

The Origin of Behavior*

Thomas J. Brennan[†] and Andrew W. Lo[‡]

This Draft: April 5, 2011

Abstract

We propose a single evolutionary explanation for the origin of several behaviors that have been observed in organisms ranging from ants to human subjects, including risk-sensitive foraging, risk aversion, loss aversion, probability matching, randomization, and diversification. Given an initial population of individuals, each assigned a purely arbitrary behavior with respect to a binary choice problem, and assuming that offspring behave identically to their parents, only those behaviors linked to reproductive success will survive, and less reproductively successful behaviors will disappear at exponential rates. When the uncertainty in reproductive success is systematic, natural selection yields behaviors that may be individually sub-optimal but are optimal from the population perspective; when reproductive uncertainty is idiosyncratic, the individual and population perspectives coincide. This framework generates a surprisingly rich set of behaviors, and the simplicity and generality of our model suggest that these derived behaviors are primitive and nearly universal within and across species.

Keywords: Probability Matching; Loss Aversion; Risk Aversion; Risk Preferences; Behavioral Finance; Evolution, Adaptive Markets Hypothesis.

JEL Classification: G00, D81, D01, D03, C73

*The views and opinions expressed in this article are those of the authors only, and do not necessarily represent the views and opinions of AlphaSimplex Group, MIT, Northwestern University, or any of their affiliates and employees. The authors make no representations or warranty, either expressed or implied, as to the accuracy or completeness of the information contained in this article, nor are they recommending that this article serve as the basis for any investment decision—this article is for information purposes only. Research support from AlphaSimplex Group, the MIT Laboratory for Financial Engineering, and the Northwestern University School of Law Faculty Research Program is gratefully acknowledged. We thank Elena Asparouhova, Nittai Bergman, Henry Cohn, Jayna Cummings, Arnout Eikeboom, Doyne Farmer, Nobu Kiyotaki, Simon Levin, Cecilia Lo, Arthur Robson, Jon Wilkins, Fernando Zapatero, and participants at the MIT Sloan Finance Lunch Seminar, the American Philosophical Society’s Fall 2010 meeting, the JOIM 2010 Spring Conference, the Q Group Fall 2010 Conference, and the 2010 Western Finance Association Conference for helpful comments and discussion.

[†]Assistant Professor, School of Law, Northwestern University, 375 East Chicago Avenue, Chicago, IL 60611–3069, t-brennan@law.northwestern.edu (email).

[‡]Harris & Harris Group Professor, MIT Sloan School of Management, and Chief Investment Strategist, AlphaSimplex Group, LLC. Please direct all correspondence to: Andrew Lo, MIT Sloan School, 100 Main Street, E62–618, Cambridge, MA 02142–1347, (617) 253–0920 (voice), alo@mit.edu (email).

Contents

1	Introduction	1
2	Literature Review	5
3	The Binary Choice Model	9
3.1	The Role of $\Phi(x_a, x_b)$	10
3.2	Individual Behavior	11
3.3	Population Dynamics	11
3.4	Asymptotic Properties	15
4	Probability Matching	17
4.1	Exact Probability Matching	17
4.2	The General Case	18
4.3	Individually Optimal vs. Growth-Optimal Behavior	21
5	Risk Preferences	21
5.1	Growth-Optimal Risk Preferences	22
5.2	Risk Aversion	26
5.3	Loss Aversion	27
6	Idiosyncratic vs. Systematic Risk	32
6.1	Idiosyncratic Risk	33
6.2	The General Case	37
7	Qualifications and Extensions	39
8	Conclusion	41
A	Appendix	43
A.1	Proof of Proposition 1	44
A.2	Proof of Corollary 1	44
A.3	Proof of Proposition 2	45
A.4	Proof of Proposition 3	45
A.5	Proof of Corollary 2	46
A.6	Proof of Proposition 4	46
A.7	Proof of Proposition 5	46
A.8	Proof of Corollary 4	47
A.9	Proof of Corollary 5	47

1 Introduction

One of the most influential ideas in all the social sciences is the efficient markets hypothesis of Samuelson (1965) and Fama (1970), the notion that market prices “fully reflect all available information”. This disarmingly simple but profound theory has been an object of intense study and debate, launching literally thousands of empirical investigations of its many specific implications for human behavior. These implications have been confirmed so often across so many contexts that Jensen (1978) concluded that “there is no other proposition in economics which has more solid empirical evidence supporting it than the Efficient Market Hypothesis”.

However, in recent years, accumulating evidence from psychology, behavioral economics and finance, the cognitive neurosciences, and biology have highlighted significant and abiding inconsistencies between *Homo economicus* and *Homo sapiens*.¹ These inconsistencies have called into question the predicates of expected utility theory, rational expectations, equilibrium, and even the Law of One Price, which are the very foundations of efficient markets, and much of the rest of neoclassical economics and finance. The Financial Crisis of 2007–2009 has only added more fuel to the many fires now threatening the efficient markets edifice.

The juxtaposition of compelling empirical evidence in support of market efficiency with equally compelling empirical and experimental evidence of apparent irrationality and behavioral anomalies suggests that human behavior may not be solely determined by economic considerations, but is, instead, the amalgam of multiple decision-making faculties—including instinct, emotion, and logic—that yield observed actions. And because the relative importance of these faculties varies across time and circumstances, even for a given individual, it is no wonder that despite centuries of intense analysis and debate, there is still remarkably little consensus among economists, psychologists, and biologists as to how to model human behavior.

In this paper, we propose an evolutionary explanation for the origin of behavior that is simple enough to solve analytically, but general enough to explain commonly observed behaviors in animal species ranging from ants to human subjects. Specifically, we show that risk aversion, risk-sensitive foraging, loss aversion, probability matching, and more general and previously inexplicable forms of randomizing behavior can all be derived from evolutionary forces acting on an arbitrary set of behaviors over an extended period of time.

¹See, for example, Kahneman, Slovic, and Tversky (1982) and Thaler (1993).

Some of these behaviors have been the subject of significant controversy, and these so-called “behavioral biases” are especially pronounced when elements of risk and probability are involved. Two of the most ubiquitous biases are loss aversion (Tversky and Kahneman, 1974; Kahneman and Tversky, 1979)—the tendency to take greater risk when choosing between two potential losses, and less risk when choosing between two potential gains—and probability matching, also known as the “matching law” or “Herrnstein’s Law” (Grant, Hake, and Hornseth, 1951; Hake and Hyman, 1953; Herrnstein, 1997; Vulcan, 2000)—the tendency to choose randomly between heads and tails when asked to guess the outcomes of a series of biased-coin tosses, where the randomization matches the probability of the biased coin. The idea of randomizing behavior is especially difficult to reconcile with the standard economic paradigm of expected utility theory in which individual behavior is non-stochastic and completely determined by utility functions, budget constraints, and the probability laws governing the environment. Both types of biases clearly imply irrationality, i.e., individually sub-optimal choices, yet these behaviors have been observed in thousands of geographically diverse human subjects over several decades, as well as in other animal species.

Our model consists of an initial population of individuals (not necessarily human) that live for one period of unspecified length, and engage in a single binary decision that has implications for the random number of offspring they will generate. To the extent that their behavior is linked to fecundity, only the most reproductively successful behaviors will flourish due to the forces of natural selection. Although obvious from an evolutionary biologist’s perspective, this observation yields surprisingly specific implications for the types of behavior that are sustainable over time, behaviors that are likely to be innate to most living organisms.

A simple numerical example of one of our results will illustrate our approach. Consider a population of individuals, each facing a binary choice between one of two possible actions, a and b . 60% of the time, environmental conditions are positive, and action a leads to reproductive success, generating 3 offspring for the individual. 40% of the time, environmental conditions are negative, and action a leads to 0 offspring. Suppose action b has exactly the opposite outcomes—whenever a yields 3 offspring, b yields 0, and whenever a yields 0, b yields 3. From the individual’s perspective, always choosing a , which has the higher probability of reproductive success, will lead to more offspring on average. However, if all individuals in the population behaved in this “rational” manner, the first time that a negative environmental condition occurs, the entire population will become extinct. Assuming

that offspring behave identically to their parents, the behavior “always choose a ” cannot survive over time. For the same reason, “always choose b ” is also unsustainable. In fact, we show below that in this special case, the behavior with the highest reproductive success over time is for each individual to choose a 60% of the time and b 40% of the time, matching the probabilities of reproductive success and failure. Eventually, this behavior will dominate the entire population.

Probability matching has long puzzled economists and psychologists because of its apparent inconsistency with basic self interest.² However, probability matching is perfectly consistent with evolution, arising purely from the forces of natural selection and population growth. Moreover, for more general environmental conditions, i.e., more general assumptions for reproductive success, we derive more general types of behavior that involve randomization but not necessarily probability matching. These results may explain the inconsistency with which such behavior is observed: whether or not randomizing behavior matches environmental probabilities depends on the relative reproductive success of the outcomes, and our framework yields a simple and specific condition for such behavior. Our results do not depend on how individuals arrive at their choices, whether they learn over time, or whether individuals possess a theory of mind or a self-awareness of the consequences of their actions. In fact, our results do not even require individuals to possess central nervous systems.

We also show that the concepts of risk aversion and risk-sensitive foraging behavior emerge from the same framework. Because populations grow geometrically, a sequence of 50/50 gambles yielding 2 or 4 offspring each generation will yield a slower average growth rate than sure bets of 3 offspring (the product of 2 and 4 is smaller than the product of 3 and 3). While this principle of “geometric mean fitness” is well known among population biologists, its implications for risk-bearing activity in economic settings has not been fully explored. For example, the fact that the preferences most likely to survive over time are those that require a higher expected fecundity in return for taking risk implies the existence of a positive evolutionary “risk premium”, which we are able to derive explicitly and quantify as a function of environmental conditions.

Our model also generates asymmetric risk preferences for gains and losses, i.e., loss aver-

²One of the earliest papers to document this phenomenon is Grant, Hake, and Hornseth (1951), and as recently as 2007, Kogler and Kühberger (2007) report that: “Experimental research in simple repeated risky choices shows a striking violation of rational choice theory: the tendency to match probabilities by allocating the frequency of response in proportion to their relative probabilities”.

sion, resolving a longstanding debate among disciples of von Neumann and Morgenstern expected-utility theory (von Neumann and Morgenstern, 1944; Schoemaker, 1982) and proponents of behavioral alternatives such as prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1981). Each perspective is capturing somewhat different aspects of the same behavior shaped by natural selection and population growth. Moreover, our framework provides an explanation for the apparent variability in the experimental evidence regarding loss aversion, and the nature of the “reference point” that demarcates risk-seeking and risk-averse behavior: evolution shapes an individual’s decision-making mechanism or preferences to enhance reproductive success which, in turn, is determined by total wealth, not incremental wealth. Accordingly, two individuals with the same decision-making process but different levels of net worth may behave differently when offered the same incremental prospects.

Finally, and perhaps most significantly, our framework illustrates the role of the environment in shaping behavior, and why certain populations appear to exhibit irrational behavior while others do not. The difference can be traced to a single feature of the environment’s impact on individual biology: systematic vs. idiosyncratic sources of randomness in reproduction. In populations where environmental factors are largely systematic, i.e., they affect the reproduction rates of all individuals in the same manner, any form of synchronization in behavior may lead to extinction, hence such synchrony is unlikely to be perpetuated. In other words, if environmental risks are systematic, survival depends on the population diversifying its behavior so that some fraction will survive to reproduce no matter what the environment is like. In such cases, it may seem as if certain individuals are acting irrationally since they may not be behaving optimally for a given environment. But such heterogeneous behavior is, in fact, optimal from the perspective of the population. If, on the other hand, environmental risks are largely idiosyncratic, then individuals engaging in identical behavior, e.g., individually optimal behavior, will not expose the population to extinction. Moreover, in this case, natural selection will favor the individually optimal behavior; hence, a population of individually rational individuals will emerge. Thus, rationality is not necessarily in the eyes of the individual beholder, but is sometimes in the hands of systematic environmental factors.

In Section 2 we provide a review of the literature, and in Section 3, we present our binary choice framework. Using this framework, in Section 4 we derive necessary and sufficient conditions under which probability matching emerges and also show how probability match-

ing breaks down under different environments. By making one of the two binary choices riskless, in Section 5, we show how risk preferences evolve, derive an evolutionary equity risk premium, and show how loss aversion arises naturally from these preferences. In Section 6, we develop the implications of systematic and idiosyncratic risk for behavior, and show that the former yields populations in which individuals do not always act rationally, but not the latter. We propose several extensions of our binary choice framework in Section 7 and conclude in Section 8.

2 Literature Review

The literature on evolution and behavior is overwhelming, spanning the disciplines of evolutionary biology, ecology, evolutionary and social psychology, and economics, with myriad branches of relevant citations within each of these broad fields. While a comprehensive survey is well beyond the scope of this section, we attempt to provide a representative sampling of the many related strands of this vast body of research.

Evolutionary principles are now routinely used to derive implications for animal behavior. While each species may have developed unique responses for addressing particular environmental challenges, the most critical of these have been shaped by the forces of mutation, competition, and natural selection. Although such forces operate at the genetic level, as described so compellingly by Dawkins (1976), the pathbreaking work of Hamilton (1964), Trivers (1971, 1985, 2002), Wilson (1975), and Maynard Smith (1982, 1984) show that evolutionary mechanisms may also explain a variety of counter-intuitive behaviors including altruism, cooperation, kin selection, reciprocity, and other social customs. More recently, the field of evolutionary psychology (Cosmides and Tooby, 1994; Barkow, Cosmides, and Tooby, 1992; Tooby and Cosmides, 1995; Pinker, 1979; Pinker, 1991; Pinker, 1994; Gigerenzer, 2000; Buss, 2004; Ehrlich and Levin, 2005) has expanded the reach of evolution to even broader domains such as language, culture, and religion.

Evolutionary ideas have also played an important role in economics. Thomas Malthus (1826) used a simple biological argument—the fact that populations increase at geometric rates whereas natural resources increase at only arithmetic rates (at least in the nineteenth century)—to arrive at the dire economic consequences that earned the field the moniker “dismal science”. Both Darwin and Wallace were aware of and apparently influenced by these arguments (see Hirshleifer (1977) for further details). Also, Schumpeter’s (1939) view

of business cycles, entrepreneurs, and capitalism has an unmistakable evolutionary flavor to it; in fact, his ideas of “creative destruction” and “bursts” of entrepreneurial activity bear a striking resemblance to natural selection and Eldredge and Gould’s (1972) notion of “punctuated equilibrium”.

More recently, economists and biologists have begun to explore these connections in several veins: economic extensions of sociobiology (Becker, 1976; Hirshleifer, 1977); evolutionary game theory (Maynard Smith 1982, 1984, Weibull, 1995); an evolutionary interpretation of economic change (Nelson and Winter, 1982); economies as complex adaptive systems (Anderson, Arrow, and Pines, 1988); and the impact of uncertainty regarding the number of offspring on current consumption patterns (Arrow and Levin, 2009).

Evolutionary concepts have also appeared in the finance literature. For example, Luo (1995) explores the implications of natural selection for futures markets, Hirshleifer and Luo (2001) consider the long-run prospects of overconfident traders in a competitive securities market, and Kogan, Ross, Wang, and Westerfield (2006) show that irrational traders can influence market prices even when their wealth becomes negligible. The literature on agent-based modeling pioneered by Arthur, Holland, LeBaron, Palmer, and Tayler (1997), in which interactions among software agents programmed with simple heuristics are simulated, relies heavily on evolutionary dynamics. And at least two prominent investment professionals have proposed Darwinian alternatives to explain market behavior. In a chapter titled “The Ecology of Markets”, Niederhoffer (1997, Ch. 15) likens financial markets to an ecosystem with dealers as “herbivores”, speculators as “carnivores”, and floor traders and distressed investors as “decomposers”. And Bernstein (1998) makes a compelling case for active management by pointing out that the notion of equilibrium is rarely realized in practice and that market dynamics are better explained by evolutionary processes.

