An Evolutionary Perspective on Updating Risk and Ambiguity Preferences*

Philipp Sadowski† Todd Sarver‡

Current Draft:
December 17, 2024

Abstract

Using an evolutionary approach, we address the prominent tensions in the literature that updating of ambiguity and non-expected-utility preferences cannot, in general, be both dynamically consistent and consequentialist. Perhaps not surprisingly, evolutionary optimality requires dynamic consistency. The more subtle insight is the evolutionary optimality of systematic violations of consequentialism. We base our investigation on the model of *adaptive preferences* (Sadowski and Sarver (2024)), which generalizes the model of Robson (1996) and nests variants of many well known models in the literature.

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

JEL Classification: D81, D83, D84

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

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

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

1 Introduction

It is well known in the literature on ambiguity aversion 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?[1] The tension between consequentialism and dynamic consistency in models of ambiguity aversion, which also arises in non-expected-utility models of choice under objective risk, and the subsequent disagreement over which property to give priority are impediments to applying these models in dynamic contexts such as macroeconomics and finance where information plays a central role.

This paper provides an evolutionary perspective on the issue, based on the notion that natural selection not only can influence physical traits, but can also shape choice behavior. In a seminal paper, Robson (1996) showed that evolutionary optimality generates a preference for idiosyncratic uncertainty over common uncertainty.[2] Sadowski and Sarver (2024) generalized this insight and showed that the class they called *adaptive preferences* is evolutionarily optimal when individuals simultaneously make multiple decisions, some of which are observable and others which are hidden from the modeler. They further argued that in many contexts ambiguity coincides with common uncertainty and risk with idiosyncratic uncertainty. Based on this notion, they established that adaptive preferences include rank-dependent expected utility preferences 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. Here, we take choice under ambiguity and risk according to adaptive preferences as a starting point, so that our results on updating apply to this rich class of preferences. We find that preferences will be dynamically consistent and demonstrate why violations of consequentialism should be neither surprising nor concerning.

A central feature of ambiguity aversion is that randomization can provide a hedge against ambiguity. Consequently, self-randomization may be optimal when exogenous randomization is unavailable. We observe that this is true also in our model and, going further, that adaptive preferences many lead to self-randomization even when exogenous randomization is available.

2 Evolutionary Setting

The basic idea behind the evolutionary approach is that a large population of individuals is initially made up of subpopulations with different genotypes, where a genotype specifies the physical traits as well as the programmed behavior (choices) of an organism. These

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

[2] Oprea and Robalino (2024) recently documented the prevalence of such preferences experimentally.

choices lead to a possibly uncertain outcome, and this outcome together with the physical traits of the organism determine its evolutionary fitness, that is, its number of offspring. The offspring inherit the parent's genotype and will face a choice of their own, and so on. In this way, the number of individuals who share a particular genotype may shrink or grow over time, relative to the whole population. A genotype is evolutionarily optimal among those initially present if the relative size of its subpopulation does not vanish over time.

2.1 Uncertainty and Information

The formal setup is as in Sadowski and Sarver (2024), but with the addition of signals. In particular, there is a state space Ω for which the realization of $\omega \in \Omega$ is common to all individuals in the population. As mentioned before, we can also interpret Ω as capturing ambiguity. In addition, given ω, idiosyncratic uncertainty is captured via a state space S, where each individual in the population receives an independent draw of the state $s \in S$. We also refer to uncertainty from S as risk. 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).[3] 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 $\mathcal{B}_\Omega$, $\mathcal{B}_S$, and $\mathcal{B}_\Sigma$, respectively, and we endow the product of these spaces with the product σ-algebra $\mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma$.

Given any measurable space $(Y, \mathcal{Y})$, let $\triangle(Y)$ denote the set of countably additive probability measures on Y, and let $\triangle_s(Y)$ denote the set of all simple probability measures on Y (i.e., measures with finite support). Each period, there is a draw of the state ω that is common to all individuals in the population according to a measure $\mu \in \triangle(\Omega)$. Conditional on ω, both the s dimension of the state and the signal σ are drawn independently for each individual according to the probability measure $\nu_\omega \in \triangle(S \times \Sigma)$.[4] For example, if ω represents the unknown efficacy of a vaccine (which is common to all individuals and potentially ambiguous), then s might capture the vaccine response for a particular individual (so the distribution ν_ω of s would naturally depend on ω). Finally, combining the marginal μ on Ω with the conditional probability distribution ν_ω on $S \times \Sigma$ gives a joint distribution over $\Omega \times S \times \Sigma$, and the information content of an individual's signal σ is captured by conditioning this distribution on σ. In our example, σ could represent lab work on a particular patient. This information structure is quite general and includes, among other things, the partitional structures often used in the literature on ambiguity and updating.

[3] Since S describes idiosyncratic risk, it is natural to consider private signals. In Section S4 of the Online Supplement, we discuss how behavior differs between common and private signals when both are informative only about the common component Ω.

[4] Assume also that the mapping $\omega \mapsto \nu_\omega(E)$ is $\mathcal{B}_\Omega$-measurable for every $E \in \mathcal{B}_S \otimes \mathcal{B}_\Sigma$.

2.2 Consumption and Fitness

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

Individuals face the task of choosing acts in each period contingent on the observed signal $\sigma \in \Sigma$, but before learning the realization of the state (ω, s). Each genotype determines preferences that are used for this choice, contingent on σ. In addition to the observable choice of act f, we assume that individuals might also take hidden actions, that is, actions that are unobservable to the modeler. Incomplete data of this sort is pervasive in economic analysis, as data sets often contain only a snapshot of one dimension of the full spectrum of decisions being made by individuals. We model hidden actions in a simple and tractable reduced form by allowing individuals to select a fitness function ψ from some feasible set Ψ in each period, which we also refer to as adaptation. As we discuss in Section S6 of the Supplement, our use of multiple fitness functions can also be interpreted in terms of phenotypic flexibility in the context of evolutionary biology.

Putting the pieces of our model together, the timing within each period of the infinitely repeated reproductive process is as follows:

$$\text{observe } \sigma \;\longrightarrow\; \text{choose } \psi \text{ and } f \;\longrightarrow\; \text{observe } (\omega, s) \;\longrightarrow\; \text{fitness } \psi(f(\omega, s))$$

3 Evolutionarily Optimal Choice

To analyze dynamic choice in general—and dynamic consistency in particular—it is necessary to compare ex post behavior after the arrival of information with the ex ante plan that would be formed if the individual committed to signal-contingent choices prior to the realization of the signal. We begin our analysis in Section 3.1 by deriving the evolutionarily optimal value function over (possibly random) ex ante *plans* of action and fitness function selection. In Section 3.2, we discuss the role that self-randomization plays in this formula. In Section 3.3, we consider evolutionarily optimal ex post behavior and establish that choice is dynamically consistent.

3.1 Ex Ante Plans

Definition 1. A *random plan* is a function $\pi \in \mathcal{R}(\mathcal{F}, \Psi) \equiv (\triangle_s(\mathcal{F} \times \Psi))^{\Sigma}$ from the space of signals to the set of simple probability measures over the space of acts and feasible fitness functions. The probability π assigns to (f, ψ) following signal σ is denoted by $\pi_\sigma(f, \psi)$.

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

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

Since only the random choice of act is observed, while the choice of fitness function corresponds to some unobservable action, it will be useful to decompose π_σ into its (observable) marginal distribution over acts and (unobservable) conditional distribution over fitness functions given the act.

Definition 2. A *(random) action plan* is a function $\rho \in \mathcal{R}(\mathcal{F}) \equiv (\triangle_s(\mathcal{F}))^\Sigma$ from the space of signals to the set of simple probability measures over acts, where $\rho_\sigma(f)$ is the probability assigned to f following signal σ. A *(random) adaptation plan* is a function $\tau \in \mathcal{R}(\Psi|\mathcal{F}) \equiv (\triangle_s(\Psi))^{\Sigma\times\mathcal{F}}$ from the space of signals and acts to the set of simple probability measures over the feasible fitness functions, where $\tau_\sigma(\psi|f)$ denotes the probability assigned to fitness function ψ following signal σ and the observable choice of act f.

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

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

Given an action plan ρ and adaptation plan τ, the corresponding joint plan over both actions and adaptation is $\pi = \tau \otimes \rho \equiv (\tau_\sigma \otimes \rho_\sigma)_{\sigma\in\Sigma}$.

A *decision problem* $A = (A_\sigma)_{\sigma\in\Sigma}$ specifies a nonempty and finite set of available acts A_σ following each signal σ. The resulting set of feasible action plans is

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

Finally, 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 $x \leq 0$.

Our first result concerns the evolutionarily optimal action plan. It is well known that evolutionary optimality requires maximizing the long-run asymptotic growth rate,[5] which differs from expected growth in the case of common uncertainty.

[5]This insight is revisited in Section 2.3 of Sadowski and Sarver (2024).

Theorem 1 (Adaptive Preferences with Signals)**.** *Fix* Ψ *and fix measures* $\mu \in \triangle(\Omega)$ *and* $\nu_\omega \in \triangle(S \times \Sigma)$ *for each* $\omega \in \Omega$*. Then, for every infinitely repeated decision problem* A*, the genotype that chooses an action plan in* $\operatorname{argmax}_{\rho \in \mathcal{R}(A)} V(\rho)$ *achieves a weakly higher long-run growth rate than all others, where*

$$V(\rho) = \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \int_\Omega \ln\left(\int_{S\times\Sigma} \mathbb{E}_{\tau_\sigma \otimes \rho_\sigma} \left[\psi(f(\omega, s))\right] d\nu_\omega(s,\sigma) \right) d\mu(\omega). \tag{1}$$

The concavity of the logarithm implies that V is more adversely affected by common uncertainty about ω (ambiguity) than by idiosyncratic uncertainty about s (risk). The proof of the theorem closely parallels that for the case without signals in Sadowski and Sarver (2024), and is omitted for brevity. At its heart is the same logic that is behind the seminal result of Robson (1996), who considered the special case with a single fitness function ψ, and without signals or random choice. In an application to status and relative consumption effects, Nöldeke and Samuelson (2005) also considered an environment with signals, but without adaptation ($\Psi = \{\psi\}$) and without random choice. The survey by Robson and Samuelson (2011) summarizes these and other recent developments in the literature on evolution of preferences.

