the general case of both common and idiosyncratic uncertainty. Therefore, we begin our analysis in this section with a duality result to recast our representation in a form that facilitates the analysis of this and other special cases. We then proceed to study several special cases in detail in Sections 6.1 and 6.2. In order to streamline the exposition, we focus in this section on the simplified setting with a trivial signal structure $\Sigma = \{\sigma\}$, which allows signals to be dropped from the model. Therefore, the set of action plans becomes $\mathcal{R}(\mathcal{F}) = \Delta^s(\mathcal{F})$ instead of $(\Delta^s(\mathcal{F}))^\Sigma$, the set of adaptation plans becomes $\mathcal{R}(\Psi|\mathcal{F}) = (\Delta^s(\Psi))^\mathcal{F}$ instead of $(\Delta^s(\Psi))^{\Sigma \times \mathcal{F}}$, and Equation (2) reduces to

$$V(\rho) = \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \int_{\Omega} \ln \left(\int_S \mathbb{E}_{\tau \otimes \rho} [\psi(f(\omega, s))] d\mu(s|\omega) \right) d\mu(\omega). \quad (7)$$

The analysis of signals in the applications in Sections 4 and 5 should make clear both the importance of updating and how it operates within our model. While we leave out signals for ease of exposition in this section, we provide the corresponding dual representation result for the general case with signals in Appendix A.

Our results will involve the relative entropy (or Kullback–Leibler divergence) of one probability measure with respect to another, defined as follows:

$$R(p \| q) = \begin{cases} \int \ln \left(\frac{dp}{dq} \right) dp & \text{if } p \ll q, \\ \infty & \text{otherwise.} \end{cases}$$

The notation $p \ll q$ indicates that p is *absolutely continuous* with respect to q , that is, for any measurable set A , $q(A) = 0$ implies $p(A) = 0$. The term $\frac{dp}{dq}$ denotes the Radon–Nikodym derivative (density) of p with respect to q , which exists if and only if p is absolutely continuous with respect to q .³² It is a standard result that $R(p \| q) \geq 0$, with equality if and only if $p = q$. In what follows, for any probability measure $p \in \Delta(\Omega)$, let

$$M(p) = \{q \in \Delta(\Omega) : q \ll p \text{ and } R(p \| q) < \infty\}.$$

In particular, since $R(p \| q) < \infty$ requires that $p \ll q$, if $q \in M(p)$ then the measures p and

³¹There are relatively few models in the axiomatic decision theory literature that combine ambiguity aversion and non-expected utility for risk; see, for example, Segal (1987), Dean and Ortleva (2017) and Izhakian (2017).

³²Formally, $\frac{dp}{dq}$ is the integrable function that satisfies $p(A) = \int_A \frac{dp}{dq} dq$ for any measurable set A .

q are mutually absolutely continuous, that is, both $p \ll q$ and $q \ll p$.³³

When necessary to avoid confusion, we will denote the marginal distribution of μ on Ω by μ_Ω . Also, recall that we take $\ln(x) = -\infty$ for all $x \leq 0$. Finally, in order to accommodate certain special cases, it will be technically convenient to permit the fitness functions ψ to take the value $-\infty$, so throughout this section we assume that Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$. We maintain our previous assumption that Ψ is pointwise bounded above, that is, $\sup_{\psi \in \Psi} \psi(z) < \infty$ for every $z \in Z$. The only assumption we add in the next theorem is that Ψ is closed in the topology of pointwise convergence on the extended reals.

Theorem 3. *Suppose $\Sigma = \{\sigma\}$, suppose Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$ that is pointwise bounded above and closed in the topology of pointwise convergence (on the extended reals), and fix $\mu \in \Delta(\Omega \times S)$. For any random action plan $\rho \in \mathcal{R}(\mathcal{F})$, the function V defined by Equation (7) can be equivalently expressed as*

$$V(\rho) = \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\mathbb{E}_\rho \left[\sup_{\psi \in \Psi} \int_\Omega \int_S \psi(f(\omega, s)) d\mu(s|\omega) dq(\omega) \right] \right) + R(\mu_\Omega \| q) \right]. \quad (8)$$

See Appendix A for a more general version of this result for the case with signals. For intuition, we highlight the key steps in the proof: First, using duality techniques related to those employed in the literature on large deviations in statistics (cf. Dupuis and Ellis (1997)), we show that Equation (7) can be equivalently expressed as

$$V(\rho) = \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_\Omega \int_S \mathbb{E}_{\tau \otimes \rho} [\psi(f(\omega, s))] d\mu(s|\omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right].$$

This expression is not yet amenable to analysis, as we would like to reverse the order of the supremum and infimum in order to further simplify it and connect with existing functional forms. The next step in the proof is to do just that by leveraging a particular version of the von Neumann–Sion minimax theorem (von Neumann (1928), Sion (1958)) that is due to Tuy (2004). Then, after we switch the order of the supremum and infimum, the supremum over τ applies to the expression inside the logarithm, which is linear in τ . Therefore, optimization over adaptation plans τ can be reduced to the deterministic selection of a fitness function ψ following every act f that realizes under ρ , giving Equation (8). This final observation will greatly simplify the analysis of the model since it eliminates randomization over ψ from the formula for long-run growth rates.

Despite the resemblance, the functional in Equation (8) with a single fitness function $\Psi = \{\psi\}$ is not a variational representation (Maccheroni, Marinacci, and Rustichini (2006)). The distinction is the logarithm around the integral in the first term. In fact, in the case of a single fitness function, taking the exponential transformation of the representation in Equation (8)

³³Note that it is possible to have $R(p \| q) = \infty$ even if $p \ll q$, so $M(p)$ may be a strict subset of the set of all measures that are mutually absolutely continuous with respect to p .

establishes it as a special case of the confidence preferences studied by Chateauneuf and Faro (2009), where confidence in a prior q is measured by $\exp(R(\mu_\Omega \parallel q))$. More generally, this no-adaptation case is also nested by the general representation for uncertainty-averse preferences proposed by Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011).

Turning to the specifics of our functional form, relative entropy has appeared in a number of representations for ambiguity-averse preferences, perhaps most notably in the multiplier preferences introduced by Hansen and Sargent (2001) and studied axiomatically by Strzalecki (2011),³⁴ and also within a version of confidence preferences in Chateauneuf and Faro (2012). However, in these models, the entropy term used is $R(q \parallel \mu_\Omega)$ rather than $R(\mu_\Omega \parallel q)$. While relative entropy is often interpreted as a “distance” between the two distributions involved, it is not a distance function in the metric sense, because it is not symmetric. To interpret the subtle difference in the context of the representation in Equation (8), suppose the decision-maker takes as the reference measure μ_Ω the empirical frequencies in a large sample of independently realized states $\omega \in \Omega$, but worries that the data is actually generated by the measure q on Ω . Of course, the larger the sample, the closer to zero the probability that it would be generated by $q \neq \mu_\Omega$. The theory of large deviations establishes that the rate at which this probability vanishes increases in $R(\mu_\Omega \parallel q)$ (see, e.g., Cover and Thomas (2006, Section 11.4)). The representation suggests, therefore, that the decision-maker is less confident in a measure q the faster it becomes implausible with growing sample size.

In order to describe the special cases of the next two subsections, it will be convenient to define a measure $\mu \otimes q$ on $\Omega \times S$ with marginal q on Ω and conditional distribution $\mu(\cdot|\omega)$ on S . That is, for any event E in the product σ -algebra $\mathcal{B}_\Omega \otimes \mathcal{B}_S$, let

$$\mu \otimes q(E) = \int_\Omega \int_S \mathbf{1}[(\omega, s) \in E] d\mu(s|\omega) dq(\omega).$$

With this definition in hand, Equation (8) can be written as

$$V(\rho) = \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\mathbb{E}_\rho \left[\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d(\mu \otimes q)(\omega, s) \right] \right) + R(\mu_\Omega \parallel q) \right]. \quad (9)$$

6.1 Nesting Rank-Dependent Utility

Proposition 1 linked our adaptive model to RDU preferences in the special case of $\Omega = \{\omega\}$, in which case the state space was effectively S . The next corollary follows from the same duality arguments (see Equation (6)) after replacing S with $\Omega \times S$ and replacing the measure $\mu \in \Delta(S)$ with $\mu \otimes q \in \Delta(\Omega \times S)$. Note that this application is only possible because we first apply Theorem 3 to remove the logarithm from between the two layers of integration.

³⁴Hansen and Sargent (2001) interpret their representation in terms of a concern about robustness to model misspecification. Our approach provides a related perspective on concern for robustness in contexts where uncertainty about ω can be interpreted as model uncertainty.

Corollary 3. *Suppose $Z \subset \mathbb{R}$. Fix μ , and fix any bounded nondecreasing function $u : Z \rightarrow \mathbb{R}$ and any function $\varphi : [0, 1] \rightarrow [0, 1]$ that is nondecreasing, concave, and onto. Then, there exists a set Ψ of functions $\psi : Z \rightarrow \mathbb{R}$ satisfying the assumptions of Theorem 3 such that, for any act $f \in \mathcal{F}$ and any $q \in \Delta(\Omega)$,*

$$\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d(\mu \otimes q)(\omega, s) = \int_Z u(z) d(\varphi \circ F_{f, \mu \otimes q})(z),$$

where

$$F_{f, \mu \otimes q}(z) = \int_{\Omega \times S} \mathbf{1}[f(\omega, s) \leq z] d(\mu \otimes q)(\omega, s)$$

is the cumulative distribution function of f given $\mu \otimes q$. Therefore, for that set Ψ , the function V defined by Equation (7) can be equivalently expressed as

$$V(\rho) = \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\mathbb{E}_\rho \left[\int_Z u(z) d(\varphi \circ F_{f, \mu \otimes q})(z) \right] \right) + R(\mu_\Omega \| q) \right]$$

This representation illustrates the simplicity of analyzing the combination of ambiguity aversion, non-expected-utility risk preferences, and random choice when working with the dual formula in Equation (9) and its special cases. In this application, the RDU representation inside the logarithm generates aversion to any kind of uncertainty, while ambiguity aversion (roughly speaking, the *additional* aversion to uncertainty from Ω) is captured by the outer part of the representation—the confidence preferences within which the RDU representation is embedded. The outer part is fixed across genotypes, even if those differ in terms of Ψ and hence in terms of their attitudes towards risk.³⁵ Random choice of acts is also easy to analyze in this representation, since the expectation with respect to ρ appears inside the confidence preferences (reflecting the hedging benefits of self-randomization) but outside of the RDU formula. In contrast, as observed in Section 3.1, exogenous sources of randomization (mixtures) would be treated differently: Exogenous randomization would be subject to the RDU probability distortion function, thereby lessening its hedging benefit.

6.2 Nesting Divergence Preferences

Definition 6. Fix a continuous convex function $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $\phi(1) = 0$. The ϕ -divergence of p with respect to q is given by

$$D_\phi(p \| q) = \begin{cases} \int \phi\left(\frac{dp}{dq}\right) dq & \text{if } p \ll q, \\ \infty & \text{otherwise.} \end{cases}$$

³⁵There is some empirical evidence that risk aversion and additional aversion to ambiguity indeed have little correlation in the population (Chapman et al. (2019)).

Kullback–Leibler relative entropy is a special case of ϕ –divergence where $\phi(t) = t \ln(t) - t + 1$. Maccheroni, Marinacci, and Rustichini (2006) observed that variational preferences with a divergence cost function are probabilistically sophisticated. Ben-Tal and Teboulle (1987, 2007) provided an explicit dual characterization of these variational divergence preferences as the supremum of a set of expected utilities under the reference measure, where the supremum is taken over a set of possible Bernoulli utility indices. The following proposition extends their result to permit a nondecreasing transformation k of the divergence term.

Proposition 2 (Divergence Duality). *Fix any ϕ –divergence $D_\phi(\cdot \| \cdot)$ and any function $u : Z \rightarrow \mathbb{R}$. Also, fix any nondecreasing, convex, and lower semicontinuous function $k : \mathbb{R} \rightarrow (-\infty, \infty]$ such that $k(0) = 0$ and k is finite on some interval $(-\varepsilon, \varepsilon)$. Then, there exists a set Ψ satisfying the assumptions of Theorem 3 such that, for any $f \in \mathcal{F}$ and any $p \in \Delta(\Omega \times S)$,*³⁶

$$\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) dp(\omega, s) = \inf_{r \in \Delta(\Omega \times S)} \left[\int_{\Omega \times S} u(f(\omega, s)) dr(\omega, s) + k(D_\phi(r \| p)) \right].$$

The following corollaries apply Proposition 2 to our representation for adaptive preferences from Theorem 3 by taking $p = \mu \otimes q$. The first corollary considers the special case of $k(x) = \theta x$ for some scalar $\theta > 0$.

Corollary 4. *Fix any ϕ –divergence $D_\phi(\cdot \| \cdot)$, any scalar $\theta > 0$, and any function $u : Z \rightarrow \mathbb{R}$. Then, there exists a set Ψ of functions $\psi : Z \rightarrow [-\infty, \infty)$ such that the function V defined by Equation (7) can be equivalently expressed as*

$$V(\rho) = \inf_{q \in \mathcal{M}(\mu_\Omega)} \left[\ln \left(\mathbb{E}_\rho \left[\inf_{r \in \Delta(\Omega \times S)} \int_{\Omega \times S} u(f(\omega, s)) dr(\omega, s) + \theta D_\phi(r \| \mu \otimes q) \right] \right) + R(\mu_\Omega \| q) \right].$$

This value function embeds a general divergence representation inside confidence preferences. To see how it captures ambiguity aversion, note that the measure r ultimately used to evaluate an act may be more pessimistic than $\mu \otimes q$ on $\Omega \times S$, which in turn may be more pessimistic than μ only on Ω . Hence, compared to μ , there is more “opportunity” for r to be pessimistic about Ω than about S .

The next corollary considers the special case of Proposition 2 where we fix a scalar $\kappa > 0$ and take $k(x) = 0$ if $x \leq \kappa$, and $k(x) = +\infty$ if $x > \kappa$.

Corollary 5. *Fix any ϕ –divergence $D_\phi(\cdot \| \cdot)$ and any function $u : Z \rightarrow \mathbb{R}$. Fix a scalar $\kappa > 0$, and for any $p \in \Delta(\Omega \times S)$ define*

$$\mathcal{D}(p, \kappa) = \{r \in \Delta(\Omega \times S) : D_\phi(r \| p) \leq \kappa\}.$$

³⁶We adopt the convention that $k(\infty) = \infty$. Thus, for any function k as in the statement of the proposition, if $D_\phi(r \| p) = \infty$ then $k(D_\phi(r \| p)) = \infty$.

Then, there exists a set Ψ of functions $\psi : Z \rightarrow [-\infty, \infty)$ such that the function V defined by Equation (7) can be equivalently expressed as

$$V(\rho) = \inf_{q \in \mathcal{M}(\mu_\Omega)} \left[\ln \left(\mathbb{E}_\rho \left[\inf_{r \in \mathcal{D}(\mu \otimes q, \kappa)} \int_{\Omega \times S} u(f(\omega, s)) dr(\omega, s) \right] \right) + R(\mu_\Omega \| q) \right].$$

In this value function, the multiple prior representation (Gilboa and Schmeidler (1989)) inside the logarithm generates aversion to any kind of uncertainty, while ambiguity aversion is again captured by the confidence preferences that this representation is embedded within.

6.3 Characterizing the Set of Fitness Functions

In this paper, we focused on exploring the scope of our model (and hence of our insights on dynamically consistent updating) by identifying a range of special cases that can be nested by adaptive preferences. A natural next step in this line of research is to examine how properties of preferences connect to restrictions on the set Ψ . We conclude this section by mentioning examples of the types of results one might obtain.

On the one hand, one could try to characterize particular special cases of our model in terms of Ψ , as in the previous subsections. In the context of pure risk, the Supplementary Material of Sarver (2018) provides another example that easily lends itself to economic interpretation, where each ψ is piecewise linear around a different target consumption level and the resulting preferences are RDU with a piecewise linear distortion function.

On the other hand, one could compare adaptive preferences for different Ψ .³⁷ Suppose, for instance, that all conceivable genotypes perform equally well when facing deterministic outcomes (no uncertainty). In terms of the model of adaptive preferences, this means that the upper envelope of Ψ is the same for all those genotypes. In this case, one can show that individual A with adaptive preferences for Ψ_A is more risk averse than an individual B with Ψ_B if and only if individual A is also more uncertainty averse than B . For example, in the representations of Corollaries 3, 4 and 5 the upper envelope of Ψ is u , and hence holding fixed u , individuals with any of these three types of preferences who can be ranked in terms of risk aversion will be ranked the same way in terms of overall uncertainty aversion.

7 Realism of the Evolutionary Model

Section 7.1 discusses two assumptions that are implicit in our formulation of the evolutionary model and that are commonly made in economic contexts. Section 7.2 concludes by

³⁷We have taken the set Ψ as given throughout. To compare individuals with different Ψ , it is important to understand how Ψ is determined. One possibility is that different choice situations involve different sets of hidden actions. Another possibility is that Ψ itself is subject to constrained evolutionary optimization.

discussing the interpretation of adaptive preferences in the context of phenotypic flexibility.

7.1 Simplifying Assumptions

Corollary 1 shows that the long-run growth rate is optimized by choosing the action plan $\rho \in \mathcal{R}(A)$ that maximizes V , assuming the decision problem A is faced by the genotype repeatedly *in every period*. In fact, this assumption is unnecessarily strong and is made solely for ease of exposition. As can be seen in the proof of Theorem 1, aggregate fitness in each period affects the population size multiplicatively, which provides a degree of separability for choice problems that appear at different times. For example, if the genotype faces an infinite sequence of decision problems $(A_t)_{t \in \mathbb{N}}$, then attaining the highest possible long-run growth rate requires that individuals maximize adaptive preferences from any decision problem A that repeats with fixed frequency within this sequence.³⁸

The second assumption in our model is that time is divided into discrete time periods. Robatto and Szentes (2017) made the surprising observation that correlation aversion disappears in the continuous-time limit of this basic model. Further extending this line of research, Robson and Samuelson (2019) allowed fertility and mortality rates to vary with age in order to separate the assumption of continuous time from the assumption that new organisms can reproduce immediately after birth, and they found that correlation aversion can be recovered even in continuous time. Investigating the implications of different timing and age structures in our context of hidden actions and updating could be an interesting avenue for future research. In this paper, we stick to discrete time with age-independent fertility and mortality rates as is common in evolutionary models in economics.

7.2 Phenotypic Flexibility in Evolutionary Biology

While our approach is inspired by evolutionary biology, we hope that our insights might in turn also be useful in biological contexts where phenotypic flexibility plays a role, as we now explain in more detail. Evolutionary success appears to be greatly enhanced by the ability of organisms of a particular genotype to adapt their phenotype to the environment. Adopting the terminology proposed by Piersma and Drent (2003), we use *phenotypic flexibility* to refer to the rapid and apparently purposeful variation in phenotype expressed by individual reproductively mature organisms throughout their life. This is in contrast to *developmental plasticity*, environmentally induced variations that occur only during development.³⁹

³⁸The assumption that all individuals of the genotype face the same decision problem at the same time is also implicit in our model, and this assumption can be relaxed as well. If, instead, there is a distribution of decision problems within the population, then this uncertainty can be encoded into the state spaces in our model (similar to the way we incorporate signals and contingent plans).

³⁹Piersma and Drent (2003) use *phenotypic plasticity* as an umbrella term that includes both phenotypic flexibility and developmental plasticity.

While developmental plasticity has long been a focus of evolutionary biologists, the role of phenotypic flexibility in the evolutionary process has only recently attracted significant attention. According to Piersma and Drent (2003):

When environmental conditions change rapidly [...] individuals that can show continuous but reversible transformations in behaviour, physiology and morphology might incur a selective advantage. There are now several studies documenting substantial but reversible phenotypic changes within adult organisms.

