

An Evolutionary Perspective on Updating Risk and Ambiguity Preferences*

Philipp Sadowski[†] Todd Sarver[‡]

April 6, 2020

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 many established ambiguity and risk preferences as special cases, while also providing a tight connection between ambiguity attitudes and violations of expected utility. We 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, phenotypic flexibility

JEL CLASSIFICATION: D81, D83, D84

*This research was supported by the National Science Foundation through Grant SES-1851664. We thank Jiayun Dong, José Heleno Faro, Cosmin Ilut, Frank Johannes, Bart Lipman, Wolfgang Pesendorfer, Arthur Robson, 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? Second, a plethora of models of ambiguity aversion have been proposed that differ in subtle ways in the behavior they predict. What criterion should be used to select among those? Third, is there a connection between ambiguity aversion (Ellsberg-type behavior) and violations of expected utility in the context of objective risk (Allais-type behavior)?

This paper provides an evolutionary perspective on these issues. Evolution of preferences refers to the notion that natural selection not only can influence physical traits, but can also shape preferences. Using this approach, we develop a foundation for non-expected-utility and ambiguity-averse preferences and study updating of such preferences in response to information. A key finding is that evolutionarily optimal preferences must be dynamically consistent, but may violate consequentialism.¹ Another important contribution of this paper is to expand the scope of the evolutionary approach: Taking inspiration from recent developments in evolutionary biology, we enrich the standard model of asexual reproduction often used in evolutionary economics and introduce phenotypic flexibility, the rapid and reversible adaptation of mature organisms to different conditions. Our result on dynamically consistent updating can then be applied to the resulting rich class of evolutionarily optimal preferences, which we call *adaptive preferences*. We show that adaptive preferences include 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. Importantly, while ambiguity-averse preferences are typically assumed to reduce to expected utility when facing objective risk, our model closely links different uncertainty attitudes to violations of expected utility.

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 behavior (preferences) of an organism. First, those preferences guide the choice of an action that leads to a possibly uncertain outcome. Next, the outcome together with the physical traits of the organism determine its evolutionary fitness, that is, its number of offspring. Rather than modeling physical traits explicitly, we directly specify the mapping that translates outcomes to offspring, which is referred to as a fitness function. 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

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

initially present if the relative size of its subpopulation does not vanish over time.

In economic applications involving evolution of preferences or evolutionary game theory, the fitness function is usually taken as fixed. As noted, in this paper we instead allow phenotypic flexibility: A genotype may express different *phenotypes*, and the individual phenotype can be rapidly and reversibly adapted. We model this rapid adaptation by describing a genotype as a set of possible fitness functions from which one is selected at the time of choosing an action. Naturally, a genotype that takes into account this flexibility when choosing an action has an evolutionary advantage.²

Before describing the contributions of this paper in greater detail, it is worth highlighting three key insights that will be central in our analysis:

1. 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.
2. Phenotypic flexibility leads to non-expected-utility preferences over both idiosyncratic and common uncertainty.
3. Even with all of these components present (idiosyncratic and common uncertainty, phenotypic flexibility), preferences following the arrival of information will be dynamically consistent.

The starting point of our analysis is observation 1 above, which dates back to the seminal paper by Robson (1996). The main innovations of our paper are the incorporation of phenotypic flexibility and the study of preferences following the arrival of information, which yield observations 2 and 3. Importantly, these two innovations are not independent of each other, as the introduction of phenotypic flexibility 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. This, in turn, allows our observation concerning dynamic consistency to be applied to a much wider class of models. In the following paragraphs, we touch briefly on each of the three insights as we introduce the elements of our model.

The intuition behind observation 1 is quite simple. To see how evolution can generate aversion to common uncertainty, suppose there are two actions between which all individuals must choose in every period. For both actions, individual growth (under the best available fitness function) 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

²We discuss our modeling assumptions and how they are motivated by the evolutionary biology literature in more detail in Section 1.1. Specifically, adopting the terminology from Piersma and Drent (2003), phenotypic flexibility refers to rapid and reversible adaptation of a mature individual, which is different from developmental plasticity (irreversible adaptation of an organism to its environment during development).

(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 higher long-run growth, as previously observed by Robson (1996). This implies it will almost surely dominate in the long run (Lemma 1).³

The second part of observation 1 regards the connection between ambiguity and common uncertainty. Note that in many applications of ambiguity, the unknown uncertainty indeed concerns a common factor, which we will refer to as the *environment*, that affects all individuals in the population. Examples include the value of a macroeconomic variable, the efficacy of a medical treatment for a population of individuals, or even the composition of an urn in a lab experiment. In addition, the literature on model misspecification suggests that individuals typically do not know the correct model of the world and therefore treat the “model” itself as a common unknown factor. We discuss and justify this close connection between ambiguity and correlated uncertainty in more detail in Section 1.2. However, strictly speaking, our results will concern the differential treatment of common versus idiosyncratic uncertainty, and 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 key insight to keep in mind is that if ambiguity is frequently associated with common uncertainty—and we would argue that it is—then the evolutionary optimality of correlation aversion provides one possible mechanism through which ambiguity aversion may have originated. To simplify the exposition, we will use the two notions interchangeably for the bulk of the paper.

Before discussing observations 2 and 3, we briefly sketch the key elements of our model and describe our representation for adaptive preferences. Formally, a genotype determines how to evaluate an act f that depends on both the common component of the state of the world, $\omega \in \Omega$, and a component $s \in S$ that is idiosyncratic (distributed identically and independently across individuals) conditional on ω . The joint distribution of ω and s is given by μ . Phenotypic flexibility allows individuals to select a fitness function ψ from a feasible set Ψ , prior to the realization of the state. Recall that the evolutionarily optimal genotype evaluates the act f according to the long-run growth its repeated choice would generate for the population as a whole. Suppose for now that no signal has arrived at the time of choice of f and ψ . We show that the logarithm of the long-run annualized growth rate is then

$$V(f) = \max_{\psi \in \Psi} \int_{\Omega} \ln \left(\int_S \psi(f(\omega, s)) d\mu(s|\omega) \right) d\mu(\omega). \quad (1)$$

³The 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 3.

We refer to the preferences represented by this equation as *adaptive preferences*. If the space of outcomes is ordered and each ψ is concave, then growth is affected negatively by either type of uncertainty—the optimal genotype is uncertainty averse. Note that there are two types of ambiguity in this model, as common sources of ambiguity could enter either through the direct influence of ω on the outcome of the act f or through the effect of ω on the distribution of s , capturing model uncertainty. In either case, the logarithm between the two integrals in Equation (1) means that aversion to objective idiosyncratic risk (a component of s whose distribution does not depend on ω) is less severe than aversion to ambiguity—the optimal genotype is ambiguity averse. The formula obtained by Robson (1996) is the special case where there is a single fitness function, $\Psi = \{\psi\}$ (no phenotypic flexibility).

To gain some sense of the effects of phenotypic flexibility as summarized in observation 2, consider the genotype that has two fitness functions, with the first yielding peak reproduction only for outcome a and the second yielding the same peak reproduction only for outcome b . Suppose further that there are three acts, f , g , and h . Acts f and g give outcomes a and b with certainty, respectively, while act h gives a or b , each with strictly positive probability.⁴ If this genotype has evolutionarily optimal preferences, then it is indifferent between choosing acts f and g (with the corresponding simultaneous choice of fitness function), but it does not like act h , because there is no phenotype that is equally well adapted to outcomes a and b at the same time. This genotype therefore does not maximize an expected-utility preference.⁵

Turning to observation 3, it is well-known that except in special circumstances, non-expected-utility preferences (in particular, ambiguity-averse preferences) cannot satisfy both dynamic consistency and consequentialism. This tension, and the subsequent disagreement in the literature over which property to prioritize, is an impediment to applying these preferences in dynamic contexts where information plays a central role, such as in macroeconomics and finance. Observation 3 leverages the evolutionary perspective that we develop to provide clear guidance on this issue. Our general representation extends the formula for evolutionarily optimal preferences in Equation (1) to evaluate ex ante contingent plans for the choice of act following the arrival of a private signal. We also consider evolutionarily optimal ex post preferences after the actual arrival of the signal. Since evolutionary optimality of both ex ante and ex post preferences requires maximization of the long-run growth rate of the genotype, it becomes almost tautological 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,

⁴Whether h depends on common or idiosyncratic uncertainty is unimportant for this example.

⁵In the context of pure risk, this type of violation of expected utility is explored axiomatically in Sarver (2018). For us, an important question is how these violations of expected utility for idiosyncratic uncertainty interact with attitudes toward common uncertainty, thereby generating a link between risk and ambiguity attitudes.

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 fitness 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.”

The remainder of the paper is structured as follows. The next two subsections discuss in more detail the biological evidence and economic interpretation of phenotypic flexibility and the relationship between correlated uncertainty and ambiguity, respectively. 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.

To illustrate our results and to demonstrate the connection between our evolutionary model and other established models of ambiguity and risk preferences, Section 4 highlights two canonical special cases of adaptive preferences: In Section 4.1, 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 phenotypic flexibility ($\Psi = \{\psi\}$). This special case was previously studied by Robson (1996) for preferences without signals. In Section 4.2, we explore an alternative special case where phenotypic flexibility is permitted but there is no common uncertainty ($\Omega = \{\omega\}$), and we show that the evolutionarily optimal preferences in this case correspond to the optimal risk attitude preferences studied by Sarver (2018). In particular, a dual version of this representation nests rank-dependent expected utility (RDU) as one special case, and we describe the evolutionarily optimal updating of RDU preferences.

In Section 5, we analyze other special cases of our general model when phenotypic flexibility and common uncertainty are simultaneously permitted, including a special case that merges our two leading examples. Via a number of duality results, we show that our adaptive model nests a class of representations where: (i) Beliefs about common uncertainty are first pessimistically distorted using either confidence or variational preferences with an entropic confidence function; and (ii) Beliefs about both common and idiosyncratic uncertainty are then additionally distorted using, among other things, the well-established formulas for rank-dependent utility, divergence preferences, or multiple priors expected utility. Since in all of these representations the second stage of pessimistic belief distortion is applied to both types of uncertainty, each closely links Allais- and Ellsberg-type behaviors. 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.

1.1 Phenotypic Flexibility

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.⁶

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. 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 data only consists of observable choices between outcome-relevant actions (f), while the phenotype (ψ) is not directly observable.

⁶Piersma and Drent (2003) use *phenotypic plasticity* as an umbrella term that includes both phenotypic flexibility and developmental plasticity.

As a consequence, an economic analyst can typically not distinguish whether adaptation is due to a biological change or a strategic but hidden choice of action that affects the reproductive value of different outcomes. For instance, an economic agent who is observed choosing between more or less uncertain investments might, unobserved by the analyst, be able to buy additional insurance. Being insured comes at a cost, but naturally reduces the observed aversion towards uncertainty—in the language of evolutionary biology, the individual becomes more of a generalist and less of a specialist.⁷ Since the analyst cannot distinguish between adaptation via biological changes and via hidden actions, our model allows both interpretations.

1.2 Ambiguity as Common Uncertainty

As noted earlier in the introduction, in many examples and applications of ambiguity, the unknown probabilities concern common factors that affect all individuals in the population. We will refer to these common factors as the *environment*. 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.⁸ 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 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

⁷Other examples of hidden actions that take the form of physical commitments (prior to the resolution of uncertainty) include mortgage agreements or purchases of durable consumption goods. A large literature details the impact of such commitments on risk preferences, for example, 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).

⁸See Epstein and Schneider (2010) for a survey of the large body of subsequent research in this area.

ambiguous, then all individuals share in the resulting common uncertainty.⁹ 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. The main thrust of the preceding discussion is that there are indeed many situations in which ambiguity and risk are tightly linked to common and idiosyncratic uncertainty, respectively, 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.

2 Evolutionary Setting

2.1 Uncertainty and Information

Common uncertainty about the environment is modeled via a state space Ω . The realization of the environment ω is common to all individuals in the population. In addition, given the environment, 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) .¹⁰ 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 $\mathcal{E} = \mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma$.

Uncertainty is described by a measure μ in the set $\Delta(\Omega \times S \times \Sigma)$ of (countably additive)

⁹This interpretation is closely connected to the macroeconomic literature on robustness to model uncertainty (Hansen and Sargent (2001, 2008)). It is further discussed in the context of the evolutionary model in Robson (1996), who noted that in many examples “aggregate uncertainty might be reinterpreted as a lack of precision concerning the risk facing an individual.”

¹⁰Since S describes idiosyncratic risk, it is natural to consider private signals. Samuelson (2004) and Nöldeke and Samuelson (2005) also incorporated private signals in an evolutionary framework with common uncertainty. 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 .

probability measures on the measurable space $(\Omega \times S \times \Sigma, \mathcal{E})$. The marginal distribution of μ on Ω assigns probability $\mu(E)$ to any measurable event $E \in \mathcal{B}_\Omega$. When necessary to avoid confusion, we will sometimes denote this marginal distribution more explicitly by μ_Ω . 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 $S \times \Sigma$.¹¹ 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 the set of outcomes. We assume that Z is a convex subset of a vector space. 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 functions $f : \Omega \times S \rightarrow Z$ that take a finite number of possible values. An evolutionary fitness function $\psi : Z \rightarrow \mathbb{R}$ specifies the (expected) individual growth rate associated with each outcome.¹² Given an act $f \in \mathcal{F}$, the individual growth rate in state (ω, s) is then $\psi(f(\omega, s))$. For example, 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 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 evaluating acts 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 evaluation, contingent on ψ , μ , and σ . We refer to a particular fitness function ψ (together with these preferences) as a phenotype. Building on insights from evolutionary biology, we assume that a genotype does not necessarily fully determine the

¹¹More 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).

