

INFORMATION-BASED RELATIVE CONSUMPTION EFFECTS*

Larry Samuelson
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, Wisconsin 53706-1320 U.S.A.

July 26, 2001

Abstract Relative consumption effects, in which agents' preferences depend upon others' consumption, are often said to be the result of contests to secure resources that are allocated on the basis of one's status. This paper argues that ill nature may induce relative consumption effects in order to compensate for incomplete environmental information. Status-based and information-based relative consumption effects can lead to quite different comparative static properties and quite different policy prescriptions.

Contents

1	Introduction	1
2	Information and Consumption	3
2.1	The Environment	4
2.2	Full-Information Optimum	6
2.3	Ill-Information Optimum	8
3	Relative consumption effects	10
3.1	Observable-Consumption Optimum	10
3.2	Imperfect Information Processing	14
3.3	Example	18
4	A Comparison with Rank-Based Effects	22
5	Conclusion	25
6	Appendix: Proofs	28

*This paper originated in joint work with Jeroen Swinkels. I thank Jeroen, Andrew Postlewaite, John Kennan, Phil Reny, and Bill Sandholm for helpful comments and discussions. Financial support from the National Science Foundation is gratefully acknowledged.

INFORMATION-BASED RELATIVE CONSUMPTION EFFECTS

by Larry Samuelson

1 Introduction

Economists typically assume that agents' preferences depend only on their own consumption. In contrast, a variety of evidence points to the existence of relative consumption effects, in which person i 's preferences regarding i 's consumption depend upon person j 's consumption. "Happiness" studies, in which the correlation between measures of consumption and measures of happiness are weaker than expected, are often interpreted as evidence that utility depends largely upon how one's consumption compares to that of one's peers.¹ Eumark and Postlewaite [34] find evidence consistent with relative consumption effects in the labor-supply decisions of sisters. Frank [24] offers examples of consumption that appears to be motivated by relative considerations.²

Why would one care about relative consumption effects? Because they can lead to inequitable allocations. The standard foundation for relative consumption effects supposes that "prizes" (such as access to good schools or mates) are allocated to those who rank high on a measure of status. If one's rank is particularly sensitive to the conspicuous consumption of certain goods (such as housing), then consumption will be inequidly concentrated on these goods at the expense of others.³ Frank [24] argues that such considerations lead Americans to undersave, overconsume luxury goods, and underconsume leisure and public goods.⁴ Cole, Mailath and Postlewaite [14]

¹For example, Blanchflower and Oswald [9], Easterlin [20, 21], Inglehart [28], Oswald [35], Veenhoven [41].

²The relative income hypothesis (Dusenberry [19], Pollak [36]) similarly explains aggregate consumption data by positing a link between desired consumption and a reference level based on previous consumption.

³Inefficiency is not implied by the mere existence of what Hirsch [26] called positional goods, i.e., goods in fixed supply that can be consumed only by a (perhaps high ranking) subset of an economy's agents. Competitive market prices capture the externality that agent i cannot consume such a good if agent j does.

⁴The claim that Americans underconsume leisure often begins with an argument that hours worked have increased in recent decades, fueled primarily by increased female labor supply (because "we both have to work to make ends meet"). Cox and Alm [16] argue that leisure time has actually *increased* once one accounts for the technology-induced reduction in time required for domestic production.

model the ability of status considerations to boost consumption of status-producing activities.

Concerns about relative consumption effects are heightened by the fact that income inequality has worsened in the United States in recent decades (Katz and Autor [30]). Does the resulting consumption pattern exacerbate the inefficiency caused by relative consumption effects? If so, are there appropriate policy responses?

This paper pursues an argument of Samuelson and Swinkels [39]: that relative consumption effects may have nothing to do with status-allocated goods. Instead, relative consumption effects may be a tool nature uses to equip us with preferences that are more effectively matched to the uncertainties of our environment.

Nature must design us, through the trial-and-error process of evolution, to make decisions in a wide variety of circumstances. Giving us preferences over absolute consumption levels can be risky, causing us to exert too much effort in low-yield environments and too little effort in high-yield environments. Nature responds by designing us to condition our behavior on information gleaned from our environment. But the consumption of one's contemporaries is often an important source of environmental information. High consumption levels reflect a relatively bountiful environment in which one might be well advised to increase one's own consumption. Low consumption levels reflect meager environments in which fanatical attempts at consumption are ill advised. Relative consumption effects, in the form of a tendency to push consumption levels in the direction of those observed in one's contemporaries, thus make optimal use of scarce information.

In the absence of constraints, nature could induce such relative consumption effects by coupling ordinary utility functions, defined only over one's own consumption, with a sufficiently sophisticated ability to process information. Information would then be used optimally. The policy implications of relative consumption effects could consist entirely of measures to ensure that information is also produced optimally.

Unfortunately, it is beyond our capabilities to process all of the useful information we receive. As a result, nature must make some compromises in designing our preferences and information-processing rules. Samuelson and Swinkels [39] examine nature's ability to mitigate information problems with utility functions that seem anomalous from a classical point of view. Building relative consumption effects directly into preferences is one of these complexity-cutting shortcuts, substituting an approximate but simple decision rule for a more precise but complicated inference problem.

Two risks arise in working with such preferences. For nature, there is the

danger that her agents will sometimes react to others' consumption levels without good informational grounds for doing so. Nature must balance this danger against the complexity costs of more precise information processing. In the course of doing so, relative consumption effects spill into utility functions, giving rise to policy implications that go beyond the mere production and use of information. For an analyst modeling the preferences, there is the danger of building an ever-increasing menagerie of arbitrary effects or "just so stories" into preferences, a new one for every behavioral anomaly (cf. Postlewaite [37]). A useful safeguard against abuse in this respect is to provide an evolutionary foundation for such preference effects, as we do in Sections 2-3.

This paper concentrates on information-driven relative consumption effects. Status is undoubtedly also important in resource allocation. But status and information can have quite different implications that can be untangled only by carefully modeling each.

Section 2 examines a simple model of the problem Nature faces when designing agents to live in a fluctuating environment. Section 3 derives Nature's optimal solution, notes that information-processing constraints may prompt Nature to implement this solution by building relative consumption effects into the utility function, and presents an example. Section 4 contrasts the implications of information-based relative consumption effects with those arising out of concerns for status. Concentrating the distribution of productivities alleviates relative-consumption distortions in the first case, but amplifies them in the second. Section 5 concludes with a very brief discussion of how information-based and status-based relative consumption effects might be empirically distinguished.

2 Information and Consumption

This section presents a model in which agents' decisions are guided by inferences they draw from observations of others' decisions. Beginning with the work of Banerjee [5] and Bhikhchandani, Hirshleifer and Welch [7], a variety of herding models have captured similar considerations. In a typical herding model, each agent uses Bayes' rule to assess the signal he has observed directly and the information contained in his observations of his predecessors' decisions. Crucial to the latter assessment is the agent's ability to infer the Bayesian decision problems faced by each of his predecessors.

In our case, this updating process is complicated by two features of Nature's design problem. First, agents in our model observe their predecessors

through the filter of natural selection. Agents who have chosen strategies well-suited to the current environment are more likely to survive than are others, biasing the mix of survivors whose decisions can be observed by subsequent agents. An agent's observed behavior thus mixes clues about the agent's information with clues about his evolutionary experience, both of which enter the observer's inference problem. The incorporation of the latter effect makes the problem more similar to that of Banerjee and Fudenberg [6] and Ellison and Fudenberg [22, 23] than to herding models.

Secondly, the state of the environment itself fluctuates in our model, forcing the agents to draw inferences about a moving target. This disrupts the simple recursive structure typical of herding models. Instead, we examine a stationary equilibrium in which the optimal decision rule shapes the inference and decision problems faced by one's predecessors, which in turn affects the current inference problem and hence the optimal decision rule. A sequence of recursive optimization problems is thus replaced by a fixed-point problem.

2.1 The Environment

Let time be divided into discrete periods. At the beginning of each period t , the environment is characterized by a variable $\mu_t \in [\underline{\mu}, \bar{\mu}]$. The events within a period proceed as follows:

1. Each member of a continuum of surviving agents gives birth, with each surviving agent giving birth to the same, exogenously fixed number of offspring. Each offspring is characterized by a parameter θ . The realized values of θ are uniformly distributed on $[0, 1]$. We interpret each offspring as (not necessarily independently) drawing θ from a uniform distribution.
2. Each newly-born agent observes n randomly selected surviving agents from the previous generation, discerning whether each chose \underline{z} or \bar{z} , and also observes a signal $\sigma \in \mathbb{R}$ drawn from a distribution $G(\sigma; \mu_t)$ with density $g(\sigma; \mu_t)$.
3. All parents then die. Each member of the new generation chooses a consumption strategy $z \in \{\underline{z}, \bar{z}\}$.
4. Nature then conducts survival lotteries, where $h(z; \theta; \mu) \in [0, 1] \in [\underline{\mu}, \bar{\mu}] \in [0, 1]$ gives the probability that an agent with strategy z and characteristic θ survives when the state of the environment is μ . As a gain, we assume no aggregate uncertainty.

5. Nature draws a value $\mu_{t+1} \in [\underline{\mu}, \bar{\mu}]$.

We interpret \underline{z} and \bar{z} as denoting low-consumption and high-consumption strategies. It may appear as if higher consumption levels, by reducing the chance of falling below a starvation barrier, must be preferred in a biological context (with free disposal obviating the danger of excess consumption). However, the high-consumption strategy may demand more effort and energy and may expose the agent to greater danger, opening the possibility that the low-consumption alternative is superior.

We allow the ranking of these strategies to depend upon both individual characteristics and the state of the environment. Some agents may be better-endowed with the skills that reduce the risk of procuring consumption than others. Some environments may feature more plentiful and less risky consumption opportunities than others. These effects appear in the specification of the survival probabilities $h(z; \mu)$, given by

$$h(\underline{z}; \mu) = \frac{1}{2}$$

$$h(\bar{z}; \bar{\mu}) = \frac{1}{2} + b(\alpha - q) \tag{1}$$

$$h(\bar{z}; \underline{\mu}) = \frac{1}{2} + b(\alpha - (1 - q)); \tag{2}$$

where $0 < q < 1/2$. Figure 1 depicts these probabilities. The low-consumption strategy \underline{z} yields a survival probability of $\frac{1}{2}$, regardless of the agent's characteristic or state of the environment. The high-consumption strategy \bar{z} yields a higher survival probability for agents with higher values of α and yields a higher survival probability when the state is $\bar{\mu}$. To ensure the probabilities are well defined, we assume

$$0 < b < 1/(2(1 - q));$$

The environmental parameter μ follows a Markov process, with transition rule

$$\begin{array}{c|cc} & \bar{\mu} & \underline{\mu} \\ \hline \bar{\mu} & 1 - \zeta & \zeta \\ \hline \underline{\mu} & \zeta & 1 - \zeta \end{array} \tag{3}$$

where

$$\zeta < \frac{1}{2};$$

so that next period's state is most likely to match this period's.⁵

⁵The symmetry of this Markov process simplifies the calculations, but is not crucial for the results.




Figure1: Survival probabilities

We assume that the density $g(\gg; \mu)$ is continuous and strictly positive on \mathbb{R} , for both $\underline{\mu}$ and $\bar{\mu}$, and satisfies the monotone likelihood ratio property (Milgrom [33]):

$$\gg'' > \gg' \Rightarrow \frac{g(\gg''; \bar{\mu})}{g(\gg''; \underline{\mu})} > \frac{g(\gg'; \bar{\mu})}{g(\gg'; \underline{\mu})} \quad (4)$$

so that higher values of \gg lead the agent to believe that state $\bar{\mu}$ is more likely. We assume that the proportion $G(\gg'; \mu)$ of the agents characterized by each value of μ draw values of $\gg \cdot \gg'$.⁶

2.2 Full-Information Optimum

We assume throughout that each agent knows the value of his characteristic μ . Much then hinges on what the agents and nature can observe concerning the state μ .

If the agent can observe μ (i.e., $G(\gg; \mu)$ is perfectly informative), then observations of the previous generation are irrelevant and the optimal decision

⁶Again, we assume there is no aggregate uncertainty. In each case, the fact that we do not require independence could be exploited to construct an explicit measure-theoretic foundation.

is straightforward, given by:⁷

$$z = \begin{cases} \mu = \bar{\mu} \text{ and } \sigma^2 > q \\ \mu = \underline{\mu} \text{ and } \sigma^2 > 1 - q \end{cases}$$

If $\underline{\mu}$ denotes winter and $\bar{\mu}$ summer, then the environmental characteristic seems sufficiently unambiguous as to be observed by the agents. Even here, however, mistakes can occur. Plants can produce buds at an inopportune time, sometimes disastrously, on the strength of an uncharacteristically warm spell. Animals can make ill-advised migration and hibernation decisions. People can plant or harvest crops too early or too late. Hence, even with something so obvious as the season, we can expect agents' abilities to observe the state to fall short of the desired precision.

More generally, the environmental variable may be something less obvious than climate. For example, $\underline{\mu}$ and $\bar{\mu}$ may index the relative availability of edible plants and slow-moving animals, which in turn may move in noisy cycles, leaving the agents with quite imprecise beliefs about the state of the environment.

This uncertainty on the part of the agents is no problem if Nature can observe the state of the environment. But Nature "observes" the state of the environment only by allowing the trial-and-error process of evolution to produce a decision rule that is optimal given the state, and hence can only observe states that are stationary for long periods of time. If the relevant information contained in $\underline{\mu}$ and $\bar{\mu}$ is only whether the edible animals in one's environment are mostly dinosaurs or mostly mammals, then the state changes slowly enough that Nature can simply count on adapting her agents to the environment, equipping them with the $\underline{\mu}$ or $\bar{\mu}$ decision rule (and causing them to ignore σ) as relevant. There will be some disruption whenever the state changes, but such transitions are sufficiently infrequent and idiosyncratic that Nature cannot endow her agents with the capability to respond.

We are interested in cases in which the state μ exhibits fluctuations which are not perfectly observed by the agents (so that $G(\sigma; \mu)$ is nontrivial) and which are sufficiently transitory that Nature cannot observe them. We assume that the process generating the state, given by (3), remains constant for a sufficiently long time as to be observed by Nature.⁸

⁷Throughout, we ignore measure-zero cases in which decisions are arbitrary.