Striking examples among vertebrates include various species of amphibious fish that adjust to life on land with reversible and rapid (sometimes within minutes) changes to their muscle tissue, breathing organs, and skin properties (Wright and Turko (2016) provide a survey), or marine iguanas on the Galapagos islands that can shrink their overall body length by up to 20% (6.8 cm) in what appears to be a reversible, rapid, and strategic response to food scarcity during an El Niño weather pattern (Wikelski and Thom (2000)). A familiar example that can be viewed as phenotypic flexibility in humans and other mammals is the adjustment of the makeup of muscle tissue in response to changes in functional demands (Flück (2006)), for instance, from a more or less active lifestyle.

Of course, the evolutionary benefit of phenotypic flexibility is that different phenotypes may perform better in different situations, and hence have different fitness functions ψ . For instance, each possible phenotype might be tailored to a specific range of outcomes, such as the amount of available food for the iguanas in the example above. Or one phenotype might be a specialist with high fitness for a small range of outcomes, while the other is a generalist, with lower peak fitness that is more robust to the outcome.

Biologists in the studies above directly observe variations in individual phenotypes over time. In economic applications, in contrast, phenotypes, such as the determinants of risk and ambiguity preferences in our model, are notoriously hard to observe. Economists instead rely on preferences that are revealed from observable choice data. Respecting this limitation, our model predictions concern only observable choices between outcome-relevant actions (f), treating the phenotype and resulting fitness function (ψ) as unobservable. As a consequence, our model does not distinguish between the case where adaptation is due to a biological change (phenotypic flexibility) or a strategic but hidden choice of action, and it is equally relevant and applicable under either interpretation of the set of fitness functions Ψ .

A Duality: General Treatment

In this section, we generalize Theorem 3 from Section 6 to allow for a nondegenerate signal structure. For any $q \in \Delta(\Omega)$, define the measure $\mu \otimes q$ on $\Omega \times S \times \Sigma$ to have marginal q on Ω and conditional distribution $\mu(\cdot|\omega)$ on $S \times \Sigma$. That is, for any event E in the product σ -algebra $\mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma$, let

$$\mu \otimes q(E) = \int_\Omega \int_{S \times \Sigma} \mathbf{1}[(\omega, s, \sigma) \in E] d\mu(s, \sigma|\omega) dq(\omega).$$

Theorem 4. *Suppose Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$ that is pointwise bounded above and closed in the topology of pointwise convergence (on the extended reals), and fix $\mu \in \Delta(\Omega \times S \times \Sigma)$. For any random action plan $\rho \in \mathcal{R}(\mathcal{F})$, the function V defined by Equation (2) can be equivalently expressed as*

$$V(\rho) = \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_\Sigma \mathbb{E}_{\rho_\sigma} \left[\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d(\mu \otimes q)(\omega, s|\sigma) \right] d(\mu \otimes q)(\sigma) \right) + R(\mu_\Omega \| q) \right]. \quad (10)$$

Special cases such as rank-dependent utility or divergence preferences can be embedded in this general representation with signals analogously to our analysis of special cases in Section 6.

B Proofs

B.1 Proof of Lemma 1

Note that

$$\ln(N^i(T)) = \ln(N^i(0)) + \sum_{t=1}^T \ln(\lambda_t^i),$$

and therefore

$$\ln \left(\frac{N^A(T)}{N^B(T)} \right) = \ln \left(\frac{N^A(0)}{N^B(0)} \right) + \sum_{t=1}^T \ln(\lambda_t^A) - \sum_{t=1}^T \ln(\lambda_t^B).$$

Since α^A and α^B are the long-run growth rates of these two genotypes, we have

$$\frac{1}{T} \left[\sum_{t=1}^T \ln(\lambda_t^A) - \sum_{t=1}^T \ln(\lambda_t^B) \right] \rightarrow \alpha^A - \alpha^B \quad \text{a.s.}$$

Since $\alpha^A - \alpha^B > 0$, this implies

$$\ln \left(\frac{N^A(T)}{N^B(T)} \right) \rightarrow \infty \quad \text{a.s.}$$

Therefore, $N^A(T)/N^B(T) \rightarrow \infty$ almost surely as $T \rightarrow \infty$. This completes the proof.

B.2 Proof of Theorem 4

The following two propositions will be central in our proof of Theorem 4.

Proposition 3. *Suppose Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$ that is pointwise bounded above, and fix $\mu \in \Delta(\Omega \times S \times \Sigma)$. For any random action plan $\rho \in (\Delta^s(\mathcal{F}))^\Sigma$, the function V defined by Equation (2) can be equivalently expressed as*

$$V(\rho) = \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_{\Omega} \int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right].$$

Proposition 4. *Fix a measure $\mu \in \Delta(\Omega \times S \times \Sigma)$, and suppose Ξ is a nonempty set of functions $\xi : \Omega \times S \times \Sigma \rightarrow [-\infty, \infty)$ with the following properties:*

1. *Closedness: When the set of extend reals $[-\infty, \infty]$ is endowed with its usual topology and $[-\infty, \infty]^{\Omega \times S \times \Sigma}$ is endowed with the product topology (i.e., the topology of pointwise convergence), Ξ is a closed subset of this space.*
2. *Finite measurability: There exists a finite partition $\mathcal{E} \subset \mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma$ of $\Omega \times S \times \Sigma$ such that every $\xi \in \Xi$ is measurable with respect to \mathcal{E} .*
3. *Pointwise boundedness: $\sup_{\xi \in \Xi} \xi(\omega, s, \sigma) < \infty$ for every $(\omega, s, \sigma) \in \Omega \times S \times \Sigma$.*

Then,

$$\begin{aligned} & \sup_{\xi \in \text{co}(\Xi)} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_{\Omega} \int_{S \times \Sigma} \xi(\omega, s, \sigma) d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right] \\ &= \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\sup_{\xi \in \Xi} \int_{\Omega} \int_{S \times \Sigma} \xi(\omega, s, \sigma) d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right]. \end{aligned}$$

Proposition 3 is based on dual formulas for relative entropy that are related to those commonly invoked in the theory of large deviations (e.g., Dupuis and Ellis (1997)). Proposition 4 is based on an application of an extension of the von Neumann–Sion Minimax Theorem due to Tuy (2004). Despite the reliance on these established tools and techniques, the complete proofs of these propositions are quite involved and are therefore relegated to Section S5 of the Online Appendix.

Proceeding with the proof of Theorem 4, fix any $\rho \in (\Delta^s(\mathcal{F}))^\Sigma$. For each $\sigma \in \Sigma$, let $B_\sigma = \text{supp}(\rho_\sigma)$. Since ρ_σ is a simple lottery, B_σ is finite. Let $B = \bigcup_{\sigma \in \Sigma} B_\sigma$. Since Σ is finite, B is a finite set of acts. We will define Ξ to be the set of individual expected fitness functions that are attainable given the fixed random choice of act under the action plan ρ and together with some deterministic adaptation plan. That is, we are focusing for now on adaptations plans τ that place probability one on some fitness function $\psi_{\sigma, f} \in \Psi$ following each $\sigma \in \Sigma$ and $f \in B$.

Formally, deterministic adaptation plans are denoted by $(\psi_{\sigma, f})_{\sigma \in \Sigma, f \in B} \in \Psi^{\Sigma \times B}$, or $(\psi_{\sigma, f})$ for

short.⁴⁰ Define a mapping $J : \Psi^{\Sigma \times B} \rightarrow [-\infty, \infty]^{\Omega \times S \times \Sigma}$ by

$$J\left[(\psi_{\hat{\sigma}, f})_{\hat{\sigma} \in \Sigma, f \in B}\right](\omega, s, \sigma) = \int_B \psi_{\sigma, f}(f(\omega, s)) d\rho_{\sigma}(f) \quad (11)$$

for $(\omega, s, \sigma) \in \Omega \times S \times \Sigma$. Define Ξ to be the range of J , that is,

$$\Xi = \left\{ J[(\psi_{\sigma, f})] \in [-\infty, \infty]^{\Omega \times S \times \Sigma} : (\psi_{\sigma, f}) \in \Psi^{\Sigma \times B} \right\}. \quad (12)$$

In other words, Ξ is the set of all functions ξ that take the form

$$\xi(\omega, s, \sigma) = \int_B \psi_{\sigma, f}(f(\omega, s)) d\rho_{\sigma}(f)$$

for some deterministic adaptation plan $(\psi_{\sigma, f})_{\sigma \in \Sigma, f \in B}$. The next two lemmas show that taking the convex hull of Ξ generates precisely the set of individual expected fitness functions that can be attained through random adaptation plans and that the set Ξ is closed. Indeed, the use of deterministic action plans above was precisely in order to ensure that Ξ is closed. The proofs of these two lemmas are based on standard arguments and are relegated to Section S5 of the Online Appendix.

Lemma 2. *Define Ξ as in Equation (12). For any random adaptation plan $\tau \in \mathcal{R}(\Psi|\mathcal{F})$, define $\xi^{\tau} : \Omega \times S \times \Sigma \rightarrow [-\infty, \infty]$ by*

$$\xi^{\tau}(\omega, s, \sigma) = \mathbb{E}_{\tau_{\sigma} \otimes \rho_{\sigma}}[\psi(f(\omega, s))] = \int_{\mathcal{F}} \int_{\Psi} \psi(f(\omega, s)) d\tau_{\sigma}(\psi|f) d\rho_{\sigma}(f).$$

Then,

$$\text{co}(\Xi) = \{\xi^{\tau} : \tau \in \mathcal{R}(\Psi|\mathcal{F})\}.$$

Lemma 3. *The set Ξ defined in Equation (12) is a closed subset of $[-\infty, \infty]^{\Omega \times S \times \Sigma}$.*

We now verify that the set Ξ defined in Equation (12) satisfies the three conditions from Proposition 4. Lemma 3 already showed that this set is closed, which establishes first condition. We now show that Ξ satisfies the second condition (finite measurability) from Proposition 4. Since each $f \in \mathcal{F}$ is a simple act, and since the set of acts B in the support of ρ is finite, there exists a finite partition $\hat{\mathcal{E}} \subset \mathcal{B}_{\Omega} \otimes \mathcal{B}_S$ of $\Omega \times S$ such that every act $f \in B$ is measurable with respect to $\hat{\mathcal{E}}$. Let

$$\mathcal{E} = \{\hat{E} \times \{\sigma\} : \hat{E} \in \hat{\mathcal{E}} \text{ and } \sigma \in \Sigma\}.$$

Since $\hat{\mathcal{E}}$ and Σ are finite, \mathcal{E} is a finite partition of $\Omega \times S \times \Sigma$. We claim that every function in Ξ is measurable with respect to \mathcal{E} . To see this, fix any $\xi \in \Xi$. Then, there exists $(\psi_{\sigma, f}) \in \Psi^{\Sigma \times B}$ such that

$$\xi(\omega, s, \sigma) = \int_B \psi_{\sigma, f}(f(\omega, s)) d\rho_{\sigma}(f).$$

⁴⁰Note that since $(\psi_{\sigma, f})$ is an element of $\Psi^{\Sigma \times B}$ rather than $\Psi^{\Sigma \times \mathcal{F}}$, the value of $\psi_{\sigma, f}$ is unspecified for $f \in \mathcal{F} \setminus B$. However, since acts $f \notin B$ are chosen with probability zero, expected individual fitness is fully determined by the values of $\psi_{\sigma, f}$ for $f \in B$.

Fix any $E \in \mathcal{E}$ and $(\omega, s, \sigma), (\omega', s', \sigma') \in E$. By construction of the partition \mathcal{E} , we must have $\sigma' = \sigma$ and $f(\omega, s) = f(\omega', s')$ for any $f \in \text{supp}(\rho_\sigma)$. Therefore,

$$\xi(\omega, s, \sigma) = \int_B \psi_{\sigma, f}(f(\omega, s)) d\rho_\sigma(f) = \int_B \psi_{\sigma, f}(f(\omega', s')) d\rho_\sigma(f) = \xi(\omega', s', \sigma'),$$

as claimed. Thus, the second condition of Proposition 4 is satisfied.

To verify the third condition (pointwise boundedness) in Proposition 4, note that since B is a finite set of simple acts, there is a finite set $\widehat{Z} \subset Z$ such that $f(\omega, s) \in \widehat{Z}$ for all $f \in B$ and $(\omega, s) \in \Omega \times S$. Recall that the set Ψ is pointwise bounded above, so $\sup_{\psi \in \Psi} \psi(z) < \infty$ for all $z \in Z$. Therefore, and any $(\omega, s, \sigma) \in \Omega \times S \times \Sigma$,

$$\begin{aligned} \sup_{\xi \in \Xi} \xi(\omega, s, \sigma) &= \sup_{(\psi_{\sigma, f}) \in \Psi^{\Sigma \times B}} \int_B \psi_{\sigma, f}(f(\omega, s)) d\rho_\sigma(f) \\ &\leq \int_B \sup_{\psi \in \Psi} \psi(f(\omega, s)) d\rho_\sigma(f) \leq \max_{z \in \widehat{Z}} \sup_{\psi \in \Psi} \psi(z) < \infty, \end{aligned}$$

where the last inequality follows from the finiteness of \widehat{Z} . Thus, Ξ satisfies condition 3.

We are now ready to apply Propositions 3 and 4. Define V as in Equation (2). Then, we have

$$\begin{aligned} V(\rho) &= \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_\Omega \int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right] \\ &= \sup_{\xi \in \text{co}(\Xi)} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_\Omega \int_{S \times \Sigma} \xi(\omega, s, \sigma) d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right] \\ &= \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\sup_{\xi \in \Xi} \int_\Omega \int_{S \times \Sigma} \xi(\omega, s, \sigma) d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right] \\ &= \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\sup_{\xi \in \Xi} \int_{\Omega \times S \times \Sigma} \xi(\omega, s, \sigma) d(\mu \otimes q)(\omega, s, \sigma) \right) + R(\mu_\Omega \| q) \right], \end{aligned}$$

where the first equality follows from Proposition 3, the second from Lemma 2, the third from Proposition 4, and the fourth from the definition of the measure $\mu \otimes q$. Simple manipulations of the term inside the logarithm yield

$$\begin{aligned} &\sup_{\xi \in \Xi} \int_{\Omega \times S \times \Sigma} \xi(\omega, s, \sigma) d(\mu \otimes q)(\omega, s, \sigma) \\ &= \sup_{(\psi_{\sigma, f}) \in \Psi^{\Sigma \times B}} \int_\Sigma \int_{\Omega \times S} \int_B \psi_{\sigma, f}(f(\omega, s)) d\rho_\sigma(f) d(\mu \otimes q)(\omega, s|\sigma) d(\mu \otimes q)(\sigma) \\ &= \sup_{(\psi_{\sigma, f}) \in \Psi^{\Sigma \times B}} \int_\Sigma \int_B \int_{\Omega \times S} \psi_{\sigma, f}(f(\omega, s)) d(\mu \otimes q)(\omega, s|\sigma) d\rho_\sigma(f) d(\mu \otimes q)(\sigma) \\ &= \int_\Sigma \int_B \sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d(\mu \otimes q)(\omega, s|\sigma) d\rho_\sigma(f) d(\mu \otimes q)(\sigma) \\ &= \int_\Sigma \mathbb{E}_{\rho_\sigma} \left[\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d(\mu \otimes q)(\omega, s|\sigma) \right] d(\mu \otimes q)(\sigma), \end{aligned}$$

and hence

$$V(\rho) = \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_{\Sigma} \mathbb{E}_{\rho_\sigma} \left[\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d(\mu \otimes q)(\omega, s | \sigma) \right] d(\mu \otimes q)(\sigma) \right) + R(\mu_\Omega \| q) \right].$$

Since this is true for any $\rho \in (\Delta^s(\mathcal{F}))^\Sigma$, the proof is complete.

B.3 Proof of Proposition 1 and Corollary 3

Proposition 1 and Corollary 3 follow immediately from the next general duality result.

Proposition 5 (Rank-Dependent Utility Duality). *Suppose $Z \subset \mathbb{R}$. Fix any bounded nondecreasing function $u : Z \rightarrow \mathbb{R}$ and any function $\varphi : [0, 1] \rightarrow [0, 1]$ that is nondecreasing, concave, and onto. Then, there exists a set Ψ of bounded, nondecreasing functions $\psi : Z \rightarrow \mathbb{R}$ that is pointwise bounded above and closed in the topology of pointwise convergence such that, for any $f \in \mathcal{F}$ and $\mu \in \Delta(\Omega \times S)$,*

$$\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d\mu(\omega, s) = \int_Z u(z) d(\varphi \circ F_{f, \mu})(z).$$

Since u is bounded, there exists $a, b \in \mathbb{R}$ such that $u(Z) \subset [a, b]$. The following two lemmas provide key steps in our construction.

Lemma 4. *Suppose $\varphi : [0, 1] \rightarrow [0, 1]$ is nondecreasing, concave, and onto. Define a function $W : \Delta([a, b]) \rightarrow \mathbb{R}$ by*

$$W(\eta) = \int_a^b x d(\varphi \circ F_\eta)(x),$$

where $F_\eta(x) = \eta([a, x])$ is the cumulative distribution function for the measure η . Then, there exists a set Φ of nondecreasing and concave continuous functions $\phi : [a, b] \rightarrow \mathbb{R}$ such that

$$W(\eta) = \sup_{\phi \in \Phi} \int_Z \phi(z) d\eta(z).$$

Proof. It can be shown that W is convex using similar arguments to those in Section S.2.1 of the Supplementary Material of Sarver (2018) (alternatively, see Wakker (1994) or Chatterjee and Krishna (2011)). It is also not difficult to show that W is continuous in the topology of weak convergence. Finally, since φ is concave, the function W respects second-order stochastic dominance by Theorem 2 in Yaari (1987).⁴¹ In light of these conditions, we can apply Proposition 1 from Sarver (2018) to obtain a set Φ with the claimed properties. ■

Lemma 5. *Fix a set Ψ of functions $\psi : Z \rightarrow [-\infty, \infty)$ that is pointwise bounded above. Then, for any $f \in \mathcal{F}$ and $\mu \in \Delta(\Omega \times S)$,*

$$\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d\mu(\omega, s) = \sup_{\psi \in \text{cl}(\Psi)} \int_{\Omega \times S} \psi(f(\omega, s)) d\mu(\omega, s),$$

⁴¹This was also proved by Chew, Karni, and Safra (1987) in the special case where φ is Lipschitz continuous.

where the closure is taken with respect to the product topology (i.e., the topology of pointwise convergence) on $[-\infty, \infty]^Z$.