¹²Realized individual growth must, of course, amount to an integer number of offspring, but since reproduction may be uncertain given the outcome $z \in Z$, the expected individual growth rate 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 individual growth rates $\psi(z)$ its individuals attain from each outcome z .

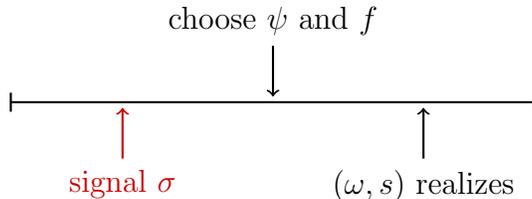


Figure 1: Timeline: after-signal adaptation

expression of a phenotype, but is instead constrained by a set of fitness functions Ψ . For instance, different phenotypes may attain their peak reproduction for different outcomes, or some may be specialists with high peak reproduction while others are generalists with lower peak reproduction that is more robust to the outcome. Since preferences over acts can condition on ψ , we model phenotypic flexibility simply as the selection of ψ .

Throughout the main text, we assume that the expressed phenotype can be rapidly adapted to changing environmental conditions, which in the context of our model means the phenotype can be adjusted following the realization of the signal. Thus, both the phenotype ψ and the act f are chosen after observing the signal σ . The timing of information and the choice of phenotype and act are illustrated in Figure 1. In Section S2 of the Online Appendix, we consider the alternative case where adaptation is a slower process and the phenotype must be selected prior to learning the realization of the signal.

A central theme of our analysis will be the connection between the set Ψ of available fitness functions and the corresponding preferences under uncertainty. The appropriate set Ψ itself will depend on the choice context. For instance, it may reflect the range of biological changes that are feasible for an organism, or it may reflect the set of hidden actions that are available to each individual.¹³ For now, we only impose the technical assumptions that Ψ is a nonempty and convex set and that $\sup_{\psi \in \Psi} \psi(z)$ is finite for every $z \in Z$.¹⁴ Further restrictions on Ψ will determine the exact model of choice under uncertainty.

2.3 Growth Rates

The choice of act f and fitness function ψ by a genotype will determine the overall growth rate of its subpopulation in a given time period. We assume each decision problem is faced repeatedly, leading to a stochastic sequence of growth rates for each genotype. Our analysis

¹³In the former case, Ψ itself may be subject to evolutionary selection, and thus should be optimal for the (possibly changing) choice setting, given the relevant physical constraints. Investigating the evolutionary selection of Ψ and conditions for heterogeneity in steady state are topics for future research.

¹⁴In Section S1 of the Online Appendix, we discuss how when the set of phenotypes Ψ is not convex, it may be optimal for the genotype to program the individuals in its subpopulation to engage in idiosyncratic self-randomization in their choice of ψ . We show that such idiosyncratic randomization effectively convexifies the set of phenotypes Ψ . In this sense, the assumption that the set Ψ is convex is without loss of generality.

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 evolutionary model, the long-run growth rate exists for any act f and phenotype 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 .¹⁵ 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

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

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.

¹⁵Using 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.

3 Evolutionarily Optimal Preferences and Updating

It is well-known that for preferences that violate Savage’s Sure Thing Principle—including both ambiguity-averse preferences and probabilistically sophisticated non-expected-utility preferences—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 preferences should respond to new information: Hanany and Klibanoff (2007, 2009) 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).

The tension between consequentialism and dynamic consistency in models of ambiguity aversion or non-expected utility for risk, and the subsequent disagreement over which property to give priority, is an impediment to applying these models in dynamic contexts such as macroeconomics and finance where information plays a central role. In this section, we use our evolutionary model in order to study updating of preferences in response to information. We will show, in particular, that evolutionarily optimal behavior will be dynamically consistent, even at the expense of consequentialism. Importantly, the evolutionary approach will provide a novel rationale for violations of consequentialism, showing that such violations should be neither surprising nor concerning. In Section 4, we further illustrate these results using two special cases of our model, the smooth model of ambiguity aversion and rank-dependent expected utility. Section 5 explores additional special cases.

Signals arrive before the choice of an act. To analyze dynamic choice in general, and dynamic consistency in particular, it is necessary to compare ex post behavior after the arrival of information to the ex ante plan that was formed prior to the realization of the signal. In other words, we need to first analyze the plan that an individual would form if she could commit ex ante to her signal-contingent choices, and then compare this plan to her actual ex post choices to see if the individual deviates from her intended actions. We therefore begin our analysis by describing the evolutionarily optimal ex ante plans of action, after which we proceed to study the evolutionarily optimal ex post updating of preferences.

3.1 Ex Ante Plans and Preferences

Since the signal resolves prior to the choice of act and phenotype (see Figure 1), the individual can select both of them conditional on the signal.

Definition 2. An *action plan* is a function $\mathbf{f} : \Sigma \rightarrow \mathcal{F}$ from the space of signals to the set of acts. Given a set Ψ of available fitness functions, an *adaptation plan* is a function $\psi : \Sigma \rightarrow \Psi$ from the space of signals to that set.

An action plan \mathbf{f} specifies a path through a decision tree, where the act $f = \mathbf{f}_\sigma$ is selected following the signal $\sigma \in \Sigma$. Likewise, an adaptation plan ψ selects the phenotype $\psi = \psi_\sigma$ after the signal σ . Let \mathcal{F}^Σ and Ψ^Σ denote the set of all action and adaptation plans, respectively. From the ex ante perspective, before the realization of a signal, the individual forms a joint plan (\mathbf{f}, ψ) over acts and phenotypes, which achieves a fitness of $\psi_\sigma(\mathbf{f}_\sigma(\omega, s))$ after the realization of (ω, s, σ) .

We begin by formalizing the evolutionarily optimal ex ante preferences over action plans. 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. *Suppose Ψ and μ are fixed, and consider a genotype with an (infinitely) large subpopulation of individuals. If the adaptation plan $\psi \in \Psi^\Sigma$ is chosen optimally, then the long-run growth rate of the genotype from choosing the action plan $\mathbf{f} \in \mathcal{F}^\Sigma$ in every period is*

$$V(\mathbf{f}) = \sup_{\psi \in \Psi^\Sigma} \int_{\Omega} \ln \left(\int_{S \times \Sigma} \psi_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma | \omega) \right) d\mu(\omega). \quad (2)$$

We refer to the ex ante preferences represented by the value function V in Equation (2) as *adaptive preferences*. The proof of Theorem 1 is based on the same logic that is behind the seminal result of Robson (1996), who considered the special case of no signals ($\Sigma = \{\sigma\}$) and no phenotypic flexibility ($\Psi = \{\psi\}$). In that special case of our model, signals and phenotypes can be dropped from the representation, and action plans reduce to acts.

Corollary 1 (Robson (1996)). *Suppose $\Sigma = \{\sigma\}$, and suppose $\Psi = \{\psi\}$ and μ are fixed. The long-run growth rate of a genotype with an (infinitely) large subpopulation of individuals from choosing the act f in every period is*

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

In a context with signals, Nöldeke and Samuelson (2005) considered a representation of ex ante preferences similar to that in Equation (2), but also in the absence of adaptation ($\Psi = \{\psi\}$). The survey by Robson and Samuelson (2011) summarizes these papers and much of the other recent literature on evolution of preferences.

Proof of Theorem 1. Recall that, conditional on the common environment ω , both the s dimension of the state and the signal σ are independently distributed for each individual in the population. Therefore, by the law of large numbers, conditional on the environment ω_t at time t , the average growth rate of a large population of individuals choosing a particular action plan \mathbf{f} and adaptation plan $\boldsymbol{\psi}$ is approximately

$$\lambda_t(\omega_t) = \int_{S \times \Sigma} \boldsymbol{\psi}_\sigma(\mathbf{f}_\sigma(\omega_t, s)) d\mu(s, \sigma | \omega_t).$$

Since we consider infinite subpopulations in our model, we can treat this approximation as exact.¹⁶ Taking the product over a sequence of realized environments $\omega_1, \dots, \omega_T$ and raising to the power $1/T$ gives the realized annualized growth rate over this sequence of periods:

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

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

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \ln \left(\int_{S \times \Sigma} \boldsymbol{\psi}_\sigma(\mathbf{f}_\sigma(\omega_t, s)) d\mu(s, \sigma | \omega_t) \right) \\ \rightarrow \int_{\Omega} \ln \left(\int_{S \times \Sigma} \boldsymbol{\psi}_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma | \omega) \right) d\mu(\omega) \quad \text{a.s.}, \end{aligned} \quad (4)$$

by the law of large numbers.¹⁷ Optimizing over $\boldsymbol{\psi} \in \Psi^\Sigma$ gives precisely Equation (2). ■

Since the individual selects both the adaptation plan $\boldsymbol{\psi}$ and the action plan \mathbf{f} , the long-run growth rate of the population is optimized by choosing both to maximize the limit in Equation (4). However, if only the choice of act is observed—and the choice of fitness function corresponds to an unobservable selection of phenotype or some other hidden action—then the highest possible long-run growth rate associated with each different choice of action plan \mathbf{f} is given by Equation (2). Note that since V expresses the long-run average growth rate in log terms, $V(\mathbf{f}) = -\infty$ corresponds to extinction and $V(\mathbf{f}) = 0$ corresponds to constant population size. By Lemma 1, the evolutionarily optimal genotype is the one that maximizes

¹⁶Note that an approximate (limiting) version of this theorem 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 $V(\mathbf{f})$ conditional on non-extinction. Moreover, it can be shown that when $V(\mathbf{f}) > 0$, the probability of extinction converges to zero as the initial population becomes large.

¹⁷Note that the integral in Equation (4) may be either finite or $-\infty$, but it cannot take the value $+\infty$ since \mathbf{f} takes finitely many possible values and $\sup_{\boldsymbol{\psi} \in \Psi} \boldsymbol{\psi}(z)$ is assumed to be finite for every outcome $z \in Z$. It is a standard result that the strong law of large numbers continues to hold for random variables that are bounded above, with almost sure convergence to $-\infty$ in the case where the expectation is $-\infty$. See, for example, the corollary to Theorem 22.1 in Billingsley (1995).

the long-run growth rate, so the most successful genotype will be the one with adaptive preferences, that is, preferences over ex ante action plans given by $\mathbf{f} \succsim \mathbf{f}' \iff V(\mathbf{f}) \geq V(\mathbf{f}')$.

There is one important consideration that we have thus far ignored: When choosing between two action plans \mathbf{f} and \mathbf{f}' , a genotype may be able to attain a greater long-run growth rate than either $V(\mathbf{f})$ or $V(\mathbf{f}')$ by programming the individuals in its subpopulation to engage in idiosyncratic randomization between these two action plans. We address the possibility of self-randomization as a form of hedging formally in Section S1 of the Online Appendix. For the moment, it is worth observing that when the choice set of acts F is convex and each $\psi \in \Psi$ is concave, there is no evolutionary benefit from self-randomization in the choice of action plan.

Corollary 2. *Suppose Ψ and μ are fixed and each $\psi \in \Psi$ is a concave function. Then, the genotype that maximizes adaptive preferences achieves a weakly higher long-run growth rate than all others for every infinitely repeated convex choice set $F \subset \mathcal{F}$.*

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 1.2, 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 . The following no-information example illustrates.

Example 1 (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 a bet that pays \$1 if the ball drawn is black, BY indicates a bet that pays \$1 if the ball is either black or yellow, and so on. Formally, B is the act defined by $B(b) = 1$, $B(r) = 0$, and $B(y) = 0$, and similarly for the other bets. 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 of environments $\Omega = \{\omega_1, \omega_2, \omega_3\}$, 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 environments, this correlated uncertainty is treated differently than uncorrelated uncertainty. For ease of illustration, suppose μ assigns equal probability to each urn composition, and assume for expositional simplicity that there is a single phenotype ψ that takes values $\psi(0) = 0$ and $\psi(1) = 1$. Since Equation (3) is a smooth model 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

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

In Example 1, 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. If one is not convinced that the Ellsberg urn is a perfect fit for this interpretation, the objects in the example can of course 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, with B being a better understood treatment than R .¹⁸

In Section 4, we provide further analysis of two canonical examples of our model. In Section 4.1, we revisit Example 1 and use it to illustrate evolutionarily optimal updating of the smooth model of ambiguity aversion following signals. In Section 4.2, we show that the richness of our general model with phenotypic flexibility allows it to also nest rank-dependent expected utility, and we describe the evolutionarily optimal updating of RDU preferences. Section 5 then expands on these examples by providing a pair of duality results

¹⁸Specifically, if the treatment B has been well-studied, then its success rate might be known to be $1/3$. If the treatments R and Y are typically administered together, their joint success rate might be known to be $2/3$, yet their individual success rates remain unknown. This is obviously stylized, since the events where each treatment is successful are typically not disjoint and there may be interactions between treatments.

that permit a rich class of additional special cases to be characterized, including one that merges ambiguity-averse preferences with rank-dependent utility over risk.

Before proceeding to the analysis of ex post preferences, we should first highlight several assumptions that are implicit in our formulation of the evolutionary model. Corollary 2 shows that the long-run growth rate is optimized by choosing the action plan $\mathbf{f} \in F^\Sigma$ that maximizes V , assuming the choice set F 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 choice sets $(F_t)_{t \in \mathbb{N}}$, then attaining the highest possible long-run growth rate requires that individuals maximize adaptive preferences from any set F that repeats with fixed frequency within this sequence.¹⁹

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 phenotypic flexibility 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.

3.2 Ex Post Preferences

Denote the ex post preferences after observing the signal σ by $\succsim_{\sigma, \mathbf{f}}$. Note that we allow for the possibility that ex post preferences depend on the ex ante action plan. In the case where these preferences do not depend on the ex ante plan, we say that preferences are consequentialist.