3.2 The Role of Random Choice

A central behavioral feature of models of ambiguity aversion is that exogenous randomization that resolves after the resolution of ambiguity can serve as a hedge against said ambiguity. The potential of random choice to similarly serve as a hedging device against ambiguity was first discussed by Raiffa (1961), and has been explored axiomatically more recently by Saito (2015) and Ke and Zhang (2020).[6] Additionally, in the literature on evolutionary biology, several studies have highlighted the potential benefits of randomization in behavior or in the assignment of physical characteristics to organisms (e.g., Cooper and Kaplan (1982), Bergstrom (2014)).

One implication of our model that has not received attention in this prior literature is the possibility of strict preferences to self-randomize rather than use some exogenous source of randomization.[7] The key conceptual point is that, contrary to the intuition that late randomization more convincingly hedges against ambiguity, in our model early randomization has an advantage, as it allows individuals to coordinate between observed actions (choice of act) and unobserved actions (adaptation). If, as is usually assumed, exogenous random-

[6]For a related discussion of random choice induced by quasiconcave non-expected-utility preferences for risk, see Machina (1985) and Cerreia-Vioglio et al. (2018).

[7]The aforementioned axiomatic models invoke the opposite preference. In the experimental literature, there is some evidence that subjects have limited or even negative willingness to pay for mixtures of acts (see Dominiak and Schnedler (2011) and Agranov and Ortoleva (2017)), indicating a potential preference for self-randomization.

ization takes place after the resolution of ambiguity and hence after the choice of act, then the freedom to choose when to perform self-randomization makes it the preferred hedging mechanism. We provide an illustrative example in Section S1 of the Supplement.

3.3 Updating and Dynamic Consistency

Since it is impossible to directly observe preferences over random action plans,[8] we focus on the choice of action plan given a decision problem $A = (A_\sigma)_{\sigma \in \Sigma}$. As a natural extension of the standard choice correspondence used for deterministic choice, let $\mathcal{C}(A)$ denote the set of all random action plans ρ that the individual is willing to choose ex ante. Similarly, let $\mathcal{C}(A_\sigma|\sigma, \rho)$ denote the set of all random actions $\widehat{\rho}_\sigma \in \triangle_s(A_\sigma)$ that the individual is willing to choose after observing the signal σ and given the ex ante plan ρ.

Definition 3. Random choices satisfy *consequentialism* if ex post choices do not depend on ex ante plans: $\mathcal{C}(A_\sigma|\sigma, \rho) = \mathcal{C}(A_\sigma|\sigma, \widehat{\rho})$ for all A and all $\rho, \widehat{\rho} \in \mathcal{R}(A)$ and $\sigma \in \Sigma$.[9]

Definition 4. Random choices satisfy *dynamic consistency* if ex ante plans are carried out ex post: For any A and $\rho, \widehat{\rho} \in \mathcal{R}(A)$ such that $\rho_{\sigma'} = \widehat{\rho}_{\sigma'}$ for all $\sigma' \neq \sigma$,

1. $\rho \in \mathcal{C}(A)$ and $\widehat{\rho} \notin \mathcal{C}(A)$ together imply $\rho_\sigma \in \mathcal{C}(A_\sigma|\sigma, \rho)$ and $\widehat{\rho}_\sigma \notin \mathcal{C}(A_\sigma|\sigma, \rho)$.
2. $\rho \in \mathcal{C}(A)$ implies $\rho_\sigma \in \mathcal{C}(A_\sigma|\sigma, \rho)$ whenever σ occurs with positive probability.[10]

These conditions extend the standard definitions used in the special case of partitional information structures and deterministic choice (e.g., Machina and Schmeidler (1992), Epstein and Le Breton (1993), or Hanany and Klibanoff (2007)) to our more general framework. Specifically, in the case of deterministic choice, the random action plan ρ reduces to a deterministic plan $(f_\sigma)_{\sigma \in \Sigma}$. Partitional learning corresponds to the special case where Σ is a partition of $\Omega \times S$, so each signal σ is a subset of $\Omega \times S$ and, conditional on the signal σ, states outside of the event σ are assigned probability zero. In this case, a deterministic action plan can be reduced to an act by defining $f(\omega, s) = f_\sigma(\omega, s)$ for $(\omega, s) \in \sigma \in \Sigma$.

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

[9] One could also weaken the definition of consequentialism so that it only applies to plans that are *optimal* for some off-path option sets without changing any of our results: For any $\sigma \in S$ and any decision problems A, A' such that $A_\sigma = A'_\sigma$, we have $\mathcal{C}(A_\sigma|\sigma, \rho) = \mathcal{C}(A'_\sigma|\sigma, \widehat{\rho})$ for all $\rho \in \mathcal{C}(A)$, $\widehat{\rho} \in \mathcal{C}(A')$.

[10] Formally, $\int_\Omega \nu_\omega(\sigma)\, d\mu(\omega) > 0$.

Finally, translating choice to preferences the standard way, if $f, g \in A$, then $f \in \mathcal{C}(A)$ and $g \notin \mathcal{C}(A)$ means that $f \succ g$, and similarly for ex post preferences $\succsim_{\sigma,f}$. Our version of dynamic consistency then implies the standard definition: If $f(s) = g(s)$ for all $s \notin \sigma$,

$$\begin{aligned} f \succ g &\implies f \succ_{\sigma,f} g \quad \text{and} \\ f \succsim g &\implies f \succsim_{\sigma,f} g \quad \text{whenever } \sigma \text{ occurs with positive probability.} \end{aligned}$$

As noted above, ambiguity-aversion and violations of expected-utility in general imply that consequentialism and dynamic consistency cannot be simultaneously satisfied. Fortunately, the evolutionary approach gives clear guidance about which property to favor: The evolutionarily optimal ex ante plans are precisely those that maximize the long-run growth rate of the genotype. The evolutionarily optimal ex post plans are those that achieve the same objective. Thus, dynamic consistency is necessarily satisfied, as the following results demonstrate.

Theorem 2 (Ex Post Long-Run Growth)**.** *Fix Ψ and measures $\mu \in \triangle(\Omega)$ and $\nu_\omega \in \triangle(S \times \Sigma)$ for each $\omega \in \Omega$. Suppose the genotype forms an action plan $\rho \in \mathcal{R}(\mathcal{F})$ ex ante, which it follows after every signal $\sigma' \neq \sigma$, but it deviates from this plan after signal σ by instead implementing the ex post random action $\widehat{\rho}_\sigma \in \triangle_s(\mathcal{F})$. Then, its long-run growth rate is*

$$\begin{aligned} V(\widehat{\rho}_\sigma|\sigma, \rho) = \sup_{\tau \in \mathcal{R}(\Psi|\mathcal{F})} \int_\Omega \ln\Bigg(\nu_\omega(\sigma) \int_S \mathbb{E}_{\tau_\sigma \otimes \widehat{\rho}_\sigma} \big[\psi(f(\omega, s))\big] \, d\nu_\omega(s|\sigma) \\ + \int_{S \times \Sigma \setminus \{\sigma\}} \mathbb{E}_{\tau_{\sigma'} \otimes \rho_{\sigma'}} \big[\psi(f(\omega, s))\big] \, d\nu_\omega(s, \sigma') \Bigg) d\mu(\omega). \end{aligned} \tag{2}$$

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

Corollary 1 (Dynamic Consistency)**.** *Given a decision problem $A = (A_\sigma)_{\sigma \in \Sigma}$, the long-run growth rate is optimized if individuals maximize Equation (1) ex ante and Equation (2) ex post, so $\mathcal{C}(A) = \operatorname{argmax}_{\rho \in \mathcal{R}(A)} V(\rho)$ and $\mathcal{C}(A_\sigma|\sigma, \rho) = \operatorname{argmax}_{\widehat{\rho}_\sigma \in \triangle_s(A_\sigma)} V(\widehat{\rho}_\sigma|\sigma, \rho)$. Thus, evolutionarily optimal random choice is dynamically consistent.*

Given the tension between dynamic consistency and consequentialism, one implication of Corollary 1 is that choice may violate consequentialism. In Section 4.1, we discuss such a violation in the context of the three-color Ellsberg urn. The evolutionary approach provides a natural interpretation for why consequentialism may be violated. For expositional clarity, consider deterministic plans. Consequentialism states that preferences between acts f and g following a signal σ do not depend on what act would have been obtained following other signals σ'. This property could therefore be interpreted as preferences not depending on "what might have been." In our model, the genotype consists of a large subpopulation of individuals. Even if one individual receives the signal σ, other members of this subpopulation are simultaneously receiving different signals. From the individual perspective, choice after updating can be thought of as the best response to other individuals who are all playing the Pareto optimal equilibrium of the game that has long-run population growth as the payoff. In other words, individuals may violate consequentialism because they care about the correlation between their own fitness and the fitness others with the same genotype are experiencing.

Hanany and Klibanoff (2007, 2009) similarly studied dynamically consistent (and hence non-consequentialist) conditional preferences. 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. Their conditional preferences may violate consequentialism since their updating rule for beliefs is typically non-Bayesian and depends on the original choice set and the ex ante plan, sacrificing the separation of tastes and beliefs. In contrast, updated beliefs in Equation (2) are derived using standard Bayesian updating and hence are independent of the decision problem. The violation of consequentialism in this expression comes instead from an externality—in the sense that each individual is programmed to care about correlation with other individuals—which requires the ex ante plan to be a part of the ex post value function. In the context of our evolutionary model, this strikes us as the most natural formulation of the conditional growth rate, as it emphasizes the underlying reason for the dependence of the optimal ex post choices on the ex ante plan.

4 Updating Special Cases

4.1 Ellsberg with Signals

We now consider signals and updating in the special case of a single fitness function, $\Psi = \{\psi\}$. For expositional ease, we will focus on an example where deterministic choice is optimal, and we therefore restrict attention to deterministic action plans $(f_\sigma)_{\sigma\in\Sigma}$. In this case, the long-run growth rate in Equation (1) becomes

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

The long-run growth rate in Equation (2) from deviating from the plan $(f_\sigma)_{\sigma\in\Sigma}$ by instead selecting g following the signal $\bar\sigma$ becomes

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

In particular, choice is dynamically consistent (this is a special case of Corollary 1).

To illustrate the tension between consequentialism and dynamic consistency, consider the following example where acts depend only on s.