Proof. Fix any $f \in \mathcal{F}$ and $\mu \in \Delta(\Omega \times S)$. Since f is a simple act, there exists a finite partition $\mathcal{E} \subset \mathcal{B}_\Omega \otimes \mathcal{B}_S$ such that f is measurable with respect to \mathcal{E} . For each $E \in \mathcal{E}$, let $z_E = f(\omega, s)$ for some $(\omega, s) \in E$. Since f is \mathcal{E} -measurable, the value z_E does not depend on the exact choice of $(\omega, s) \in E$. Define a function $G : [-\infty, \infty]^Z \rightarrow \mathbb{R}$ by

$$G(\psi) = \int_{\Omega \times S} \psi(f(\omega, s)) d\mu(\omega, s) = \sum_{E \in \mathcal{E}} \psi(z_E) \mu(E),$$

and let $\gamma = \sup_{\psi \in \Psi} G(\psi)$. Note that γ is finite since the functions in Ψ are pointwise bounded above. Now, fix any $\psi \in \text{cl}(\Psi)$. By the definition of the closure, there exists a net $(\psi_\alpha)_{\alpha \in D}$ in Ψ that converges to ψ .⁴² Note that since $\psi_\alpha \in \Psi$ for each α , we must have $G(\psi_\alpha) \leq \gamma$. Since convergence is preserved under scalar multiples and finite sums, $\psi_\alpha \rightarrow \psi$ implies that $G(\psi_\alpha) \rightarrow G(\psi)$ and hence $G(\psi) \leq \gamma$. Since this is true for all $\psi \in \text{cl}(\Psi)$, we have

$$\sup_{\psi \in \text{cl}(\Psi)} \int_{\Omega \times S} \psi(f(\omega, s)) d\mu(\omega, s) = \sup_{\psi \in \text{cl}(\Psi)} G(\psi) = \gamma,$$

as desired. ■

Proof of Proposition 5. Take Φ as in Lemma 4 for the function φ , and let $\Psi = \{\phi \circ u : \phi \in \Phi\}$. Fix any $f \in \mathcal{F}$ and $\mu \in \Delta(\Omega \times S)$, and let η be the distribution of utility values induced by μ , f , and u . Formally,

$$\eta = \mu \circ f^{-1} \circ u^{-1} \in \Delta([a, b]).$$

Then, we have

$$\begin{aligned} \sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d\mu(\omega, s) &= \sup_{\phi \in \Phi} \int_{\Omega \times S} \phi(u(f(\omega, s))) d\mu(\omega, s) \\ &= \sup_{\phi \in \Phi} \int_a^b \phi(x) d\eta(x) && \text{(change of variables)} \\ &= \int_a^b x d(\varphi \circ F_\eta)(x) && \text{(Lemma 4)} \\ &= \int_Z u(z) d(\varphi \circ F_{f, \mu})(z). \end{aligned}$$

The last equality is essentially another application of the change of variables formula, but there are a few subtleties. One needs to show that if ν^u is the probability measure over utility values with cumulative distribution function $\varphi \circ F_\eta$ and if ν^z is the probability measure over outcomes in Z with cumulative distribution function $\varphi \circ F_{f, \mu}$, then $\nu^u = \nu^z \circ u^{-1}$. This is not true for arbitrary u , but it can be shown to hold whenever u is nondecreasing.

⁴²It is well known that the product topology on an uncountable product space cannot be completely described by sequential convergence, as such spaces are not metrizable. Hence, we must use nets.

Note that since $W(\eta) = x$ when $\eta(\{x\}) = 1$, we must have $\phi(x) \leq x$ for all $x \in [a, b]$ and $\phi \in \Phi$. Now, for any $\psi \in \Psi$ there exists $\phi \in \Phi$ such that $\psi = \phi \circ u$. Thus, $\psi(z) = \phi(u(z)) \leq b$ for all $z \in Z$, so the set Ψ is bounded above. Moreover, taking the closure of Ψ does not alter the values in the equality above by Lemma 5, so we can assume that Ψ is closed without loss of generality. ■

B.4 Proof of Proposition 2

Some basic definitions and results from functional analysis will be used frequently in this proof. If X is a Banach space, we use X^* to denote the space of all continuous linear functionals on X (the norm dual of X). For $x \in X$ and $x^* \in X^*$, we use $\langle x^*, x \rangle$ to denote the duality pairing $x^*(x)$.

Given a function $F : X \rightarrow (-\infty, \infty]$, the *effective domain* of F is the set

$$\text{dom}(F) = \{x \in X : F(x) < \infty\}.$$

The function F is *proper* if $\text{dom}(F) \neq \emptyset$, that is, if it is not identically equal to ∞ . The (*Fenchel*) *conjugate* of F is the function $F^* : X^* \rightarrow [-\infty, \infty]$ defined by

$$F^*(x^*) = \sup_{x \in X} [\langle x^*, x \rangle - F(x)].$$

Note that if F is proper, then $F^*(x^*) > -\infty$ for all $x^* \in X^*$. Finally, given a set $C \subset X$, we define δ_C by $\delta_C(x) = 0$ if $x \in C$ and $\delta_C(x) = \infty$ if $x \notin C$. This is the indicator function commonly used in functional analysis. Note that

$$(\delta_C)^*(x^*) = \sup_{x \in C} \langle x^*, x \rangle.$$

In this proof, we will work with the L^1 and L^∞ spaces of functions. That is, given a probability space $(\Omega, \mathcal{B}_\Omega, p)$, the space $L^1(\Omega, \mathcal{B}_\Omega, p)$ is the set of all (equivalence classes of) integrable functions, and the space $L^\infty(\Omega, \mathcal{B}_\Omega, p)$ is the set of all (equivalence classes of) essentially bounded functions. When the reference probability space is understood, we will sometimes denote these spaces simply as L^1 and L^∞ , respectively. It is a standard result that these are Banach spaces (when endowed with the L^1 and L^∞ norms, respectively) and that $(L^1)^* = L^\infty$, with the duality pairing

$$\langle X, Y \rangle = \int_{\Omega} X(\omega)Y(\omega) dp(\omega)$$

for $Y \in L^1$, $X \in L^\infty$.

Proposition 6. *Fix any probability space $(\Omega, \mathcal{B}_\Omega, p)$. Let $D_\phi(\cdot \| \cdot)$ be a ϕ -divergence, and fix any nondecreasing, convex, and lower semicontinuous function $k : \mathbb{R} \rightarrow (-\infty, \infty]$ such that $k(0) = 0$ and k is finite on some interval $(-\varepsilon, \varepsilon)$. Then, for any random variable $X \in L^\infty(\Omega, \mathcal{B}_\Omega, p)$,*

$$\inf_{q \in \Delta(\Omega)} \left[\int_{\Omega} X(\omega) dq(\omega) + k(D_\phi(q \| p)) \right] = \max_{\gamma \in \mathbb{R}} \max_{\alpha \geq 0} \int_{\Omega} \psi_{\gamma, \alpha}(X(\omega)) dp(\omega),$$

where

$$\psi_{\gamma,\alpha}(x) = \begin{cases} \gamma - \alpha\phi^*\left(\frac{\gamma-x}{\alpha}\right) - k^*(\alpha) & \text{if } \alpha > 0 \\ \gamma - \delta_{\mathbb{R}_-}(\gamma - x) - k^*(0) & \text{if } \alpha = 0. \end{cases}$$

Recall that $\delta_{\mathbb{R}_-}$ denotes the indicator function for \mathbb{R}_- , so $\delta_{\mathbb{R}_-}(t) = 0$ if $t \leq 0$ and $\delta_{\mathbb{R}_-}(t) = \infty$ if $t > 0$. Also, note that our definition of a divergence requires ϕ to be a continuous convex function mapping from \mathbb{R}_+ to \mathbb{R}_+ . However, we can treat ϕ as lower semicontinuous convex function defined on all of \mathbb{R} by taking $\phi(y) = \infty$ for $y < 0$, and hence

$$\phi^*(x) = \sup_{y \in \mathbb{R}_+} [xy - \phi(y)].$$

Proposition 2 follows as a special case of this result where the state space is $\widehat{\Omega} = \Omega \times S$, the probability measure is $p \in \Delta(\Omega \times S)$, and $X : \Omega \times S \rightarrow \mathbb{R}$ is defined by

$$X(\omega, s) = u(f(\omega, s)).$$

Note that since f is a simple act and u is real-valued, X is bounded. Thus, by Proposition 6,

$$\inf_{r \in \Delta(\Omega \times S)} \left[\int_{\Omega \times S} u(f(\omega, s)) dr(\omega, s) + k(D_\phi(r \| p)) \right] = \max_{\gamma \in \mathbb{R}} \max_{\alpha \geq 0} \int_{\Omega \times S} \psi_{\gamma,\alpha}(u(f(\omega, s))) dp(\omega, s).$$

Take Ψ to be the closure of the set

$$\{\psi_{\gamma,\alpha} \circ u : \gamma \in \mathbb{R}, \alpha \geq 0\},$$

where the closure is taken with respect to the topology of pointwise convergence on the extended reals. Then, Ψ satisfies all of the properties asserted in the statement of Proposition 2, and the arguments above together with Lemma 5 (which allows us to take the closure) establish that the equality in the statement of the proposition holds.

Therefore, all that remains is to prove Proposition 6. Our proof will be based on the following three lemmas. The first two lemmas closely parallel the proof strategy used by Ben-Tal and Teboulle (1987, Theorem 4.2) who provide a similar result for the case when $k(x) = x$, that is, when there is no transformation of the divergence term.

Lemma 6. *Fix any probability space $(\Omega, \mathcal{B}_\Omega, p)$. Let $H : L^1 \rightarrow (-\infty, \infty]$ be a convex and lower semicontinuous function, and suppose there exist $\alpha < 1 < \beta$ such that $Y \in L^1$ and $\alpha \leq Y(\omega) \leq \beta$ for all $\omega \in \Omega$ implies $H(Y) < \infty$. Then, for any $X \in L^\infty$,*

$$\inf_{\substack{Y \in L^1: \\ \int Y(\omega) dp(\omega) = 1}} \left[\int_{\Omega} X(\omega) Y(\omega) dp(\omega) + H(Y) \right] = \max_{\gamma \in \mathbb{R}} [\gamma - H^*(\gamma - X)]$$

Proof. The proof of this result replicates the first steps in the proof of Theorem 4.2 in Ben-Tal and Teboulle (2007), but we include it for completeness. Denote by v the value of the left side of the

equation in the statement of the lemma:

$$v \equiv \inf_{\substack{Y \in L^1: \\ \int Y(\omega) dp(\omega) = 1}} \left[\int_{\Omega} X(\omega)Y(\omega) dp(\omega) + H(Y) \right].$$

The Lagrangian dual of this convex minimization problem is given by

$$\begin{aligned} w &\equiv \sup_{\gamma \in \mathbb{R}} \inf_{Y \in L^1} \left[\int_{\Omega} X(\omega)Y(\omega) dp(\omega) + H(Y) + \gamma \left(1 - \int_{\Omega} Y(\omega) dp(\omega) \right) \right] \\ &= \sup_{\gamma \in \mathbb{R}} \left[\gamma + \inf_{Y \in L^1} \left(H(Y) + \int_{\Omega} (X(\omega) - \gamma)Y(\omega) dp(\omega) \right) \right] \\ &= \sup_{\gamma \in \mathbb{R}} \left[\gamma - \sup_{Y \in L^1} \left(\int_{\Omega} (\gamma - X(\omega))Y(\omega) dp(\omega) - H(Y) \right) \right] \\ &= \sup_{\gamma \in \mathbb{R}} [\gamma - H^*(\gamma - X)]. \end{aligned}$$

It remains only to show that $v = w$, that is, there is no duality gap. The convex duality result in Corollary 4.8 of Borwein and Lewis (1992) shows that there is no duality gap and there is attainment of a solution in the dual problem if the following constraint qualification condition is satisfied:⁴³

(CQ) There exist $\alpha < \beta$ such that $\alpha \leq Y(\omega) \leq \beta$ implies $H(Y) < \infty$, and there exists some $Y \in L^1$ with $\alpha < Y(\omega) < \beta$ that satisfies the constraint $\int_{\Omega} Y(\omega) dp(\omega) = 1$.

Given the assumptions in the statement of the lemma, this condition is satisfied by taking Y identically equal to 1. This completes the proof. \blacksquare

Lemma 7. Fix any probability space $(\Omega, \mathcal{B}_{\Omega}, p)$, and fix any proper convex and lower semicontinuous function $\phi : \mathbb{R} \rightarrow (-\infty, \infty]$. Define a functional $J : L^1 \rightarrow (-\infty, \infty]$ by

$$J(Y) = \int_{\Omega} \phi(Y(\omega)) dp(\omega).$$

Then, J is a proper convex and lower semicontinuous functional, and the Fenchel conjugate $J^* : L^{\infty} \rightarrow (-\infty, \infty]$ of J is given by

$$J^*(X) = \int_{\Omega} \phi^*(X(\omega)) dp(\omega).$$

Proof. See the corollary to Theorem 2 in Rockafellar (1968). \blacksquare

⁴³Borwein and Lewis (1992) define the *quasi relative interior* of a set C to be the set of all points $x \in C$ such that the closure of the cone generated by $C - x$ is a subspace. In the context of our minimization problem, their constraint qualification condition requires that there is a function Y in the quasi relative interior of the set $\text{dom}(H) \equiv \{Y \in L^1 : H(Y) < \infty\}$ that satisfies the constraint $\int_{\Omega} Y(\omega) dp(\omega) = 1$. It can be shown that if $\{Y \in L^1 : \alpha \leq Y \leq \beta\} \subset \text{dom}(H)$ then any $Y \in L^1$ with $\alpha < Y(\omega) < \beta$ is in the quasi relative interior of $\text{dom}(H)$ (see Example 3.11(i) in Borwein and Lewis (1992)).

Fix any proper convex and lower semicontinuous function $\phi : \mathbb{R} \rightarrow (-\infty, \infty]$ that is finite on an open interval containing 1. Then, defining J as in Lemma 7 and setting $H = J$ in Lemma 6, we obtain the following dual formula:

$$\inf_{\substack{Y \in L^1: \\ \int Y(\omega) dp(\omega) = 1}} \left[\int_{\Omega} X(\omega)Y(\omega) dp(\omega) + J(Y) \right] = \max_{\gamma \in \mathbb{R}} \int_{\Omega} [\gamma - \phi^*(\gamma - X(\omega))] dp(\omega).$$

This is precisely Theorem 4.2 in Ben-Tal and Teboulle (2007). To extend their result to $H = k \circ J$, we need the following lemma.

Lemma 8. *Fix any probability space $(\Omega, \mathcal{B}_{\Omega}, p)$, and fix any continuous convex function $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ that satisfies $\phi(1) = 0$. Also, fix any nondecreasing, convex, and lower semicontinuous function $k : \mathbb{R} \rightarrow (-\infty, \infty]$ such that k is finite on some interval $(-\varepsilon, \varepsilon)$. Define $J : L^1 \rightarrow (-\infty, \infty]$ by*

$$J(Y) = \int_{\Omega} \phi(Y(\omega)) dp(\omega),$$

and define $H : L^1 \rightarrow (-\infty, \infty]$ by $H = k \circ J$. Then, for any $X \in L^{\infty}$,

$$H^*(X) = \min_{\alpha \geq 0} [(\alpha J)^*(X) + k^*(\alpha)], \quad (13)$$

where

$$(\alpha J)^*(X) = \begin{cases} \int_{\Omega} \alpha \phi^*\left(\frac{X(\omega)}{\alpha}\right) dp(\omega) & \text{if } \alpha > 0 \\ \int_{\Omega} \delta_{\mathbb{R}_-}(X(\omega)) dp(\omega) & \text{if } \alpha = 0. \end{cases}$$

Proof. To obtain the formula for the conjugate of the composition of two functions, we appeal to Theorem 2 of Hiriart-Urruty (2006):⁴⁴ Since k and J are both lower semicontinuous and convex, k is nondecreasing, and there exists a function $Y \in L^1$ such that $J(Y) \in \text{int}(\text{dom}(k))$ (namely, Y identically equal to 1), his theorem implies that the Fenchel conjugate of $k \circ J$ is given by Equation (13), when one sets $(0J) = \delta_{\text{dom}(J)}$. For $\alpha > 0$, we therefore have

$$(\alpha J)^*(X) = \int_{\Omega} (\alpha \phi)^*(X(\omega)) dp(\omega) = \int_{\Omega} \alpha \phi^*\left(\frac{X(\omega)}{\alpha}\right) dp(\omega),$$

where the first equality follows from Lemma 7 and the second equality follows directly from the definition of the conjugate.

It remains only to establish the formula for $(0J)^*$. By the definition of the conjugate,

$$(0J)^*(X) = \sup_{Y \in L^1} [\langle X, Y \rangle - \delta_{\text{dom}(J)}(Y)] = \sup_{Y \in \text{dom}(J)} \int_{\Omega} X(\omega)Y(\omega) dp(\omega).$$

Now, fix any $X \in L^{\infty}$ and let $E = \{\omega \in \Omega : X(\omega) > 0\}$. We will show that if $p(E) = 0$ then $(0J)^*(X) = 0$, and if $p(E) > 0$ then $(0J)^*(X) = \infty$. Consider first the case of $p(E) = 0$. Recall

⁴⁴Hiriart-Urruty (2006) provides a concise treatment of this problem, but earlier, more general results about conjugates of compositions of convex functions exist, e.g., Kutateladze (1979, Theorem 3.7.1) or Combari, Laghdir, and Thibault (1996, Theorem 3.4(ii)).

that since ϕ is defined on \mathbb{R}_+ , we can treat it as a lower semicontinuous function on all of \mathbb{R} such that $\phi(y) = \infty$ for $y < 0$. Therefore, if the set of all ω such that $Y(\omega) < 0$ has positive probability under p , then $J(Y) = \infty$. Thus, $\text{dom}(J)$ includes only functions Y that are nonnegative almost surely, so for any $Y \in \text{dom}(J)$ and $X \leq 0$, $\langle X, Y \rangle \leq 0$. Therefore, when $p(E) = 0$, the supremum of $\langle X, Y \rangle$ over $Y \in \text{dom}(J)$ is attained by $Y = 0$, and $(0J)^*(X) = 0$. Next, consider the case of $p(E) > 0$. Define Y_n by $Y_n(\omega) = n$ for $\omega \in E$ and $Y_n(\omega) = 0$ for $\omega \notin E$. Since ϕ is finite and continuous on \mathbb{R}_+ , we have $Y_n \in \text{dom}(J)$ for all n . Note that

$$\int_E X(\omega) dp(\omega) > 0,$$

and therefore

$$\langle X, Y_n \rangle = n \int_E X(\omega) dp(\omega) \rightarrow \infty$$

as $n \rightarrow \infty$. Thus, $(0J)^*(X) = \infty$.

We have shown that $(0J)^*(X) = 0$ if $X \leq 0$ a.s., and $(0J)^*(X) = \infty$ otherwise. Recall that the indicator function $\delta_{\mathbb{R}_-}$ satisfies $\delta_{\mathbb{R}_-}(x) = 0$ if $x \leq 0$ and $\delta_{\mathbb{R}_-}(x) = \infty$ if $x > 0$. Therefore, we have

$$(0J)^*(X) = \int_{\Omega} \delta_{\mathbb{R}_-}(X(\omega)) dp(\omega).$$

This completes the proof. ■

Proof of Proposition 6. Note that $D_{\phi}(q \| p) = \infty$ whenever q is not absolutely continuous with respect to p . Thus, we can restrict attention to $q \ll p$, and we can therefore express the divergence using Radon–Nikodym derivatives $Y = \frac{dq}{dp} \in L^1(\Omega, \mathcal{B}_{\Omega}, p)$:

$$\begin{aligned} & \inf_{q \in \Delta(\Omega)} \left[\int_{\Omega} X(\omega) dq(\omega) + k(D_{\phi}(q \| p)) \right] \\ &= \inf_{q \ll p} \left[\int_{\Omega} X(\omega) \frac{dq}{dp}(\omega) dp(\omega) + k \left(\int_{\Omega} \phi \left(\frac{dq}{dp}(\omega) \right) dp(\omega) \right) \right] \\ &= \inf_{\substack{Y \in L^1: \\ \int Y(\omega) dp(\omega) = 1}} \left[\int_{\Omega} X(\omega) Y(\omega) dp(\omega) + k \left(\int_{\Omega} \phi(Y(\omega)) dp(\omega) \right) \right]. \end{aligned}$$

Note that for $Y \in L^1$ to be a Radon–Nikodym derivative, we must have $\int_{\Omega} Y(\omega) dp(\omega) = 1$ and $Y \geq 0$ a.s. The first constraint is stated explicitly in the equation above, and since $\phi(y) = \infty$ for $y < 0$, the second constraint becomes superfluous.