Definition 3. Preferences satisfy *consequentialism* if the ex post ranking of any acts f and g is independent of the ex ante action plan: $f \succsim_{\sigma, \mathbf{f}} g \iff f \succsim_{\sigma, \mathbf{f}'} g$.

A second important property of preferences is dynamic consistency. In what follows, we write $g\sigma\mathbf{f}$ to denote the action plan that selects the act g following the signal σ and selects the act $\mathbf{f}_{\sigma'}$ following each signal $\sigma' \neq \sigma$.

¹⁹The assumption that all individuals of the genotype face the same choice set 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 choice sets 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).

Definition 4. Preferences satisfy *dynamic consistency* if ex ante plans are not reversed ex post: $\mathbf{f} \succ g\sigma\mathbf{f} \implies \mathbf{f}_\sigma \succ_{\sigma,\mathbf{f}} g$ and, in addition, $\mathbf{f} \succsim g\sigma\mathbf{f} \implies \mathbf{f}_\sigma \succsim_{\sigma,\mathbf{f}} g$ whenever $\mu(\sigma) > 0$.

Much of the literature on the updating of ambiguity preferences has focused on partitional information structures. Note that partitional learning fits easily within our general framework, and in this special case, our definition of dynamic consistency is similar to the definitions proposed previously in the literature, for example, by Machina and Schmeidler (1992), Epstein and Le Breton (1993), or Hanany and Klibanoff (2007).²⁰

As noted above, ambiguity-averse preferences or non-expected-utility preferences for risk cannot in general satisfy both consequentialism and dynamic consistency, as we will illustrate in Section 4.1. Fortunately, the evolutionary approach gives clear guidance about which property to favor: The evolutionarily optimal ex ante preferences are precisely those that maximize the long-run growth rate of the genotype. The evolutionarily optimal ex post preferences have the same objective. Thus, dynamic consistency is necessarily satisfied, as the following results demonstrate.

Corollary 3. *Suppose Ψ and μ are fixed, and suppose the genotype forms an action plan $\mathbf{f} \in \mathcal{F}^\Sigma$ ex ante, which it follows after every signal $\sigma' \neq \sigma$, but it deviates from this plan after signal σ by instead selecting the act g . Then, its long-run growth rate is*

$$V(g|\sigma, \mathbf{f}) = \sup_{\psi \in \Psi^\Sigma} \int_{\Omega} \ln \left(\mu(\sigma|\omega) \int_S \psi_\sigma(g(\omega, s)) d\mu(s|\omega, \sigma) + \int_{S \times \Sigma \setminus \{\sigma\}} \psi_{\sigma'}(\mathbf{f}_{\sigma'}(\omega, s)) d\mu(s, \sigma'|\omega) \right) d\mu(\omega). \quad (5)$$

Corollary 3 follows directly from Theorem 1. The growth rate formula in Equation (5) simply evaluates g (following σ) in conjunction with the ex ante plan \mathbf{f} (following other signals) according to Equation (2), that is, $V(g|\sigma, \mathbf{f}) = V(g\sigma\mathbf{f})$. From this it follows that preferences are dynamically consistent.

Corollary 4. *Suppose that ex ante preferences are represented by Equation (2) and ex post preferences are represented by Equation (5), so $\mathbf{f} \succsim \mathbf{f}' \iff V(\mathbf{f}) \geq V(\mathbf{f}')$ and $f \succsim_{\sigma,\mathbf{f}} g \iff V(f|\sigma, \mathbf{f}) \geq V(g|\sigma, \mathbf{f})$. Then, preferences satisfy dynamic consistency.*

Given the tension between dynamic consistency and consequentialism, one implication of Corollary 4 is that preferences may violate consequentialism. In Section 4.1, we discuss such a violation in the context of the Ellsberg example. Since consequentialism is generally considered to be an appealing property, such violations may seem counterintuitive. However,

²⁰Partitional 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, an action plan can be reduced to an act by defining $f(s) = \mathbf{f}_\sigma(s)$ for $s \in \sigma \in \Sigma$. Our definition of dynamic consistency then reads as $f \succ g\sigma f \implies f \succ_{\sigma,f} g$ and $f \succsim g\sigma f \implies f \succsim_{\sigma,f} g$.

the evolutionary approach not only predicts that dynamic consistency will be maintained at the possible expense of consequentialism, it also provides some perspective on why consequentialism may be violated: 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 \mathbf{f} . 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 (5) 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 representation of conditional preferences, as it brings out the underlying reason for the dependence of updated preferences on the plan.

4 Illustrative Examples

In this section, we describe two canonical special cases of our evolutionary model that help to demonstrate its scope and further clarify its connection with the existing literature on ambiguity and risk preferences: In Section 4.1, we describe evolutionarily optimal updating for the smooth model of ambiguity aversion. This corresponds to the case of no phenotypic flexibility ($\Psi = \{\psi\}$) that was previously studied in Example 1. In Section 4.2, we consider the case with phenotypic flexibility but where there is only idiosyncratic uncertainty ($\Omega = \{\omega\}$). We show that rank-dependent utility is nested by this special case of our representation and we describe the optimal updating of RDU preferences.

4.1 Smooth Model of Ambiguity Aversion

In the special case of a single fitness function, $\Psi = \{\psi\}$, the ex ante and ex post formulas for growth rates given above specialize as follows. The long-run growth rate in Equation (2) from choosing the action plan $\mathbf{f} \in \mathcal{F}^\Sigma$ becomes

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

The long-run growth rate in Equation (5) from deviating from the plan \mathbf{f} by instead selecting g following the signal σ becomes

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

In particular, since evolutionarily optimal ex ante preferences are represented by $V(\mathbf{f})$ and optimal ex post preferences are represented by $V(g|\sigma, \mathbf{f})$, dynamic consistency is satisfied.

We now return to the special case of the smooth model from Example 1 (where acts f depend only on s) to illustrate the tension between consequentialism and dynamic consistency.

Example 2 (Ellsberg—with signals). *As in Example 1, let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ denote the possible urn compositions, so $\omega_1 = (b, r, r)$, $\omega_2 = (b, r, y)$, and $\omega_3 = (b, y, y)$. Recall that when μ assigns equal probability to each urn composition, ex ante preferences satisfy the typical Ellsberg pattern:*

$$B \succ R \quad \text{and} \quad BY \prec RY.$$

Now, suppose that individuals each receive a private signal that tells them whether the ball drawn for them is yellow (y) or not yellow ($\neg y$).²¹ 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:

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

Therefore, dynamic consistency requires that

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

²¹Formally, for each $\omega \in \Omega$, we have $\mu(y|s, \omega) = 1$ if $s = y$ and $\mu(\neg y|s, \omega) = 1$ if $s = b, r$.

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 together with this specific information structure cannot satisfy both dynamic consistency and consequentialism.²² Our model and results imply that individuals with the ex ante preferences $B \succ R$ and $BY \prec RY$ will exhibit the conditional ex post preferences listed above. Thus, individuals will be dynamically consistent but will violate consequentialism. The reason that consequentialism is violated is that 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 in the environments ω in which the probability of signal y (and hence state y) is lower, as the first two columns of the following table 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	$1 \cdot \frac{1}{3} = \frac{1}{3}$
$\omega_2 = (b, r, y)$	$\frac{2}{3} \cdot \frac{1}{2} = \frac{1}{3}$	$\frac{1}{3}$	$\frac{2}{3} \cdot \frac{1}{2} = \frac{1}{3}$
$\omega_3 = (b, y, y)$	$\frac{1}{3} \cdot 0 = 0$	$\frac{2}{3}$	$\frac{1}{3} \cdot 1 = \frac{1}{3}$

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 environment, as shown in the last column of the table. When the ex ante plan is instead

²²That is, a violation of dynamic consistency follows directly from Ellsberg behavior and the assumption that $B \succ R$ implies $B\neg yB \succ R\neg yB$ and $BY \prec RY$ implies $B\neg yY \prec R\neg yY$. Note, 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).

$B \neg y B$, 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:

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

4.2 Rank-Dependent 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.

To focus on risk preferences, we restrict attention for now to the special case of our model with non-degenerate after-signal adaptation, $\psi \in \Psi^\Sigma$, but without common uncertainty, $\Omega = \{\omega\}$. In this case, the expectation over Ω can be dropped from the ex ante value function in Equation (2), which now becomes

$$\begin{aligned} V(\mathbf{f}) &= \sup_{\psi \in \Psi^\Sigma} \ln \left(\int_{S \times \Sigma} \psi_\sigma(\mathbf{f}_\sigma(s)) d\mu(s, \sigma) \right) \\ &= \ln \left[\int_\Sigma \sup_{\psi \in \Psi} \left(\int_S \psi(\mathbf{f}_\sigma(s)) d\mu(s|\sigma) \right) d\mu(\sigma) \right]. \end{aligned} \tag{6}$$

Note that in this case the logarithm can also be dropped by taking a monotone transformation, but we will retain it for consistency of 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.²³

Proposition 1 (Rank-Dependent Utility Duality). *Suppose $\Omega = \{\omega\}$ and $Z \subset \mathbb{R}$. 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*

²³The optimal risk attitude preferences from Sarver (2018) correspond to the representation in Equation (6) without signals and without the log. Proposition 1 can be found (with minor differences in assumptions) in Chatterjee and Krishna (2011) or in the supplementary appendix of Sarver (2018). The key steps to this result also appear in several earlier papers, including Machina (1984) and Wakker (1994).

functions $\psi : Z \rightarrow \mathbb{R}$ that is convex, pointwise bounded above,²⁴ and closed in the topology of pointwise convergence such that, for any act $f \in \mathcal{F}$ and any $\mu \in \Delta(S)$,

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

where

$$F_{f,\mu}(z) = \int_S \mathbf{1}[f(s) \leq z] d\mu(s)$$

is the cumulative distribution function of f given μ . Moreover, if Z is an interval and the function u is concave, then there exists a set Ψ satisfying the conditions above such that each $\psi \in \Psi$ is a concave function.

To interpret this result, consider first the case of no signals. Then, for the set Ψ as in Proposition 1, the ex ante value function for adaptive preferences can be written as a monotone transformation of a rank-dependent utility function:

$$V(f) = \ln \left(\int_Z u(z) d(\varphi \circ F_{f,\mu})(z) \right).$$

Since μ only captures idiosyncratic uncertainty in this section and since we identify idiosyncratic uncertainty with pure risk, $F_{f,\mu}$ amounts to the cumulative distribution of an unambiguous risky prospect. The rank-dependent utility representation with concave φ thus 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.

This connection readily extends to the case with signals. Applying Proposition 1 to the acts \mathbf{f}_σ and the measures $\mu(\cdot|\sigma)$ in Equation (6) immediately yields the following corollary.

Corollary 5 (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 (6) can be equivalently expressed as*

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

and ex post adaptive preferences are represented by

$$\hat{V}(g|\sigma, \mathbf{f}) = \int_Z u(z) d(\varphi \circ F_{g, \mu(\cdot|\sigma)})(z),$$

²⁴As discussed in Section 2, exact dual characterizations of some special cases of our model, such as this one, require some $\psi \in \Psi$ to take negative values for some outcomes. Here, each ψ is a bounded function, but the set Ψ need not be pointwise bounded below and we may have $\inf_{\psi \in \Psi} \psi(z) = -\infty$ for some $z \in Z$.

where

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

is the cumulative distribution function of g given μ and σ .

We denote the ex post value function in the corollary by $\hat{V}(g|\sigma, \mathbf{f})$ rather than $V(g|\sigma, \mathbf{f})$ to indicate that it not only differs from the value function in Equation (5) 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 since it is not necessary to consider what happens after signals $\sigma' \neq \sigma$ in the case without common uncertainty.

In contrast to the examples considered in Machina (1989), ex post preferences in Corollary 5 are actually independent of the plan \mathbf{f} , 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, any idiosyncratic risk that resolves before the selection of the phenotype is evaluated in accordance with expected utility. This is reflected by the ex ante value function, where only the cdf $F_{\mathbf{f}_\sigma, \mu(\cdot|\sigma)}$ over outcomes *given* σ is distorted by φ , rather than the cdf that also incorporates uncertainty about the realization of σ itself.

In Section S2 of the Online Appendix, we show that if adaptation is slower, so that ψ has to be chosen before the realization of σ , then the distortion function φ applies to all uncertainty, including the signal realization. In that case, the rank of an outcome depends on the entire plan \mathbf{f} , and by dynamic consistency the ex post value of act g must also depend on \mathbf{f} , violating consequentialism. This is the approach to modeling rank-dependent utility with information suggested by Machina (1989).

5 Duality Results and Other Special Cases

Other well-known models of choice under uncertainty are typically founded on behavioral axioms rather than evolution. We already observed in Section 4 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 use duality results to recast our representation in a form that permits direct comparisons with other behaviorally founded utility representations. 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, action plans \mathbf{f} reduce to acts f and the value function in Equation (2) reduces to that of Equation (1). Appendix A contains the corresponding representations and theorems for the general case with signals.

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 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 two theorems that reformulate our representation in a way that facilitates the analysis of this and other special cases. These theorems follow as special cases of their counterparts for the case of general signal structures, which can be found in Appendix A. We then proceed to study several special cases in detail in Sections 5.1 and 5.2.

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

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

The notation $q \ll p$ indicates that q is *absolutely continuous* with respect to p , that is, for any measurable set A , $p(A) = 0$ implies $q(A) = 0$. The term $\frac{dq}{dp}$ denotes the Radon–Nikodym derivative (density) of q with respect to p , which exists if and only if q is absolutely continuous with respect to p .²⁵ It is a standard result that $R(q \parallel p) \geq 0$, with equality if and only if $q = p$.

