Evolution with Private Information: Caution, Contrarianism and Herding*

Fernando M. Louge
Department of Economics
University of Wisconsin - Madison

Abstract

This paper considers a model where agents receive private signals correlated with the unknown state of the world. The standard approach to this problem is to assume that agents maximize their (objective) expected utility based on their Bayesian posteriors. We present a repeated, non-strategic version of this model and show that the expected utility rule is evolutionarily suboptimal. We provide a characterization of the evolutionarily optimal rule. Compared to the behavior rule that maximizes the expected utility, our evolutionary criterion provides more 'smoothing' of the population growth rate across states of the world. This translates into two properties of the optimal behavior rule: contrarian behavior and caution. Contrarian behavior consists of a probabilistic bias towards actions that defy the 'common wisdom' embedded in the prior beliefs. Agents exhibit caution when, compared to expected utility maximizers, a more extreme prior is required before disregarding their private information. We extend the model of social learning of Smith and Sørensen (2000) to a general class of behavior rules that includes the evolutionary and the expected utility behavior rules. We show that the qualitative properties of the model are preserved within this class. In particular, herds eventually arise. The limit distributions of public beliefs, however, are different. We find that our evolutionary-founded rule induces herding on the optimal action with higher probability than the expected utility rule.

JEL Classification: D81, D82, D83

1 Introduction

Economic agents often rely on private information to make decisions when the state of the world is unknown. Some well known examples include bidders in common

*I am indebted to Bill Sandholm for his guidance and support in all stages of this project. I am also grateful to Dan Quint, Ricardo Serrano-Padial and seminar participants at the University of Wisconsin - Madison for their comments. All remaining errors are mine.

value auctions (Milgrom and Weber, 1982), informed voters choosing among candidates (Feddersen and Pesendorfer, 1996, 1997), traders of financial assets in competitive markets (Grossman, 1976) and oligopolists with unknown common costs (Palfrey, 1985). These models share and important feature: the state of the world not only affects the payoffs of a particular agent, but it rather affects a large group of them. In other words, the realization of a particular state of the world is aggregate risk.

The standard approach to this problem in economics and other disciplines is to assume that agents maximize their (objective) expected utility. The probability distribution is obtained from the Bayesian update of the prior beliefs after observing the private signal. We refer to this approach as the expected utility criterion.

In this paper we asses whether this approach can be justified using evolutionary arguments. More concretely, we ask whether the behavior rule induced by the expected utility criterion is evolutionarily successful. We find that this is not the case. Our next question is then, what *is* the behavior rule selected by evolution?

We will consider a sequence of generations, each facing an (ex-ante) identical decision problem with uncertainty as the ones described above. The initial society consists of agents that behave according to several behavior rules. Higher payoffs for a particular agent imply more agents carrying the same behavior trait in the subsequent generation. We keep track of the number of agents following each behavior rule.

We find that there is a unique behavior rule that dominates in the long run. That is, the fraction of the society that follows one particular rule converges to one. Moreover, this rule is not consistent with the expected utility criterion.

In the presence of aggregate risk, the population growth rate depends on the realization of a random variable even in the long run, when the population is large. Therefore a bad realization of this random variable has a multiplicative effect on the population size for all subsequent generations. Evolution 'penalizes' behavior rules that allow for large differences in the population growth rates across states. Put differently, evolution selects a pattern of population growth rates across states that is 'smooth'.

Because of the multiplicative nature of population dynamics, the appropriate measure of evolutionary success is the geometric mean across states. That is, agents that maximize this objective function are evolutionarily successful. Conditional on a given state, however, evolution selects according to the standard arithmetic expectation. This is consistent with the findings of Blume and Easley (1992) and

Robson (1996)[1].

Once the correct measure of evolutionary success is established, we turn to one of the main questions of this paper: what is the evolutionarily optimal behavior after receiving private information about the underlying state of the world? Building on the ideas described above, we characterize the evolutionary-founded behavior rule. We show that the evolutionary rule provides more smoothing across state-contingent payoffs than the expected utility rule. This is achieved through two mechanisms: *contrarian behavior* and *caution*. If the prior beliefs indicate that a particular state is relatively more likely, a contrarian behavior rule exhibits a bias towards actions that are optimal in *other* states. If agents are too keen on following the 'common wisdom', then there is a high correlation of choices across the population. As a consequence, there is a large difference in state-contingent expected payoffs. Contrarian behavior partially insures the population against this form of correlation.

The evolutionarily optimal rule exhibits also caution. In our model, the informativeness of private signals is bounded. In this case, it is well known that under the expected utility behavior rule, agents disregard their private information if the prior is informative enough. This is also true under the evolutionary behavior rule. However, the smoothing introduced by the geometric mean fitness implies that priors have to be even more informative before making private signals obsolete.

As it was discussed above, the evolutionary rule in the presence of aggregate uncertainty does not have an expected utility representation. In addition to this, we find another departure from the standard expected utility model: the evolutionarily optimal behavior rule cannot be fully expressed in terms of Bayesian posteriors. This follows from the fact that the evolutionary rule depends directly on the state-contingent expected payoffs. Therefore, the complete distribution functions for private signals are necessary to compute the optimal behavior rule. In contrast, Bayesian posteriors are only 'local' statistics: they depend only on the probabilities of the signal that was actually received and they ignore any other information from the distribution function.

Despite the mentioned differences, the evolutionary rule shares many properties with the expected utility maximizing rule. In some applications, these similarities permit studying the consequences of the evolutionary rule using tools similar to those developed in the literature for the expected utility rule. In this paper we take this approach to analyze evolutionary rules in a well-known social learning model.

[1] There is also a significant literature in biology that uses this notion. See for example Crow and Kimura (1970), Cooper and Kaplan (1982) and McNamara (1995).

Section 4 presents a model of sequential social learning à la Bikhchandani, Hirshleifer and Welch (1992), Banerjee (1992) and Smith and Sørensen (2000). In this model, agents make a single decision sequentially, with the order of decisions exogenously given. As in our original setting, agents receive private signals correlated with the unknown state of the world. In addition, agents observe the decisions made by previous agents, but not their signals. While our previous analysis was concerned with a infinite generation problem, in this section we will consider only one generation and the state is realized once.

The main result of this literature is that eventually, the history of past play carries more information than a single private signal, so agents disregard the latter. When this occurs, the history of play stops incorporating new information and all subsequent agents choose the same actions. This phenomenon is called *herding*.

In our social learning model we consider monomorphic populations. That is, populations whose members follow the same rule from a class of behavior rules that we call CM-threshold rules. Within this class, that agents disregard their private signals when the informativeness of the public signal is higher than some value. The expected utility and the evolutionary rules are particular cases of this class. To start our analysis, we extend some of the results of Smith and Sørensen (2000) to any CM-threshold rule. In particular, we show that public beliefs converge in the long run and herds eventually arise.

Notice that in this context, public information is now endogenous: the choices made by agents affect the public signal received by subsequent agents. In other words, the social learning model presents an informational externality. Therefore, the evolutionary rule is not necessarily optimal in the social learning model. Nevertheless, we find that the probability of herding on the optimal action is higher for the evolutionary rule than for the expected utility rule.

The intuition behind our result is the following. There are three main differences between the evolutionary and the expected utility rule that have to be taken into account. First, the evolutionary rule exhibits caution. Therefore it requires more information from the history of play in order to disregard the private signal. This effect increases the probability of herding on the correct action under the evolutionary rule. Another property of the evolutionary rule is contrarianism. The second difference is that with contrarian behavior, it is less likely that individuals follows suit when the history of past play suggests that a particular state is more likely. This effect also increases the probability of herding on the correct action under the evolutionary rule. The third effect is the following. In a population of agents that follow the evolutionary rule, when agents make inferences from the actions of others, they take into account the contrarianism embedded in their behavior. Even in a population that follows the evolutionary rule there is a positive chance that agents will choose the action suggested by the history of play. A contrarian agent that follows suit provides much information to other agents, since it takes a more revealing private signal to do so. Therefore, their beliefs are updated more under the evolutionary rule than with the expected utility rule when this choice is observed. This effect decreases the probability of herding on the correct action under the evolutionary rule. Since public beliefs are martingales, the second and third effects cancel each other. Therefore, caution is the only remaining force and the evolutionary rule induces herding on the optimal action with higher probability.

The paper is organized as follows. Section 2 discusses some related work. Section 3 studies the model of evolution with private information. Section 4 analyzes a sequential social learning problem. Section 5 concludes. All the proofs are collected in the Appendix, unless noted otherwise.

2 Related literature

There is a long tradition in economics of analyzing the evolutionary fitness of standard behavior assumptions, starting with Alchian (1950) and Friedman (1953) and continuing with Becker (1976) and Hirshleifer (1977).

The population dynamics in our model (section 3.2) are closely related to Robson (1996). He presents a decision-making model where agents have to choose over lotteries. He argues that evolution selects for agents that have different attitudes towards risk, depending of the source of the uncertainty. In particular, he notes that individuals should be more risk averse towards aggregate uncertainty than idiosyncratic uncertainty. The measure of evolutionary success presented in Robson (1996) is the same as the one presented here. The novelty introduced in our paper is to study the implications of this criterion to behavior rules in a particular model of private information.

The literature on evolution of expectations in competitive markets has used also the geometric mean fitness criterion. In particular, Blume and Easley (1992, 2006), Sandroni (2000) show that evolution may select against agents with rational expectations.

Two other applications of the geometric mean fitness in economics are Curry (2001) and Robson and Samuelson (2008). Curry (2001) considers a model were agents choose over lotteries and shows that the adaptive behavior consists in maximizing the expected payoff relative to the population mean. Robson and Samuelson (2008) study a life-cycle model in the presence of aggregate uncertainty. Using the geometric mean fitness criterion, they obtain a present bias in intertemporal preferences.

In our model, the results are driven in part by the fact that decisions are made in groups of agents that follow the same behavior rule[2]. Eliaz, Ray & Razin (2006) also find that decision-making in groups alters individuals' attitudes towards risk and that these 'shifts' are consistent with the Allais paradox.

We show the evolutionarily rule does not involve randomization. This contrasts with the large literature in biology that followed the seminal paper of Cooper and Kaplan (1982). Using a simple model where the only source of uncertainty is aggregate, Cooper and Kaplan (1982) show that in order to achieve diversification at the population level, the optimal behavior rule involves randomization at the individual level. Our results are different in appearance only. In our model, individuals choose a deterministic action almost surely, but the heterogeneity in private information guarantees the appropriate randomization at the population level.

A similar question to the one posed in this paper has been studied in the literature of evolution of preferences in games. When preferences are unobservable, Ok and Vega-Redondo (2001), Ely and Yilankaya (2001) and Dekel, Ely and Yilankaya (2007) show that evolution selects for expected utility maximizers. In this literature, all uncertainty is introduced through some random matching mechanism and the (potential) strategy mixing of the opponent. Since all uncertainty is idiosyncratic, the expected fitness is the appropriate measure of evolutionary performance.

The model of social learning that is considered in this paper is similar to the ones in Bikhchandani et al (1992), Banerjee (1992) and Smith and Sørensen (2000). In these models, agents make decisions sequentially. In contrast, the models of Ellison and Fudenberg (1993, 1995) and Banerjee and Fudenberg (2004) consider large groups of agents that make decisions simultaneously. Our conclusions on the performance of the evolutionary behavior rule in the social learning model are similar to the ones presented in Bernardo & Welch (2001) for overconfident individuals. In our model, however, we consider a particular behavior rule that is optimal in a specific sense, while in their paper different types of overconfident individuals are introduced exogenously.

[2] In biological terms, these agents share the same genes.

3 Evolution with private information

3.1 The model

Time is discrete and infinite. In each period $t = 0, 1, 2, \ldots$ there is only one generation alive. The set of states of the world is $S = \{A, B\}$. In each period, the state $s_t \in S$ is realized independently. The probability of state A being drawn is $p \in (0, 1)$. Agents have to choose an action $x \in \{a, b\}$. The number of offspring left by an agent who chooses x in state s is given by $w(s, x) \in \mathbb{N}$. It is assumed that $w(A, a) > w(A, b)$ and $w(B, b) > w(B, a)$. In other words, action a is optimal if the true state is A while action b is optimal if the true state is B. For simplicity, we will assume $w(s, x) \geq 1$. Since each agent leaves at least one offspring, this assumption rules out the possibility of extinction regardless of the action chosen by the agents. In section 3.2 we discuss how this assumption can be relaxed without affecting the results of the paper based on the findings of Robson (1996). We will sometimes need the following additional assumption:

Assumption 1 (Undominated states) *The offspring function satisfies* $w(A, a) > w(B, a)$ *and* $w(B, b) > w(A, b)$.

Assuming undominated states eliminates the cases where offspring is higher in one of the states irrespective of the action chosen. In other words, it is better to be right than to be wrong in a good environment.

As is customary in many economic models of evolution, we identify offspring with utility. Offspring (or fitness) is the currency of evolution. Thus, we interpret that agents that are successful in an evolutionary sense behave as if they were maximizing offspring. We will use the terms payoff, utility, offspring and fitness interchangeably.

Each agent receives a private signal, $\sigma \in [\underline{\sigma}, \overline{\sigma}] \subset \mathbb{R}$. Private signals in state s are realized independently according to the density function g^s. It is assumed that g^s has full support on $[\underline{\sigma}, \overline{\sigma}]$ and has no atoms. The cumulative distribution is $G^s(\sigma) \equiv \int_{\underline{\sigma}}^{\sigma} g^s(\tau)\, d\tau$. For simplicity, we will assume that the likelihood ratio $g(\sigma) \equiv \frac{g^A(\sigma)}{g^B(\sigma)}$ is continuous and strictly increasing. The informativeness of the signal is bounded: $g(\underline{\sigma}) = \underline{g} > 0$ and $g(\overline{\sigma}) = \overline{g} < \infty$. The strict monotonicity of the likelihood ratio implies that $\underline{g} < 1 < \overline{g}$.

