

Information, evolution and utility

LARRY SAMUELSON

Department of Economics, University of Wisconsin–Madison

JEROEN SWINKELS

Olin School of Business, Washington University in St. Louis

Human utility embodies a number of seemingly irrational aspects. The leading example in this paper is that utilities often depend on the presence of salient unchosen alternatives. Our focus is to understand *why* an evolutionary process might optimally lead to such seemingly dysfunctional features in our motivations and to derive implications for the nature of our utility functions.

KEYWORDS. Evolution, information, internal conflict, self control, temptation, time consistency, utility.

JEL CLASSIFICATION. C70, C72, D80, D82.

1. INTRODUCTION

Individual choice is often accompanied by internal conflict. For example, presently-biased preferences lead to a “multiple selves” setting rich with self-control problems.¹ Psychologists have studied the phenomenon on which this paper focuses: the utility of a choice can depend importantly on the set of salient *alternative* choices. A salad may be less attractive when presented within the sight and smell of a grilling steak. An affair that one would never contemplate in the cold light of day becomes difficult to avoid in the heat of the moment. Cashews are hard to resist on the coffee table, but easy to leave in the pantry. Such *choice-set dependent* preferences also generate self-control problems.²

Larry Samuelson: LarrySam@ssc.wisc.edu

Jeroen Swinkels: Swinkels@wustl.edu

First draft, September 2000. We thank Douglas Bernheim, Ken Binmore, Jeff Ely, Faruk Gul, John Nachbar, Robert Pollak, Antonio Rangel, Phil Reny, Bill Sandholm, and several referees for helpful comments and discussions. We thank the National Science Foundation and the Boeing Center for Technology and Information Management for financial support.

¹See [Ainslie \(1992\)](#), [Loewenstein and Prelec \(1992\)](#), and [Loewenstein and Thaler \(1989\)](#) for treatments of presently-biased preferences. See [Rubinstein \(2003\)](#) for an alternative perspective. Early studies of present bias and self control by [Pollak \(1968\)](#), [Schelling \(1984\)](#), and [Strotz \(1956\)](#) have engendered a large literature. For a few examples, see [Elster \(1985\)](#), [O'Donoghue and Rabin \(1999a,b\)](#), and [Thaler and Shefrin \(1981\)](#).

²[Gardner and Lowinson \(1993\)](#), [Loewenstein \(1996\)](#), [Mischel et al. \(1992\)](#), and [Siegel \(1979\)](#) examine the importance of salient alternatives. The resulting choice-set dependent preferences have been modeled and studied experimentally by [Tversky and Simonson \(1993\)](#) and [Shafir et al. \(1993\)](#). [Gul and Pesendorfer \(2001\)](#) present a model centered on the assumption that resisting tempting alternatives is costly. [Laibson \(2001\)](#) examines a model in which instantaneous utilities adjust in response to external cues. The focus of

Why would evolution lead to such seemingly irrational preferences? Wouldn't a more "rational" agent be more likely to survive and reproduce? Couldn't evolution just give the agent utilities for steak and salad that reflect the evolutionary benefits of their consumption, and then let him figure out what to do? Why build an agent who avoids situations that might lead to an affair, but falls into temptation when one is available? If an intense desire for cashews sitting in a bowl is evolutionarily beneficial, why build an agent who is content to leave them in the pantry? At a more basic level, given that successful descendants are the currency of evolutionary success, why do people have utility for anything else?

This paper addresses these questions. Our guiding principle is that to understand our utility function, we must think through the constraints on what evolution can do in designing us to make good decisions. We view one such constraint as paramount: *the difficulty in equipping us with an accurate prior understanding of the causal and statistical structure of the world.*³

Using a stylized model (presented in [Section 2](#)), [Section 3.1](#) begins with the question of why agents attach utility to intermediate actions rather than simply to the evolutionary goal of successful descendants. We show that if the agent fully understands the causal and statistical structure of the world, the utility function "maximize the expected number of your descendants" does strictly better than one that puts weight on intermediate actions like eating and having sex. In the absence of such a perfect prior understanding of the world, however, there is evolutionary value in placing utility on intermediate actions.⁴

What determines the utilities that evolution attaches to intermediate actions? One might think of the utility of a steak as reflecting the average evolutionary impact of its consumption. [Section 3.2](#) explores how this intuition is misleading. First, the utility of an action should reflect not its evolutionary value, but rather that part of the evolutionary value not already captured by the agent's understanding of (and utility for) the consequences of that action. Utility need not be attached to courtship if the agent likes sex and knows that certain behaviors lead to sex.⁵ Second, what matters is not how much of the value of a particular action the agent fails to understand on *average*, but rather on

this paper is different. We ask *why* we might have such preferences in the first place and why they take the form they do.