In what follows, for any probability measure $q \in \Delta(\Omega)$, let

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

In particular, since $R(q \parallel p) < \infty$ requires that $q \ll p$, if $p \in M(q)$ then the measures p and q are mutually absolutely continuous, that is, both $p \ll q$ and $q \ll p$.²⁶ Also, recall that we take $\ln(x) = -\infty$ for all $x \leq 0$, and we use μ_Ω to denote the marginal of μ on Ω . Finally, in order to accommodate certain special cases, it will be technically convenient to permit the fitness functions ψ to take the value $-\infty$. That is, throughout this section we assume that Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$.

Theorem 2. *Suppose Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$, and fix μ . For any act $f \in \mathcal{F}$, the function V defined by Equation (1) can be equivalently expressed as*

$$V(f) = \sup_{\psi \in \Psi} \inf_{p \in M(\mu_\Omega)} \left[\ln\left(\int_\Omega \int_S \psi(f(\omega, s)) d\mu(s|\omega) dp(\omega)\right) + R(\mu_\Omega \parallel p) \right]. \quad (7)$$

Despite the resemblance, the functional in Equation (7) with a single phenotype $\Psi = \{\psi\}$

²⁵Formally, $\frac{dq}{dp}$ is the integrable function that satisfies $q(A) = \int_A \frac{dq}{dp} dp$ for any measurable set A .

²⁶Note that it is possible to have $R(q \parallel p) = \infty$ even if $q \ll p$, so $M(q)$ may be a strict subset of the set of all measures that are mutually absolutely continuous with respect to q .

is not a variational representation (Maccheroni, Marinacci, and Rustichini (2006)), although we will see in Section 5.2 that special cases of variational preferences can be accommodated by our general model. The distinction is the logarithm around the integral in the first term. In fact, in the case of a single phenotype, taking the exponential transformation of the representation in Equation (7) establishes it as a special case of the confidence preferences studied by Chateauneuf and Faro (2009), where confidence in a prior p is measured by $\exp(R(\mu_\Omega \| p))$. More generally, this single-phenotype 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),²⁷ and also within a version of confidence preferences in Chateauneuf and Faro (2012). However, in these models, the entropy term used is $R(p \| \mu_\Omega)$ rather than $R(\mu_\Omega \| p)$. 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 (7), 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 p on Ω . Of course, the larger the sample, the closer to zero the probability that it would be generated by $p \neq \mu_\Omega$. The theory of large deviations establishes that the rate at which this probability vanishes increases in $R(\mu_\Omega \| p)$.²⁸ The representation suggests, therefore, that the decision-maker is less confident in a measure p the faster it becomes implausible with growing sample size.

The limitation of Equation (7) is that we would like to reverse the order of the supremum and infimum in order to further simplify it and connect with existing functional forms. Fortunately, this is possible in some cases. Our next result builds on Theorem 2 together with a version of the von Neumann–Sion minimax theorem (von Neumann (1928), Sion (1958)) that is due to Tuy (2004). Roughly speaking, the minimax theorem permits the order of the supremum and infimum to be reversed, provided the objective function satisfies suitable continuity and quasiconcavity/quasiconvexity properties in the respective variables. Most of these necessary conditions for the minimax theorem follow directly from the functional form of our representation in Equation (7) together with our previous assumptions that the set Ψ is convex and pointwise bounded above, that is, $\sup_{\psi \in \Psi} \psi(z) < \infty$ for every $z \in Z$. The

²⁷Hansen and Sargent (2001) interpret their representation in terms of a concern about robustness to model misspecification. Our approach can provide a different perspective on this type of robustness in contexts where uncertainty about ω can be interpreted as model uncertainty.

²⁸If μ_Ω is not rational, it cannot be an empirical frequency in any finite sample, but one can instead consider a sequence of frequencies that converge to μ_Ω as the sample size increases. See, for example, the treatment of Sanov’s theorem in Cover and Thomas (2006, Section 11.4) or Rassoul-Agha and Seppäläinen (2015, Section 5.2).

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 Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$ that is convex, pointwise bounded above, and closed in the topology of pointwise convergence (on the extended reals), and fix μ . For any act $f \in \mathcal{F}$, the function V defined by Equation (1) can be equivalently expressed as*

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

Using this result, we show in the following subsections that adaptive preferences are quite general by discussing a non-exclusive list of special cases they nest. The generality of our model, which reflects that we imposed only minimal restrictions on the set Ψ of available fitness functions, has two main benefits. First, our arguments for dynamically consistent updating are broadly applicable. Second, adaptive preferences can be used to compare different models of choice under uncertainty via the corresponding constraints on Ψ . Thus, imposing restrictions on Ψ turns our evolutionary approach into a tool for model selection.

Since different models that can be nested as special cases imply different connections between risk and ambiguity attitudes, restrictions on Ψ also generate predictions about this connection. Suppose, for instance, that all possible 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 .²⁹

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

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

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

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

²⁹We have taken the set Ψ as given throughout. In order to compare individuals with adaptive preferences who have different sets of phenotypes, it is important to understand the determinants of Ψ . For instance, suppose Ψ itself is subject to evolutionary selection. Heuristically, if the environment were such that evolutionary pressure gave a severe disadvantage to genotypes who perform worse for deterministic outcomes, then the restriction that surviving genotypes should all have the same upper envelope of Ψ would be plausible.

5.1 Special Cases with Confidence Preferences

Note that the argument of the logarithm in the representation of Equation (9) affects the attitude towards uncertainty from S and Ω alike. We now apply another round of duality to rewrite this term in more familiar forms. For now, we leave untouched the outer part of the representation, which determines additional aversion to uncertainty from Ω via entropic confidence preferences.

5.1.1 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 following corollary follows immediately from that same duality result by replacing S with $\Omega \times S$ and replacing the measure $\mu \in \Delta(S)$ with $\mu \otimes p \in \Delta(\Omega \times S)$. Note that this application of Proposition 1 is only possible because we first apply Theorem 3 to remove the logarithm from between the two layers of integration.

Corollary 6. *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 $\psi : Z \rightarrow \mathbb{R}$ such that, for any act $f \in \mathcal{F}$ and any $p \in \Delta(\Omega)$,*

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

where

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

is the cumulative distribution function of f given $\mu \otimes p$. Therefore, for that set Ψ , the function V defined by Equation (1) can be equivalently expressed as

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

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 that the RDU representation is embedded in. The outer part is fixed across genotypes, even if those differ in terms of Ψ , and hence in terms of their attitude towards risk.³⁰ It is easy to see that the upper envelope of Ψ is u . Thus, if the upper envelope of Ψ is fixed, the only free parameter in the representation is the distortion function φ , and it follows easily that more risk aversion (aversion to uncertainty that does not depend on ω) implies more uncertainty aversion (aversion to any type of uncertainty, including both risk and ambiguity).

³⁰There is some empirical evidence that risk aversion and additional aversion to ambiguity indeed have little correlation in the population (Chapman et al. (2019)).

5.1.2 Divergence Preferences

Definition 5. 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

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

Kullback–Leibler relative entropy corresponds to the 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 $r \in \Delta(\Omega \times S)$,*³¹

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

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

Corollary 7. *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, for any $f \in \mathcal{F}$, the function V defined by Equation (1) can be equivalently expressed as*

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

This value function embeds a general divergence representation inside confidence preferences. To see how it captures ambiguity aversion, note that the measure q ultimately used to evaluate an act may be more pessimistic than $\mu \otimes p$ on $\Omega \times S$, which in turn may be more

³¹We adopt the convention that $k(\infty) = \infty$. Thus, for any function k as in the statement of the proposition, if $D_\phi(q \| r) = \infty$ then $k(D_\phi(q \| r)) = \infty$.

pessimistic than μ only on Ω . Hence, compared to μ , there is more “opportunity” for q 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 8. *Fix any ϕ -divergence $D_\phi(\cdot \| \cdot)$ and any function $u : Z \rightarrow \mathbb{R}$. Fix a scalar $\kappa > 0$, and for any $r \in \Delta(\Omega \times S)$ define*

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

Then, there exists a set Ψ of functions $\psi : Z \rightarrow [-\infty, \infty)$ such that, for any $f \in \mathcal{F}$, the function V defined by Equation (1) can be equivalently expressed as

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

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. As in the representation of Corollary 6, in both Corollaries 7 and 8 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.

5.2 Special Cases with Variational Preferences

In the previous subsection we provided special cases of the representation in Equation (8) that modify the argument of the logarithm, but kept the representation of confidence preferences surrounding it untouched. We now show how to drop the logarithm from the value function by transforming the set Ψ , so that ambiguity aversion is captured by a variational representation with entropy costs.

Proposition 3. *Fix any function $u : Z \rightarrow [-\infty, \infty)$. For each $\gamma \geq 0$, define $u_\gamma : Z \rightarrow [-\infty, \infty)$ as follows:*

$$u_\gamma(z) = \begin{cases} \gamma u(z) - \gamma \ln(\gamma) + \gamma & \text{if } \gamma > 0 \\ 0 & \text{if } \gamma = 0. \end{cases}$$

Then, for any $f \in \mathcal{F}$ and any $q \in \Delta(\Omega \times S)$,

$$\max_{\gamma \geq 0} \int_{\Omega \times S} u_\gamma(f(\omega, s)) dq(\omega, s) = \exp \left(\int_{\Omega \times S} u(f(\omega, s)) dq(\omega, s) \right).$$

The following corollary applies Proposition 3 to our representation for adaptive preferences from Theorem 3 by letting $q = \mu \otimes p$ and taking Ψ to be the closed convex hull of the set of functions $\{u_\gamma : \gamma \geq 0\}$.

Corollary 9. *Fix μ , and fix any function $u : Z \rightarrow \mathbb{R}$. Let Ψ be the closed convex hull of $\{u_\gamma : \gamma \geq 0\}$. Then, for any $f \in \mathcal{F}$, the function V defined by Equation (1) can be equivalently expressed as*

$$V(f) = \inf_{p \in M(\mu_\Omega)} \left[\int_{\Omega \times S} u(f(\omega, s)) d(\mu \otimes p)(\omega, s) + R(\mu_\Omega \| p) \right].$$

The value function in Corollary 9 is a special case of the representation for variational preferences developed by Maccheroni, Marinacci, and Rustichini (2006). Note that Proposition 3 can also be used to drop the logarithm in the other special cases considered in the previous subsections, which allows us to nest RDU or divergence preferences in the inner layer within a variational (rather than confidence) preference in the outer layer. Specifically, given any set Ψ , expand the set to $\tilde{\Psi} = \{\psi_\gamma : \psi \in \Psi, \gamma \geq 0\}$. Then, Equation (8) becomes

$$\begin{aligned} V(f) &= \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\sup_{\tilde{\psi} \in \tilde{\Psi}} \int_{\Omega \times S} \tilde{\psi}(f(\omega, s)) d(\mu \otimes p)(\omega, s) \right) + R(\mu_\Omega \| p) \right] \\ &= \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\sup_{\psi \in \Psi} \max_{\gamma \geq 0} \int_{\Omega \times S} \psi_\gamma(f(\omega, s)) d(\mu \otimes p)(\omega, s) \right) + R(\mu_\Omega \| p) \right] \\ &= \inf_{p \in M(\mu_\Omega)} \left[\sup_{\psi \in \Psi} \int_{\Omega \times S} \psi(f(\omega, s)) d(\mu \otimes p)(\omega, s) + R(\mu_\Omega \| p) \right]. \end{aligned}$$

This equality can be used in conjunction with the duality result for rank-dependent utility from Proposition 1 to obtain a version of the RDU preferences from Corollary 6 without the logarithm around the inner term. Similarly, it can be used with Proposition 2 to obtain versions of the divergence preferences from Corollaries 7 and 8 that drop the logarithm.

In this section, we considered the case without signals, $\Sigma = \{\sigma\}$. It is straightforward to generalize the value functions derived here to capture the ex ante value of any action plan $\mathbf{f} \in \mathcal{F}^\Sigma$ when there are signals (see Appendix A). We note that the corresponding ex post value functions would typically not look as clean as in the special cases of Section 4 because they depend on \mathbf{f} indirectly through the minimizing $p \in M(\mu_\Omega)$. It should now be clear, however, that ex post preferences follow immediately from ex ante preferences over plans (and hence ex ante value functions), based on the central insight that they must be dynamically consistent.

A Special Cases: General Treatment

The following theorems generalize Theorems 2 and 3 from Section 5 to allow for a nondegenerate signal structure. Proofs of Theorems 4 and 5 are contained in Section S5 of the Online Appendix.

Theorem 4. *Suppose Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$, and fix μ . For any action plan $\mathbf{f} \in \mathcal{F}^\Sigma$, the function V defined by Equation (2) can be equivalently expressed as*

$$V(\mathbf{f}) = \sup_{\psi \in \Psi^\Sigma} \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\int_{\Omega} \int_{S \times \Sigma} \psi_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma|\omega) dp(\omega) \right) + R(\mu_\Omega \| p) \right]. \quad (10)$$

Theorem 5. *Suppose Ψ is a nonempty set of functions $\psi : Z \rightarrow [-\infty, \infty)$ that is convex, pointwise bounded above, and closed in the topology of pointwise convergence (on the extended reals), and fix μ . For any action plan $\mathbf{f} \in \mathcal{F}^\Sigma$, the function V defined by Equation (2) can be equivalently expressed as*

$$V(\mathbf{f}) = \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\sup_{\psi \in \Psi^\Sigma} \int_{\Omega} \int_{S \times \Sigma} \psi_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma|\omega) dp(\omega) \right) + R(\mu_\Omega \| p) \right]. \quad (11)$$

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

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

With this definition in hand, note that Equation (11) can be written as

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

A.1 Rank-Dependent Utility

The next corollary follows directly from Proposition 1 and Theorem 5.