A behavior rule is a function $\xi : [\underline{\sigma}, \overline{\sigma}] \to \Delta\{a, b\}$, where $\xi_x(\sigma)$ is the probability of choosing action x after observing signal σ. Conditional on the state of the world being s, the action chosen by an agent that follows behavior rule ξ is a random variable denoted $X^{\xi}(s)$ with a distribution given by:

$$\Pr\left[X^{\xi}(s)=x\right]=\int_{\underline{\sigma}}^{\overline{\sigma}}\xi_x(\sigma)\,g^s(\sigma)\,d\sigma \tag{1}$$

There are N_t^{ξ} agents in period t who follow behavior rule ξ, with $N_0^{\xi}\geq 1$ given. It is assumed that there are no mutations (i.e. offspring follow the same behavior rule as their parents). Since there are no strategic interactions in our model, there is no loss of generality in studying the performance of each behavior rule independently. Also, the lack of externalities of any kind makes the timing of decisions within a generation irrelevant.

3.2 Population dynamics

In period t, conditional on state s_t, the choices of agents $n=1,\ldots,N_t^{\xi}$ are i.i.d. random variables $X_n^{\xi}(s_t)$, each of them distributed according to (1). The number of agents that follow rule ξ is a stochastic process $\left\{N_t^{\xi}\right\}$with

$$N_{t+1}^{\xi}=\sum_{n=1}^{N_t^{\xi}}w\left(s_t,X_n^{\xi}(s_t)\right) \tag{2}$$

Since the offspring of all agents following rule ξ in the same generation have the same distribution, $\left\{N_t^{\xi}\right\}$ is a branching (or Galton-Watson) process with a random environment. In particular, it is a single-type branching process, as there are no mutations. This class of processes was introduced by Smith and Wilkinson (1969) for i.i.d. environments, as the one considered here.

The population growth rate for rule ξ in period t is denoted by $\nu_t^{\xi}(s_t)$, where

$$\nu_t^{\xi}(s_t)\equiv\frac{N_{t+1}^{\xi}}{N_t^{\xi}}(s_t)=\frac{1}{N_t^{\xi}}\sum_{n=1}^{N_t^{\xi}}w\left(s_t,X_n^{\xi}(s_t)\right) \tag{3}$$

The population growth rate has two sources of randomness. First, there is idiosyncratic risk through $X_n^{\xi}(s_t)$. This source of risk disappears in the long run by the law of large numbers, since $N_t^{\xi}\to\infty$ as $t\to\infty$ almost surely for all ξ. As the population becomes large, the population growth rate conditional on state s converges to $E\left[\nu^{\xi}(s)\right]\equiv E\left[w\left(s,X^{\xi}(s)\right)\mid s\right]$.

Second, there is aggregate risk introduced by s_t. Every period, the reproductive success of all agents depends on the realization of this random variable. Even in the long run, the population growth rate $E\left[\nu^{\xi}(s)\right]$ depends only on s_t. The evolutionarily optimal behavior rule will maximize an average of the state-contingent population growth rates. The main result in this section is that appropriate averaging in this model is geometric.

As it was pointed out before, every rule ξ gives $N_t^{\xi} \to \infty$ a.s.. We define the evolutionarily optimal rule as the one that makes the population size grow *fastest* in the limit. The rationale behind this criterion is that if multiple rules are present, even though there will be an infinite number of agents following each rule, the fraction of agents using the evolutionary rule converges to one.

To formalize this idea, we define the geometric average fitness as follows:

$$W^*(\xi) = \left(E\left[\nu^{\xi}(A)\right]\right)^p \left(E\left[\nu^{\xi}(B)\right]\right)^{1-p}$$

Proposition 1 *Consider the process $\left\{N_t^{\xi}\right\}_{t=0}^{\infty}$ defined by (1) and (2). If $W^*(\xi) > 1$, then $\frac{1}{t}\ln N_t^{\xi} \to \ln W^*(\xi)$ as $t \to \infty$ a.s.*

Proof. See Tanny (1977), Theorem 5.5. ■

Proposition 1 is a consequence of the law of large numbers. We know that $N_t^{\xi} \to \infty$ as $t \to \infty$. Therefore, for a fixed state s, $\nu_t^{\xi}(s) \to E\left[\nu^{\xi}(s)\right]$. Since the state of the world is i.i.d., the evolutionarily successful behavior rule has to maximize the average expected population growth rate. Proposition 1 establishes that the appropriate average measure is the geometric mean. To understand why, consider for a moment the case in which the population growth rate is exactly equal to $E\left[\nu^{\xi}(s)\right]$ in every period. Then,

$$\begin{aligned} N_{t+1}^{\xi} &= N_0^{\xi}\prod_{i=0}^{t} E\left[\nu^{\xi}(s_i)\right] \\ &= N_0^{\xi}\left(E\left[\nu^{\xi}(A)\right]\right)^{k_t(A)}\left(E\left[\nu^{\xi}(B)\right]\right)^{t-k_t(A)} \end{aligned}$$

where $k_t(A) = \sum_{i=0}^{t} 1_{\{s_i = A\}}$. The law of large numbers gives $\frac{1}{t}k_t(A) \to p$ as $t \to \infty$. Therefore, when t is large,

$$N_t^{\xi} \approx N_0^{\xi}\left[W^*(\xi)\right]^t$$

Therefore, $W^*(\xi)$ is the average population growth rate in the long run. In the general case, the state-contingent population growth rate is random in the short run due to idiosyncratic uncertainty. However, since $\nu_t^{\xi}(s) \to E\left[\nu^{\xi}(s)\right]$, our previous argument still applies in the long run.

The geometric average is a consequence of the multiplicative nature of reproduction. Instead of maximizing the population growth rate of any given period, evolution selects the behavior that maximizes the expected compounded growth rate. The behavior rule that dominates in the long run is the one that maximizes $W^*(\xi)$.

In the same vein as in Robson (1996), $W^*(\xi)$ is not an expected utility representation. This is because there are different attitudes towards risk, depending on its source. In particular, conditional on the state of the world, $W^*(\xi)$ evaluates risk through the standard expectation. Across states, however, $W^*(\xi)$ exhibits a higher degree of risk aversion.

In our model, the process $\left\{N_t^\xi\right\}$ is a single-type branching process. In contrast, the population dynamics in Robson (1996) consider the possibility of extinction. If arbitrarily unlikely mutations are allowed, the evolutionarily optimal behavior rule dominates in the long run. Even in the case where the optimal rule becomes temporarily extinct, mutations guarantee that agents that follow the optimal rule eventually are born again. This analysis is based in the more complex theory of multi-type branching processes. Since the population dynamics presented here are the same as in Robson (1996), our no extinction assumption is without loss of generality for obtaining the long run average population growth rate.

3.3 The expected utility criterion

The expected utility criterion is the standard approach in economics and other disciplines. It is assumed that agents update their beliefs after observing their private signals according to Bayes' rule. These beliefs are used to compute the (objective) expected payoff. The expected utility criterion states that agents follow a behavior rule that maximizes this expected utility. We show below that in the present setting, the evolutionarily optimal rule is different from the one induced by the expected utility criterion.

Let $l = \frac{1-p}{p}$ be the prior likelihood ratio. Define the Bayesian posterior probability of A being the true state of the world after observing signal σ as

$$\begin{aligned} q(\sigma) &= \frac{pg^A(\sigma)}{pg^A(\sigma) + (1-p)\, g^B(\sigma)} \\ &= \frac{g(\sigma)}{g(\sigma) + l} \end{aligned} \tag{4}$$

The (ex-post) expected payoff is

$$q(\sigma)\left[\xi_a w(A,a)+\xi_b w(A,b)\right]+\left[1-q(\sigma)\right]\left[\xi_a w(B,a)+\xi_b w(B,b)\right]$$

The (ex-post) expected payoff is related to the unconditional expectation of the population growth rate. The expected population growth rate is

$$\begin{aligned} W^{EU}(\xi) &= E\left[\frac{N_{t+1}^{\xi}}{N_t^{\xi}}\right] \\ &= pE\left[\nu^{\xi}(A)\right]+(1-p)E\left[\nu^{\xi}(B)\right] \\ &= pE\left[w\left(s,X^{\xi}(s)\right)\mid s=A\right]+(1-p)E\left[w\left(s,X^{\xi}(s)\right)\mid s=B\right] \quad (5) \end{aligned}$$

It can be seen from (5) that the expected population growth rate is the ex-ante expected utility. Therefore, maximizing the unconditional expectation of the population growth rate is equivalent to maximizing the (ex-post) expected payoff.

3.4 Level curves

It is immediate that $W^*(\xi)$ is a strictly quasiconcave function of the state-contingent expected population growth rates. In contrast, level curves of $W^{EU}(\xi)$ are linear. Figure 1 shows level curves for $W^*(\xi)$ and $W^{EU}(\xi)$. The slope of a level curve for $W^*(\xi)$ is $-\frac{E\left[\nu^{\xi}(A)\right]}{E\left[\nu^{\xi}(B)\right]}l$ while the slope of a level curve for $W^{EU}(\xi)$ is $-l$. Therefore, level curves of $W^*(\xi)$ are steeper if and only if $E\left[\nu^{\xi}(A)\right]>E\left[\nu^{\xi}(B)\right]$. As we will see in section 3.6, this property drives an important result of this paper: evolution selects behavior rules that smooth the population growth rates across states more than the expected utility rule.

3.5 State-contingent expected payoff frontier

In this section we will construct the frontier of the feasible set of state-contingent expected offspring. In other words, we want to find all the efficient pairs $\left(E\left[\nu^{\xi}(A)\right],E\left[\nu^{\xi}(B)\right]\right)$ that can be obtained from a behavior rule ξ. The characterization of this set is useful in order to find the behavior rules prescribed by the expected utility and evolutionary criteria above.

Sections 3.2 and 3.3 deal with issues that are common to any evolutionary model with aggregate risk. In contrast, this section is concerned with the particular aspects of the private information model at hand. Given our assumptions on the signal

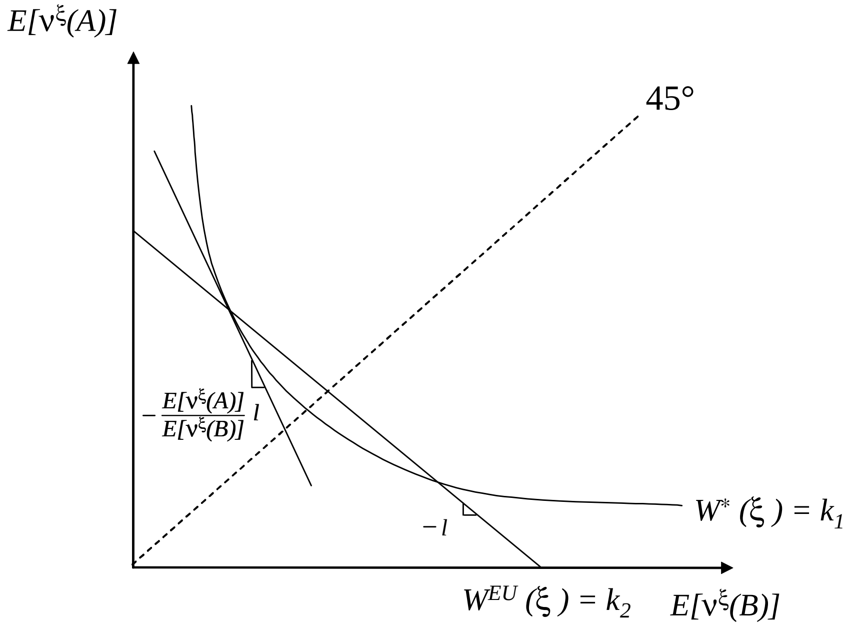

Figure 1: Level curves for $W^*(\xi)$ and $W^{EU}(\xi)$. The former are steeper if and only if $E\left[\nu^{\xi}(A)\right] > E\left[\nu^{\xi}(B)\right]$.

structure, we are able to focus on the class of threshold behavior rules.

Definition 1 *A behavior rule ξ is a* threshold rule *if there exists $\theta \in [\underline{\sigma}, \overline{\sigma}]$ such that:*

$$\xi_a(\sigma) = \begin{cases} 1 & \text{if } \sigma > \theta \\ 0 & \text{if } \sigma < \theta \end{cases} \tag{6}$$

The simple structure of (6) is a consequence of the monotone likelihood ratio property of private signals. We show below that all the efficient pairs of state-contingent payoffs are of this form. Our notion of efficiency is the following:

Definition 2 *The* state-contingent expected payoff frontier *is a set $U \subset [w(A,b), w(A,a)] \times [w(B,a), w(B,b)]$ such that if $u \in U$, then there is no behavior rule ξ such that $u < \left(E\left[\nu^{\xi}(A)\right], E\left[\nu^{\xi}(B)\right]\right)$.*

Lemma 1 *Let U be the state-contingent expected payoff frontier. If $u \in U$, then there exists a unique behavior rule ξ such that $\left(E\left[\nu^{\xi}(A)\right], E\left[\nu^{\xi}(B)\right]\right) = u$. Moreover, ξ a is threshold rule. Conversely, if ξ is a threshold rule, then $\left(E\left[\nu^{\xi}(A)\right], E\left[\nu^{\xi}(B)\right]\right) \in U$.*

Lemma 1 states that the frontier U can be fully characterized by threshold rules. Therefore, we will restrict attention to this class in the remainder of the paper.