But in our specific context, the two most relevant lines of research—one from biology and the other from economics—involve direct applications of evolutionary principles to individual behavior and preferences. In the evolutionary biology literature, Maynard Smith (1982) has developed the concept of an “evolutionarily stable strategy” (ESS), specific behaviors that survive over time by conferring reproductive advantages or “fitness”, typically measured by the rate of population growth. Using this notion of fitness, Fretwell (1972), Cooper and Kaplan (1982), and Frank and Slatkin (1990) observe that randomizing behavior can be advantageous (in terms of maximizing geometric growth rates) in the face of stochastic en-

vironmental conditions. The impact of variability in reproductive success among individuals in a population has been shown to yield a kind of risk aversion (which increases average reproductive success) and “bet-hedging” (which reduces the variance of reproductive success) (Slatkin, 1974; Caraco, 1980; Real, 1980; Rubenstein, 1972; Seger and Brockmann, 1987). Frank and Slatkin (1990) propose a framework that highlights the importance of correlations among individual reproductive success in determining the path of evolution. And similar results have been derived in the behavioral ecology literature, in which the maximization of fitness via dynamic programming has been shown to yield several observed behaviors including risk-sensitive foraging in mammals (Real and Caraco, 1986; Stephens and Krebs, 1986; Mangel and Clark, 1988), and seed dispersal strategies in plants (Levin, Cohen, and Hastings, 1984; Levin, Muller-Landau, Nathan, and Chave, 2003).

In the economics literature, evolutionary principles have been used to justify the existence of utility functions and develop implications for their functional form, as in Hansson and Stuart (1990) and Robson (1996a, 2001b) (see, Robson, 2001a, and Robson and Samuelson, 2010, for comprehensive reviews of this literature). For example, in an equilibrium model of economic growth, Hansson and Stuart (1990) derive restrictions on individual preferences for consumption, savings, and labor-supply arising from the forces of natural selection. Robson (1996a) investigates expected and non-expected utility behaviors, and finds that idiosyncratic-risk-seeking may be optimal from a population perspective even though it is suboptimal from an individual perspective (see, also, Grafen, 1999, and Curry, 2001). And Robson (2001b) argues that the kind of predictable behavior capable of being captured by a utility function emerged naturally as an adaptive mechanism for individuals faced with repeated choices in a nonstationary environment. Specifically, early exploration in choice-making (which is the primary focus of our analysis), coupled with a utility-based rule of thumb for deciding when to cut exploration off and stick with a particular choice, leads to evolutionarily optimal adaptation to unknown underlying distributions of outcomes. Robson and Samuelson (2007) find that exponential discounting in utility functions is consistent with evolutionarily optimal growth of a population, and the emergence of time preference is derived by Rogers (1994), Samuelson (2001), Robson and Samuelson (2007, 2009), and Robson and Szentes (2008).

However, as Waldman (1994, p. 483) observed, individually optimal behavior predicted by expected utility—even when utility functions are derived from evolutionary principles—may

not always coincide with behavior that maximizes fitness:

Another possible outcome is that preferences do not equate utility maximization with fitness maximization, and correspondingly evolution then does not favor humans who are efficient utility-maximizers. Instead what happens in this case is that evolution favors a systematic bias in the decision-making process which moves behavior away from the maximization of utility and toward the maximization of fitness.

Waldman (1994) provides a compelling illustration of this insight through the comparison between asexual and sexual reproduction in which the latter yields evolutionarily stable second-best adaptations. In our binary choice framework, we show that even with asexual reproduction, systematic “errors” such as probability matching can persist and become dominant despite the fact that such behavior is sub-optimal from the individual’s perspective.

Our approach builds on the insights of Fretwell (1972), Maynard Smith (1982), Waldman (1994), and Robson (1996a, 2001a) in applying the well-known principle of geometric-mean fitness (Dempster, 1955) to the actions of a heterogeneous population of individuals and deriving the subset of behaviors that survive.³ However, our framework is considerably simpler, involving only a single binary choice for each individual during its lifetime, a choice that has implications for the individual’s reproductive success. One virtue of such parsimony is the universality with which this framework’s derived behaviors are likely to be found among living organisms.⁴ Our model is simple enough to solve analytically, but remarkably rich in its implications for behavior, yielding risk aversion, probability matching, loss aversion, and more general forms of randomization. Also, as with most other models in the population biology literature, the individual behaviors that survive in our framework need not be optimal from the individual’s perspective, and may appear irrational. In fact, such behavior is merely

³Geometric-mean fitness has also appeared in the financial context as the “Kelly criterion” for maximizing the geometric growth rate of a portfolio (Kelly, 1956; Cover and Thomas, 1991). However, the motivation for geometric-mean fitness in population biology is considerably more compelling than in financial investments, as Samuelson (1971) has argued (maximizing the geometric-mean return of a portfolio is optimal only for individuals with a very specific risk preference, i.e., those with logarithmic utility functions).

⁴For example, the fact that probability matching behavior has been observed in non-human subjects—including ants (Deneubourg, Aron, Goss, and Pasteels, 1987; Pasteels, Deneubourg, and Goss, 1987; Kirman, 1993; Hölldobler and Wilson, 1990), bees (Harder and Real, 1987; Thuijsman, Peleg, Amitai, and Shmida, 1995; Keasar, Rashkovich, Cohen, and Shmida, 2002), fish (Bitterman, Wodinsky, and Candland, 1958; Behrend and Bitterman, 1961), pigeons (Graf, Bullock, and Bitterman, 1964; Young, 1981), and primates (Woolverton and Rowlett, 1998)—suggests that they may have a common and ancient origin, and an evolutionary role that belies their apparent shortcomings.

adaptive, a product of natural selection that is likely to be more primitive on an evolutionary timescale than the more sophisticated learned behaviors captured by Rogers (1994), Robson (1996a, 2001a), and Robson and Samuelson (2007, 2009). In this respect, our analysis complements those of the existing literature on evolutionary foundations of utility theory, providing additional evidence for the link between behavior and natural selection at the most basic level of choice.

3 The Binary Choice Model

We begin with a population of individuals that live for one period, produce a random number of offspring asexually, and then die (“asexual semelparous” organisms, in the jargon of evolutionary biology). During their lives, individuals make only one decision: they choose one of two possible courses of action, denoted a and b , and this results in one of two corresponding random numbers of offspring, x_a and x_b , described by some well-behaved probability distribution function $\Phi(x_a, x_b)$. We assume that x_a and x_b are not perfectly correlated, otherwise for all intents and purposes, individuals have only one action available to them. We also assume:

- (A1) (x_a, x_b) and $\log(fx_a + (1-f)x_b)$ have finite moments up to order 2 for all $f \in [0, 1]$.
- (A2) (x_a, x_b) is independently and identically distributed (IID) over time, and identical for all individuals in a given generation

(A1) and (A2) are standard assumptions that allow us to derive analytically tractable results, and are not nearly as implausible in biological contexts as they are when applied to financial data.⁵ For the moment, we assume in (A2) that x_a and x_b are the same two outcomes for all individuals in the population; in other words, if two individuals choose the same action a , both will produce the same number of random offspring x_a . This implies that the variation in offspring due to behavior is wholly “systematic”, i.e., the link between action and reproductive success is the same throughout the entire population. This assumption is

⁵Both assumptions can be relaxed to some degree and at the expense of analytical simplicity. For example, (A1) can be relaxed by considering random variables with no finite moments of any order, in which case we must focus our attention on location and scale parameters. The IID assumption of (A2) can also be relaxed by imposing stationarity and ergodicity, or by allowing heterogeneity in the marginal distributions but imposing mixing conditions as in White (1984). We expect qualitatively similar results in these more general cases.

significant, and its ramifications can be better understood when we consider the alternate case of “idiosyncratic” random offspring in Section 6.

3.1 The Role of $\Phi(x_a, x_b)$

The role of Φ is critical in our framework, as it represents the entirety of the implications of an individual’s actions for reproductive success. Embedded in Φ is the biological machinery that is fundamental to evolution, i.e., genetics, but of less direct interest to economists than the link between behavior and reproductive success. If action a leads to higher fecundity than action b for individuals in a given population, the particular set of genes that predispose individuals to select a over b will be favored by natural selection, in which case these genes will survive and flourish, implying that the behavior “choose a over b ” will flourish as well. On the other hand, if a and b have identical implications for success, i.e., $x_a \equiv x_b$, then Φ is a degenerate distribution. By asserting that Φ is a non-degenerate bivariate distribution, we have essentially defined two equivalence classes of actions that have different implications for reproduction, i.e., all actions yielding the same reproductive fitness are considered equivalent in our framework.

The specification of Φ also captures the fundamental distinction between traditional models of population genetics (Levins, 1969; Wright, 1968; Wilson and Bossert, 1990; Dawkins, 1976) and more recent applications of evolution to behavior (Hamilton, 1964; Trivers, 1971; Wilson, 1975; Maynard Smith, 1982); the former focuses on the natural selection of traits (determined by genetics), whereas the latter focuses on the natural selection of behavior. Although behavior is obviously linked to genetics, the specific genes involved, their loci, and the mechanisms by which they are transmitted from one generation to the next are of less relevance to economic analysis than the ultimate implications of behavior for reproduction, which is captured by Φ . In the jargon of econometrics, Φ may be viewed as a “reduced form” representation of an individual’s biology, whereas the molecular biology of genetics corresponds to the “structural form”.⁶ This terminology is more than a simple analogy—it accurately summarizes the difference between our framework and the emerging field of behavioral genomics (Plomin, 1990; Plomin, Owen, and McGuffin, 1994; McGuffin, Riley, and Plomin, 2001) which attempts to map traits and behaviors to specific genes.

⁶Waldman (1994) uses the same terminology in his framework in which individual traits are directly linked to reproductive success rather than specific genes.

3.2 Individual Behavior

Now suppose that each individual i chooses a with some probability $f \in [0, 1]$ and b with probability $1 - f$, denoted by the Bernoulli variable I_i^f , hence i 's offspring x_i^f is given by:

$$x_i^f = I_i^f x_a + (1 - I_i^f) x_b \quad , \quad I_i^f \equiv \begin{cases} 1 & \text{with probability } f \\ 0 & \text{with probability } 1 - f \end{cases} . \quad (1)$$

We shall henceforth refer to f as the individual's "behavior" since it completely determines how the individual chooses between a and b . Note that f can be 0 or 1, hence we are not requiring individuals to randomize—this will be derived as a population-wide consequence of natural selection under certain conditions. Also, we ascribe no intelligence or volition to this behavior; we are simply providing a formal representation for it, and then investigating its evolutionary prospects. To that end, we assume that offspring behave in a manner identical to their parents, i.e., they choose between a and b according to the same f , hence the population may be viewed as being comprised of "types" of individuals indexed by f that range continuously from 0 to 1, including the endpoints. In this manner, we are able to study the evolutionary dynamics of each type of individual over many generations.

3.3 Population Dynamics

Because our analysis involves individuals making binary decisions over time and across multiple generations, the number and type of subscripts used are sometimes excessive and confusing. Therefore, before considering the dynamics of the population we have proposed, a few clarifying comments regarding our notation may be useful.

Individuals in a given generation t are indexed by i , and generations are indexed by $t = 1, \dots, T$. Because in all cases, we assume independently and identically distributed randomness over time or, equivalently, generations, on occasion we will omit the t subscript unless we wish to emphasize the temporal ordering of the variables (such as a recursive relation between two successive generations). Finally, a superscript f will denote the particular type of individual as defined by the decision rule I_i^f in (1).

With these notational conventions in mind, denote by n_t^f the total number of offspring of type f in generation t , which is simply the sum of all the offspring from the type- f individuals

of the previous generation:

$$n_t^f = \sum_{i=1}^{n_{t-1}^f} x_{i,t}^f = \left(\sum_{i=1}^{n_{t-1}^f} I_{i,t}^f \right) x_{a,t} + \left(\sum_{i=1}^{n_{t-1}^f} (1 - I_{i,t}^f) \right) x_{b,t} \quad (2)$$

$$n_t^f \stackrel{p}{=} n_{t-1}^f (f x_{a,t} + (1-f)x_{b,t}) \quad (3)$$

where we have added time subscripts to the relevant variables to clarify their temporal ordering, and “ $\stackrel{p}{=}$ ” in (3) denotes equality in probability as n_{t-1}^f increases without bound (see Definition A.1 in the Appendix), which follows from the Law of Large Numbers applied to the sum $\sum_i I_{i,t}^f/n_{t-1}^f$ (recall that $I_{i,t}^f$ is IID across i).⁷

Through backward recursion, the population size from (3) of type- f individuals in generation T is given by:

$$n_T^f \stackrel{p}{=} \prod_{t=1}^T (f x_{a,t} + (1-f)x_{b,t}) = \exp\left(\sum_{t=1}^T \log(f x_{a,t} + (1-f)x_{b,t}) \right) \quad (4)$$

$$\frac{1}{T} \log n_T^f \stackrel{p}{=} \frac{1}{T} \sum_{t=1}^T \log(f x_{a,t} + (1-f)x_{b,t}) \xrightarrow{p} \mathbb{E}[\log(f x_a + (1-f)x_b)] \quad (5)$$

where “ \xrightarrow{p} ” in (5) denotes convergence in probability (see Definition A.1 in the Appendix), and follows from the Kolmogorov Law of Large Numbers applied to the sum $\sum_t \log(f x_{a,t} + (1-f)x_{b,t})/T$ as T increases without bound (Lewontin and Cohen, 1969) (recall that $(x_{a,t}, x_{b,t})$ is assumed to be IID over time, hence we have dropped the t subscripts in the expectation in (5)), and we have assumed that $n_0^f = 1$ without loss of generality.

Since the value of f that maximizes the population size n_T^f is also the value of f that maximizes $T^{-1} \log n_T^f$,⁸ (5) implies that this value converges in probability to the maximum of the following the expectation:⁹

$$\mu(f) \equiv \mathbb{E}[\log(f x_a + (1-f)x_b)] . \quad (6)$$

⁷In particular, the Kolmogorov Law of Large Numbers asserts that $\sum_i I_{i,t}^f/n_{t-1}^f$ converges almost surely to $\mathbb{E}[I_{i,t}^f] = f$ (see, for example, Serfling (1980, Ch. 1.8)).

⁸This follows from the fact that $T^{-1} \log(n_T^f)$ is a monotone transformation n_T^f .

⁹More precisely, the value f^* that maximizes $\mu(f)$ corresponds to a population size $n_t^{f^*}$ that is asymptotically larger than any other population n_t^f , $f \neq f^*$, in the sense that $\text{plim}_{t \rightarrow \infty} n_t^f/n_t^{f^*} = 0$. See Section 3.4 for a more detailed exposition of these asymptotic properties.