Example 1 (Ellsberg—with 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 bets on the color of the ball drawn. The act B pays \$1 if the ball drawn is black and \$0 otherwise, BY pays \$1 if the ball is either black or yellow, and so on. The typical preference pattern documented by Ellsberg (1961) is $B\succ R$ and $BY\prec RY$, in violation of Savage's sure-thing principle.*

To understand such preferences within our evolutionary model, 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 take $\Omega=\{\omega_1,\omega_2,\omega_3\}$, where $\omega_1=(b,r,r)$, $\omega_2=(b,r,y)$, and $\omega_3=(b,y,y)$. Even if individuals form subjective probability assessments on the possible urn compositions, this correlated uncertainty is treated differently than uncorrelated uncertainty. For ease of illustration, suppose that μ assigns equal weight to each urn composition and that $\psi(0)=0$ and $\psi(1)=1$.

Replicating calculations in Sadowski and Sarver (2024), we first verify 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

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

Further simple calculations show that randomizations $\alpha\delta_B+(1-\alpha)\delta_R$ for $0<\alpha<1$ yield long-run growth rates strictly between those of B and R, and likewise for the second choice scenario. Thus, deterministic choice is indeed optimal in this situation and we have $\mathcal{C}(\{B,R\})=\{B\}$ and $\mathcal{C}(\{BY,RY\})=\{RY\}$.[11]

[11] For other decision problems random choice may be optimal. For instance, if individuals are able to bet on red or yellow, then an independent 50–50 randomization will yield an independent 1/3 probability of winning for each urn composition. An example where random choice even outperforms exogenous randomization is in Section S1 of the Supplement.

Now suppose each individual receives a private signal that tells them whether the ball drawn for them is yellow (y) or not yellow ($\neg y$).[12] *Under partitional learning, preferences over signal-contingent action plans 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:*[13]

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

Therefore, dynamic consistency requires that

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

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

Note that the tension illustrated in this example depends neither on a particular choice of updating rule nor on the specific model of ambiguity aversion: Ellsberg behavior of the form $B\neg yB \succ R\neg yB$ and $B\neg yY \prec R\neg yY$ together with this specific information structure cannot satisfy both dynamic consistency and consequentialism.[14] Our model and results imply that individuals with these ex ante preferences will exhibit the conditional ex post preferences listed above. Thus, individuals will be dynamically consistent but will violate consequentialism.

As noted previously, consequentialism is violated because evolutionarily optimal preferences include an "externality" that incorporates the growth rate of other individuals in the population who are simultaneously receiving different signals. For example, given the ex ante action plan $R\neg yY$, following signal $\neg y$ there is a complementarity between r and y: The probability of seeing signal $\neg y$ and then state r is higher for urn compositions ω where the probability of signal y (and hence state y) is lower. 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,

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

[13] We are making two implicit assumptions in this argument (both are implied by our model). The first is that individuals are indifferent between action plans that differ only on measure zero events, such as $R\neg yB$ and R. The second is that preferences between any two acts (trivial action plans) f and g do not change after the introduction of signals when there are no hidden actions. The discussion of Proposition 1 in Section 4.2 makes clear that this may not be true in the general case with hidden actions.

[14] 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).

and as a result

$$V(R|\neg y, R\neg yY) = \sum_{\omega\in\Omega} \mu(\omega) \ln \Big(\underbrace{\nu_\omega(\neg y)}_{\substack{\text{fraction}\\ \text{getting}\\ \text{signal } \neg y}} \ \underbrace{\nu_\omega(r|\neg y)}_{\substack{\text{average fitness}\\ \text{from } R \text{ after}\\ \text{signal } \neg y}} \ + \underbrace{\nu_\omega(y)}_{\substack{\text{fraction}\\ \text{getting}\\ \text{signal } y}} \ \underbrace{1}_{\substack{\text{fitness}\\ \text{from } Y\\ \text{after } y}} \Big)$$

$$> \sum_{\omega\in\Omega} \mu(\omega) \ln \Big(\nu_\omega(\neg y)\ \nu_\omega(b|\neg y) \ + \ \nu_\omega(y)\ 1 \Big) = V(B|\neg y, R\neg yY).$$

In contrast, choosing B does not hedge against this aggregate growth-rate risk, because the probability of state b is independent of the urn composition. When the ex ante plan is instead $B\neg yB$, the hedging motive for the choice of R following $\neg y$ disappears, as now the growth rate following signal y is zero. In this case, we have the opposite conditional preference: $V(B|\neg y, B\neg yB) > V(R|\neg y, B\neg yB)$.

Our evolutionary approach to ambiguity aversion applies directly to situations where ambiguity can be identified with common uncertainty, that is, where the same uncertainty is faced by a large subpopulation. In contrast, the number of subjects in lab experiments (for instance, in Ellsberg experiments) is small. Whether the adaptive model nonetheless provides a good description of behavior in experiments, perhaps due to heuristics developed in the case of common ambiguity, is an empirical question. Section S5 of the Supplement discusses evidence from experiments, as well as lessons from our analysis for their design.

4.2 Updating Rank-Dependent Expected Utility Preferences

The tension between dynamic consistency and consequentialism is not exclusive to environments with ambiguity, but can also arise when updating models of non-expected utility for risk. Machina (1989) prominently argued that those models should be updated in a way that is dynamically consistent, even at the cost of consequentialism. The adaptive model accommodates violations of expected utility, and since updating in the adaptive model is dynamically consistent, our results support this general position for the models that it nests as special cases. To illustrate, in this section we consider another canonical special case of our model: rank-dependent utility. Perhaps surprisingly, evolution can generate a version of dynamic RDU that is both dynamically consistent and consequentialist.[15]

To focus on risk preferences, we restrict attention to the special case of our model without common uncertainty ($\Omega = \{\omega\}$), but we now permit non-degenerate after-signal adaptation. In this case, there is no strict benefit to self-randomization, so it is without loss of generality

[15] This result depends crucially on the timing in the model. In Section S2 of the Supplement we discuss how the non-consequentialist version of dynamic RDU suggested by Machina (1989) is obtained when adaptation (the hidden action) instead takes place before the signal arrives.

to restrict attention to deterministic action plans $(f_\sigma)_{\sigma\in\Sigma}$ and adaptation plans $(\psi_\sigma)_{\sigma\in\Sigma}$.[16] Therefore, Equation (1) becomes

$$V\big((f_\sigma)_{\sigma\in\Sigma}\big) = \sup_{(\psi_\sigma)_{\sigma\in\Sigma}\in\Psi^\Sigma} \ln\bigg(\int_{S\times\Sigma} \psi_\sigma(f_\sigma(s))\, d\nu(s,\sigma)\bigg) \\ = \ln\bigg(\int_\Sigma \sup_{\psi\in\Psi}\bigg[\int_S \psi(f_\sigma(s))\, d\nu(s|\sigma)\bigg] d\nu(\sigma)\bigg). \tag{3}$$

Note that in this case the logarithm can also be dropped by taking a monotone transformation, but we will retain it for consistency. Although the connection is nontrivial, the following result shows that rank-dependent utility with a pessimistic probability distortion function can be expressed as a special case of our model.

Proposition 1 (Updating Rank-Dependent Utility)**.** *Suppose $\Omega = \{\omega\}$ and $Z \subset \mathbb{R}$. Fix ν, and fix any bounded nondecreasing function $u : Z \to \mathbb{R}$ and any function $\varphi : [0,1] \to [0,1]$ that is nondecreasing, concave, and onto. Then, there exists a set Ψ of functions $\psi : Z \to \mathbb{R}$ such that the ex ante value function V defined by Equation (3) can be equivalently expressed as*

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

and ex post adaptive preferences following any signal $\bar\sigma$ that occurs with positive probability are represented by

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

where

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

denotes the cumulative distribution function of g given ν and $\bar\sigma$.

Proof. Proposition 1 of Sadowski and Sarver (2024) implies that for any probability distortion function φ as in the statement of the proposition, there exists a set of fitness functions Ψ such that, for any act g and any $\eta \in \triangle(S)$,

$$\sup_{\psi\in\Psi}\int_S \psi(g(s))\, d\eta(s) = \int_Z u(z)\, d(\varphi\circ F_{g,\eta})(z). \tag{4}$$

Applying Equation (4) to the acts $g = f_\sigma$ and measures $\eta = \nu(\cdot|\sigma)$ for each signal σ in

[16] Formally, after dropping the expectation over Ω from Equation (1), we have

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

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

Equation (3) yields the ex ante value function in the proposition. The formula for the ex post value function follows by applying this same duality to the expression in Theorem 2. However, the value function $\hat{V}(g|\bar{\sigma},(f_\sigma))$ in this proposition differs from the function $V(g|\bar{\sigma},(f_\sigma))$ in Equation (2) not only because it is expressed as a rank-dependent utility, but also because we drop the logarithm and the conditional fitness associated with other signals. This is possible because ex post preferences do not depend on what happens after signals $\sigma \neq \bar{\sigma}$ in the case without common uncertainty. ■

Since we identify idiosyncratic uncertainty over S with pure risk, the distribution of outcomes $F_{f_\sigma,\nu(\cdot|\sigma)}$ amounts to an unambiguous risky prospect. Thus, given the appropriate set of fitness functions Ψ, conditional on each signal realization, adaptive preferences are equivalent to rank-dependent utility where individuals overweight the probability assigned to worse outcomes.[17]

In sharp contrast to the examples considered in Machina (1989), ex post preferences in Proposition 1 are independent of the plan $(f_\sigma)_{\sigma\in\Sigma}$. That is, consequentialism is not violated by this dynamically consistent version of the rank-dependent utility model with information. The intuition behind this result is that σ realizes prior to adaptation, and in our model, idiosyncratic risk that resolves before the selection of the hidden action is evaluated in accordance with expected utility. This is reflected by the ex ante value function, where each cdf $F_{f_\sigma,\nu(\cdot|\sigma)}$ over outcomes *given* σ is distorted by φ, rather than the unconditional distribution that also incorporates uncertainty about the realization of σ itself. In Section S2 of the Supplement, we show that if the hidden action ψ has to be chosen before the realization of σ, then the RDU distortion function will be applied to all uncertainty, including the signal realization, in which case consequentialism will be violated.

4.3 Ambiguity Aversion and Non-Expected Utility

Theorem 2 in Sadowski and Sarver (2024) provides a duality result for the representation of adaptive preferences that simultaneously incorporates both ambiguity aversion and non-expected utility for risk. The dual representation greatly facilitates the analysis of additional special cases. In this section, we generalize their theorem to allow for a nondegenerate signal structure.