⁸If not, then we would embed the current process in a larger one which generates short-term transition matrices like (3) according to a rule that is sufficiently stationary as to be observed by Nature.

2.3 No-Information Optimum

Suppose $n = 0$ (no members of the previous generation can be observed) and the agents know only their characteristic z and signal s and that the state of the environment is generated by a Markov process given by (3). The stationary distribution of this Markov process attaches a probability of $\frac{1}{2}$ to each of the states $\underline{\mu}$ and $\bar{\mu}$. Hence, the probability that the state is $\bar{\mu}$, conditional on having observed signal s , is

$$\frac{1}{2}(\bar{\mu}|s) = \frac{\frac{1}{2}g(s; \bar{\mu})}{\frac{1}{2}g(s; \bar{\mu}) + \frac{1}{2}g(s; \underline{\mu})}.$$

An agent with no additional information about the state maximizes his probability of withstanding the survival lottery he faces by choosing

$$z = \begin{cases} \bar{\mu} & \text{if } \frac{1}{2}(\bar{\mu}|s)h(z; \bar{\mu}) + (1 - \frac{1}{2}(\bar{\mu}|s))h(z; \underline{\mu}) > \frac{1}{2}; \\ \underline{\mu} & \text{otherwise.} \end{cases} \quad (5)$$

Using (1)–(2) to fill in the survival probabilities, this inequality becomes

$$z = \begin{cases} \bar{\mu} & \text{if } z > \frac{1}{2}(\bar{\mu}|s)q + (1 - \frac{1}{2}(\bar{\mu}|s))(1 - q); \\ \underline{\mu} & \text{otherwise.} \end{cases}$$

For each posterior probability $\frac{1}{2}(\bar{\mu}|s) = \frac{1}{2}$ there is thus a critical characteristic $z^*(\frac{1}{2}) = \frac{1}{2}q + (1 - \frac{1}{2})(1 - q)$, with larger values of z prompting a choice of $\bar{\mu}$ and lower values prompting $\underline{\mu}$. The function $z^*(\frac{1}{2})$ is decreasing in $\frac{1}{2}$, meaning that the agent is more likely to choose $\bar{\mu}$ when the prior probability of state $\bar{\mu}$ is larger. Notice that $z^*(1) = q$ and $z^*(0) = 1 - q$ so that the agent's actions coincide with the full-information optimum when the prior is degenerate.

The decision rule embedded in $z^*(\frac{1}{2})$ maximizes the probability that an agent survives a single period. However, evolution will select for the decision rule that conditions the agent's behavior on his posterior $\frac{1}{2}$ (only, since it is a sufficient statistic for the problem facing the agent) so as to choose⁹

⁹The evolutionarily most successful strategy in this setting will (as nearly as possible) maximize the probability of surviving for T periods, for all sufficiently large values of T . The probability of surviving for T periods is given by $\sum_{k=0}^T \bar{p}(\epsilon^*)^k \underline{p}(\epsilon^*)^{T-k} \text{prob}_T(k)$, where $\text{prob}_T(k)$ is the probability that, over the course of T periods, exactly k of these periods are characterized by state $\bar{\theta}$ and $T-k$ by state $\underline{\theta}$. As T gets large, this maximization problem converges to (7), in the sense that the maximizer of (7) comes arbitrarily close to maximizing the T -period survival probability for all large T , and hence evolution will select for the maximizer of (7). (The key to establishing this equivalence is to note that over long periods of time, the proportion of $\bar{\theta}$ states is very nearly $\rho(\bar{\theta}|\xi)$. The probability of survival until (large) period T is then very nearly $\bar{p}(\epsilon^*)^{\rho(\bar{\theta}|\xi)} (1 - \bar{p}(\epsilon^*))^{\rho(\underline{\theta}|\xi)}$, which is equivalent to (7).) The equivalence of (7) and (6) then follows from noting that (6) is necessary and sufficient (excepting deviations on sets of measure zero) for maximizing (7). See Robson [38]. Results of this type are standard in evolutionary biology (Charlesworth [11, pp. 47–50], Tuljapurkar [40]).

$$z^* = \frac{1}{2} \ln \left(\frac{p(\bar{\mu})}{p(\underline{\mu})} \right) \quad (6)$$

Equivalently, let $p(z^*; \bar{\mu})$ be the expected probability that an agent carrying a gene inducing decision cutoff z^* ($\frac{1}{2}$) survives nature's survival lottery in a period in which the state of the environment is $\bar{\mu}$, with the expectation taken over the value of z . Let $p(z^*; \underline{\mu})$ be similarly defined for state $\underline{\mu}$. Then (6) is equivalent to choosing the rule z^* ($\frac{1}{2}$) to maximize:

$$\frac{1}{2} p(z^*; \bar{\mu}) + \left(1 - \frac{1}{2}\right) p(z^*; \underline{\mu}) \quad (7)$$

In contrast, the decision rule given by (5) maximizes the average survival probability $\frac{1}{2} p(\bar{\mu}) + \left(1 - \frac{1}{2}\right) p(\underline{\mu})$. To see the difference, consider an environment in which state $\underline{\mu}$ is quite unlikely. Maximizing the probability of surviving a single period (as in (5)) may then call for a strategy that entails a zero survival probability in state $\underline{\mu}$ (in return for a higher survival probability in the much more likely state $\bar{\mu}$). But over a longer period of time, this strategy is a sure recipe for extinction. A superior strategy increases the survival probability in state $\underline{\mu}$ at the expense of state $\bar{\mu}$, as does the maximizer of (7).

When $\frac{1}{2} = \frac{1}{2}$, (5) will set $z^* = \frac{1}{2}$. In contrast, (7) calls for a value of z^* that exceeds $\frac{1}{2}$. In particular, the functions $p(z^*; \bar{\mu})$ and $p(z^*; \underline{\mu})$ are given by:

$$p(z^*; \bar{\mu}) = \int_{\epsilon^*}^{\frac{1}{2}} h(z; \bar{\mu}) dz + \int_{\frac{1}{2}}^1 h(z; \bar{\mu}) dz$$

$$p(z^*; \underline{\mu}) = \int_0^{\frac{1}{2}} h(z; \underline{\mu}) dz + \int_{\frac{1}{2}}^{\epsilon^*} h(z; \underline{\mu}) dz$$

These functions are strictly concave and have slopes of equal absolute value but opposite sign when $z^* = \frac{1}{2}$.¹⁰ To maximize the product $\frac{1}{2} p(z^*; \bar{\mu}) + \frac{1}{2} p(z^*; \underline{\mu})$, one then increases z^* above $\frac{1}{2}$, trading off a decrease in $p(z^*; \bar{\mu})$ for a smaller increase in the value of $p(z^*; \underline{\mu})$. Because $p(z^*; \underline{\mu})$ is relatively small, this increases the sum $\frac{1}{2} p(z^*; \bar{\mu}) + \frac{1}{2} p(z^*; \underline{\mu})$ and hence increases the long-term survival probability.

¹⁰The slope characterization follows immediately from the observation that $h(\bar{z}, \frac{1}{2}, \bar{\theta}) - \frac{1}{2} = \frac{1}{2} - h(\bar{z}, \frac{1}{2}, \underline{\theta})$, as shown in Figure 1. It is this property which ensures that setting $\epsilon^* = \frac{1}{2}$ maximizes the single-period survival probability $\frac{1}{2}(p(\epsilon^*, \bar{\theta}) + p(\epsilon^*, \underline{\theta}))$.

As $s^{1/2}$ increases, the optimal value z^* decreases and $p(z^*; \bar{\mu})$ increases while $p(z^*; \underline{\mu})$ decreases. The investment in survival probability is thus shifted toward the more likely state.

3 Relative consumption effects

3.1 Observable-Consumption Optimum

We now turn to the case in which agents can infer information about μ from the noisy signal provided by observing n (> 0) agents from the preceding generation (as well as from the signal θ). If the state affects the relative payoff of a high-consumption hunting strategy, for example, then an agent may be able to infer information about μ by watching whether his surviving neighbors appear to be mostly those who hunt or those who plant.

Nature's optimum incorporates the agent's information into his decision rule.¹¹ A decision rule for an agent is now a collection of functions $f^2(n; \theta); f^2(n; 1; \theta); \dots; f^2(0; \theta)$, where $f^2(k; \theta)$ is the cutoff value of z for choosing Z (Z being chosen for larger values of z) when the agent has observed k cases of Z and $n - k$ cases of \bar{Z} in the previous generation and has observed signal θ . Figure 2 illustrates such a strategy for the case in which $n = 1$. We let E denote a collection $f^2(n; \theta); \dots; f^2(0; \theta)$.

It is clear that it can never be optimal for $f^2(k; \theta)$ to lie outside the bounds $[q; 1 - q]$. Our intuition is that higher values of θ and more observations of Z (rather than \bar{Z}) will indicate that state $\bar{\mu}$ is more likely, making the agent more inclined to choose Z . Hence, we expect equilibrium decision rules to be at least weakly decreasing in θ and to satisfy, for all θ ,

$$q \leq f^2(n; \theta) \leq \dots \leq f^2(0; \theta) \leq 1 - q$$

We refer to a collection of functions $f^2(n; \theta); \dots; f^2(0; \theta)$ that satisfies these inequalities as being admissible.

Let \bar{A}_t be the proportion of strategy Z among those agents who survived period $t; 1$. Then \bar{A}_t describes the distribution from which period- t agents draw their observations of Z , with each period- t newborn observing Z on each draw with probability \bar{A}_t and observing \bar{Z} with probability $1 - \bar{A}_t$. Let $\varepsilon(\bar{A}_t; \mu_t)$ be the proportion of surviving period- t agents who chose Z , given that (i) these agents, as newborns, drew observations from the distribution described by \bar{A}_t , (ii) the period- t state of the environment relevant for

¹¹Any decision rule or mix of rules that does not make use of this information, whether tailored to state $\bar{\theta}$, $\underline{\theta}$, or some combination of the two, is evolutionarily dominated by a single rule that responds optimally to the information.




Figure 2: A dismissible decision rule $(\geq(1; \gg); \geq(0; \gg))$ when $n = 1$, conditional on a value $\gg = \gg'$

If nature's survival lottery is μ_t , and (iii) the agents' decision rule is given by $E(\cdot)$ given that nature first runs the period- t survival lottery and then determines the period- $(t+1)$ state, we can describe our system as a Markov process $(\tilde{A}_t; \mu_t)$ defined on the state space $[0; 1] \times \mathcal{F}(\bar{\mu})$. Letting $E(\mu)$ denote the transition rule given by (3), $(\tilde{a}_\varepsilon; E)$ denotes the transition rule for the process $(\tilde{A}_t; \mu_t)$:

$$\begin{aligned}\tilde{A}_{t+1} &= \tilde{a}_\varepsilon(\tilde{A}_t; \mu_t) \\ \mu_{t+1} &= E(\mu_t)\end{aligned}$$

Let $\mathbb{P}_\varepsilon((\tilde{A}_t; \mu_t); (\tilde{A}_0; \mu_0))$ be the probability distribution induced by the transition rule $(\tilde{a}_\varepsilon; E)$ in period t given initial condition $(\tilde{A}_0; \mu_0)$. Then define, for $k = 0; \dots; n$ (suppressing the dependence of \mathbb{P}_ε on E , the dependence on the initial condition, and some time subscripts on μ)

$$\mathbb{P}_\varepsilon(\bar{\mu}; j; k; \gg; t) = \frac{\int_{\psi_t=0}^1 g(\gg; \bar{\mu}) r(k; \tilde{A}_t) d\mathbb{P}(\tilde{A}_t; \bar{\mu})}{\int_{\psi_t=0}^1 g(\gg; \bar{\mu}) r(k; \tilde{A}_t) d\mathbb{P}(\tilde{A}_t; \bar{\mu}) + \int_{\psi_t=0}^1 g(\gg; \underline{\mu}) r(k; \tilde{A}_t) d\mathbb{P}(\tilde{A}_t; \underline{\mu})} \quad (8)$$

where

$$r(k; \tilde{A}_t) = \binom{n}{k} \tilde{A}_t^k (1 - \tilde{A}_t)^{n-k}$$

The function $\mathbb{P}_\varepsilon(\bar{\mu}; j; k; \gg; t)$ thus gives the probability that the state in time t is $\bar{\mu}$, given a time- t observation of signal \gg and k values of $\bar{\mu}$.¹² This

¹²The numerator of the right side is the joint probability that the state was $\bar{\theta}$ and the

probability reflects two considerations: the extent to which an observation of \bar{z} indicates that the previous-period state was relatively favorable for \bar{x} or strategy Z (i.e., was $\bar{\mu}$), and the probability that the state may have changed since the previous period.

Our first observation establishes conditions under which these probabilities have well-behaved limits:

Lemma 1 There exists a value $q^* \in (0, \frac{1}{2})$ such that for any $q \in (q^*, \frac{1}{2})$ and any admissible E , there exist probabilities $\frac{1}{2}_\varepsilon(\bar{\mu}^k; \bar{x})$ ($k = 0, \dots, n$) satisfying for all initial conditions,

$$\lim_{t \rightarrow \infty} \frac{1}{2}_\varepsilon(\bar{\mu}_t^k; \bar{x}) = \frac{1}{2}_\varepsilon(\bar{\mu}^k; \bar{x}):$$

In addition, the $\frac{1}{2}_\varepsilon(\bar{\mu}^k; \bar{x})$ are strictly increasing in k and satisfy $\frac{1}{2}_\varepsilon(\bar{\mu}^{k+1}; \bar{x}) > \frac{1}{2}_\varepsilon(\bar{\mu}^k; \bar{x})$.