³For example, it is relatively easy for random mutations and selection to adjust the intensity of our sexual urges or to create a preference for mates with symmetric features, but it is vastly more difficult to create an agent who knows that the probability of a successful birth from a random sexual encounter is about 2% ([Einon 1998](#)), that this probability varies systematically with health, age, and other observable features of the mate, and that asymmetric features are sometimes caused by genetic damage that may transmit to offspring.

⁴[Robson \(2001\)](#) provides a complementary rationale for such utilities, based on the value of being able to cope with evolutionarily novel situations. Agents in Robson's model must acquire, through experience, such information as what foods are most useful in a particular environment. Attaching utilities only to descendants gives rise to a small-sample problem that hinders effective learning, and that can be attenuated by attaching utilities to intermediate goals such as the consumption of certain nutrients.

⁵Moths seem to be programmed with a utility for "flying toward the light," presumably reflecting their inability to assess the benefits of the induced contact with other moths.

the *margin*. A small increase in the utility for an action is relevant to the action chosen only if the existing utility function and beliefs of the agent happen to make him close to indifferent about which action to choose. Evolutionarily optimal utility thus reflects not just the (misunderstood) value of the action, but also the information-processing problem the agent faces in assessing this action in his particular circumstances. This opens the door to all sorts of seemingly counterintuitive effects in our utility functions.

Once we realize that evolution uses intermediate utilities to compensate for faulty information processing, it is a small step to think there might be settings where these intermediate utilities are a better or worse proxy for what really counts. Indeed, it is hard to imagine a coherent story under which we routinely process information incorrectly, but the direction and magnitude of our errors is the same across different settings.

Section 4 shows that this naturally leads to choice-set dependence. Imagine that in evolutionary time, migratory patterns, optimal hunting techniques, and other environmental factors have fluctuated *relatively* quickly and unpredictably (compared to the pace of evolution), but in ways the agent might well be able to observe and learn about. Suppose further that the value of eating and the probabilities of famine have been much more stable over the course of evolution—even if there has not been a famine for a long period, it is still quite possible that there may be one next year.⁶ Then from “evolution’s” point of view, the agent has more trustworthy information about hunting than about eating. If the agent has come to the conclusion that a particular hunting trip is unusually dangerous or unlikely to succeed, there is a relatively good chance that the true state of the world is indeed one where hunting is ill-advised. If the agent has come to the conclusion that food is plentiful, and so the risks involved in eating a particular readily available steak make it not worthwhile, it is still quite likely that the right decision from an evolutionary point of view is to go ahead.

Now consider the utility that has evolved for a steak. A classical utility function involves a difficult trade-off. A high utility of steak gives the agent incentives to go hunting even when he thinks it quite dangerous (and so it is relatively likely to be), while a low utility of steak makes the agent too quick to forgo a readily available steak. Consider an agent who gets a moderate utility from eating steak, but is made worse off by an uneaten steak. The utility gain from eating a readily-available steak reflects both forces, and so only strong beliefs about the inadvisability of eating the steak will dissuade the agent. A decision to go hunting is also partly driven by the anticipated pleasure of eating the steak. But, if he chooses not to hunt, his utility is not lowered by the presence of an uneaten steak. As a result, the hunting decision responds more readily to the agent’s beliefs. Choice-set dependence thus leads to an agent who better matches his decisions to his environment. The choice set provides useful clues as to how reliable the agent’s information is likely to be, and hence the relevant balance between what the agent currently thinks is a good idea and what has worked well in the evolutionary past.

⁶In contrast, a change in migratory patterns for even a couple years may be good evidence that next year may be different as well. This emphasizes the point that the fundamental difficulty is the inability to give the agent the right prior knowledge that would tell him to update slowly about famine, but quickly about migration.

Section 4.2 presents examples exploring the intuition that the effects of choice-set dependence will be especially strong when the agent's information is relatively poor. Because the optimal degree of choice-set dependence is determined by conditional expectations on the margin, we find a rich dependence of optimal utilities on details of the settings for which the utilities are designed.