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

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

[17] Sarin and Wakker (1998) considered a similar two-stage version of rank-dependent utility.

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

For any probability measure $p \in \triangle(\Omega)$, let

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

Since $R(p \,\|\, q) < \infty$ requires that $p \ll q$, if $q \in M(p)$ then the measures p and q are mutually absolutely continuous, that is, both $p \ll q$ and $q \ll p$.

Finally, for any $q \in \triangle(\Omega)$, define the measure $\nu \otimes q$ on $\Omega \times S \times \Sigma$ to have marginal q on Ω and conditional distribution ν_ω on $S \times \Sigma$. That is, for any event E in the product σ-algebra $\mathcal{B}_\Omega \otimes \mathcal{B}_S \otimes \mathcal{B}_\Sigma$, let

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

Theorem 3. *Suppose Ψ is a nonempty set of functions $\psi : Z \to [-\infty, \infty)$ that is pointwise bounded above and closed in the topology of pointwise convergence (on the extended reals), and fix measures $\mu \in \triangle(\Omega)$ and $\nu_\omega \in \triangle(S \times \Sigma)$ for each $\omega \in \Omega$. For any random action plan $\rho \in \mathcal{R}(\mathcal{F})$, the function V defined by Equation (1) can be equivalently expressed as*

$$V(\rho) = \inf_{q \in M(\mu)} \Bigg[\ln\Bigg(\int_\Sigma \mathbb{E}_{\rho_\sigma}\Bigg[\sup_{\psi\in\Psi} \int_{\Omega\times S} \psi(f(\omega, s))\, d(\nu \otimes q)(\omega, s|\sigma) \Bigg] d(\nu \otimes q)(\sigma) \Bigg) + R(\mu \,\|\, q) \Bigg]. \quad (5)$$

The proof is in Appendix A. The dual formula in Equation (5) is useful for analyzing special cases. For instance, the RDU duality used in Proposition 1 can also be applied to express the inner term in Equation (5),

$$\sup_{\psi\in\Psi} \int_{\Omega\times S} \psi(f(\omega, s))\, d(\nu \otimes q)(\omega, s|\sigma),$$

as rank-dependent expected utility with a pessimistic distortion function, given a suitable class of fitness functions Ψ. Other dual formulations of this inner expression are also possible for other specific classes of fitness functions Ψ, such as divergence preferences or certain cases of (probabilistically sophisticated) multiple prior preferences. Some of these dual forms are established and analyzed in Sadowski and Sarver (2024), and thanks to Theorem 3, they can also be applied in our current setting with signals.

By Theorem 2, preferences must be dynamically consistent. Note, however, that unlike in Proposition 1 where RDU preferences could satisfy both consequentialism and dynamic consistency, embedding RDU within our general model in Equation (5) leads to violations of consequentialism in the case of nontrivial common uncertainty for the same reasons discussed in Section 4.1.

Finally, individuals may still wish to randomize to hedge against common uncertainty, as discussed in Section 3.2. Interestingly, the idiosyncratic randomization created by conditioning choice on private signals in the general model in Equation (5) can sometimes alleviate the need for self-randomization, as we discuss in more detail in Section S3 of the Supplement.

A Appendix: Proof of Theorem 3

The first step in the proof of Theorem 3 is the following duality result.

Proposition 2. *Suppose Ψ is a nonempty set of functions $\psi : Z \to [-\infty, \infty)$ that is pointwise bounded above, and fix measures $\mu \in \triangle(\Omega)$ and $\nu_\omega \in \triangle(S \times \Sigma)$ for each $\omega \in \Omega$. For any random action plan $\rho \in \mathcal{R}(\mathcal{F})$, the function V defined by Equation (1) can be equivalently expressed as*

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

The proof of this proposition is based on applying the duality result from Proposition 4 in Sadowski and Sarver (2024) to the random variable $X : \Omega \to [-\infty, \infty)$ defined by

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

for a given $\rho \in (\triangle_s(\mathcal{F}))^\Sigma$ and $\tau \in \mathcal{R}(\Psi|\mathcal{F})$. The arguments are similar to the proof of Lemma 7 in Sadowski and Sarver (2024). We therefore omit the details for brevity.

The second step in the proof of the theorem involves switching the order of the supremum and infimum in the formula from Proposition 2. This accomplished using an extension of the von Neumann-Sion Minimax Theorem due to Tuy (2004). The following result from Sadowski and Sarver (2024) sets up the necessary technical machinery for our desired application.[18]

Proposition 3 (Sadowski and Sarver (2024))**.** *Fix measures $\mu \in \triangle(\Omega)$ and $\nu_\omega \in \triangle(S \times \Sigma)$ for each $\omega \in \Omega$. and suppose Ξ is a nonempty set of functions $\xi : \Omega \times S \times \Sigma \to [-\infty, \infty)$ with the following properties:*

[18]The statement of this result differs from the original statement in Proposition 5 of Sadowski and Sarver (2024), as that result does not include the signal space Σ. However, by taking the state space S to be $S \times \Sigma$ in their result, it is immediate that our current Proposition 3 follows.

1. Closedness: *When the set of extend reals $[-\infty,\infty]$ is endowed with its usual topology and $[-\infty,\infty]^{\Omega\times S\times\Sigma}$ is endowed with the product topology (i.e., the topology of pointwise convergence), Ξ is a closed subset of this space.*

2. Finite measurability: *There exists a finite partition $\mathcal{E}\subset\mathcal{B}_\Omega\otimes\mathcal{B}_S\otimes\mathcal{B}_\Sigma$ of $\Omega\times S\times\Sigma$ such that every $\xi\in\Xi$ is measurable with respect to $\mathcal{E}$.*

3. Pointwise boundedness: $\sup_{\xi\in\Xi}\xi(\omega,s,\sigma)<\infty$ *for every* $(\omega,s,\sigma)\in\Omega\times S\times\Sigma$.

Then,

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

Proceeding with the proof of Theorem 3, fix any $\rho\in\mathcal{R}(\mathcal{F})$. For each $\sigma\in\Sigma$, let $B_\sigma=\mathrm{supp}(\rho_\sigma)$. Since ρ_σ is a simple lottery, B_σ is finite. Let $B=\bigcup_{\sigma\in\Sigma}B_\sigma$. Since Σ is finite, B is a finite set of acts. Focus for the moment on adaptations plans τ that place probability one on some fitness function $\psi_{\sigma,f}\in\Psi$ following each $\sigma\in\Sigma$ and $f\in B$. Such deterministic adaptation plans are denoted by $(\psi_{\sigma,f})_{\sigma\in\Sigma,f\in B}\in\Psi^{\Sigma\times B}$, or $(\psi_{\sigma,f})$ for short.

Define Ξ to be the set of individual expected fitness functions that are attainable given the fixed random choice of act under the action plan ρ and together with some deterministic adaptation plan, that is, the set of all functions $\xi:\Omega\times S\times\Sigma\to[-\infty,\infty]$ that take the form

$$\xi(\omega,s,\sigma)=\int_B\psi_{\sigma,f}(f(\omega,s))\,d\rho_\sigma(f)\tag{6}$$

for some $(\psi_{\sigma,f})_{\sigma\in\Sigma,f\in B}$. Thus,

$$\Xi=\Big\{\xi\in[-\infty,\infty]^{\Omega\times S\times\Sigma}:\xi\text{ is defined by Equation (6) for some }(\psi_{\sigma,f})\in\Psi^{\Sigma\times B}\Big\}.\tag{7}$$

The next two lemmas show that taking the convex hull of Ξ generates precisely the set of individual expected fitness functions that can be attained through random adaptation plans and that the set Ξ satisfies the three conditions from Proposition 3.

Lemma 1. *Define Ξ as in Equation (7). For any random adaptation plan $\tau\in\mathcal{R}(\Psi|\mathcal{F})$, define $\xi^\tau:\Omega\times S\times\Sigma\to[-\infty,\infty)$ by*

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

Then,

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

Lemma 2. *The set Ξ defined in Equation (7) satisfies conditions 1–3 in Proposition 3.*

Lemma 1 is similar to Lemma 8 from Sadowski and Sarver (2024), but modified to incorporate signals. Lemma 2 is based on standard arguments. Therefore, the proofs of both of these lemmas are relegated to Section S8 of the Supplement.

Define V as in Equation (1). Then, we have

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

where the first equality follows from Proposition 2, the second from Lemma 1, the third from Proposition 3, and the fourth from the definition of the measure $\nu\otimes q$. Simple manipulations of the term inside the logarithm yield

$$\begin{aligned} &\sup_{\xi\in\Xi} \int_{\Omega\times S\times\Sigma} \xi(\omega,s,\sigma)\, d(\nu\otimes q)(\omega,s,\sigma) \\ &\quad= \sup_{(\psi_{\sigma,f})\in\Psi^{\Sigma\times B}} \int_\Sigma \int_{\Omega\times S} \int_B \psi_{\sigma,f}(f(\omega,s))\, d\rho_\sigma(f)\, d(\nu\otimes q)(\omega,s|\sigma)\, d(\nu\otimes q)(\sigma) \\ &\quad= \sup_{(\psi_{\sigma,f})\in\Psi^{\Sigma\times B}} \int_\Sigma \int_B \int_{\Omega\times S} \psi_{\sigma,f}(f(\omega,s))\, d(\nu\otimes q)(\omega,s|\sigma)\, d\rho_\sigma(f)\, d(\nu\otimes q)(\sigma) \\ &\quad= \int_\Sigma \int_B \sup_{\psi\in\Psi} \int_{\Omega\times S} \psi(f(\omega,s))\, d(\nu\otimes q)(\omega,s|\sigma)\, d\rho_\sigma(f)\, d(\nu\otimes q)(\sigma) \end{aligned}$$

and hence

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

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

References

Agranov, M. and P. Ortoleva (2017): "Stochastic Choice and Preferences for Randomization," *Journal of Political Economy*, 125, 40–68 (page 5).

Bergstrom, T. C. (2014): "On the Evolution of Hoarding, Risk-Taking, and Wealth Distribution in Nonhuman and Human Populations," *Proceedings of the National Academy of Sciences*, 111, 10860–10867 (page 5).

Cerreia-Vioglio, S., D. Dillenberger, P. Ortoleva, and G. Riella (2018): "Deliberately Stochastic," working paper (page 5).