This expression is simply the expectation of the log-geometric-average growth rate of the population, and the value f^* that maximizes it—which we shall call the “growth-optimal” behavior to distinguish it from behavior that may be optimal for the individual—is given by (all proofs are relegated to the Appendix):

Proposition 1 *Under Assumptions (A1)–(A2), the growth-optimal behavior f^* is:*

$$f^* = \begin{cases} 1 & \text{if } E[x_a/x_b] > 1 \text{ and } E[x_b/x_a] < 1 \\ \text{solution to (8)} & \text{if } E[x_a/x_b] \geq 1 \text{ and } E[x_b/x_a] \geq 1 \\ 0 & \text{if } E[x_a/x_b] < 1 \text{ and } E[x_b/x_a] > 1 \end{cases} \quad (7)$$

where f^* is defined implicitly in the second case of (7) by:

$$0 = E \left[\frac{x_a - x_b}{f^* x_a + (1 - f^*) x_b} \right] \quad (8)$$

and the expectations in (5)–(8) are with respect to the joint distribution $\Phi(x_a, x_b)$.

The three possible behaviors in (7) reflect the relative reproductive success of the two choices, and is a generalization of the “adaptive coin-flipping” strategies of Cooper and Kaplan (1982). Choosing a deterministically will be optimal if choice a exhibits unambiguously higher expected relative fecundity; choosing b deterministically will be optimal if the opposite is true; and randomizing between a and b will be optimal if neither choice has a clear-cut reproductive advantage. This last outcome is perhaps the most counter-intuitive because it is sub-optimal from an individual’s perspective, but the population perspective implies that in such cases, the individuals that have the most reproductive success over time will be those that choose randomly according to probability f^* .¹⁰ If, however, one choice is significantly better than the other in terms of the expected ratio of offspring, then over time, the behavior that will survive is deterministic choice, not randomization. In these extreme cases, because of the unambiguous implications for fecundity, the deterministic choice that leads to higher reproductive success will quickly dominate the population.

¹⁰Cooper and Kaplan (1982) interpret this behavior as a form of altruism because individuals seem to be acting in the interest of the population at the expense of their own fitness. However, Grafen (1999) provides a different interpretation by proposing an alternate measure of fitness, one that reflects the growth rate of survivors.

The behavior f^* that emerges through the forces of natural selection is quite distinct from the neoclassical economic framework of expected utility in one important respect: expected utility theory implies deterministic behavior.¹¹ Given the same environmental conditions and preference parameters, the action that maximizes expected utility will be the same. In our framework, there are many circumstances in which f^* is strictly greater than 0 and less than 1, hence even if an individual's circumstances are identical, the behavior shaped by natural selection will not be.¹²

The deterministic choices $f^* = 0, 1$ in Proposition 1 may be viewed as a primitive form of herding behavior—where all individuals in the population choose to act in the identical manner—especially if the relative fecundities $E[x_a/x_b]$ and $E[x_b/x_a]$ shift suddenly from the intermediate state in (7) to one of the deterministic states due to rapid environmental changes, and if there is sufficient diversity of behavior left in the population after the change occurs.¹³ To an outside observer, behaviors among individuals in this population may seem heterogenous before the shift (because individuals are randomizing), but will become increasingly similar after the shift—as selective pressures begin to favor deterministic behavior over randomization—creating the appearance (but not the reality) of intentional coordination, communication, and synchronization. If the reproductive cycle is sufficiently short, this change in population-wide behavior may seem highly responsive to environmental changes, giving the impression that individuals are learning about their environment. This is indeed a form of learning, but it occurs at the population level, not at the individual level, and not within an individual's lifespan. Considerably more sophisticated adaptations are necessary to generate true herding and synchronization behavior as in Hamilton (1971), Mirollo and Strogatz (1990), and Strogatz and Stewart (1993), including sensory inputs,

¹¹Although random utility models have been proposed by Thurstone (1927), McFadden (1973), Manski (1975), and others in the discrete-choice literature (see Manski and McFadden (1981)), the source of randomness in these models is assumed to be measurement error, not behavior itself.

¹²In this respect, our framework differs in a fundamental way from the evolutionary models of Robson (1996a, 2001b), Robson (1996a) Robson (2001b), Robson and Samuelson (2007), and Robson and Szentes (2008) in which evolutionary arguments are used to derive specific utility functions that individuals optimize to yield particular behaviors. Although Robson (2001b) proposes the emergence of utility functions as an adaptation to repeated choices in a nonstationary environment, our framework shows that the concept of utility is not necessarily primitive to behavior or natural selection, and there is at least one type of behavior—randomization—that cannot be captured by standard expected utility theory.

¹³This last qualification is critical due to our assumption that offspring behave exactly as their parents, hence if individuals with behavior $f = 1$ no longer exist in the population, such behavior cannot emerge even if it becomes optimal from the population perspective. This observation underscores the importance of random mutations and the evolutionary advantages of sexual reproduction in the face of stochastic environmental conditions.

conditional behavior (conditioned on additional state variables), and neuroplasticity.

Proposition 1 may also be interpreted as a primitive form of group selection, in which natural selection appears to operate at the group level instead of, or in addition to, the level of individuals, traits, or genes (Wynne-Edwards, 1962; Sober and Wilson, 1998). However, in this case, the notion of a “group” is determined by the interaction between behavior and the environment—those individuals with behavior f^* will appear to be favored, and those with other behaviors $f \neq f^*$ will be disadvantaged.

3.4 Asymptotic Properties

The growth-optimal behavior described in (7) is simple, but a direct corollary is that f^* leads to a “winner-take-all” outcome in which individuals of all other sub-optimal types f' will be rapidly overrun by individuals of type f^* , since the ratio of the population sizes of f' and f^* converges exponentially fast to 0 (due to the optimality of f^*):

Corollary 1 *Under Assumptions (A1)–(A2), as T increases without bound, the geometric average growth rate $(n_T^f)^{1/T}$ of the population of individuals with behavior f converges in probability to $\exp(\mu(f))$, and the growth-optimal behavior f^* will dominate the population exponentially fast since:*

$$\left(\frac{n_T^{f'}}{n_T^{f^*}}\right)^{1/T} \stackrel{P}{\underset{P}{\rightarrow}} \exp([\mu(f') - \mu(f^*)]) \stackrel{P}{\rightarrow} 0 \quad (9)$$

which implies that $n_T^{f'}/n_T^{f^*} \stackrel{P}{\rightarrow} 0$ at an exponential rate.

This corollary confirms that the behavior produced by the forces of natural selection is indeed given by (7), and the exponential rate of convergence underlies the winner-take-all phenomenon that seems to characterize so many competitive situations (see, for example, Frank and Cook (1995)). Whether such behavior is deterministic or random depends entirely on the implications of such behavior for reproductive success.

Note that our asymptotic approach to studying the evolutionary properties of behavior is different from the typical biologist’s perspective in that the main object of interest for us is f^* , not the population itself. In particular, the actual size of the population of type- f individuals is of less concern to us than the fact that selection will favor one particular f^* . In fact, in contrast to standard models in population dynamics that are formulated to

have stable equilibria, the population size in our framework approaches 0 or infinity in the limit, depending on the parameters of $\Phi(x_a, x_b)$.¹⁴ Nevertheless, we can fully characterize the statistical properties of the population via the Central Limit Theorem. Specifically, as T grows without bound, it can be shown that the suitably normalized population size n_T^f approaches a lognormal distribution:

Proposition 2 *Under Assumptions (A1)–(A2), as T increases without bound, the geometric average $(n_T^f)^{1/T}$ of the population size of individuals with behavior f converges in distribution to a lognormal random variable:*

$$\sqrt{T} \left(T^{-1} \log n_T^f - \mu(f) \right) \stackrel{a}{\sim} \mathcal{N} \left(0, \sigma^2(f) \right) \quad , \quad \sigma^2(f) \equiv \text{Var}[\log(fx_a + (1-f)x_b)] \quad (10)$$

where ‘ $\stackrel{a}{\sim}$ ’ denotes asymptotic equivalence in distribution.

This evolutionary basis of behavior is the same as Seger and Brockmann’s (1987) geometric-mean fitness criterion, but which is applied directly to reproductive success x , not to specific genes. The basic logic of (7) is similar to the ESS of Maynard Smith (1982) in which mixed strategies are shown to be evolutionarily stable, but our approach is more parsimonious and, by design, yields broader implications for behavior as represented by Φ . Our derived behaviors are also distinct from those of Hansson and Stuart (1990), Robson (1996a, 1996b, 2001a,b), Grafen (1999), Curry (2001), Samuelson (2001), and Robson and Samuelson (2007) in which evolutionary arguments are used to justify specific types of utility functions and risk preferences—in the latter cases, two individuals with identical utility functions will behave identically, whereas in our case, two individuals with identical $f^* \in (0, 1)$ may make different choices at any point in time (see, also, Waldman (1994)).

In the sections to follow, we show that this simple binary choice framework and growth-optimal behavior can explain a surprisingly rich set of behavioral anomalies that have been a source of controversy in economics, psychology, and evolutionary biology, including more general forms of probability matching, loss aversion, and risk aversion.

¹⁴To see why, observe that the Kolmogorov Law of Large Numbers implies that $T^{-1} \log n_T^f$ converges almost surely to $\mu(f) \equiv \text{E}[\log(fx_a + (1-f)x_b)]$. If $\mu(f)$ is positive, the population grows without bound, and if $\mu(f)$ is negative, the population becomes extinct.

4 Probability Matching

Using the binary choice framework, we can easily derive probability matching behavior as emerging solely through the forces of natural selection. In Section 4.1, we derive conditions that yield exact probability matching. In Section 4.2, we generalize this result considerably, deriving conditions for approximate probability matching, as well as conditions under which probability matching does not arise.

4.1 Exact Probability Matching

To develop further intuition for the binary choice model, consider the special case in which the number of offspring (x_a, x_b) are simply Bernoulli random variables that are perfectly out of phase in each of two possible environmental states:

	State 1	State 2	
Action	(prob. p)	(prob. $1-p$)	
a	$x_a = m$	$x_a = 0$	(11)
b	$x_b = 0$	$x_b = m$	

With probability p , one environmental state is realized in which choice a yields $m > 0$ offspring and choice b yields none, and with probability $1-p$ the other environmental state is realized in which the reverse is the case. Without loss of generality, assume that $p \in (\frac{1}{2}, 1]$ so that the environmental state in which choice a produces offspring is more likely.

In this simple case of 0 or m offspring, the expectation in (6) can be evaluated explicitly as:

$$\mu(f) = \log m + p \log f + (1-p) \log(1-f) \tag{12}$$

and the value of f that maximizes this expression is p . Despite the fact that setting $f = 1$ maximizes the likelihood that an individual is able to reproduce, such “selfish” behavior is not sustainable from a population perspective because if all individuals were to behave in this way, the entire population would be wiped out the first time $x_a = 0$. Since such an extinction is almost sure to occur eventually,¹⁵ the behavior $f = 1$ will ultimately be eliminated from the

¹⁵If all individuals always choose a , and if (x_a, x_b) is independently and identically distributed over time, the probability of extinction by time t is $1-p^t$, which approaches 1 as t increases without bound.

population. In contrast, the randomizing behavior $f^* = p$ yields the highest possible growth rate $I(p) = \log(mp^p(1-p)^{(1-p)})$. This is classic probability matching, first documented over half a century ago (Grant, Hake, and Hornseth, 1951). Since then, this behavior has become so well-established among both human and non-human subjects that it is now referred to as the “matching law” or “Herrnstein’s law”.¹⁶ However, there have been several notable departures from this behavioral pattern (Baum, 1974; Horne and Lowe, 1993; Kogler and Kühberger, 2007), hence this “law” may not be as consistent as its moniker suggests. In the next section, we provide a general explanation for both matching behavior and departures from it.

4.2 The General Case

The Bernoulli example above can be easily generalized to any arbitrary number of offspring for both choices:

$$\begin{aligned} \text{Prob}(x_a = c_{a1}, x_b = c_{b1}) &= p \in [0, 1] \\ \text{Prob}(x_a = c_{a2}, x_b = c_{b2}) &= 1 - p \equiv q \end{aligned} \tag{13}$$

where we assume that $c_{ij} \geq 0$ and $c_{aj} + c_{bj} \neq 0$, $i = a, b$ and $j = 1, 2$. The condition $c_{aj} + c_{bj} \neq 0$ rules out the case where both c_{aj} and c_{bj} are 0, in which case the binary choice problem becomes degenerate because both actions lead to extinction hence the only choice that has any impact on fecundity is in the non-extinction state, and the only behavior that is sustainable is to select the action with the higher number of offspring.

The growth-optimal behavior in this case will depend on the relation between the probability p and the relative-fecundity variables $r_j \equiv c_{aj}/c_{bj}$ for each of the two possible states of the world $j = 1, 2$.¹⁷ Specifically, we have:

Proposition 3 *Under Assumptions (A1)–(A2), and if (x_a, x_b) satisfies (13), then the growth-*

¹⁶See Herrnstein (1961, 1970, 1997), Bradshaw, Szabadi, and Bevan (1976), Davison and McCarthy (1988), Herrnstein and Prelec (1991), and the references cited in footnote 4.

¹⁷Since c_{ij} may be 0, the ratios r_j may be infinite if a finite numerator is divided by 0, which poses no issues for any of the results in this paper as long as the usual conventions involving infinity are followed. The ambiguous case of $r_j = 0/0$ is ruled out by the condition $c_{aj} + c_{bj} \neq 0$.

optimal behavior f^* is given by:

$$f^* = \begin{cases} 1 & \text{if } r_2 \in [q + \frac{pq}{r_1-p}, \infty) \text{ and } r_1 > p \\ \frac{p}{1-r_2} + \frac{q}{1-r_1} & \text{if } \begin{cases} r_2 \in (\frac{1}{q} - \frac{p}{q}r_1, q + \frac{pq}{r_1-p}) \text{ and } r_1 > p, \text{ or} \\ r_2 \in (\frac{1}{q} - \frac{p}{q}r_1, \infty) \text{ and } r_1 \leq p \end{cases} \\ 0 & \text{if } r_2 \in [0, \frac{1}{q} - \frac{p}{q}r_1] \end{cases} \quad (14)$$

Figure 1 illustrates the values of r_1 and r_2 that yield each of the three types of behaviors in (14). If r_1 and r_2 are not too different—implying that the ratio of fecundities of choices a and b is not that different between the two states of the world—then random behavior yields no evolutionary advantage over deterministic choice. In this case, the individually optimal behavior ($f^* = 0$ or 1) will prevail in the population. If, on the other hand, one of the r variables is large while the other is small, then random behavior will be more advantageous from the population perspective than a deterministic one. In such cases, there are times in which each choice performs substantially better than the other, hence it is evolutionarily optimal for a population to diversify between the two choices rather than to always choose the outcome with the highest probability of progeny in a single generation. This case is summarized in:

Corollary 2 *Suppose there exists a large difference between r_1 and r_2 ; without loss of generality, let $r_1 \gg 0$, $r_2 \ll 1$, and $p > \frac{1}{2}$. Then under Assumptions (A1)–(A2), and if (x_a, x_b) satisfies (13), the growth-optimal behavior is given by:*

$$f^* = p(1 + O(1/r_1) + O(r_2)) \approx p. \quad (15)$$

Equation (15) shows that if one choice is much worse than the other choice p -percent of the time, and if the other choice is much worse than the first $(1-p)$ -percent of the time, then the first choice should be chosen with probability p and the second choice should be chosen with probability $1-p$. The definition of “much worse” is made precise by specifying that the values of $1/r_1$ and r_2 are both close to zero—over time, the individuals that flourish in such a world are precisely those that engage in approximate probability matching behavior.