The restriction that $q > q^*$ ensures that the population can never get too heavily concentrated on a single consumption strategy, either \bar{z} or \underline{z} . This in turn ensures that changes in the state of the environment are reflected relatively quickly in the observed distribution of consumption strategies, and hence that the latter is informative. To see how this could fail, consider the extreme case of $q = 0$. In this environment, there is no value of q for which \bar{z} is a dominant strategy. As a result, it is possible that virtually the entire population chooses \underline{z} . A change from state $\underline{\mu}$ to $\bar{\mu}$ will then not produce a noticeable change in the distribution of consumption strategies for an extraordinarily long time, causing this distribution to be relatively uninformative.¹³

The inequality $\frac{1}{2}_\varepsilon(\bar{\mu}^{k+1}; \bar{x}) > \frac{1}{2}_\varepsilon(\bar{\mu}^k; \bar{x})$ indicates that observations of high consumption enhance the posterior probability that the state of the environment is $\bar{\mu}$. This is the foundation of relative consumption effects.

The obvious approach to proving Lemma 1 is to show that the Markov process $(\xi; E)$ has an ergodic measure $\pi^*(\bar{A}; \mu; \bar{x})$ which can be inserted in (8) to calculate the probabilities $\frac{1}{2}_\varepsilon(\bar{\mu}^k; \bar{x})$. However, the Markov process

agent drew ξ and k observations of \bar{z} , while the denominator is the probability of drawing ξ and k observations of \underline{z} .

¹³If $q = 0$, our model would also be sensitive to the specification of the Markov process governing the state θ . If this process were asymmetric, then the possibility would arise that the population could be absorbed in a state in which all agents always choose either \bar{z} or \underline{z} . A similar possibility arises in Ellison and Fudenberg [23].

($\epsilon; E$) is not irreducible, failing the standard route to ergodicity.¹⁴ The proof constructs a candidate measure μ^* and establishes an appropriate convergence result.

Definition 1 The functions $E = f^2(\eta; \kappa); \dots; f^2(\theta; \kappa)$ are an equilibrium if $f^2(\kappa; \kappa)$ solves, for each value of κ and κ (cf. (7)),

$$\max_{\mu} \text{prcb}(\bar{\mu} | \kappa; \kappa) \text{Inp}(\epsilon; \bar{\mu}) + (1 - \text{prcb}(\bar{\mu} | \kappa; \kappa)) \text{Inp}(\epsilon; \mu) \quad (9)$$

where

$$\text{prcb}(\bar{\mu} | \kappa; \kappa) = \frac{1}{2} \epsilon (\bar{\mu} | \kappa; \kappa) \quad (10)$$

In defining an equilibrium, we use the limiting probabilities $\frac{1}{2} \epsilon (\bar{\mu} | \kappa; \kappa)$ to evaluate the payoff of a strategy. This reflects two timing assumptions. First, the state of the environment changes rapidly relative to the frequency with which nature can respond via mutation and selection, forcing nature to rely on agents' observations of others' consumption to infer information about the environment. Second, the process governing the state of the environment persists for a sufficiently long time that (i) nature can adapt her agents to this process, in particular tuning their information-updating rules to this process, and (ii) the limiting probabilities $\frac{1}{2} \epsilon (\bar{\mu} | \kappa; \kappa)$ are useful approximations of the information-updating problem facing the agents.

Proposition 1 Let $\epsilon > \frac{1}{2}$. Then an equilibrium with admissible functions $f^2(\eta; \kappa); \dots; f^2(\theta; \kappa)$ exists. In any such equilibrium, $f^2(\kappa; \kappa)$ is strictly decreasing in κ and satisfies $f^2(\kappa + 1; \kappa) < f^2(\kappa; \kappa)$.

Figure 3 depicts the resulting equilibrium. Agents are more likely to choose high consumption, i.e., choose Z for a wider range of ϵ , when κ and κ are large. Observations of high consumption, by increasing the expectation that the environment is in a state favorable to high consumption, increase an agent's propensity to choose high consumption. Relative consumption thus becomes useful as a response to environmental information that neither the agents nor nature can observe.¹⁵

¹⁴In particular, ψ_t is drawn from the uncountable set $[0, 1]$, while fixing a sequence of values of θ and iterating the function $\Psi_{\epsilon}(\psi, \theta)$ can generate only a countable set of values of ψ (though different initial conditions lead to different sets of such values). If we set $\tau = \frac{1}{2}$, sacrificing the (essential, for our purposes) persistence of the process Θ , then the period- t state would be independent of its predecessor and we could appeal to Futia's results on random contractions [25, Defs. 5.1, 6.2] to establish ergodicity.

¹⁵Given our restriction to admissible \mathcal{E} (e.g., to cases in which $\epsilon(k + 1, \xi) \leq \epsilon(k, \xi)$), showing that consumption observations are informative entails only precluding equality. We could strengthen our results to show that there is no equilibrium in which $\epsilon(k + 1, \xi) > \epsilon(k, \xi)$, but the argument is tedious.




Figure 3: Equilibrium consumption strategies

It is important to note that an agent's survival in our model depends only on the agent's own consumption. The route to genetic success is to choose optimal consumption levels, regardless of others' choices. The consumption levels of others are relevant (only) because they serve as valuable indicators of the state of the environment.

3.2 Imperfect Information Processing

The agents in our model optimally react to the uncertainty in their environment by being more likely to choose high consumption strategies when they observe others who have chosen high consumption. Nature can ensure such behavior by equipping her agents with an understanding of Bayes' rule and their environment and with a utility function over decision outcomes z^* given by

$$u(z^*; \mu) = \ln p(z^*; \mu):$$

An agent's expected-utility maximization problem, conditional on an observation $(k; \mathcal{E})$, is then given by the counterpart of (7):

$$\max_{\epsilon^*} \frac{1}{2} \epsilon (\mu | k; \mathcal{E}) \ln p(z^*; \bar{\mu}) + (1 - \frac{1}{2} \epsilon) (\mu | k; \mathcal{E}) \ln p(z^*; \mu): \quad (11)$$

Information is then used optimally, but is not produced optimally. Instead, agents make decisions concerning the production of information that ignore the value of this information to others.¹⁶ As a result, policy interventions designed to increase the production of information can yield Pareto superior outcomes. However, an appropriate policy could consist solely of providing information. If the resulting information was a sufficient statistic for observation of others' consumption, then relative consumption effects would disappear.

Unfortunately for Nature, she faces formidable constraints on her ability to design agents who process information perfectly. Cognitive resources are costly (Clark [13, Chapter 4]). Nature cannot ensure we have sufficient resources to flawlessly process the information contained in our environment, instead being forced to work with agents whose information processing is noisy. To capture this, we expand the model by assuming that, when calculating the posterior probability of state $\bar{\mu}$, there is probability $\lambda \in [0, 1)$ that the agent mistakenly conditions on an uninformative signal s^3 rather than s . Hence, the agent's posterior belief, now denoted $\mathbb{P}_\varepsilon(\mu|k; s; s^3)$, is given by

$$\mathbb{P}_\varepsilon(\mu|k; s; s^3) = \begin{cases} \mathbb{P}_\varepsilon(\mu|k; s) & \text{with probability } 1 - \lambda; \\ \mathbb{P}_\varepsilon(\mu|k; s^3) & \text{with probability } \lambda. \end{cases}$$

where the function \mathbb{P}_ε is given by (8). The agent calculates the posterior correctly with probability $1 - \lambda \in [0, 1)$. A miscalculating agent not only fails to recognize when his information is based on the signal s^3 , but fails to recognize the possibility of conditioning on such an uninformative signal (hence continuing to apply (8)). This is the essence of the agent's imperfect information processing.

We let $f(s^3)$ denote the density of s^3 . We assume this density is positive on \mathbb{R} and does not depend upon the state μ , ensuring that s^3 is uninformative. However, we ensure that s^3 is not too idiosyncratically distributed by assuming

$$s'' > s' \Rightarrow \begin{cases} \frac{g(s'', \bar{\mu})}{f(s'')} > \frac{g(s', \bar{\mu})}{f(s')} & ; \\ \frac{f(s'')}{g(s'', \bar{\mu})} > \frac{f(s')}{g(s', \bar{\mu})} \end{cases} \quad (12)$$

so that large signals (for example) are more likely to come from the informative distribution (given state $\bar{\mu}$) than the uninformative distribution. The

¹⁶An efficient solution would spread the values of $\{\varepsilon(n, \xi), \dots, \varepsilon(0, \xi)\}$ further apart, trading second-order losses in individual optimization for the production of more information. In the absence of group selection, the restriction of Nature's design mechanism to individual survival puts such a solution beyond her grasp.

value of ϵ^3 and whether $\frac{1}{2}\epsilon(\mu|k; \gg)$ or $\frac{1}{2}\epsilon(\mu|k; \epsilon^3)$ is the applicable posterior are drawn independently of one another and the other random variables in the model.

Inserting the posterior $\frac{1}{2}\epsilon(\mu|k; \gg; \epsilon^3)$ into (11), the agent's maximization problem is now

$$\max_{\epsilon^*} \frac{1}{2}\epsilon(\mu|k; \gg; \epsilon^3) \ln p(z^*; \bar{\mu}) + (1 - \frac{1}{2}\epsilon(\mu|k; \gg; \epsilon^3)) \ln p(z^*; \mu) \quad (13)$$

The noise captured by ϵ^3 pushes the agent's values of $\epsilon^2(k; \gg)$, away from their optimal values, so that information is no longer used optimally. Nature can always improve these choices by enhancing the agent's information-processing ability, which corresponds to decreasing ϵ , in this context, but eventually finds the opportunity cost of further improvements prohibitive.¹⁷ A reflective policy response to the resulting relative consumption effects now calls for the provision of information as well as the provision of assistance in processing information, with appropriately provided and processed information again banishing relative consumption effects and their attendant inefficiencies.

Given that she must work with the noisy beliefs $\frac{1}{2}$ rather than $\frac{1}{2}$, however, Nature will no longer find it optimal to let the utility function be given by $u(z^*; \mu) = \ln p(z^*; \mu)$ (as in (13)). Instead, Nature can improve the agent's performance by building the observation k into the utility function. Intuitively, she uses "distortions" of the utility function to correct for the noise in the agent's information processing.

Let $\bar{\mu}(\bar{\mu}^{z^*})$ be the probability of state $\bar{\mu}$ that causes z^* to solve the maximization problem (11) (or equivalently, (13)). Then:

Proposition 2 Nature can achieve her optimum, given that the agents' information updating is given by $\frac{1}{2}$, with a utility function $u(z^*; k; \mu)$ for which, for each $k = 0; \dots; n$, there is a value $z^*(k)$ such that

$$\bar{A}(\bar{\mu}^{z^*}) \frac{du(z^*; k; \bar{\mu})}{d\epsilon^*} + (1 - \bar{A}(\bar{\mu}^{z^*})) \frac{du(z^*; k; \mu)}{d\epsilon^*}$$

¹⁷ Among other difficulties, a perfect ability to process information runs the risk of being too cumbersome to be useful. We lack the time to analyze every piece of information. Nature again responds by building reactions to common information-processing and decision problems into our behavior. If Nature's agents habitually calculate that it is optimal to hedge consumption levels in the direction of those observed in their peers, then Nature will dispense with the calculation and embed relative considerations into the utility function. LeDoux [31] discusses the incentives for Nature to arm us with a mix of "hard-wired" and cognitive responses to our environment, arguing that many of our seemingly hard-wired reactions are engineered to ensure that we do not tarry to process information before acting.

$$\begin{aligned}
 & > (<) \bar{A}(\bar{\mu}^j; z^*) \frac{d \ln p(z^*; \bar{\mu})}{d p^*} + (1 - \bar{A}(\bar{\mu}^j; z^*)) \frac{d \ln p(z^*; \bar{\mu})}{d p^*} \\
 & \text{Nature's } z^* < (>) z^*(k): \tag{14}
 \end{aligned}$$

Recognizing that both $\bar{\mu}$ and z^* play a role in shaping the agent's beliefs, Nature regards the agent's posterior beliefs as being less informative than does the agent. She responds by adjusting the utility function to reduce the sensitivity of the agent's actions to his beliefs.¹⁸ She does so by designing the utility function to discourage the agent from straying too far from an action $z^*(k)$ (for each observation k), increasing (decreasing) the derivative of expected utility (relative to the utility function $\ln p(z^*; k)$) for smaller (larger) actions. Hence, Nature requires that the agent observe more persuasive information before straying too far from a target consumption plan, where the latter depends upon observations of others' consumption. The key observation is that Nature's manipulation of the utility function, and hence the agent's utility, depends upon the value of k . Relative consumption effects now take the form of incorporating others' consumption levels directly into the utility function.

Nature's problem would be particularly simple if $\bar{\mu}$ were uninformative. In this case, the agents have no useful information beyond that which is conveyed by the number of observations k , upon which Nature can condition utilities. She could then abandon all thoughts of allowing her agents to process information, having them maximize a utility function that depends upon k and that induces the optimal strategy $(z^*(n); \dots; z^*(1))$. Nature is pushing her agents toward this outcome when designing the utility function u of Proposition 2, and will do so more vigorously the less informative is $\bar{\mu}$. Similarly, Nature would not encounter difficulties if she could simply condition utilities directly on $\bar{\mu}$ (but not z^*), again essentially achieving perfect information processing. Nature's difficulties arise when she would like her agents to condition their actions on information such as $\bar{\mu}$ that she cannot 'observe', but which her agents cannot process perfectly.¹⁹

¹⁸Nature would prefer to make the agents' beliefs less sensitive to ζ , and will do so to the extent possible. But as long as she cannot purge ζ from the agent's beliefs, there are gains from manipulating the utility function. Why doesn't Nature instead make the agent less responsive to his beliefs by giving him a stronger prior or otherwise tinkering with the function p ? Because this makes the agent less responsive to k as well as ζ .

¹⁹What if the agent can also be mistaken concerning the information conveyed by k ? Then Nature's design problem is more difficult and the details of the solution will differ, but the utility function will again be manipulated to compensate for shortcomings in information processing.