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 (page 5).

Dominiak, A. and W. Schnedler (2011): "Attitudes Toward Uncertainty and Randomization: An Experimental Study," *Economic Theory*, 48, 289–312 (page 5).

Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75, 643–669 (page 9).

Epstein, L. G. and M. Le Breton (1993): "Dynamically Consistent Beliefs Must Be Bayesian," *Journal of Economic Theory*, 61, 1–22 (page 6).

Epstein, L. G. and M. Schneider (2003): "Recursive Multiple-Priors," *Journal of Economic Theory*, 113, 1–31 (page 10).

Hanany, E. and P. Klibanoff (2007): "Updating Preferences with Multiple Priors," *Theoretical Economics*, 2, 261–298 (pages 6, 8, 10).

Hanany, E. and P. Klibanoff (2009): "Updating Ambiguity Averse Preferences," *The B.E. Journal of Theoretical Economics*, 9 (page 8).

Ke, S. and Q. Zhang (2020): "Randomization and Ambiguity Aversion," *Econometrica*, 88, 1159–1195 (page 5).

Machina, M. J. (1985): "Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries," *The Economic Journal*, 95, 575–594 (page 5).

Machina, M. J. (1989): "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty," *Journal of Economic Literature*, 27, 1622–1668 (pages 11, 13).

Machina, M. J. and D. Schmeidler (1992): "A More Robust Definition of Subjective Probability," *Econometrica*, 60, 745–780 (page 6).

Nöldeke, G. and L. Samuelson (2005): "Information-Based Relative Consumption Effects: Correction," *Econometrica*, 73, 1383–1387 (page 5).

Oprea, R. and N. Robalino (2024): "Are Humans More Averse to Aggregate Risk? Testing an Evolutionary Economic Theory," working paper (page 1).

Raiffa, H. (1961): "Risk, Ambiguity, and the Savage Axioms: Comment," *Quarterly Journal of Economics*, 75, 690–694 (page 5).

Robson, A. J. (1996): "A Biological Basis for Expected and Non-Expected Utility," *Journal of Economic Theory*, 68, 397–424 (pages 1, 5).

Robson, A. J. and L. Samuelson (2011): "The Evolutionary Foundations of Preferences," in *Handbook of Social Economics*, vol. 1, Elsevier, 221–310 (page 5).

Sadowski, P. and T. Sarver (2024): "Adaptive Preferences: An Evolutionary Model of Non-Expected Utility and Ambiguity Aversion," *Journal of Economic Theory*, 218, 105840 (pages 1, 2, 4, 5, 9, 12–15, 17).

Saito, K. (2015): "Preferences for Flexibility and Randomization under Uncertainty," *American Economic Review*, 105, 1246–1271 (pages 5, 6).

Sarin, R. and P. Wakker (1998): "Dynamic Choice and Nonexpected Utility," *Journal of Risk and Uncertainty*, 17, 87–119 (page 13).

Tuy, H. (2004): "Minimax Theorems Revisited," *Acta Mathematica Vietnamica*, 29, 217–229 (page 15).

Online Supplement to: An Evolutionary Perspective on Updating Risk and Ambiguity Preferences

Philipp Sadowski Todd Sarver

December 17, 2024

In this online supplement, we explore several alternative assumptions and extensions of the analysis in the main text. Section S1 shows that self-randomization is a better hedging device than exogenous randomization that takes place after the resolution of ambiguity (common uncertainty). Section S2 shows that the specifics of our representation change when adaptation is slower and must be undertaken before the realization of the signal, yet evolutionarily optimal preferences remain dynamically consistent. Section S3 shows that responding to private signals results in idiosyncratic randomization in choice that can lessen, or in some cases even eliminate, the need for self-randomization by members of the population of a genotype. Section S4 examines how the optimal responses of a genotype to public and private signals differ. Section S5 draws lessons for our model from the lab, and vice versa. The interpretation of our model in terms of phenotypic flexibility is discussed in Section S6. Proofs of results in this supplement are contained in Section S7, and the proofs of Lemmas 1 and 2 from the main paper are contained in Section S8.

S1 Exogenous Versus Self-Randomization

For ease of exposition, we will focus in this section of the case of a trivial information structure $\Sigma = \{\sigma\}$, so that signals can be dropped from the model. However, the definitions below could be extended to the case of nontrivial signal structures, and the main result of this section would continue to hold.

A central behavioral feature of models of ambiguity aversion is that exogenous randomization can serve as a hedge against ambiguity. In the context of those models, this randomization is typically assumed to take place after ambiguity resolves, and hence after the choice of act, as otherwise the individual would be exposed to ambiguity in the interim.

It would be most natural to formalize this late exogenous randomization as a mixture in the Anscombe-Aumann framework where an act maps each common (ambiguous) state in Ω to a lottery over outcomes, $f : \Omega \to \triangle(Z)$. In keeping with the setup of our paper, however, we continue to model risk via the idiosyncratic state space S, and not via lotteries.

The following definition captures that, on the event where $f(\omega, s) = z$ and $g(\omega, s) = z'$, the mixture $\alpha f + (1-\alpha)g$ assigns probabilities α and $1-\alpha$ to z and z', respectively.

Definition S1. Suppose S is infinitely divisible and ν_ω is non-atomic for each $\omega \in \Omega$. For $z \in \text{supp}(f)$ and $z' \in \text{supp}(g)$, let

$$S_{f,g}(\omega, z, z') \equiv \{s \in S : f(\omega, s) = z,\ g(\omega, s) = z'\}.$$

For $\alpha \in (0,1)$, there is $S^\alpha_{f,g}(\omega, z, z') \subset S_{f,g}(\omega, z, z')$ such that

$$\nu_\omega(S^\alpha_{f,g}(\omega, z, z')) = \alpha \nu_\omega(S_{f,g}(\omega, z, z')).$$

Define the *mixture* $\alpha f + (1-\alpha)g$ by

$$(\alpha f + (1-\alpha)g)(\omega, s) = \begin{cases} f(\omega, s) & \text{if } s \in S^\alpha_{f,g}(\omega, f(\omega, s), g(\omega, s)) \\ g(\omega, s) & \text{otherwise .} \end{cases}$$

Note that there could be multiple acts h that satisfy this definition of $\alpha f + (1-\alpha)g$, since the sets $S^\alpha_{f,g}(\omega, z, z')$ are not uniquely determined. However, any act h that satisfies this definition will have the same long-run growth according to the value function V, so we allow any element of the resulting equivalence class of mixtures to be selected as $\alpha f + (1-\alpha)g$. Since the mixture of two acts is itself an act, the mixture $\sum_{f\in A} \rho(f) f$ according to a measure ρ over any decision problem A can be constructed recursively.

Contrary to the intuition that late randomization more convincingly hedges against ambiguity, in our model early randomization has an advantage, as it allows individuals to coordinate between observed actions (choice of act) and unobserved actions (adaptation). If, as formalized above, exogenous randomization takes place late, then the freedom to perform self-randomization early makes it the preferred hedging mechanism.

Proposition S1. *Suppose S is infinitely divisible and ν_ω is non-atomic for each $\omega \in \Omega$. For any decision problem A and random action plan $\rho \in \mathcal{R}(A)$,*

$$V(\rho) \geq V\left(\sum_{f\in A} \rho(f) f\right).$$

Proof. For $h = \sum_{f\in A} \rho(f) f$, the act h is deterministic. Hence, there exists a deterministic adaptation choice $\hat{\psi}$ that is optimal when evaluating h. Thus,

$$\begin{aligned}
V(h) &= \int_\Omega \ln\left(\int_S \hat{\psi}(h(\omega,s))\, d\nu_\omega(s)\right) d\mu(\omega) \\
&= \int_\Omega \ln\left(\int_S \mathbb{E}_\rho\big[\hat{\psi}(f(\omega,s))\big]\, d\nu_\omega(s)\right) d\mu(\omega) \\
&\leq \sup_{\tau\in\mathcal{R}(\Psi|\mathcal{F})} \int_\Omega \ln\left(\int_S \mathbb{E}_{\tau\otimes\rho}\big[\psi(f(\omega,s))\big]\, d\nu_\omega(s)\right) d\mu(\omega) = V(\rho).
\end{aligned}$$

where the second equality follows because h is constructed to generate the same distribution over outcomes as ρ. ■

We show that the inequality can be strict via the following example, which also serves to illustrate the result in an applied context.

Example S1. *The farmers in a community have to choose between planting one of two crops, f and g. There are two common states of the world, a rainy state and a dry state, denoted $\Omega = \{r, d\}$, where $\mu(r) = \mu(d) = 0.5$. Crop f produces high yield $\bar{f}$ in state r and low yield $\underline{f}$ in d. Crop g instead produces high yield $\bar{g}$ in d and low yield $\underline{g}$ in r. Suppose that there are two harvesting technologies, and each farmer makes an unobserved (hidden) investment in harvesting equipment at the time of choosing a crop and before the state resolves. Denote the two resulting fitness functions by ψ_1 and ψ_2. The first technology is suited to crop f: $\psi_1(\bar{f}) = 2$ and $\psi_1(z) = 1$ for $z \neq \bar{f}$. The second technology is suited to crop g: $\psi_2(\bar{g}) = 2$ and $\psi_2(z) = 1$ for $z \neq \bar{g}$. The individual reproductive fitness associated with each crop and technology combination in the two states $\omega = r, d$ is summarized in Table 1.*

	ψ_1	ψ_2
f	$2, 1$	$1, 1$
g	$1, 1$	$1, 2$

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

The signal space Σ and the idiosyncratic state space S play no role in this example and can be dropped from the long-run growth formula in Equation (1). We now consider deterministic action plans, action plans involving self-randomization, and exogenous randomization over acts (mixtures), and we compare the long-run growth rates associated with each.