When r_1 and r_2 satisfy the condition:

$$0 = p \frac{r_2}{1-r_2} + q \frac{1}{1-r_1}, \quad (16)$$

exact probability matching behavior arises, and the solid black curve in Figure 1 illustrates the locus of values for which this condition holds. The horizontal asymptote of the curve occurs at $r_2 = 0$, so as r_2 tends toward zero and r_1 becomes relatively large, exact probability matching will be optimal (note that the asymmetry between r_1 and r_2 is due entirely to our requirement that $f^* = p$ and $p \neq \frac{1}{2}$). However, values of (r_1, r_2) off this curve but still within the shaded region imply random behavior that is approximately—but not exactly—probability matching, providing a potential explanation for more complex but non-deterministic foraging patterns observed in various species (Deneubourg, Aron, Goss, and Pasteels, 1987; Pasteels, Deneubourg, and Goss, 1987; Kirman, 1993; Thuijsman, Peleg, Amitai, and Shmida, 1995; Keesar, Rashkovich, Cohen, and Shmida, 2002).




Figure 1: Regions of the (r_1, r_2) -plane that imply deterministic ($f^* = 0$ or 1) or randomizing ($0 < f^* < 1$) behavior, where $r_j = c_{aj}/c_{bj}$ measures the relative fecundities of action a to action b in the two states $j = 1, 2$. The asymptotes of the curved boundary line occur at $r_1 = p$ and $r_2 = q$. Values of r_1 and r_2 for which exact probability matching is optimal is given by the solid black curve.

4.3 Individually Optimal vs. Growth-Optimal Behavior

It is instructive to compare the growth-optimal behavior of (14) with the behavior that maximizes an individual's reproductive success, denoted by \hat{f} :

Proposition 4 *Under Assumptions (A1)–(A2), and if (x_a, x_b) satisfies (13), then the individually optimal behavior \hat{f} is deterministic and given by:*

$$\hat{f} = \begin{cases} 1 & \text{if } r_2 > 1 + \frac{p}{q}(1 - r_1)r_3 \\ 0 & \text{if } r_2 < 1 + \frac{p}{q}(1 - r_1)r_3 \end{cases}, \quad r_3 \equiv \frac{c_{b1}}{c_{b2}}. \quad (17)$$

For a fixed value of r_3 , the threshold in (17) that determines the optimal individual behavior is a line that divides the (r_1, r_2) -plane into two regions. For values of (r_1, r_2) above this line, $\hat{f} = 1$, and for values below this line, $\hat{f} = 0$. When $r_3 = 1$, this implies that any time the growth-optimal behavior involves randomization, it will always be at odds with the individually optimal behavior $\hat{f} = 1$. We shall return to this important special case below. Of course, from a population perspective, the growth-optimal behavior f^* is independent of r_3 , and depends only on the relative performance of the two possible choices as measured by the relative fecundities r_1 and r_2 .

We shall revisit this distinction between growth-optimal and individually optimal behavior in Section 6 when we consider the case where the randomness in offspring is idiosyncratic, i.e., (x_a, x_b) is independently and identically distributed across individuals, as well as across time. Under this alternate environment, we will show that the growth-optimal and individually optimal behaviors are identical.

5 Risk Preferences

Our binary choice model can also be used to study the evolution of risk preferences by making one of the two choices riskless. In particular, in the Bernoulli case (13), suppose that choice b yields a non-random outcome $c_{b1} = c_{b2} = c_b$ (or $r_3 = 1$). In this case, each individual is choosing between a random outcome and a certain one, where we maintain the assumption that:

$$\begin{aligned} \text{Prob}(x_a = c_{a1}, x_b = c_b) &= p \in [0, 1] \\ \text{Prob}(x_a = c_{a2}, x_b = c_b) &= 1 - p \equiv q. \end{aligned} \quad (18)$$

Without loss of generality, assume that $c_{a1} < c_{a2}$. To ensure that the choice between a and b is not trivial, we require that $c_{a1} < c_b < c_{a2}$, otherwise one choice will always dominate the other trivially. For convenience, we shall parametrize c_b as a convex combination of the risky outcomes c_{a1} and c_{a2} :

$$c_b \equiv \theta c_{a1} + (1-\theta)c_{a2} \quad , \quad \theta \in (0, 1) \quad . \quad (19)$$

When $\theta = 0$, the riskless outcome is c_{a2} , which dominates the risky choice, and when $\theta = 1$, the riskless outcome is c_{a1} which is dominated by the risky choice, hence $\theta \in (0, 1)$ covers the entire spectrum of possible risk/reward trade-offs in which neither choice dominates the other.

In Section 5.1, we derive the growth-optimal risk preferences by applying the results of Section 3 to (18). Using these preferences, in Section 5.2 we show how risk aversion emerges purely through the forces of natural selection, and derive an implied evolutionary “risk premium” that is completely independent of any notion of economic equilibria. And in Section 5.3, we show how growth-optimal behavior f^* can explain loss aversion.

5.1 Growth-Optimal Risk Preferences

Applying Proposition 3 to (18) yields:

Corollary 3 *Under Assumptions (A1)–(A2), and if (x_a, x_b) satisfies (18), then the growth-optimal behavior f^* is given by:*

$$f^* = \begin{cases} 1 & \text{if } \theta \in [\theta_o, 1) \\ (1 - \frac{p}{\theta}) \left(1 + \frac{1}{(1-\theta)(\sigma-1)}\right) & \text{if } \theta \in (p, \theta_o) \\ 0 & \text{if } \theta \in (0, p] \end{cases} \quad (20)$$

where $\theta_o \equiv p\sigma / (p\sigma + q)$, and $\sigma \equiv c_{a2} / c_{a1} > 1$.

Note that σ can be viewed as a crude measure of a 's risk, hence our choice to use the symbol typically reserved for standard deviation. As θ increases from 0 to 1, the number of offspring produced by the riskless choice b decreases from c_{a2} to c_{a1} , making the risky choice a relatively more attractive. However, for values of θ from 0 to p , the risky choice is not sufficiently attractive and the optimal behavior is to select the sure thing ($f^* = 0$). As θ

increases from p to θ_o , the optimal behavior f^* from the population perspective is to select choice a with increasingly higher probability, and when θ exceeds the threshold θ_o , the risky choice becomes so attractive relative to the sure thing that individuals always choose the risky alternative ($f^*=1$). Figure 2 graphs this relation between f^* and θ for three values of σ , and shows that it is close to piece-wise linear for $\sigma=2$ but nonlinear for $\sigma=100$. Figure 3 provides a more complete depiction of the trade-off between σ and θ in determining the growth-optimal behavior f^* .




Figure 2: The growth-optimal behavior f^* as a function of θ for $p=1/3$ and various levels of the ratio $\sigma = c_{a2}/c_{a1}$ of the two outcomes of the risky alternative a . The parameter θ determines the magnitude of the payoff $c_b = \theta c_{a1} + (1 - \theta) c_{a2}$, of the riskless alternative b .

A more transparent version of the growth-optimal behavior in Corollary 3 can be derived by restating (20) in terms of the outcomes of each of the two choices a and b . As before, let c_p and c_o denote the arithmetic and harmonic means of the a outcomes, respectively, so that

$$c_p \equiv pc_{a1} + qc_{a2} \quad \text{and} \quad c_o = \frac{1}{\frac{p}{c_{a1}} + \frac{q}{c_{a2}}} . \quad (21)$$

The values c_o and c_p correspond to the values $\theta = \theta_o$ and $\theta = p$ in Corollary 3, hence the growth-optimal behavior of an individual depends entirely on where the sure thing c_b lies in the range (c_{a1}, c_{a2}) : if $c_b \in [c_{a1}, c_o]$, then the risky choice is always growth-optimal; if $c_b \in (c_o, c_p)$, randomizing between the risky choice and the safe choice is growth-optimal; and




Figure 3: Radial plot of the growth-optimal behavior f^* for $p=1/3$ as a function of θ , which varies from 0 to 1 clockwise around the semi-circle, and $\sigma = c_{a2}/c_{a1}$, which varies from 1 to infinity from the center of the semi-circle to its perimeter. The parameter θ determines the magnitude of the payoff $c_b = \theta c_{a1} + (1 - \theta) c_{a2}$, of the riskless alternative b .

if $c_b \in [c_p, c_{a2})$, the safe choice is always growth-optimal.¹⁸ These outcomes are illustrated in Figure 4, in which the green range for c_b yields the safe choice as the growth-optimal behavior, the red range yields the risky choice, and the gray range yields randomization.

Of course, the growth-optimal behavior depends also on p , the probability of the higher-yielding outcome in a , and this dependence is implicit in the relative widths of the three ranges in Figure 4. In particular, the relative width of the interval in which the riskless choice is always growth-optimal is given by p —the more likely is the lower risky outcome c_{a1} , the greater the range of values for which the certain outcome c_b dominates the risky outcome c_a . The relative width of the interval in which the risky choice is always growth-optimal is given by $1/(1 + \sigma p/(1 - p))$, which is decreasing in both σ and p —as the riskiness of a or the likelihood of the lower outcome increases, the region in which the risky choice is always growth-optimal becomes smaller. And for a fixed probability p , the risk parameter σ determines the relative magnitude of the randomization interval as compared to the risky choice interval. Not surprisingly, for larger σ , the randomization interval is wider, implying a greater range of values of c_b for which randomizing between the risky and riskless choices is growth-optimal. We shall see another manifestation of this behavior toward risk in the next section.

¹⁸Note that the ordering $c_{a1} < c_o < c_p < c_{a2}$ always holds, since the harmonic mean is less than the arithmetic mean.




Figure 4: Growth-optimal behavior of risky/riskless choice as a function of the magnitude of the riskless outcome c_b . The green range for c_b indicates that the riskless choice is growth-optimal, the red range indicates that the risky choice is growth-optimal, and the gray region indicates that randomizing between the risky and riskless choices is growth-optimal. The relative widths of the three regions, relative to the total range $c_{a2} - c_{a1}$, are displayed to the right of each region.

5.2 Risk Aversion

Our binary choice model also shows that the property of risk aversion—the need to compensate individuals with a positive payment, i.e., a “risk premium”, to induce them to accept a fair gamble—arises quite naturally from natural selection. To see how, consider the risky/riskless case of (18) in Section 5, but now let the payoffs for the risky option a be defined relative to the riskless payoff c_b of b :

$$c_{a1} \equiv c_b - d \quad , \quad c_{a2} \equiv c_b + u \quad , \quad u, d > 0 \quad (22)$$

and consider the case in which the optimal behavior f^* is exactly $\frac{1}{2}$, so that neither choice is selected more frequently than the other. This value implies that the growth-optimal behavior is indifferent between the riskless payoff of b and the risky payoff of a , which, in turn, implies that the two choices must have the same implications for population growth. In this respect, b 's sure payoff may be viewed as the “certainty equivalent” of a , an economic concept used to measure the dollar value of random payoffs.

Assuming that $p=q=\frac{1}{2}$ so that u and d are equally likely outcomes, we then derive the implications of these parameter settings for u and d :

$$u = d + \frac{d^2}{c_b - d} . \quad (23)$$

When $u=d$, behavior is said to be “risk neutral” because the expected value of a is identical to the sure payoff of b . However, equation (23) shows that u must exceed d by a positive amount $d^2/(c_b - d)$ to be consistent with the behavior $f^* = \frac{1}{2}$. The difference between the expected values of a and b is:

$$\pi \equiv \frac{d^2}{2(c_b - d)} \quad (24)$$

which can be considered an “evolutionary risk premium”. However, unlike the risk premia of economic models of rational markets (Merton, 1980; Mehra and Prescott, 1985), which depend on the equalization of supply and demand, π arises from the fact that populations grow geometrically (see equation 6), and the factor by which the population grows, $\exp(\mathbb{E}[\log(x)])$, is always less than or equal to the expected number of offspring in a single generation, $\mathbb{E}[x]$,

due to Jensen’s Inequality.¹⁹ Therefore, risky choices will always yield lower population growth than the corresponding riskless choices with identical expected values. From an evolutionary perspective, the only sustainable behavior in which a and b are equally likely to be chosen is if the risky choice a yields a larger expected number of offspring than the riskless choice b , i.e., u must always be larger than d .

This result may also explain risk aversion in non-human animal species, often called “risk-sensitive foraging behavior” by ecologists, who have observed this behavior in organisms from bacteria to primates (Deneubourg, Aron, Goss, and Pasteels, 1987; Harder and Real, 1987; Pasteels, Deneubourg, and Goss, 1987; Hölldobler and Wilson, 1990; Kirman, 1993; Thuijsman, Peleg, Amitai, and Shmida, 1995; Smallwood, 1996; Keasar, Rashkovich, Cohen, and Shmida, 2002; Ben-Jacob, 2008). Regardless of the species, (23) shows that when c_b is very large relative to d , the evolutionary risk premium π becomes negligible since a bad outcome for a has very little impact on growth rates given the magnitude of c_b . However, when d is close to c_b , a bad outcome for a implies near sterility for that individual, hence a substantial risk premium is required to maintain the individual’s indifference between a and b .

5.3 Loss Aversion

The growth-optimal behavior can also generate loss aversion, the tendency of human subjects to take less risk when choosing between two potential gains, and to take more risk when choosing between two potential losses (Tversky and Kahneman, 1974; Kahneman and Tversky, 1979). For example, when offered the choice between investment opportunities A and B, where A is a lottery ticket paying \$1 million with a 25% probability and \$0 with 75% probability and B generates a sure profit of \$240,000, the vast majority of MIT Sloan School of Management MBA students have chosen B in many trials over the past two decades.²⁰

¹⁹ Jensen’s Inequality states that the expected value of a convex function $g(\cdot)$ of a random variable x is greater than or equal to the function of the expected value of x , or $E[g(x)] \geq g(E[x])$. The opposite inequality holds for concave functions. In a study on risk-sensitive foraging behavior, Smallwood (1996) applies Jensen’s Inequality to derive risk averse behavior by first assuming that an animal’s fitness is an increasing but concave function of its energy level, and then applying Jensen’s Inequality to foraging behavior that yields random energy levels.

²⁰ This example is a modification of original experiments conducted by Tversky and Kahneman (1974) with Stanford undergraduate students using actual cash payoffs. The MIT Sloan MBA versions were based exclusively on hypothetical classroom surveys, and the primary modification was an increase in the dollar values of the outcomes, a very telling change that was necessitated by the fact that MBA students did not exhibit loss aversion with smaller payoffs. See the discussion in the remainder of Section 5.3 for an

However, when these same subjects were offered the choice between investment opportunities A' and B', where A' is a lottery ticket yielding a loss of \$1,000,000 with 75% probability and \$0 with 25% probability and B' generates a sure loss of \$750,000, virtually all of them chose A', the more risky alternative. Such inconsistent risk attitudes have material adverse implications: the most popular choices of B and A' yield a combined outcome that is \$10,000 less than A and B', no matter how the lotteries turn out.²¹

To see how loss aversion arises in our framework, we must first consider the relation between monetary payoffs and reproductive success, i.e., we must link an individual's financial wealth to the number of offspring such wealth affords. Only when this relation is specified can we deduce the impact of natural selection on decisions involving financial gain or loss. Therefore, denote by $c(w)$ a "reproduction function" that yields the number of offspring produced by an individual with total wealth w . Basic biological and economic considerations suggest the following three properties for $c(w)$:

(A3) $c(w)$ is a continuous non-decreasing function of wealth w .

(A4) $c(w) = 0$ for all levels of wealth w below a subsistence level w_0 .

(A5) $c(w)$ is bounded above by some finite number $\bar{c} > 0$.

Assumption (A3) states that more wealth leads to more offspring, Assumption (A4) acknowledges the existence of a subsistence level of wealth below which an individual cannot produce any offspring, and Assumption (A5) reflects environmental resource constraints that place an upper limit on the number of offspring any individual can generate, irrespective of wealth.