Lemma 1 implies that we can characterize the frontier in parametric form by $u(\theta) = \left(E\left[\nu^{\xi}(A)\right], E\left[\nu^{\xi}(B)\right]\right)$, with $\theta \in [\underline{\sigma}, \overline{\sigma}]$ and

$$E\left[\nu^{\xi}(A)\right] = w(A,a) - [w(A,a) - w(A,b)]\, G^{A}(\theta) \quad (7)$$

$$E\left[\nu^{\xi}(B)\right] = w(B,a) + [w(B,b) - w(B,a)]\, G^{B}(\theta) \quad (8)$$

Since the distribution of signals is given by atomless density functions, the frontier is differentiable. The slope of the frontier is given by

$$\frac{dE\left[\nu^{\xi}(A)\right]}{dE\left[\nu^{\xi}(B)\right]} = \frac{\frac{\partial E\left[\nu^{\xi}(A)\right]}{\partial\theta}}{\frac{\partial E\left[\nu^{\xi}(B)\right]}{\partial\theta}} = -Kg(\theta)$$

where $K \equiv \frac{w(A,a)-w(A,b)}{w(B,b)-w(B,a)}$.

Let ξ be a threshold behavior rule with threshold θ and define:

$$Y(\theta) \equiv \frac{E\left[\nu^{\xi}(A)\right]}{E\left[\nu^{\xi}(B)\right]}$$

where $E\left[\nu^{\xi}(A)\right]$ and $E\left[\nu^{\xi}(B)\right]$ are given by (7) and (8) respectively. It is immediate that $Y : [\underline{\sigma}, \overline{\sigma}] \to \mathbb{R}$ is continuous and strictly decreasing. Under assumption 1, we have $Y(\underline{\sigma}) = \frac{w(A,a)}{w(B,a)} > 1$ and $Y(\overline{\sigma}) = \frac{w(A,b)}{w(B,b)} < 1$. Moreover, there exists a unique $\theta_0 \in (\underline{\sigma}, \overline{\sigma})$ such that $Y(\theta_0) = 1$.

Figure 2 depicts the frontier U. When θ travels along $[\underline{\sigma}, \overline{\sigma}]$, $\left(E\left[\nu^{\xi}(A)\right], E\left[\nu^{\xi}(B)\right]\right)$ moves along the frontier from $(w(A,a), w(B,a))$ to $(w(A,b), w(B,b))$[3]. The slope of the frontier is $-Kg(\theta)$. For any $\theta \in [\underline{\sigma}, \overline{\sigma}]$, $Y(\theta)$ gives the slope of the line that connects the origin with the corresponding point in the frontier. Threshold θ_0 gives the intersection of the frontier and the 45° line.

3.6 Evolutionary and expected utility behavior rules: caution and contrarianism

We now characterize the behavior rules that are optimal according to the expected utility and evolutionary criteria. We denote these rules ξ^{EU} and ξ^{*}, respectively. Since the expected utility and evolutionary criteria prescribe maximization of objec-

[3] The first coordinate corresponds to the vertical axis.

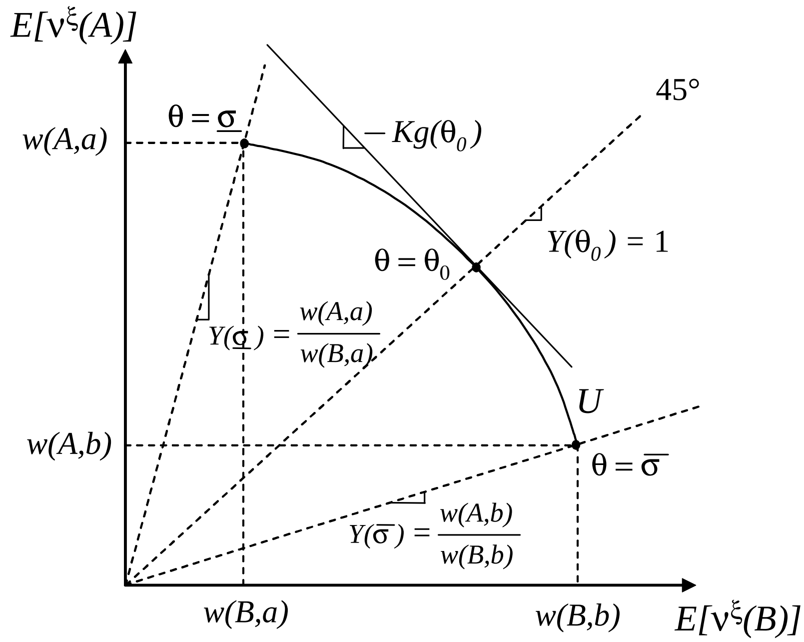

Figure 2: The state-contingent expected payoff frontier U parametrized by θ with undominated states. The slope of the frontier is $-Kg(\theta)$. $Y(\theta)$ gives the slope of the line that connects the origin with the corresponding point in the frontier. Threshold θ_0 gives the intersection of the frontier and the 45° line.

tive functions that are strictly monotonic in state-contingent payoffs, both criteria will choose an element of the frontier U. By Lemma 1, this implies that ξ^{EU} and ξ^* are threshold rules, with thresholds θ^{EU} and θ^*, respectively. Propositions 2 and 3 characterize these thresholds for all possible prior likelihood ratios. Define $\underline{l}^{EU} \equiv K\underline{g}$ and $\bar{l}^{EU} \equiv K\overline{g}$.

Proposition 2 *Let ξ^{EU} be the expected utility behavior rule. Then ξ^{EU} is a threshold rule with threshold $\theta^{EU} : (0, +\infty) \to [\underline{\sigma}, \overline{\sigma}]$, where*

1. *If $\underline{l}^{EU} < l < \bar{l}^{EU}$, then $\theta^{EU}(l)$ is defined implicitly by $Kg\left(\theta^{EU}\right) = l$.*
2. *If $l \leq \underline{l}^{EU}$, then $\theta^{EU}(l) = \underline{\sigma}$. If $l \geq \bar{l}^{EU}$, then $\theta^{EU}(l) = \overline{\sigma}$.*

Analogously, define $\bar{l}^* \equiv \frac{w(B,b)}{w(A,b)} K\overline{g}$ and $\underline{l}^* \equiv \frac{w(B,a)}{w(A,a)} K\underline{g}$.

Proposition 3 *Let ξ^* be the evolutionary behavior rule. Then ξ^* is a threshold rule with threshold $\theta^* : (0, +\infty) \to [\underline{\sigma}, \overline{\sigma}]$, where*

1. *If* $\underline{l}^* < l < \overline{l}^*$*, then* $\theta^*(l)$ *is defined implicitly by* $Kg(\theta^*) = Y(\theta^*)\, l$.

2. *If* $l \leq \underline{l}^*$*, then* $\theta^*(l) = \underline{\sigma}$*. If* $l \geq \overline{l}^*$*, then* $\theta^*(l) = \overline{\sigma}$.

The Implicit Function Theorem implies that threshold functions $\theta^{EU}(l)$ and $\theta^*(l)$ are continuous and strictly increasing on $l \in \left(\underline{l}^{EU}, \overline{l}^{EU}\right)$ and $l \in \left(\underline{l}^*, \overline{l}^*\right)$, respectively.

Propositions 2 and 3 allow for the possibility that agents disregard their private information when one of the states is *a priori* very likely. This is a consequence of the bounded informativeness of private signals. When this happens, actions are perfectly correlated in the population. Even if agents do take into account their private signals, having one of the states *a priori* more likely than the other induces *some* form of correlation. More correlation of actions across the population imply that the state-contingent expected payoffs are further apart.

The difference in the curvature of W^* and W^{EU} gives the main intuition for comparing both rules: the evolutionary criterion is more averse to correlation of actions. This aversion is manifested in the properties of the evolutionary rule. Our next result formalizes the comparison of rules ξ^{EU} and ξ^*.

Proposition 4 *Assume that states are undominated. Then,* $\overline{l}^* > \overline{l}^{EU}$ *and* $\underline{l}^* < \underline{l}^{EU}$*. Additionally, there exists* $\lambda \in \left(\underline{l}^{EU}, \overline{l}^{EU}\right)$ *such that* $\theta^{EU}(l) < \theta^*(l)$ *for all* $l \in (\underline{l}^*, \lambda)$ *and* $\theta^{EU}(l) > \theta^*(l)$ *for all* $l \in \left(\lambda, \overline{l}^*\right)$.

Proposition 4 establishes two main differences between the expected utility rule to the evolutionarily optimal rule. First, the evolutionary rule exhibits *caution*: a more revealing prior than the expected utility criterion is required for agents to disregard their private information.

Second, the evolutionarily optimal rule is *contrarian*. When agents do not disregard their private information, behavior rules ξ^{EU} and ξ^* differ in the threshold private signal for almost every prior. This 'bias' is the mechanism in which the evolutionary criterion introduces smoothing of the state-contingent expected payoffs. Contrarianism operates through the ratio $Y = \frac{E\left[\nu^{\xi}(A)\right]}{E\left[\nu^{\xi}(B)\right]}$. If the expected payoff in state B is larger than the expected payoff in state A (i.e. $Y < 1$), the optimal rule prescribes choosing action a (the optimal action in state A) more often than the expected utility criterion. Figure 3(a) illustrates this point. The consequence of this bias is that the difference between $E\left[\nu^{\xi}(A)\right]$ and $E\left[\nu^{\xi}(B)\right]$ is reduced. Figure 3(b) shows the same mechanism for the case in which $Y > 1$.

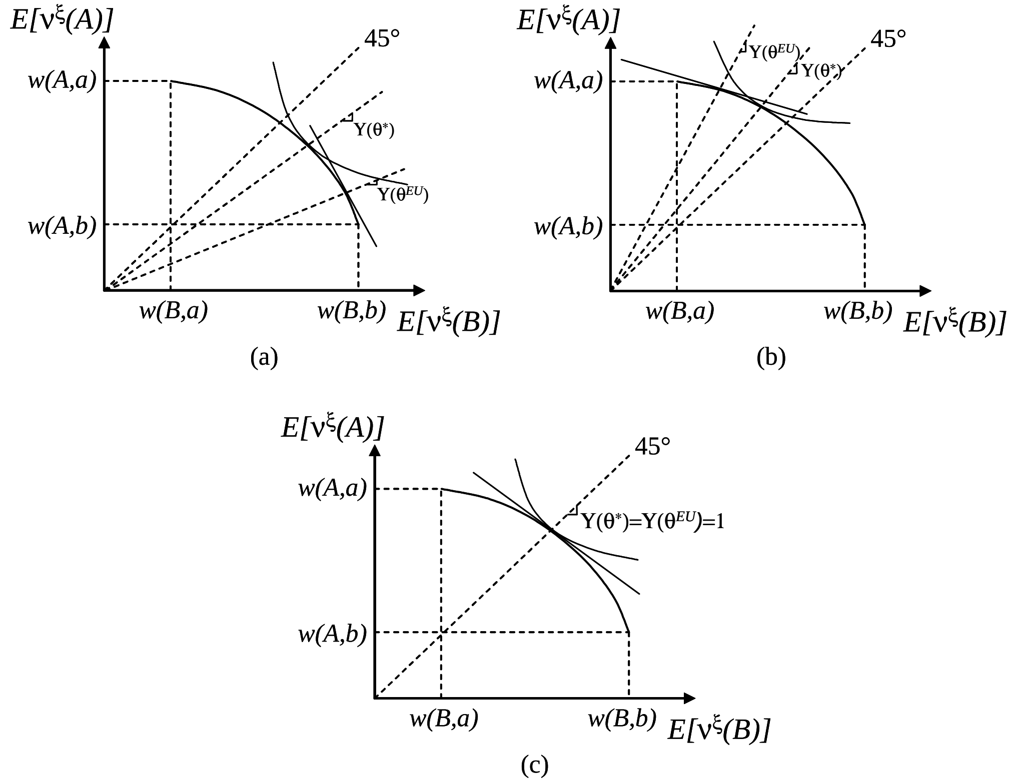

Figure 3: Population growth rate smoothing: the evolutionary rule reduces the differences between $E\left[\nu^{\xi}(A)\right]$ and $E\left[\nu^{\xi}(B)\right]$ compared to the expected utility rule. Panels: (a) $l > \lambda$, (b) $l < \lambda$ and (c) $l = \lambda$ (perfect insurance).

The evolutionarily optimal rule induces a ratio Y closer to 1 than the expected utility rule. In the borderline case when $l = \lambda$, depicted in Figure 3(c), both rules give $E\left[\nu^{\xi}(A)\right] = E\left[\nu^{\xi}(B)\right]$ (i.e. perfect insurance). Notice that $\theta^{EU}(\lambda) = \theta^{*}(\lambda) = \theta_0$, the threshold at which the utility frontier and the 45° line intersect (see Figure 2).

As Proposition 4 shows, if the prior indicates that a particular state is likely the optimal rule has a bias towards actions that are optimal in the *other* state. For instance, if the prior suggests state A is relatively more likely (i.e. $l < \lambda$), then $\theta^{EU} < \theta^{*}$. This implies that the population that follows the optimal rule chooses action b with higher probability. Figure 4 illustrates this point.

If assumption 1 is not verified, then either $Y > 1$ or $Y < 1$ for all thresholds θ. Therefore, only one of the situations depicted in Figure 3 applies (panel (a) or (b), respectively).