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

$$V(f) = \int_\Omega \ln\big(\psi_1(f(\omega))\big) d\mu(\omega) = \frac{1}{2}\ln(2) + \frac{1}{2}\ln(1) \approx 0.3466$$
$$V(g) = \int_\Omega \ln\big(\psi_2(g(\omega))\big) d\mu(\omega) = \frac{1}{2}\ln(1) + \frac{1}{2}\ln(2) \approx 0.3466.$$

- ***Self-randomization:*** *To analyze random choice, we first determine the optimal joint plan π over crops and technology, and then deduce from it the optimal (observable) action plan ρ over crops. Let $\pi = \frac{1}{2}\delta_{(f,\psi_1)} + \frac{1}{2}\delta_{(g,\psi_2)}$ be the plan that randomizes uniformly over (f, ψ_1) and (g, ψ_2), pairing each crop with the appropriate harvest technology.*[19] *From a quick examination of Table 1, we see that this equal weight randomization eliminates common uncertainty and gives an average individual fitness of 1.5 in each state. The resulting action plan $\rho = \frac{1}{2}\delta_f + \frac{1}{2}\delta_g$ is the marginal distribution of π, and*

$$V\left(\frac{1}{2}\delta_f + \frac{1}{2}\delta_g\right) = \int_\Omega \ln\left(\frac{1}{2}\psi_1(f(\omega)) + \frac{1}{2}\psi_2(g(\omega))\right) d\mu(\omega) = \ln(1.5) \approx 0.4055.$$

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

$$\int_\Omega \ln\Big(\alpha\psi_1(h(\omega)) + (1-\alpha)\psi_2(h(\omega))\Big) d\mu(\omega) = \frac{1}{2}\ln\big(1 + 0.5\alpha\big) + \frac{1}{2}\ln\big(1.5 - 0.5\alpha\big).$$

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

$$V\left(\frac{1}{2}f + \frac{1}{2}g\right) = \ln(1.25) \approx 0.2231.$$

  Intuitively, since exogenous uncertainty about the realized crop hinders the farmer's ability to coordinate on the complementary harvesting equipment, the mixture of f and g is strictly worse than self-randomization between these two acts.

In the example, not only does self-randomization serve as a better hedge against common uncertainty or ambiguity than mixtures of acts, but those mixtures perform even worse than the original acts since the hedging benefit of the mixture is outweighed by the loss of fitness associated with sometimes mis-coordinating the hidden action (harvesting equipment) with the realized act (crop).

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

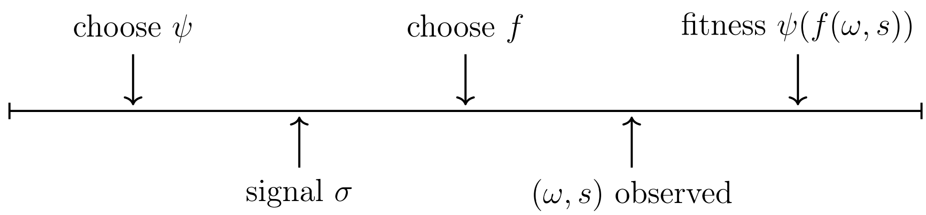

Figure S1: Within-period timeline: before-signal adaptation

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 fitness function also happens after the realization of a signal, reflecting the implicit assumption either that adaptation via selection of the hidden action can be undertaken rapidly or that signals arrive sufficiently early to allow time for such adaptation. We now consider the alternative of before-signal adaptation, where adaptation of the fitness function through the choice of hidden action is still fast enough to take into account the action plan, but too slow to react to the realization of a signal and the subsequent final choice of act. This alternative timing was mentioned briefly in Section 4.2 of the main paper in the context of rank-dependent utility.

For ease of illustration, we will focus on the long-run growth rates from deterministic action and adaptation plans. In the case of no common uncertainty (as in Section 4.2), deterministic plans will be optimal and our analysis is therefore sufficient for determining optimal choice. However, the reader should keep in mind in the case of common uncertainty, self-randomization may be optimal; extending our analysis accordingly is relatively straightforward and not central to the intuitions of this section.

Formally, the signal σ arrives after the choice of fitness function ψ, as illustrated in Figure S1. From the ex ante perspective, the individual thus selects an action plan $(f_\sigma)_{\sigma\in\Sigma}$ together with a fixed fitness function ψ, which achieves a fitness of $\psi(f_\sigma(\omega, s))$ after the realization of (ω, s, σ). Clearly, the growth rate will be lower than under after-signal adaptation, since fitness functions can no longer be optimized based on the signal realization. This will also generate subtle but important differences in the representation of evolutionarily optimal preferences over action plans. The following characterization follows from identical logic to Theorem 1. We therefore omit the proof.

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

in every period is

$$V\big((f_\sigma)_{\sigma\in\Sigma}\big) = \sup_{\psi\in\Psi}\int_\Omega \ln\bigg(\int_{S\times\Sigma}\psi(f_\sigma(\omega,s))\,d\nu_\omega(s,\sigma)\bigg)d\mu(\omega). \tag{S1}$$

The optimal fitness function ψ^* for plan $(f_\sigma)_{\sigma\in\Sigma}$ satisfies[20]

$$\psi^* \in \arg\max_{\psi\in\Psi}\int_\Omega \ln\bigg(\int_{S\times\Sigma}\psi(f_\sigma(\omega,s))\,d\nu_\omega(s,\sigma)\bigg)d\mu(\omega), \tag{S2}$$

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

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

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

Ex post adaptive preferences after learning signal $\bar\sigma$ now have to take into account not only the plan $(f_\sigma)_{\sigma\in\Sigma}$, but also the fitness function ψ^*, which is given at the time of choosing an act, as it was chosen optimally in conjunction with $(f_\sigma)_{\sigma\in\Sigma}$ prior to the realization of $\bar\sigma$. When Equation (S2) uniquely pins down ψ^*, ex post preferences are fully determined by $\bar\sigma$ and $(f_\sigma)_{\sigma\in\Sigma}$ alone, and so can be derived from ex ante preferences.

Consider three plans $(f_\sigma)_{\sigma\in\Sigma}$, $(g_\sigma)_{\sigma\in\Sigma}$, and $(h_\sigma)_{\sigma\in\Sigma}$ such that $f_\sigma = g_\sigma = h_\sigma$ for all $\sigma\neq\bar\sigma$. Suppose $(f_\sigma)_{\sigma\in\Sigma}$ is strictly optimal ex ante, and suppose ψ^* is the corresponding uniquely optimal fitness function. Since ex ante adaptive preferences incorporate the optimal choice of ψ while ex post preferences take ψ^* as given, it is possible to have $(g_\sigma)_{\sigma\in\Sigma}\succ(h_\sigma)_{\sigma\in\Sigma}$ ex ante and $g_{\bar\sigma}\prec_{\bar\sigma,(f_\sigma)} h_{\bar\sigma}$ ex post. In other words, the ranking of two *suboptimal* plans can change ex post. This is not a violation of our notion of dynamic consistency, which only requires no deviations from the *optimal* plan, and hence only applies when comparing $(f_\sigma)_{\sigma\in\Sigma}$ to the other plans. However, it does violate stronger notions commonly found in the literature, for instance, the definitions found in Machina and Schmeidler (1992) and Epstein and Le Breton (1993).[21] The following example illustrates that those violations do not depend on the arrival

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

[21] Preferences in the case of after-signal adaptation that we considered in the main text will satisfy this stronger notion of dynamic consistency: For plans $(f_\sigma)_{\sigma\in\Sigma}$, $(g_\sigma)_{\sigma\in\Sigma}$, and $(h_\sigma)_{\sigma\in\Sigma}$ such that $f_\sigma = g_\sigma = h_\sigma$ for all $\sigma\neq\bar\sigma$, we have $(g_\sigma)_{\sigma\in\Sigma}\succ(h_\sigma)_{\sigma\in\Sigma} \implies g_{\bar\sigma}\succ_{\bar\sigma,(f_\sigma)} h_{\bar\sigma}$ (and $(g_\sigma)_{\sigma\in\Sigma}\succsim(h_\sigma)_{\sigma\in\Sigma} \implies g_{\bar\sigma}\succsim_{\bar\sigma,(f_\sigma)}$

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

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

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

The special case of rank-dependent expected utility also demonstrates the importance of the timing of adaptation.

Corollary S1 (RDU with Before-Signal Adaptation)**.** *Suppose* $\Omega = \{\omega\}$ *and* $Z \subset \mathbb{R}$*. Fix* $\nu \in \triangle(S \times \Sigma)$*, and fix any bounded nondecreasing function* $u : Z \to \mathbb{R}$ *and any function* $\varphi : [0, 1] \to [0, 1]$ *that is continuous, nondecreasing, concave, and onto. Then, there exists a set* Ψ *of functions* $\psi : Z \to \mathbb{R}$ *such that the ex ante value function* V *defined by Equation (S1) can be equivalently expressed as*

$$V\big((f_\sigma)_{\sigma\in\Sigma}\big) = \ln\bigg(\int_Z u(z)\, d(\varphi \circ F_{(f_\sigma),\nu})(z)\bigg)$$

where

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

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

In contrast to the results from the main paper where the distortion function only applied to the distribution of s conditional on each signal realization σ, in the case of 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). Of course, ex

$h_{\bar\sigma}$ whenever $\bar\sigma$ occurs with strictly positive probability). This is because for after-signal adaptation, the conditional preference $\succsim_{\bar\sigma,(f_\sigma)}$ does not depend on $f_{\bar\sigma}$, only on f_σ for $\sigma \neq \bar\sigma$. Note that in terms of observable behavior, the two notions are typically equivalent, as choice can only reveal whether or not an individual prefers deviating from the ex ante optimal plan.

post preferences will still satisfy our notion of dynamic consistency, but will now in general violate consequentialism.[22]

S3 Signal Response in lieu of Self-Randomization

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

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

With slight abuse of notation, let $\rho_\sigma(\omega)$ denote the probability that an individual bets on state ω after observing σ.[23] If the state on which the individual bets realizes then their payoff is 1; otherwise, their payoff is 0. Assume that $\psi(1) > \psi(0) \geq 0$. The long-run growth rate is now given by

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

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

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

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

Proposition S2. *Fix two states $\omega, \omega' \in \Omega$ and two signals $\sigma, \sigma' \in \Sigma$. If $\mu(\omega) > 0$ and $\mu(\omega') > 0$, and if*

$$\nu_\omega(\sigma)\, \nu_{\omega'}(\sigma') > \nu_\omega(\sigma')\, \nu_{\omega'}(\sigma),$$

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

The proof of Proposition S2 is in Section S7.1. In the case where the probabilities in the proposition are strictly positive, the inequality in the proposition can be written as

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

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 S3. *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 in the main paper, 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, with*

$$\nu_r(\sigma) = \frac{4}{5},\ \nu_r(\sigma') = \frac{1}{5} \quad \textit{and} \quad \nu_y(\sigma) = 1,\ \nu_y(\sigma') = 0$$

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

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

which is maximized by taking

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

Thus, there is no randomization contingent on signal σ'. There is, however, randomization contingent on σ. Intuitively, since σ' is much less likely, exclusively conditioning on the two informative signals by taking $\rho_\sigma(Y) = 1$ would lead to excess correlation in outcomes across individuals.[24]

[24] In some cases, conditioning on informative signals may completely eliminate self-randomization. In the example, if instead $\nu_r(\sigma) = 2/5$, $\nu_r(\sigma') = 3/5$, $\nu_y(\sigma) = 1$, and $\nu_y(\sigma') = 0$, then $\rho_\sigma(Y) = 1$ and $\rho_{\sigma'}(R) = 1$, so there is no randomization following either signal. In this case, removing residual correlation through randomization is not worth the cost of worsening the expected individual outcomes.