These assumptions seem obvious and almost trivially true,²² yet they have surprisingly sharp implications for the properties of $c(w)$: they imply that $c(w)$ necessarily resembles the

explanation of this interesting difference.

²¹The combined outcome of choices B and A' is $-\$760,000$ with probability 75% and $\$240,000$ with probability 25%, while in contrast the combined outcome of choices A and B' is $-\$750,000$ with probability 75% and $\$250,000$ with probability 25%. Thus, a subject choosing B and A' expects to have an outcome that is \$10,000 worse than that corresponding to the choice of A and B'. Moreover, if the results of the two lotteries are perfectly correlated, then a subject choosing B and A' obtains a combined payoff that is \$10,000 lower in each possible state of the outcomes.

²²Perhaps the least obvious of the three is the continuity assumption in (A3), which depends, of course, on the numeraire with which wealth is measured, and how wealth interacts with the biology of reproduction. With the advent of fiat money, and given current *in vitro* fertilization technologies, we believe that continuity is a reasonable approximation to human reproduction.

S-shaped utility functions documented experimentally by Kahneman, Slovic, and Tversky (1982) and other behavioral economists in a sense made precise by the following proposition.²³

Proposition 5 *If $c(w)$ satisfies (A3)–(A5) and is twice continuously differentiable, then $c(w)$ is concave for sufficiently large values of w and convex for sufficiently small values of w . If $c(w)$ satisfies (A3)–(A5) but is not continuously differentiable, then a slightly weaker result holds: there exist values $w_1 < w_2$ such that*

$$c(\lambda w + (1 - \lambda)w_1) \geq \lambda c(w) + (1 - \lambda)c(w_1), \quad \text{for } w \ll w_1, \text{ and} \quad (25)$$

$$c(\lambda w + (1 - \lambda)w_2) \leq \lambda c(w) + (1 - \lambda)c(w_2), \quad \text{for } w \gg w_2, \quad (26)$$

where $\lambda \in [0, 1]$.

With our assumptions and Proposition 5 in place, we can now translate monetary payoffs into number of offspring through $c(w)$ and consider the impact of evolution on the behavior of individuals choosing between dollar-denominated choices in the experimental setting described above.

Consider an experiment along the lines of Kahneman and Tversky (1979) in which an individual is asked to choose between a risky investment A yielding one of two final wealth levels w_{a1} and w_{a2} which imply reproductive outcomes $c_{a1} \equiv c(w_{a1})$ and $c_{a2} \equiv c(w_{a2})$, and riskless investment B yielding a guaranteed final wealth level w_b which translates into reproductive outcome $c_b \equiv c(w_b)$. Suppose that the initial wealth of the individual is modest relative to the incremental payoffs of A, so that the total wealth obtained in either outcome of A is large. Then we have:

Corollary 4 *Under Assumptions (A1)–(A5), and if w_{a1} is sufficiently large so that $c(\cdot)$ is concave throughout the interval $[w_{a1}, w_{a2}]$, we have $c^{-1}(c_p) < w_p = pw_{a1} + qw_{a2}$, and the growth-optimal behavior is to take the safe bet whenever w_b is in the region $[c^{-1}(c_p), w_{a2}]$,*

²³An even more fundamental implication of our model of behavior is that an individual’s relative wealth—relative to others—should be more important than absolute wealth, as suggested by Duesenberry (1949) and Frank (2000). This follows trivially from Corollary 1 and the existence of a mapping $c(w)$ between wealth and fecundity. Because evolutionary success is determined solely by relative growth rates, a monotonic relation between wealth and fecundity implies that, ceteris paribus, and assuming reasonably similar functions $c(w)$ across individuals (i.e., similar biological specifications and constraints), natural selection will favor those individuals with higher relative wealth. Note that Duesenberry’s (1949) “relative income hypothesis” refers to income, not wealth. However, because wealth is highly correlated with cumulative income flows, empirical phenomena associated with relative wealth should also manifest themselves in relative income, ceteris paribus.

which is strictly larger than $[w_p, w_{a2}]$ by the incremental interval $[c^{-1}(c_p), w_p]$ that depends on the concavity of $c(w)$. For values of w_b below $c^{-1}(c_p)$, the growth optimal behavior is either to randomize or to take the risky bet according to whether w_b is larger or smaller than $c^{-1}(c_0)$, respectively, where c_0 is defined in (21).

Corollary 4 shows that if an individual is confronted with a risky investment A that involves sufficiently large gains—so large that the relevant portion of $c(w)$ is concave—and the riskless investment B is not too small relative to the payoffs from A, but possibly smaller than the expected payoff, w_p , from A, then the growth-optimal behavior will be to select the sure thing. These results show that an individual is risk averse when it comes to gains, preferring sure bets that pay less than the expected value of risky bets, if the individual begins with a modest level of wealth, if A and B both move the individual's total wealth into the concave region of $c(w)$, and if w_b is in the non-empty range $[c^{-1}(c_p), w_p]$.

Now consider two investment alternatives A' and B', where A' is a risky investment yielding one of two final wealth levels w'_{a1} and w'_{a2} which translate into reproductive outcomes $c'_{a1} \equiv c(w'_{a1})$ and $c'_{a2} \equiv c(w'_{a2})$, and B' is a riskless investment with guaranteed final wealth level w'_b which translates into reproductive outcome $c'_b \equiv c(w'_b)$, and suppose that w'_{a2} is small, implying incremental losses relative to initial wealth. Then we have:

Corollary 5 *Under Assumptions (A1)–(A5), and if $c(\cdot)$ is convex throughout the interval $[w'_{a1}, w'_{a2}]$, we have $c^{-1}(c'_p) > w'_p = pw'_{a1} + qw'_{a2}$, and the growth-optimal behavior is either to randomize or always to take the risky bet whenever w'_b is in the region $[w'_{a1}, c^{-1}(c'_p)]$, which is strictly larger than $[w'_{a1}, w'_p]$. Moreover, the region in which the growth-optimal behavior is always to take the risky bet, namely $[w'_{a1}, c^{-1}(c'_0)]$, may also be strictly larger than $[w'_{a1}, w'_p]$. A sufficient condition for this to occur is that the function $c(w)$ is invertible in $[w'_{a1}, w'_{a2}]$ and also sufficiently convex so that the function $1/c(w)$ is concave in this region.*

Corollary 5 shows that for a sufficiently extreme sure loss w'_b , individuals will always choose the risky option A', despite the fact that this choice may lead to an even greater loss w'_{a1} with probability p . In addition, individuals will always choose the risky option with some positive probability even though the expected payoff of the risky option is less than that of the safe bet, provided that the safe payoff w'_b is in the nonempty region $[w'_p, c^{-1}(c'_p)]$. Moreover, they will choose the risky option with 100% probability even if the expected payoff of the risky option is less than that of the safe bet, provided that $c(w)$ is sufficiently convex so

that $1/c(w)$ is concave and w'_b is in the nonempty region $[w'_p, c^{-1}(c'_o)]$. This type of behavior for $c(w)$ will occur, for example, whenever $c(w)$ follows a power law of the form $\alpha(\beta - w)^{-\gamma}$, for $0 < \gamma < 1$.

Loss aversion arises when more risk is taken when choosing between a sure and risky loss, and less risk is taken when choosing between a sure and risky gain. Corollaries 4 and 5 show that such behavior is clearly growth-optimal under Assumptions (A3)–(A5). However, our results go beyond this basic pattern, and imply that individuals will not simply choose always to take the safe bet or always to take the risky choice. Rather, there is also a region of w_b values, namely $[c^{-1}(c_o), c^{-1}(c_p)]$ or $[c^{-1}(c'_o), c^{-1}(c'_p)]$, where randomization is optimal. In addition, for suitably constructed extreme losses, i.e., tail risk, taking the risky bet deterministically is rarely optimal, and instead some degree of randomization is generally best. These results are illustrated graphically in Figure 5.




Figure 5: Kinked (a) and smooth (b) subsistence thresholds in reproduction functions $c(w)$ (in blue) relating the number of offspring c to monetary wealth w . In both graphs, if w_b is in the green region, then always choosing the riskless choice b is growth-optimal ($f^* = 0$); if w_b is in the red region, then always choosing the risky choice a is growth-optimal ($f^* = 1$); and if w_b is in the gray region between red and green, randomization is growth-optimal. Wealth values with primes are equal to those without primes less a common fixed amount, and the color coding is identical.

Our framework may also explain some of the inconsistencies and instabilities in the experimental evidence for loss aversion. One potential source of instability is the fact that the typical experimental design offers the same fixed set of choices to all subjects in a given experiment, a standard protocol that maintains identical “treatments” across subjects so as

to render the outcomes comparable across individuals. However, if subjects make decisions based on reproductive success, then the proper method for controlled experimentation is to offer all subjects the same choices denominated in reproductive success, not the same incremental dollar wagers. Of course, without knowledge of each subject’s net worth and reproduction function, it is impossible to construct such a controlled experiment even if it were economically viable. But by offering the same fixed-dollar wagers, two sources of variation across subjects are injected into the experimental outcomes: variation in wealth levels and variation in the reproduction function $c(w)$. These factors may explain some of the variability in the findings of loss aversion studies across venues and subject pools.

Of course, the fact that loss aversion emerges from an evolutionary process does not imply that it is optimal from an individual investor’s perspective. In fact, loss-averse behavior is routinely singled out by professional traders as counter-productive to their objectives. For example, the phenomenon of “doubling down” in the face of mounting losses is a common behavioral pattern among inexperienced traders, and the adage to “cut your losses and ride your gains” is time-honored Wall Street wisdom that is meant to correct for loss aversion. In this respect, loss aversion is yet another example of the evolutionary principle that individually optimal behavior need not coincide with the growth-optimal behavior, a distinction we address explicitly in Section 6.

6 Idiosyncratic vs. Systematic Risk

So far we have assumed that the number of offspring from actions a and b are given by the same two random variables x_a and x_b , respectively, for all individuals, i.e., fecundity is systematic, implying differences between growth-optimal and individually optimal behavior under certain conditions. In this section, we show that if uncertainty in reproduction is, instead, idiosyncratic to each individual, the growth-optimal behavior always coincides with individually optimal behavior. This distinction points to the central role that aggregate uncertainty plays in shaping the evolution of behavior and preferences.²⁴ In Section 6.1, we consider the case where fecundity is purely idiosyncratic, so that $(x_{a,i}, x_{b,i})$ are independently

²⁴We thank Arthur Robson for encouraging us to explore the distinction between systematic and idiosyncratic risk in our framework. In Robson and Samuelson (2009), they show that evolution selects for agents who maximize discounted expected utility, discounting at the sum of the population growth rate and mortality rate, when risk is idiosyncratic, but in the face of aggregate risk, the growth-optimal set of preferences involve higher discount rates that imply non-exponential discounting.

and identically distributed across individuals in a given generation. The general case of both idiosyncratic and systematic fecundity is developed in Section 6.2.

6.1 Idiosyncratic Risk

As before, denote by x_i^f the random number of offspring produced by individual i , but now suppose that:

$$x_i^f = I_i^f x_{a,i} + (1 - I_i^f) x_{b,i} \quad , \quad I_i^f \equiv \begin{cases} 1 & \text{with probability } f \\ 0 & \text{with probability } 1-f \end{cases} . \quad (27)$$

where

(A2') $(x_{a,i}, x_{b,i})$ is independently and identically distributed over time and across individuals i in a given generation

In contrast to the systematic case (1) and Assumption (A2), $(x_{a,i}, x_{b,i})$ are now assumed to be independently and identically distributed across individuals as well as across time, and I_i^f are the same Bernoulli 0/1 random variables summarizing the behavior of individual i as defined in Section 3.2.

Assumption (A2') is a seemingly small change, but it has dramatic consequences for the evolutionary dynamics of the population and growth-optimal behavior. In this case, the randomness in the number of offspring is strictly idiosyncratic in the sense that the correlation between the number of offspring for two individuals i and j is 0, even if both individuals choose the same course of action. Recall from (1) that in the systematic case, if two individuals choose the same action a , both will generate the same number of random offspring x_a , i.e., their reproductive success is perfectly correlated. Idiosyncratic fecundity implies that even if all individuals in a given population choose the same action, there will still be considerable cross-sectional variability in the number of offspring produced. This more diversified outcome has a very different set of implications for growth-optimal behavior.

In particular, in this case the total number of offspring n_t^f across all type- f individuals in generation t is:

$$n_t^f = \sum_{i=1}^{n_{t-1}^f} x_{i,t}^f = \sum_{i=1}^{n_{t-1}^f} I_{i,t}^f x_{a,i,t} + \sum_{i=1}^{n_{t-1}^f} (1 - I_{i,t}^f) x_{b,i,t} \quad (28)$$

$$n_t^f \stackrel{p}{=} n_{t-1}^f (f\mu_a + (1-f)\mu_b) \quad , \quad \mu_a \equiv \mathbb{E}[x_{a,i}] \quad , \quad \mu_b \equiv \mathbb{E}[x_{b,i}] \quad (29)$$

where (29) follows from the Law of Large Numbers applied to the sums $\sum_i I_i^f x_{a,i}/n^f$ and $\sum_i I_i^f x_{b,i}/n^f$.²⁵ The key difference between (3) and (29) is that in the latter case, both the individual's choice and the number of offspring are idiosyncratic, hence both are subject to the Law of Large Numbers. This implies that in a large population of n^f individuals, even if all individuals choose the same action a , the outcomes will vary across individuals ($x_{a,i}$) whereas in the systematic case, all individuals will receive the identical number of offspring x_a . Alternatively, in the systematic case, the number of offspring of individuals i and j are perfectly correlated, but in the idiosyncratic case, they are perfectly uncorrelated.

This difference has significant consequences when we consider the behavior of this group of type- f individuals over time. Adding a time subscript t as before to denote the population size of this group n_t^f at time t , we have:

$$n_t^f \stackrel{p}{=} n_{t-1}^f (f\mu_a + (1-f)\mu_b) \quad . \quad (30)$$

Without loss of generality, we normalize $n^f = 1$ at $t = 1$, hence the total population size in generation T is:

$$n_T^f \stackrel{p}{=} \prod_{t=1}^T n_t^f = (f\mu_a + (1-f)\mu_b)^T \quad (31)$$

$$(n_T^f)^{1/T} \stackrel{p}{=} \exp\left(\log(f\mu_a + (1-f)\mu_b)\right) \quad (32)$$

where (32) differs from (5) in that there is no expectation over x_a and x_b because the idiosyncratic nature of the cross section of individuals has eliminated the randomness in the

²⁵By construction I_i^f and $x_{k,i}$ are independent, $k = a, b$, hence their expected product is the product of their expectations, and the Kolmogorov Law of Large Numbers guarantees almost sure convergence as long as they are IID and their first moments are finite.

population. This simple expression attains its optimum at the extremes of f , hence we have:

Proposition 6 *Under Assumptions (A1) and (A2'), the growth-optimal behavior f^* is given by:*

$$f^* = \begin{cases} 1 & \text{if } \mu_a > \mu_b \\ 0 & \text{if } \mu_a \leq \mu_b \end{cases} \quad (33)$$

which coincides with individually-optimal behavior \hat{f} .

Proposition 6 states that when the randomness in fecundity is purely idiosyncratic, the growth-optimal behavior coincides with the individually optimal behavior. In contrast to the case of systematic fecundity, because the outcome of each individual's decision is independent of the outcomes of other individuals' decisions, it is exceedingly improbable for the entire population to become extinct, even if all individuals engage in identical behavior.²⁶ However, when the randomness is purely systematic, identical behavior among individuals does lead to extinction with positive probability, hence growth-optimal behavior differs from individually optimal behavior. If environmental challenges to reproductive success are systematic, the only type of behavior that can survive in the long run in our framework is some form of randomization.