As before, define $J : L^1 \rightarrow (-\infty, \infty]$ by

$$J(Y) = \int_{\Omega} \phi(Y(\omega)) dp(\omega),$$

and define $H : L^1 \rightarrow (\infty, \infty]$ by $H = k \circ J$. Note that J is convex and lower semicontinuous by Lemma 7, and therefore H is convex and lower semicontinuous given our assumptions on k . We also assumed that there is an interval $(-\varepsilon, \varepsilon)$ on which k is finite. Since $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous

and satisfies $\phi(1) = 0$, there exists $\alpha < 1 < \beta$ such that $\alpha \leq y \leq \beta$ implies $0 \leq \phi(y) < \varepsilon$. Thus, $\alpha \leq Y(\omega) \leq \beta$ for all $\omega \in \Omega$ implies $0 \leq J(Y) < \varepsilon$ and hence $H(Y) < \infty$. Therefore,

$$\begin{aligned} & \inf_{\substack{Y \in L^1: \\ \int Y(\omega) dp(\omega) = 1}} \left[\int_{\Omega} X(\omega)Y(\omega) dp(\omega) + k\left(\int_{\Omega} \phi(Y(\omega)) dp(\omega)\right) \right] \\ &= \max_{\gamma \in \mathbb{R}} [\gamma - H^*(\gamma - X)] \\ &= \max_{\gamma \in \mathbb{R}} \max_{\alpha \geq 0} [\gamma - (\alpha J)^*(\gamma - X) - k^*(\alpha)], \end{aligned}$$

where the first equality follows from Lemma 6 and the second equality follows from Lemma 8. Then, using the formula for $(\alpha J)^*$ from Lemma 8, we have that for any $X \in L^\infty$, $\gamma \in \mathbb{R}$, and $\alpha \geq 0$,

$$\begin{aligned} \gamma - (\alpha J)^*(\gamma - X) - k^*(\alpha) &= \begin{cases} \gamma - \int_{\Omega} \alpha \phi^*\left(\frac{\gamma - X(\omega)}{\alpha}\right) dp(\omega) - k^*(\alpha) & \text{if } \alpha > 0 \\ \gamma - \int_{\Omega} \delta_{\mathbb{R}_-}(\gamma - X(\omega)) dp(\omega) - k^*(0) & \text{if } \alpha = 0 \end{cases} \\ &= \int_{\Omega} \psi_{\gamma, \alpha}(X(\omega)) dp(\omega), \end{aligned}$$

where $\psi_{\gamma, \alpha}(x)$ is defined as in the statement of the proposition. This completes the proof. \blacksquare

References

- Agranov, M. and P. Ortoleva (2017): “Stochastic Choice and Preferences for Randomization,” *Journal of Political Economy*, 125, 40–68.
- Al-Najjar, N. I. and J. Weinstein (2009): “The Ambiguity Aversion Literature: A Critical Assessment,” *Economics & Philosophy*, 25, 249–284.
- Athreya, K. B. and P. E. Ney (1972): *Branching Processes*, New York: Springer-Verlag.
- Ben-Tal, A. and M. Teboulle (1987): “Penalty Functions and Duality in Stochastic Programming via φ -Divergence Functionals,” *Mathematics of Operations Research*, 12, 224–240.
- Ben-Tal, A. and M. Teboulle (2007): “An Old-New Concept of Convex Risk Measures: The Optimized Certainty Equivalent,” *Mathematical Finance*, 17, 449–476.
- Bergstrom, T. C. (2014): “On the Evolution of Hoarding, Risk-Taking, and Wealth Distribution in Nonhuman and Human Populations,” *Proceedings of the National Academy of Sciences*, 111, 10860–10867.
- Bianchi, F., C. L. Ilut, and M. Schneider (2018): “Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle,” *Review of Economic Studies*, 85, 810–854.
- Bleichrodt, H., J. Eichberger, S. Grant, D. Kelsey, and C. Li (2020): “Testing Dynamic Consistency and Consequentialism under Ambiguity,” working paper.
- Borwein, J. M. and A. S. Lewis (1992): “Partially Finite Convex Programming, Part I: Quasi Relative Interiors and Duality Theory,” *Mathematical Programming*, 57, 15–48.

- Cerreia-Vioglio, S., D. Dillenberger, P. Ortoleva, and G. Riella (2018): “Deliberately Stochastic,” working paper.
- Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011): “Uncertainty Averse Preferences,” *Journal of Economic Theory*, 146, 1275–1330.
- Chapman, J., M. Dean, P. Ortoleva, E. Snowberg, and C. Camerer (2019): “Econographics,” working paper.
- Chateauneuf, A. and J. H. Faro (2009): “Ambiguity Through Confidence Functions,” *Journal of Mathematical Economics*, 45, 535–558.
- Chateauneuf, A. and J. H. Faro (2012): “On the Confidence Preferences Model,” *Fuzzy Sets and Systems*, 188, 1–15.
- Chatterjee, K. and R. V. Krishna (2011): “A Nonsmooth Approach to Nonexpected Utility Theory under Risk,” *Mathematical Social Sciences*, 62, 166–175.
- Chetty, R. and A. Szeidl (2007): “Consumption Commitments and Risk Preferences,” *Quarterly Journal of Economics*, 122, 831–877.
- Chetty, R. and A. Szeidl (2016): “Consumption Commitments and Habit Formation,” *Econometrica*, 84, 855–890.
- Chew, S. H., E. Karni, and Z. Safra (1987): “Risk Aversion in the Theory of Expected Utility with Rank Dependent Probabilities,” *Journal of Economic Theory*, 42, 370–381.
- Combari, C., M. Laghdir, and L. Thibault (1996): “A Note on Subdifferentials of Convex Composite Functionals,” *Archiv der Mathematik*, 67, 239–252.
- Cooper, W. S. and R. H. Kaplan (1982): “Adaptive ‘Coin-Flipping’: a Decision-Theoretic Examination of Natural Selection for Random Individual Variation,” *Journal of Theoretical Biology*, 94, 135–151.
- Cover, T. M. and J. A. Thomas (2006): *Elements of Information Theory*, 2nd ed., Hoboken, New Jersey: John Wiley & Sons.
- Cubitt, R., C. Starmer, and R. Sugden (1998): “Dynamic Choice and the Common Ratio Effect: An Experimental Investigation,” *The Economic Journal*, 108, 1362–1380.
- Dean, M. and P. Ortoleva (2017): “Allais, Ellsberg, and Preferences for Hedging,” *Theoretical Economics*, 12, 377–424.
- Dominiak, A., P. Duersch, and J.-P. Lefort (2012): “A Dynamic Ellsberg Urn Experiment,” *Games and Economic Behavior*, 75, 625–638.
- Dominiak, A. and W. Schnedler (2011): “Attitudes Toward Uncertainty and Randomization: An Experimental Study,” *Economic Theory*, 48, 289–312.
- Dow, J. and S. R. Werlang (1992): “Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio,” *Econometrica*, 60, 197–204.
- Dudley, R. M. (2002): *Real Analysis and Probability*, 2nd ed., Cambridge, United Kingdom: Cambridge University Press.

- Dupuis, P. and R. S. Ellis (1997): *A Weak Convergence Approach to the Theory of Large Deviations*, New York: John Wiley & Sons.
- Ellsberg, D. (1961): “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics*, 75, 643–669.
- Epstein, L. G. and Y. Halevy (2020): “Hard-to-Interpret Signals,” working paper.
- Epstein, L. G. and M. Le Breton (1993): “Dynamically Consistent Beliefs Must Be Bayesian,” *Journal of Economic Theory*, 61, 1–22.
- Epstein, L. G. and M. Schneider (2003): “Recursive Multiple-Priors,” *Journal of Economic Theory*, 113, 1–31.
- Epstein, L. G. and T. Wang (1994): “Intertemporal Asset Pricing under Knightian Uncertainty,” *Econometrica*, 62, 283–322.
- Ergin, H. and F. Gul (2009): “A Theory of Subjective Compound Lotteries,” *Journal of Economic Theory*, 144, 899–929.
- Ergin, H. and T. Sarver (2015): “Hidden Actions and Preferences for Timing of Resolution of Uncertainty,” *Theoretical Economics*, 10, 489–541.
- Flück, M. (2006): “Functional, Structural and Molecular Plasticity of Mammalian Skeletal Muscle in Response to Exercise Stimuli,” *Journal of Experimental Biology*, 209, 2239–2248.
- Gabaix, X. and D. Laibson (2001): “The 6D Bias and the Equity-Premium Puzzle,” *NBER Macroeconomics Annual*, 16, 257–312.
- Ghirardato, P. (2002): “Revisiting Savage in a Conditional World,” *Economic Theory*, 20, 83–92.
- Gilboa, I. and D. Schmeidler (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- Grossman, S. J. and G. Laroque (1990): “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods,” *Econometrica*, 58, 25–51.
- Halevy, Y. (2007): “Ellsberg Revisited: An Experimental Study,” *Econometrica*, 75, 503–536.
- Halevy, Y. and V. Feltkamp (2005): “A Bayesian Approach to Uncertainty Aversion,” *Review of Economic Studies*, 72, 449–466.
- Hammond, P. J. (1988): “Consequentialist Foundations for Expected Utility,” *Theory and Decision*, 25, 25–78.
- Hammond, P. J. (1989): “Consistent Plans, Consequentialism, and Expected Utility,” *Econometrica*, 57, 1445–1449.
- Hanany, E. and P. Klibanoff (2007): “Updating Preferences with Multiple Priors,” *Theoretical Economics*, 2, 261–298.
- Hanany, E. and P. Klibanoff (2009): “Updating Ambiguity Averse Preferences,” *The B.E. Journal of Theoretical Economics*, 9.
- Hansen, L. P. and T. J. Sargent (2001): “Robust Control and Model Uncertainty,” *American Economic Review*, 91, 60–66.

- Hansen, L. P. and T. J. Sargent (2008): *Robustness*, Princeton University Press.
- Hey, J. D. and L. Panaccione (2011): “Dynamic Decision Making: What Do People Do?,” *Journal of Risk and Uncertainty*, 42, 85–123.
- Hey, J. D. and M. Paradiso (2006): “Preferences over Temporal Frames in Dynamic Decision Problems: An Experimental Investigation,” *The Manchester School*, 74, 123–137.
- Hiriart-Urruty, J.-B. (2006): “A Note on the Legendre-Fenchel Transform of Convex Composite Functions,” in *Nonsmooth Mechanics and Analysis*, ed. by P. Alart, O. Maisonneuve, and R. T. Rockafellar, Springer, 35–46.
- Ilut, C. L. and M. Schneider (2014): “Ambiguous Business Cycles,” *American Economic Review*, 104, 2368–99.
- Izhakian, Y. (2017): “Expected Utility with Uncertain Probabilities Theory,” *Journal of Mathematical Economics*, 69, 91–103.
- Ke, S. and Q. Zhang (2020): “Randomization and Ambiguity Aversion,” *Econometrica*, 88, 1159–1195.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005): “A Smooth Model of Decision Making Under Ambiguity,” *Econometrica*, 73, 1849–1892.
- Kreps, D. M. and E. L. Porteus (1979): “Temporal von Neumann-Morgenstern and Induced Preferences,” *Journal of Economic Theory*, 20, 81–109.
- Kutateladze, S. S. (1979): “Convex Operators,” *Russian Mathematical Surveys*, 34, 181–214.
- Kuzmics, C. (2017): “Abraham Wald’s Complete Class Theorem and Knightian Uncertainty,” *Games and Economic Behavior*, 104, 666–673.
- Maccheroni, F., M. Marinacci, and A. Rustichini (2006): “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica*, 74, 1447–1498.
- Machina, M. J. (1984): “Temporal Risk and the Nature of Induced Preferences,” *Journal of Economic Theory*, 33, 199–231.
- Machina, M. J. (1985): “Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries,” *The Economic Journal*, 95, 575–594.
- Machina, M. J. (1989): “Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty,” *Journal of Economic Literature*, 27, 1622–1668.
- Machina, M. J. and D. Schmeidler (1992): “A More Robust Definition of Subjective Probability,” *Econometrica*, 60, 745–780.
- Marinacci, M. (2015): “Model Uncertainty,” *Journal of the European Economic Association*, 13, 1022–1100.
- Nau, R. F. (2006): “Uncertainty Aversion with Second-Order Utilities and Probabilities,” *Management Science*, 52, 136–145.
- von Neumann, J. (1928): “Zur Theorie der Gesellschaftsspiele,” *Mathematische Annalen*, 100, 295–320.

- Nöldeke, G. and L. Samuelson (2005): “Information-Based Relative Consumption Effects: Correction,” *Econometrica*, 73, 1383–1387.
- Piersma, T. and J. Drent (2003): “Phenotypic Flexibility and the Evolution of Organismal Design,” *Trends in Ecology & Evolution*, 18, 228–233.
- Raiffa, H. (1961): “Risk, Ambiguity, and the Savage Axioms: Comment,” *Quarterly Journal of Economics*, 75, 690–694.
- Robatto, R. and B. Szentes (2017): “On the Biological Foundation of Risk Preferences,” *Journal of Economic Theory*, 172, 410–422.
- Robson, A. J. (1996): “A Biological Basis for Expected and Non-Expected Utility,” *Journal of Economic Theory*, 68, 397–424.
- Robson, A. J. and L. Samuelson (2011): “The Evolutionary Foundations of Preferences,” in *Handbook of Social Economics*, vol. 1, Elsevier, 221–310.
- Robson, A. J. and L. Samuelson (2019): “Evolved Attitudes to Idiosyncratic and Aggregate Risk in Age-Structured Populations,” *Journal of Economic Theory*, 181, 44–81.
- Rockafellar, R. T. (1968): “Integrals Which Are Convex Functionals,” *Pacific Journal of Mathematics*, 24, 525–539.
- Saito, K. (2015): “Preferences for Flexibility and Randomization under Uncertainty,” *American Economic Review*, 105, 1246–1271.
- Sarver, T. (2018): “Dynamic Mixture-Averse Preferences,” *Econometrica*, 86, 1347–1382.
- Segal, U. (1987): “The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach,” *International Economic Review*, 28, 175–202.
- Shishkin, D. and P. Ortoleva (2020): “Ambiguous Information and Dilation: An Experiment,” working paper.
- Siniscalchi, M. (2009): “Two Out Of Three Ain’t Bad: A Comment On ‘The Ambiguity Aversion Literature: A Critical Assessment’,” *Economics & Philosophy*, 25, 335–356.
- Siniscalchi, M. (2011): “Dynamic Choice under Ambiguity,” *Theoretical Economics*, 6, 379–421.
- Sion, M. (1958): “On General Minimax Theorems,” *Pacific Journal of Mathematics*, 8, 171–176.
- Strzalecki, T. (2011): “Axiomatic Foundations of Multiplier Preferences,” *Econometrica*, 79, 47–73.
- Tuy, H. (2004): “Minimax Theorems Revisited,” *Acta Mathematica Vietnamica*, 29, 217–229.
- Wakker, P. (1994): “Separating Marginal Utility and Probabilistic Risk Aversion,” *Theory and Decision*, 36, 1–44.
- Wikelski, M. and C. Thom (2000): “Marine Iguanas Shrink to Survive El Niño,” *Nature*, 403, 37.
- Wright, P. A. and A. J. Turko (2016): “Amphibious Fishes: Evolution and Phenotypic Plasticity,” *Journal of Experimental Biology*, 219, 2245–2259.
- Yaari, M. E. (1987): “The Dual Theory of Choice under Risk,” *Econometrica*, 55, 95–115.

SUPPLEMENTARY APPENDIX

Abstract

In this supplementary appendix, we explore several alternative assumptions and extensions of the analysis in the main text. Section S1 shows that the specifics of our representation change when adaptation is slower and must be undertaken before the realization of the signal, yet evolutionarily optimal preferences remain dynamically consistent. Section S2 shows that responding to private signals results in idiosyncratic randomization in choice that can lessen, or in some cases even eliminate, the need for self-randomization by members of the population of a genotype. Section S3 examines how the optimal responses of a genotype to public and private signals differ. Proofs of results in this online appendix are contained in Section S4, and the proofs of Propositions 3 and 4 and Lemmas 2 and 3 from Appendix B.2 of the main paper are contained in Section S5.

S1 Adaptation Before Information

We assume throughout that signals resolve prior to the choice of act. So far, we further assumed after-signal adaptation, where the choice of fitness function also happens after the realization of a signal, reflecting the implicit assumption either that adaptation via selection of the hidden action can be undertaken rapidly or that signals arrive sufficiently early to allow time for such adaptation. We now consider the alternative of before-signal adaptation, where adaptation of the fitness function through the choice of hidden action is still fast enough to take into account the action plan, but too slow to react to the realization of a signal and the subsequent final choice of act. This alternative timing was discussed briefly in Section 5.2 in the context of rank-dependent utility. In this section, we provide the formal results behind that discussion.

For ease of illustration, we will focus on the long-run growth rates from deterministic action and adaptation plans. In the case of no common uncertainty (as in Section 5), deterministic plans will be optimal and our analysis is therefore sufficient for determining optimal choice. However, the reader should keep in mind in the case of common uncertainty, self-randomization may be optimal; extending our analysis accordingly is relatively straightforward and not central to the intuitions of this section.

Formally, the signal σ arrives after the choice of fitness function ψ , as illustrated in Figure S1. From the ex ante perspective, the individual thus selects an action plan $(f_\sigma)_{\sigma \in \Sigma}$




Figure S1: Within-period timeline: before-signal adaptation

together with a fixed fitness function ψ , which achieves a fitness of $\psi(f_\sigma(\omega, s))$ after the realization of (ω, s, σ) . Clearly, the growth rate will be lower than under after-signal adaptation, since fitness functions can no longer be optimized based on the signal realization. This will also generate subtle but important differences in the representation of evolutionarily optimal preferences over action plans. The following characterization follows from identical logic to Theorem 1. We therefore omit the proof.

Theorem S1. *Suppose Ψ and μ are fixed, and individuals can engage in slow (before-signal) adaptation. If the fitness function $\psi \in \Psi$ is chosen optimally, then the long-run growth rate of a genotype from choosing the deterministic action plan $(f_\sigma)_{\sigma \in \Sigma} \in \mathcal{F}^\Sigma$ in every period is*

$$V((f_\sigma)_{\sigma \in \Sigma}) = \sup_{\psi \in \Psi} \int_{\Omega} \ln \left(\int_{S \times \Sigma} \psi(f_\sigma(\omega, s)) d\mu(s, \sigma | \omega) \right) d\mu(\omega). \quad (\text{S1})$$

The optimal fitness function ψ^* for plan $(f_\sigma)_{\sigma \in \Sigma}$ satisfies⁴⁵

$$\psi^* \in \arg \max_{\psi \in \Psi} \int_{\Omega} \ln \left(\int_{S \times \Sigma} \psi(f_\sigma(\omega, s)) d\mu(s, \sigma | \omega) \right) d\mu(\omega), \quad (\text{S2})$$

and if plan $(f_\sigma)_{\sigma \in \Sigma}$ is followed for all signals $\sigma \neq \bar{\sigma}$ and the ex ante choice of fitness function is ψ^* , then the long-run growth from choosing g following $\bar{\sigma}$ is

$$V(g | \bar{\sigma}, (f_\sigma)_{\sigma \in \Sigma}, \psi^*) = \int_{\Omega} \ln \left(\mu(\bar{\sigma} | \omega) \int_S \psi^*(g(\omega, s)) d\mu(s | \omega, \bar{\sigma}) \right. \\ \left. + \int_{S \times \Sigma \setminus \{\bar{\sigma}\}} \psi^*(f_\sigma(\omega, s)) d\mu(s, \sigma | \omega) \right) d\mu(\omega).$$

The preferences that maximize these ex ante and ex post long-run growth rates are dynamically consistent.

Ex post adaptive preferences after learning signal $\bar{\sigma}$ now have to take into account not

⁴⁵We directly assume for this result that the optimal fitness function ψ^* exists for each plan $(f_\sigma)_{\sigma \in \Sigma}$. Alternatively, one could impose additional assumptions directly on the set Ψ to ensure that this is the case; for example, requiring that Ψ be compact in the topology of pointwise convergence would guarantee the existence of an optimal fitness function.

only the plan $(f_\sigma)_{\sigma \in \Sigma}$, but also the fitness function ψ^* , which is given at the time of choosing an act, as it was chosen optimally in conjunction with $(f_\sigma)_{\sigma \in \Sigma}$ prior to the realization of $\bar{\sigma}$. When Equation (S2) uniquely pins down ψ^* , ex post preferences are fully determined by $\bar{\sigma}$ and $(f_\sigma)_{\sigma \in \Sigma}$ alone, and so can be derived from ex ante preferences.