S4 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 $\mu \in \triangle(\Omega)$ and measures $\nu_\omega \in \triangle(\Sigma)$ for $\omega \in \Omega$, let $V^{\mathrm{Pr}}(\rho)$ denote the now familiar long-run growth rate for the action plan ρ under private signals:

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

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

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

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

$$\rho^{\mathrm{Pr}} \in \operatorname*{argmax}_{\widehat{\rho}\in\mathcal{R}(A)} V^{\mathrm{Pr}}(\widehat{\rho}) \qquad \text{and} \qquad \rho^{\mathrm{Pu}} \in \operatorname*{argmax}_{\widehat{\rho}\in\mathcal{R}(A)} V^{\mathrm{Pu}}(\rho)$$

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

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

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

A more subtle question is how ρ^{Pu} and ρ^{Pr} differ. We already saw in Section S3 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 next briefly discuss an illustrative example to portfolio choice.

Application: Portfolio Choice

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

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

Proposition S3. *Let σ_* and σ^* be the signals that induce the lowest and highest investment in the risky asset under private signals, respectively, that is, $\alpha_{\sigma_*}^{\mathrm{Pr}} \le \alpha_\sigma^{\mathrm{Pr}} \le \alpha_{\sigma^*}^{\mathrm{Pr}}$ for all $\sigma \in \Sigma$. If $\alpha_{\sigma_*}^{\mathrm{Pr}} \neq \alpha_{\sigma^*}^{\mathrm{Pr}}$, then the following must be true:*

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

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 S3 is contained in Section S7.2.

S5 Lessons from and for the Lab

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

S5.1 Testing Dynamic Consistency

When investigating dynamic consistency between ex ante signal-contingent plans and ex post choice after a particular signal realization, it is often implicitly assumed that the information structure that gives rise to that signal realization is irrelevant for ex post choice. For instance, Dominiak, Duersch, and Lefort (2012) considered a decision situation similar to Example 1 of our main paper. To efficiently collect ex post preferences contingent on the signal $\neg y$, their experimental design redraws the ball from the urn until the ball comes up black or red, so that no individual ever learns y.[25] While this protocol may be adequate in the context of some theoretical models, our model of adaptive preferences predicts that dynamic consistency and consequentialism are at odds with each other only when there are complementarities in payoffs between individuals who receive different signals. If no individuals receive the counterfactual signal (in this case y), then our model will not predict any violations of consequentialism.

To illustrate, recall that in the context of Example 1 it is optimal to follow the plan $R\neg yY$ after learning $\neg y$ only because this choice hedges against fluctuations in population growth from individuals who learn y. However, if subjects never learn y, then this hedging motive disappears and adaptive preferences will instead favor B over R following $\neg y$, in line with the experimental findings in Dominiak, Duersch, and Lefort (2012):

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

for any ex ante action plan $(f_\sigma)_{\sigma\in\Sigma}$. Note that our model predicts this ranking both ex ante and ex post, so preferences are both dynamically consistent and consequentialist given this decision setting. This demonstrates that in order to test our model's prediction of dynamic consistency, the entire event tree must be implemented.

A recent experiment by Bleichrodt et al. (2020) that did implement the entire game tree found that consequentialism was satisfied slightly more often than dynamic consistency

[25] Dominiak, Duersch, and Lefort (2012) used different colors for the balls in their experiment. We have translated to the colors used in our Example 1 in the main text for ease of exposition.

(73.2% versus 66.2% of subjects). Interestingly, although not the focus of their experiment, some of their experimental evidence seems to suggest that violations of dynamic consistency may be connected with cognitive constraints and narrow bracketing. In particular, the primary violation of dynamic consistency observed in their experiment also constitutes a violation of monotonicity (stochastic dominance), and in many cases of consequentialism, if subjects properly integrate payoffs across the different questions in the experiment.[26]

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

Table 2: Composition of Ambiguous Bag (Table 1 from Bleichrodt et al. (2020))

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

These experimental findings draw to light an interesting parallel and potential future research question: Narrow bracketing constitutes a failure to integrate payoffs from different aspects of one's overall choice situation when making decisions. Consequentialism by definition involves not integrating other unrealized branches of the decision tree when making choices. In the experiment by Bleichrodt et al. (2020), there seems to be a connection between the two, since making the purportedly consequentialist choice of red after learning the parity of the card could only be viewed as optimal if subjects narrowly bracket the

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

payments from different questions in the experiment. This begs the empirical question of whether these two behavioral patters—narrow bracketing and consequentialism (at the expense of dynamic consistency)—correlate across individuals, and if so, whether both should be considered mistakes as our model would suggest.

S5.2 Ambiguous Signals

While many models of updating with ambiguity consider an ambiguous prior with unambiguous signals, it is also possible that the information content of the signals themselves is ambiguous. For instance, there may be no prior ambiguity at all until the result of a poorly understood test becomes available. In a recent laboratory experiment, Epstein and Halevy (2024) examined the response of subjects to signals that have ambiguous precision, and they documented violations of the martingale property of beliefs.

Shishkin and Ortoleva (2023) subsequently tested one striking implication of all common consequentialist approaches to updating models of ambiguity aversion, namely that ambiguous information can have negative value. This implication seems counterintuitive, and indeed their evidence casts doubt on it: Ambiguity-averse subjects appear to ignore information unless it is valuable. Reacting to information only when this adds value is precisely what maximizes ex ante expected utility, and is hence the dynamically consistent course of action. In other words, their experimental findings are in line with our model of adaptive preferences, and its evolutionary foundation provides a rationale for them.

S6 Phenotypic Flexibility in Evolutionary Biology

While our approach is inspired by evolutionary biology, we hope that our insights might in turn also be useful in biological contexts where phenotypic flexibility plays a role, as we now explain in more detail. Evolutionary success appears to be greatly enhanced by the ability of organisms of a particular genotype to adapt their phenotype to the environment. Adopting the terminology proposed by Piersma and Drent (2003), we use *phenotypic flexibility* to refer to the rapid and apparently purposeful variation in phenotype expressed by individual reproductively mature organisms throughout their life. This is in contrast to *developmental plasticity*, environmentally induced variations that occur only during development.[27]

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):

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

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

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

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

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

S7 Proofs of Results in the Online Appendix

S7.1 Proof of Proposition S2

As in Section S3, with slight abuse of notation let $\rho_\sigma(\omega)$ denote the probability that an individual bets on state ω following signal σ. Let $\overline{\rho}(\omega)$ denote the probability of betting on state ω when the actual state is ω, given $\rho_\sigma(\omega)$ and $\nu_\omega(\sigma)$. That is,

$$\overline{\rho}(\omega) = \sum_{\sigma \in \Sigma} \rho_\sigma(\omega)\nu_\omega(\sigma).$$

Simple direct computation yields the partial derivative of V with respect to $\rho_\sigma(\omega)$:[28]

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

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

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

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

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

Multiplying these two expressions, we obtain

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

Using the formula for the partial derivative and the assumption that $\mu(\omega)>0$ and $\mu(\omega')>0$ and rearranging terms, this implies that

$$\nu_\omega(\sigma')\,\nu_{\omega'}(\sigma) \geq \nu_\omega(\sigma)\,\nu_{\omega'}(\sigma').$$

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

S7.2 Proof of Proposition S3

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

$$\begin{aligned}
V^{\mathrm{Pu}}((\alpha_\sigma)_{\sigma\in\Sigma}) &= \int_\Omega \int_\Sigma \ln\Big(\psi\big(\alpha_\sigma f(\omega)-(1-\alpha_\sigma)c\big)\Big)\,d\nu_\omega(\sigma)\,d\mu(\omega)\\
V^{\mathrm{Pr}}((\alpha_\sigma)_{\sigma\in\Sigma}) &= \int_\Omega \ln\bigg(\int_\Sigma \psi\big(\alpha_\sigma f(\omega)-(1-\alpha_\sigma)c\big)\,d\nu_\omega(\sigma)\bigg)d\mu(\omega).
\end{aligned}$$

Observe first that for any $(\alpha_\sigma)_{\sigma\in\Sigma}$ and any $\sigma\in\Sigma$,

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

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

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

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
$$\frac{\psi'(\alpha f(\omega)-(1-\alpha)c)}{\psi(\alpha f(\omega)-(1-\alpha)c)}(f(\omega)-c) > \frac{\psi'(\alpha f(\omega)-(1-\alpha)c)}{\sum_{\sigma\in\Sigma}\nu_\omega(\sigma)\psi(\alpha_\sigma f(\omega)-(1-\alpha_\sigma)c)}(f(\omega)-c).$$
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 claim in the proposition, that $\alpha^{\mathrm{Pr}}_{\sigma_*}\geq\alpha^{\mathrm{Pu}}_{\sigma_*}$ and $\alpha^{\mathrm{Pr}}_{\sigma_*}>0$. Then, we have
$$\frac{\partial V^{\mathrm{Pu}}((\alpha^{\mathrm{Pu}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma_*}}\geq\frac{\partial V^{\mathrm{Pu}}((\alpha^{\mathrm{Pr}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma_*}}>\frac{\partial V^{\mathrm{Pr}}((\alpha^{\mathrm{Pr}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma_*}},$$
where the first inequality follows from observation 1 since $\alpha^{\mathrm{Pr}}_{\sigma_*}\geq\alpha^{\mathrm{Pu}}_{\sigma_*}$, and the second inequality follows from observation 2 with $\alpha=\alpha^{\mathrm{Pr}}_{\sigma_*}$ since $\alpha^{\mathrm{Pr}}_{\sigma_*}\leq\alpha^{\mathrm{Pr}}_{\sigma}$ for all $\sigma\in\Sigma$ (with strict inequality for at least one σ). Since, by assumption, $\alpha^{\mathrm{Pu}}_{\sigma_*}\leq\alpha^{\mathrm{Pr}}_{\sigma_*}<\alpha^{\mathrm{Pr}}_{\sigma^*}\leq 1$, the optimality of α^{Pu} requires that
$$\frac{\partial V^{\mathrm{Pu}}((\alpha^{\mathrm{Pu}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma_*}}\leq 0,$$
and hence
$$\frac{\partial V^{\mathrm{Pr}}((\alpha^{\mathrm{Pr}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma_*}}<0.$$
Since $(\alpha^{\mathrm{Pr}}_\sigma)_{\sigma\in\Sigma}$ is optimal, this requires that $\alpha^{\mathrm{Pr}}_{\sigma_*}=0$, a contradiction. This establishes the first claim in the proposition.