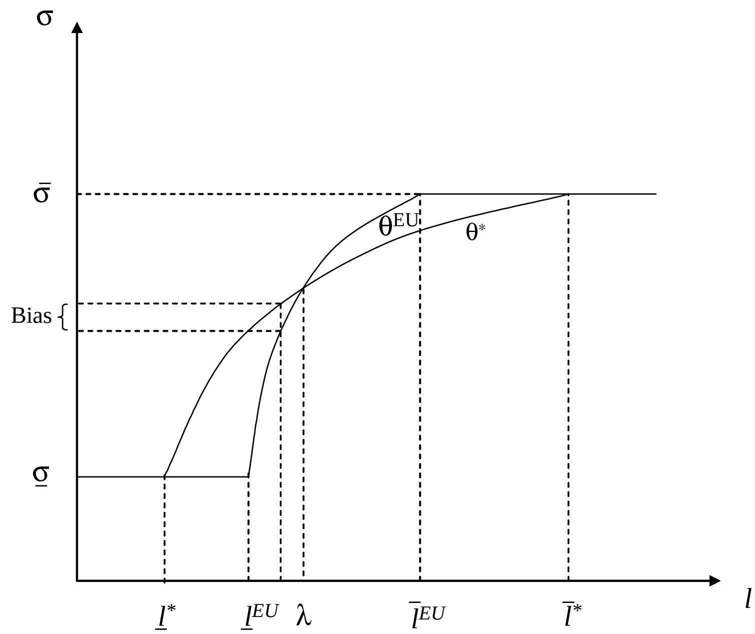

Figure 4: Threshold functions for the evolutionary and expected utility behavior rules. If state A is more likely *a priori* ($l < \lambda$), then the evolutionary rule exhibits a bias towards action b.

Notice that caution and contrarianism are two different aspects of the same phenomenon. While we have distinct interpretations for each of them, it is not possible to have one without the other in our model.

Finally, notice that Proposition 2 implies that the expected utility rule can be fully determined by the likelihood ratios g and l. In contrast, Proposition 3 shows that this is not true for the evolutionary rule. In particular, the evolutionary rule cannot be expressed as a function of the Bayesian posterior of the private signal observed given by (4). This is because the threshold signal θ^* is determined implicitly by

$$Y\left(\theta^*\right) = \frac{w\left(A,a\right) - \left[w\left(A,a\right) - w\left(A,b\right)\right] G^A\left(\theta\right)}{w\left(B,a\right) + \left[w\left(B,b\right) - w\left(B,a\right)\right] G^B\left(\theta\right)}$$

The Bayesian posterior $q\left(\sigma\right)$ as defined in (4) is only a 'local' measure, since it depends on the the density functions $g^A\left(\sigma\right)$ and $g^B\left(\sigma\right)$ only at the signal observed. In contrast, $Y\left(\theta^*\right)$ depends on the cumulative distributions $G^A\left(\sigma\right)$ and $G^B\left(\sigma\right)$. Therefore, in order to compute ξ^* not only we need to know the values of the density function at the signal received, but also how much probability is placed to the left and right of that signal.

3.7 A numerical example

In this section, we present a particular case for our evolutionary model from sections 3.1-3.6. The private signal set is $[\underline{\sigma}, \overline{\sigma}] = [0,1]$. The density functions are $g^A(\sigma) = 1$ and $g^B(\sigma) = \frac{3}{2} - \sigma$. Therefore, the cumulative distribution functions are $G^A(\sigma) = \sigma$ and $G^B(\sigma) = \frac{3\sigma - \sigma^2}{2}$. The private likelihood ratio is $g(\sigma) = \frac{2}{3-2\sigma}$, with $\underline{g} = \frac{2}{3}$ and $\overline{g} = 2$.

Payoffs are given by $w(A,a) = w(B,b) = 3$ and $w(A,b) = w(B,a) = 2$. These payoffs satisfy Assumption 1. Thus we obtain, $\underline{l}^{EU} = \frac{2}{3}, \bar{l}^{EU} = 2, \underline{l}^* = \frac{4}{9}$ and $\bar{l}^* = 3$, which satisfy $\bar{l}^* > \bar{l}^{EU}$ and $\underline{l}^* < \underline{l}^{EU}$ as it was established in Proposition 4.

We can compute the private signal threshold functions for the expected utility and evolutionary behavior rules for these values. We obtain

$$\theta^{EU}(l) = \frac{3}{2} - \frac{1}{l} \tag{9}$$

$$\theta^*(l) = \frac{3 + 9l - \left(25 + 50l + 9l^2\right)^{\frac{1}{2}}}{2 + 4l} \tag{10}$$

The threshold functions (9) and (10) satisfy $\theta^{EU}\left(\underline{l}^{EU}\right) = \underline{\sigma}$, $\theta^{EU}\left(\bar{l}^{EU}\right) = \overline{\sigma}$, $\theta^*\left(\underline{l}^*\right) = \underline{\sigma}$ and $\theta^*\left(\bar{l}^*\right) = \overline{\sigma}$ as established in Propositions 2 and 3. Additionally, it can be checked that they are continuous and strictly increasing functions

Using the cumulative distribution functions, we get $Y(\theta) = \frac{6-2\theta}{4+3\theta-\theta^2}$. By Proposition 4, we know that there is a unique λ such that $\theta^{EU}(\lambda) = \theta^*(\lambda) = \theta_0 \in (\underline{\sigma}, \overline{\sigma})$ and therefore $Y(\theta_0) = 1$. Solving gives $\theta_0 = \frac{5-\sqrt{17}}{2} \approx .438$. Using either (9) or (10) we obtain $\lambda = \frac{2}{\sqrt{17}-2} \approx .942$. Figure 5 shows the private signal thresholds $\theta^{EU}(l)$ and $\theta^*(l)$.

In line with Proposition 4, we obtain $\theta^{EU}(l) < \theta^*(l)$ if $l < \frac{2}{\sqrt{17}-2}$ and $\theta^{EU}(l) > \theta^*(l)$ if $l > \frac{2}{\sqrt{17}-2}$.

3.8 Public signals

Our previous analysis assumes that the underlying properties of the environment are summarized by a common prior p that is also invariant over time. We also implicitly assumed that all public information was summarized in the prior. The latter might look surprising, since we have shown that the evolutionarily optimal behavior rule cannot be summarized by a private belief when agents have private information. In

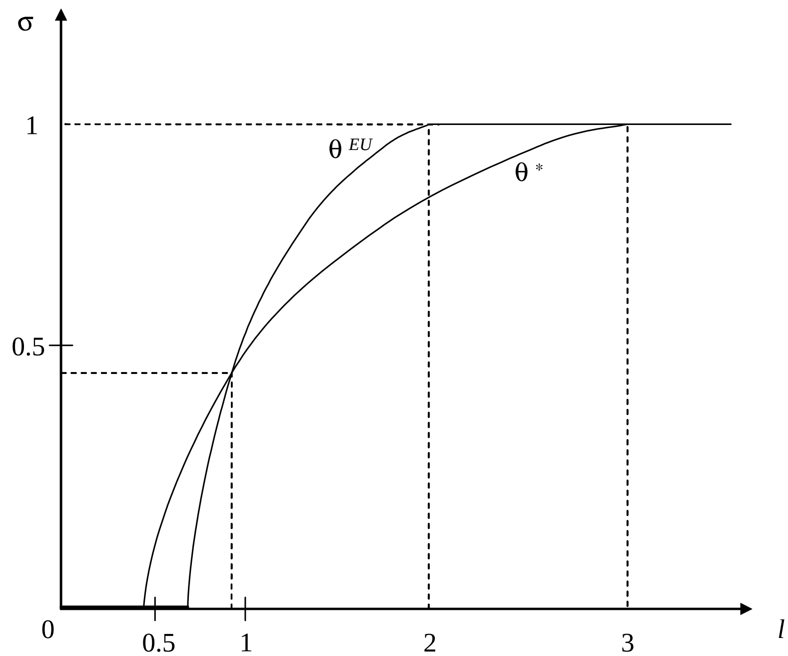

Figure 5: Numerical example for private signal thresholds for ξ^* and ξ^{EU}. Parameter values: $[\underline{\sigma}, \overline{\sigma}] = [0, 1]$, $g^A(\sigma) = 1$, $g^B(\sigma) = \frac{3}{2} - \sigma$, $w(A, a) = w(B, b) = 3$ and $w(A, b) = w(B, a) = 2$.

this section we show that relaxing these simplifying assumptions does not affect our results.

We modify our model in the following sense. In addition to private information, all agents receive a public signal ω correlated with the state of the world. The key assumption that we need to maintain is that public and private signals are independent, conditional on the state of the world.

As before, the prior probability of state A in any given period is p. In each period all agents receive (the same) public signal, $\omega \in \Omega$ where Ω is an arbitrary set. The probability of signal ω being drawn in state s is $f^s(\omega)$. Behavior rules are functions $\xi : \Omega \times [\underline{\sigma}, \overline{\sigma}] \to \Delta\{a, b\}$, where $\xi_x(\omega, \sigma)$ is the probability of choosing action x after observing signals ω and σ.

The main observation is that in this context, the pair (s, ω) is the effective 'state', even though ω is irrelevant in terms of payoffs. Define the random variable $\widetilde{X}^{\xi}(s, \omega)$, where

$$\Pr\left[\widetilde{X}^{\xi}(s,\omega)=x\right]=\int_{\underline{\sigma}}^{\overline{\sigma}}\xi_x(\omega,\sigma)\,g^s(\sigma)\,d\sigma \tag{11}$$

The population growth rate is given by

$$\widetilde{\nu}_t^{\xi}(s_t,\omega_t)=\frac{1}{N_t^{\xi}}\sum_{n=1}^{N_t^{\xi}} w\left(s_t, X_n^{\xi}(s_t,\omega_t)\right) \tag{12}$$

Define the geometric average fitness:

$$\widetilde{W}(\xi)=\prod_{\omega\in\Omega}\left[E\left[\widetilde{\nu}^{\xi}(A,\omega)\right]\right]^{pf^A(\omega)}\left[E\left[\widetilde{\nu}^{\xi}(B,\omega)\right]\right]^{(1-p)f^B(\omega)}$$

Once the 'state' is appropriately redefined as (s,ω), the limit result is a simple extension of Proposition 1.

Corollary 1 *Consider the process* $\left\{N_t^{\xi}\right\}_{t=0}^{\infty}$ *defined by (11) and (12). If* $\widetilde{W}(\xi)>1$, *then* $\frac{1}{T}\ln N_T^{\xi}\to\ln\widetilde{W}^P(\xi)$ *as* $T\to\infty$ *a.s.*

Define the Bayesian posterior (of state A being true) after observing signal ω:

$$r(\omega)=\frac{pf^A(\omega)}{pf^A(\omega)+(1-p)\,f^B(\omega)}$$

Let $\widetilde{\xi}\in\arg\max\widetilde{W}(\xi)$ and fix ω. Since $\widetilde{W}(\xi)$ is a product of independent factors,

$$\begin{aligned}\widetilde{\xi}|_{\omega} &\in \arg\max\left[E\left[\widetilde{\nu}^{\xi}(A,\omega)\right]\right]^{pf^A(\omega)}\left[E\left[\widetilde{\nu}^{\xi}(B,\omega)\right]\right]^{(1-p)f^B(\omega)}\\ &= \arg\max\left[E\left[\widetilde{\nu}^{\xi}(A,\omega)\right]\right]^{r(\omega)}\left[E\left[\widetilde{\nu}^{\xi}(B,\omega)\right]\right]^{(1-r(\omega))}\end{aligned}$$

Therefore, $\widetilde{\xi}(\omega,\sigma)$ is equivalent to rule $\xi^*(\sigma)$ if the prior is adjusted to account for the public signal. In sum, the Bayes' posterior of public signal is a sufficient statistic and our previous approach was without loss of generality.

Finally, the extension presented in this section is not the most general version of the model that we could think of. In fact, the results in Tanny (1977) establish that the convergence rate for single-type branching processes is the geometric mean across states when s follows any stationary and ergodic process.

4 Sequential social learning

This section presents a sequential social learning model, similar to the ones introduced by Banerjee (1992) and Bikhchandani et al (1992). The setting considered here follows Smith and Sørensen (2000) closely. The literature focuses on the expected utility rule. We show here that many well-known results can be extended to subclass of threshold rules that satisfy some additional regularity conditions. This subclass includes the expected utility and the evolutionary behavior rules from section 3.6. This result will allow us to easily compare the performance of rules within that class.

Section 3.1 considers a model where there are no externalities of any kind. In contrast, informational externalities are present in social learning models. While each agent maximizes its own payoff based on its private information, they also observe the decisions made and make inferences about the signals received by others. Therefore, an agent's action affects the information received by others.

We have shown in section 3.6 that the evolutionary rule is optimal in an environment without externalities of any kind. However, there is no reason to expect it to be optimal in a model of social learning where informational externalities are present. We show, nevertheless, that the evolutionary rule outperforms the expected utility rule.

We will not provide a characterization the (constrained) efficient behavior rule according to the evolutionary criterion presented in section 3.2. Smith and Sørensen (2006) address this issue using the expected utility criterion and find that the (constrained) efficient behavior rule also exhibits contrarianism. Further connections between their approach and ours are left for future research.

While section 3 was concerned with a infinite generation problem, in this section we will consider only one generation in which the state is realized once.

4.1 The model

The state of the world s is drawn from $\{A, B\}$, each with equal probability. There is an infinite sequence of individuals ordered exogenously. An agent in stage $n = 0, 1, \ldots$ receives an independent signal $\sigma_n \in [\underline{\sigma}, \overline{\sigma}]$ according to the density functions $g^s(\sigma)$. Agents choose an action $x \in \{a, b\}$ after observing their own signal and the actions of the agents that preceded them. Past actions in stage n are an element of $H^n = \{a, b\}^{n-1}$. The set of histories is $H = \bigcup_{n=0}^{\infty} H^n$, with $H^0 = \{\varnothing\}$ and typical element h.