Consider three plans $(f_\sigma)_{\sigma \in \Sigma}$, $(g_\sigma)_{\sigma \in \Sigma}$, and $(h_\sigma)_{\sigma \in \Sigma}$ such that $f_\sigma = g_\sigma = h_\sigma$ for all $\sigma \neq \bar{\sigma}$. Suppose $(f_\sigma)_{\sigma \in \Sigma}$ is strictly optimal ex ante, and suppose ψ^* is the corresponding uniquely optimal fitness function. Since ex ante adaptive preferences incorporate the optimal choice of ψ while ex post preferences take ψ^* as given, it is possible to have $(g_\sigma)_{\sigma \in \Sigma} \succ (h_\sigma)_{\sigma \in \Sigma}$ ex ante and $g_{\bar{\sigma}} \prec_{\bar{\sigma}, (f_\sigma)} h_{\bar{\sigma}}$ ex post. In other words, the ranking of two *suboptimal* plans can change ex post. This is not a violation of our notion of dynamic consistency, which only requires no deviations from the *optimal* plan, and hence only applies when comparing $(f_\sigma)_{\sigma \in \Sigma}$ to the other plans. However, it does violate stronger notions commonly found in the literature, for instance, the definitions found in Machina and Schmeidler (1992) and Epstein and Le Breton (1993).⁴⁶ The following example illustrates that those violations do not depend on the arrival of actual information, but only on the fact that ex ante preferences are elicited before the commitment to a particular ψ , while ex post preferences apply after ψ is chosen.

Example S1. Let $S = \{s, s'\}$, $\Omega = \{\omega\}$, $\Sigma = \{\sigma\}$, $\mu(s) = \mu(s') = 1/2$, and $\Psi = \{\psi_1, \psi_2\}$ where $\psi_1(x) = x$ and $\psi_2(x) = x^{1/2}$. That is, there is no common uncertainty and only one uninformative signal. Consider the acts $f = (4, 4)$, $g = (1/25, 1/25)$, and $h = (0, 1/9)$. The following table lists these acts and displays their values under ψ_1 and ψ_2 , respectively:

	s	s'	$V(\cdot \psi_1)$	$V(\cdot \psi_2)$
f	4	4	4	2
g	$\frac{1}{25}$	$\frac{1}{25}$	$\frac{1}{25}$	$\frac{1}{5}$
h	0	$\frac{1}{9}$	$\frac{1}{18}$	$\frac{1}{6}$

Ex ante, each act is evaluated under the optimal ψ , so that $V(f) = 4 > V(g) = 1/5 > V(h) = 1/6$, or $f \succ g \succ h$. However, $V(h|\psi_1) = 1/18 > 1/25 = V(g|\psi_1)$. For the optimal plan f with optimal fitness function $\psi^* = \psi_1$, this means $h \succ_{\sigma, f} g$.

The special case of rank-dependent expected utility serves well to demonstrate the importance of the timing of adaptation.

Corollary S1 (RDU with Before-Signal Adaptation). Suppose $\Omega = \{\omega\}$ and $Z \subset \mathbb{R}$. Fix μ , and fix any bounded nondecreasing function $u : Z \rightarrow \mathbb{R}$ and any function $\varphi : [0, 1] \rightarrow [0, 1]$

⁴⁶Preferences in the case of after-signal adaptation that we considered in the main text will satisfy this stronger notion of dynamic consistency: For plans $(f_\sigma)_{\sigma \in \Sigma}$, $(g_\sigma)_{\sigma \in \Sigma}$, and $(h_\sigma)_{\sigma \in \Sigma}$ such that $f_\sigma = g_\sigma = h_\sigma$ for all $\sigma \neq \bar{\sigma}$, we have $(g_\sigma)_{\sigma \in \Sigma} \succ (h_\sigma)_{\sigma \in \Sigma} \implies g_{\bar{\sigma}} \succ_{\bar{\sigma}, (f_\sigma)} h_{\bar{\sigma}}$ (and $(g_\sigma)_{\sigma \in \Sigma} \succsim (h_\sigma)_{\sigma \in \Sigma} \implies g_{\bar{\sigma}} \succsim_{\bar{\sigma}, (f_\sigma)} h_{\bar{\sigma}}$ whenever $\mu(\bar{\sigma}) > 0$). This is because for after-signal adaptation, the conditional preference $\succ_{\bar{\sigma}, (f_\sigma)}$ does not depend on $f_{\bar{\sigma}}$, only on f_σ for $\sigma \neq \bar{\sigma}$. Note that in terms of observable behavior, the two notions are typically equivalent, as choice can only reveal whether or not an individual prefers deviating from the ex ante optimal plan.

that is continuous, nondecreasing, concave, and onto. Then, there exists a set Ψ of functions $\psi : Z \rightarrow \mathbb{R}$ such that the ex ante value function V defined by Equation (S1) can be equivalently expressed as

$$V((f_\sigma)_{\sigma \in \Sigma}) = \ln \left(\int_Z u(z) d(\varphi \circ F_{(f_\sigma), \mu})(z) \right)$$

where

$$F_{(f_\sigma), \mu}(z) = \int_{S \times \Sigma} \mathbf{1}[f_\sigma(s) \leq z] d\mu(s, \sigma)$$

is the cumulative distribution function of $(f_\sigma)_{\sigma \in \Sigma}$ given μ .

According to the corollary, for before-signal adaptation, the transformation function φ affects all uncertainty, including the realization of σ . This is the model considered in the literature following Machina (1989) and is in contrast to the case of after-signal adaptation. Of course, ex post preferences will still satisfy our notion of dynamic consistency, but will now in general violate consequentialism.⁴⁷

S2 Signal Response in lieu of Self-Randomization

Recall that the motive for self-randomization in our model is to reduce the correlation of outcomes across individuals, thereby reducing the aggregate risk faced by the population. Notice that if a completely uninformative idiosyncratic signal existed, then responding to that signal would simply amount to self-randomization. In other words, an uninformative private signal is nothing more than a private randomization device. An informative signal can play a similar role in alleviating—although not perfectly—the need for self-randomization, as we illustrate in this section.

Consider a simple discrete choice setting where Ω is finite, $S = \{s\}$, and $\Psi = \{\psi\}$. Suppose that individuals have to bet on any one state $\omega \in \Omega$ and can randomize over the possible bets. When there is no information ($\Sigma = \{\sigma\}$), then for any prior with support Ω , optimal choice involves randomization that places positive probability on all available bets. However, as soon as there are even minimally informative signals, there is at least one signal for which this is no longer the case.

With slight abuse of notation, let $\rho_\sigma(\omega)$ denote the probability that an individual bets on state ω after observing σ .⁴⁸ If the state on which the individual bet realizes then their

⁴⁷As noted above, ex post preferences may also violate the slightly stronger notion of dynamic consistency considered by Machina and Schmeidler (1992), Epstein and Le Breton (1993), and much of the subsequent literature. Hanany and Klibanoff (2007) proposed a weaker definition that is similar to ours in the context of partitional learning.

⁴⁸If ψ is strictly concave, then the genotype would clearly benefit if individuals could diversify by averaging these bets to obtain an act that pays a smaller but strictly positive amount in every state. Such diversification is prohibited here, as individuals must ultimately place a bet on a single state, but individuals may nonetheless prefer to randomize over bets on different states in order to replace aggregate uncertainty with idiosyncratic.

payoff is 1; otherwise, their payoff is 0. Assume that $\psi(1) > \psi(0) \geq 0$. The long-run growth rate is now given by

$$V(\rho) = \int_{\Omega} \ln \left(\int_{\Sigma} \left(\rho_{\sigma}(\omega) \psi(1) + (1 - \rho_{\sigma}(\omega)) \psi(0) \right) d\mu(\sigma|\omega) \right) d\mu(\omega). \quad (\text{S3})$$

The following proposition shows that if the likelihood ratio between states ω and ω' is higher after signal σ than σ' , then individuals will either not bet with positive probability on state ω' following signal σ , or they will not bet with positive probability on state ω following signal σ' . Note that this result includes the possibility that the conditional probability of one of these states is much higher than that of the other following both of these signals, in which case individuals might never bet on the other state with positive probability.

Proposition S1. *Fix two states $\omega, \omega' \in \Omega$ and two signals $\sigma, \sigma' \in \Sigma$. If*

$$\mu(\omega, \sigma) \mu(\omega', \sigma') > \mu(\omega, \sigma') \mu(\omega', \sigma),$$

then $\rho_{\sigma}(\omega') = 0$ or $\rho_{\sigma'}(\omega) = 0$, or both.

The proof of Proposition S1 is in Section S4.1. In the case where the probabilities in the proposition are strictly positive, the inequality in the proposition can be written as

$$\frac{\mu(\omega|\sigma)}{\mu(\omega'|\sigma)} > \frac{\mu(\omega|\sigma')}{\mu(\omega'|\sigma')}.$$

This extreme individual reaction to information reflects not only “updating”, but also the need to reduce the correlation between individual outcomes. The following example illustrates.

Example S2. *There is an ambiguous urn in which all balls are either red or yellow, which we model by taking the common component of the state space to be $\Omega = \{r, y\}$. Suppose $\mu(r) = \mu(y) = 1/2$ and $\psi(1) = 1 > \psi(0) = 0$. As in Example 2, R and Y are the bets on a ball drawn from the urn being red or yellow, respectively, so that choice between R and Y amounts to betting on $\omega \in \Omega$. Signals in $\Sigma = \{\sigma, \sigma'\}$ are informative, as $\mu(y, \sigma) = 5/10$, $\mu(r, \sigma) = 4/10$, $\mu(y, \sigma') = 0$, and hence $\mu(r, \sigma') = 1/10$, which yields the conditional probabilities*

$$\mu(\sigma|r) = \frac{4}{5} \quad \text{and} \quad \mu(\sigma|y) = 1.$$

Let $\rho_{\sigma}(R)$ denote the probability of choosing R following signal σ , and define $\rho_{\sigma}(Y)$, $\rho_{\sigma'}(R)$, and $\rho_{\sigma'}(Y)$ similarly. Then,

$$V(\rho) = \frac{1}{2} \ln \left(\frac{4}{5} \rho_{\sigma}(R) + \frac{1}{5} \rho_{\sigma'}(R) \right) + \frac{1}{2} \ln(\rho_{\sigma}(Y)),$$

which is maximized by taking

$$\begin{aligned} \rho_\sigma(R) &= \frac{3}{8} & \rho_{\sigma'}(R) &= 1 \\ \rho_\sigma(Y) &= \frac{5}{8} & \rho_{\sigma'}(Y) &= 0. \end{aligned}$$

Thus, there is no randomization contingent on signal σ' . There is, however, randomization contingent on σ . Intuitively, since σ' is much less likely, exclusively conditioning on the two informative signals by taking $\rho_\sigma(Y) = 1$ would lead to excess correlation in outcomes across individuals.⁴⁹

S3 Public versus Private Signals

When signals are informative only about the common component, Ω , then they can either be public (so that all individuals receive the same signal) or private as in the analysis thus far (so signals are independent across individuals contingent on ω). This distinction does not arise when updating beliefs in most preference-based models of individual decision-making, but it may matter for behavior in our evolutionary model. To streamline exposition, consider acts that depend only on Ω and suppress S for the remainder of this section, and let $\Psi = \{\psi\}$.

Not surprisingly, private signals are preferred over public signals because public signals introduce correlation which is harmful to long-run growth. Formally, given a signal space Σ and a measure μ on $\Omega \times \Sigma$, let $V^{\text{Pr}}(\rho)$ denote the now familiar long-run growth rate for the action plan ρ under private signals:

$$V^{\text{Pr}}(\rho) = \int_{\Omega} \ln \left(\int_{\Sigma} \mathbb{E}_{\rho_\sigma} [\psi(f(\omega))] d\mu(\sigma|\omega) \right) d\mu(\omega).$$

Let $V^{\text{Pu}}(\rho)$ denote the growth rate for ρ under public signals:

$$V^{\text{Pu}}(\rho) = \int_{\Omega \times \Sigma} \ln \left(\mathbb{E}_{\rho_\sigma} [\psi(f(\omega))] \right) d\mu(\omega, \sigma).$$

Fix a decision problem $A = (A_\sigma)_{\sigma \in \Sigma}$, and let

$$\rho^{\text{Pr}} \in \operatorname{argmax}_{\hat{\rho} \in \mathcal{R}(A)} V^{\text{Pr}}(\hat{\rho}) \quad \text{and} \quad \rho^{\text{Pu}} \in \operatorname{argmax}_{\hat{\rho} \in \mathcal{R}(A)} V^{\text{Pu}}(\hat{\rho})$$

⁴⁹In some cases, conditioning on informative signals may completely eliminate self-randomization. In the example, if instead $\mu(y, \sigma) = 5/10$, $\mu(r, \sigma) = 2/10$, $\mu(y, \sigma') = 0$, and $\mu(r, \sigma') = 3/10$, then $\rho_\sigma(R) = 0$ and $\rho_{\sigma'}(Y) = 0$, so there is no randomization following either signal. In this case, removing residual correlation through randomization is not worth the cost of worsening the expected individual outcomes.

be optimal plans under private and public signals, respectively. Then,

$$V^{\text{Pr}}(\rho^{\text{Pr}}) \geq V^{\text{Pr}}(\rho^{\text{Pu}}) \geq V^{\text{Pu}}(\rho^{\text{Pu}}),$$

where the second inequality is strict whenever $\mathbb{E}_{\rho_{\sigma}^{\text{Pu}}}[\psi(f(\omega))]$ is not constant in σ for some $\omega \in \Omega$.

A more subtle question is how ρ^{Pu} and ρ^{Pr} differ. We already saw in Section S2 that the reaction to private signals may be extreme, because they may serve as a randomization device. To gain some intuition, note that when $\omega \in \Omega$ becomes more likely upon learning a signal $\sigma \in \Sigma$, then there must also be some signal σ' where it becomes less likely. Intuitively, when signals are private it may be possible to bet on ω under σ and against ω under σ' without creating much correlation, because both signals will be present in the population at the same time. In contrast, if the same signals are public, then the entire population receives σ or σ' at the same time, and reacting to information will lead to additional correlation in outcomes across individuals. Based on this rough intuition, we would expect there to be a stronger reaction to private information than to public information, which may provide a different perspective on the often-discussed overconfidence that agents appear to have in their private information, for instance when investing in financial markets, as in Daniel, Hirshleifer, and Subrahmanyam (1998). We now briefly discuss an illustrative example.

Application: Portfolio Choice

Let μ be a positive prior on a finite space of states and signals $\Omega \times \Sigma$. We continue to suppress S , and we assume there is a single fitness function ψ that is increasing, strictly concave, and differentiable. Consider a simple portfolio-choice problem consisting of a risk-free asset with deterministic return c and a single risky asset with return $f(\omega)$ in state ω , where f is nonconstant and $\sum_{\omega \in \Omega} \mu(\omega) f(\omega) > c$. In this domain, there will be a deterministic solution since averaging the state-dependent monetary outcomes of two acts via their portfolio weights provides a superior hedging benefit to self-randomizing over the acts whenever ψ is concave. We will therefore focus on deterministic portfolio decisions in what follows.

Suppose that each individual has unit wealth, and let the plan $(\alpha_{\sigma})_{\sigma \in \Sigma}$ specify for each signal $\sigma \in \Sigma$ the proportion $\alpha_{\sigma} \in [0, 1]$ of wealth invested in the risky asset, so that an individual holds act $f_{\sigma} = \alpha_{\sigma} f + (1 - \alpha_{\sigma})c$ upon learning σ . Holding fixed $\mu \in \Delta(\Omega \times \Sigma)$, let $(\alpha_{\sigma}^{\text{Pu}})_{\sigma \in \Sigma}$ and $(\alpha_{\sigma}^{\text{Pr}})_{\sigma \in \Sigma}$ denote the optimal portfolio plans for the case where the signals in Σ are public and private, respectively.

Proposition S2. *Let σ_* and σ^* be the signals that induce the lowest and highest investment in the risky asset under private signals, respectively, that is, $\alpha_{\sigma_*}^{\text{Pr}} \leq \alpha_{\sigma}^{\text{Pr}} \leq \alpha_{\sigma^*}^{\text{Pr}}$ for all $\sigma \in \Sigma$. If $\alpha_{\sigma_*}^{\text{Pr}} \neq \alpha_{\sigma^*}^{\text{Pr}}$, then the following must be true:*

1. $\alpha_{\sigma_*}^{\text{Pr}} < \alpha_{\sigma_*}^{\text{Pu}}$ or $\alpha_{\sigma_*}^{\text{Pr}} = \alpha_{\sigma_*}^{\text{Pu}} = 0$.

2. $\alpha_{\sigma^*}^{\text{Pu}} < \alpha_{\sigma^*}^{\text{Pr}}$ or $\alpha_{\sigma^*}^{\text{Pu}} = \alpha_{\sigma^*}^{\text{Pr}} = 1$.

In particular, when there are only two signals, reaction to private signals is unambiguously stronger than to public signals in the sense that asset holdings react more to the signal realization. The proof of Proposition S2 is contained in Section S4.2.

S4 Proofs of Results in the Online Appendix

S4.1 Proof of Proposition S1

As in Section S2, with slight abuse of notation let $\rho_\sigma(\omega)$ denote the probability that an individual bets on state ω following signal σ . Let $\bar{\rho}(\omega)$ denote the probability of betting on state ω when the actual state is ω , given $\rho_\sigma(\omega)$ and $\mu(\sigma|\omega)$. That is,

$$\bar{\rho}(\omega) = \sum_{\sigma \in \Sigma} \rho_\sigma(\omega) \mu(\sigma|\omega).$$

Simple direct computation yields the partial derivative of V with respect to $\rho_\sigma(\omega)$:⁵⁰

$$\frac{\partial V(\rho)}{\partial \rho_\sigma(\omega)} = \frac{(\psi(1) - \psi(0)) \mu(\omega, \sigma)}{\bar{\rho}(\omega)\psi(1) + (1 - \bar{\rho}(\omega))\psi(0)}.$$

The proof proceeds by contrapositive. We will show that if $\rho_\sigma(\omega') > 0$ and $\rho_{\sigma'}(\omega) > 0$, then the inequality in the statement of the proposition cannot be satisfied. First, note that if $\rho_\sigma(\omega') > 0$, then it must be the case that

$$\frac{\partial V(\rho)}{\partial \rho_\sigma(\omega')} \geq \frac{\partial V(\rho)}{\partial \rho_\sigma(\omega)},$$

for otherwise it would be a strict improvement to reduce $\rho_\sigma(\omega')$ by some $\varepsilon > 0$ and increase $\rho_\sigma(\omega)$ by ε . Similarly, $\rho_{\sigma'}(\omega) > 0$ implies that

$$\frac{\partial V(\rho)}{\partial \rho_{\sigma'}(\omega)} \geq \frac{\partial V(\rho)}{\partial \rho_{\sigma'}(\omega')}.$$

Multiplying these two expressions, we obtain

$$\frac{\partial V(\rho)}{\partial \rho_{\sigma'}(\omega)} \frac{\partial V(\rho)}{\partial \rho_\sigma(\omega')} \geq \frac{\partial V(\rho)}{\partial \rho_\sigma(\omega)} \frac{\partial V(\rho)}{\partial \rho_{\sigma'}(\omega')}.$$

Using the formula for the partial derivative and rearranging terms, this implies that

$$\mu(\omega, \sigma') \mu(\omega', \sigma) \geq \mu(\omega, \sigma) \mu(\omega', \sigma').$$

⁵⁰The choice of ρ by individuals is clearly subject to the constraint that $\sum_{\omega \in \Omega} \rho_\sigma(\omega) = 1$ for all $\sigma \in \Sigma$. This partial derivative treats $\rho_\sigma(\omega)$ as any real number to consider marginal utility independently of feasibility.