We now have an expected utility function of the form

$$\max_{z^*} \frac{1}{2} \mathbb{E}(\mu | k; \gg; \mathbb{3}) u(z^*; k; \bar{\mu}) + (1 - \frac{1}{2}) \mathbb{E}(\mu | k; \gg; \mathbb{3}) u(z^*; k; \underline{\mu}) \quad (15)$$

or, equivalently, we can write this in terms of a utility function defined over consumption levels $z \in [z; z_g]$ rather than over consumption outcomes z^* (just as (6) and (7) are equivalent):

$$\max_z \frac{1}{2} \mathbb{E}(\mu | k; \gg; \mathbb{3}) u(z; k; \bar{\mu}) + (1 - \frac{1}{2}) \mathbb{E}(\mu | k; \gg; \mathbb{3}) u(z; k; \underline{\mu}) \quad (16)$$

Relative consumption effects now arise not only through the processing of information but also directly through utility externalities. An appropriate policy response may then call for more than simply providing and processing information. In particular, removing uncertainty from the economy will not banish relative consumption effects and their attendant inefficiencies.

"Notice that there is no question of Nature's designing us to solve some problems of inordinate complexity. The ability to read these words reflects a triumph of biological engineering. Our argument requires only that Nature cannot ensure that we can solve every complex problem we encounter, and that she will accordingly adopt information-processing shortcuts whenever she can. \ In general, evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand-in for the information-processing operations concerned." (Clark [13, p. 64].) We have developed a particularly simple model of the limitations Nature faces in designing her agents and the resulting information-processing shortcuts. The details of the shortcuts will depend upon the specific form of the limitations, but the general principle remains: information processing complexities can prompt relative consumption effects to spill over into utility externalities.

3.3 Example

We can illustrate the implications of (16). To focus attention on the persistence of relative consumption effects in the absence of uncertainty, we assume that the distribution of the state μ and the signals \gg and $\mathbb{3}$, and hence the agents' beliefs captured in \mathbb{E} , are degenerate. To simplify the derivations, we take a step toward realism in assuming that the consumption choice z is drawn from \mathbb{R}_+ (rather than $[z; z_g]$). The information received by the agent is then a vector of n consumption levels. We denote this observation by $(z_1, \dots, z_n) \in \mathbb{Z}$ and let $m(\mathbb{Z})$ be the resulting mean observed consumption level.

Let the utility function corresponding to (15) be given by

$$u(z; \mathbf{z}^j) = z \left[1 - \frac{\alpha}{2} (z - m(\mathbf{z}))^2 \right] \frac{z^2}{2} \quad (17)$$

We interpret $z^2/2$ as the disutility associated with securing consumption z and interpret α (distributed uniformly among the continuum of agents) as an index of productivity. As suggested by Proposition 2, relative consumption effects (captured by the term $1 - \frac{\alpha}{2} (z - m(\mathbf{z}))^2$) take the form of discouraging the agent from choosing a consumption level too different from a reference level that depends upon observations of others' consumption. Higher values of the parameter $\alpha > 0$ indicate that relative consumption effects are more important.

Let F be a distribution function with $F(z)$ identifying the proportion of the agents in the economy choosing a consumption level less than or equal to z . Let \hat{A}_F^n be the measure over n -tuples (z_1, \dots, z_n) induced by the distribution F . An equilibrium is then a distribution of consumption levels with the property that, if each agent is allowed to draw n observations from the previous generation's distribution and then choose an optimal consumption level, we reproduce the distribution.²⁰

Definition 2 An information-based equilibrium is a distribution $F^*(z)$ and a specification of utility-maximizing choices $z^*(z; \mathbf{z})$ such that

$$F^*(z) = \int \mathbb{1}_{F^*(z)}(\mathbf{z}) \cdot z^*(z; \mathbf{z}) \cdot \mu(\mathbf{z}) \quad (18)$$

where μ is Lebesgue measure.

Establishing the existence of such an equilibrium is a straightforward fixed-point argument. Solving the first-order condition for utility maximization, we obtain the optimal choice function

$$z^*(z; \mathbf{z}) = \frac{z \left[1 + \alpha m(\mathbf{z}) \right]}{1 + \alpha z} \quad (19)$$

If $\alpha > 0$, then relative consumption effects push the allocation away from the 'no distortion' solution of $z^*(z; m(\mathbf{z})) = z$, with larger values of α prompting large distortions. The optimal choice function $z^*(z; m(\mathbf{z}))$ is increasing in both z and $m(\mathbf{z})$, so that larger characteristics and larger observations of

²⁰We could equivalently interpret this as a rational-expectations equilibrium in which agents sample their own generation.

others' consumption lead to larger consumption. The equilibrium is inefficient if $\beta > 0$, with gains to be had from pushing the consumption of each agent for whom $\beta > 0$ closer to the mean consumption.

Information-based relative consumption effects thus induce some agents to increase consumption and others to decrease consumption, compared to the base case of no such effects. The former will be those who have observed other agents with particularly high consumption levels, while the latter will have observed small consumption levels.

We can calculate the support of the equilibrium distribution of consumption levels by first noticing that

$$z^*(0; m(\hat{z})|j) = 0;$$

so that agents with the lowest characteristic choose their no-distortion consumption level of zero, regardless of what they observe. Next, notice that

$$z^*(2; z^*(2; m(\hat{z})|j)) = z^*(2; m(\hat{z})|j) = 2. \quad (20)$$

Hence, an individual who observes a mean consumption equal to the consumption the agent would have been chosen in the absence of relative consumption effects acts as if there are no relative consumption effects. It follows from (20) that the least upper bound on the range of equilibrium values of $z^*(2; m(\hat{z})|j)$ is unity.²¹ As β increases from zero, the range of realized values thus remains fixed at $[0; 1]$, while the distortion created by relative consumption effects grows stronger.

We now consider the effects of a policy intervention that replaces β in the utility function given by (17) with $\alpha + (1 - \alpha)\beta$ for some $\alpha \in [0; 1]$.²² Figure 4 depicts this change. We can think of the policy intervention as a program

²¹To verify this, let $\tilde{z}(\alpha)$ be the upper bound. It suffices to consider $\epsilon = 1$ (since $z^*(\epsilon, m(\hat{z})|\alpha)$ is increasing in ϵ). If the range of $z^*(\epsilon, m(\hat{z})|\alpha)$ is finite, it suffices to observe that $z^*(1, \tilde{z}(\alpha)|\alpha) = \tilde{z}(\alpha)$ must hold, which can be satisfied only if $\tilde{z}(\alpha) = 1$. To see that an unbounded range of values of $z^*(1, m(\hat{z})|\alpha)$ is impossible, let $z^*(1, m(\hat{z})|\alpha) = h(m(\hat{z}))$. Then note that for $m(\hat{z}) > 1$, we have $m(\hat{z}) < h(m(\hat{z}))$ and $dh(m(\hat{z}))/dm(\hat{z}) < 1$. A positive measure of observations in the interval $[1 + \delta, 1 + \delta + \eta]$ would then require an identical measure in each of the infinite sequence of ascending, disjoint (for sufficiently small η) intervals $[h^{-1}(1 + \delta), h^{-1}(1 + \delta + \eta)]$, $[h^{-1}(h^{-1}(1 + \delta)), h^{-1}(h^{-1}(1 + \delta + \eta))]$, $[h^{-1}(h^{-1}(h^{-1}(1 + \delta))), h^{-1}(h^{-1}(h^{-1}(1 + \delta + \eta)))]$, and so on, an impossibility.

²²A growing literature addresses the policy implications of relative consumption effects. For example, Corneo [15] shows that in the presence of rank-based relative income effects, redistributive taxation may be most useful in societies with relatively egalitarian distributions of income. Ljungqvist and Uhlig [32] show that if current utilities depend upon past consumption, then countercyclical fiscal policy can be effective in dampening the relative effects.




Figure 4: Productivity enhancement

enhancing productivities in the population, concentrating on those at the bottom end of the productivity scale, perhaps by increasing access to educational resources. As $a \rightarrow 1$, the productivity of every agent in the economy approaches the maximum productivity of unity. We began this example with the case of $a = 0$.

Utility maximization now yields an optimal choice function of (cf. (19)):

$$z^*(z; m(z)j^\theta; a) = \frac{(a + (1-a)^2)(1 + \theta m(z))}{1 + \theta(a + (1-a)^2)};$$

Once again, for any value of θ , an agent of characteristic z who observes $m(z) = z^*(z; m(z)j^\theta; a) = a + (1-a)^2$ will choose $z = a + (1-a)^2$, so that an observation of a mean consumption level equal to the choice that would have been chosen in the absence of relative consumption effects leaves the optimal choice unaltered.

An argument analogous to the case of $a = 0$ allows us to conclude that, for any value of $\theta > 0$, the range of the equilibrium consumption levels is now $[a, 1]$ (compared to $[0, 1]$ in the original case of $a = 0$). As $a \rightarrow 1$, equilibrium consumption levels $z^*(z; m(z)j^\theta; a)$ thus converge on a mass point at unity (for all θ):

$$z^*(z; m(z)j^\theta; 1) = 1 \quad (21)$$

$$u(z^*(z; m(z)^{\theta}; 1); z; m(z)^{\theta}; 1) = \frac{1}{2}; \quad (22)$$

Hence, relative consumption effects and their attendant inefficiency disappear as productivity variation is eliminated from the population. Utilities converge to a value of $\frac{1}{2}$, yielding an efficient outcome in which each agent receives the maximum utility level in the population before the policy intervention.

4 A Comparison with Rank-Based Effects

We contrast the previous example with one in which relative consumption effects arise out of rank considerations.²³ We again consider a continuum of agents with characteristics z^2 distributed uniformly on $[0; 1]$ and who must choose consumption or effort levels $z \in \mathbb{R}_+$. Let the utility function now be given by

$$u(z; z; r(z)) = r(z) \cdot \frac{z^2 (1 + \theta z)^3}{2^2 (1 + \theta m)^2}; \quad (23)$$

where θ matches its value from Section 3.3 and m is the mean consumption level for the information-based equilibrium of Section 3.3 when $a = 0$ (both for reasons that will soon be apparent), and $r(z) : \mathbb{R}_+ \rightarrow [0; 1]$ identifies the rank of an agent who chooses consumption level z . We interpret rank as being valuable because there are prizes allocated on the basis of rank. The variable z may thus be valued for its own sake as well as for its ability to secure rank, where (23) captures the extreme case in which only rank is important.

Let $z^*(z; r)$ be a utility-maximizing consumption choice given characteristic z^2 and rank function r .

Definition 3 A rank-based equilibrium is a strictly increasing function $r^*(z) : \mathbb{R}_+ \rightarrow [0; 1]$ and specification of utility-maximizing consumption levels $z^*(z; r^*)$ such that

$$r^*(z) = \int_{\mathbb{R}_+} \mathbb{1}(z^*(z; r^*) \cdot z) \, d\mu;$$

where μ is Lebesgue measure.

A rank-based equilibrium is thus a distribution of consumption levels that is reproduced when each agent makes an optimal consumption choice, given

²³See Hopkins and Kornienko [27] for a similar analysis of rank-based effects.

the equilibrium choices of others and hence the translation of consumption levels to ranks.

In equilibrium, we must have $r(z^2) = z^2$, and hence the efficient solution maximizes

$$u(z^2; z^2) = z^2; \frac{z^2 (1 + \theta z^2)^3}{2^2 (1 + \theta m)^2};$$

giving

$$z^*(z^2; z^2) = 0;$$

Because only rank matters for allocation, efficient effort levels are zero

The first-order condition for the maximization of (23) is

$$\frac{dr(z)}{dz} = \frac{z (1 + \theta z)^3}{z^2 (1 + \theta m)^2} = 0;$$

Using the equilibrium condition $r(z^2) = z^2$, we have

$$\frac{dz}{d\theta} = \frac{1}{d\theta} = \frac{2 (1 + \theta m)^2}{z (1 + \theta z)^3}; \quad (24)$$

yielding a differential equation that we can solve for equilibrium consumption levels:

$$z^*(z^2; r^*) = z^2 \frac{1 + \theta m}{1 + \theta z^2};$$

Hence, this rank-based equilibrium gives rise to consumption levels identical to the mean consumption levels of the information-based equilibrium of Example 1 (cf. (19)). Average behavior in the two economies is thus indistinguishable.²⁴

Rank concerns prompt agents to choose higher equilibrium effort levels than does the efficient solution.²⁵ It would be individually valuable to do so, but the equilibrium effect of all agents doing so is to decrease utility. In equilibrium, the rank of agent 2 is given by z^2 . The efficient solution calls for agent 2 to simply assume that his rank is fixed at z^2 , and then choose an optimal consumption plan contingent on this rank. The realized equilibrium produces the same distribution of ranks, but only after every agent has increased consumption in a collectively vain effort to increase rank.

²⁴Brock and Durlauf [10] and Durlauf [18] discuss the issues that arise when studying behavior in the presence of relative consumption effects.

²⁵If we replaced $r(z)$ with $(1 - \gamma)z + \gamma r(z)$ in (23), so that effort level z leads both to direct consumption of $(1 - \gamma)z$ and to rank-based consumption of $\gamma r(z)$, then efficient effort levels would be positive, but equilibrium effort levels would still exceed efficient levels, with the discrepancy increasing in γ .

As before, let the characteristic z be replaced by $a + (1 - a)^2$, so that an increase in a corresponds to an enhancement of productivities that is concentrated among low-productivity agents. As $a \rightarrow 1$, every agent's productivity increases to the maximum level of unity. The efficient level of effort remains unaltered at zero. Taking the first-order condition for maximizing utility and again using the equilibrium relationship that $r(z^2) = z$, we can solve for (cf. (24))

$$\frac{dz}{d^2} = \frac{a + (1 - a)^2}{z} \frac{(1 + \theta m)^2}{(1 + \theta (a + (1 - a)^2))^3} = \frac{1 (1 + \theta m)^2}{z (1 + \theta)^3};$$

where the second equality holds for the limiting case of $a = 1$. Confining attention to this case and using the boundary condition that $z(0) = 0$, this differential equation can be solved for

$$z^*(z; r^*j1) = \frac{\bar{A}}{2^2} \frac{(1 + \theta m)^2}{(1 + \theta)^3} \frac{1}{2} > \frac{1 + \theta m}{1 + \theta^2}$$

$$u(z^*(z; r^*j1); z; r(z^*(z; r^*j1))) = 0:$$

Raising the productivity of every agent in the economy to the maximum tends to tighten the distribution of consumption, enhancing the payoff to seeking status and unleashing a countervailing tendency for every agent to compete more vigorously by increasing effort. Rank-based inequalities are thus exacerbated, with the entire population exerting more effort while being reduced to the zero utility level garnered by only the least fortunate agent before the policy intervention.²⁶

Inequality policies can thus have quite different effects depending upon the nature of the relative consumption effects that appear in utility functions. If these effects are nature's reaction to an inference problem, then smoothing productivities can alleviate the distortions of relative income effects. If the effects arise because rank or status is an important allocation device, then smoothing productivities can be ruinous. By increasing social mobility, inequality smoothing in the latter case ensures that more resources are wasted in the zero-sum quest for status. The inequality of relative consumption effects might be reduced in this case by entrenching consumption differences more solidly, making social mobility more difficult. Perhaps paradoxically, everyone in our deliberately extreme example might prefer social rigidity to social mobility, including those on the bottom.