The difference between idiosyncratic and systematic risk affects more than behavior; populations subjected to idiosyncratic environmental risk grow much more quickly. To see why, recall from Corollary 1 that in the systematic-risk case, the geometric-average growth rate of the growth-optimal population f^* is given by:

$$(n_T^{f^*})^{1/T} \stackrel{p}{=} \exp(\mu(f^*)) = \exp(\mathbb{E}[\log(f^*x_a + (1-f^*)x_b)]) \quad (34)$$

where f^* is given by (7). In the case of idiosyncratic risk, Proposition 6 implies that the geometric-average growth rate of the growth-optimal population f^* is given by the larger of μ_a and μ_b . Without loss of generality, assume that $\mu_a > \mu_b$. We then have the following

²⁶If we assume two states of nature as in the exact probability matching case of (11), the probability of extinction is the probability that all individuals acting individually optimally will simultaneously experience the 0-offspring outcome. Without loss of generality, if we assume $p < 1/2$, then $\hat{f} = f^* = 0$ (see Proposition 6), hence the probability of extinction is p^n which approaches 0 quickly even for modest population sizes (for example, $0.45^{50} = 4.6 \times 10^{-18}$).

inequality:

$$\exp(\mathbb{E}[\log(f^*x_a + (1-f^*)x_b)]) \leq \exp(\log \mathbb{E}[f^*x_a + (1-f^*)x_b]) \quad (35)$$

$$= f^*\mu_a + (1-f^*)\mu_b \leq \mu_a \quad (36)$$

where (35) follows from Jensen’s Inequality and (36) is a strict inequality unless $f^* = 1$ in the systematic-risk case.

In two otherwise identical populations—both of which have the same marginal distribution for (x_a, x_b) , but where (x_a, x_b) is systematic in one and idiosyncratic in the other—the idiosyncratic-risk population will grow at least as fast, and typically faster. This holds despite the fact that the expected number of offspring (μ_a, μ_b) is the same in both populations, and is driven by two distinct factors: Jensen’s Inequality and the possibility of randomizing behavior in the systematic-risk case. The former is the same mechanism at work in generating risk aversion (see Section 5.2 and footnote 19); an IID cross-section of individuals is able to approximate a “riskless” rate of growth for large populations because the Law of Large Numbers applies within a single generation, which is not the case for a population with systematic risk. The latter is simply a consequence of the fact that in populations where randomizing behavior is growth-optimal, a fraction of the population selects the lower-expected-value choice, hence this population will, by definition, grow at a slower rate than the idiosyncratic-risk population in which every individual selects the higher-expected-value choice.

To develop an appreciation for the magnitude of these effects, consider the case of exact probability matching in Section 4.1, in which $f^* = p$ where p is the probability of state 1 in (11), and without loss of generality, let $p > \frac{1}{2}$. The ratio of the geometric-average growth rates of the idiosyncratic and systematic cases is given by:

$$\left(\frac{n_T^{f^*=1}}{n_T^{f^*=p}}\right)^{1/T} \stackrel{a}{=} \frac{mp}{mp^p(1-p)^{1-p}} = \left(\frac{p}{1-p}\right)^{1-p}. \quad (37)$$

For $p = 0.6$, this ratio is 1.176, implying that the geometric-average growth rate of the idiosyncratic-risk population is 17.6% greater than the systematic-risk population. While this may seem small, it implies that after 50 generations, the population size in the idiosyncratic-risk case will be approximately 3,325 times that of the systematic-risk case.

6.2 The General Case

We now consider the general case in which the random mechanism determining fecundity contains both systematic and idiosyncratic components. Denote by x_i^f the random number of offspring produced by individual i , and now suppose that:

$$x_i^f = I_i^f x_{a,i} + (1 - I_i^f) x_{b,i} \quad , \quad I_i^f \equiv \begin{cases} 1 & \text{with probability } f \\ 0 & \text{with probability } 1 - f \end{cases} \quad (38)$$

$$x_{k,i} = y_{k,i} + z_k \quad , \quad k = a, b \quad (39)$$

where:

(A2'') $(y_{a,i}, y_{b,i})$ is independently and identically distributed across individuals and over time, and z_k is independently and identically distributed over time and identical across individuals in the same generation

In Assumption (A2''), the idiosyncratic component $(y_{a,i}, y_{b,i})$ is independently and identically distributed across individuals as well as across time, whereas the systematic component z_k is the same across individuals but independently and identically distributed across time, and I_i^f are the same Bernoulli 0/1 random variables summarizing the behavior of individual i as defined in Section 3.2.

Given that I_i^f are independently and identically distributed across the population of individuals of type f , we have the following expression for the total number of offspring n_t^f across all type- f individuals in generation t :

$$n_t^f = \sum_{i=1}^{n_{t-1}^f} x_{i,t}^f = \sum_{i=1}^{n_{t-1}^f} I_i^f (y_{a,i} + z_a) + \sum_{i=1}^{n_{t-1}^f} (1 - I_i^f) (y_{b,i} + z_b) \quad (40)$$

$$\stackrel{p}{=} n_{t-1}^f (f(\mu_a + z_a) + (1-f)(\mu_b + z_b)) \quad , \quad \mu_k \equiv \text{E}[y_{k,i}] \quad , \quad k = a, b \quad (41)$$

Note that (41) is a combination of (3) and (29), containing both stochastic and non-stochastic components, which reflects the impact of the systematic and idiosyncratic components in reproduction. This implies that as T increases without bound, the geometric-average growth

rate of the population of type- f individuals is well-approximated by:

$$(n_T^f)^{1/T} \stackrel{p}{=} \exp\left(\mathbb{E}\log(f(\mu_a + z_a) + (1-f)(\mu_b + z_b))\right) \quad (42)$$

which follows from the Law of Large Numbers as before. From this expression, we see that the value of f that maximizes the population size n_T^f is the value that maximizes the expectation:

$$\mu(f) \equiv \mathbb{E}[\log(f(\mu_a + z_a) + (1-f)(\mu_b + z_b))] , \quad (43)$$

which is virtually identical to the systematic case (6), the only difference coming from the shift in mean μ_k from the idiosyncratic component $y_{k,i}$. We then have:

Proposition 7 *Under Assumptions (A1) and (A2'), the growth-optimal behavior f^* is given by:*

$$f^* = \begin{cases} 1 & \text{if } \mathbb{E}[w_a/w_b] > 1 \text{ and } \mathbb{E}[w_b/w_a] < 1 \\ \text{solution to (45)} & \text{if } \mathbb{E}[w_a/w_b] \geq 1 \text{ and } \mathbb{E}[w_b/w_a] \geq 1 \\ 0 & \text{if } \mathbb{E}[w_a/w_b] < 1 \text{ and } \mathbb{E}[w_b/w_a] > 1 \end{cases} \quad (44)$$

where f^* is defined implicitly in the second case of (44) by:

$$0 = \mathbb{E}\left[\frac{w_a - w_b}{f^*w_a + (1 - f^*)w_b}\right] \quad (45)$$

and the expectations in (5)–(8) are with respect to the joint distribution $\Phi(w_a, w_b)$, where

$$w_k \equiv \mu_k + z_k , \quad k = a, b . \quad (46)$$

Proposition 7 shows that the general case in which an individual's action leads to random offspring with idiosyncratic and systematic components is mathematically equivalent to the systematic case with an additional shift in mean from the idiosyncratic component.

Our comparison of systematic and idiosyncratic environmental risk yields an interesting implication for populations in which the environment is heterogeneous. Consider an envi-

ronment that consists of several distinct regions, some with idiosyncratic risk and others with systematic risk, and suppose that these regions are initially populated uniformly, i.e., all possible behaviors $f \in [0, 1]$ are represented within each region. Over time, the regions that will reproduce most rapidly are those with the greatest proportion of idiosyncratic risk. This effect may give the appearance of individuals “seeking out” environmental conditions and niches in which the risks are more idiosyncratic, but is, of course, merely a consequence of Jensen’s inequality and requires no self-awareness, volition, or intelligence.

The differences across regions may also be interpreted as yet another form of group selection (Wynne-Edwards, 1962; Sober and Wilson, 1998), in which groups defined by regions of idiosyncratic risk are favored. As in the case of Proposition 1, it is the environment—specifically, the type of reproductive risk—that determines the definition of the group on which natural selection operates.

This result also provides an intriguing normative implication: to the extent that survival is the objective and the future environment is uncertain, policies that lead to greater idiosyncratic fecundity such as bio-diversity may be desirable, leading to improvements in population growth rates by $f^* \mu_a + (1 - f^*) \mu_b$.

7 Qualifications and Extensions

The link between behavior and reproductive success in a binary choice model is the key to an evolutionary explanation for several commonly observed behaviors in many animal species. While risk aversion, loss aversion, probability matching, and more general forms of randomizing behavior may seem sub-optimal for the individual, these behaviors persist over time and across many species precisely because they are optimal from a population perspective. Moreover, as environmental conditions change, the growth-optimal behavior may also change in response to new selective pressures, hence the inexplicably erratic actions of certain species may well be adaptive rather than simply irrational.

These considerations may seem more relevant for the foraging behavior of ants and bees than for human challenges such as investing in the stock market. After all, most decisions we face each day have little to do with our reproductive success, hence the relevance of f^* for economic and social behavior may be questioned. The answer lies in the degree to which evolutionary pressures have any bearing on behavior, which is summarized by the specification of $\Phi(x_a, x_b)$. For example, if the choice between a and b has no impact on

fecundity, e.g., monetary prizes that are small in comparison to an individual's net worth, then x_a and x_b will be statistically identical because either choice leaves the individual with the same reproductive prospects. In this case, the bivariate distribution $\Phi(x_a, x_b)$ reduces to a univariate distribution $\Phi(x)$, natural selection is indifferent to the choice between a and b , and f^* is indeterminate. Therefore, whether f^* is applicable to a given context is fully captured by the relation between the individuals' behaviors and their impact on reproductive success, i.e., Φ .

One concrete illustration of how Φ plays this role is the case of loss aversion. When the dollar amounts offered to subjects in the standard loss-aversion experiments are too low, loss aversion is not observed (see footnote 20). This result can be easily understood in the context of Φ : when the stakes are high, the outcomes have measurable impact on reproductive fitness (broadly defined), hence one aspect of loss aversion may be observed ($f^* = 1$). But when the stakes are too low, the implications of both choices for fitness are identical, hence the growth-optimal behavior has nothing to say about the outcome of the experiment.

It should be emphasized that the behaviors derived in our simple framework are primitive, both conceptually and from an evolutionary perspective. We have purposefully abstracted from more realistic aspects of biology and behavior such as sexual reproduction, random mutations, neuroplasticity, learning, communication, and strategic behavior to focus on those behaviors that are primordial and common to most living organisms. By definition, those common behaviors must confer significant advantages in the face of stochastic environmental conditions, otherwise they would not have survived over time and become so widespread within and across species.

Of course, more complex behaviors will arise as new species emerge and evolve. Although the specific biological manifestations of behavior are beyond the scope of our analysis, recent imaging and neurophysiological studies of decision-making under uncertainty in humans and primates (Breiter, Aharon, Kahneman, Anders, and Shizgat, 2001; Smith, Dickhaut, McCabe, and Pardo, 2002; Gold and Shadlen, 2007; Yang and Shadlen, 2007; Fehr and Camerer, 2007; Spitzer, Fischbacher, Herrnberger, Grön, and Fehr, 2007; Rangel, Camerer, and Montague, 2008; Bossaerts, 2009; Resulaj, Kiani, Wolpert, and Shadlen, 2009; Wunderlich, Rangel, and O'Doherty, 2009; Hare, Camerer, Knoepfle, O'Doherty, and Rangel, 2010), including studies of loss aversion (Kuhnen and Knutson, 2005; De Martino, Kumaran, Seymour, and Dolan, 2006; Tom, Fox, Trepel, and Poldrack, 2007), are beginning to identify

the neural mechanisms involved in these adaptations. From these studies, it is not difficult to see how sensory inputs, memory, and other neural substrates can yield a much greater variety of behaviors from which Nature selects the most advantageous, including the ability to avoid probability matching altogether (Shanks, Tunney, and Mccarthy, 2002).

To see how such mechanisms might have arisen through the forces of natural selection, our simple binary-choice model can easily be extended to allow for sexual reproduction, random mutations, and an arbitrary number n of possible actions, each with its own implications for the number of offspring, and these actions can, in turn, depend on a vector of auxiliary “state variables” \mathbf{Z} . This more complex framework can generate considerably more sophisticated types of behavior, including learning, memory and, ultimately, the emergence of intelligence. We are currently exploring such extensions, however, economists have implicitly incorporated these more sophisticated decision-making mechanisms through utility theory (Robson, 1996a; Grafen, 1999; Robson, 2001a; Robson, 2001b; Bernheim and Rangel, 2009) and have argued convincingly that utility functions are also shaped by the pressures of natural selection (Hansson and Stuart, 1990; Rogers, 1994; Robson, 1996b; Curry, 2001; Samuelson, 2001; Robson and Samuelson, 2007; Robson and Samuelson, 2009). Our current framework may be viewed as a bridge between the higher-level utility-based models of human behavior and the more primitive decision-making components that we share with other animal species.

8 Conclusion

The evolutionary origin of behavior has important implications for economics, not only in resolving the efficient-markets/behavioral-finance debate, but in providing a broader framework in which conflicts between rationality and human behavior can be resolved in an intellectually consistent manner. Specifically, much of neoclassical economic theory is devoted to deriving the aggregate implications of individually optimal behavior, i.e., maximization of expected utility or profits subject to budget or production constraints. By documenting departures from individual rationality, behavioral critics argue that rational expectations models are invalid and irrelevant. Both perspectives make valid points but are incomplete.

Animal behavior is, in fact, the outcome of multiple decision-making faculties—in many cases involving different neural substrates—that each species has developed through the course of evolution. What economists consider to be individually rational behavior is likely to emanate from the prefrontal cortex, a relatively new component of the brain on the evolu-

tionary timescale, and one that exists only in *Homo sapiens* and certain great apes. However, the human brain also contains other components such as the amygdala, a considerably older structure that is responsible for the “fight-or-flight” response. Faced with life-threatening circumstances, even the most disciplined individual may not be able to engage in individually rational behavior thanks to adaptive “hard-wired” neural mechanisms that conferred survival benefits to the species (and not necessarily to any given individual). Our analysis of idiosyncratic vs. systematic shocks to the number of offspring provides an explanation for these apparent contradictions in behavior: in the face of systematic factors in fecundity, growth-optimal behavior may differ from individually optimal behavior; in the presence of purely idiosyncratic fecundity, the two types of behavior converge.

Our framework may be useful in differentiating primitive behaviors from more refined decision-making faculties, providing a clearer map of the boundaries of rational economic theory versus instinctive behavior. For example, our results show that loss aversion is not a stable phenomenon, but depends on the relation between incremental risks and total net worth. As aggregate wealth in the economy declines, loss-averse behavior is likely to be more prevalent in the population, but during periods of prosperity, other behaviors will emerge. A better understanding of this pattern may allow consumers, investors, and policymakers to manage their risks more effectively.