Corollary 10. *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 $\psi : Z \rightarrow \mathbb{R}$ such that, for any action plan $\mathbf{f} \in \mathcal{F}^\Sigma$, the function V defined by Equation (2) can be equivalently expressed as*

$$V(\mathbf{f}) = \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\int_{\Sigma} \int_Z u(z) d(\varphi \circ F_{\mathbf{f}_\sigma, \mu \otimes p(\cdot|\sigma)})(z) d(\mu \otimes p)(\sigma) \right) + R(\mu_\Omega \| p) \right],$$

where

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

is the cumulative distribution function of \mathbf{f}_σ given $\mu \otimes p$ and σ .

A.2 Divergence Preferences

The following corollaries apply Proposition 2 to our representation for adaptive preferences from Theorem 5 by taking $r = (\mu \otimes p)(\cdot|\sigma)$. The first corollary considers the special case of $k(x) = \theta x$ for some scalar $\theta > 0$. The second corollary considers the special case where we fix a scalar $\kappa > 0$ and take $k(x) = 0$ if $x \leq \kappa$, and $k(x) = +\infty$ if $x > \kappa$.

Corollary 11. *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, for any $\mathbf{f} \in \mathcal{F}^\Sigma$, the function V defined by Equation (2) can be equivalently expressed as*

$$V(\mathbf{f}) = \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\int_\Sigma \inf_{q \in \Delta(\Omega \times S)} \left[\int_{\Omega \times S} u(\mathbf{f}_\sigma(\omega, s)) dq(\omega, s) + \theta D_\phi(q \| (\mu \otimes p)(\cdot|\sigma)) \right] d(\mu \otimes p)(\sigma) \right) + R(\mu_\Omega \| p) \right].$$

Corollary 12. *Fix any ϕ -divergence $D_\phi(\cdot\|\cdot)$ and any function $u : Z \rightarrow \mathbb{R}$. Fix a scalar $\kappa > 0$, and for any $r \in \Delta(\Omega \times S)$ define*

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

Then, there exists a set Ψ of functions $\psi : Z \rightarrow [-\infty, \infty)$ such that, for any $\mathbf{f} \in \mathcal{F}^\Sigma$, the function V defined by Equation (2) can be equivalently expressed as

$$V(\mathbf{f}) = \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\int_\Sigma \inf_{q \in \mathcal{D}((\mu \otimes p)(\cdot|\sigma), \kappa)} \left[\int_{\Omega \times S} u(\mathbf{f}_\sigma(\omega, s)) dq(\omega, s) \right] d(\mu \otimes p)(\sigma) \right) + R(\mu_\Omega \| p) \right].$$

A.3 Pseudo-Variational Preferences

The following corollary applies Proposition 3 to our representation for adaptive preferences from Theorem 5 by letting $q = (\mu \otimes p)(\cdot|\sigma)$ and taking Ψ to be the closed convex hull of the set $\{u_\gamma : \gamma \geq 0\}$. Note that the logarithm is not eliminated in the case of a nondegenerate signal structure, but the \ln and \exp operations cancel each other in the special case of $\Sigma = \{\sigma\}$, as we observed in Section 5.2.

Corollary 13. *Fix μ , and fix any function $u : Z \rightarrow \mathbb{R}$. Then, there exists a set Ψ of functions $\psi : Z \rightarrow \mathbb{R}$ such that, for any $\mathbf{f} \in \mathcal{F}^\Sigma$, the function V defined by Equation (2) can be equivalently expressed as*

$$V(\mathbf{f}) = \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\int_\Sigma \exp \left(\int_{\Omega \times S} u(\mathbf{f}_\sigma(\omega, s)) d(\mu \otimes p)(\omega, s|\sigma) \right) d(\mu \otimes p)(\sigma) \right) + R(\mu_\Omega \| p) \right].$$

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

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

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

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

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

B.2 Proof of Proposition 1

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

Lemma 2. *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).³² In light of these conditions, we can apply Proposition 1 from Sarver (2018) to obtain a set Φ with the claimed properties. ■

³²This was also proved by Chew, Karni, and Safra (1987) in the special case where φ is Lipschitz continuous.

Fix any $f \in \mathcal{F}$ and $\mu \in \Delta(S)$, and let η be the distribution of utility values induced by μ , f , and u . Formally,

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

Take Φ as in Lemma 2 for the function φ , and let $\Psi = \{\phi \circ u : \phi \in \Phi\}$. Then, we have

$$\begin{aligned} \sup_{\psi \in \Psi} \int_S \psi(f(s)) d\mu(s) &= \sup_{\phi \in \Phi} \int_S \phi(u(f(s))) d\mu(s) \\ &= \sup_{\phi \in \Phi} \int_a^b \phi(x) d\eta(x) && \text{(change of variables)} \\ &= \int_a^b x d(\varphi \circ F_\eta)(x) && \text{(Lemma 2)} \\ &= \int_Z u(z) d(\varphi \circ F_{f,\mu})(z). \end{aligned}$$

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.

Note that since $W(\eta) = x$ when $\eta(\{x\}) = 1$, we must have $\phi(x) \leq x$ for all $x \in [a, b]$ and $\phi \in \Phi$. Thus, $\psi(z) = \phi(u(z)) \leq b$ for all $z \in Z$ and $\psi \in \Psi$, so the set Ψ is bounded above. Moreover, taking the closed convex hull (in the topology of pointwise convergence) of Ψ does not alter the values in the equality above, so we can assume that Ψ is closed and convex without loss of generality.

It remains only to prove the last claim in the proposition. Suppose Z is an interval. If u is concave, then since each function in $\psi \in \Psi$ takes the form $\psi = \phi \circ u$ for some $\phi \in \Phi$, and since the functions in Φ are all nondecreasing and concave, each function in Ψ is also concave. Since concavity is preserved under convex combinations and pointwise limits, the conclusion still holds for the the closed convex hull of Ψ .

B.3 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 4. *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 $r \in \Delta(\Omega \times S)$, 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 4,

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

Take Ψ to be the closed convex hull 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.

Therefore, all that remains is to prove Proposition 4. 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 3. *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:³³

³³Borwein 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)).

(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. ■

Lemma 4. 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). ■

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 4 and setting $H = J$ in Lemma 3, 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 5. 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 $X \in L^{\infty}$,

$$H^*(X) = \min_{\alpha \geq 0} [(\alpha J)^*(X) + k^*(\alpha)], \tag{12}$$

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):³⁴ 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 (12), when one sets $(0J) = \delta_{\text{dom}(J)}$.³⁵ 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 4 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 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 4. Note that $D_\phi(q \| p) = \infty$ whenever q is not absolutely continuous with

³⁴Hiriart-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)).

³⁵This definition reflects the convention that $0J(Y) = \infty$ if $J(Y) = \infty$ and $0J(Y) = 0$ otherwise.

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 4, 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 3 and the second equality follows from Lemma 5. Then, using the formula for $(\alpha J)^*$ from Lemma 5, 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. ■

B.4 Proof of Proposition 3

We begin with a useful preliminary observation. Define $H : [-\infty, \infty) \times \mathbb{R}_+ \rightarrow [-\infty, \infty)$ by

$$H(x, \gamma) = \begin{cases} \gamma x - \gamma \ln(\gamma) + \gamma & \text{if } \gamma > 0 \\ 0 & \text{if } \gamma = 0. \end{cases}$$

We will show that for any $x \in [-\infty, \infty)$,

$$\max_{\gamma \geq 0} H(x, \gamma) = \exp(x). \quad (13)$$

Consider first the case of $x = -\infty$. Then, $H(x, \gamma) = -\infty$ for all $\gamma > 0$ and $H(x, \gamma) = 0$ for $\gamma = 0$. Thus,

$$\max_{\gamma \geq 0} H(x, \gamma) = 0 = \exp(x).$$

(Recall that we adopt the convention that $\exp(-\infty) = 0$.) Next, consider the case of $x \in \mathbb{R}$. Taking $\bar{\gamma} = \exp(x) > 0$, it follows directly from the definition of H that $H(x, \bar{\gamma}) = \exp(x)$. We will argue that $\bar{\gamma}$ is the maximizing value of γ . First, observe that $H(x, 0) = 0 < H(x, \bar{\gamma})$. Second, observe that for any $\gamma > 0$,

$$\frac{\partial}{\partial \gamma} H(x, \gamma) = x - \ln(\gamma) \quad \text{and} \quad \frac{\partial^2}{\partial \gamma^2} H(x, \gamma) = -\frac{1}{\gamma}.$$

Thus, H is strictly concave and the unique solution to the first-order condition for optimality is $\gamma = \bar{\gamma}$. This proves Equation (13).

Now, fix any $u : Z \rightarrow [-\infty, \infty)$, $f \in \mathcal{F}$, and $q \in \Delta(\Omega \times S)$, and define u_γ as in the statement of the proposition. Note that

$$\begin{aligned} \int_{\Omega \times S} u_\gamma(f(\omega, s)) dq(\omega, s) &= \begin{cases} \gamma \int_{\Omega \times S} u(f(\omega, s)) dq(\omega, s) - \gamma \ln(\gamma) + \gamma & \text{if } \gamma > 0 \\ 0 & \text{if } \gamma = 0 \end{cases} \\ &= H\left(\int_{\Omega \times S} u(f(\omega, s)) dq(\omega, s), \gamma\right). \end{aligned}$$

Thus, by Equation (13), we have

$$\begin{aligned} \max_{\gamma \geq 0} \int_{\Omega \times S} u_\gamma(f(\omega, s)) dq(\omega, s) &= \max_{\gamma \geq 0} H\left(\int_{\Omega \times S} u(f(\omega, s)) dq(\omega, s), \gamma\right) \\ &= \exp\left(\int_{\Omega \times S} u(f(\omega, s)) dq(\omega, s)\right). \end{aligned}$$

This completes the proof.

References

- 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.
- 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.
- Billingsley, P. (1995): *Probability and Measure*, New York: John Wiley and Sons, Inc.
- 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., 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.
- Cover, T. M. and J. A. Thomas (2006): *Elements of Information Theory*, 2nd ed., Hoboken, New Jersey: John Wiley & Sons.

- 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.
- Ellsberg, D. (1961): “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics*, 75, 643–669.
- 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 M. Schneider (2010): “Ambiguity and Asset Markets,” *Annual Review of Financial Economics*, 2, 315–346.
- 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.
- 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.
- 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. Maiconneuve, and R. T. Rockafellar, Springer, 35–46.
- Ilut, C. L. and M. Schneider (2014): “Ambiguous Business Cycles,” *American Economic Review*, 104, 2368–99.
- 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.
- 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. (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.
- 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.
- Rassoul-Agha, F. and T. Seppäläinen (2015): *A Course on Large Deviations with an Introduction to Gibbs Measures*, vol. 162, Graduate Studies in Mathematics, American Mathematical Society.

- 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.
- Samuelson, L. (2004): “Information-Based Relative Consumption Effects,” *Econometrica*, 72, 93–118.
- Sarver, T. (2018): “Dynamic Mixture-Averse Preferences,” *Econometrica*, 86, 1347–1382.
- 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 considers the possibility of idiosyncratic self-randomization by members of the population of a genotype. Section S2 shows how 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 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 proofs of Theorems 4 and 5 from the main paper are contained in Section S5.

S1 Self-Randomization

One observation about the Ellsberg examples dating back as early as Raiffa (1961) is that individuals may be able to hedge against ambiguity by self-randomizing when choosing an act $f \in F$. In our model with phenotypic flexibility, it is important to recognize that if different parts of the overall population select different acts, they may also select different corresponding fitness functions $\psi \in \Psi$.³⁶

To illustrate the potential benefit of randomizing when choosing an act, suppose for a moment that there is no phenotypic flexibility, $\Psi = \{\psi\}$, and consider a slight variation of Example 1, where now individuals have the option to bet on any single color of their choosing: black, red, or yellow. As argued previously, the value function described in Equation (3) will strictly prefer betting on black to betting on either red or yellow (and will be indifferent between the latter two bets). However, consider an independent 50–50 randomization between betting on red or yellow. This randomization 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.³⁷

³⁶Even for a degenerate choice set $F = \{f\}$, randomization over fitness functions might be beneficial.

³⁷Raiffa (1961) took this logic a step further and used a similar thought experiment as the basis for a normative argument that individuals should satisfy the Savage axioms. While our model can generate strict benefits from randomization, it should be clear that it does not support Raiffa’s normative argument against ambiguity aversion. In particular, in the original version of Example 1, where it was not possible to bet on yellow, the optimal genotype strictly preferred betting on black over betting on red, and the same would have been true when replacing red with yellow.

In the literature on evolutionary biology, the benefits of independent randomization were first observed by Cooper and Kaplan (1982). In the economics and decision theory literature, randomization as a type of hedging mechanism was explored by Machina (1985) in the context of non-expected-utility preferences for risk,³⁸ and by Saito (2015) in the case of ambiguity-averse preferences.

To incorporate self-randomization into our model, suppose each individual in the population has access to an independently distributed random variable θ that is drawn from the uniform distribution on $\Theta = [0, 1]$. The choice of act from a set F and of a fitness function from the set Ψ can depend on this independently drawn θ . Incorporating randomization into the value function in Equation (2) is straightforward, as one can think of θ as an uninformative private signal. Formally, expand the signal space to be $\Sigma \times \Theta$, and let $\mathcal{E}_\Theta = \mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma \otimes \mathcal{B}_\Theta$, where \mathcal{B}_Θ is the Borel σ -algebra on Θ . Uncertainty is now captured by a measure μ in $\Delta(\Omega \times S \times \Sigma \times \Theta)$ with the feature that $\mu(\cdot|\omega, s, \sigma)$ is uniform on Θ for all $(\omega, s, \sigma) \in \Omega \times S \times \Sigma$.