²⁶For similar reasons, equilibria in rent-seeking games or wars of attrition are typically most wasteful when contestants are evenly matched.

5 Conclusion

The message of this paper is that concerns about rank and status are not the only possible origins of relative consumption effects. Instead, relative consumption effects may be an important information-processing tool for making better decisions in a fluctuating environment. In addition, the relative consumption effects arising out of these two motivations can have quite different properties. Smoothing the distribution of productivities can push equilibrium consumption patterns toward the base case of no relative consumption effects when the latter reflect information concerns, suggesting that equalization policies may be surprisingly effective in boosting utilities. Smoothing the productivity distribution can push consumption patterns away from the base case of no relative consumption effects when the latter arise out of status considerations, suggesting that attempts to tighten the distribution of status may counterproductively redouble the quest for high status. Relative-consumption inefficiencies are alleviated in the former case, but exacerbated in the latter.

Is there any evidence for information-based relative consumption effects, or for the limitations in our ability to process information that would prompt ill-attitude to build relative consumption effects into our utility functions? Psychologists report experimental evidence suggesting that people are poor Bayesians (Kahneman and Tversky [29]). Psychologists also report an inclination to conform to the behavior of others (Aronson [1, Chapter 2], Cialdini [12, ch. 4]), even in situations in which one would be extremely hard-pressed to identify an information-based reason for doing so²⁷

Information-based relative consumption effects induce pressure for people to conform to the consumption decisions of others. Low productivity agents will strive to increase consumption, while the high productivity agents will attenuate their consumption in order to not be too conspicuously different. The latter finding contrasts with the popular view of relative consumption effects as creating incessant incentives to consume more. Attempts to "keep up with the Joneses" appear to be ubiquitous, but no one conspicuously strives to "keep down with the Joneses." Does this suggest that status

²⁷Among the most striking is the classic work of Asch [2, 3], in which an apparent desire to conform prompted experimental subjects to make obviously incorrect choices when matching the lengths of lines, while denying that they were influenced by the choices of others. A committed Bayesian could still rationalize this behavior with a pure information-processing story, centered around subjects' uncertainty as to the consequences of the various choices they faced. But a pure information processing model is not useful if it can be preserved only with the help of complicated updating models that depend too finely on the circumstances of the choice.

is more important than information as a source of relative consumption effects?

Two responses are in order. First, it is likely that the observations which motivate relative consumption effects are stratified, with people influenced more by the consumption of others who appear to be "like them" than people whose lifestyles are quite different. We may be unfazed by comparisons with internet billionaires, but may be much more conscious of how our consumption compares with that of our colleagues. Similarly, people with high consumption levels may then compare themselves primarily to people in similar circumstances, freeing the internet billionaire from desperate attempts to reduce consumption and opening the possibility of an upward bias in consumption even at the high end of the consumption scale.²⁸

In addition, the concept of likeness on which such stratification is based is both endogenous and liable to manipulation. Much like Sherlock Holmes' dog that didn't bark in the night, it is easier to observe what someone has than what they don't have, making high consumption levels inherently more obvious than low levels. More importantly, the advent of modern advertising and mass communication may especially contribute to this asymmetry.²⁹ Relative consumption effects may then push primarily upwards.

Second, and more importantly, information-based relative consumption effects do not imply that we should observe people anxious to reduce consumption levels to those of their peers (however broadly or narrowly defined). Instead, individual consumption patterns cannot distinguish information-based and status-based effects. The utility function given by (17), incorporating information-based effects, is concave in consumption z and has the comparative-static implication

$$\frac{dz^*(z; z_j^R)}{dz_j} > 0;$$

so that higher values of observed consumption induce higher consumption. But the latter is what one means by "keeping up with the Joneses," and these are precisely the properties one expects from a utility function with status-based relative consumption effects. Information-based relative consumption effects imply not that we must observe people trying to keep down

²⁸Notice that Nature has incentives to encourage such stratification, since information about the environment is more relevant when it comes from observing those in similar circumstances. At the same time, Nature must beware of building too much taste for stratification into our comparisons, since excessively narrow comparison groups may ensure that virtually no information gets transferred.

²⁹Aronson [1, Chapter 3] discusses the persuasive ability of modern communications.

with the Joneses, but rather that people whose characteristics lead to high consumption levels should strive less vigorously to keep ahead of the Joneses than they would to catch up if the Joneses were ahead. This ability to coast while ahead is reflected in the longstanding notion of the idle rich. "...we must make allowances for the rich men's failings, and recollect that we, too, were very likely indolent and voluptuous, had we ... the daily temptation of a large income."³⁰ The salient observation is not that wealth may induce decreased effort, but rather that such effects should continue to characterize the relatively wealthy even as absolute wealth levels change.³¹

An ability to observe the qualitative features of utility functions is thus insufficient to distinguish the two models. Suppose instead we could observe utility functions exactly. In the case of (17), we would observe

$$z(1 + 2^{\alpha} m(z))_i z^{\mu} + \frac{1}{2^2} i^{\alpha} m(z)^2: \quad (25)$$

One could rearrange this function as in (17) and conclude that information-based effects are at work. But one could also interpret this as a version of (23), in which the rank of consumption level z is given by $z + 2^{\alpha} z m(z)_i m(z)^2$ and conclude that status-based effects are at work. A gain, individual-level data may be insufficient to distinguish the two.

Sharp differences in information-based and status-based relative consumption effects appear in the restrictions they place on the collection of consumption levels and utility functions in the economy (rather than on any particular utility function). These are most obviously reflected in the contrasting comparative static implications sketched in Sections 3.3 and 4.

Distinguishing information-based and status-based relative consumption effects thus promises to be a formidable task. But given the sensitivity of economic policy outcomes to the nature of relative consumption effects, it is clearly an important task. Further theoretical modeling is likely to be useful in this respect, as are carefully designed experiments.

³⁰William Madepeace Thackeray, "George the Third," in *The Four Georges and The English Humorists* (London, Collins' Clear-Type Press 1910 (originally 1855)).

³¹Two centuries before Thackeray, Robert Burton would comment that "Idleness is an appendix to nobility" (*Anatomy of Melancholy* (East Lansing, MI, Michigan State University Press 1965 (originally 1621), part 1, section 2, member 2, subsection 6), while nearly a century later Andrew Carnegie would view the disincentive effects of a fortune as sufficiently deleterious as to comment that "I would as soon leave my son a curse as the almighty dollar" (quoted in Burton J. Hendrick, *Life of Andrew Carnegie* (Garden City, NJ, Doubleday, Doran and Company, 1932), volume 1, chapter 17).

6 Appendix: Proofs

6.1 Proof of Lemma 1

6.1.1 Survival Probabilities

We first calculate $a_{\varepsilon}(\bar{A}_t; \bar{\mu})$. Fix an admissible $E = (z^2(n, \gg); \dots; z^2(0, \gg))g$, let the state be $\bar{\mu}$, and fix the proportion \bar{A}_t . Then proportion $r(k; \bar{A}_t)$ of the agents in period t will observe k values of Z , where

$$r(k; \bar{A}_t) = \binom{\bar{A}_t}{k} \bar{A}_t^k (1 - \bar{A}_t)^{n-k}.$$

These agents will choose $z^2(k, \gg)$ as their decision auto. Hence, proportion

$$S(z^2(k, \gg); z; \bar{\mu}) = \int_{-\infty}^{\infty} \int_{\epsilon(k, \xi)}^{Z_{\infty} - Z_1} \left(\frac{1}{2} + b(z^2(k, \gg) - q)\right) g(\gg; \bar{\mu}) d\phi d\psi$$

of these agents will choose strategy Z and survive, while proportion

$$S(z^2(k, \gg); z; \bar{\mu}) = \int_{-\infty}^{\infty} \int_0^{Z_{\infty} - Z_{\epsilon(k)}} \frac{1}{2} g(\gg; \bar{\mu}) d\phi d\psi$$

will choose z and survive. We then have

$$a_{\varepsilon}(\bar{A}_t; \bar{\mu}) = \frac{\sum_{k=0}^n r(k; \bar{A}_t) S(z^2(k, \gg); z; \bar{\mu})}{\sum_{k=0}^n r(k; \bar{A}_t) S(z^2(k, \gg); z; \bar{\mu}) + \sum_{k=0}^n r(k; \bar{A}_t) S(z^2(k, \gg); z; \bar{\mu})}.$$

Evaluating this equation at $\bar{A}_t = 0$ and $\bar{A}_t = 1$ gives

$$\begin{aligned} a_{\varepsilon}(0; \bar{\mu}) &> 0 \\ a_{\varepsilon}(1; \bar{\mu}) &< 1; \end{aligned}$$

and hence that probabilities remain within the unit interval. We next establish that $a_{\varepsilon}(\bar{A}_t; \bar{\mu})$ is increasing in \bar{A}_t . To do so, we note that

$$\frac{S(z^2(k); z; \bar{\mu})}{S(z^2(k); z; \bar{\mu}) + S(z^2(k); z; \bar{\mu})}$$

is increasing in k , since more agents choose Z as k increases, increasing $S(z^2(k); z; \bar{\mu})$ and decreasing $S(z^2(k); z; \bar{\mu})$. The result that $a_{\varepsilon}(\bar{A}_t; \bar{\mu})$ is increasing in \bar{A}_t then follows from noting that the distribution $r(k; \bar{A}_t)$ increases, in the sense of first-order stochastic dominance, as \bar{A}_t increases. We illustrate $a_{\varepsilon}(\bar{A}_t; \bar{\mu})$ in Figure 5.




Figure 5: Functions ${}^a \varepsilon(\bar{A}_t; \bar{\mu})$ and ${}^a \varepsilon(\bar{A}_t; \underline{\mu})$

Let $\{q_i\}_{i=1}^{\infty}$ be a sequence of values of q and let $E(q)$ be admissible given q . Then as $q \rightarrow \frac{1}{2}$, the sequence of functions ${}^a \varepsilon(q)(\bar{A}_t; \bar{\mu})$ converges uniformly (over admissible E and $\bar{A}_t \in [0; 1]$) to a constant function. Hence, there exists a value $q^* \in [0; \frac{1}{2})$ such that for all $q \in [q^*; \frac{1}{2}]$, the slope of ${}^a \varepsilon(\bar{A}_t; \bar{\mu})$ does not exceed $\frac{1}{2}$.

We now let the state be $\underline{\mu}$ and perform a similar examination of the function ${}^a \varepsilon(\bar{A}_t; \underline{\mu})$. Here, we have

$${}^a \varepsilon(\bar{A}_t; \underline{\mu}) = \frac{\mathbf{P} \sum_{k=0}^n r(k; \bar{A}_t) S(\mathcal{Z}(k); \underline{z}; \underline{\mu})}{\mathbf{P} \sum_{k=0}^n r(k; \bar{A}_t) S(\mathcal{Z}(k); \underline{z}; \underline{\mu}) + \mathbf{P} \sum_{k=0}^n r(k; \bar{A}_t) S(\mathcal{Z}(k); \underline{z}; \underline{\mu})};$$

where

$$S(\mathcal{Z}(k); \underline{z}; \underline{\mu}) = \int_{-\infty}^{\mathbf{Z}_{\infty}} \int_{\mathbf{Z}_{\epsilon(k)}}^{\mathbf{Z}_1} (\frac{1}{2} + b^2 i (1 - q)) g(\mathcal{Z}; \underline{\mu}) d\mathcal{Z} db$$

$$S(\mathcal{Z}(k); \underline{z}; \underline{\mu}) = \int_{-\infty}^{\mathbf{Z}_{\infty}} \int_{\mathbf{Z}_{\epsilon(k)}}^{\mathbf{Z}_{\epsilon(k)}} \frac{1}{2} g(\mathcal{Z}; \underline{\mu}) d\mathcal{Z} db;$$

Once again, ${}^a \varepsilon(\bar{A}_t; \underline{\mu})$ is strictly contained within the unit interval. The fraction

$$\frac{S(\mathcal{Z}(k); \underline{z}; \underline{\mu})}{S(\mathcal{Z}(k); \underline{z}; \underline{\mu}) + S(\mathcal{Z}(k); \underline{z}; \underline{\mu})}$$

is increasing in k , as a smaller value of $z^2(k)$ induces more agents to choose z (increasing $S(z^2(k); z; \underline{\mu})$ and decreasing $S(z^2(k); z; \bar{\mu})$), and hence ${}^a \varepsilon(\bar{A}_t; \bar{\mu})$ is increasing in \bar{A}_t . An argument analogous to that for the case of ${}^a \varepsilon(\bar{A}_t; \bar{\mu})$ establishes that there exists q^* such that for all $q \geq q^*$ ($q^* \geq \frac{1}{2}$), both ${}^a \varepsilon(\bar{A}_t; \bar{\mu})$ and ${}^a \varepsilon(\bar{A}_t; \underline{\mu})$ have slopes less than $\frac{1}{2}$. We hereafter assume that $q \geq q^*$. Finally, we note that, for all k ,

$$\frac{S(z^2(k); z; \bar{\mu})}{S(z^2(k); z; \bar{\mu}) + S(z^2(k); z; \underline{\mu})} > \frac{S(z^2(k); z; \underline{\mu})}{S(z^2(k); z; \underline{\mu}) + S(z^2(k); z; \bar{\mu})}$$

and hence

$${}^a \varepsilon(\bar{A}_t; \bar{\mu}) > {}^a \varepsilon(\bar{A}_t; \underline{\mu}):$$

Figure 5 illustrates both functions.