Finally suppose, contrary to the second part of the proposition, that $\alpha^{\mathrm{Pr}}_{\sigma^*}\leq\alpha^{\mathrm{Pu}}_{\sigma^*}$ and $\alpha^{\mathrm{Pr}}_{\sigma^*}<1$. Then, we have
$$\frac{\partial V^{\mathrm{Pu}}((\alpha^{\mathrm{Pu}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma^*}}\leq\frac{\partial V^{\mathrm{Pu}}((\alpha^{\mathrm{Pr}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma^*}}<\frac{\partial V^{\mathrm{Pr}}((\alpha^{\mathrm{Pr}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma^*}},$$
where the first inequality follows from observation 1 since $\alpha^{\mathrm{Pr}}_{\sigma^*}\leq\alpha^{\mathrm{Pu}}_{\sigma^*}$, and the second inequality follows from observation 2 with $\alpha=\alpha^{\mathrm{Pr}}_{\sigma^*}$ since $\alpha^{\mathrm{Pr}}_{\sigma}\leq\alpha^{\mathrm{Pr}}_{\sigma^*}$ for all $\sigma\in\Sigma$ (with strict inequality for at least one σ). Since, by assumption, $0\leq\alpha^{\mathrm{Pr}}_{\sigma_*}<\alpha^{\mathrm{Pr}}_{\sigma^*}\leq\alpha^{\mathrm{Pu}}_{\sigma^*}$, the optimality of $(\alpha^{\mathrm{Pu}}_\sigma)_{\sigma\in\Sigma}$ requires that
$$\frac{\partial V^{\mathrm{Pu}}((\alpha^{\mathrm{Pu}}_\sigma)_{\sigma\in\Sigma})}{\partial\alpha_{\sigma^*}}\geq 0,$$

and hence

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

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

S8 Omitted Proofs from the Main Paper

In this section, we prove Lemmas 1 and 2 from the main paper. We restate the results here for ease of reference.

Lemma 1. *Define* Ξ *as in Equation (7). For any random adaptation plan* $\tau \in \mathcal{R}(\Psi|\mathcal{F})$*, define* $\xi^\tau : \Omega \times S \times \Sigma \to [-\infty, \infty)$ *by*

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

Then,

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

Lemma 2. *The set* Ξ *defined in Equation (7) satisfies conditions 1–3 in Proposition 3.*

S8.1 Proof of Lemma 1

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

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

where we define $\tau \in \mathcal{R}(\Psi|\mathcal{F})$ for each $\sigma \in \Sigma$ and $f \in B$ by[29]

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

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

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

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

has finite support. That is, there exists a product measure η on $\Psi^{\Sigma\times B}$ with finite support, defined by

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

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

$$\operatorname{supp}(\eta) = \big\{(\psi^1_{\sigma,f}),\ldots,(\psi^n_{\sigma,f})\big\}.$$

Thus,

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

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

S8.2 Proof of Lemma 2

Condition 1 — Closedness: To establish that Ξ is closed, it is helpful to first formulate an alternative, equivalent definition of this set. Denoting deterministic adaptation plans by $(\psi_{\sigma,f})_{\sigma\in\Sigma,f\in B}\in\Psi^{\Sigma\times B}$, or $(\psi_{\sigma,f})$ for short, we can define a mapping $J:\Psi^{\Sigma\times B}\to[-\infty,\infty]^{\Omega\times S\times\Sigma}$ by

$$J\Big[(\psi_{\hat\sigma,\hat f})_{\hat\sigma\in\Sigma,\hat f\in B}\Big](\omega,s,\sigma) = \int_B \psi_{\sigma,f}(f(\omega,s))\,d\rho_\sigma(f) \tag{S4}$$

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

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

Next, note that the set $[-\infty,\infty]$ is a compact Hausdorff space when endowed with its usual topology.[30] By the Tychonoff Product Theorem (Theorem 2.61 in Aliprantis and Border (2006)), the set $[-\infty,\infty]^Z$ endowed with the product topology (also know as the topology of pointwise convergence) is compact. Since $\Psi\subset[-\infty,\infty]^Z$ is closed, it is also compact. Applying the Tychonoff Product Theorem again, the set $\Psi^{\Sigma\times B}$ is compact in the product topology.

We next show that the mapping $J:\Psi^{\Sigma\times B}\to[-\infty,\infty]^{\Omega\times S\times\Sigma}$ defined in Equation (S4) is continuous when $[-\infty,\infty]^{\Omega\times S\times\Sigma}$ is endowed with the product topology. To see this, fix any net $(\psi^\alpha_{\sigma,f})_{\alpha\in D}$ in $\Psi^{\Sigma\times B}$ that converges to some $(\psi_{\sigma,f})\in\Psi^{\Sigma\times B}$. We will show that $J[(\psi^\alpha_{\sigma,f})]$ converges

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

to $J[(\psi_{\sigma,f})]$.[31] First, by the definition of the product topology, convergence of the net $(\psi^\alpha_{\sigma,f})$ implies that $\psi^\alpha_{\sigma,f}(z) \to \psi_{\sigma,f}(z)$ for all σ, f, and z. In particular, $\psi^\alpha_{\sigma,f}(f(\omega,s)) \to \psi_{\sigma,f}(f(\omega,s))$ for all σ, f, ω, and s. Therefore, since convergence is preserved under scalar multiples and finite sums,

$$\sum_{f\in B} \psi^\alpha_{\sigma,f}(f(\omega,s))\rho_\sigma(f) \to \sum_{f\in B} \psi_{\sigma,f}(f(\omega,s))\rho_\sigma(f)$$

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

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

Condition 2 — Finite Measurability: We now show that Ξ satisfies the second condition (finite measurability) from Proposition 3. Since each $f \in \mathcal{F}$ is a simple act, and since the set of acts B in the support of ρ is finite, there exists a finite partition $\widehat{\mathcal{E}} \subset \mathcal{B}_\Omega \otimes \mathcal{B}_S$ of $\Omega \times S$ such that every act $f \in B$ is measurable with respect to $\widehat{\mathcal{E}}$. Let

$$\mathcal{E} = \left\{\widehat{E} \times \{\sigma\} : \widehat{E} \in \widehat{\mathcal{E}} \text{ and } \sigma \in \Sigma\right\}.$$

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

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

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

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

as claimed. Thus, the second condition of Proposition 3 is satisfied.

Condition 3 — Pointwise Boundedness: To verify the third condition (pointwise boundedness) in Proposition 3, note that since B is a finite set of simple acts, there is a finite set $\widehat{Z} \subset Z$ such that $f(\omega,s) \subset \widehat{Z}$ for all $f \in B$ and $(\omega,s) \in \Omega \times S$. Recall that the set Ψ is pointwise bounded above, so $\sup_{\psi\in\Psi}\psi(z) < \infty$ for all $z \in Z$. Therefore, and any $(\omega,s,\sigma) \in \Omega\times S\times\Sigma$,

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

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

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

References for Supplementary Appendix

Aliprantis, C. and K. Border (2006): *Infinite Dimensional Analysis*, 3rd ed., Berlin, Germany: Springer-Verlag (pages S19, S20).

Bleichrodt, H., J. Eichberger, S. Grant, D. Kelsey, and C. Li (2020): "Testing Dynamic Consistency and Consequentialism under Ambiguity," working paper (pages S12, S13).

Daniel, K., D. Hirshleifer, and A. Subrahmanyam (1998): "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance*, 53, 1839–1885 (page S11).

Dominiak, A., P. Duersch, and J.-P. Lefort (2012): "A Dynamic Ellsberg Urn Experiment," *Games and Economic Behavior*, 75, 625–638 (page S12).

Epstein, L. G. and Y. Halevy (2024): "Hard-to-Interpret Signals," *Journal of the European Economic Association*, 22, 393–427 (page S14).

Epstein, L. G. and M. Le Breton (1993): "Dynamically Consistent Beliefs Must Be Bayesian," *Journal of Economic Theory*, 61, 1–22 (pages S6, S8).

Flück, M. (2006): "Functional, Structural and Molecular Plasticity of Mammalian Skeletal Muscle in Response to Exercise Stimuli," *Journal of Experimental Biology*, 209, 2239–2248 (page S15).

Hanany, E. and P. Klibanoff (2007): "Updating Preferences with Multiple Priors," *Theoretical Economics*, 2, 261–298 (page S8).

Kuzmics, C. (2017): "Abraham Wald's Complete Class Theorem and Knightian Uncertainty," *Games and Economic Behavior*, 104, 666–673 (page S13).

Machina, M. J. (1989): "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty," *Journal of Economic Literature*, 27, 1622–1668 (page S7).

Machina, M. J. and D. Schmeidler (1992): "A More Robust Definition of Subjective Probability," *Econometrica*, 60, 745–780 (pages S6, S8).

Piersma, T. and J. Drent (2003): "Phenotypic Flexibility and the Evolution of Organismal Design," *Trends in Ecology & Evolution*, 18, 228–233 (page S14).

Shishkin, D. and P. Ortoleva (2023): "Ambiguous Information and Dilation: An Experiment," *Journal of Economic Theory*, 208, 105610 (page S14).

Wikelski, M. and C. Thom (2000): "Marine Iguanas Shrink to Survive El Niño," *Nature*, 403, 37 (page S15).

Wright, P. A. and A. J. Turko (2016): "Amphibious Fishes: Evolution and Phenotypic Plasticity," *Journal of Experimental Biology*, 219, 2245–2259 (page S15).