The history of past play is public information and can be interpreted as a public signal. The Bayesian posterior after observing history h. This posterior will serve as the prior distribution when we apply the behavior rules from section 3.6. The argument presented in section 3.8 guarantees that we can focus on this posterior without loss of generality.

We restrict attention to monomorphic populations playing a fixed behavior rule ξ. An agent in position n follows the behavior prescribed by $\xi(\sigma_n)$. Moreover, this fact is common knowledge among agents. This assumption is needed for agents to make inferences about previous agents' private signals after observing histories.

The decision of an agent in state s is a random variable $X^{\xi}(s)$, where

$$\Pr\left[X^{\xi}(s) = x\right] = \int_{\underline{\sigma}}^{\overline{\sigma}} \xi_x(\sigma) g^s(\sigma) d\sigma$$

Define

$$\varphi^{\xi}(x, l) \equiv \frac{\Pr\left[X^{\xi}(B) = x\right]}{\Pr\left[X^{\xi}(A) = x\right]} \tag{13}$$

where l is the public likelihood ratio used to compute behavior rule ξ. Bayes' rule implies that, if action x is observed and $\Pr\left[X^{\xi}(s) = x\right] \in (0, 1)$, the public likelihood ratio is updated according to

$$l^{\xi}_{n+1}\left(x, l^{\xi}_n\right) = \varphi^{\xi}\left(x, l^{\xi}_n\right) l^{\xi}_n$$

If $\Pr\left[X^{\xi}(s) = x\right] \in \{0, 1\}$, then $l^{\xi}_{n+1}(x, l_n) = l_n$.

Notice that conditional on A, $\left\{l^{\xi}_n\right\}$ is a martingale, since when $\Pr\left[X^{\xi}(s) = x\right] \in (0, 1)$ we have

$$\begin{aligned} E\left[l^{\xi}_{n+1} \mid l^{\xi}_n\right] &= \Pr\left[X^{\xi}(A) = a\right] \varphi^{\xi}\left(a, l^{\xi}_n\right) l^{\xi}_n + \Pr\left[X^{\xi}(A) = b\right] \varphi^{\xi}\left(b, l^{\xi}_n\right) l^{\xi}_n \\ &= \Pr\left[X^{\xi}(B) = a\right] l^{\xi}_n + \Pr\left[X^{\xi}(B) = b\right] l^{\xi}_n = l^{\xi}_n \end{aligned}$$

By the same reasoning, $\left\{\left(l^{\xi}_n\right)^{-1}\right\}$ is a martingale conditional on state B. This property will drive many of the results below.

4.2 CM-threshold behavior rules

Even though the evolutionarily optimal rule and the expected utility rule are different in many respects, they share a common structure. In particular, section 3.6

shows that these rules are threshold rules. In this section, we show that if threshold rules satisfy some regularity conditions, then the outcome of the social learning model exhibits the same qualitative properties that were established in the literature for expected utility rules.

We will focus on a subset of the threshold rules from Definition 1.

Definition 3 *A behavior rule ξ is a* CM-threshold rule *if there are some public likelihood ratio values $\underline{l} \in (0,1)$ and $\bar{l} \in (1,+\infty)$ and a private signal threshold function $\theta : (0,+\infty) \to [\underline{\sigma}, \overline{\sigma}]$ such that:*

$$\xi_a(\sigma) = \begin{cases} 1 & \text{if } \sigma > \theta(l) \\ 0 & \text{if } \sigma < \theta(l) \end{cases}$$

where $\theta(l) = \underline{\sigma}$ for all $l \leq \underline{l}$ and $\theta(l) = \overline{\sigma}$ for all $l \geq \bar{l}$. Additionally, $\theta(l)$ is continuous and strictly increasing on $\left(\underline{l}, \bar{l}\right)$.

A CM-threshold rule requires that the threshold is a continuous and strictly monotonic function of the public likelihood ratio. The expected utility and evolutionary rules are also CM-threshold rules. In this section we will extend many of the results from Smith and Sørensen (2000) for this class of behavior rules.

Let ξ be a CM-threshold behavior rule with thresholds $\underline{l}$, $\bar{l}$ and $\theta(l)$. By definition, $\Pr\left[X^{\xi}(s) = x\right] \in (0,1)$ if and only if $l_n \in \left(\underline{l}, \bar{l}\right)$. Since $l_0 = 1$, we have that $\left\{l_n^{\xi}\right\}$ is a stochastic process with the transition law defined by:

$$l_{n+1}^{\xi}\left(x, l_n^{\xi}\right) = \begin{cases} \varphi^{\xi}\left(x, l_n^{\xi}\right) l_n^{\xi} & \text{if } l_n^{\xi} \in \left(\underline{l}, \bar{l}\right) \\ l_n^{\xi} & \text{if } l_n^{\xi} \notin \left(\underline{l}, \bar{l}\right) \end{cases} \tag{14}$$

If $l_n^{\xi} \notin \left(\underline{l}, \bar{l}\right)$ for some n, then behavior rule ξ prescribes that agent n should disregard its private information and base its decision only on the public likelihood ratio. Consider now the agent in position $n+1$. There is no information to infer from agent n's choice, so the public likelihood ratio is not updated. By induction, we can conclude that if $l_n^{\xi} \notin \left(\underline{l}, \bar{l}\right)$, then $l_m^{\xi} = l_n^{\xi}$ for every $m > n$. Notice also that this implies that actions must also coincide (i.e. $\rho\left(x \mid \xi, s, l_m^{\xi}\right) = \rho\left(x \mid \xi, s, l_n^{\xi}\right)$ for every $m > n$). When $l_n^{\xi} \notin \left(\underline{l}, \bar{l}\right)$ for some n, we say that the process $\left\{l_n^{\xi}\right\}$ has entered a *cascade*. Additionally, it is possible that the public likelihood ratio never enters a cascade, but approaches it in the limit. In this case, CM-threshold rules prescribe that the probability of one of the actions approaches unity. We call the latter phenomenon a *herd*[4]. Cascades imply herds, but the converse is not true.

If ξ is a threshold rule, then $\Pr\left[X^{\xi}(s)=b\right] = 1 - \Pr\left[X^{\xi}(s)=a\right] = G^{s}(\theta)$. Therefore, if $l \in \left(\underline{l}, \bar{l}\right)$, (13) becomes:

$$\varphi^{\xi}(x,l) = \begin{cases} \frac{1-G^{B}(\theta(l))}{1-G^{A}(\theta(l))} & \text{if } x = a \\ \frac{G^{B}(\theta(l))}{G^{A}(\theta(l))} & \text{if } x = b \end{cases} \tag{15}$$

Proposition 5 extends some facts about the continuation functions φ that are well known in the social learning literature for expected utility maximizers to CM-threshold behavior rules.

Proposition 5 *Let ξ be a CM-threshold behavior rule. Then, for all $l \in \left(\underline{l}, \bar{l}\right)$,*

1. $\varphi^{\xi}(a,l) < 1 < \varphi^{\xi}(b,l)$.

2. $\varphi^{\xi}(x,l)$ *is continuous and strictly decreasing in* l.

3. $\lim_{l\to\underline{l}} \varphi^{\xi}(a,l) = \lim_{l\to\bar{l}} \varphi^{\xi}(b,l) = 1$.

The monotone private likelihood ratio implies that the cumulative distribution function G^{A} first order stochastically dominates G^{B}. This gives part 1 of Proposition 5. This condition guarantees $l^{\xi}_{n+1}\left(a, l^{\xi}_{n}\right) < l^{\xi}_{n} < l^{\xi}_{n+1}\left(b, l^{\xi}_{n}\right)$. In other words, barring a cascade, observing a makes the public likelihood ratio in the subsequent stage indicate that state A is more likely. Conversely,.observing b makes the public likelihood ratio in the subsequent stage strictly larger, an consequently, public information indicates that state B more likely. The monotone likelihood ratio of private signals is not necessary for obtaining this result, although it simplifies the argument.

The second part of Proposition 5 states that less likely actions are more informative. For CM-threshold rules, as l increases, action b becomes more likely and, consequently, action a less likely. Proposition 5 part 2 says that as l increases, $\varphi^{\xi}(b,l)$ becomes closer to 1 while $\varphi^{\xi}(a,l)$ moves away from 1. As l increases, action b is more frequent and therefore provides less information, so the public likelihood ratio is updated by a smaller amount. On the other hand, action a becomes more rare. If action a is observed, the public likelihood ratio is updated by a large amount.

Part 3 of Proposition 5 establishes that $l^{\xi}_{n+1}(a,l)$ is continuous at $\underline{l}$ and that $l^{\xi}_{n+1}(b,l)$ is continuous at $\bar{l}$. This is a consequence of having a continuum of private signals together with a continuous threshold function $\theta(l)$. As the public likelihood

[4] This distinction was introduced by Smith and Sørensen (2000).

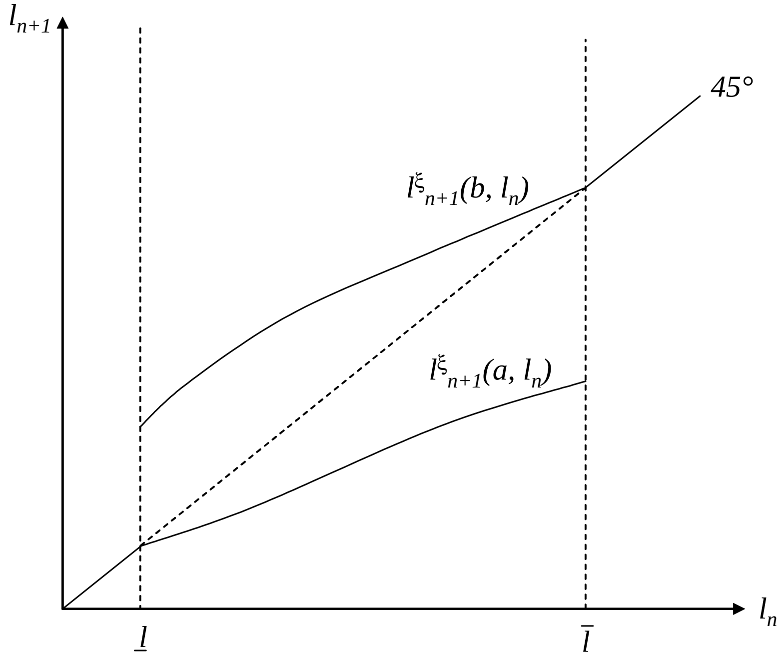

Figure 6: Public likelihood dynamics where $l_{n+1}^{\xi}\left(a, l_n^{\xi}\right) < l_n^{\xi} < l_{n+1}^{\xi}\left(b, l_n^{\xi}\right)$ (implied by Proposition 5 part 1), $l_{n+1}^{\xi}(a, l) \to \underline{l}$ as $l \to \underline{l}$ and $l_{n+1}^{\xi}(b, l) \to \bar{l}$ as $l \to \bar{l}$ (implied by Proposition 5 part 3) and $l_{n+1}^{\xi}\left(x, l_n^{\xi}\right)$ is increasing in l_n^{ξ}.

approaches say, $\underline{l}$, the probability of observing a signal that would induce choosing action b approaches zero. In earlier papers from this literature, where discrete private signals were considered, this probability is bounded away from zero. If $l_{n+1}^{\xi}(a, l)$ is increasing in l, this continuity result precludes cascade formation (see Proposition 8).

Figure 6 shows the basic properties implied in Proposition 5. First, $l_{n+1}^{\xi}\left(a, l_n^{\xi}\right) < l_n^{\xi} < l_{n+1}^{\xi}\left(b, l_n^{\xi}\right)$ (Proposition 5 part 1). Second, $l_{n+1}^{\xi}(a, l) \to \underline{l}$ as $l \to \underline{l}$ and $l_{n+1}^{\xi}(b, l) \to \bar{l}$ as $l \to \bar{l}$(Proposition 5 part 3). Finally, since $l_{n+1}^{\xi}\left(x, l_n^{\xi}\right) = \varphi^{\xi}\left(x, l_n^{\xi}\right) l_n^{\xi}$ if $l_n^{\xi} \in \left(\underline{l}, \bar{l}\right)$, Proposition 5 part 2 allows for the possibility that $l_{n+1}^{\xi}\left(x, l_n^{\xi}\right)$ is decreasing at some l_n^{ξ}. Figure 6 depicts the case in which the effect of Proposition 5 part 2 is dominated and $l_{n+1}^{\xi}\left(x, l_n^{\xi}\right)$ is increasing in l_n^{ξ}. We will focus on the case in which $l_{n+1}^{\xi}\left(x, l_n^{\xi}\right)$ is increasing in l_n^{ξ}, since many private signal distribution functions satisfy this property.

Denote by $\rho^{\xi}(s)$ the probability of herding on the optimal action in state s. Then, we have:

$$\begin{aligned}\rho^{\xi}(A) &= \Pr\left[l_{\infty}^{\xi}(A) \in (0, \underline{l}] \mid A\right] \\ \rho^{\xi}(B) &= \Pr\left[l_{\infty}^{\xi}(B) \in \left[\bar{l}, \infty\right) \mid B\right]\end{aligned}$$

Our next result is well known in the literature. It establishes that public beliefs converge to a cascade with probability one (although they might never enter the cascade). We present a version of this result close to the one in Smith and Sørensen (2000), again extending the arguments to the class for CM-threshold behavior rules.