6.1.2 An Algebra of Sets

Let π^t denote the process generated by the transition rule $({}^a \varepsilon; E)$. Let $\pi^t((\bar{A}_t; \mu_t); (\bar{A}_0; \mu_0))$ be the probability measure over the time- t state $(\bar{A}_t; \mu_t)$ given initial condition $(\bar{A}_0; \mu_0)$. We construct a measure π^{1*} and show π^t converges to π^{1*} as t gets large.

Because q has been chosen so that $q > q^*$, and hence the slopes of ${}^a \varepsilon(\bar{A}_t; \bar{\mu})$ and ${}^a \varepsilon(\bar{A}_t; \underline{\mu})$ fall short of $\frac{1}{2}$, it must be that

$${}^a \varepsilon(\bar{A}; \bar{\mu}) > {}^a \varepsilon(\bar{A}; \underline{\mu}):$$

Figure 6 illustrates this property, while footnote 32 describes how it is used.

We now define a pair of countably infinite collections of sets,

$$\{\bar{A}_k\}_{k=1}^{\infty}; \quad \{\underline{A}_k\}_{k=1}^{\infty}; \quad (26)$$

where each of the sets \bar{A}_k or \underline{A}_k is a subset of the state space $[\underline{A}; \bar{A}] \in \mathcal{F}; \underline{\mu}; \bar{\mu}$. We will adopt the convention that for any k , $\underline{A}_k \cap \bar{A}_k = \emptyset$ and $\bar{A}_k \cap \underline{A}_k = \emptyset$, and hereafter suppress the notation for the values of μ when specifying a set \bar{A}_k or \underline{A}_k . We also drop the time subscripts and drop the subscript E on the function ${}^a \varepsilon$.

The collections $\{\bar{A}_k\}_{k=1}^{\infty}$ and $\{\underline{A}_k\}_{k=1}^{\infty}$ are defined recursively. First, we define two infinite collections of the level-one sets by letting

$$\begin{aligned} \bar{A}_{11} &= \bar{A}; \quad {}^a \varepsilon(\bar{A}; \bar{\mu}) \\ \bar{A}_{1i+1} &= {}^a \varepsilon(\bar{A}_{1i}; \bar{\mu}); \end{aligned}$$




Figure 6: Functions $f^a(\bar{A}_t; z)$ and $f^a(\underline{A}_t; z)$, with initial first-level sets

and

$$\begin{aligned} \underline{A}_{11} &= (f^a(\bar{A}; \underline{\mu}); \bar{A}) \\ \underline{A}_{1i+1} &= f^a(\underline{A}_{1i}; \underline{\mu}) \end{aligned}$$

Figure 6 illustrates the initial steps in this construction. Notice that the collections $\{f^a(\bar{A}; \underline{\mu})\}_{i \in \mathbb{N}}$ and $\{f^a(\underline{A}; \bar{\mu})\}_{i \in \mathbb{N}}$ are disjoint and that the union of each collection equals $[\underline{A}; \bar{A}]$, except for a countable collection of endpoints that have been omitted. As we define each level of sets, we shall similarly exclude a countable collection of endpoints. Let M denote the union of these collections of excluded sets and note that M will itself be countable. It will be clear from the construction that the sets M and $[0; 1] \setminus M$ are both forward invariant under the Markov process (i.e., it never leaves M if it begins in M , and that no path originating outside M can enter M).

We now define the level-two sets. For each first-level set \underline{A}_{1i} or \bar{A}_{1i} , we define a countably infinite collection of second-level sets, given by

$$\begin{aligned} \bar{A}_{21} &= f^a(\underline{A}_{1i}; \bar{\mu}) \in f^a(\bar{\mu}) \\ \bar{A}_{2j+1} &= f^a(\bar{A}_{2j}; \bar{\mu}) \\ \underline{A}_{21} &= f^a(\bar{A}_{1i}; \underline{\mu}) \in f^a(\underline{\mu}) \end{aligned}$$

$$\underline{A}_{2j+1} = \text{a}(\underline{A}_{2j}; \underline{\mu}):$$

We continue in this fashion, defining for each level- $(n-1)$ set $\underline{A}_{(n-1)i}$ or $\overline{A}_{(n-1)i}$, a countably infinite collection of level- n sets, given by

$$\begin{aligned} \overline{A}_{n1} &= \text{fa}(\underline{A}_{(n-1)i}; \underline{\mu})g \in \overline{f\bar{\mu}g} \\ \overline{A}_{nj+1} &= \text{a}(\overline{A}_{nj}; \overline{\mu}) \end{aligned}$$

$$\begin{aligned} \underline{A}_{n1} &= \text{fa}(\overline{A}_{(n-1)i}; \overline{\mu})g \in \underline{f\mu g} \\ \underline{A}_{nj+1} &= \text{a}(\underline{A}_{nj}; \underline{\mu}): \end{aligned}$$

We thus have a countably infinite collection of levels $\{1; 2; \dots\}$, with the sets defined at each level $n > 1$ containing a countably infinite collection of sets for each of the sets occurring at a previous level, of which there are countably many. The entire collection of such sets is thus countably infinite, and the sets can be placed in an order captured by the index k of (26).

Let \mathfrak{A}_i denote the collection of level- i sets, and let \mathfrak{A} denote the collection of all such sets along with the empty set, or $\mathfrak{A} = [\cup_{i=1}^{\infty} \mathfrak{A}_i; \emptyset; f; g]$. Let $\overline{\mathfrak{A}}$ be the collection of countable unions of disjoint sets in \mathfrak{A} . Then we note that $\overline{\mathfrak{A}}$ is an algebra on $[\mathfrak{A} = [0; 1]] \in \overline{f\bar{\mu}; \mu}g \mathbb{M}$.³²

6.1.3 Candidate Limiting Measure^{1*}

We now define a sequence of measures $\{f^k g_{k=1}^{\infty}$.

To define the first-level measure $f^1 g_1$, consider a hypothetical Markov process π_1 , in which the states are taken to be the first-level sets, so that the state space is \mathfrak{A}_1 (where each set is taken to be a state, rather than the union of the sets taken to be the state space). Let the transition probability between sets A and A' , under process π_1 , be the probability that the process

³² An algebra on a set A is a collection of subsets of A that includes A and the empty set and that is closed under the taking of complements and finite unions. Our assumption that the slopes of $\Psi(\psi_t, \bar{\theta})$ and $\Psi(\psi_t, \underline{\theta})$ are bounded below one-half, implied by $q > q^*$, makes its first appearance here. By ensuring that $\Psi_{\mathcal{E}}(\psi, \bar{\theta}) > \Psi_{\mathcal{E}}(\bar{\psi}, \underline{\theta})$, this assumption ensures that each level- i set is contained in a unique level- j set for every $j < i$ (rather than intersecting more than one level- j set) and that each level- j set is a countable union of level- i sets for $j < i$. This allows us to conclude that $\overline{\mathfrak{A}}$ is an algebra, with the complement of any finite union of elements in $\overline{\mathfrak{A}}$ being a countable union of elements in $\overline{\mathfrak{A}}$. In addition, it simplifies subsequent calculations. It allows us to recursively calculate the measures μ_i in Section 6.1.3 and verify in each case that the resulting measures are consistent, in that the level- i measure of a set A is the sum of the level- k measures of its subsets, for any $k > i$. It also simplifies the calculation of the stationary distribution in Section 6.1.6.

... moves from A to A' . Our sets are constructed so that this probability is independent of the particular state in A in which the process ... finds itself, and hence this transition probability is well defined. The Markov process ... has a countably infinite state space, and is an aperiodic, irreducible, positive recurrent process. It is accordingly ergodic, with a unique stationary distribution, which we denote by μ_1 . We can view μ_1 as a measure on the countable set \mathfrak{A}_1 , endowed with the discrete \mathfrak{A} -algebra. We can also view μ_1 as a measure on the space $[0; 1] \in \mathfrak{F}; \bar{\mu}g$, with the \mathfrak{A} -algebra generated by the members of \mathfrak{A}_1 , each interpreted as a subset of $[0; 1] \in \mathfrak{F}; \bar{\mu}g$.

At the first level, it is straightforward to calculate the stationary distribution μ_1 . The measure μ_1 is given by

$$\begin{aligned} \mu_1(\bar{A}_{11}) &= \frac{1}{2}c \\ \mu_1(\bar{A}_{1i+1}) &= (1-i-c)\mu_1(\bar{A}_{1i}) \\ \mu_1(A_{11}) &= \frac{1}{2}c \\ \mu_1(A_{1i+1}) &= (1-i-c)\mu_1(A_{1i}); \end{aligned}$$

Similarly, we can define a Markov process ... whose state space is the collection of second-level sets \mathfrak{A}_2 . This Markov process is again aperiodic, irreducible, and positive recurrent. It is accordingly ergodic, with a unique stationary distribution which we call μ_2 , which we can interpret as a measure on the discrete \mathfrak{A} -algebra on the elements of \mathfrak{A}_2 , or as a measure on $[0; 1] \in \mathfrak{F}; \bar{\mu}g$ with the \mathfrak{A} -algebra generated by the subsets of $[0; 1] \in \mathfrak{F}; \bar{\mu}g$ contained in \mathfrak{A}_2 . In addition, every level-one set is a countable union of level-two sets (minus a countable collection of endpoints), with a measure μ_1 that is the sum of the μ_2 -measures of the level-two sets in contains. Hence, the \mathfrak{A} -algebra generated by \mathfrak{A}_2 matches that generated by \mathfrak{A}_1 [\mathfrak{A}_2 , and μ_1 and μ_2 agree on any set from this \mathfrak{A} -algebra for which μ_1 is defined.

Continuing in this way, we define the entire collection

$$\mathfrak{F}^1_k \mathfrak{G}_{k=1}^\infty;$$

where each μ_k is defined on the \mathfrak{A} -algebra generated by $\bigcup_{i=1}^k \mathfrak{A}_i$.

For any $A \in \mathfrak{A}$ and μ_k , either A is not contained in \mathfrak{A}_i for any $i \leq k$, in which case $\mu_k(A)$ is undefined, or A is contained in some \mathfrak{A}_i with $i \leq k$, in which case $\mu_k(A)$ is defined and is equal to $\mu_{k'}(A)$ for any $k' \geq i$. We can then define μ^* on \mathfrak{A} by taking $\mu^*(A) = \mu_k(A)$ for some k for which $\mu_k(A)$ is defined.

We can think of μ^* as being the limit of the sequence $\mathfrak{F}^1_k \mathfrak{G}_{k=1}^\infty$. However, instead of a sequence of changing measures on a fixed domain, we have a

sequence of measures on an ever-expanding sequence of domains, with the measures being constant along the sequence, in the sense that once a set enters the domain, its measure remains fixed along the remainder of the sequence.

6.1.4 μ^* is a measure

Intuitively, our next step is to show that μ^* is a measure. More precisely, we embed \mathfrak{A} in a \mathfrak{B} -algebra and then show that μ^* extends uniquely to a measure on that \mathfrak{B} -algebra. Recall that $\bar{\mathfrak{A}}$ is the collection of countable unions of disjoint sets in \mathfrak{A} , and is an algebra on Ω . We then proceed in three steps.

First, we extend μ^* from \mathfrak{A} to $\bar{\mathfrak{A}}$. For each set A in $\bar{\mathfrak{A}}$, we take a minimal representation of A to be the smallest collection of sets in \mathfrak{A} whose union is A . Because of the recursive structure of $\bar{\mathfrak{A}}$, minimal representations are uniquely defined and consist of disjoint sets. Then $\mu^*(A)$ is defined for any A in the minimal representation of A , and we take $\mu^*(A)$ to be the sum of the measures $\mu^*(A')$ of the sets in the minimal representation. It is then straightforward that μ^* is finite on $\bar{\mathfrak{A}}$ (being bounded above by unity) and is finitely additive.

Second, we note that μ^* is countably subadditive on $\bar{\mathfrak{A}}$. Suppose not, so we can find a case in which

$$\left[\bigcup_{i=1}^{\infty} A_i = A; \quad \sum_{i=1}^{\infty} \mu^*(A_i) > \mu^*(A); \right]$$

where the union is disjoint. Because $\mu^*(A)$ is finite, there must be a finite n such that

$$\sum_{i=1}^n \mu^*(A_i) > \mu^*(A);$$

Because μ^* is finitely additive, we then have

$$\mu^*\left(\bigcup_{i=1}^n A_i\right) > \mu^*(A); \quad \bigcup_{i=1}^n A_i \subsetneq A;$$

with strict set inclusion in the second case, which is precluded by the construction of μ^* .

Third, we argue that μ^* is continuous from above at the empty set (see and Doleans-Dade [4, p. 10]), which in turn implies that it is countably additive ([4, Theorem 1.2.8]). To verify this, let $\{B_n\}_{n=0}^{\infty}$ be a sequence of sets in $\bar{\mathfrak{A}}$ with

$$B_n \supset B_{n+1}; \quad \bigcap_{n=0}^{\infty} B_n = \emptyset;$$

We need to show that

$$1^*(B_n) = 0;$$

We can take the sequence B_n to be decreasing without sacrificing any generality. (If necessary, redefine each B_n to be the union $\bigcap_{k=n}^{\infty} B_k$.) Given that $\bar{\alpha}$ consists of countable unions of sets in α , there must be a sequence of disjoint sets $\{C_m\}_{m=0}^{\infty}$ with each C_m in $\bar{\alpha}$ and, for each B_n , a sequence of numbers k_{ni} , $i = 1, \dots, n$, such that

$$B_n = \bigcup_{i=1}^{\infty} C_{k_{ni}};$$

for each set B_n . Because 1^* is countably subadditive, we have

$$\sum_{m=1}^{\infty} 1^*(C_m) < 1;$$

and hence,

$$\lim_{k \rightarrow \infty} \sum_{m=k}^{\infty} 1^*(C_m) = 0;$$

but for every k , there is an $h(k)$ such that

$$B_n = \bigcup_{i=1}^{\infty} C_{k_{ni}} \cap \bigcup_{m=h(k)}^{\infty} C_m;$$

and the fact that $B_n \downarrow \emptyset$ implies that $h(k) \rightarrow \infty$, giving $1^*(B_n) = 0$, as required.