Thus, the inequality in the statement of the proposition can only be satisfied if either $\rho_\sigma(\omega') = 0$ or $\rho_{\sigma'}(\omega) = 0$, or both. This completes the proof.

S4.2 Proof of Proposition S2

Since f and c are fixed and deterministic portfolio plans are optimal, we will slightly abuse notation and denote $V^{\text{Pu}}(\rho)$ simply by $V^{\text{Pu}}((\alpha_\sigma)_{\sigma \in \Sigma})$, and similarly denote $V^{\text{Pr}}(\rho)$ by $V^{\text{Pr}}((\alpha_\sigma)_{\sigma \in \Sigma})$. Observe first that for any $(\alpha_\sigma)_{\sigma \in \Sigma}$ and any $\sigma \in \Sigma$,

$$\frac{\partial V^{\text{Pu}}((\alpha_\sigma)_{\sigma \in \Sigma})}{\partial \alpha_\sigma} = \sum_{\omega \in \Omega} \mu(\omega, \sigma) \frac{\psi'(\alpha_\sigma f(\omega) - (1 - \alpha_\sigma)c)}{\psi(\alpha_\sigma f(\omega) - (1 - \alpha_\sigma)c)} (f(\omega) - c)$$

and

$$\frac{\partial V^{\text{Pr}}((\alpha_\sigma)_{\sigma \in \Sigma})}{\partial \alpha_\sigma} = \sum_{\omega \in \Omega} \mu(\omega, \sigma) \frac{\psi'(\alpha_\sigma f(\omega) - (1 - \alpha_\sigma)c)}{\sum_{\sigma' \in \Sigma} \mu(\sigma' | \omega) \psi(\alpha_{\sigma'} f(\omega) - (1 - \alpha_{\sigma'})c)} (f(\omega) - c).$$

Since ψ is positive, increasing, and strictly concave, we can make two straightforward observations that will be useful in the remainder of the proof:

1. The term

$$\frac{\psi'(\alpha f(\omega) - (1 - \alpha)c)}{\psi(\alpha f(\omega) - (1 - \alpha)c)} (f(\omega) - c)$$

is nonincreasing in $\alpha \in [0, 1]$.

2. If $f(\omega) \neq c$ and $\alpha \leq \alpha_\sigma$ for all $\sigma \in \Sigma$, with strict inequality for at least one σ , then

$$\frac{\psi'(\alpha f(\omega) - (1 - \alpha)c)}{\psi(\alpha f(\omega) - (1 - \alpha)c)} (f(\omega) - c) > \frac{\psi'(\alpha f(\omega) - (1 - \alpha)c)}{\sum_{\sigma \in \Sigma} \mu(\sigma | \omega) \psi(\alpha_\sigma f(\omega) - (1 - \alpha_\sigma)c)} (f(\omega) - c).$$

The opposite inequality holds if $\alpha_\sigma \leq \alpha$ for all $\sigma \in \Sigma$, with strict inequality for at least one σ .

Now suppose, contrary to the first part of the proposition, that $\alpha_{\sigma_*}^{\text{Pr}} \geq \alpha_{\sigma_*}^{\text{Pu}}$ and $\alpha_{\sigma_*}^{\text{Pr}} > 0$. Then, we have

$$\frac{\partial V^{\text{Pu}}((\alpha_\sigma^{\text{Pu}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma_*}} \geq \frac{\partial V^{\text{Pu}}((\alpha_\sigma^{\text{Pr}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma_*}} > \frac{\partial V^{\text{Pr}}((\alpha_\sigma^{\text{Pr}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma_*}},$$

where the first inequality follows from observation 1 since $\alpha_{\sigma_*}^{\text{Pr}} \geq \alpha_{\sigma_*}^{\text{Pu}}$, and the second inequality follows from observation 2 with $\alpha = \alpha_{\sigma_*}^{\text{Pr}}$ since $\alpha_{\sigma_*}^{\text{Pr}} \leq \alpha_{\sigma_*}^{\text{Pr}}$ for all $\sigma \in \Sigma$ (with strict inequality for at least one σ). Since, by assumption, $\alpha_{\sigma_*}^{\text{Pu}} \leq \alpha_{\sigma_*}^{\text{Pr}} < \alpha_{\sigma_*}^{\text{Pr}} \leq 1$, the optimality of α^{Pu} requires that

$$\frac{\partial V^{\text{Pu}}((\alpha_\sigma^{\text{Pu}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma_*}} \leq 0,$$

and hence

$$\frac{\partial V^{\text{Pr}}((\alpha_\sigma^{\text{Pr}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma_*}} < 0.$$

Since $(\alpha_{\sigma}^{\text{Pr}})_{\sigma \in \Sigma}$ is optimal, this requires that $\alpha_{\sigma^*}^{\text{Pr}} = 0$, a contradiction. This establishes the first claim in the proposition.

Finally suppose, contrary to the second part of the proposition, that $\alpha_{\sigma^*}^{\text{Pr}} \leq \alpha_{\sigma^*}^{\text{Pu}}$ and $\alpha_{\sigma^*}^{\text{Pr}} < 1$. Then, we have

$$\frac{\partial V^{\text{Pu}}((\alpha_{\sigma}^{\text{Pu}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma^*}} \leq \frac{\partial V^{\text{Pu}}((\alpha_{\sigma}^{\text{Pr}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma^*}} < \frac{\partial V^{\text{Pr}}((\alpha_{\sigma}^{\text{Pr}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma^*}},$$

where the first inequality follows from observation 1 since $\alpha_{\sigma^*}^{\text{Pr}} \leq \alpha_{\sigma^*}^{\text{Pu}}$, and the second inequality follows from observation 2 with $\alpha = \alpha_{\sigma^*}^{\text{Pr}}$ since $\alpha_{\sigma}^{\text{Pr}} \leq \alpha_{\sigma^*}^{\text{Pr}}$ for all $\sigma \in \Sigma$ (with strict inequality for at least one σ). Since, by assumption, $0 \leq \alpha_{\sigma^*}^{\text{Pr}} < \alpha_{\sigma^*}^{\text{Pr}} \leq \alpha_{\sigma^*}^{\text{Pu}}$, the optimality of $(\alpha_{\sigma}^{\text{Pu}})_{\sigma \in \Sigma}$ requires that

$$\frac{\partial V^{\text{Pu}}((\alpha_{\sigma}^{\text{Pu}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma^*}} \geq 0,$$

and hence

$$\frac{\partial V^{\text{Pr}}((\alpha_{\sigma}^{\text{Pr}})_{\sigma \in \Sigma})}{\partial \alpha_{\sigma^*}} > 0.$$

Since $(\alpha_{\sigma}^{\text{Pr}})_{\sigma \in \Sigma}$ is optimal, this requires that $\alpha_{\sigma^*}^{\text{Pr}} = 1$, a contradiction. This establishes the second claim in the proposition.

S5 Omitted Proofs from the Main Paper

In this section, we provide proofs of Propositions 3 and 4 and Lemmas 2 and 3 from Appendix B.2 of the main paper. We restate the results here for ease of reference.

Proposition 3. *Suppose Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$ that is pointwise bounded above, and fix $\mu \in \Delta(\Omega \times S \times \Sigma)$. For any random action plan $\rho \in (\Delta^s(\mathcal{F}))^\Sigma$, the function V defined by Equation (2) can be equivalently expressed as*

$$V(\rho) = \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_{\Omega} \int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right].$$

Proposition 4. *Fix a measure $\mu \in \Delta(\Omega \times S \times \Sigma)$, and suppose Ξ is a nonempty set of functions $\xi : \Omega \times S \times \Sigma \rightarrow [-\infty, \infty)$ with the following properties:*

1. *Closedness: When the set of extend reals $[-\infty, \infty]$ is endowed with its usual topology and $[-\infty, \infty]^{\Omega \times S \times \Sigma}$ is endowed with the product topology (i.e., the topology of pointwise convergence), Ξ is a closed subset of this space.*
2. *Finite measurability: There exists a finite partition $\mathcal{E} \subset \mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma$ of $\Omega \times S \times \Sigma$ such that every $\xi \in \Xi$ is measurable with respect to \mathcal{E} .*
3. *Pointwise boundedness: $\sup_{\xi \in \Xi} \xi(\omega, s, \sigma) < \infty$ for every $(\omega, s, \sigma) \in \Omega \times S \times \Sigma$.*

Then,

$$\begin{aligned} & \sup_{\xi \in \text{co}(\Xi)} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_\Omega \int_{S \times \Sigma} \xi(\omega, s, \sigma) d\mu(s, \sigma | \omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right] \\ &= \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\sup_{\xi \in \Xi} \int_\Omega \int_{S \times \Sigma} \xi(\omega, s, \sigma) d\mu(s, \sigma | \omega) dq(\omega) \right) + R(\mu_\Omega \| q) \right]. \end{aligned} \quad (\text{S4})$$

Lemma 2. Define Ξ as in Equation (12). For any random adaptation plan $\tau \in \mathcal{R}(\Psi | \mathcal{F})$, define $\xi^\tau : \Omega \times S \times \Sigma \rightarrow [-\infty, \infty)$ by

$$\xi^\tau(\omega, s, \sigma) = \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] = \int_{\mathcal{F}} \int_{\Psi} \psi(f(\omega, s)) d\tau_\sigma(\psi | f) d\rho_\sigma(f).$$

Then,

$$\text{co}(\Xi) = \{\xi^\tau : \tau \in \mathcal{R}(\Psi | \mathcal{F})\}.$$

Lemma 3. The set Ξ defined in Equation (12) is a closed subset of $[-\infty, \infty]^{\Omega \times S \times \Sigma}$.

S5.1 Proof of Proposition 3

We begin with a useful proposition. As in the main text, let $(\Omega, \mathcal{B}_\Omega)$ be any measurable space, and let $\Delta(\Omega)$ be the set of all countably additive probability measures on this space. Recall that $M(p) = \{q \in \Delta(\Omega) : q \ll p \text{ and } R(p \| q) < \infty\}$. In particular, since $R(p \| q) < \infty$ requires that $p \ll q$, the measures q and p are mutually absolutely continuous whenever $q \in M(p)$.

Proposition S3. Suppose $X : \Omega \rightarrow [-\infty, \infty)$ is measurable and bounded above, and let $p \in \Delta(\Omega)$. Then,

$$\int_\Omega \ln(X(\omega)) dp(\omega) = \inf_{q \in M(p)} \left[\ln \left(\int_\Omega X(\omega) dq(\omega) \right) + R(p \| q) \right]. \quad (\text{S5})$$

In addition, if X is bounded away from zero, that is, if $X(\omega) \geq \varepsilon > 0$ for all $\omega \in \Omega$, then the infimum in Equation (S5) is uniquely attained by the measure q_0 with Radon–Nikodym derivative

$$\frac{dq_0}{dp}(\omega) = \frac{1}{X(\omega) \int_\Omega \frac{1}{X(\hat{\omega})} dp(\hat{\omega})}. \quad (\text{S6})$$

Proposition S3 restricts to $q \in M(p)$, thereby ensuring that we do not encounter terms of the form $-\infty + \infty$. That is, while the first term inside the infimum in Equation (S5) could take the value $-\infty$, the second term $R(p \| q)$ will necessarily be finite.

Proof. The proof proceeds in three steps. We first prove Equation (S5) for random variables X that are bounded above and satisfy $X(\omega) \geq \varepsilon > 0$ for all $\omega \in \Omega$. We then extend the result to all bounded $X \geq 0$. Finally, we extend to any X that is bounded above.⁵¹

⁵¹The first two steps in our proof employ similar techniques to the proofs of Propositions 1.4.2 and 4.5.1 in Dupuis and Ellis (1997), although the details are quite different.

Step 1: Suppose that X that is bounded above and satisfies $X(\omega) \geq \varepsilon > 0$ for all $\omega \in \Omega$. Then, $\ln(X)$ is a bounded function, and it is therefore integrable. Fix any measures $p, q \in \Delta(\Omega)$ with $p \ll q$ and define a measure p_0 by its Radon–Nikodym derivative

$$\frac{dp_0}{dq}(\omega) = \frac{X(\omega)}{\int_{\Omega} X(\hat{\omega}) dq(\hat{\omega})}. \quad (\text{S7})$$

Since X is strictly positive, p_0 and q are mutually absolutely continuous. In particular, since $p \ll q$, this implies $p \ll p_0$. Thus, $\frac{dp}{dp_0}$ exists and $\frac{dp}{dq} = \frac{dp}{dp_0} \cdot \frac{dp_0}{dq}$. Note that

$$\begin{aligned} & \int_{\Omega} \ln(X) dp - R(p \parallel q) \\ &= \int_{\Omega} \ln(X) dp - \int_{\Omega} \ln\left(\frac{dp}{dq}\right) dp \\ &= \int_{\Omega} \ln(X) dp - \int_{\Omega} \ln\left(\frac{dp}{dp_0}\right) dp - \int_{\Omega} \ln\left(\frac{dp_0}{dq}\right) dp \\ &= \int_{\Omega} \ln(X) dp - \int_{\Omega} \ln\left(\frac{dp}{dp_0}\right) dp - \int_{\Omega} \ln(X) dp + \ln\left(\int_{\Omega} X dq\right) \\ &= -R(p \parallel p_0) + \ln\left(\int_{\Omega} X dq\right). \end{aligned}$$

By Lemma 1.4.1 in Dupuis and Ellis (1997), $R(p \parallel p_0) \geq 0$, with equality if and only if $p = p_0$. Therefore,

$$\int_{\Omega} \ln(X) dp \leq \ln\left(\int_{\Omega} X dq\right) + R(p \parallel q),$$

with equality if and only if $p = p_0$. It is not difficult to show that Equations (S6) and (S7) are dual in the sense that $p = p_0$ if and only if $q = q_0$. Therefore, given p , if we set $q = q_0$ then the above holds with equality. Moreover, since X is bounded and $1/X \leq 1/\varepsilon$,

$$R(p \parallel q_0) = \int_{\Omega} \ln\left(\frac{dp}{dq_0}\right) dp = \int_{\Omega} \ln(X) dp + \ln\left(\int_{\Omega} \frac{1}{X} dp\right) < \infty,$$

which implies $q_0 \in M(p)$. Hence the infimum in Equation (S5) is attained at q_0 .

Step 2: Consider now any bounded $X \geq 0$. Define a sequence of random variables $(X_n)_{n \in \mathbb{N}}$ by $X_n(\omega) = \max\{X(\omega), 1/n\}$. By step 1, we know that Equation (S5) holds for each X_n and for any p . Using this, together with the fact that $X_n \geq X$ for all n , we have

$$\begin{aligned} \int_{\Omega} \ln(X_n) dp &= \inf_{q \in M(p)} \left[\ln\left(\int_{\Omega} X_n dq\right) + R(p \parallel q) \right] \\ &\geq \inf_{q \in M(p)} \left[\ln\left(\int_{\Omega} X dq\right) + R(p \parallel q) \right]. \end{aligned}$$

Since $\int \ln(X_1) dp < \infty$ and $\ln(X_n) \downarrow \ln(X)$, the monotone convergence theorem for extended real-

valued functions (e.g., Theorem 4.3.2 of Dudley (2002)) implies

$$\begin{aligned} \int_{\Omega} \ln(X) dp &= \lim_{n \rightarrow \infty} \int_{\Omega} \ln(X_n) dp \\ &\geq \inf_{q \in M(p)} \left[\ln \left(\int_{\Omega} X dq \right) + R(p \| q) \right]. \end{aligned}$$

Note that these terms could take the value $-\infty$.

To prove the opposite inequality, note that for any n and any $q \in M(p)$, Equation (S5) applied to the function X_n implies

$$\int_{\Omega} \ln(X_n) dp \leq \ln \left(\int_{\Omega} X_n dq \right) + R(p \| q).$$

Since both sides of this inequality are finite for all n , we can again take the limit as $n \rightarrow \infty$ and apply the monotone convergence theorem to obtain

$$\int_{\Omega} \ln(X) dp \leq \ln \left(\int_{\Omega} X dq \right) + R(p \| q).$$

Since this is true for all $q \in M(p)$, we have

$$\int_{\Omega} \ln(X) dp \leq \inf_{q \in M(p)} \left[\ln \left(\int_{\Omega} X dq \right) + R(p \| q) \right].$$

Thus, Equation (S5) holds for any bounded $X \geq 0$.

Step 3: Finally, consider any X that is bounded above. Let $X^+(\omega) = \max\{X(\omega), 0\}$. Since we have adopted the standard convention that $\ln(x) = -\infty$ for any $x \leq 0$, we have $\ln(X^+(\omega)) = \ln(X(\omega))$ for all ω . Therefore, since Equation (S5) holds for X^+ by step 2,

$$\begin{aligned} \int_{\Omega} \ln(X) dp &= \int_{\Omega} \ln(X^+) dp \\ &= \inf_{q \in M(p)} \left[\ln \left(\int_{\Omega} X^+ dq \right) + R(p \| q) \right] \\ &\geq \inf_{q \in M(p)} \left[\ln \left(\int_{\Omega} X dq \right) + R(p \| q) \right]. \end{aligned}$$

To establish the opposite inequality, we consider two cases. Let $A = \{\omega \in \Omega : X(\omega) \leq 0\}$. The first case is when $p(A) > 0$. Then, $\int_{\Omega} \ln(X) dp = -\infty$, so the above must hold with equality. The second case is when $p(A) = 0$. Then, $q(A) = 0$ for all $q \in M(p)$, since any $q \in M(p)$ must be absolutely continuous with respect to p . Therefore, $\int_{\Omega} X dq = \int_{\Omega} X^+ dq$ for all $q \in M(p)$ and hence

$$\inf_{q \in M(p)} \left[\ln \left(\int_{\Omega} X dq \right) + R(p \| q) \right] = \inf_{q \in M(p)} \left[\ln \left(\int_{\Omega} X^+ dq \right) + R(p \| q) \right].$$

Thus, the equality is established for both cases, which completes the proof. ■

We now proceed with the proof of Proposition 3. For a given $\rho \in (\Delta^s(\mathcal{F}))^\Sigma$ and $\tau \in \mathcal{R}(\Psi|\mathcal{F})$, define $X : \Omega \rightarrow [-\infty, \infty)$ by

$$X(\omega) = \int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma|\omega).$$

To verify that X is bounded above, recall that for each σ , $\rho_\sigma \in \Delta^s(\mathcal{F})$ has finite support and each $f \in \text{supp}(\rho_\sigma)$ is a simple act. Moreover, since Σ is finite, this implies that only finitely many realizations of z occur with positive probability. Since the set Ψ is pointwise bounded above, this implies that there exists $\kappa \in \mathbb{R}$ such that $\psi(f(\omega, s)) \leq \kappa$ for all ω, s, σ , and $f \in \text{supp}(\rho_\sigma)$. Therefore, $X(\omega) \leq \kappa$ for all ω . Applying Proposition S3 to this function, we obtain

$$\begin{aligned} & \int_{\Omega} \ln \left(\int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma|\omega) \right) d\mu(\omega) \\ &= \int_{\Omega} \ln(X(\omega)) d\mu_{\Omega}(\omega) \\ &= \inf_{q \in M(\mu_{\Omega})} \left[\ln \left(\int_{\Omega} X(\omega) dq(\omega) \right) + R(\mu_{\Omega} \| q) \right] \\ &= \inf_{q \in M(\mu_{\Omega})} \left[\ln \left(\int_{\Omega} \int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_{\Omega} \| q) \right]. \end{aligned}$$

Thus, when V is defined by Equation (2), we have

$$\begin{aligned} V(\rho) &= \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \int_{\Omega} \ln \left(\int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma|\omega) \right) d\mu(\omega) \\ &= \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \inf_{q \in M(\mu_{\Omega})} \left[\ln \left(\int_{\Omega} \int_{S \times \Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} [\psi(f(\omega, s))] d\mu(s, \sigma|\omega) dq(\omega) \right) + R(\mu_{\Omega} \| q) \right]. \end{aligned}$$

This completes the proof.