Proposition 6 *Let ξ be a CM-threshold behavior rule. Then,*

1. *Conditional on state s, $l_n^{\xi} \to l_{\infty}^{\xi}(s)$ almost surely as $n \to \infty$, where $l_{\infty}^{\xi}(s)$ is a random variable and $l_{\infty}^{\xi}(s) \notin \left(\underline{l}, \bar{l}\right)$*
2. *$\rho^{\xi}(s) < 1$*
3. *If $l_{n+1}^{\xi}\left(x, l_n^{\xi}\right)$ is increasing in l_n^{ξ}, then $l_n^{\xi}(s) \in \left(\underline{l}, \bar{l}\right)$ for all n and $l_{\infty}^{\xi}(s) \in \left\{\underline{l}, \bar{l}\right\}$*

Proposition 6 exploits the Markov-martingale nature of the stochastic process $\left\{l_n^{\xi}\right\}$. Part 1 establishes that beliefs converge to a cascade with probability one. In other words, learning eventually stops. The history of past play becomes so informative that agents start disregarding their private information almost surely. The Martingale Convergence Theorem guarantees that $l_n^{\xi} \to l_{\infty}^{\xi}$ almost surely, since conditional on s, either $\left\{l_n^{\xi}\right\}$ or $\left\{\left(l_n^{\xi}\right)^{-1}\right\}$ is a martingale. The Markov property implies that every l in the support of l_{∞}^{ξ} has to be pointwise stationary. Finally, Proposition 5 implies that l is pointwise stationary if and only if $l \in (0, \underline{l}] \cup \left[\bar{l}, +\infty\right)$.

Part 2 establishes that there is a positive probability of herding on the inefficient action. Since $\{l_n\}$ is bounded, $E\left[l_{\infty}^{\xi} \mid l_0, A\right] = l_0 = 1 \in \left(\underline{l}, \bar{l}\right)$. If $l_{\infty}^{\xi}(s)$ puts all the mass either on $(0, \underline{l}]$ or $\left[\bar{l}, +\infty\right)$, then we could not have $E\left[l_{\infty}^{\xi} \mid l_0, A\right] \in \left(\underline{l}, \bar{l}\right)$.

Part 3 establishes that even though public beliefs converge to a cascade, they may never enter one. That is, it is possible that $l_n^{\xi} \to l_{\infty}^{\xi} \in \left(\underline{l}, \bar{l}\right)$ with $l_n^{\xi} \in \left(\underline{l}, \bar{l}\right)$ for all n. Notice that the continuity of $\theta(l)$ implies that if $l_n^{\xi} \to l_{\infty} \notin \left(\underline{l}, \bar{l}\right)$, then $\Pr\left[X^{\xi}(s) = x\right] \to \{0, 1\}$. Therefore, action convergence (herding) eventually arises. When $l_{n+1}^{\xi}\left(x, l_n^{\xi}\right)$ is increasing in l_n^{ξ}, the only elements of the cascade sets that are accessible from $\left(\underline{l}, \bar{l}\right)$ are $\left\{\underline{l}, \bar{l}\right\}$. Therefore, cascades cannot occur. Herding, on the other hand, occurs with probability one.

4.3 Evolutionary versus expected utility behavior rules

As it was established in section 3.6, the evolutionary rule introduces smoothing across states through two different mechanisms. First, caution implies the cascade sets are smaller. Second, when the public likelihood ratio gets closer to a cascade set, the evolutionary rule has a bias towards the opposite action. These properties are established in Proposition 4. Together with Proposition 5, we obtain the following result as a corollary:

Proposition 7 *There exists* $\lambda \in \left(\underline{l}^{EU}, \bar{l}^{EU}\right)$ *such that* $l_{n+1}^{*}(a,l) < l_{n+1}^{EU}(a,l)$ *for all* $l < \lambda$ *and* $l_{n+1}^{*}(b,l) > l_{n+1}^{EU}(b,l)$ *for all* $l > \lambda$.

Figure 7 illustrates the result from Proposition 7. If $l < \lambda$, then the past history indicates that state A is relatively more likely. Contrarian behavior implies that the evolutionary rule exhibits a bias towards action b. Therefore, action a is more rare in a population of agents that follow the evolutionary rule than the expected utility rule. Proposition 5 part 2 implies that if action a is observed, the public likelihood ratio adjusts more under the expected utility rule than under the evolutionary rule. Symmetrically, if $l > \lambda$ is observed, action b is more frequent in a population of expected utility maximizers. If action b is observed, the public likelihood adjusts more in a population that follows the evolutionary rule.

We are interested in comparing the performance of both rules in the social learning model. Since herding occurs almost surely, we compare the probability of herding on the optimal action. Since prior probabilities assign equal probability to each state, this probability is given by

$$\gamma(\xi) = \frac{1}{2}\rho^{\xi}(A) + \frac{1}{2}\rho^{\xi}(B)$$

Proposition 8 *Let* $l_{n+1}^{\xi}\left(x, l_n^{\xi}\right)$ *be increasing in* l_n^{ξ} *for* $\xi \in \left\{\xi^{*}, \xi^{EU}\right\}$. *Then,* $\gamma(\xi^{*}) > \gamma\left(\xi^{EU}\right)$.

Proposition 8 establishes that, under some technical assumptions, the evolutionary rule induces herding on the correct action with a higher probability than the expected utility rule. When making this comparison, there are three different effects that have to be taken into account. First, caution requires more information from the history of play in order to enter a cascade. Since private signals are correlated with the true state of the world, caution increases the probability that this information points in the correct direction.

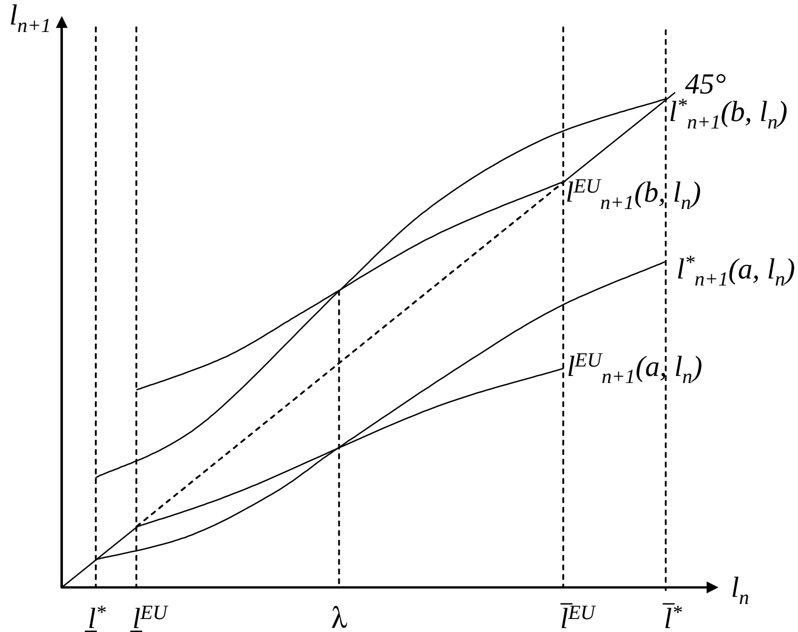

Figure 7: Public likelihood ratio dynamics for ξ^* and ξ^{EU}. Caution implies $\bar{l}^* > \bar{l}^{EU}$ and $\underline{l}^* < \underline{l}^{EU}$. Contrarian behavior implies $l^*_{n+1}(a, l) < l^{EU}_{n+1}(a, l)$ for all $l < \lambda$ and $l^*_{n+1}(b, l) > l^{EU}_{n+1}(b, l)$ for all $l > \lambda$.

Second, contrarian behavior makes less likely for individuals to follow suit when the public likelihood ratio approaches the cascade set. This effect also increases the probability of herding on the correct action under the evolutionary rule.

Third, when agents make inferences from the actions of others, they take into account the contrarianism embedded in the behavior rule. Even in a population that follows the evolutionary rule there is a positive chance that agents will choose the action suggested by the history of play. If this happens, Proposition 5 part 2 implies that beliefs are updated more under the evolutionary rule than with the expected utility rule. This effect decreases the probability of herding on the correct action under the evolutionary rule.

Proposition 8 can establish the desired result because the martingale property of public beliefs make the second and third effect to perfectly cancel each other. Therefore, caution is the only remaining force that drives our result.

Proposition 8 does not imply that the probability of herding on the optimal action is higher for the evolutionarily rule in every state. Nevertheless, it is higher when taking the average probability across states. Also, this implies that while the probability may be lower in some state, it cannot be lower in both states.

We are also interested in comparing the long run population growth rates if populations were to face this sequential social learning problem in every generation. Under some assumptions on the payoff structure, the result from Proposition 8 implies that the our evolutionary rule has a higher average long run population growth rate than the expected utility rule.

Following Banerjee (1992), from an ex-ante point of view if all agents have the same probability of being the nth decision-maker for $n = 0, 1, \ldots$, the appropriate measure of 'welfare' is the average payoff. Since herds eventually arise, it is natural to compare the evolutionary performance in the limit. As in our analysis in section 3.8, there are four effective states, determined by the two underlying states of the world and herding on two possible actions. The average long run population growth rate for ξ is given by

$$W^{\infty}(\xi) = [w(A,a)]^{\frac{\rho^{\xi}(A)}{2}} [w(A,b)]^{\frac{1-\rho^{\xi}(A)}{2}} [w(B,b)]^{\frac{\rho^{\xi}(B)}{2}} [w(B,a)]^{\frac{1-\rho^{\xi}(B)}{2}} \tag{16}$$

If the payoff function $w(s,x)$ is symmetric (i.e. $w(A,a) = w(B,b)$ and $w(A,b) = w(B,a)$), then $W^{\infty}(\xi^*) > W^{\infty}(\xi^{EU})$ if and only if $\gamma(\xi^*) > \gamma(\xi^{EU})$. Therefore, Proposition 8 implies that the average long run population growth rate for ξ^* is higher than for ξ^{EU}.

4.4 A numerical example

In this section, we continue the numerical example from section 3.7. Recall that the model's parameters are given by $[\underline{\sigma}, \overline{\sigma}] = [0, 1]$, $g^A(\sigma) = 1$ and $g^B(\sigma) = \frac{3}{2} - \sigma$. The cumulative distribution functions give:

$$\begin{aligned} \frac{1 - G^B(\sigma)}{1 - G^A(\sigma)} &= 1 - \frac{\sigma}{2} \\ \frac{G^B(\sigma)}{G^A(\sigma)} &= \frac{3}{2} - \frac{\sigma}{2} \end{aligned}$$

Offspring is given by $w(A,a) = w(B,b) = 3$ and $w(A,b) = w(B,a) = 2$.

Therefore, since $\underline{l}^{EU} = \frac{2}{3}, \overline{l}^{EU} = 2, \underline{l}^* = \frac{4}{9}$ and $\overline{l}^* = 3$, we obtain:

$$\begin{aligned}
\varphi^{EU}(a,l) &= \frac{1}{4} + \frac{1}{2l} \\
\varphi^{EU}(b,l) &= \frac{3}{4} + \frac{1}{2l} \\
\varphi^*(a,l) &= \frac{1 - l + \left(25 + 50l + 9l^2\right)^{\frac{1}{2}}}{4 + 8l} \\
\varphi^*(b,l) &= \frac{3(1+l) + \left(25 + 50l + 9l^2\right)^{\frac{1}{2}}}{4 + 8l}
\end{aligned}$$

Consistently with Proposition 5 part 1 , $\varphi^{EU}(a,l) < 1 < \varphi^{EU}(b,l)$ for all $l \in \left(\frac{2}{3}, 2\right)$ and $\varphi^*(a,l) < 1 < \varphi^*(b,l)$ for all $l \in \left(\frac{4}{9}, 3\right)$. Clearly, $\varphi^{EU}(a,l)$ and $\varphi^{EU}(b,l)$ are decreasing in l. The same can be checked for $\varphi^*(a,l)$ and $\varphi^*(b,l)$ (Proposition 5 part 2). Finally, $\lim_{l \to \frac{2}{3}} \varphi^{EU}(a,l) = \lim_{l \to \frac{4}{9}} \varphi^*(a,l) = \lim_{l \to 2} \varphi^{EU}(b,l) = \lim_{l \to 3} \varphi^*(b,l) = 1$, as established in Proposition 5 part 3.

The results of Proposition 7 are verified since, for $\lambda = \frac{2}{\sqrt{17}-2} \approx .942$, we obtain $\varphi^{EU}(a,\lambda) = \varphi^*(a,\lambda) = \frac{\sqrt{17}-1}{4}$ and $\varphi^{EU}(b,\lambda) = \varphi^*(b,\lambda) = \frac{\sqrt{17}+1}{4}$. Moreover, $\varphi^{EU}(x,l) > \varphi^*(x,l)$ for all $l < \frac{2}{\sqrt{17}-2}$ and $\varphi^{EU}(x,l) < \varphi^*(x,l)$ for all $l > \frac{2}{\sqrt{17}-2}$.

It can be checked that the continuation functions $l^{\xi}_{n+1}\left(x, l^{\xi}_n\right)$ are strictly increasing. Figure 8 depicts the dynamics for the public likelihood ratio for the parameters considered in this example.

Since $\varphi^*(x,l)\,l$ and $\varphi^{EU}(x,l)\,l$ are strictly increasing, Proposition 6 implies that if A is the true state, then $l^*_n \to l^*_\infty(A) \in \left\{\frac{4}{9}, 3\right\}$ and $l^{EU}_n \to l^{EU}_\infty(A) \in \left\{\frac{2}{3}, 2\right\}$ and if B is the true state, then $l^*_n \to l^*_\infty(B) \in \left\{\frac{1}{3}, \frac{9}{4}\right\}$ and $l^{EU}_n \to l^{EU}_\infty(B) \in \left\{\frac{1}{2}, \frac{3}{2}\right\}$.