While species with more highly developed nervous systems exhibit greater behavioral variation, even in these cases, primitive behaviors are still likely to be available, if not always chosen. We conjecture that such behaviors are most readily actuated under conditions similar to those of our binary choice model, namely, when outcomes are significant enough to impact reproductive fitness, broadly defined, and the effects of other variables on fitness is relatively small. These primitive behaviors may also be the basis of more modern adaptations such as boredom, thrill-seeking behavior, rebellion, innovation, and most recently, financial market bubbles and crashes.

From an evolutionary perspective, financial markets are neither efficient nor irrational—they are merely adaptive (Farmer and Lo, 1999; Lo, 2004; Lo, 2005). In short, the behaviors derived in our evolutionary framework may well be the “animal spirits” that Keynes (Keynes, 1936) singled out seven decades ago, and which is apparently still a force to be reckoned with today.

A Appendix

In this Appendix, we provide proofs for the main results of the paper and the less obvious corollaries. Before doing so, we present a brief summary of the basic properties of stochastic convergence that are used throughout our analysis; see Serfling (1980) for further details.

We begin with formal definitions of convergence in probability and distribution:

Definition A.1 *A sequence of random variables $\{X_n\}$ is said to converge in probability to X if and only if for any $\epsilon > 0$:*

$$\lim_{n \rightarrow \infty} \text{Prob}(|X_n - X| > \epsilon) = 0 \quad (\text{A.1})$$

and we denote this type of convergence by the expressions:

$$\text{plim}_{n \rightarrow \infty} X_n = X \quad \text{or} \quad X_n \xrightarrow{p} X . \quad (\text{A.2})$$

If two sequences of random variables $\{X_n\}$ and $\{Y_n\}$ satisfy the relation $X_n - Y_n \xrightarrow{p} 0$, they are said to be “equal in probability” and we denote this relation as $X_n \stackrel{p}{=} Y_n$. Under this definition, we can write $\text{plim}_{n \rightarrow \infty} X_n = X$ as $X_n \stackrel{p}{=} X$.

Definition A.2 *A sequence of random variables $\{X_n\}$ with distribution functions $\{F_n(x)\}$ is said to converge in distribution to X with distribution function $F(x)$ if and only if:*

$$\lim_{n \rightarrow \infty} F_n(x) = F(x) \quad (\text{A.3})$$

for each continuity point x of $F(\cdot)$, and we denote this type of convergence by the expression:

$$X_n \xrightarrow{d} X . \quad (\text{A.4})$$

These notions of convergence satisfy the following properties:

Proposition A.1 *If $X_n \xrightarrow{p} X$, $Y_n \xrightarrow{p} Y$, and $Z_n \xrightarrow{d} Z$, then:*

$$X_n + Y_n \xrightarrow{p} X + Y \quad (\text{A.5})$$

$$X_n Y_n \xrightarrow{p} XY \quad (\text{A.6})$$

$$g(X_n) \xrightarrow{p} g(X) \quad , \quad \text{for any } g(\cdot) \text{ continuous at } X \quad . \quad (\text{A.7})$$

If X is a finite constant, then:

$$X_n + Z_n \xrightarrow{d} X + Z \quad (\text{A.8})$$

$$X_n Z_n \xrightarrow{d} XZ \quad (\text{A.9})$$

$$Z_n/X_n \xrightarrow{d} Z/X \quad , \quad \text{if } X \neq 0 \quad . \quad (\text{A.10})$$

A.1 Proof of Proposition 1

This follows from the first and second derivatives of (6). Because the second derivative is strictly negative, there is exactly one maximum value obtained in the interval $[0, 1]$. The values of the first derivative of $\mu(f)$ at the endpoints are given by:

$$\mu'(0) = E[x_a/x_b] - 1 \quad , \quad \mu'(1) = 1 - E[x_b/x_a] \quad . \quad (\text{A.11})$$

If $\mu'(0)$ and $\mu'(1)$ are both positive or both negative, then $\mu(f)$ increases or decreases, respectively, throughout the interval and the maximum value is attained at $f=1$ or $f=0$, respectively. Otherwise, $f=f^*$ is the unique point in the interval for which $\mu'(f)=0$, where f^* is defined in (8), and it is at this point that $\mu(f)$ attains its maximum value. The expression (7) summarizes the results of these observations for the various possible values of $E[x_a/x_b]$ and $E[x_b/x_a]$. Note that the case $E[x_a/x_b] \leq 1$ and $E[x_b/x_a] \leq 1$ is not considered because this set of inequalities implies that $\mu'(0) \leq 0$ and $\mu'(1) \geq 0$, which is impossible since $\mu''(f)$ is strictly negative.

A.2 Proof of Corollary 1

The Kolmogorov Law of Large Numbers (Serfling, 1980, Ch. 1.8) implies that $T^{-1} \log n_T^f$ converges almost surely to its expectation $\mu(f)$, hence the exponential of the former converges almost surely to the exponential of the latter (see White (1984, Proposition 2.11)). Note that

almost-sure convergence is stronger than convergence in probability, but we use the latter concept in this and later results in anticipation of generalizations in which $\{x_{a,t}, x_{b,t}\}$ are not necessarily independently and identically distributed over time. As long as $\{x_{a,t}, x_{b,t}\}$ is stationary and ergodic, a Weak Law of Large Numbers applies, implying convergence in probability (White, 1984, Ch. 5).

The second part of the corollary follows from the fact that

$$(n_T^f)^{1/T} = \exp\left(T^{-1} \sum_{t=1}^T \log(fx_{a,t} + (1-f)x_{b,t})\right) \stackrel{p}{=} \exp(\mu(f)) \quad (\text{A.12})$$

$$\left(\frac{n_T^{f'}}{n_T^{f^*}}\right)^{1/T} \stackrel{p}{=} \exp\left((\mu(f') - \mu(f^*))\right) \rightarrow 0 \quad (\text{A.13})$$

where the equality in probability in (A.12) follows from the Kolmogorov Law of Large Numbers and the continuity of the exponential function, and (A.13) follows from applying Proposition A.1 to the ratio $(n_T^{f'})^{1/T}/(n_T^{f^*})^{1/T}$ and the fact that $\mu(f') < \mu(f^*)$ for any $f' \neq f^*$ due to the optimality of f^* .

A.3 Proof of Proposition 2

This result follows directly from the Lindeberg-Levy central limit theory applied to $T^{-1} \log n_T^f$ (Serfling, 1980, Ch. 1.9).

A.4 Proof of Proposition 3

This result follows from the fact that the expectation (6) is given by:

$$\mu(f) = p \log(fc_{a1} + (1-f)c_{b1}) + q \log(fc_{a2} + (1-f)c_{b2}) \quad (\text{A.14})$$

and taking the derivative of this function and solving for $\mu'(f)=0$ yields the unique solution $f=f^*$, where f^* is as defined in the second case of (14). Whenever this value of f^* lies in the interval $[0, 1]$, it is the optimal value of f . Otherwise, analysis of the sign of the first derivative of $\mu(f)$ at each endpoint of the interval $[0, 1]$ shows that the optimal value f^* is either 0 or 1 as described in the first and third cases of (14).

A.5 Proof of Corollary 2

Observe that the intermediate expression for f^* in (14) can be rewritten as:

$$f^* = p + p \left(\frac{r_2}{1-r_2} \right) + q \left(\frac{1}{1-r_1} \right) .$$

The second term on the right is $O(r_2)$, and the third term on the right is $O(1/r_1)$.

A.6 Proof of Proposition 4

In particular, because an individual's expected number of offspring $E[x_i]$:

$$E[x_i] = p(fc_{a1} + (1-f)c_{b1}) + q(fc_{a2} + (1-f)c_{b2}) \quad (\text{A.15})$$

is a monotone function of f , an individual seeking to maximize $E[x_i]$ will select f to be 0 or 1, depending on which of these two extremes yields a higher expectation, as specified in (17).

A.7 Proof of Proposition 5

If $c(w)$ is twice continuously differentiable, then the derivative of $c(w)$ is everywhere non-negative because of assumption (A3), and tends toward zero as w becomes sufficiently large or sufficiently small because of assumptions (A4) and (A5). It must also be positive in some region so that $c(w)$ increases from 0 to a positive number. As a result, the derivative of $c(w)$ is increasing for sufficiently small values of w and decreasing for sufficiently large values, which, in turn, implies that $\tilde{c}(w)$ is concave for sufficiently large w and convex for sufficiently small w .

If $c(w)$ is not continuously differentiable, then pick any two values w'_1 and w'_2 such that $c(w'_1) < c(w'_2)$ and such that $c(w)$ is not linear and increasing in a neighborhood of either point. (It must be possible to pick such points, because if not $c(w)$ would always be linear and increasing as w becomes small or large without bound, and so either assumption (A4) or (A5) would necessarily be violated.) Consider the segment connecting the point on the graph of $c(w)$ at $w = w'_1$ to the point at $w = w'_2$, and let w_1 be the smallest value of w greater than w'_1 at which the segment intersects the graph, and let w_2 be the largest value of w less than w'_2 at which the segment intersects the graph. By construction, $w'_1 < w_1 \leq w_2 < w'_2$,

and w_1 satisfies (25) for all $w \leq w'_1$, and w_2 satisfies (26) for all $w \geq w'_2$.

A.8 Proof of Corollary 4

It was shown in Section 5.1 that the growth-optimal behavior is to take the risky bet, to randomize, or to take the safe bet, according to whether c_b is in $[c_{a1}, c_0]$, $[c_0, c_p]$, or $[c_p, c_{a2}]$, respectively. The corollary follows if these c values are transformed into w values via the function $c^{-1}(\cdot)$, and the fact that $c^{-1}(c_p) < w_p$ follows because $c(w)$ is concave throughout the interval $[w_{a1}, w_{a2}]$.

A.9 Proof of Corollary 5

As with the proof of Corollary 4, we note that the growth-optimal behavior is to take the risky bet, to randomize, or to take the safe bet, according to whether c'_b is in $[c'_{a1}, c'_0]$, $[c'_0, c'_p]$, or $[c'_p, c'_{a2}]$, respectively. These values can be transformed into w values via the function $c^{-1}(\cdot)$, and the fact that $c^{-1}(c'_p) > w'_p$ follows because $c(w)$ is concave throughout the interval $[w'_{a1}, w'_{a2}]$. If $c(w)$ is invertible in this region, and if $1/c(w)$ is concave throughout this region as well, then

$$\frac{1}{c'_0} = \frac{p}{c'_{a1}} + \frac{q}{c'_{a2}} \leq \frac{1}{c(w'_p)}.$$

Taking reciprocals of this inequality and applying the function c^{-1} , we have $c^{-1}(c'_0) \geq w'_p$. Therefore, the interval $[w'_{a1}, c^{-1}(c'_0)]$, which is the interval for which the risky choice B' is always optimal, contains the point w'_p . As a result, the final assertion of the corollary follows.

References

- Anderson, P., K. Arrow, and D. Pines (eds.), 1988, *The Economy as an Evolving Complex System*. Addison-Wesley, Reading, MA.
- Arrow, K. J., and S. A. Levin, 2009, “Intergenerational resource transfers with random offspring numbers,” *Proc. Nat. Acad. Sci.*, 106, 13702–13706.
- Arthur, B., J. Holland, B. LeBaron, R. Palmer, and P. Tayler, 1997, “Asset Pricing Under Endogenous Expectations in an Artificial Stock Market,” in *The Economy as an Evolving Complex System II*, ed. by B. Arthur, S. Durlauf, and D. Lane. Addison Wesley, Reading, MA.
- Barkow, J. H., L. Cosmides, and J. Tooby, 1992, *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford University Press, New York.
- Baum, W. M., 1974, “On two types of deviation from the matching law: Bias and under-matching,” *Journal of the Experimental Analysis of Behavior*, 22, 231–242.
- Becker, G., 1976, “Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology,” *Journal of Economic Literature*, 14, 817–826.
- Behrend, E. R., and M. E. Bitterman, 1961, “Probability-Matching in the Fish,” *American Journal of Psychology*, 74, 542–551.
- Ben-Jacob, E., 2008, “Social Behavior of Bacteria: From Physics to Complex Organizations,” *European Physics Journal B*, pp. 1–8.
- Bernheim, B., and A. Rangel, 2009, “Beyond Revealed Preference: Theoretic Foundations for Behavioral Economics,” *Quarterly Journal of Economics*, 124, 51–104.
- Bernstein, P., 1998, “Why The Efficient Market Offers Hope To Active Management,” in *Economics and Portfolio Strategy*. Peter Bernstein, Inc., New York.
- Bitterman, M. E., J. Wodinsky, and D. K. Candland, 1958, “Some Comparative Psychology,” *Am. J. Psychol.*, 71, 94–110.
- Bossaerts, P., 2009, “What Decision Neuroscience Teaches Us About Financial Decision Making,” *Annual Review of Financial Economics*, 1, 383–404.
- Bradshaw, C. M., E. Szabadi, and P. Bevan, 1976, “Behavior of humans in variable-interval schedules of reinforcement,” *Journal of the Experimental Analysis of Behavior*, 26, 135–141.
- Breiter, H. C., I. Aharon, D. Kahneman, D. Anders, and P. Shizgat, 2001, “Functional Imaging of Neural Responses to Expectancy and Experience of Monetary Gains and Losses,” *Neuron*, 30, 619–639.

- Buss, D. M., 2004, *Evolutionary Psychology: The New Science of the Mind*. Pearson, Boston, MA.
- Caraco, T., 1980, "On foraging time allocation in a stochastic environment," *Ecology*, 61, 119–128.
- Cooper, W. S., and R. H. Kaplan, 1982, "Adaptive coin-flipping: A decision theoretic examination of natural selection for random individual variation," *Journal of Theoretical Biology*, 94, 135–151.
- Cosmides, L., and J. Tooby, 1994, "Better Than Rational: Evolutionary Psychology and the Invisible Hand," *American Economic Review*, 84, 327–332.
- Cover, T., and J. Thomas, 1991, *Elements of Information Theory*. John Wiley & Sons, New York.
- Curry, P. A., 2001, "Decision making under uncertainty and the evolution of interdependent preferences," *Journal of Economic Theory*, 98, 357–369.
- Davison, M., and D. McCarthy, 1988, *The matching law: A research review*. Lawrence Erlbaum, Hillsdale, NJ.
- Dawkins, R., 1976, *The Selfish Gene*. Oxford University Press, Oxford, UK.
- De Martino, B., D. Kumaran, B. Seymour, and R. J. Dolan, 2006, "Frames, Biases, and Rational Decision-Making in the Human Brain," *Science*, 313, 684.
- Dempster, E. R., 1955, "Maintenance of Genetic Heterogeneity," *Cold Spring Harbor Symp. Quant. Biol.*, 20, 25–32.
- Deneubourg, J. L., S. Aron, S. Goss, and J. M. Pasteels, 1987, "Error, communication and learning in ant societies," *European Journal of Operational Research*, 30(2), 168–172.
- Duesenberry, J. S., 1949, *Income, Saving, and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, MA.
- Ehrlich, P. R., and S. A. Levin, 2005, "The Evolution of Norms," *PLoS Biology*, 3, e194.
- Eldredge, N., and S. Gould, 1972, "Punctuated Equilibria: An Alternative to Phyletic Gradualism," in *Models in Paleobiology*, ed. by T. Schopf. Freeman Cooper, San Francisco, CA, pp. 82–115.
- Fama, E., 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance*, 25, 383–417.
- Farmer, J. D., and A. W. Lo, 1999, "Frontiers of Finance: Evolution and Efficient Markets," *Proc. Nat. Acad. Sci.*, 96, 9991–9992.