For a finite set of acts F , let $\mathcal{R}(F)$ denote the set of all measurable functions $\mathbf{f} : \Sigma \times \Theta \rightarrow F$, which we refer to as *randomized action plans* from F . Let $\mathcal{R}(\mathcal{F})$ denote the set of all measurable simple functions $\mathbf{f} : \Sigma \times \Theta \rightarrow \mathcal{F}$, that is, $\mathcal{R}(\mathcal{F})$ is the union of $\mathcal{R}(F)$ over all finite sets $F \subset \mathcal{F}$. For $\mathbf{f} \in \mathcal{R}(\mathcal{F})$, we denote the act $f \in \mathcal{F}$ selected by \mathbf{f} for a given (σ, θ) using $\mathbf{f}_{\sigma, \theta}$ rather than $\mathbf{f}(\sigma, \theta)$, and we write $\mathbf{f}_{\sigma, \theta}(\omega, s)$ to denote the outcome generated by this act in state (ω, s) . Correspondingly, let $\mathcal{R}(\Psi)$ denote the set of all measurable functions $\boldsymbol{\psi} : \Sigma \times \Theta \rightarrow \Psi$, which we will refer to as *randomized adaptation plans*. We write $\boldsymbol{\psi}_{\sigma, \theta}$ to denote the fitness function $\psi \in \Psi$ selected by $\boldsymbol{\psi}$ for a given (σ, θ) .

The objective function that maximizes evolutionary fitness in Equation (2) then extends to randomized plans in $\mathcal{R}(\mathcal{F})$ by simply replacing Σ with $\Sigma \times \Theta$:

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

Note that the variable θ is not observable by an analyst. Indeed, we introduced θ merely as a convenient and canonical means of modeling idiosyncratic self-randomization by individuals. Therefore, for a finite set of acts F , if the optimal $\mathbf{f} \in \mathcal{R}(F)$ depends nontrivially on θ , then from the perspective of the analyst who does not observe θ , choices of each individual are stochastic and there will be cross-sectional variation in the choices in the population at the aggregate level. Specifically, the probability of choosing act f from F following signal realization σ is given by

$$\rho(f|F, \sigma) = \mu(\{\theta \in \Theta : \mathbf{f}_{\sigma, \theta} = f\}).$$

Moreover, if the signal σ is also not observable to the analyst, then the (unconditional) choice

³⁸See also Agranov and Ortoleva (2017) and Cerreia-Vioglio, Dillenberger, Ortoleva, and Riella (2018) for more recent research on non-expected-utility preferences and random choice.

probabilities observed by the analyst would be

$$\rho(f|F) = \mu(\{(\sigma, \theta) \in \Sigma \times \Theta : \mathbf{f}_{\sigma, \theta} = f\}).$$

Note that in the value function in Equation (S1), self-randomization in the choice of phenotype effectively convexifies the set of fitness functions Ψ . In the main text of the paper, we explicitly assumed that the set Ψ was convex. We see now that this assumption was without loss of generality, provided the genotype can program the members of its subpopulation to engage in idiosyncratic self-randomization in the choice of phenotype.

As was also mentioned in the main text, there is no strict benefit from randomization when the choice set $F \subset \mathcal{F}$ and the set of phenotypes Ψ are convex, and when each $\psi \in \Psi$ is concave. In some parts of the main text, we assumed directly that this was the case, which allowed us to focus on the value function in Equation (2) and its special cases, thereby simplifying the exposition of ideas. Without those assumptions, there could be a strict benefit to the genotype from self-randomization in the choice of act or phenotype, or both.

That said, even when F and Ψ are not convex, conditioning choice on informative private signals not only improves the expected outcomes for individuals, but it has the additional benefit of reducing the correlation of outcomes across individuals, thereby reducing the aggregate risk faced by the population. We now illustrate the role of signals themselves as a means of randomization within a simplified choice environment.

Application: Signal Response in lieu of Self-Randomization

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.

Let $\mathbf{f}_{\sigma, \theta}(\omega) \in \{0, 1\}$ be the action plan that describes whether an individual bets on state ω after observing (σ, θ) . Since each individual must ultimately place a bet on a single state, the action plans are constrained in that $\mathbf{f}_{\sigma, \theta}(\omega) = 1$ implies that $\mathbf{f}_{\sigma, \theta}(\omega') = 0$ for all $\omega' \neq \omega$.³⁹ Thus,

$$\rho(\omega|\sigma) = \mu(\{\theta \in \Theta : \mathbf{f}_{\sigma, \theta}(\omega) = 1\})$$

is the probability of betting on state ω following signal σ . Assume that $\psi(1) > \psi(0) \geq 0$.

³⁹If 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, so that $\mathbf{f}_{\sigma, \theta}(\omega) > 0$ for all ω . Such diversification is prohibited here, but individuals may nonetheless prefer to randomize over bets on different states in order to replace aggregate uncertainty with idiosyncratic.

The long-run growth rate is now given by

$$\begin{aligned} V(\mathbf{f}) &= \int_{\Omega} \ln \left(\int_{\Sigma \times \Theta} \psi(\mathbf{f}_{\sigma, \theta}(\omega)) d\mu(\sigma, \theta | \omega) \right) d\mu(\omega) \\ &= \int_{\Omega} \ln \left(\int_{\Sigma} \left(\rho(\omega | \sigma) \psi(1) + (1 - \rho(\omega | \sigma)) \psi(0) \right) d\mu(\sigma | \omega) \right) d\mu(\omega). \end{aligned} \quad (\text{S2})$$

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 on state ω' with positive probability following signal σ , or they will not bet on state ω with positive probability 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 either 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 $\rho(\omega' | \sigma) = 0$ or $\rho(\omega | \sigma') = 0$, 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 S1. *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 1, 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 $\rho(R | \sigma)$ denote the probability of choosing R following signal σ , and define $\rho(Y | \sigma)$, $\rho(R | \sigma')$, and $\rho(Y | \sigma')$ similarly. Then,

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

which is maximized by taking

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

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 $\rho(Y|\sigma) = 1$ would lead to excess correlation in outcomes across individuals.⁴⁰

S2 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 phenotype also happens after the realization of a signal, reflecting the rapid nature of this adaptation. We now consider the alternative of before-signal adaptation, where adaptation of the phenotype 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 action.

Formally, signal σ arrives after the choice of phenotype ψ , as illustrated in Figure S1. From the ex ante perspective, the individual thus selects (\mathbf{f}, ψ) , a signal-contingent action plan together with a fixed phenotype, which achieves a fitness of $\psi(\mathbf{f}_\sigma(\omega, s))$ after the realization of (ω, s, σ) . Clearly, the growth rate will be lower than under after-signal adaptation, since the phenotype can no longer be optimized based on the signal realization. There will also be 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 action plan $\mathbf{f} \in \mathcal{F}^\Sigma$ in every period is*

$$V(\mathbf{f}) = \sup_{\psi \in \Psi} \int_{\Omega} \ln \left(\int_{S \times \Sigma} \psi(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma|\omega) \right) d\mu(\omega) \quad (\text{S3})$$

⁴⁰In some cases, conditioning on informative signals may completely eliminate randomization based on uninformative signals. 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 $\rho(R|\sigma) = 0$ and $\rho(Y|\sigma') = 0$, so that 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.

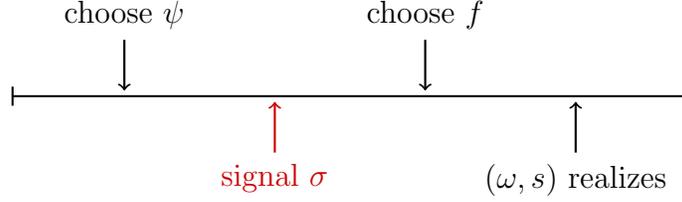


Figure S1: Timeline: before-signal adaptation

The optimal fitness function $\psi_{\mathbf{f}}$ for plan \mathbf{f} satisfies⁴¹

$$\psi_{\mathbf{f}} \in \arg \max_{\psi \in \Psi} \int_{\Omega} \ln \left(\int_{S \times \Sigma} \psi(\mathbf{f}_{\sigma}(\omega, s)) d\mu(s, \sigma | \omega) \right) d\mu(\omega), \quad (\text{S4})$$

and if plan \mathbf{f} is followed for all signals $\sigma' \neq \sigma$ and the ex ante choice of fitness function is $\psi_{\mathbf{f}}$, then long-run growth from choosing g following σ is

$$V(g | \sigma, \mathbf{f}, \psi_{\mathbf{f}}) = \int_{\Omega} \ln \left(\mu(\sigma | \omega) \int_S \psi_{\mathbf{f}}(g(\omega, s)) d\mu(s | \omega, \sigma) + \int_{S \times \Sigma \setminus \{\sigma\}} \psi_{\mathbf{f}}(\mathbf{f}_{\sigma'}(\omega, s)) d\mu(s, \sigma' | \omega) \right) d\mu(\omega). \quad (\text{S5})$$

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

Ex-post adaptive preferences after learning signal σ now have to take into account not only the plan \mathbf{f} , but also the fitness function $\psi = \psi_{\mathbf{f}}$, which is given at the time of choosing an act, as it was chosen optimally in conjunction with \mathbf{f} prior to the realization of σ . When Equation (S4) uniquely pins down $\psi_{\mathbf{f}}$, ex post preferences are fully determined by σ and \mathbf{f} alone, and so can be derived from ex ante preferences.

Since ex ante adaptive preferences are based on the optimal choice of ψ while ex post preferences take $\psi_{\mathbf{f}}$ as given, it is possible to have $g\sigma\mathbf{f} \succ h\sigma\mathbf{f}$ and $g \prec_{\sigma, \mathbf{f}} h$. This is not a violation of our notion of dynamic consistency, which only applies when $g = \mathbf{f}_{\sigma}$, but 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).⁴² The following exam-

⁴¹We directly assume for this result that the optimal fitness function $\psi_{\mathbf{f}}$ exists for each plan \mathbf{f} . 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.

⁴²Preferences in the case of after-signal adaptation that we considered in the main text will satisfy this stronger notion of dynamic consistency: $g\sigma\mathbf{f} \succ h\sigma\mathbf{f} \implies g \succ_{\sigma, \mathbf{f}} h$ (and $g\sigma\mathbf{f} \succeq h\sigma\mathbf{f} \implies g \succeq_{\sigma, \mathbf{f}} h$ whenever $\mu(\sigma) > 0$). This is because for after-signal adaptation, the conditional preference $\succeq_{\sigma, \mathbf{f}}$ does not depend on \mathbf{f}_{σ} , only on $\mathbf{f}_{\sigma'}$ for $\sigma' \neq \sigma$. 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 optimal plan.

ple 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 S2. 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 $f \succ g \sigma f = g \succ h = h \sigma f$. However, $V(h|\psi_1) = 1/18 > 1/25 = V(g|\psi_1)$. For the optimal plan $\mathbf{f} = f$ with optimal fitness function ψ_1 , this means $h \succ_{\sigma, \mathbf{f}} g$.

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 $\varphi : [0, 1] \rightarrow [0, 1]$ 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 adaptive preferences over action plans \mathbf{f} are represented by

$$\hat{V}(\mathbf{f}) = \int_Z u(z) d(\varphi \circ F_{\mathbf{f}, \mu})(z)$$

where

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

is the cumulative distribution function of \mathbf{f} given μ , and $V(\mathbf{f}) = \ln \hat{V}(\mathbf{f})$ is the long-run growth rate defined in Equation (S3).

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.⁴³

⁴³As 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 partitioned learning.

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 $V^{\text{Pr}}(\mathbf{f})$ denote the now familiar long-run growth rate for plan \mathbf{f} under private signals

$$V^{\text{Pr}}(\mathbf{f}) = \int_{\Omega} \ln \left(\int_{\Sigma} \psi(\mathbf{f}_{\sigma}(\omega)) d\mu(\sigma|\omega) \right) d\mu(\omega).$$

Let $V^{\text{Pu}}(\mathbf{f})$ denote the growth rate for \mathbf{f} under public signals,

$$V^{\text{Pu}}(\mathbf{f}) = \int_{\Omega \times \Sigma} \ln \left(\psi(\mathbf{f}_{\sigma}(\omega)) \right) d\mu(\omega, \sigma).$$

For a given menu of acts F , let $\mathbf{f}^{\text{Pr},F} \in \text{argmax}_{\mathbf{f} \in F} V^{\text{Pr}}(\mathbf{f})$ be an optimal plan under private signals, and $\mathbf{f}^{\text{Pu},F} \in \text{argmax}_{\mathbf{f} \in F} V^{\text{Pu}}(\mathbf{f})$ under public signals. When the menu F is fixed, we will write \mathbf{f}^{Pr} and \mathbf{f}^{Pu} without risk of confusion. Then,

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

where the second inequality is strict whenever $\psi(\mathbf{f}_{\sigma}^{\text{Pu}}(\omega))$ is not constant on Σ for some $\omega \in \Omega$.

A more subtle question is how \mathbf{f}^{Pu} and \mathbf{f}^{Pr} differ. We already saw in Section S1 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 $\Omega \times \Sigma$. 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 $\sum_{\omega \in \Omega} \mu(\omega) f(\omega) > c$. Suppose that each individual has unit wealth, and let the plan α specify for each signal $\sigma \in \Sigma$ the proportion $\alpha_\sigma \in [0, 1]$ of the risky asset in the portfolio, so that an individual holds act $\mathbf{f}_\sigma = \alpha_\sigma f + (1 - \alpha_\sigma)c$ upon learning σ . Holding fixed $\mu \in \Delta(\Omega \times \Sigma)$, let α^{Pu} and α^{Pr} denote the optimal portfolio plans for the case where the signals in Σ are public and private, respectively. Finally, assume that the fitness function ψ is increasing, strictly concave, differentiable, and positive on some interval that contains the range of f and c .