The probabilities of herding on the correct action satisfy

$$\begin{aligned}
\rho^*(A)\frac{4}{9} + [1 - \rho^*(A)]\,3 &= 1 \\
\rho^{EU}(A)\frac{2}{3} + \left[1 - \rho^{EU}(A)\right]2 &= 1 \\
\rho^*(B)\frac{1}{3} + [1 - \rho^*(B)]\frac{9}{4} &= 1 \\
\rho^{EU}(B)\frac{1}{2} + \left[1 - \rho^{EU}(B)\right]\frac{3}{2} &= 1
\end{aligned}$$

Solving gives $\rho^*(A) = \frac{18}{23}$ while $\rho^{EU}(A) = \frac{3}{4} < \frac{18}{23}$.and $\rho^*(B) = \frac{15}{23}$ while $\rho^{EU}(B) = \frac{1}{2} < \frac{15}{23}$. Therefore, the evolutionary rule induces herding on the correct action with higher probability. Therefore, Proposition 8 is verified, as $\gamma(\xi^*) = \frac{33}{46} \approx 0.717 > 0.625 = \frac{5}{8} = \gamma\left(\xi^{EU}\right)$.

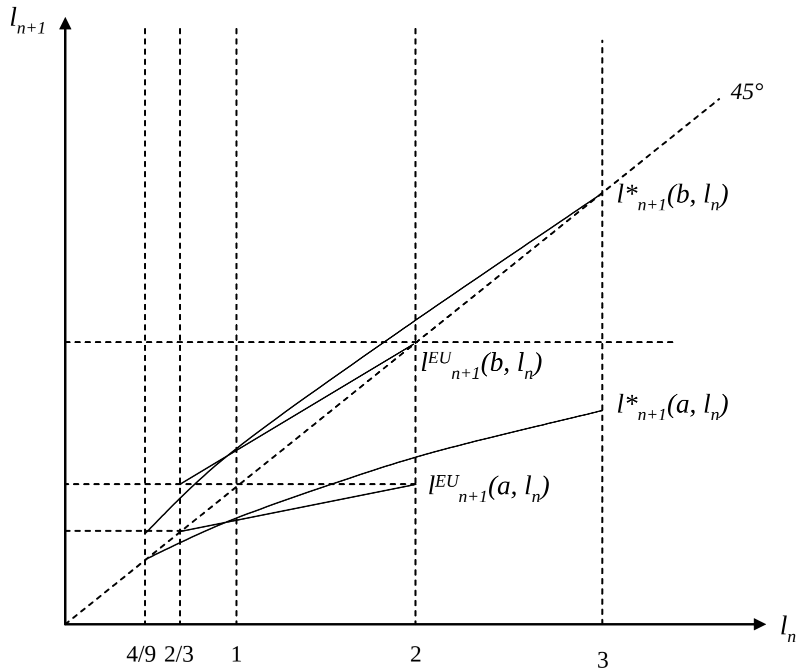

Figure 8: Numerical example for public likelihood ratio dynamics. Parameter values: $\sigma \in [0,1]$, $g^A = 1$, $g^B = \frac{3}{2} - \sigma$, $w(A,a) = w(B,b) = 3$ and $w(A,b) = w(B,a) = 2$.

Finally, since payoffs are symmetric, Proposition 8 implies that the population growth rate defined in (16) is higher for the evolutionary rule. We obtain $W^\infty\left(\xi^{EU}\right) = 3^{\gamma\left(\xi^{EU}\right)}2^{1-\gamma\left(\xi^{EU}\right)} \approx 2.577 < 2.675 \approx 3^{\gamma(\xi^*)}2^{1-\gamma(\xi^*)} = W^\infty(\xi^*)$.

5 Summary and further research

We presented a model where agents receive private signals correlated with the underlying state of the world. We found the appropriate measure of evolutionary performance and characterized the optimal behavior rule using this criterion. Compared to the standard expected utility rule, the evolutionary rules exhibits contrarianism and caution. Contrarian behavior consists of a probabilistic bias towards actions that defy the 'common wisdom' embedded in the prior beliefs. Agents exhibit caution when, compared to expected utility maximizers, a more extreme prior is required before disregarding their private information.

We introduced a class of behavior rules (CM-threshold rules) that include the evolutionary and the expected utility rules. We presented an application to social learning and showed that the basic properties of CM-threshold rules are sufficient to extend many well known results in the literature. In particular, public beliefs converge in the long run and herds eventually arise. Additionally, we showed that the evolutionary rule outperforms the expected utility. In this particular application, caution drives the result.

There are many issues that are left for future research. First, it would be interesting to study the consequences of the evolutionary rules to other models with private information, such as common value auctions, voting and information aggregation in markets, among others.

Second, the evolutionary model from section 3 can be extended in more than one direction. We assumed that private signals were received exogenously. However, the same reasoning used in this paper can be used in an information acquisition model. Intuitively, we expect that more revealing public information will increase incentives to acquire private information.

Also, it would be interesting to consider models in which there are strategic interactions between agents. The specific measure of evolutionary success would probably depend on the application at hand. In our model, there is a unique behavior rule that eventually dominates the population. Strategic interactions may provide conditions under which more than one behavior rule can be a nonvanishing fraction of the population in the long run.

Finally, the characterization of the constrained efficient behavior rule using the evolutionary criterion of section 3.1 remains an open question. Further research should only focus on this rule and compare it to the rules presented in this paper and in Smith & Sørensen (2006).

Appendix

Proof of Lemma 1. By definition,

$$\begin{aligned} E\left[\nu^{\xi}(s)\right] &= E\left[w\left(s, X^{\xi}(s)\right) \mid s\right] \\ &= w(s,b) + \Pr\left[X^{\xi}(s) = a\right]\left[w(s,a) - w(s,b)\right] \end{aligned}$$

Fix $E\left[\nu^{\xi}(A)\right]$. This is equivalent to setting $\Pr\left[X^{\xi}(A) = a\right] = z$, where z is a constant. We want to find the behavior rule ξ that maximizes $E\left[\nu^{\xi}(B)\right]$ from the set of rules that satisfy $\Pr\left[X^{\xi}(A)=a\right]=z$. In other words, we need to solve the following linear program:

$$\min_{\xi} \int_{\underline{\sigma}}^{\overline{\sigma}} \xi_a(\sigma)\, g^B(\sigma)\, d\sigma \tag{A.1}$$

$$\text{s.t.} \qquad \int_{\underline{\sigma}}^{\overline{\sigma}} \xi_a(\sigma)\, g^A(\sigma)\, d\sigma = z$$

$$\xi_a(\sigma) \in [0,1] \text{ for all } \sigma \in [\underline{\sigma}, \overline{\sigma}]$$

Define the multipliers η for the $\Pr\left[X^{\xi}(A)=a\right]=z$ constraint, $\overline{\beta}(\sigma)$ for the $\xi_a(\sigma)\leq 1$ constraint and $\underline{\beta}(\sigma)$ for the $\xi_a(\sigma)\geq 0$ constraint. The first order conditions are, for every σ,

$$g^B(\sigma) - \eta g^A(\sigma) + \underline{\beta}(\sigma) - \overline{\beta}(\sigma) = 0 \tag{A.2}$$

where $\eta, \underline{\beta}(\sigma), \overline{\beta}(\sigma) \geq 0$.

Suppose that $\xi_a(\sigma)\in(0,1)$. Then, $\underline{\beta}(\sigma)=\overline{\beta}(\sigma)=0$ and in consequence $g(\sigma)=\eta^{-1}$. Therefore, there can be only one $\sigma\in[\underline{\sigma},\overline{\sigma}]$ such that $\xi_a(\sigma)\in(0,1)$. We will call this signal θ. Suppose that $\sigma<\theta$. The monotone likelihood ratio property implies $g(\sigma)<g(\theta)=\eta^{-1}$. Therefore, (A.2) gives necessarily $\underline{\beta}(\sigma)>0$ and consequently $\xi_a(\sigma)=0$. Symmetrically, $\sigma<\theta$ implies $\overline{\beta}(\sigma)>0$ and $\xi_a(\sigma)=1$.

Therefore, threshold rules with threshold θ attains the frontier U. Moreover, since the solution to the program (A.1) is unique, only threshold rules can attain the frontier U.

For the converse statement, notice that if ξ is a threshold rule with threshold θ, then

$$E\left[\nu^{\xi}(A)\right] = w(A,a) - w(A,b)\, G^A(\theta)$$

Since g^A is an atomless density function, the distribution function $G^A(\theta)$ is continuous. Therefore, for every $z\in(w(A,b), w(A,a))$ there is a $\theta\in[\underline{\sigma},\overline{\sigma}]$ such that $z=w(A,a)-w(A,b)\,G^A(\theta)$. ■

Proof of Proposition 2. The expected utility rule ξ^{EU} is a behavior rule that maximizes (5). Since (5) is strictly increasing in $E\left[\nu^{\xi}(s)\right]$ for $s=A,B$, then ξ^{EU} is an element of the frontier U. By Lemma 1, ξ^{EU} is a threshold rule, with threshold θ^{EU}. The monotone likelihood ratio property of private signals implies that U can is a concave function in the space $\left(E\left[\nu^{\xi}(A)\right], E\left[\nu^{\xi}(B)\right]\right)$. Notice that $\theta^{EU}=\underline{\sigma}$ only if $-Kg(\underline{\sigma})=-K\underline{g}<-l$. Alternatively, if $l\leq K\underline{g}\equiv \underline{l}^{EU}$ then $\theta^{EU}(l)=\underline{\sigma}$. Symmetrically, if $l\geq K\overline{g}\equiv \overline{l}^{EU}$ then $\theta^{EU}(l)=\overline{\sigma}$. Finally, if $l\in\left(\underline{l}^{EU},\overline{l}^{EU}\right)$, then the solution to the optimization problem is interior and the tangency condition $-Kg\left(\theta^{EU}\right)=-l$ must apply. ■

Proof of Proposition 3. The evolutionary rule ξ^* is a behavior rule that maximizes $W^*(\xi)$. Since $W^*(\xi)$ is strictly increasing in $E\left[\nu^{\xi}(s)\right]$ for $s=A,B$, then ξ^* is an element of the frontier U. By Lemma 1, ξ^* is a threshold rule, with threshold θ^*. The monotone likelihood ratio property of private signals implies that U can is a concave function in the space $\left(E\left[\nu^{\xi}(A)\right], E\left[\nu^{\xi}(B)\right]\right)$. Notice that $\theta^*=\underline{\sigma}$ only if $-Kg(\underline{\sigma})=-K\underline{g}<-\frac{E\left[\nu^{\xi}(A)\right]}{E\left[\nu^{\xi}(B)\right]}l=-\frac{w(A,a)}{w(B,a)}l$. Alternatively, if $l\leq\frac{w(B,a)}{w(A,a)}K\underline{g}\equiv\underline{l}^*$ then $\theta^*(l)=\underline{\sigma}$. Symmetrically, if $l\geq$

$\frac{w(B,b)}{w(A,b)}K\overline{g} \equiv \bar{l}^*$ then $\theta^*(l) = \overline{\sigma}$. Finally, if $l \in \left(\underline{l}^*, \bar{l}^*\right)$, then the solution to the optimization pproblem is interior and the tangency condition $-Kg(\theta^*) = -\frac{E[\nu^\xi(A)]}{E[\nu^\xi(B)]}l$ must apply. ■

Proof of Proposition 4. First notice that for a continuous strictly increasing private signal threshold function $\theta^{EU}(l)$, we have that $Y\left(\theta^{EU}(l)\right)$ is continuous and strictly decreasing function of l. Also, $Y\left(\theta^{EU}\left(\underline{l}^{EU}\right)\right) = Y(\underline{\sigma}) > 1$ and $Y\left(\theta^{EU}\left(\bar{l}^{EU}\right)\right) = \Upsilon(\overline{\sigma}) < 1$. Therefore, there exists a unique $\lambda \in \left(\underline{l}^{EU}, \bar{l}^{EU}\right)$ such that $Y\left(\theta^{EU}(\lambda)\right) = 1$.

If $l \in \left(\underline{l}^*, \underline{l}^{EU}\right]$, then $\theta^{EU}(l) = \underline{\sigma} < \theta^*(l)$. Suppose now that $l \in \left(\underline{l}^{EU}, \lambda\right)$. Then θ^{EU} and θ^* are determined by (**??**) and (**??**), respectively. We want to show that $\theta^{EU}(l) < \theta^*(l)$. Assume in negation that $\theta^*(l) \leq \theta^{EU}(l)$. Then,

$$
\begin{aligned}
g\left(\theta^{EU}(l)\right) &= \frac{w(B,b) - w(B,a)}{w(A,a) - w(A,b)}l \\
&< Y\left(\theta^{EU}(l)\right)\frac{w(B,b) - w(B,a)}{w(A,a) - w(A,b)}l \\
&\leq Y(\theta^*(l))\frac{w(B,b) - w(B,a)}{w(A,a) - w(A,b)}l \\
&= g(\theta^*(l))
\end{aligned}
$$

where the first inequality follows from $Y\left(\theta^{EU}(l)\right) > 1$ and the second one from the fact that Υ is decreasing. Since g is strictly increasing, we get $\theta^*(l) > \theta^{EU}(l)$, a contradiction. This proves that $\theta^{EU}(l) < \theta^*(l)$ for all $l \in (\underline{l}^*, \lambda)$.

Using the symmetrical argument, it follows that $\theta^{EU}(l) > \theta^*(l)$ for all $l \in \left(\lambda, \bar{l}^*\right)$.