Fifth, we have now shown that 1^* is a countably additive, positive set function on the algebra $\bar{\alpha}$. The Carathéodory-Halm Extension Theorem (Wheeden and Zygmund [42, p. 206]) then ensures that 1^* extends uniquely to a measure that is defined on the σ -algebra $\mathcal{M}(\bar{\alpha})$ generated by $\bar{\alpha}$ and that agrees with 1^* on α . Hence, 1^* is a measure (on $\mathcal{M}(\bar{\alpha})$).

6.1.5 Ergodicity

We now establish the required limiting property. We construct a pair of sequences of Markov processes $\{f_{-i}^{\leftarrow} g_{i=1}^{\infty}\}$ and $\{f_{-i}^{\rightarrow} g_{i=1}^{\infty}\}$ from \mathbb{P} by replacing $(\bar{A}_t; \mu_t)$ in each sample path generated by \mathbb{P} with $(\bar{A}; \mu_t)$, where \bar{A} is given by

$$\begin{aligned} \min\{\bar{A} \geq \bar{A}_{ij}\} & \quad \text{if } (\bar{A}_t; \mu_t) \geq \bar{A}_{ij} \\ \max\{\bar{A} \leq \bar{A}_{ij}\} & \quad \text{if } (\bar{A}_t; \mu_t) \leq \bar{A}_{ij} \end{aligned}$$

in the case of f_{-i}^{∞} and

$$\begin{aligned} \max \bar{A} & \geq \bar{A}_{ij} & \text{if } (\bar{A}_t; \mu_t) \geq \bar{A}_{ij} \\ \min \bar{A} & \geq \bar{A}_{ij} & \text{if } (\bar{A}_t; \mu_t) \geq \bar{A}_{ij} \end{aligned}$$

in the case of f_{-i}^{∞} . The processes \bar{x}_i and \bar{y}_i are the counterparts of the Markov process \bar{x}_i , constructed in Section 6.1.3, obtained by having each of \bar{x}_i and \bar{y}_i assign a value of \bar{A} to every element in the state space of \bar{x}_i .

Letting $\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i)$ and $\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i)$ be the counterparts of $h_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i)$ for the processes \bar{x}_i and \bar{y}_i , we have

$$\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i) \cdot \bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i) \cdot \bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i):$$

We obtain this result by noting that the process \bar{x}_i chooses the minimum (maximum) possible value of \bar{A} (consistent with staying within the element of \mathcal{X}_i occupied by the process \bar{x}_i) whenever the state is $\bar{\mu}$ ($\bar{\mu}$), and hence minimizes the extent to which an observation of Z signals state $\bar{\mu}$. The process \bar{y}_i does just the opposite, and hence maximizes the extent to which Z indicates that the state is $\bar{\mu}$.

In addition, the processes \bar{x}_i and \bar{y}_i inherit the ergodicity of \bar{x}_i . Hence, $\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i)$ and $\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i)$ converge (as $\tau \rightarrow \infty$) to limits

$$\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; i) \cdot \bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; i):$$

Furthermore, the former sequence is increasing in i and the latter sequence decreasing in i , with the former approaching the latter as $i \rightarrow \infty$. The probability $\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i)$, being bounded by $\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i)$ and $\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i)$, must accordingly converge to this common limit. This allows us to establish the desired conclusion, namely that there exist $\bar{h}_{\mathcal{E}}(\bar{\mu}; j, k)$, $k = 0, \dots, n$, such that

$$\lim_{\tau \rightarrow \infty} \bar{h}_{\mathcal{E}}(\bar{\mu}; j, k; \tau; i) = \bar{h}_{\mathcal{E}}(\bar{\mu}; j, k):$$

This establishes the result for any initial condition not contained in the set \mathcal{M} . To extend the result to initial conditions in \mathcal{M} , we note that because $\bar{h}_{\mathcal{E}}(\bar{A}_t; \bar{\mu})$ and $\bar{h}_{\mathcal{E}}(\bar{A}_t; \bar{\mu})$ have slopes bounded below $\frac{1}{2}$, we have, for any fixed sequence of realizations of values of μ and any two initial conditions \bar{A} and \bar{A}' , that

$$|\bar{A}_{t+1} - \bar{A}'_{t+1}| \leq \frac{1}{2} |\bar{A}_t - \bar{A}'_t|$$

Hence, for any realization of the sequence of random variables μ_t , the sample paths of realized values of \bar{A}_t generated by any two initial conditions converge

at an exponential rate.³³ Furthermore, the sample paths of realized values of \bar{A}_t have the monotonicity property that $\bar{A}_0 < \bar{A}'_0 \Rightarrow \bar{A}_t < \bar{A}'_t$ for all $t > 0$. This in turn ensures that the limiting probability $\frac{1}{2} \varepsilon(\bar{\mu}; k; \gg)$ given an initial state in \mathbb{M} lies between the limiting probabilities for some pair of larger and smaller initial states not contained in \mathbb{M} , and hence must converge to the common limit attained by the latter pair.

6.1.6 $\frac{1}{2} \varepsilon(\bar{\mu}; k; \gg)$ is increasing in \gg with $\frac{1}{2} \varepsilon(\bar{\mu}; k+1; \gg) > \frac{1}{2} \varepsilon(\bar{\mu}; k; \gg)$

That $\frac{1}{2} \varepsilon(\bar{\mu}; k; \gg)$ is increasing in \gg follows from (8) and the monotone likelihood ratio property. Next, to establish $\frac{1}{2} \varepsilon(\bar{\mu}; k+1; \gg) > \frac{1}{2} \varepsilon(\bar{\mu}; k; \gg)$, it suffices (from (8)) to show that

$$\int_{\psi=0}^{\mathbf{Z}_1} \bar{A} d\pi^*(\bar{A}; \bar{\mu}) > \int_{\psi=0}^{\mathbf{Z}_1} \bar{A} d\pi^*(\bar{A}; \underline{\mu}) \quad (27)$$

First, we recall that

$$\underline{\bar{A}} < {}^a \varepsilon(\bar{A}; \underline{\mu}) \cdot {}^a \varepsilon(\underline{\bar{A}}; \bar{\mu}) < \bar{A};$$

which we use to construct the intervals in the following expression. Next, notice that the construction of the measure π^* ensures, for any $A \in \mathcal{A}$

$$\begin{aligned} A \cap \frac{1}{2} [\underline{\bar{A}}; {}^a \varepsilon(\bar{A}; \underline{\mu})] & \cap \pi^*(A; \bar{\mu}) = \emptyset \\ A \cap \frac{1}{2} [{}^a \varepsilon(\bar{A}; \underline{\mu}); {}^a \varepsilon(\underline{\bar{A}}; \bar{\mu})] & \cap \pi^*(A; \bar{\mu}) = \pi^*(A; \underline{\mu}) \\ A \cap \frac{1}{2} [{}^a \varepsilon(\underline{\bar{A}}; \bar{\mu}); \bar{A}] & \cap \pi^*(A; \bar{\mu}) = \pi^*(A; \underline{\mu}) \end{aligned}$$

Since $\varepsilon < \frac{1}{2}$, this implies (27). jj

6.2 Proof of Proposition 1

Let \mathcal{E} be the set of admissible E , i.e., the set of n_j tuples of decreasing functions from \mathbb{R} into $[q_1; q_j]$, denoted by $\mathcal{F}^2(n; \gg); \dots; \mathcal{F}^2(n; \gg)$, with the property that $\mathcal{F}^2(k+1; \gg) \cdot \mathcal{F}^2(k; \gg)$. We construct a function $H(E)$, defined on \mathcal{E} . For any admissible collection E , calculate the probabilities $p_{\varepsilon}(\bar{\mu}; k)$ (whose existence is established in Lemma 1), insert these probabilities into the maximization problem given by (9)-(10), solve for equilibrium values $\mathcal{F}^2(k; \gg)$, and let these new values be $H(E)$.

We recall from Lemma 1 that, for any admissible collection, the posterior probability attached to $\bar{\mu}$ after an observation of $k+1$ values of Z must exceed

³³It would suffice for this last step that the slopes of $\Psi_{\varepsilon}(\psi_t, \bar{\theta})$ and $\Psi_{\varepsilon}(\psi_t, \underline{\theta})$ have an upper bound less than one.

that after an observation of k such values ($k = 0; \dots; n_j - 1$), conditional on \gg , and that this posterior is increasing in \gg . This ensures that the mapping we have constructed will always choose new cut-offs that are decreasing in \gg and satisfy $z^2(k+1; \gg) \cdot z^2(k; \gg)$, and hence the map is from the set \mathbb{Y} into itself.

We next note that, for fixed probabilities $p_{\mathcal{E}}(\bar{\mu} | k; \gg)$, the objective contained in (9)-(10) are strictly concave, and hence the maximizer $z^2(k; \gg)$ is unique, ensuring that our mapping is a function.

The set \mathbb{Y} is a convex subset of a locally convex linear topological space (the metric space of n_j -tuples of decreasing bounded (by $[q_1; q]$) functions, with the metric induced by the L^1 norm:

$$\|j\|_{\mathbb{Y}} = \int_{\mathbb{X}} \sum_{k=1}^{n_j} \int_{\gg=-\infty}^{\infty} |z^2(k; \gg) - z^2(k; \gg')| d\bar{\nu}$$

where $\bar{\nu}$ is the measure induced by the density $\frac{1}{2}g^2(\bar{\mu}; \mu) + \frac{1}{2}g^2(\bar{\mu}; \bar{\mu})$, which is the density of \gg conditional on μ being drawn from the stationary distribution.

The set \mathbb{Y} is also compact. In particular, a subset of a metric space is compact if it is closed and sequentially compact. The set \mathbb{Y} is clearly closed. To establish sequential compactness, fix a value of k and let $f_j^2; g_{j=1}^{\infty}$ be a sequence of decreasing functions $z_j^2(k; \gg)$. We show that there exists a converging subsequence. By successively taking subsequences, we can find a subsequence $f_j^2; g_{j=1}^{\infty}$ that converges at every rational value of \gg , and let z^* , defined only on the rationals, be the (pointwise) limit of this subsequence. Then z^* must be decreasing. Extend z^* to the reals in \mathbb{R} by letting (for $\gg^* \in \mathbb{R}$) $z^*(k; \gg^*) = \lim_{n \rightarrow \infty} z^*(k; \gg_n^*)$, where \gg_n^* is a sequence of rationals approaching \gg^* from below. Then $z^*(k; \gg)$ is decreasing.

We next note that the sequence of functions $z_j^2(k; \gg)$ must converge to $F^*(k; \gg)$ almost everywhere. In particular, $F_j^2(k; \gg)$ must converge to $F^*(k; \gg)$ at any value \gg at which $F^*(k; \gg)$ is continuous. Since $F^*(k; \gg)$ is decreasing it is continuous at all but countably infinitely many values, ensuring that the sequence $F_j^2(k; \gg)$ converges to $F^*(k; \gg)$ almost everywhere. Because each F_j^2 and F^* are bounded above by the $\bar{\nu}$ -integrable function $z^2(\gg) = [q_1; q]$, Lebesgue's dominated convergence theorem (Billingsley [8, Theorem 16.4]) ensures that $\|j\|_{\mathbb{Y}} \rightarrow 0$.

It then remains to verify the continuity of the function H . To do so, we must show that the probabilities $\frac{1}{2}p_{\mathcal{E}}(\bar{\mu} | k; \gg)$ are continuous in E , for which it suffices to note that the functions $\int_{\mathcal{E}} \bar{A}_t; \bar{\mu}$ and $\int_{\mathcal{E}} \bar{A}_t; \mu$ move continuously (in the sup norm) in E .

This gives a continuous mapping from a convex subset of a locally convex linear topological space into itself. The Schauder-Tychonoff theorem

(Dunford and Schwartz [17, p. 456]) ensures that this mapping has a fixed point E^* , and hence that an equilibrium exists.

It follows from the full-support assumptions and monotone-likelihood-ratio-property assumptions on $g(\cdot; \bar{\mu})$ and $g(\cdot; \underline{\mu})$ that $z^*(k; \cdot)$ is strictly decreasing in \cdot . It follows that $z^*(k+1; \cdot) = z^*(k; \cdot)$. In order for these values to be optimal, it must be that the $k+1$ st observation of Z does not change the likelihood that the state is $\bar{\mu}$. But this can occur only if no additional observations of Z are informative, i.e., only if $p_{\varepsilon^*}(\bar{\mu}|k+1; \cdot) = p_{\varepsilon^*}(\bar{\mu}|k; \cdot)$ for $k = 0, \dots, n_j - 1$. But then identical proportions of the agents choose Z and \bar{z} , regardless of their observations, ensuring that their exit numbers \bar{p} and $\underline{p} < \bar{p}$ such that proportion of surviving agents choosing Z is given by \bar{p} when the state is $\bar{\mu}$ and \underline{p} when the state is $\underline{\mu}$. But then an observation of Z yields a higher posterior probability of state $\bar{\mu}$, or $p_{\varepsilon^*}(\bar{\mu}|k+1; \cdot) > p_{\varepsilon^*}(\bar{\mu}|k; \cdot)$, a contradiction. jj

6.3 Proof of Proposition 2

First, fix an admissible E and hence information updating rules $\frac{1}{2}_{\varepsilon}(\bar{\mu}|k; \cdot)$ and $\frac{1}{2}_{\varepsilon}(\underline{\mu}|k; \cdot)$ ³. Consider an agent who has observed k values of Z and a signal $s \in \mathbb{R}$, where s may be a realization of either \bar{z} or \underline{z} . Let ${}_{\varepsilon}^{\circ}(k)$ be the posterior probability that the state is $\bar{\mu}$, given k observations of Z (but ignoring any other signals). Then the agent's posterior $\frac{1}{2}_{\varepsilon}(\bar{\mu}|k; \cdot)$ ³ is given by (8), which we can rewrite as:

$$\frac{1}{2}_{\varepsilon}(\bar{\mu}|k; s) = \frac{g(s; \bar{\mu}) {}_{\varepsilon}^{\circ}(k)}{g(s; \bar{\mu}) {}_{\varepsilon}^{\circ}(k) + g(s; \underline{\mu}) (1 - {}_{\varepsilon}^{\circ}(k))};$$

Moreover, recognizing the possibility that the signal s is an uninformative realization of the random variable \bar{z} ³, would instead assign posterior

$$\frac{1}{2}_{\varepsilon}^N(\bar{\mu}|k; s) = \frac{[(1 - \lambda)g(s; \bar{\mu}) + \lambda f(s)] {}_{\varepsilon}^{\circ}(k)}{[(1 - \lambda)g(s; \bar{\mu}) + \lambda f(s)] {}_{\varepsilon}^{\circ}(k) + [(1 - \lambda)g(s; \underline{\mu}) + \lambda f(s)] (1 - {}_{\varepsilon}^{\circ}(k))};$$

Dividing the numerator and denominator of this expression by $[(1 - \lambda)g(s; \bar{\mu}) + \lambda f(s)]$, we obtain an expression to which (4) and (12) can be applied to conclude that³⁴

$$0 < \frac{d\frac{1}{2}_{\varepsilon}^N(\bar{\mu}|k; s)}{ds} < \frac{d\frac{1}{2}_{\varepsilon}(\bar{\mu}|k; s)}{ds}; \quad (28)$$

³⁴In particular, (4) and (12) imply that the term $[(1 - \lambda)g(s, \underline{\theta}) + \lambda f(s)] / [(1 - \lambda)g(s, \bar{\theta}) + \lambda f(s)]$ is decreasing in s , but at a slower rate than $g(s, \underline{\theta}) / g(s, \bar{\theta})$, which suffices for (28).