S5.2 Proof of Proposition 4

Our proof will rely on a version of the von Neumann–Sion Minimax Theorem. von Neumann (1928) proved that when $F : C \times D \rightarrow \mathbb{R}$ is a bilinear function and C and D are finite-dimensional simplexes,

$$\sup_{x \in C} \inf_{y \in D} F(x, y) = \inf_{y \in D} \sup_{x \in C} F(x, y).$$

Perhaps the most important and well-known extension of von Neumann’s result is due to Sion (1958), who showed that the same conclusion can be derived under the weaker assumptions that C and D are convex subsets of topological vector spaces, one of these sets is compact, F is quasiconcave and upper semicontinuous in x , and F is quasiconvex and lower semicontinuous in y . Sion’s result is not quite strong enough for our purposes, since in our application it may be that neither C nor D is compact and since F may not be lower semicontinuous in y . We will therefore rely on the following generalization of the von Neumann–Sion Theorem, which is due to Tuij (2004).

Theorem S2 (von Neumann–Sion–Tuy Minimax Theorem). *Let C be a closed and convex subset of a topological vector space, and let D be a convex subset of a topological vector space. Suppose $F : C \times D \rightarrow \mathbb{R}$ satisfies the following conditions:*

1. *For every $y \in D$, the function $x \mapsto F(x, y)$ is quasiconcave and upper semicontinuous on C .*
2. *For every $x \in C$ and $y, y' \in D$, the function $\lambda \mapsto F(x, \lambda y + (1 - \lambda)y')$ is quasiconvex and lower semicontinuous on $[0, 1]$.*
3. *There exists some $\eta < \inf_{y \in D} \sup_{x \in C} F(x, y)$ and a nonempty finite set $L \subset D$ such that the set $C_\eta^L = \{x \in C : \min_{y \in L} F(x, y) \geq \eta\}$ is compact.*

Then,

$$\sup_{x \in C} \inf_{y \in D} F(x, y) = \inf_{y \in D} \sup_{x \in C} F(x, y).$$

Proof. This result is a special case of Theorem 2 in Tuy (2004). His result requires that F be what he calls α -connected. This condition is implied by our assumptions that C is closed and convex, D is convex, F is quasiconcave and upper semicontinuous in x , and $\lambda \mapsto F(x, \lambda y + (1 - \lambda)y')$ is quasiconvex in λ for all x, y, y' . His result also requires the lower semicontinuity property that we assumed in condition 2.⁵² The final assumption needed for his result is condition 3.⁵³ ■

Note that the theorem of Sion (1958) follows as a corollary to this result: If F is quasiconvex and lower semicontinuous in y then condition 2 is implied, and if D is compact then condition 3 is implied (given that F is upper semicontinuous in x).

We now proceed with the proof of Proposition 4. Fix any measure $\mu \in \Delta(\Omega \times S \times \Sigma)$, and fix any convex set Ξ satisfying the properties described in the statement of the proposition. We proceed in several steps. Using the second property of Ξ from the statement of the proposition, we know that there exists a finite partition \mathcal{E} of $\Omega \times S \times \Sigma$ such that every $\xi \in \Xi$ is measurable with respect to \mathcal{E} . We can enumerate the elements of this partition as

$$\mathcal{E} = \{E_i : i \in N\},$$

where N is a finite index set. For each $i \in N$, fix an arbitrary element $(\omega_i, s_i, \sigma_i) \in E_i$. Since each $\xi \in \Xi$ is measurable with respect \mathcal{E} , we know that $\xi(\omega, s, \sigma) = \xi(\omega_i, s_i, \sigma_i)$ for all $i \in N$ and $(\omega, s, \sigma) \in E_i$. Consider the mapping

$$\xi \mapsto \theta^\xi = (\xi(\omega_i, s_i, \sigma_i))_{i \in N}$$

⁵²Note that the assumption of lower semicontinuity in y in every line segment (that is, lower semicontinuity of the mapping $\lambda \mapsto F(x, \lambda y + (1 - \lambda)y')$ for all x, y, y') in condition 2 is in general weaker than assuming lower semicontinuity in y . However, the assumption of quasiconvexity in y in every line segment (that is, quasiconvexity of the mapping $\lambda \mapsto F(x, \lambda y + (1 - \lambda)y')$ for all x, y, y') in condition 2 is equivalent to quasiconvexity in y . Also, note that we have switched the roles of C and D compared to Tuy (2004).

⁵³Strictly speaking, Theorem 2 in Tuy (2004) assumes that C_η^L is compact for $\eta = \sup_{x \in C} \inf_{y \in D} F(x, y)$ and shows that $\eta < \inf_{y \in D} \sup_{x \in C} F(x, y)$ leads to a contradiction. As is evident from his proof, our condition 4 is sufficient to obtain the same result.

from Ξ into $[-\infty, \infty]^N$. It is easy to see that this mapping is a homeomorphism from Ξ to the set

$$\Theta = \{\theta^\xi : \xi \in \Xi\} \subset [-\infty, \infty]^N.$$

In other words, the set of functions Ξ is topologically equivalent to the set of vectors Θ .

As in the main paper, for any $q \in M(\mu_\Omega)$, define the measure $\mu \otimes q$ on $\Omega \times S \times \Sigma$ to have marginal q on Ω and conditional distribution $\mu(\cdot|\omega)$ on $S \times \Sigma$. That is, for any event E in the product σ -algebra $\mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma$, let

$$\mu \otimes q(E) = \int_\Omega \int_{S \times \Sigma} \mathbf{1}[(\omega, s, \sigma) \in E] d\mu(s, \sigma|\omega) dq(\omega).$$

Define a function $H : [-\infty, \infty]^N \times M(\mu_\Omega) \rightarrow \mathbb{R}_+$ by

$$H(\theta, q) = \max \left\{ 0, \sum_{i \in N} \theta_i \cdot \mu \otimes q(E_i) \right\} \exp(R(\mu_\Omega \| q)).$$

Lemma S1. *The set Θ and function H satisfy the following conditions:*

1. *When $[-\infty, \infty]^N$ is endowed with the product topology (i.e., the topology of pointwise convergence), Θ is compact.*
2. *There exists $\kappa \in \mathbb{R}$ such that $\theta_i \leq \kappa$ for all $\theta \in \Theta$ and $i \in N$.*
3. *Equation (S4) from the statement of the proposition is equivalent to the following:*

$$\sup_{\theta \in \text{co}(\Theta)} \inf_{q \in M(\mu_\Omega)} H(\theta, q) = \inf_{q \in M(\mu_\Omega)} \sup_{\theta \in \Theta} H(\theta, q). \quad (\text{S8})$$

Proof. Since Ξ is a closed subset of $[-\infty, \infty]^{\Omega \times S \times \Sigma}$ by the first property in the statement of the proposition and since Ξ and Θ are homeomorphic, Θ is closed. In addition, since $[-\infty, \infty]^N$ is a compact space when endowed with the product topology,⁵⁴ this implies that Θ is compact. Since Ξ is pointwise bounded above by the third property in the statement of the proposition, we have

$$\sup_{\theta \in \Theta} \theta_i = \sup_{\xi \in \Xi} \xi(\omega_i, s_i, \sigma_i) < \infty$$

for all $i \in N$. In particular, since N is finite, there exists $\kappa \in \mathbb{R}$ such that $\theta_i \leq \kappa$ for all $\theta \in \Theta$ and

⁵⁴It is easy to see that the set of extended reals $[-\infty, \infty]$ is compact in its usual topology (see Example 2.75 in Aliprantis and Border (2006)), and hence $[-\infty, \infty]^N$ endowed with the product topology is compact by the Tychonoff Product Theorem (Theorem 2.61 in Aliprantis and Border (2006)).

$i \in N$. To establish the third condition, note that⁵⁵

$$\begin{aligned}
& \ln \left[\sup_{\theta \in \text{co}(\Theta)} \inf_{q \in M(\mu_\Omega)} H(\theta, q) \right] \\
&= \sup_{\theta \in \text{co}(\Theta)} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\sum_{i \in N} \theta_i \cdot \mu \otimes q(E_i) \right) + R(\mu_\Omega \| q) \right] \\
&= \sup_{\xi \in \text{co}(\Xi)} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\sum_{i \in N} \xi(\omega_i, s_i, \sigma_i) \cdot \mu \otimes q(E_i) \right) + R(\mu_\Omega \| q) \right] \\
&= \sup_{\xi \in \text{co}(\Xi)} \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\int_{\Omega \times S \times \Sigma} \xi(\omega, s, \sigma) d(\mu \otimes q)(\omega, s, \sigma) \right) + R(\mu_\Omega \| q) \right].
\end{aligned}$$

Similarly,

$$\begin{aligned}
& \ln \left[\inf_{q \in M(\mu_\Omega)} \sup_{\theta \in \Theta} H(\theta, q) \right] \\
&= \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\sup_{\theta \in \Theta} \sum_{i \in N} \theta_i \cdot \mu \otimes q(E_i) \right) + R(\mu_\Omega \| q) \right] \\
&= \inf_{q \in M(\mu_\Omega)} \left[\ln \left(\sup_{\xi \in \Xi} \int_{\Omega \times S \times \Sigma} \xi(\omega, s, \sigma) d(\mu \otimes q)(\omega, s, \sigma) \right) + R(\mu_\Omega \| q) \right].
\end{aligned}$$

Thus, Equation (S4) is equivalent to Equation (S8). ■

Next, we show that we can remove any indices $i \in N$ that correspond to probability zero events. By definition, q and μ_Ω must be mutually absolutely continuous for any $q \in M(\mu_\Omega)$, and hence $\mu \otimes q$ and μ are also mutually absolutely continuous. Thus, for any $i \in N$ and $q \in M(\mu_\Omega)$,

$$\mu \otimes q(E_i) = 0 \iff \mu(E_i) = 0.$$

We can therefore remove any events $E_i \in \mathcal{E}$ that occur with zero probability under μ , since such events must also occur with zero probability under $\mu \otimes q$ for any $q \in M(\mu_\Omega)$. That is, consider the index set $M \subset N$ given by

$$M = \{i \in N : \mu(E_i) > 0\}.$$

Define the projection function $P_M : [-\infty, \infty]^N \rightarrow [-\infty, \infty]^M$ by $P_M(\theta) = (\theta_i)_{i \in M}$, and set

$$\Theta' = P_M(\Theta) = \{\theta' = P_M(\theta) : \theta \in \Theta\}.$$

Define a function $F : [-\infty, \infty]^M \times M(\mu_\Omega) \rightarrow \mathbb{R}_+$ by

$$F(\theta, q) = \max \left\{ 0, \sum_{i \in M} \theta_i \cdot \mu \otimes q(E_i) \right\} \exp(R(\mu_\Omega \| q)).$$

Lemma S2. *The set Θ' and function F satisfy the following conditions:*

⁵⁵To deal with vectors θ and functions ξ that can take the value $-\infty$, we adopt the notational convention throughout that $\ln(x) = -\infty$ for any $x \in [-\infty, 0]$. Hence $\ln(\max\{0, x\}) = \ln(x)$ for all $x \in [-\infty, \infty)$.

1. When $[-\infty, \infty]^M$ is endowed with the product topology, Θ' is compact (hence closed).
2. There exists $\kappa \in \mathbb{R}$ such that $\theta_i \leq \kappa$ for all $\theta \in \Theta'$ and $i \in M$.
3. Equation (S8) is equivalent to the following:

$$\sup_{\theta \in \text{co}(\Theta')} \inf_{q \in M(\mu_\Omega)} F(\theta, q) = \inf_{q \in M(\mu_\Omega)} \sup_{\theta \in \Theta'} F(\theta, q). \quad (\text{S9})$$

Proof. The projection function P_M is continuous when $[-\infty, \infty]^N$ and $[-\infty, \infty]^M$ are endowed with their product topologies. Therefore, the set Θ' is compact, as it is the image of the compact set Θ under the continuous function P_M . Since $[-\infty, \infty]^N$ is a Hausdorff space, compact subsets of this space are closed (Lemma 2.32 in Aliprantis and Border (2006)). Hence, Θ' is closed. The second condition follows directly from the second condition in Lemma S1. To establish the third condition, recall from above that μ and $\mu \otimes q$ are mutually absolutely continuous for any $q \in M(\mu_\Omega)$. This implies that for any $\theta \in [-\infty, \infty]^N$, if we take $\theta' = P_M(\theta) \in [-\infty, \infty]^M$, then $H(\theta, q) = F(\theta', q)$ for all $q \in M(\mu_\Omega)$. Therefore, Equations (S8) and (S9) are equivalent. ■

We now show that we can remove any $\theta \in \Theta'$ such that $\theta_i = -\infty$ for some $i \in M$, thereby reducing this set to a subset of the Euclidean space \mathbb{R}^M . Formally, let

$$\Theta'' = \{\theta \in \Theta' : \theta_i > -\infty, \forall i \in M\}.$$

Note that it is possible to have $\Theta'' = \emptyset$.

Lemma S3. *The set Θ'' and function F satisfy the following conditions:*

1. When \mathbb{R}^M is endowed with the Euclidean topology, Θ'' is closed.
2. There exists $\kappa \in \mathbb{R}$ such that $\theta_i \leq \kappa$ for all $\theta \in \Theta''$ and $i \in M$.
3. Equation (S9) holds either if $\Theta'' = \emptyset$, or if $\Theta'' \neq \emptyset$ and

$$\sup_{\theta \in \text{co}(\Theta'')} \inf_{q \in M(\mu_\Omega)} F(\theta, q) = \inf_{q \in M(\mu_\Omega)} \sup_{\theta \in \Theta''} F(\theta, q). \quad (\text{S10})$$

4. Fix any $q \in M(\mu_\Omega)$. When restricted to \mathbb{R}^M (endowed with the Euclidean topology), the mapping $\theta \mapsto F(\theta, q)$ is continuous, nondecreasing, quasiconcave, and quasiconvex.

Proof. Since Θ' is a closed subset of $[-\infty, \infty]^M$ (endowed with the product topology of the extended reals), it is easy to verify that Θ'' is a closed subset of \mathbb{R}^M (endowed with the Euclidean topology). Note, however, that Θ'' need not be compact. Next, the second condition follows directly from the second condition in Lemma S2. To establish the third condition, note that if $\theta \in \text{co}(\Theta')$ has $\theta_i = -\infty$ for some $i \in M$, then for any $q \in M(\mu_\Omega)$,

$$\sum_{i \in M} \theta_i \cdot \mu \otimes q(E_i) = -\infty,$$

and hence $F(\theta, q) = 0$. Thus, if $\Theta'' = \emptyset$, then $F(\theta, q) = 0$ for all $\theta \in \text{co}(\Theta')$ and $q \in M(\mu_\Omega)$, so Equation (S9) holds trivially. In the alternative case of $\Theta'' \neq \emptyset$, it is immediate that Equations (S9) and (S10) are equivalent.

To verify the fourth condition, fix any $q \in M(\mu_\Omega)$. Note that the mapping

$$\theta \mapsto \sum_{i \in M} \theta_i \cdot \mu \otimes q(E_i)$$

is continuous, nondecreasing, and linear. Therefore, the mapping $\theta \mapsto F(\theta, q)$ is continuous, non-decreasing, quasiconcave, and quasiconvex (though it is obviously no longer linear). ■

To apply the minimax theorem, we need the set over which the supremum is being taking to be closed and convex. That is, we will want to show that we can replace $\text{co}(\Theta'')$ with $\text{cl}(\text{co}(\Theta''))$ on the left side of Equation (S10) and replace Θ'' with $\text{cl}(\text{co}(\Theta''))$ on the right side without affecting either of these values. The next two lemmas show that this is possible for the set Θ'' and function F in question.

Lemma S4. *Suppose $Y \subset \mathbb{R}^M$ is closed, and suppose there exists $\kappa \in \mathbb{R}$ such that $y_i \leq \kappa$ for all $y \in Y$ and $i \in M$. Then, for any $y \in \text{cl}(\text{co}(Y))$ there exists $y' \in \text{co}(Y)$ such that $y' \geq y$ (that is, $y'_i \geq y_i$ for all $i \in M$).*

Proof. Suppose $y \in \text{cl}(\text{co}(Y))$. Then there exists a sequence (y_n) in $\text{co}(Y)$ such that $y_n \rightarrow y$. Let m be the cardinality of the set M . By Caratheodory's Convexity Theorem (Theorem 5.32 in Aliprantis and Border (2006)), every element of $\text{co}(Y)$ can be written as a convex combination of at most $m + 1$ vectors from Y . Therefore, each y_n can be written as

$$y_n = \sum_{j=1}^{m+1} \alpha_n^j y_n^j,$$

where $y_n^j \in Y$ for all $n \in \mathbb{N}$ and $j \in \{1, \dots, m + 1\}$, and $\alpha_n = (\alpha_n^1, \dots, \alpha_n^{m+1}) \in [0, 1]^{m+1}$ satisfies $\alpha_n^1 + \dots + \alpha_n^{m+1} = 1$ for all $n \in \mathbb{N}$. Since $[0, 1]^{m+1}$ is compact, (α_n) has a convergent subsequence. With slight abuse of notation, denote this subsequence again by (α_n) . That is, we can assume without loss of generality that $\alpha_n \rightarrow \alpha$ for some $\alpha = (\alpha^1, \dots, \alpha^{m+1}) \in [0, 1]^{m+1}$.

We claim that the sequence (y_n^j) in Y is bounded for all j such that $\alpha^j > 0$. For suppose to the contrary that (y_n^j) is unbounded. Then, since Y is bounded above by κ , this would imply there exists some subsequence $(y_{n_k}^j)$ and some dimension $i \in M$ such that $y_{i, n_k}^j \rightarrow -\infty$. However, since $\alpha^j > 0$ and $y_{i, n_k}^j \leq \kappa$ for all j' , this implies $y_{i, n_k} \rightarrow -\infty$, contradicting the fact that this subsequence converges to $y_i \in \mathbb{R}$. Thus, (y_n^j) must be bounded.

Therefore, by passing to subsequences if necessary, it is without loss of generality to assume that (y_n^j) converges for all j for which $\alpha^j > 0$. Denote the limits of these sequences by y^j , respectively, and let

$$y' = \sum_{\substack{j \in \{1, \dots, m+1\}: \\ \alpha^j > 0}} \alpha^j y^j.$$

Since Y is closed, each of these y^j is in Y , and hence $y' \in \text{co}(Y)$. Now, for every $n \in \mathbb{N}$ and $i \in M$,

$$y_{i,n} = \sum_{\substack{j \in \{1, \dots, m+1\}: \\ \alpha^j > 0}} \alpha_n^j y_{i,n}^j + \sum_{\substack{j \in \{1, \dots, m+1\}: \\ \alpha^j = 0}} \alpha_n^j y_{i,n}^j \leq \sum_{\substack{j \in \{1, \dots, m+1\}: \\ \alpha^j > 0}} \alpha_n^j y_{i,n}^j + \sum_{\substack{j \in \{1, \dots, m+1\}: \\ \alpha^j = 0}} \alpha_n^j \kappa,$$

since $y_{i,n}^j \leq \kappa$. Taking limits, the left side of this inequality converges to y_i and the right side converges to y'_i . Thus, $y \leq y'$, as claimed. \blacksquare

Lemma S5. *If $\Theta'' \neq \emptyset$, Equation (S10) is equivalent to the following:*

$$\sup_{\theta \in \text{cl}(\text{co}(\Theta''))} \inf_{q \in M(\mu_\Omega)} F(\theta, q) = \inf_{q \in M(\mu_\Omega)} \sup_{\theta \in \text{cl}(\text{co}(\Theta''))} F(\theta, q). \quad (\text{S11})$$

Proof. The function F is nondecreasing in θ by Lemma S3. Therefore, for any $\theta, \theta' \in \mathbb{R}^M$ and $q \in M(\mu_\Omega)$, $\theta' \geq \theta$ implies $F(\theta', q) \geq F(\theta, q)$. Therefore,

$$\theta' \geq \theta \implies \inf_{q \in M(\mu_\Omega)} F(\theta', q) \geq \inf_{q \in M(\mu_\Omega)} F(\theta, q).$$

Also, since Θ'' is closed and bounded above by Lemma S3, Lemma S4 implies for any $\theta \in \text{cl}(\text{co}(\Theta''))$ there exists $\theta' \in \text{co}(\Theta'')$ such that $\theta' \geq \theta$. Therefore,

$$\sup_{\theta \in \text{co}(\Theta'')} \inf_{q \in M(\mu_\Omega)} F(\theta, q) = \sup_{\theta \in \text{cl}(\text{co}(\Theta''))} \inf_{q \in M(\mu_\Omega)} F(\theta, q).$$

This establishes that the left sides of Equations (S10) and (S11) are the same.