Choice-set dependence, of course, leads to internal conflict. An agent in a restaurant may choose a steak, but feel he would have been happier had he never entered the restaurant. He might avoid particular settings even knowing that he will not fall prey to temptation, simply to avoid the utility effects of a tempting but unchosen alternative. The point of this paper is that this internal conflict is *not* some evolutionary error or driven simply by a mismatch between the environment in which we were formed and the one we now inhabit. Rather, choice-set dependence is an evolutionarily useful device for tailoring the responsiveness of the agent's decisions to the value of his information in different settings.

Section 5 sketches extensions of this information-processing view of utility functions. For example (and returning to ideas related to those of **Robson 2001**), we suggest that evolution is likely to make us especially fearful of dangers whose nature we may be unable to safely evaluate through experience.

Psychologists have suggested that our behavior is driven partly by a collection of utility-altering visceral urges (**Loewenstein 1996**). It is straightforward to appreciate why we have urges reflecting direct evolutionary consequences such as hunger, thirst, or fatigue (**Plutchik 1984**). Evolutionary psychologists have explored in some detail how our behavior is shaped by our evolutionary past, and how modern society is rife with situations in which current circumstance (abundant food), current wisdom (information on diet and health), and anachronistic utility functions create internal conflict. This explains why many of us struggle not to eat too much. But it does not answer the questions that motivate this paper. Our results help explain why things of no evident evolutionary value affect our utility, and why the resulting utilities have the patterns they do.⁷

2. A MODEL

Situations An agent enters the environment and is presented with a *situation* in which he can either accept or reject an option (see **Figure 1**). Accepting the option leads to a lottery whose outcome is a success with probability p and a failure with probability $1 - p$. Rejecting the option leads to a success with probability q and a failure with probability $1 - q$. This is the only decision the agent makes.

We will often think of the option as involving an opportunity to eat and success as surviving long enough to reproduce. The parameters p and q are random variables, reflecting the benefits of eating and the risks required to do so in any given setting.⁸ The

⁷For example, why is it successively harder to exert willpower at the stages "Should I buy this?", "Should I serve it?", and "Should I eat the food in front of me?"

⁸If p and q are unchanging, then the problem from an evolutionary approach is trivial. Reward the right action and/or penalize the wrong one.




FIGURE 1. A situation. The success probabilities p and q are randomly generated by the environment, with the agent receiving signals s_p and s_q about these probabilities. A success yields utility x , a failure utility 0, and accepting the option yields utility y .

probability of success may be either increased ($p > q$) or decreased ($p < q$) by accepting the option.

Information The agent observes noisy signals of each of the probabilities p and q . These might reflect information the agent has about how dangerous is the pursuit of food or how dangerous is the decision to not consume. The agent may have information as to whether game is plentiful, whether the food is nearby but guarded by a jealous rival, or whether a drought makes it particularly dangerous to pass up this opportunity. However, the agent is unlikely to know these probabilities precisely.

We model this by assuming that each situation is accompanied by a pair of scalar signals s_p about p and s_q about q . The probabilities p and q are independent, p is a sufficient statistic for s_p , and q is a sufficient statistic for s_q . The joint distribution of p , q , s_p , and s_q is given by F . We assume that s_p and s_q satisfy the monotone likelihood ratio property with respect to p and q respectively, so that (for example) $E\{p \mid s_p\}$ is increasing in s_p .

Information processing As a result of an evolutionary process, the agent is equipped with a rule ϕ for transforming signals into estimates of the probability of success. It will be convenient to assume that ϕ is continuous and strictly increasing. It seems likely that evolution would induce this basic consistency. The crucial restriction in our model is that the agent must use the *same* rule ϕ for evaluating all signals. If, for example, p and q come from different processes and with information of varying reliability, proper updating requires that different updating rules be applied to s_p and s_q . Our assumption is that evolution cannot build this information about the prior or signal-generation process into the agent's beliefs. We return to this assumption at the end of this section.

Utility Utility may be derived both from the outcome of the agent's action and from the action itself. A success leads to an *outcome* (e.g., successful reproduction) that yields a utility of x . A failure gives the agent a utility that we can normalize to zero. The *act* of accepting the option (e.g., eating the food) yields a utility of y .⁹ The agent maximizes

⁹Attaching another utility to the act of rejecting the option opens no new degrees of freedom at this stage.

expected utility given his beliefs about the probabilities of success that pertain to his particular situation.