Proposition S2. *Let σ_* and σ^* be private signals with the lowest and highest investment in the risky asset, 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

Note that the value function in Equation (S2) can be expressed directly in terms of the probabilities $\rho(\omega|\sigma)$ of betting on state ω following signal σ —the exact dependence of $\mathbf{f}_{\sigma,\theta}$ on the private randomization device θ is not important, so long as the same conditional choice probabilities are generated. We will therefore denote the value function in Equation (S2) as $V(\rho)$. Also, let $\bar{\rho}(\omega)$ denote the probability of betting on state ω when the actual state is ω , given $\rho(\omega|\sigma)$ and $\mu(\sigma|\omega)$. That is,

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

Simple direct computation yields the partial derivative of V with respect to $\rho(\omega|\sigma)$:⁴⁴

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

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

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

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

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

Multiplying these two expressions, we obtain

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

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

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

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

S4.2 Proof of Proposition S2

Since f and c are fixed, we will slightly abuse notation and denote $V^{\text{Pu}}(\mathbf{f})$ for $\mathbf{f}_\sigma = \alpha_\sigma f + (1 - \alpha_\sigma)c$ simply by $V^{\text{Pu}}(\alpha)$, and similarly denote $V^{\text{Pr}}(\mathbf{f})$ by $V^{\text{Pr}}(\alpha)$. Observe first that for any α and any $\sigma \in \Sigma$,

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

⁴⁴The choice of ρ by individuals is clearly subject to the constraint that $\sum_{\omega \in \Omega} \rho(\omega|\sigma) = 1$ for all $\sigma \in \Sigma$. This partial derivative treats $\rho(\omega|\sigma)$ as any real number to consider marginal utility independently of feasibility.

and

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

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

$$\frac{\partial V^{\text{Pu}}(\boldsymbol{\alpha}^{\text{Pu}})}{\partial \alpha_{\sigma_*}} \geq \frac{\partial V^{\text{Pu}}(\boldsymbol{\alpha}^{\text{Pr}})}{\partial \alpha_{\sigma_*}} > \frac{\partial V^{\text{Pr}}(\boldsymbol{\alpha}^{\text{Pr}})}{\partial \alpha_{\sigma_*}},$$

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 $\boldsymbol{\alpha}^{\text{Pu}}$ requires that

$$\frac{\partial V^{\text{Pu}}(\boldsymbol{\alpha}^{\text{Pu}})}{\partial \alpha_{\sigma_*}} \leq 0,$$

and hence

$$\frac{\partial V^{\text{Pr}}(\boldsymbol{\alpha}^{\text{Pr}})}{\partial \alpha_{\sigma_*}} < 0.$$

Since $\boldsymbol{\alpha}^{\text{Pr}}$ 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

$$\frac{\partial V^{\text{Pu}}(\boldsymbol{\alpha}^{\text{Pu}})}{\partial \alpha_{\sigma_*}} \leq \frac{\partial V^{\text{Pu}}(\boldsymbol{\alpha}^{\text{Pr}})}{\partial \alpha_{\sigma_*}} < \frac{\partial V^{\text{Pr}}(\boldsymbol{\alpha}^{\text{Pr}})}{\partial \alpha_{\sigma_*}},$$

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 α^{Pu} requires that

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

and hence

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

Since α^{Pr} 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

S5.1 Proof of Theorem 4

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 $M(q) = \{p \in \Delta(\Omega) : p \ll q \text{ and } R(q \| p) < \infty\}$. In particular, since $R(q \| p) < \infty$ requires that $q \ll p$, the measures p and q are mutually absolutely continuous whenever $p \in M(q)$.

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

$$\int_{\Omega} \ln(X(\omega)) dq(\omega) = \inf_{p \in M(q)} \left[\ln \left(\int_{\Omega} X(\omega) dp(\omega) \right) + R(q \| p) \right]. \quad (\text{S6})$$

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 (S6) is uniquely attained by the measure p_0 with Radon–Nikodym derivative

$$\frac{dp_0}{dq}(\omega) = \frac{1}{X(\omega) \int_{\Omega} \frac{1}{X(\hat{\omega})} dq(\hat{\omega})}. \quad (\text{S7})$$

Proposition S3 restricts to $p \in M(q)$, thereby ensuring that we do not encounter terms of the form $-\infty + \infty$. That is, while the first term inside the infimum in Equation (S6) could take the value $-\infty$, the second term $R(q \| p)$ will necessarily be finite.

Proof. The proof proceeds in three steps. We first prove Equation (S6) 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.

Step 1: Suppose that X that is bounded above and satisfies $X(\omega) \geq \varepsilon > 0$ for all $\omega \in \Omega$.⁴⁵ Then, $\ln(X)$ is a bounded function, and it is therefore integrable. Fix any measures $p, q \in \Delta(\Omega)$ with $p \gg q$ and define a measure q_0 by its Radon–Nikodym derivative

$$\frac{dq_0}{dp}(\omega) = \frac{X(\omega)}{\int_{\Omega} X(\hat{\omega}) dp(\hat{\omega})}. \quad (\text{S8})$$

Since X is strictly positive, q_0 and p are mutually absolutely continuous. In particular, since $p \gg q$, this implies $q_0 \gg q$. Thus, $\frac{dq}{dq_0}$ exists and $\frac{dq}{dp} = \frac{dq}{dq_0} \cdot \frac{dq_0}{dp}$. Note that

$$\begin{aligned} & \int_{\Omega} \ln(X) dq - R(q \| p) \\ &= \int_{\Omega} \ln(X) dq - \int_{\Omega} \ln\left(\frac{dq}{dp}\right) dq \\ &= \int_{\Omega} \ln(X) dq - \int_{\Omega} \ln\left(\frac{dq}{dq_0}\right) dq - \int_{\Omega} \ln\left(\frac{dq_0}{dp}\right) dq \\ &= \int_{\Omega} \ln(X) dq - \int_{\Omega} \ln\left(\frac{dq}{dq_0}\right) dq - \int_{\Omega} \ln(X) dq + \ln\left(\int_{\Omega} X dp\right) \\ &= -R(q \| q_0) + \ln\left(\int_{\Omega} X dp\right). \end{aligned}$$

By Lemma 1.4.1 in Dupuis and Ellis (1997), $R(q \| q_0) \geq 0$, with equality if and only if $q = q_0$. Therefore,

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

with equality if and only if $q = q_0$. It is not difficult to show that Equations (S7) and (S8) are dual in the sense that $q = q_0$ if and only if $p = p_0$. Therefore, given q , if we set $p = p_0$ then the above holds with equality. Moreover, since X is bounded and $1/X \leq 1/\varepsilon$,

$$R(q \| p_0) = \int_{\Omega} \ln\left(\frac{dq}{dp_0}\right) dq = \int_{\Omega} \ln(X) dq + \ln\left(\int_{\Omega} \frac{1}{X} dq\right) < \infty,$$

which implies $p_0 \in M(q)$. Hence the infimum in Equation (S6) is attained at p_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 (S6) holds for

⁴⁵Our proof of this step employs similar techniques to the proof of Proposition 1.4.2 in Dupuis and Ellis (1997), although the details are quite different.

each X_n and for any q . Using this, together with the fact that $X_n \geq X$ for all n , we have

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

Since $\int \ln(X_1) dq < \infty$ 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

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

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

To prove the opposite inequality, note that for any n and any $p \in M(q)$, Equation (S6) applied to the function X_n implies

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

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

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

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

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

Thus, Equation (S6) 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 (S6) holds for X^+ by step 2,

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

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

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

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

We now proceed with the proof of Theorem 4. For a given $\mathbf{f} \in \mathcal{F}^{\Sigma}$ and $\psi \in \Psi^{\Sigma}$, define $X_{\psi} : \Omega \rightarrow [-\infty, \infty)$ by

$$X_{\psi}(\omega) = \int_{S \times \Sigma} \psi_{\sigma}(\mathbf{f}_{\sigma}(\omega, s)) d\mu(s, \sigma | \omega).$$

To verify that X_{ψ} is bounded above, recall that for each σ , $\mathbf{f}_{\sigma} \in \mathcal{F}$ is a simple act, that is, it takes finitely many possible values. Since, in addition, Σ is finite, this implies that there exists $\kappa \in \mathbb{R}$ such that $\psi_{\sigma}(\mathbf{f}_{\sigma}(\omega, s)) \leq \kappa$ for all ω, s, σ . Therefore, $X_{\psi}(\omega) \leq \kappa < \infty$ for all ω . Applying Proposition S3 to this function, we obtain

$$\begin{aligned} & \int_{\Omega} \ln \left[\int_{S \times \Sigma} \psi_{\sigma}(\mathbf{f}_{\sigma}(\omega, s)) d\mu(s, \sigma | \omega) \right] d\mu(\omega) \\ &= \int_{\Omega} \ln(X_{\psi}(\omega)) d\mu_{\Omega}(\omega) \\ &= \inf_{p \in M(\mu_{\Omega})} \left[\ln \left(\int_{\Omega} X_{\psi}(\omega) dp(\omega) \right) + R(\mu_{\Omega} \| p) \right] \\ &= \inf_{p \in M(\mu_{\Omega})} \left[\ln \left(\int_{\Omega} \int_{S \times \Sigma} \psi_{\sigma}(\mathbf{f}_{\sigma}(\omega, s)) d\mu(s, \sigma | \omega) dp(\omega) \right) + R(\mu_{\Omega} \| p) \right]. \end{aligned}$$

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

$$\begin{aligned} V(\mathbf{f}) &= \sup_{\psi \in \Psi^{\Sigma}} \int_{\Omega} \ln \left[\int_{S \times \Sigma} \psi_{\sigma}(\mathbf{f}_{\sigma}(\omega, s)) d\mu(s, \sigma | \omega) \right] d\mu(\omega) \\ &= \sup_{\psi \in \Psi^{\Sigma}} \inf_{p \in M(\mu_{\Omega})} \left[\ln \left(\int_{\Omega} \int_{S \times \Sigma} \psi_{\sigma}(\mathbf{f}_{\sigma}(\omega, s)) d\mu(s, \sigma | \omega) dp(\omega) \right) + R(\mu_{\Omega} \| p) \right]. \end{aligned}$$

This completes the proof.

S5.2 Proof of Theorem 5

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 Tuy (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 $x \in C$, the function $y \mapsto F(x, y)$ is quasiconvex on D .*
3. *For every $x \in C$ and $y, y' \in D$, the function $\lambda \mapsto F(x, \lambda y + (1 - \lambda)y')$ is lower semicontinuous on $[0, 1]$.*
4. *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 F is quasiconvex in y . His result also requires the lower semicontinuity property that we assumed in condition 3. The final assumption needed for his result is condition 4.⁴⁶ ■

⁴⁶Strictly 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.

Note that the theorem of Sion (1958) follows as a corollary to this result: If F is lower semicontinuous in y then condition 3 is implied, and if D is compact then condition 4 is implied (given that F is upper semicontinuous in x).

We now proceed with the proof of Theorem 5. Take Ψ , μ , and \mathbf{f} as in the statement of the theorem. By Theorem 4, if V is defined by Equation (2), then it satisfies

$$V(\mathbf{f}) = \sup_{\boldsymbol{\psi} \in \Psi^\Sigma} \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\int_\Omega \int_{S \times \Sigma} \boldsymbol{\psi}_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma|\omega) dp(\omega) \right) + R(\mu_\Omega \| p) \right].$$

Fix \mathbf{f} , and define $H : \Psi^\Sigma \times M(\mu_\Omega) \rightarrow [-\infty, \infty)$ by

$$H(\boldsymbol{\psi}, p) = \ln \left(\int_\Omega \int_{S \times \Sigma} \boldsymbol{\psi}_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma|\omega) dp(\omega) \right) + R(\mu_\Omega \| p).$$

Define $J : \Psi^\Sigma \times M(\mu_\Omega) \rightarrow \mathbb{R}_+$ by

$$J(\boldsymbol{\psi}, p) = \max \left\{ 0, \int_\Omega \int_{S \times \Sigma} \boldsymbol{\psi}_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma|\omega) dp(\omega) \right\} \exp(R(\mu_\Omega \| p)).$$

Thus, we have $H(\boldsymbol{\psi}, p) = \ln(J(\boldsymbol{\psi}, p))$. Therefore,

$$V(\mathbf{f}) = \sup_{\boldsymbol{\psi} \in \Psi^\Sigma} \inf_{p \in M(\mu_\Omega)} H(\boldsymbol{\psi}, p) = \ln \left[\sup_{\boldsymbol{\psi} \in \Psi^\Sigma} \inf_{p \in M(\mu_\Omega)} J(\boldsymbol{\psi}, p) \right]$$

and

$$\begin{aligned} \ln \left[\inf_{p \in M(\mu_\Omega)} \sup_{\boldsymbol{\psi} \in \Psi^\Sigma} J(\boldsymbol{\psi}, p) \right] &= \inf_{p \in M(\mu_\Omega)} \sup_{\boldsymbol{\psi} \in \Psi^\Sigma} H(\boldsymbol{\psi}, p) \\ &= \inf_{p \in M(\mu_\Omega)} \left[\ln \left(\sup_{\boldsymbol{\psi} \in \Psi^\Sigma} \int_\Omega \int_{S \times \Sigma} \boldsymbol{\psi}_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma|\omega) dp(\omega) \right) + R(\mu_\Omega \| p) \right]. \end{aligned}$$

Given these observations, it suffices to prove that

$$\sup_{\boldsymbol{\psi} \in \Psi^\Sigma} \inf_{p \in M(\mu_\Omega)} J(\boldsymbol{\psi}, p) = \inf_{p \in M(\mu_\Omega)} \sup_{\boldsymbol{\psi} \in \Psi^\Sigma} J(\boldsymbol{\psi}, p). \quad (\text{S9})$$