Finally, we have

$$
\begin{aligned}
\bar{l}^* &= \frac{w(B,b)}{w(A,b)}\bar{l}^{EU} > \bar{l}^{EU} \\
\underline{l}^* &= \frac{w(B,a)}{w(A,a)}\underline{l}^{EU} < \underline{l}^{EU}.
\end{aligned}
$$

Under our assumptions on $w(s,x)$, $\bar{l}^* > \bar{l}^{EU}$ and $\underline{l}^* < \underline{l}^{EU}$. ■

Proof of Proposition 5. Notice that part 1 is equivalent to $G^A(\theta(l)) < G^B(\theta(l))$ for all $l \in \left(\underline{l}, \bar{l}\right)$. This is implied by the monotone likelihood ratio $g \equiv \frac{g^A}{g^B}$. Since g^A and g^B integrate to one, there exists some $\lambda \in \left(\underline{l}, \bar{l}\right)$ such that $g^A(\theta(\lambda)) = g^B(\theta(\lambda))$. Moreover, because of the strict monotonicity, $g(\theta(l)) < g(\theta(\lambda)) = 1$ for all $l < \lambda$ and $g(\theta(l)) > g(\theta(\lambda)) = 1$ for all $l > \lambda$. This implies $G^A(\theta(l)) < G^B(\theta(l))$ for all $l \in (\underline{l}, \lambda)$. Now suppose that $G^A(\theta(l)) > G^B(\theta(l))$ for some $l \in \left(\lambda, \bar{l}\right)$. Since $g^A(\theta(l)) > g^B(\theta(l))$ for all $l \in \left(\lambda, \bar{l}\right)$, then $1 = G^A(\overline{\sigma}) > G^B(\overline{\sigma}) = 1$, a contradiction.

For part 2, first notice that since G^s and $\theta(l)$ are continuous and $G^s(\theta(l)) \in (0,1)$ for all $l \in (\underline{l}, \bar{l})$, then $\varphi^{\xi}(x,l)$ is continuous. The strict monotonicity of g implies

$$\begin{aligned} G^{A}(\sigma) &= \int_{\underline{\sigma}}^{\sigma} g^{A}(\tau)\, d\tau \\ &= \int_{\underline{\sigma}}^{\sigma} g(\tau)\, g^{B}(\tau)\, d\tau \\ &< g(\sigma) \int_{\underline{\sigma}}^{\sigma} g^{B}(\tau)\, d \\ &= g(\sigma)\, G^{B}(\sigma) \end{aligned}$$

Now take $l_1, l_2 \in (\underline{l}, \bar{l})$ such that $l_1 < l_2$. We have

$$\begin{aligned} G^{A}(\theta(l_2)) - G^{A}(\theta(l_1)) &= \int_{\theta(l_1)}^{\theta(l_2)} g^{A}(\tau)\, d\tau \\ &= \int_{\theta(l_1)}^{\theta(l_2)} g(\tau)\, g^{B}(\tau)\, d\tau \\ &> g(\theta(l_1)) \int_{\theta(l_1)}^{\theta(l_2)} g^{B}(\tau)\, d\tau \\ &= g(\theta(l_1)) \left[G^{B}(\theta(l_2)) - G^{B}(\theta(l_1))\right] \\ &> \frac{G^{A}(\theta(l_1))}{G^{B}(\theta(l_1))} \left[G^{B}(\theta(l_2)) - G^{B}(\theta(l_1))\right] \end{aligned}$$

which implies $\frac{G^{B}(\theta(l_2))}{G^{A}(\theta(l_2))} < \frac{G^{B}(\theta(l_1))}{G^{A}(\theta(l_1))}$.

Analogously, $1 - G^{A}(\sigma) > g(\sigma)\left[1 - G^{B}(\sigma)\right]$, which in turn implies$\frac{1-G^{B}(\theta(l_2))}{1-G^{A}(\theta(l_2))} < \frac{1-G^{B}(\theta(l_1))}{1-G^{A}(\theta(l_1))}$.

For part 3, simply notice that $\lim_{l \to \underline{l}} \varphi^{\xi}(a,l) = \frac{1-G^{B}(\theta(\underline{l}))}{1-G^{A}(\theta(\underline{l}))} = \frac{1-G^{B}(\underline{\sigma})}{1-G^{A}(\underline{\sigma})} = 1$ and $\lim_{l \to \bar{l}} \varphi^{\xi}(b,l) = \frac{G^{B}(\theta(\bar{l}))}{G^{A}(\theta(\bar{l}))} = \frac{G^{B}(\bar{\sigma})}{G^{A}(\bar{\sigma})} = 1$. ■

Proof of Proposition 6. First, we need to show that l_n^{ξ} is a bounded random variable. It suffices to show that $l_{n+1}^{\xi}(x,l) \in (0,\infty)$ for every $l \in (\underline{l}, \bar{l})$. Since l_{n+1}^{ξ} is continuous everywhere in $(\underline{l}, \bar{l})$, we only need to check that it is bounded in a neighborhood of $l \in \{\underline{l}, \bar{l}\}$. Proposition 5 part 3 guarantees this for $l_{n+1}^{\xi}(a,\underline{l})$ and $l_{n+1}^{\xi}(b,\bar{l})$. For $l_{n+1}^{\xi}(b,\underline{l})$ and $l_{n+1}^{\xi}(a,\bar{l})$, it follows from L'Hopital rule and the bounded informativeness of the private signals.

Since l_n^{ξ} is a bounded, so there are its second moments. Therefore, the Martingale convergence theorem implies $l_n^{\xi} \to l_{\infty}^{\xi}$ a.s. See Karlin and Taylor (1975), Theorem 5.2.

Next, we have that $\{l_n^{\xi}\}$ is a Markov chain. Since l_{n+1}^{ξ} is continuous everywhere in $(\underline{l}, \bar{l})$,it follows from Smith and Sørensen (2000), Theorem B.2 that every l in the support of l_{∞} needs to satisfy $l_{n+1}^{\xi}(x,l) = l$. Therefore, by Proposition 5 part 1 if $l \in (\underline{l}, \bar{l})$, then it cannot be in the support of l_{∞}.

Since l_n^{ξ} is bounded, the Dominated Convergence Theorem gives $E[l_{\infty} \mid l_0, A] = l_0 = 1 \in (\underline{l}, \bar{l})$. If $\Pr\left[l_{\infty}^{\xi} \in (0,\underline{l}] \mid A\right] = 1$, then $E[l_{\infty} \mid l_0, A] \in (0,\underline{l}]$, a contradiction.

Finally, since $\lim_{l_n \to \underline{l}} l_{n+1}^{\xi}(a, l_n^{\xi}) = l_n^{\xi}$ and $l_{n+1}^{\xi}(a, l_n^{\xi}) < l_n^{\xi} < l_{n+1}^{\xi}(a, l_n^{\xi})$ for $l_n^{\xi} \in (\underline{l}, \bar{l})$, then the strict monotonicity implies that all $l \in (0,\underline{l})$ are not accessible from $(\underline{l}, \bar{l})$. Analogously, $l \in (\bar{l}, \infty)$ are not accessible either. ■

Proof of Proposition 8. Since $l_{n+1}^{\xi}(x, l_n^{\xi})$ is strictly increasing in l_n^{ξ} for $\xi \in \left\{\xi^{*}, \xi^{EU}\right\}$, Proposition 6 gives $l_{\infty}^{\xi}(s) \in \{\underline{l}, \bar{l}\}$. Notice that l_{∞}^{ξ} is a bounded random variable. Therefore, the Dominated Convergence Theorem implies $E\left[l_{\infty}^{\xi}(A) \mid l_0^{\xi}, A\right] = l_0^{\xi} = 1$. Since $l_{\infty}^{\xi}(A) \in \{\underline{l}, \bar{l}\}$, $\rho^{\xi}(A)\underline{l} + (1-\rho^{\xi}(A))\bar{l} = 1$. Symmetrically, $E\left[\frac{1}{l_{\infty}^{\xi}(B)} \mid l_0^{\xi}, A\right] = \frac{1}{l_0^{\xi}} = 1$ and $l_{\infty}^{\xi}(B) \in \{\underline{l}, \bar{l}\}$. Therefore, $\rho^{\xi}(B)\frac{1}{\bar{l}} + (1-\rho^{\xi}(B))\frac{1}{\underline{l}} = 1$. Solving these equations gives

$$\begin{aligned}
\rho^{\xi}(A) &= \frac{\bar{l}-1}{\bar{l}-\underline{l}} \\
\rho^{\xi}(B) &= \frac{\bar{l}-\bar{l}\underline{l}}{\bar{l}-\underline{l}}
\end{aligned}$$

Proposition 4 gives $\bar{l}^{*} > \bar{l}^{EU} > 1$ and $\underline{l}^{*} < \underline{l}^{EU} < 1$. Therefore, it would suffice to show that $\rho^{\xi}(A)+\rho^{\xi}(B)$ is strictly increasing in $\bar{l}$ and strictly decreasing in $\underline{l}$. Since $\rho^{\xi}(A)+\rho^{\xi}(B)$ is differentiable, we compute the partial derivatives:

$$\begin{aligned}
\frac{\partial}{\partial \bar{l}}\left[\rho^{\xi}(A)+\rho^{\xi}(B)\right] &= \left[\frac{1-\underline{l}}{\bar{l}-\underline{l}}\right]^2 > 0 \\
\frac{\partial}{\partial \underline{l}}\left[\rho^{\xi}(A)+\rho^{\xi}(B)\right] &= -\left[\frac{\bar{l}-1}{\bar{l}-\underline{l}}\right]^2 < 0
\end{aligned}$$

This implies $\rho^{*}(A)+\rho^{*}(B) > \rho^{EU}(A)+\rho^{EU}(B)$ and concludes the proof. ■

References

Alchian, A.A. (1950) "Uncertainty, Evolution and Economic Theory" *Journal of Political Economy* 58: 211–221.

Banerjee, A.V. (1992) "A simple model of herd behavior", *Quarterly Journal of Economic* 107, 797–817.

Banerjee, A. and D. Fudenberg (2004) "Word-of-mouth learning", *Games and Economic Behavior* 46, 1-22.

Becker, G.S. (1976) "Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology", *Journal of Economic Literature*, 14, 817-826.

Bernardo, A.E. and I. Welch (2001) "On the Evolution of Overconfidence and Entrepreneurs", *Journal of Economics and Management Strategy* 10, 301–330.

Bikhchandani, S., D. Hirshleifer and I. Welch (1992) "A theory of fads, fashion, custom, and informational cascades", *Journal of Political Economy* 100, 992–1026.

Blume, L. and D. Easley (1992) "Evolution and market behavior ", *Journal of*

Economic Theory 58, 9–40.

Blume, L. and D. Easley (2006) "If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets", *Econometrica* 74, 929–966.

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.

Crow, J.F. and M. Kimura (1970) *An Introduction to Population Genetics Theory.* New York: Harper and Row.

Curry, P.A. (2001) "Decision making under uncertainty and the evolution of interdependent preferences", *Journal of Economic Theory* 98, 357–369.

Dekel, E., J.C. Ely and O. Yilankaya (2007) "Evolution of preferences", *Review of Economic Studies* 74, 685–704.

Eliaz, K., D. Ray and R. Razin (2006) "Choice Shifts in Groups: A Decision-Theoretic Basis", *American Economic Review* 96, 1321–1332.

Ellison, G. and D. Fudenberg (1993) "Rules of thumb for social learning", *Journal of Political Economy* 101, 612–643.

Ellison, G. and D. Fudenberg (1995) "Word-of-mouth communication and social learning", *Quarterly Journal of Economic* 110, 93–125.

Ely, J.C. and O. Yilankaya (2001) "Nash equilibrium and the evolution of preferences", *Journal of Economic Theory* 97, 255–272.

Feddersen, T.J. and W. Pesendorfer (1996) "The swing voter's curse", *American Economic Review* 86, 408–424.

Feddersen, T.J. and W. Pesendorfer (1997) "Voting behavior and information aggregation in elections with private information", *Econometrica* 65, 1029–1058.

Friedman, M. (1953) *Essays in Positive Economics.* Chicago: University of Chicago Press.

Grossman, S. (1976) "On the efficiency of competitive stock markets where traders have diverse information", *Journal of Finance* 31, 573–585.

Hirshleifer, J. (1977), "Economics from a Biological Standpoint", *Journal of Law* *and Economics* 20, 1-52.

Karlin, S. and H.M. Taylor (1975) *A First Course in Stochastic Processes,* 2nd ed., Academic Press, New York, NY.

McNamara, J.M. (1995) "Implicit frequency dependence and kin selection in fluctuating environments", Evolutionary Ecology 9, 185-203.

Milgrom, P.R. and R.J. Weber (1982) "A theory of auctions and competitive bidding", *Econometrica* 50, 1089–1122.

Ok, E.A. and F. Vega-Redondo (2001) "On the evolution of individualistic preferences: an incomplete information scenario", *Journal of Economic Theory* 97, 231–254.

Palfrey, T.R. (1985) "Uncertainty resolution, private information aggregation and the Cournot competitive limit", *Review of Economic Studies* 52, 69–83.

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

Robson, A.J. and L. Samuelson (2008) "The Evolution of Time Preference with Aggregate Uncertainty", *American Economic Review*, forthcoming.

Sandroni, A. (2000) "Do markets favor agents able to make accurate predictions?", *Econometrica* 68, 1303–1341.

Smith, .L and P. N. Sørensen (2000) "Pathological Outcomes of Observational Learning", *Econometrica* 68, 371-398.

Smith, .L and P. N. Sørensen, (2006) "Informational Herding and Optimal Experimentation," Cowles Foundation Discussion Papers 1552, Cowles Foundation, Yale University.

Smith, W.L. and W.E. Wilkinson (1969) "On branching processes in random environments", *The Annals of Mathematical Statistics* 40, 814–827.

Tanny, D. (1977) "Limit Theorems for Branching Processes in a Random Environment", *The Annals of Probability* 5, 100-116.