- Fehr, E., and C. Camerer, 2007, "Social neuroeconomics: the neural circuitry of social preferences," *TRENDS in Cognitive Sciences*, 11, 419–427.
- Frank, R., 2000, *Luxury Fever: Money and Happiness in an Era of Excess*. Princeton University Press, Princeton, NJ.
- Frank, R. H., and P. J. Cook, 1995, *The Winner-Take-All Society*. Free Press, New York.
- Frank, S. A., and M. Slatkin, 1990, "Evolution in a Variable Environment," *American Naturalist*, 136, 244–260.
- Fretwell, S. D., 1972, *Populations in a Seasonal Environment*. Princeton University Press, Princeton, NJ.
- Gigerenzer, G., 2000, *Adaptive Thinking: Rationality in the Real World*. Oxford University Press, New York.
- Gold, J. I., and M. N. Shadlen, 2007, "The neural basis of decision making," *Ann Rev Neurosci*, 30, 535–574.
- Graf, V., D. H. Bullock, and M. E. Bitterman, 1964, "Further Experiments on Probability-Matching in the Pigeon," *J. Exp. Anal. Behav.*, 7, 151–157.
- Grafen, A., 1999, "Formal Darwinism, the individual-as-maximizing-agent analogy and bet-hedging," *Proceedings of the Royal Society of London, Series B*, 266, 799–803.
- Grant, D. A., H. W. Hake, and J. P. Hornseth, 1951, "Acquisition and Extinction of Verbal Conditioned Responses with Differing Percentages of Reinforcement," *J. Exp. Psychol.*, 42, 1–5.
- Hake, H. W., and R. Hyman, 1953, "Perceptions of the Statistical Structure of a Random Series of Binary Symbols," *J. Exp. Psychol.*, 45, 64–74.
- Hamilton, W. D., 1964, "The genetical evolution of social behavior I and II," *Journal of Theoretical Biology*, 7, 1–52.
- , 1971, "Geometry for the Selfish Herd," *Journal of Theoretical Biology*, 31, 295–311.
- Hansson, I., and C. Stuart, 1990, "Malthusian selection of preferences," *American Economic Review*, 80, 529–544.
- Harder, L. D., and L. A. Real, 1987, "Why are bumble bees risk averse?," *Ecology*, 68, 1104–1108.
- Hare, T., C. Camerer, D. Knoepfle, J. O'Doherty, and A. Rangel, 2010, "Value computations in ventral medial prefrontal cortex during charitable decision making incorporate input from regions involved in social cognition," *Journal of Neuroscience*, 30, 583–590.

- Herrnstein, R., and D. Prelec, 1991, "Melioration: A theory of distributed choice," *Journal of Economic Perspectives*, 5, 137–156.
- Herrnstein, R. J., 1961, "Relative and absolute strength of responses as a function of frequency of reinforcement," *Journal of the Experimental Analysis of Behaviour*, 4, 267–272.
- Herrnstein, R. J., 1997, *The Matching Law*. Harvard University Press, Cambridge, Massachusetts.
- Hirshleifer, D., and G.-Y. Luo, 2001, "On the Survival of Overconfident Traders in a Competitive Securities Market," *Journal of Financial Markets*, 4, 73–84.
- Hirshleifer, J., 1977, "Economics from a Biological Viewpoint," *Journal of Law and Economics*, 20, 1–52.
- Hölldobler, B., and E. O. Wilson, 1990, *The Ants*. Belknap Press, Cambridge, MA.
- Horne, P. J., and C. F. Lowe, 1993, "Determinants of human performance on concurrent schedules," *Journal of the Experimental Analysis of Behavior*, 59, 29–60.
- Jensen, M. C., 1978, "Some Anomalous Evidence Regarding Market Efficiency," *Journal of Financial Economics*, 6, 95–101.
- Kahneman, D., P. Slovic, and A. Tversky, 1982, *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, UK.
- Kahneman, D., and A. Tversky, 1979, "Prospect theory: An analysis of decision under risk," *Econometrica*, 47(2), 263–292.
- Keasar, T., E. Rashkovich, D. Cohen, and A. Shmida, 2002, "Bees in two-armed bandit situations: foraging choices and possible decision mechanisms," *Behavioral Ecology*, 13, 757–765.
- Kelly, J. L., 1956, "A New Interpretation of Information Rate," *Bell System Technical Journal*, 35, 917–926.
- Keynes, J. M., 1936, *The General Theory of Employment, Interest and Money*. Macmillan, London.
- Kirman, A., 1993, "Ants, rationality, and recruitment," *Quarterly Journal of Economics*, 108(1), 137–156.
- Kogan, L., S. Ross, J. Wang, and M. Westerfield, 2006, "The Price Impact Survival and Survival of Irrational Traders," *Journal of Finance*, 61, 195–229.
- Kogler, C., and A. Kühberger, 2007, "Dual Process Theories: A Key for Understanding the Diversification Bias?," *J. Risk Uncertainty*, 34, 145–154.

- Kuhnen, C. M., and B. Knutson, 2005, "The neural basis of financial risk taking," *Neuron*, 47, 763.
- Levin, S. A., D. Cohen, and A. Hastings, 1984, "Dispersal Strategies in Patchy Environments," *Theoretical Population Biology*, 26, 165–191.
- Levin, S. A., H. C. Muller-Landau, R. Nathan, and J. Chave, 2003, "The Ecology and Evolution of Seed Dispersal: A Theoretical Perspective," *Annual Review of Ecology, Evolution, and Systemics*, 34, 575–604.
- Levins, R., 1969, *Evolution in Changing Environments: Some Theoretical Explorations*. Princeton University Press, Princeton, NJ.
- Lewontin, R. C., and D. Cohen, 1969, "On population growth in a randomly varying environment," *Proc. Nat. Acad. Sci.*, 62, 1056–1060.
- Lo, A. W., 2004, "The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective," *Journal of Portfolio Management*, 30, 15–29.
- , 2005, "Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis," *Journal of Investment Consulting*, 7, 21–44.
- Luo, G.-Y., 1995, "Evolution and Market Competition," *Journal of Economic Theory*, 67, 223–250.
- Malthus, T. R., 1826, *An Essay on the Principle of Population, 6th Edition*. John Murray and Library of Economics and Liberty, London.
- Mangel, M., and C. Clark, 1988, *Dynamic Modeling in Behavioral Ecology*. Princeton University Press, Princeton, NJ.
- Manski, C., 1975, "Maximum score estimation of the stochastic utility model of choice," *Journal of Econometrics*, 3, 205–228.
- Manski, C. F., and D. McFadden (eds.), 1981, *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press, Cambridge, MA.
- Maynard Smith, J., 1982, *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK.
- , 1984, "Game Theory and the Evolution of Behaviour," *Behavioral and Brain Sciences*, 7, 95–125.
- McFadden, D., 1973, "Conditional logit analysis of qualitative choice behavior," in *Frontiers in Econometrics*, ed. by P. Zarembka. Academic Press, New York.
- McGuffin, P., B. Riley, and R. Plomin, 2001, "Toward Behavioral Genomics," *Science*, 291, 1232–1249.

- Mehra, R., and E. Prescott, 1985, "The Equity Premium Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- Merton, R. C., 1980, "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323–361.
- Mirollo, R. E., and S. H. Strogatz, 1990, "Synchronization of pulse-coupled biological oscillators," *SIAM Journal on Applied Mathematics*, 50, 1645–1662.
- Nelson, R., and S. Winter, 1982, *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.
- Niederhoffer, V., 1997, *Education of a Speculator*. John Wiley & Sons, New York.
- Pasteels, J. M., J. L. Deneubourg, and S. Goss, 1987, "Self-organization mechanisms in ant societies. I: Trail recruitment to newly discovered food sources," *Experientia. Supplementum*.
- Pinker, S., 1979, "Formal Models of Language Learning," *Cognition*, 7, 217–283.
- , 1991, "Rules of Language," *Science*, 253, 530–535.
- , 1994, *The Language Instinct: How the Mind Creates Language*. William Morrow and Company, New York.
- Plomin, R., 1990, *Nature and Nurture: An Introduction to Human Behavioral Genetics*. Brooks/Cole Publishing Company, Pacific Grove, CA.
- Plomin, R., M. Owen, and P. McGuffin, 1994, "The Genetic Basis of Complex Human Behaviors," *Science*, 264, 1733–1739.
- Rangel, A., C. Camerer, and R. Montague, 2008, "A framework for studying the neurobiology of value-based decision-making," *Nature Reviews Neuroscience*, 9, 545–556.
- Real, L., 1980, "Fitness, uncertainty, and the role of diversification in evolution and behavior," *American Naturalist*, 115, 623–638.
- Real, L., and T. Caraco, 1986, "Risk and foraging in stochastic environments," *Annual Review of Ecology and Systematics*, 17, 371–390.
- Resulaj, A., R. Kiani, D. M. Wolpert, and M. N. Shadlen, 2009, "Changes of mind in decision-making," *Nature*, advance online publication.
- Robson, A. J., 1996a, "A biological basis for expected and non-expected utility," *Journal of Economic Theory*, 68, 397–424.
- , 1996b, "The evolution of attitudes to risk: Lottery tickets and relative wealth," *Games and Economic Behavior*, 14, 190–207.

- , 2001a, “The biological basis of economic behavior,” *Journal of Economic Literature*, 39, 11–33.
- , 2001b, “Why would nature give individuals utility functions?,” *Journal of Political Economy*, 109, 900–914.
- Robson, A. J., and L. Samuelson, 2007, “The evolution of intertemporal incentives,” *American Economic Review*, 97, 492–495.
- , 2009, “The Evolution of Time Preference with Aggregate Uncertainty,” *American Economic Review*, 99, 1,925–1,953.
- Robson, A. J., and B. Szentes, 2008, “Evolution of time preference by natural selection: Comment,” *American Economic Review*, 98, 1178–1188.
- Rogers, A. R., 1994, “Evolution of time preference by natural selection,” *American Economic Review*, 84, 460–481.
- Rubenstein, D., 1972, “Risk, uncertainty and evolutionary strategies,” in *Current Problems in Sociobiology*, ed. by K. C. S. Group. Cambridge University Press, Cambridge, UK, pp. 91–111.
- Samuelson, L., 2001, “Introduction to the evolution of preferences,” *Journal of Economic Theory*, 97, 225–230.
- Samuelson, P. A., 1965, “Proof that Properly Anticipated Prices Fluctuate Randomly,” *Industrial Management Review*, 6, 41–49.
- , 1971, “The ‘fallacy’ of maximizing the geometric mean in long sequences of investing or gambling,” *Proc. Nat. Acad. Sci.*, 68, 2493–2496.
- Schoemaker, P. J. H., 1982, “The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations,” *Journal of Economic Literature*, 20, 529–563.
- Schumpeter, J., 1939, *Business Cycles: A Theoretical, Historical, And Statistical Analysis of the Capitalist Process*. McGraw-Hill, New York.
- Seger, J., and H. J. Brockmann, 1987, “What is bet-hedging?,” *Oxford Surveys of Evolutionary Biology*, 4, 182–211.
- Serfling, R., 1980, *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, New York.
- Shanks, D. R., R. J. Tunney, and J. D. McCarthy, 2002, “A re-examination of probability matching and rational choice,” *Journal of Behavioral Decision Making*, 15(3), 233–250.
- Slatkin, M., 1974, “Hedging one’s evolutionary bets,” *Nature*, 250, 704–705.

- Smallwood, P., 1996, “An introduction to risk sensitivity: The use of Jensen’s Inequality to clarify evolutionary arguments of adaptation and constraint,” *American Zoologist*, 36, 392–401.
- Smith, K., J. Dickhaut, K. McCabe, and J. Pardo, 2002, “Neuronal Substrates for Choice Under Ambiguity, Risk, Gains and Losses,” *Management Science*, 48, 711–718.
- Sober, E., and D. S. Wilson, 1998, *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Harvard University Press, Cambridge, MA.
- Spitzer, M., U. Fischbacher, B. Herrnberger, G. Grön, and E. Fehr, 2007, “The Neural Signature of Social Norm Compliance,” *Neuron*, 56, 185–196.
- Stephens, D. W., and J. R. Krebs, 1986, *Foraging Theory*. Princeton University Press, Princeton, NJ.
- Strogatz, S. H., and I. Stewart, 1993, “Coupled oscillators and biological synchronization,” *Scientific American*, 269, 102–109.
- Thaler, R. H. (ed.), 1993, *Advances in Behavioral Finance*. Russell Sage Foundation, New York.
- Thuijsman, F., B. Peleg, M. Amitai, and A. Shmida, 1995, “Automata, matching and foraging behavior of bees,” *J. Theor. Biol.*, 175, 305–316.
- Thurstone, L. L., 1927, “A law of comparative judgment,” *Psychological Review*, 34, 273–286.
- Tom, S. M., C. R. Fox, C. Trepel, and R. A. Poldrack, 2007, “The neural basis of loss aversion in decision-making under risk,” *Science*, 315, 515–518.
- Tooby, J., and L. Cosmides, 1995, “Conceptual foundations of evolutionary psychology,” in *The Handbook of Evolutionary Psychology*, ed. by J. H. Barkow, L. Cosmides, and J. Tooby. John Wiley & Sons, Hoboken, NJ, pp. 5–67.
- Trivers, R. L., 1971, “The Evolution of Reciprocal Altruism,” *Quarterly Review of Biology*, 46, 35–57.
- , 1985, *Social Evolution*. Benjamin/Cummings, Menlo Park, CA.
- , 2002, *Natural Selection and Social Theory: Selected Papers of Robert L. Trivers*. Oxford University Press, Oxford, UK.
- Tversky, A., and D. Kahneman, 1974, “Judgment under uncertainty: Heuristics and biases,” *Science*, 185(4157), 1124–1131.
- , 1981, “The Framing of Decisions and the Psychology of Choice,” *Science*, 211, 453–458.

- von Neumann, J., and O. Morgenstern, 1944, *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.
- Vulcan, N., 2000, “An Economist’s Perspective on Probability Matching,” *Journal of Economic Surveys*, 14, 101–118.
- Waldman, M., 1994, “Systematic Errors and the Theory of Natural Selection,” *American Economic Review*, 84, 482–497.
- Weibull, J. W., 1995, *Evolutionary Game Theory*. MIT Press, Cambridge, MA.
- White, H., 1984, *Asymptotic Theory for Econometricians*. Academic Press, New York.
- Wilson, E. O., 1975, *Sociobiology: The New Synthesis*. Harvard University Press, Cambridge, MA.
- Wilson, E. O., and W. H. Bossert, 1990, *A Primer of Population Biology*. Sinauer Associates, Inc., Sunderland, Massachusetts.
- Woolverton, W. L., and J. K. Rowlett, 1998, “Choice maintained by cocaine or food in monkeys: Effects of varying probability of reinforcement,” *Psychopharmacology*, 138, 102–106.
- Wright, S., 1968, *Evolution and the Genetics of Populations, Volume 1*. University of Chicago Press, Chicago, IL.
- Wunderlich, K., A. Rangel, and J. O’Doherty, 2009, “Neural computations underlying action-based decision making in the human brain,” *Proceedings of the National Academy of Sciences*, 106, 17199–17204.
- Wynne-Edwards, V., 1962, *Animal Dispersion in Relation to Social Behavior*. Oliver & Boyd, London.
- Yang, T., and M. N. Shadlen, 2007, “Probabilistic Reasoning by Neurons,” *Nature*, 447, 1075–1080.
- Young, J. S., 1981, “Discrete-Trial Choice in Pigeons: Effects of Reinforcer Magnitude,” *J. Exp. Anal. Behav.*, 35, 23–29.