Evolution of utility We focus on how evolution shapes the utilities x and y . Random mutations inject agents with various values of x and y into the population. We make no assumptions about the process generating these mutations, though we expect the most likely mutations to be small variations on existing utilities. Agents whose values of x and y make them more likely to succeed will come to dominate the population at the expense of those whose values yield lower probabilities of success. Our evolutionary environment is sufficiently simple (most notably featuring frequency independence) that we can capture the *result* of this evolutionary process by identifying the values of x and y that maximize the probability of success.¹⁰

Why not evolve better beliefs? We study the evolutionarily optimal utilities x and y for a given information-processing rule ϕ . Is it possible that evolution could render our work moot by choosing a better ϕ ? The answer is no—we show that if a single ϕ must be applied to processing all signals, then even the evolutionarily *best* ϕ will still produce the trade-offs discussed informally above and worked out more carefully below.

The constraint that belief formation cannot be tailored to every situation (captured simply in this setting by assuming a single ϕ) is critical.¹¹ Why does evolution face constraints in designing our information processing? It is fairly easy to have preferences indirectly reflect complicated causal chains. If some agents like carrots a little more, and some a little less, then evolution will select those whose choices work out best, regardless of how complicated is the process by which carrots affect evolutionary success. However, making the information behind this process available to the agent's (explicit or implicit) cognition is a vastly more complicated problem. Designing a brain to have a built-in knowledge of vitamin A and its connections to various aspects of the well-being of the organism, and knowledge of how these aspects further translate into the outcomes that really count, is itself a trial-and-error process, and one in a massively higher dimensional space.¹² Nor, for example, can we be trusted to infer this information from our environment. An agent cannot learn the relationship between specific nutrients and healthy births by trial and error quickly enough to be useful, and we certainly cannot learn quickly enough that even many generations of ample food might still be followed by famine in the next year.¹³

¹⁰Condition (3) below ensures that the evolutionary objective is single-peaked, so finding the optimal x and y subject to the constraints imposed by (for example) the agent's information processing requires only successive local mutations.

¹¹Without this restriction, the solution to the problem is trivial. Evolution need only attach a larger utility to a success than to a failure, while designing the agent to use Bayes' rule when transforming the signals he faces into posterior probabilities, to ensure that the agent's choices maximize the probability of success.

¹²We seem to be equipped with a basic but limited knowledge of grammar (see Pinker 1994). A fairly complete description of basic grammar can be conveyed in a couple dozen pages. A summary of even our current rather rudimentary knowledge of nutrition would run into thousands of pages, and would probably still be less useful than the guidance we implicitly receive through our tastes and their variations as circumstances change.

¹³This constraint is well-accepted in other areas of study. Focusing on reactions to danger, LeDoux (1996,

One evolutionary response to this problem is to provide more cognitive power and more prior information. Unfortunately, passing on additional prior information taxes the limits of trail-and-error mutations, and additional cognitive power is expensive.¹⁴ While it is obviously unrealistic to assume that we are constrained to use a single, one-size-fits-all updating rule for all situations, it seems equally obvious that information processing must be coarse compared to the set of situations we face. In our simple setting, this restriction is conveniently captured by assuming that the same rule ϕ must be used in each setting.

3. WHAT DETERMINES UTILITY?

3.1 *Utilities for actions*

The ticket to evolutionary advantage in this environment is to maximize the probability of a success. Why, then, doesn't evolution simply attach utilities to the *outcomes* of success and failure? This section shows that attaching utilities to outcomes alone will be optimal if and generically only if the agent processes his signals flawlessly. Hence, as soon as the agent is an imperfect information processor, there is evolutionary value to attaching utilities to the *actions* of accepting or rejecting the option. The fact that we have utility for so many intermediate outcomes suggests that we indeed are subject to many such imperfections.

To identify evolutionarily advantageous utility functions, we view evolution as a mechanism designer, choosing values x and y that maximize an agent's probability of success. In solving this mechanism design problem, no generality is lost by taking $x = 1$.¹⁵ The question is the choice of y . If $y = 0$, then utilities are attached only to outcomes and not to actions. In this case, we would be motivated to eat not because we enjoy food, but because we understand that eating is helpful in surviving and reproducing. If y is nonzero, then actions as well as outcomes induce utility.

Let Δ be the true difference in the success probabilities of the two alternatives:

$$\Delta = p - q.$$

The optimal decision rule from an evolutionary perspective is then

$$\text{accept iff } \Delta > 0. \quad (1)$$

Let δ be the agent's estimate of the difference in success probabilities between the two alternatives:

$$\delta = \phi(s_p) - \phi(s_q).$$

pp. 174–178) notes that evolution deliberately removes some responses from our cognitive control precisely because evolution's prior "belief" is strong. "Automatic responses like freezing have the advantage of having been test-piloted through the ages; reasoned responses do not come with this kind of fine-tuning."