To see that the right sides of these equations are also the same, first fix any $\theta \in \text{co}(\Theta'')$. Thus, $\theta = \sum_{j=1}^m \alpha^j \theta^j$ for some $m \in \mathbb{N}$ and $\theta^j \in \Theta''$, $j \in \{1, \dots, m\}$. Since for any $q \in M(\mu_\Omega)$, the mapping $\theta \mapsto F(\theta, q)$ is quasiconvex by Lemma S3, this implies that $F(\theta) \leq F(\theta^j)$ for some j . Therefore,

$$\sup_{\theta \in \Theta''} F(\theta, q) = \sup_{\theta \in \text{co}(\Theta'')} F(\theta, q)$$

for every $q \in M(\mu_\Omega)$. By the same arguments used above, it is also true that

$$\sup_{\theta \in \text{co}(\Theta'')} F(\theta, q) = \sup_{\theta \in \text{cl}(\text{co}(\Theta''))} F(\theta, q).$$

Combining these observations, we see that the right sides of Equations (S10) and (S11) are the same. \blacksquare

We are almost ready to apply the minimax theorem to prove that Equation (S11) holds whenever $\Theta'' \neq \emptyset$. First, the following lemma will be used to establish some of the necessary properties of the mapping $q \mapsto F(\theta, q)$.

Lemma S6. *Suppose $X : \Omega \rightarrow \mathbb{R}$ is measurable and bounded, and fix any $p \in \Delta(\Omega)$. Then, for any $q, q' \in M(p)$, the mapping*

$$\lambda \mapsto \max \left\{ 0, \int_{\Omega} X d(\lambda q + (1 - \lambda)q') \right\} \exp(R(p \| \lambda q + (1 - \lambda)q'))$$

is quasiconvex and lower semicontinuous on the interval $[0, 1]$.

Proof. Our proof will make use of the Donsker–Varadhan variational formula (see, for example, Lemma 1.4.3 in Dupuis and Ellis (1997)), which states that for any $p, r \in \Delta(\Omega)$,

$$R(p \parallel r) = \sup_{Y \in B_b(\Omega)} \left[\int_{\Omega} Y dp - \ln \left(\int_{\Omega} \exp(Y) dr \right) \right],$$

where $B_b(\Omega)$ denotes the space of all bounded Borel measurable real functions on Ω . Therefore,

$$\exp(R(p \parallel r)) = \sup_{Y \in B_b(\Omega)} \frac{\exp\left(\int_{\Omega} Y dp\right)}{\int_{\Omega} \exp(Y) dr},$$

and hence

$$\max \left\{ 0, \int_{\Omega} X dr \right\} \exp(R(p \parallel r)) = \max \left\{ 0, \sup_{Y \in B_b(\Omega)} \frac{\exp\left(\int_{\Omega} Y dp\right) \int_{\Omega} X dr}{\int_{\Omega} \exp(Y) dr} \right\}.$$

We will show for any $X, Y \in B_b(\Omega)$, $p \in \Delta(\Omega)$, and $q, q' \in M(p)$, the function $h : [0, 1] \rightarrow \mathbb{R}$ defined by

$$h(\lambda) = \frac{\exp\left(\int_{\Omega} Y dp\right) \int_{\Omega} X d(\lambda q + (1 - \lambda)q')}{\int_{\Omega} \exp(Y) d(\lambda q + (1 - \lambda)q')}$$

is quasiconvex and lower semicontinuous. This will establish the claim in the statement of the lemma, since the supremum of a set of quasiconvex and lower semicontinuous functions retains these properties.

Continuity of the function h in λ is immediate. To see that h is quasiconvex, fix any $\gamma \in \mathbb{R}$ and fix any $\lambda_1, \lambda_2 \in [0, 1]$ such that $h(\lambda_1) \leq \gamma$ and $h(\lambda_2) \leq \gamma$. Suppose without loss of generality that $\lambda_1 \leq \lambda_2$. We need to show that $h(\lambda) \leq \gamma$ for any $\lambda \in (\lambda_1, \lambda_2)$. Note that $h(\lambda_i) \leq \gamma$ is equivalent to

$$\exp\left(\int_{\Omega} Y dp\right) \int_{\Omega} X d(\lambda_i q + (1 - \lambda_i)q') \leq \gamma \int_{\Omega} \exp(Y) d(\lambda_i q + (1 - \lambda_i)q').$$

Any $\lambda \in (\lambda_1, \lambda_2)$ can be written as $\alpha \lambda_1 + (1 - \alpha) \lambda_2$ for $\alpha = (\lambda_2 - \lambda) / (\lambda_2 - \lambda_1)$. Therefore, we have

$$\begin{aligned} & \exp\left(\int_{\Omega} Y dp\right) \int_{\Omega} X d(\lambda q + (1 - \lambda)q') \\ &= \alpha \exp\left(\int_{\Omega} Y dp\right) \int_{\Omega} X d(\lambda_1 q + (1 - \lambda_1)q') + (1 - \alpha) \exp\left(\int_{\Omega} Y dp\right) \int_{\Omega} X d(\lambda_2 q + (1 - \lambda_2)q') \\ &\leq \alpha \gamma \int_{\Omega} \exp(Y) d(\lambda_1 q + (1 - \lambda_1)q') + (1 - \alpha) \gamma \int_{\Omega} \exp(Y) d(\lambda_2 q + (1 - \lambda_2)q') \\ &= \gamma \int_{\Omega} \exp(Y) d(\lambda q + (1 - \lambda)q'), \end{aligned}$$

which implies $h(\lambda) \leq \gamma$. This establishes that h is quasiconvex, which completes the proof. ■

The following lemma applies Theorem S2 to prove that Equation (S11) holds whenever $\Theta'' \neq \emptyset$. In light of Lemmas S1, S2, S3, and S5, this will establish Equation (S4) and complete the proof of Proposition 4.

Lemma S7. *If $\Theta'' \neq \emptyset$, then Equation (S11) is satisfied.*

Proof. We only need to establish that the assumptions of Theorem S2 are satisfied for the sets $C = \text{cl}(\text{co}(\Theta''))$, $D = M(\mu_\Omega)$, and for the function F defined above.

Note that C is a closed and convex subset of \mathbb{R}^M by definition. It is also straightforward to show that the set D is convex. To see that condition 1 is satisfied, recall that for any $q \in D$, the mapping $\theta \mapsto F(\theta, q)$ is continuous and quasiconcave on C by Lemma S3.

Next, fix any $\theta \in C$ and define $X : \Omega \rightarrow \mathbb{R}$ by⁵⁶

$$X(\omega) = \int_{S \times \Sigma} \sum_{i \in M} \theta_i \cdot \mathbf{1}[(\omega, s, \sigma) \in E_i] d\mu(s, \sigma | \omega).$$

Then, for any $q \in M(\mu_\Omega)$,

$$\begin{aligned} & \max \left\{ 0, \int_{\Omega} X dq \right\} \exp(R(\mu_\Omega \| q)) \\ &= \max \left\{ 0, \int_{\Omega \times S \times \Sigma} \sum_{i \in M} \theta_i \cdot \mathbf{1}[(\omega, s, \sigma) \in E_i] d\mu \otimes q(\omega, s, \sigma) \right\} \exp(R(\mu_\Omega \| q)) \\ &= \max \left\{ 0, \sum_{i \in M} \theta_i \cdot \mu \otimes q(E_i) \right\} \exp(R(\mu_\Omega \| q)) \\ &= F(\theta, q). \end{aligned}$$

Therefore, Lemma S6 applied to this random variable X and to $p = \mu_\Omega$ implies that for any $q, q' \in D$, the mapping $\lambda \mapsto F(\theta, \lambda q + (1 - \lambda)q')$ is quasiconvex and lower semicontinuous on $[0, 1]$. Thus, condition 2 in Theorem S2 are satisfied.

Finally, we show that either condition 3 holds for $L = \{\mu_\Omega\}$ and some $\eta > 0$, or Equation (S11) holds trivially with both sides of the equality equal to zero. Thus, there are two cases to consider. The first case is when

$$\inf_{q \in D} \sup_{\theta \in C} F(\theta, q) > 0.$$

In this case, fix any $\eta > 0$ that is strictly less than this value and take $L = \{\mu_\Omega\}$. The set

$$C_\eta^{\mu_\Omega} \equiv \{\theta \in C : F(\theta, \mu_\Omega) \geq \eta\}$$

is closed since C is closed and F is continuous in θ . Given this, and since C is a subset of the finite-dimensional Euclidean space \mathbb{R}^M , the set $C_\eta^{\mu_\Omega}$ is compact if and only if it is bounded. By

⁵⁶Note that the sets $\{(s, \sigma) : (\omega, s, \sigma) \in E_i\}$ are measurable for each $\omega \in \Omega$ and $i \in M$ by Lemma 4.46 in Aliprantis and Border (2006), and hence the function being integrated is indeed measurable.

Lemma S3, there exists $\kappa \in \mathbb{R}$ such that $\theta_i \leq \kappa$ for all $\theta \in C$ and $i \in M$. Let

$$\beta \equiv \min_{i \in M} \mu(E_i) > 0.$$

Then, for any $\theta \in C$ and $i \in M$,

$$\sum_{i' \in M} \mu(E_{i'}) \theta_{i'} \leq \mu(E_i) \theta_i + (1 - \mu(E_i)) \kappa \leq \beta \theta_i + (1 - \beta) \kappa.$$

Thus, since $R(\mu_\Omega \| \mu_\Omega) = 0$ and since $\eta > 0$, for any $\theta \in C_\eta^{\mu_\Omega}$ and $i \in M$, we have

$$\begin{aligned} 0 < \eta &\leq F(\theta, \mu_\Omega) = \sum_{i' \in M} \mu(E_{i'}) \theta_{i'} \leq \beta \theta_i + (1 - \beta) \kappa \\ \implies \theta_i &> -\frac{(1 - \beta) \kappa}{\beta}. \end{aligned}$$

Therefore, the set $C_\eta^{\mu_\Omega}$ is bounded above by κ and bounded below by $-(1 - \beta) \kappa / \beta$. This implies that $C_\eta^{\mu_\Omega}$ is bounded, hence compact. Thus, all of the assumptions of Theorem S2 are satisfied, so we can conclude that Equation (S11) holds.

The second case is when

$$\inf_{q \in D} \sup_{\theta \in C} F(\theta, q) = 0.$$

In this case, since $F \geq 0$ and since

$$\sup_{\theta \in C} \inf_{q \in D} F(\theta, q) \leq \inf_{q \in D} \sup_{\theta \in C} F(\theta, q),$$

Equation (S11) must hold with both sides equal to zero. Thus, in either case, the equation is satisfied. This completes the proof. \blacksquare

S5.3 Proof of Lemma 2

Fix any $\xi \in \text{co}(\Xi)$. By the definition of Ξ and the definition of the convex hull, there exists $n \in \mathbb{N}$ and $(\psi_{\sigma, f}^1, \dots, \psi_{\sigma, f}^n) \in \Psi^{\Sigma \times B}$ and $\alpha_1, \dots, \alpha_n \geq 0$ with $\alpha_1 + \dots + \alpha_n = 1$ such that

$$\begin{aligned} \xi(\omega, s, \sigma) &= \sum_{i=1}^n \alpha_i \int_B \psi_{\sigma, f}^i(f(\omega, s)) d\rho_\sigma(f) \\ &= \int_B \sum_{i=1}^n \alpha_i \psi_{\sigma, f}^i(f(\omega, s)) d\rho_\sigma(f) \\ &= \int_{\mathcal{F}} \int_{\Psi} \psi(f(\omega, s)) d\tau_\sigma(\psi|f) d\rho_\sigma(f), \end{aligned}$$

where we define $\tau \in \mathcal{R}(\Psi|\mathcal{F})$ for each $\sigma \in \Sigma$ and $f \in B$ by⁵⁷

$$\tau_\sigma(\psi|f) = \sum_{i=1}^n \alpha_i \mathbf{1}[\psi = \psi_{\sigma,f}^i].$$

Thus, $\xi = \xi^\tau$.

Conversely, suppose $\xi = \xi^\tau$ for some $\tau \in \mathcal{R}(\Psi|\mathcal{F})$. Since $\tau_\sigma(\cdot|f)$ has finite support for all σ and f , and since Σ and B are finite, the product measure on $\Psi^{\Sigma \times B}$ generated by these measures also has finite support. That is, there exists a product measure ν on $\Psi^{\Sigma \times B}$ with finite support, defined by

$$\nu\left((\psi_{\sigma,f})_{\sigma \in \Sigma, f \in B}\right) = \prod_{\substack{f \in B \\ \sigma \in \Sigma}} \tau_\sigma(\psi_{\sigma,f}|f).$$

We can enumerate the elements of the support of this measure as

$$\text{supp}(\nu) = \{(\psi_{\sigma,f}^1), \dots, (\psi_{\sigma,f}^n)\}.$$

Thus,

$$\begin{aligned} \xi^\tau(\omega, s, \sigma) &= \int_{\mathcal{F}} \int_{\Psi} \psi(f(\omega, s)) d\tau_\sigma(\psi|f) d\rho_\sigma(f) \\ &= \int_B \int_{\Psi^{\Sigma \times B}} \psi_{\sigma,f}(f(\omega, s)) d\nu\left((\psi_{\hat{\sigma}, \hat{f}})_{\hat{\sigma} \in \Sigma, \hat{f} \in B}\right) d\rho_\sigma(f) \\ &= \sum_{i=1}^n \nu\left((\psi_{\hat{\sigma}, \hat{f}}^i)_{\hat{\sigma} \in \Sigma, \hat{f} \in B}\right) \int_B \psi_{\sigma,f}^i(f(\omega, s)) d\rho_\sigma(f), \end{aligned}$$

and hence $\xi^\tau \in \text{co}(\Xi)$.

S5.4 Proof of Lemma 3

The set $[-\infty, \infty]$ is a compact Hausdorff space when endowed with its usual topology.⁵⁸ By the Tychonoff Product Theorem (Theorem 2.61 in Aliprantis and Border (2006)), the set $[-\infty, \infty]^Z$ endowed with the product topology (also known as the topology of pointwise convergence) is compact. Since $\Psi \subset [-\infty, \infty]^Z$ is closed, it is also compact. Applying the Tychonoff Product Theorem again, the set $\Psi^{\Sigma \times B}$ is compact in the product topology.

We next show that the mapping $J : \Psi^{\Sigma \times B} \rightarrow [-\infty, \infty]^{\Omega \times S \times \Sigma}$ defined in Equation (11) is continuous when $[-\infty, \infty]^{\Omega \times S \times \Sigma}$ is endowed with the product topology. To see this, fix any net $(\psi_{\sigma,f}^\alpha)_{\alpha \in D}$ in $\Psi^{\Sigma \times B}$ that converges to some $(\psi_{\sigma,f}) \in \Psi^{\Sigma \times B}$. We will show that $J[(\psi_{\sigma,f}^\alpha)]$ converges

⁵⁷We can define $\tau_\sigma(\cdot|f)$ arbitrarily for $f \in \mathcal{F} \setminus B$.

⁵⁸The topology on $[-\infty, \infty]$ is generated by sets of the form (a, b) , $[-\infty, c)$ and $(c, \infty]$ for $a, b, c \in \mathbb{R}$. It is easy to see that under this topology, $[-\infty, \infty]$ is Hausdorff (meaning that for any two distinct points x, y there exist neighborhoods U of x and V of y such that $U \cap V = \emptyset$) and compact. Indeed, $[-\infty, \infty]$ is often referred to as the two-point compactification of \mathbb{R} (see Example 2.75 in Aliprantis and Border (2006)).

to $J[(\psi_{\sigma,f})]$.⁵⁹ First, by the definition of the product topology, convergence of the net $(\psi_{\sigma,f}^\alpha)$ implies that $\psi_{\sigma,f}^\alpha(z) \rightarrow \psi_{\sigma,f}(z)$ for all σ, f , and z . In particular, $\psi_{\sigma,f}^\alpha(f(\omega, s)) \rightarrow \psi_{\sigma,f}(f(\omega, s))$ for all σ, f, ω , and s . Therefore, since convergence is preserved under scalar multiples and finite sums,

$$\sum_{f \in B} \psi_{\sigma,f}^\alpha(f(\omega, s)) \rho_\sigma(f) \rightarrow \sum_{f \in B} \psi_{\sigma,f}(f(\omega, s)) \rho_\sigma(f)$$

for all ω, s , and σ . Thus, $J[(\psi_{\sigma,f}^\alpha)] \rightarrow J[(\psi_{\sigma,f})]$ in the topology of pointwise convergence on $[-\infty, \infty]^{\Omega \times S \times \Sigma}$.

Therefore, the set $\Xi = J[\Psi^{\Sigma \times B}]$ is compact, since it is the image of the compact set $\Psi^{\Sigma \times B}$ under the continuous function J . Moreover, since $[-\infty, \infty]^{\Omega \times S \times \Sigma}$ is a Hausdorff space, compact subsets of this space are closed (Lemma 2.32 in Aliprantis and Border (2006)). Thus, Ξ is closed.

References for Supplementary Appendix

- Aliprantis, C. and K. Border (2006): *Infinite Dimensional Analysis*, 3rd ed., Berlin, Germany: Springer-Verlag.
- Daniel, K., D. Hirshleifer, and A. Subrahmanyam (1998): “Investor Psychology and Security Market Under- and Overreactions,” *Journal of Finance*, 53, 1839–1885.
- Dudley, R. M. (2002): *Real Analysis and Probability*, 2nd ed., Cambridge, United Kingdom: Cambridge University Press.
- Dupuis, P. and R. S. Ellis (1997): *A Weak Convergence Approach to the Theory of Large Deviations*, New York: John Wiley & Sons.
- Epstein, L. G. and M. Le Breton (1993): “Dynamically Consistent Beliefs Must Be Bayesian,” *Journal of Economic Theory*, 61, 1–22.
- Hanany, E. and P. Klibanoff (2007): “Updating Preferences with Multiple Priors,” *Theoretical Economics*, 2, 261–298.
- Machina, M. J. (1989): “Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty,” *Journal of Economic Literature*, 27, 1622–1668.
- Machina, M. J. and D. Schmeidler (1992): “A More Robust Definition of Subjective Probability,” *Econometrica*, 60, 745–780.
- von Neumann, J. (1928): “Zur Theorie der Gesellschaftsspiele,” *Mathematische Annalen*, 100, 295–320.
- Sion, M. (1958): “On General Minimax Theorems,” *Pacific Journal of Mathematics*, 8, 171–176.
- Tuy, H. (2004): “Minimax Theorems Revisited,” *Acta Mathematica Vietnamica*, 29, 217–229.

⁵⁹It is well known that the product topology on an uncountable product space cannot be completely described by sequential convergence, as such spaces are not metrizable. Although Σ and B are finite, Z could be uncountable. Hence we must use nets to establish the continuity of J .