We will prove Equation (S9) in several steps. First, since Σ is a finite set and each \mathbf{f}_σ is a simple act, the action plan \mathbf{f} can take only finitely many values. That is, for each $\sigma \in \Sigma$, there exists a finite set $Z_\sigma \subset Z$ such that

$$\{\mathbf{f}_\sigma(\omega, s) : \omega \in \Omega, s \in S\} = Z_\sigma.$$

For each $\sigma \in \Sigma$ and $z \in Z_\sigma$, let

$$E_\sigma^z = \{(\omega, s) \in \Omega \times S : \mathbf{f}_\sigma(\omega, s) = z\}.$$

We can then express J as follows:

$$\begin{aligned} J(\boldsymbol{\psi}, p) &= \max \left\{ 0, \int_{\Omega \times S \times \Sigma} \boldsymbol{\psi}_\sigma(\mathbf{f}_\sigma(\omega, s)) d(\mu \otimes p)(\omega, s, \sigma) \right\} \exp(R(\mu_\Omega \| p)) \\ &= \max \left\{ 0, \sum_{\sigma \in \Sigma} \sum_{z \in Z_\sigma} \mu \otimes p(E_\sigma^z \times \{\sigma\}) \boldsymbol{\psi}_\sigma(z) \right\} \exp(R(\mu_\Omega \| p)). \end{aligned} \quad (\text{S10})$$

Note that for any $p \in M(\mu_\Omega)$, since p and μ_Ω are mutually absolutely continuous, it follows that $\mu \otimes p$ and μ are mutually absolutely continuous. This means that $\mu \otimes p(E_\sigma^z \times \{\sigma\}) = 0$ if and only if $\mu(E_\sigma^z \times \{\sigma\}) = 0$. We can therefore assume without loss of generality that $\mu(E_\sigma^z \times \{\sigma\}) > 0$ for all $\sigma \in \Sigma$ and $z \in Z_\sigma$. For otherwise, we can change the outcome that the action plan \mathbf{f} takes on this event without altering the value of $J(\boldsymbol{\psi}, p)$ for any $\boldsymbol{\psi} \in \Psi^\Sigma$ or $p \in M(\mu_\Omega)$. In particular, if $\mu(\sigma) = 0$, then we can drop σ from Σ without altering the function J . If $\mu(\sigma) > 0$ but $\mu(E_\sigma^z \times \{\sigma\}) = 0$, then we can alter the action plan \mathbf{f} on this event by assigning any other outcome $z' \in Z_\sigma$ such that $\mu(E_\sigma^{z'} \times \{\sigma\}) > 0$.

Recall that the functions in Ψ can possibly take the value $-\infty$. This prevents us from applying Theorem S2 since the set of functions mapping into $[-\infty, \infty]$ does not form a topological vector space.⁴⁷ We now argue that we can reduce the set Ψ^Σ to include only functions that take real values for $\sigma \in \Sigma$ and $z \in Z_\sigma$. As we just observed, we can assume without loss of generality that $\mu(E_\sigma^z \times \{\sigma\}) > 0$ for all $\sigma \in \Sigma$ and $z \in Z_\sigma$. Therefore, if $\boldsymbol{\psi}_\sigma(z) = -\infty$ for some $\sigma \in \Sigma$ and $z \in Z_\sigma$, then for any $p \in M(\mu_\Omega)$,

$$\sum_{\sigma \in \Sigma} \sum_{z \in Z_\sigma} \mu \otimes p(E_\sigma^z \times \{\sigma\}) \boldsymbol{\psi}_\sigma(z) = -\infty,$$

and hence $J(\boldsymbol{\psi}, p) = 0$. Let

$$C = \{\boldsymbol{\psi} \in \Psi^\Sigma : \boldsymbol{\psi}_\sigma(z) > -\infty \text{ for all } \sigma \in \Sigma, z \in Z_\sigma\}.$$

If $C = \emptyset$, then $J(\boldsymbol{\psi}, p) = 0$ for all $\boldsymbol{\psi} \in \Psi^\Sigma$ and $p \in M(\mu_\Omega)$, so Equation (S9) holds trivially. If instead $C \neq \emptyset$, then Equation (S9) holds if and only if

$$\sup_{\boldsymbol{\psi} \in C} \inf_{p \in M(\mu_\Omega)} J(\boldsymbol{\psi}, p) = \inf_{p \in M(\mu_\Omega)} \sup_{\boldsymbol{\psi} \in C} J(\boldsymbol{\psi}, p). \quad (\text{S11})$$

To establish Equation (S11), first observe that while the functions (adaptation plans) in

⁴⁷To illustrate, if $a = -\infty$ and $b \in \mathbb{R}$, then $\alpha a + (1 - \alpha)b = -\infty$ for all $\alpha \in (0, 1]$, but $\alpha a + (1 - \alpha)b = b$ for $\alpha = 0$. Therefore, the mixture operation for functions taking values in $[-\infty, \infty]$ is not continuous.

C are defined on all of $\Sigma \times Z$, by Equation (S10) only the values that this function takes on the finite set

$$A \equiv \{(\sigma, z) : \sigma \in \Sigma, z \in Z_\sigma\} \subset \Sigma \times Z$$

are relevant for the determining the value of J . That is, $\psi|_A = \psi'|_A$ (meaning $\psi_\sigma(z) = \psi'_\sigma(z)$ for all $(\sigma, z) \in A$) implies $J(\psi, p) = J(\psi', p)$. Moreover, by the definition of C , any function $\psi \in C$ must take values in the reals for any $(\sigma, z) \in A$. Let

$$\tilde{C} = \left\{ \tilde{\psi} : A \rightarrow \mathbb{R} : \tilde{\psi} = \psi|_A \text{ for some } \psi \in C \right\}.$$

With slight abuse of notation, we can therefore treat J as function from $\tilde{C} \times M(\mu_\Omega)$ into \mathbb{R} . Thus, Equation (S11) is equivalent to

$$\sup_{\tilde{\psi} \in \tilde{C}} \inf_{p \in M(\mu_\Omega)} J(\tilde{\psi}, p) = \inf_{p \in M(\mu_\Omega)} \sup_{\tilde{\psi} \in \tilde{C}} J(\tilde{\psi}, p).$$

The proof is completed by demonstrating that the assumptions of Theorem S2 are satisfied for this set \tilde{C} , for the set $D = M(\mu_\Omega)$, and for the function $F = J$. Since Ψ is convex, it is easy to show that \tilde{C} is convex. It is also straightforward to show that the set $M(\mu_\Omega)$ is convex. Next, since Ψ was assumed to be a closed subset of $[-\infty, \infty]^Z$ (in the product topology on the extended reals), Ψ^Σ is a closed subset of $[-\infty, \infty]^{\Sigma \times Z}$ (in the product topology on the extended reals). From this, it follows that \tilde{C} is a closed subset of \mathbb{R}^A (in the standard Euclidean topology).

To see that condition 1 in Theorem S2 is satisfied, note first that for any $p \in M(\mu_\Omega)$, the mapping

$$\tilde{\psi} \mapsto \sum_{\sigma \in \Sigma} \sum_{z \in Z_\sigma} \mu \otimes p(E_\sigma^z \times \{\sigma\}) \tilde{\psi}_\sigma(z)$$

is linear and continuous on \tilde{C} . Therefore, by Equation (S10), the mapping $\tilde{\psi} \mapsto J(\tilde{\psi}, p)$ is quasiconcave and continuous on \tilde{C} . The following lemma will be used to show that conditions 2 and 3 are satisfied.

Lemma S1. *Suppose $X : \Omega \rightarrow \mathbb{R}$ is measurable and bounded, and fix any $q \in \Delta(\Omega)$. Then, for any $p, p' \in M(q)$, the mapping*

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

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

Proof. Our proof will make use of the Donsker–Varadhan variational formula (see, for ex-

ample, Lemma 1.4.3 in Dupuis and Ellis (1997)), which states that for any $q, r \in \Delta(\Omega)$,

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

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

$$\exp(R(q \| r)) = \sup_{Y \in B_b(\Omega)} \frac{\exp\left(\int_{\Omega} Y dq\right)}{\int_{\Omega} \exp(Y) dr},$$

and hence

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

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

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

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 $\kappa \in \mathbb{R}$ and fix any $\lambda_1, \lambda_2 \in [0, 1]$ such that $h(\lambda_1) \leq \kappa$ and $h(\lambda_2) \leq \kappa$. Suppose without loss of generality that $\lambda_1 \leq \lambda_2$. We need to show that $h(\lambda) \leq \kappa$ for any $\lambda \in (\lambda_1, \lambda_2)$. Note that $h(\lambda_i) \leq \kappa$ is equivalent to

$$\exp\left(\int_{\Omega} Y dq\right) \int_{\Omega} X d(\lambda_i p + (1 - \lambda_i)p') \leq \kappa \int_{\Omega} \exp(Y) d(\lambda_i p + (1 - \lambda_i)p').$$

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

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

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

Fix any $\tilde{\psi} \in \tilde{C}$ and define $X : \Omega \rightarrow \mathbb{R}$ by

$$X(\omega) = \int_{S \times \Sigma} \tilde{\psi}_\sigma(\mathbf{f}_\sigma(\omega, s)) d\mu(s, \sigma | \omega).$$

Applying Lemma S1 to this random variable X and to $q = \mu_\Omega$ implies that for any $p, p' \in M(\mu_\Omega)$, the mapping $\lambda \mapsto J(\tilde{\psi}, \lambda p + (1 - \lambda)p')$ is quasiconvex and lower semicontinuous on $[0, 1]$. This implies that both conditions 2 and 3 in Theorem S2 are satisfied.

Finally, we show that either condition 4 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

$$\inf_{p \in M(\mu_\Omega)} \sup_{\tilde{\psi} \in \tilde{C}} J(\tilde{\psi}, p) > 0.$$

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

$$\tilde{C}_\eta^{\mu_\Omega} \equiv \left\{ \tilde{\psi} \in \tilde{C} : J(\tilde{\psi}, \mu_\Omega) \geq \eta \right\}$$

is closed since J is continuous in $\tilde{\psi}$ and \tilde{C} is closed. Given this, and since \tilde{C} is a subset of the finite-dimensional Euclidean space \mathbb{R}^A , the set $\tilde{C}_\eta^{\mu_\Omega}$ is compact if and only if it is bounded. Now, recall that we assumed that $\sup_{\psi \in \Psi} \psi(z) < \infty$ for all $z \in Z$. Since A is finite, this implies that

$$\kappa \equiv \max_{(\sigma, z) \in A} \sup_{\tilde{\psi} \in \tilde{C}} \tilde{\psi}_\sigma(z) < \infty.$$

In addition,

$$\beta \equiv \min_{(\sigma, z) \in A} \mu(E_\sigma^z \times \{\sigma\}) > 0.$$

For any $\tilde{\psi} \in \tilde{C}$, $\bar{\sigma} \in \Sigma$, and $\bar{z} \in Z_{\bar{\sigma}}$,

$$\begin{aligned} \sum_{\sigma \in \Sigma} \sum_{z \in Z_\sigma} \mu(E_\sigma^z \times \{\sigma\}) \tilde{\psi}_\sigma(z) &\leq \mu(E_{\bar{\sigma}}^{\bar{z}} \times \{\bar{\sigma}\}) \tilde{\psi}_{\bar{\sigma}}(\bar{z}) + (1 - \mu(E_{\bar{\sigma}}^{\bar{z}} \times \{\bar{\sigma}\})) \kappa \\ &\leq \beta \tilde{\psi}_{\bar{\sigma}}(\bar{z}) + (1 - \beta) \kappa. \end{aligned}$$

Thus, since $\eta > 0$, $\tilde{\psi} \in \tilde{C}_\eta^{\mu_\Omega}$ implies that

$$0 < \eta \leq J(\tilde{\psi}, \mu_\Omega) \leq \beta \tilde{\psi}_{\bar{\sigma}}(\bar{z}) + (1 - \beta) \kappa \implies \tilde{\psi}_{\bar{\sigma}}(\bar{z}) \geq -\frac{(1 - \beta) \kappa}{\beta}.$$

Therefore, the functions in $\tilde{C}_\eta^{\mu_\Omega}$ are bounded above by κ and bounded below by $-(1 - \beta) \kappa / \beta$.

This implies that $\tilde{C}_\eta^{\mu_\Omega}$ is a bounded, and hence compact, subset of \mathbb{R}^A . Thus, all of the assumptions of Theorem S2 are satisfied and we can therefore conclude that Equation (S11) holds. The second case is when

$$\inf_{p \in M(\mu_\Omega)} \sup_{\tilde{\psi} \in \tilde{C}} J(\tilde{\psi}, p) = 0.$$

In this case, it is immediate that the left side of Equation (S11) must also equal zero. Thus, in either case Equation (S11) holds. This completes the proof.

References for Supplementary Appendix

- Agranov, M. and P. Ortoleva (2017): “Stochastic Choice and Preferences for Randomization,” *Journal of Political Economy*, 125, 40–68.
- Cerreia-Vioglio, S., D. Dillenberger, P. Ortoleva, and G. Riella (2018): “Deliberately Stochastic,” working paper.
- 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.
- 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. (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.
- von Neumann, J. (1928): “Zur Theorie der Gesellschaftsspiele,” *Mathematische Annalen*, 100, 295–320.

- Raiffa, H. (1961): “Risk, Ambiguity, and the Savage Axioms: Comment,” *Quarterly Journal of Economics*, 75, 690–694.
- Saito, K. (2015): “Preferences for Flexibility and Randomization under Uncertainty,” *American Economic Review*, 105, 1246–1271.
- 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.