¹⁴Humans are outliers in terms of the amount of energy required to maintain our brains (Milton 1988), the risk of maternal death in childbirth posed by infants' large heads (Leutenegger 1982), and the lengthy period of postnatal development (Harvey et al. 1987).

¹⁵If $x > 0$, then multiplying x and y by $1/x$ will leave the agent's decisions unchanged. It is strictly maladaptive to set $x \leq 0$.

The agent's utility-maximizing decision is to accept if and only if $y + \phi(s_p)x > \phi(s_q)x$. Given the definition of δ and that $x = 1$, this comes to

$$\text{accept iff } y + \delta > 0. \quad (2)$$

Consider

$$E\{\Delta \mid \delta = t\}.$$

This is the expected true success-probability difference $p - q$ conditional on the agent having received signals that lead him to assess this difference at t .¹⁶ To make our results easier to interpret, we assume throughout that

$$\frac{dE\{\Delta \mid \delta = t\}}{dt} \geq 0, \quad (3)$$

so the expected difference in success probabilities $p - q$ is weakly increasing in the agent's assessment of this difference.¹⁷

We then have the following characterization of the optimal utility function (as always, conditional on the information-processing rule ϕ).

PROPOSITION 1. *The fitness-maximizing y satisfies*

$$E\{\Delta \mid \delta = -y\} = 0. \quad (4)$$

In particular, the agent's fitness is maximized by setting $y = 0$ if and only if

$$E\{\Delta \mid \delta = 0\} = 0. \quad (5)$$

PROOF When Conditions (3) and (5) hold, setting $y = 0$ ensures that the agent's choice rule (2) coincides with the (constrained) optimal choice rule (1), with Δ replaced by its conditional expected value in the latter. There is then no way to improve on the agent's choices. More generally, let $E\{\Delta \mid \delta = -y\} > (<) 0$. Then the expected probability of success can be increased by increasing (decreasing) y . $\square$

¹⁶Formally, let F be the joint measure on p, q, s_p , and s_q . F thus reflects both the prior distribution on states of the world and the technology by which signals are generated. Then

$$E(\Delta \mid \delta = t) = \lim_{r \rightarrow 0^+} \frac{\int_{\{p, q, s_q, s_p \mid |\phi(s_p) - \phi(s_q) - t| < r\}} (p - q) dF}{\int_{\{p, q, s_q, s_p \mid |\phi(s_p) - \phi(s_q) - t| < r\}} dF},$$

where it is assumed that F and ϕ have enough regularity to make this limit well defined.

¹⁷Notice that (3) need not follow from the fact that s_p and s_q satisfy the monotone likelihood ratio property. The complications in assessing the value of information on the margin arise, for example, from the fact that the sum of two affiliated random variables need not be affiliated. Imagine the agent thinks that *accept* is a very bad idea in a particular setting, so that $t = \phi(s_q) - \phi(s_p)$ is large. It would be analytically convenient to know that such a large t typically meant that s_q was big and s_p small. But, while this is certainly intuitive and is true in many simple examples, it is *not* a general implication of the monotone likelihood ratio property.




FIGURE 2. Three illustrations of **Proposition 1**. In Case A, the agent assesses signals correctly ($E\{\Delta \mid \delta = t\} = t$), and setting $y = 0$ induces optimal choices. In Case B, the agent in general does not assess signals correctly, but does so on the margin ($E\{\Delta \mid \delta = 0\} = 0$), and setting $y = 0$ again induces optimal choices. In Case C, the agent is pessimistic on the margin—when the agent receives signals making him think that p and q are equal, expected fitness is maximized by accepting (i.e., $E\{\Delta \mid \delta = 0\} > 0$). Optimal behavior is then achieved by subsidizing the option ($y > 0$, so that $E\{\Delta \mid \delta = -y\} = 0$) to correct the marginal inference.

From (5), if the agent interprets his signals correctly, then there is no evolutionary value in attaching utilities to actions. The agent will make appropriate choices motivated by the utility of the consequences of his actions. The agent will still sometimes make mistakes, but without better information there is no way to eliminate these mistakes or improve on the expected outcome.