Hence, higher signals increase the likelihood that the state is $\bar{\mu}$, but nature reacts less sharply to the agent's information than does the agent, reflecting nature's realization that the signal may be uninformative.

Condition (12) ensures that there are unique signals \underline{s} and $\bar{s} > \underline{s}$ such that

$$g(\underline{s}; \underline{\mu}) = f(\underline{s}); \quad g(\bar{s}; \bar{\mu}) = f(\bar{s});$$

This in turn ensures that there exists a value $\hat{s} \in [\underline{s}; \bar{s}]$ such that

$$\frac{1}{2} \frac{1}{\varepsilon} (\bar{\mu} | k; \hat{s}) = \frac{1}{2} \frac{1}{\varepsilon} (\underline{\mu} | k; \hat{s});$$

Let $z_{\varepsilon}^{2^*}(k; s)$ be the value of z^* that nature would like to induce, given that the agent has observed k values of Z and signal s . Let $z_{\varepsilon}^2(k; s)$ be the value of z^* that maximizes (13), and hence the value the agent will choose given utility function $\ln p(z^*; \underline{\mu})$. Nature would choose $z_{\varepsilon}^{2^*}(k; \frac{1}{2})$ to solve

$$\max_{z^*} \frac{1}{2} \frac{1}{\varepsilon} (\bar{\mu} | k; s) \ln p(z^*; \bar{\mu}) + (1 - \frac{1}{2} \frac{1}{\varepsilon} (\bar{\mu} | k; s)) \ln p(z^*; \underline{\mu}); \quad (29)$$

This is nature's counterpart of (7), given nature's expectation conditional on k observations of Z and the agent's signal s . The objective given in (29) is strictly concave in z^* and is continuous in s (because $\frac{1}{2} \frac{1}{\varepsilon} (\bar{\mu} | k; s)$ is) and hence the maximizer $z_{\varepsilon}^{2^*}(k; s)$ is continuous and strictly decreasing in s .

Condition (28) implies that if nature allows her agents to solve (13), then they will choose values of z^* that are too small (from her point of view) whenever $s > \hat{s}$ and too large when $s < \hat{s}$. Let the value of z^* chosen by the agent (and preferred by nature) when $s = \hat{s}$ be denoted $z^*(k)$ ($= z_{\varepsilon}^2(k; \hat{s}) = z_{\varepsilon}^{2^*}(k; \hat{s})$).

Notice next that, on the interval $[q_1; q]$, the derivatives $d \ln p(z^*; \bar{\mu}) = c^*$ and $d \ln p(z^*; \underline{\mu}) = c^*$ are both strictly decreasing with

$$\frac{d \ln p(q; \underline{\mu})}{c^*} = 0 \quad \frac{d \ln p(1 - q; \bar{\mu})}{c^*} = 0;$$

Now for each signal s , define $\phi^{\circ}(z_{\varepsilon}^{2^*}(k; s); k)$ to satisfy

$$\begin{aligned} \frac{1}{2} \frac{1}{\varepsilon} (\bar{\mu} | k; s) & \frac{d \ln p(z_{\varepsilon}^{2^*}(k; s); \bar{\mu})}{c^*} + \phi^{\circ}(z_{\varepsilon}^{2^*}(k; s); k) \\ & + (1 - \frac{1}{2} \frac{1}{\varepsilon} (\bar{\mu} | k; s)) \frac{d \ln p(z_{\varepsilon}^{2^*}(k; s); \underline{\mu})}{c^*} + \phi^{\circ}(z_{\varepsilon}^{2^*}(k; s); k) = 0; \quad (30) \end{aligned}$$

Because the left side of (30) is linear in ϕ° , each $\phi^{\circ}(z_{\varepsilon}^{2^*}(k; s); k)$ is well defined. Furthermore, because $z_{\varepsilon}^{2^*}(k; s)$ is decreasing in s , we can view ϕ° as a function

$\circ(z^*; k)$ defined on the set of all values of z^* that are Nature's optimal choice $z_{\varepsilon}^N(k; s)$ for k and some signal s . It follows from (30) that this function must be continuous in z^* . Then define a function $u'(z^*; k; \mu)$ by

$$u'(z^*; k; \bar{\mu}) = \frac{d \ln p(z^*; \bar{\mu})}{dz^*} + \circ(z^*; k)$$

$$u'(z^*; k; \underline{\mu}) = \frac{d \ln p(z^*; \underline{\mu})}{dz^*} + \circ(z^*; k):$$

We will then let the utility function $u(z^*; k; \mu)$ be the integral of $u'(z^*; k; \mu)$ with respect to z^* . (Integrability follows from the continuity of $\circ(z^*; k)$ in z^* .)

The agent's first-order necessary conditions for utility maximization, given utility function $u(z^*; k; \mu)$, are given by (30) and are solved by Nature's optimum $z_{\varepsilon}^N(k; s)$. The utility function $u(z^*; k; \mu)$ thus achieves Nature's optimum if the latter uniquely solves these conditions, which we show below.

Next, let $\frac{1}{2}^N(\bar{\mu}^{jz^*})$ be the value of the posterior $\frac{1}{2}$ at which Nature would optimally induce choice z^* (and hence that causes the agent to choose z^* , given utility function $u(z^*; k; \mu)$). Recall that $\hat{A}(\bar{\mu}^{jz^*})$ is the posterior that would cause the agent to choose z^* , given utility function $\ln p(z^*; \mu)$. Comparing Nature's optimization (29) with the agent's (11), condition (28) and the fact that $z(k) = z_{\varepsilon}^N(k; s) = z_{\varepsilon}^N(k; s)$ imply:

$$\frac{1}{2}^N(\bar{\mu}^{jz^*}) > (<) \hat{A}(\bar{\mu}^{jz^*}) \quad (z^* < (>) z(k):$$

From (30), this ensures that

$$\circ(z^*) > (<) 0 \quad (z^* < (>) z(k);$$

and hence that (14) holds.

For the agent's utility-maximization problem to implement Nature's optimum given utility function $u(z^*; k; \mu)$, we need that the agent's first-order condition have a unique solution, for each posterior probability of $\bar{\mu}$, within the range of values z^* that can arise as solutions to Nature's optimization problem. (Nature can ensure other values of z^* are not chosen by attaching large utility penalties to such values.) Suppose this is not the case, so that there are values z' and z'' that both solve the first-order condition for some posterior $\frac{1}{2}$, the former being Nature's optimum for this posterior. But then z'' must be Nature's optimum for some posterior $\frac{1}{2}'$, ensuring that z'' solves the first-order condition for posteriors $\frac{1}{2}$ and $\frac{1}{2}'$, a contradiction to the fact that, conditional on z^* , the first-order condition is linear in the posterior $\frac{1}{2}$.

It remains to show that an equilibrium E exists, given that Nature responds to E with an optimal utility function $u(z^*; k; \mu)$. Given the continuity of Nature's optimum in the signal s , this argument is a straightforward adaptation of Lemma 1 and Proposition 1. jj

References

- [1] Elliott Aronson. The Social Animal. W. H. Freeman and Company, New York, 1995.
- [2] Solomon E. Asch. Effects of group pressure on the modification and distortion of judgments. In Groups, Leadership and Men. Carnegie Press, Pittsburgh, 1951.
- [3] Solomon E. Asch. Studies of independence and conformity: A minority of one against a unanimous majority. Psychological Monographs, 70(416), 1956.
- [4] Robert B. Ash and Catherine A. Downs-Doyle. Probability and Measure Theory. Academic Press, New York, 2000.
- [5] Abhijit Banerjee. A simple model of herd behavior. Quarterly Journal of Economics, 107:797-817, 1992.
- [6] Abhijit Banerjee and Drew Fudenberg. Word-of-mouth learning. mimeo, MIT and Harvard University, 1995.
- [7] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural exchange as information cascades. Journal of Political Economy, 100:992-1026, 1992.
- [8] Patrick Billingsley. Probability and Measure. John Wiley and Sons, New York, 1986.
- [9] David G. Blanchard and Andrew Oswald. Well-being over time in Britain and the USA. mimeo, University of Warwick, 1999.
- [10] William A. Brock and Steven N. Durlauf. Interactions-based models. SSRN Working Paper 9910R, University of Wisconsin, 2000.
- [11] Brian Charlesworth. Evolution in Age-Structured Populations. Cambridge University Press, Cambridge, 1994.

- [12] Robert B. Cialdini. *Influence: Science and Practice*. Scott, Foresman and Company, Boston, 1988.
- [13] Andy Clark. *Microcognition*. MIT Press, Cambridge, Massachusetts, 1993.
- [14] Hal L. Cole, George J. Mailath, and Andrew Postlewaite. social norms, savings behavior, and growth. *Journal of Political Economy*, 100:1092{1125, 1992.
- [15] Giacomo Corneo. The efficient side of progressive income taxation. Working Paper 364, CESifo Munich, Germany, 2000.
- [16] William Michael Cox and Richard Alm. *Myths of Rich and Poor*. Basic Books, New York, 1999.
- [17] Nelson Dunford and Jack T. Schwartz. *Linear Operators Part I: General Theory*. John Wiley and Sons, New York, 1988. Wiley Classics Library Edition.
- [18] Steven M. Durlauf. A framework for the study of individual behavior and social interactions. SSRN Working Paper 2016, University of Wisconsin, 2000.
- [19] J. S. Dusenberry. *Income, Saving and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, Massachusetts, 1949.
- [20] R. A. Easterlin. Does economic growth improve the human lot? Some empirical evidence. In P. David and M. Reder, editors, *Nations and Households in Economic Growth: Essays in Honour of Moses Abramovitz*. Academic Press, New York, 1974.
- [21] R. A. Easterlin. Will raising the incomes of all increase the happiness of all? *Journal of Economic Behavior and Organization*, 27:35{47, 1994.
- [22] Glenn Ellison and Drew Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 101:612{643, 1993.
- [23] Glenn Ellison and Drew Fudenberg. Word-of-mouth communication and social learning. *Quarterly Journal of Economics*, 110:95{126, 1995.
- [24] Robert H. Frank. *Luxury Fever*. Free Press, New York, 1999.
- [25] Carl A. Futia. Invariant distributions and the limiting behavior of Markovian economic models. *Econometrica*, 50:377{408, 1982.

- [26] Fred Hirsch. *Social Limits to Growth*. Harvard University Press, Cambridge, 1976.
- [27] Ed Hopkins and Tatiana Komienko. *Consumer choice as a game of status*. mimeo, University of Edinburgh, 2000.
- [28] R. Inglehart. The diminishing utility of economic growth: From maximizing security toward maximizing subjective well-being. *Critical Review*, 10:509{531, 1996.
- [29] Daniel Kahneman and Amos Tversky. Judgment under uncertainty: Heuristics and biases. In Daniel Kahneman, Paul Slovic, and Amos Tversky, editors, *On the Psychology of Prediction*, pages 48{68. Cambridge University Press, Cambridge, 1982.
- [30] Lawrence Katz and David Autor. Changes in the wage structure and earnings inequality. In Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, Volume 3A, pages 1463{1555. North Holland, New York, 1999.
- [31] Joseph LeDoux. *The Emotional Brain*. Simon and Schuster, New York, 1996.
- [32] Lars Ljungqvist and Harald Uhlig. Tax policy and aggregate demand management under catching up with the joneses. *American Economic Review*, 90:356{366, 2000.
- [33] Paul R. Milgrom. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12:380{391, 1981.
- [34] David Neumark and Andrew Postlewaite. Relative income concerns and the rise in married women's employment. *Journal of Public Economics*, 70:157{183, 1998.
- [35] Andrew J. Oswald. Happiness and economic performance. *Economic Journal*, 107:1815{1831, 1997.
- [36] Robert A. Pollak. Interdependent preferences. *American Economic Review*, 66:309{320, 1976.
- [37] Andrew Postlewaite. The social basis of interdependent preferences. *European Economic Review*, 42:779{800, 1998.
- [38] Arthur J. Robson. A biological basis for expected and non-expected utility. *Journal of Economic Theory*, 68:397{424, 1996.

- [39] Larry Samuelson and Jeroen Swinkels. Information and the evolution of the utility function. mimeo, University of Wisconsin and Washington University in St. Louis, 2000.
- [40] S. Tuljapurkar. Population Dynamics in Variable Environments. Springer-Verlag Berlin, 1990.
- [41] R. Veenhoven. Is happiness a trait? Social Indicators Research, 32:101-160, 1994.
- [42] Richard L. Wheeden and Antoni Zygmund. Measure and Integral. Marcel Dekker Inc, New York, 1977.