From (4), if the agent does not interpret his signals correctly, then evolution will attach utilities to his actions in order to correct his inferences at the *marginal* signal, i.e., at the signal at which the expected success probabilities are equal. The agent must be indifferent ($y + \delta = 0$) when his signal would lead a perfect Bayesian to be indifferent ($E\{\Delta \mid \delta = -y\} = 0$). **Figure 2** illustrates this result.¹⁸

3.2 Evolution on the margin

An implication of **Proposition 1** is that we should not expect utilities to reflect the average value of various actions to which they are attached. First, utilities are usefully attached to actions only to the extent that agents sometimes misunderstand the likelihoods of the attendant outcomes. If the outcomes are correctly assessed, then actions, no matter how valuable, need receive no utility. Optimal utilities thus reflect not the evolutionary value of an action, but the error the agent makes in assessing that evo-

¹⁸Evolution need not always find perfect solutions, just ones that are not easily improved upon. It is then noteworthy that our characterization of utility depends upon adjustments that we think should be evolutionarily “easy” to implement. If $E\{\Delta \mid \delta = -y\} > 0$, then the expected incremental value of accepting is positive when the agent deems it to be zero, and there are gains to be had by pushing the agent in the direction of accepting (increasing y), and conversely. Hence, evolution can increase fitness via incremental changes in utility, based on adjustments of the form “make the agent like steak a little more,” rather than relying on what evolutionary theorists refer to as hopeful monsters.

lutionary value. Second, one might think that fitness would be maximized by a utility function that corrected this error *on average*. As (4) makes clear, what counts is the error the agent makes in the marginal cases where he is indifferent between two actions. Once one realizes that it is the marginal case that counts, richer implications become clear. The following examples illustrate.

EXAMPLE 1. In our first example, the agent on average overestimates the value of accepting the option, but evolutionary fitness is nonetheless improved by setting $y > 0$, pushing him to accept the option more than he otherwise would. Let

$$E\{\Delta \mid \delta = t\} = a + bt,$$

with $a > 0$ and $b > 0$. Solving (4), the optimal utility is¹⁹

$$y = \frac{a}{b}. \quad (6)$$

Assume that δ is large on average and that $b < 1$. Because δ is on average large and $b < 1$, the agent on average overestimates the value of the option. However, since $y = a/b > 0$, the agent's fitness is maximized by pushing the agent even more toward acceptance. We see here the importance of the agent's marginal beliefs: When $\delta = -a/b$ (so that $E\{\Delta \mid \delta\} = 0$), the agent *underestimates* the relative value of the option (thinking it to be negative), even though he overestimates it on average.

It follows from (6) that, as one might expect, a choice with a large expected value (large a) will tend to have a large utility. It is thus no surprise that we have a powerful urge to flee dangerous animals or eat certain foods. However, there is also a second effect. The smaller is b , the larger is y . The point is that the less informative is the agent's information, holding fixed his average assessment, the more negative is the relevant marginal signal. When b is near zero, evolution effectively insists on the preferred action. While blinking is partly under conscious control, our utility functions do not allow us to go without blinking for more than a few seconds. It would seem that we are unlikely to have reliable information suggesting that this is a good idea. The experience of trying not to cough after mis-swallowing a sip of water at a dinner party has much the same aspect. $\diamond$

EXAMPLE 2. This example shows that "better beliefs" generically cannot eliminate the evolutionary value of rewarding intermediate actions. Suppose that p is distributed uniformly on $[0, 1]$, while (for transparency) q is degenerate at $q^* = \frac{1}{3}$. The signal s_p simply equals p , so that p is perfectly observed by the agent, but the signal s_q is also uniformly distributed on $[0, 1]$. The agent thus has good information about the implications of accepting the option, but useless information about the implications of rejecting.

Consider first the simple information processing rule $\phi(s) = s$, so that $\delta = \phi(s_p) - \phi(s_q) = s_p - s_q$. Then we can calculate that

$$E\{\Delta \mid \delta = t\} = \frac{1}{6} + \frac{1}{2}t.$$

¹⁹When $\delta = -(a/b)$, $E\{\Delta \mid \delta\} = a - b(a/b) = 0$.




FIGURE 3. The optimal solution (left panel) and the agent's behavior when $y = 0$, for **Example 2**. This holds no matter what rule the agent uses to transform signals into expected probabilities. Lined regions show mistaken decisions. The agent's fitness is increased by setting $y > 0$ and hence expanding the agent's acceptance region.

In this case, when the agent is indifferent ($\delta = 0$), the expected value of $p - q$ is positive ($\frac{1}{6}$). The optimal value of y is $\frac{1}{3}$. **Figure 3** shows, as a function of the signals s_p and s_q , the optimal outcome (left panel) as well as the outcome implemented (for any increasing ϕ) when $y = 0$ (right panel). The lined regions identify mistaken decisions. Because increasing y reduces the incidence of mistaken rejections more quickly than it induces mistaken acceptances, $y > 0$ is optimal.

There are other specifications of the signal-processing rule ϕ , including some (embodying a more rapid response of ϕ to s at some points along the diagonal than others) for which y is optimally zero. However, suppose the agent faces a multitude of situations of this type that differ in the value $q^* \in [0, 1]$, and where evolution can tailor utilities to each situation. For any ϕ , for only one value of q^* will the integral in **footnote 16** equal zero at $\delta = 0$, and hence $y = 0$ be optimal. Hence, in rich environments, evolution will routinely attach utilities to the actions of imperfect information processors. $\diamond$

The implication of **Example 2** is general. For convenience, assume that signals are drawn from the interval $[0, 1]$. If utilities are attached only to outcomes, then the agent's decision rule will be determined by the diagonal line in the right panel of **Figure 3** regardless of how success probabilities are distributed, signals are generated, and information is processed.

COROLLARY 1. *For any signal-processing rule ϕ , the agent's fitness is maximized by setting $y = 0$ if*

$$E\{p \mid s_p = t\} = E\{q \mid s_q = t\} \quad (7)$$

for all t .

PROOF. Let $s_p > s_q$. Then, by the monotone likelihood ratio property, $E\{p \mid s_p\} - E\{q \mid s_q\} > E\{p \mid s_q\} - E\{q \mid s_q\} = 0$. Similarly, $s_p < s_q$ implies $E\{p \mid s_p\} - E\{q \mid s_q\} < 0$. Thus,

the evolutionarily optimal policy is given by the diagonal in the right panel of **Figure 3**. For any given ϕ , this behavior is induced by setting $y = 0$. $\square$

REMARK. If $E\{p \mid s_p = t\} \neq E\{q \mid s_q = t\}$ for a nonzero-measure set of t , then the agent's behavior when $y = 0$ will generically not be evolutionarily optimal. Essentially, setting $y = 0$ will then be optimal only if, when adjusting y away from zero, the expected gains provided by making better decisions for some values of s_p and s_q are precisely balanced by payoff losses from poorer decisions at other values. In an intuitive sense, this balance is nongeneric.²⁰

In the previous example, the agent's information about rejecting the option is biased, yielding an average success probability of $\frac{1}{2}$ that exceeds the actual value of $\frac{1}{3}$. Does this play an important role in the result? The following example shows that this bias is unnecessary.

EXAMPLE 3. While some information imperfection is necessary if utilities are to be optimally attached to actions, this example shows that the agent may still be correct about both signals on average. Again let $\phi(s_p) = s_p$ and $\phi(s_q) = s_q$ (though once again the result does not depend upon the form of ϕ). Suppose that q is degenerate at $q = \frac{1}{4}$, while the agent's signal s_q is distributed uniformly on $[0, \frac{1}{2}]$. Now let p be perfectly observed by the agent, and be distributed on $[0, 1]$ with density $-\ln p$. Then both options have an expected value of $\frac{1}{4}$, and the agent's expected value is correct in both cases. However, conditional on being indifferent, the agent must have drawn a signal $p \in [0, \frac{1}{2}]$, and the expected value of $p - q$ conditional on indifference is negative.²¹ Evolution thus optimally sets $y < 0$, discouraging the option. Again, we see the importance of the agent's valuation on the margin rather than on average. $\diamond$

4. CHOICE-SET DEPENDENCE

We have shown that a utility for intermediate actions like eating a steak only makes sense because the agent makes systematic mistakes in estimating their value. The key observation is that there is no reason to believe that an agent who makes such mistakes makes the same mistakes in all contexts. In this section, we show how a setting where the agent makes different mistakes in different contexts creates evolutionary value for a utility function that depends on things that have *no* direct impact on evolutionary success. Rather, their role is to tailor utility more closely to the specific informational context at hand. How any given feature optimally affects utility thus depends both on its direct evolutionary impact *and* how it correlates with errors in information processing.

²⁰To be precise, it can be shown that for any given ϕ , the set of F for which $E\{\Delta \mid \delta = 0\}$ fails to equal zero (cf. **footnote 16**) is open and dense (in the weak topology) in the space of all distributions.

²¹This expected value is $\int_0^{\frac{1}{2}} p(-\ln p)dp - \frac{1}{4} < 0$.

4.1 Information and choice-set dependence

Suppose that the environment may place the agent in one of two situations, 1 and 2, as illustrated in [Figure 4](#). The success probability when rejecting the option is q in either case, with success probabilities p_1 and p_2 when accepting the option in Situations 1 and 2. The corresponding signals are s_q , s_{p_1} , and s_{p_2} . As before, the agent derives a utility of 1 from a success, 0 from a failure, and utility y , *the same value in both situations*, from the act of accepting the option.

For example, suppose that in Situation 2, accepting the option entails an opportunity to eat a steak. As we have shown, evolution optimally attaches a utility y to steak satisfying

$$E(\Delta \mid \delta = -y) = 0.$$

Now suppose that in Situation 1, accepting the option entails eating a steak at the end of a hunting trip. The agent is likely to have quite different sources of information about these two situations and thus to make quite different errors in processing this information. In particular, the hunter may have an idea of what hazards he will face on the hunting trip before achieving consumption and how these will affect the probability p_1 . Only coincidentally will it then be the case that

$$E(\Delta \mid \delta = -y, \text{steak on hand}) = E(\Delta \mid \delta = -y, \text{steak to be hunted}). \quad (8)$$

If (8) does not hold, the agent's expected fitness can be increased by attaching different utilities to accepting the option in the two situations.

How can evolution accomplish this? One possibility is to attach utilities to more actions. The agent can be given a taste for meat, a disutility for the physical exertion of hunting, and a fear of the predators he might encounter. However, there are limits to evolution's ability to differentiate actions and attach different utilities to them—what it means to procure food may change too quickly for evolution to keep pace—and the set of things from which we derive utility is small compared to the richness of the settings we face. As a result, evolution inevitably faces cases in which the same utility is relevant to effectively different actions. This is captured in our simple model with the extreme assumption that y must be the same in the two situations.²² The critical insight is then that the agent's overall probability of success can be boosted if utility can be conditioned on some other reliable information that is correlated with differences in the actions.

Assume that in Situation 2, a utility of z can be attached to the act of *foregoing* the option. We say that an option with this property is *salient*. In practice, an option is salient if its presence stimulates our senses sufficiently reliably that evolution can tie a utility to this stimulus, independently of our signal-processing.²³ In our example, the presence of the steak makes it salient in Situation 2. The question now concerns the

²²Once again, what we need is that evolution cannot tailor y perfectly to each of the many situations we face.

²³The importance of salient alternatives is well studied by psychologists ([Gardner and Lowinson 1993](#), [Mischel et al. 1992](#), [Siegel 1979](#)) and is familiar more generally—why else does the cookie store take pains to waft the aroma of freshly-baked cookies throughout the mall?




FIGURE 4. Two situations. Accepting the option is salient in the second situation but not the first.

value of z . If fitness is maximized by setting $z \neq 0$, then there is evolutionary advantage to tailoring the utility gradient between accepting and rejecting the option to the two situations, and we have choice-set dependence. Only if $z = 0$ do we have a classical utility function.

Let

$$\delta_1 = \phi(s_{p_1}) - \phi(s_q)$$

$$\delta_2 = \phi(s_{p_2}) - \phi(s_q)$$

$$\Delta_1 = p_1 - q$$

$$\Delta_2 = p_2 - q.$$

It is immediate from **Proposition 1** that choice-set dependence is unnecessary if and only if the agent has equivalent marginal information in Situations 1 and 2:

PROPOSITION 2. *The optimal utility function (x, y, z) does **not** exhibit choice-set dependence (sets $z = 0$) if and only if there exists t^* such that*

$$E\{\Delta_1 \mid \delta_1 = t^*\} = E\{\Delta_2 \mid \delta_2 = t^*\} = 0. \quad (9)$$

PROOF. Given (9), the agent's estimates of the success probabilities in Situations 1 and 2 are equally informative at the relevant margin. Setting $z = 0$ and $y = -t^*$ ensures that (4) holds in both situations, and there is thus no gain from choice-set dependence. Conversely, suppose that the agent's beliefs are differentially informative in the two situations (i.e., (9) fails). Then fitness can be enhanced by attaching different utility subsidies in the two situations. This can be accomplished by choosing y to induce optimal decisions in Situation 1 and $y - z$ (and hence $z \neq 0$) to induce optimal decisions in Situation 2. $\square$

For example, using choice-set dependence to boost the relative attractiveness of steak when it is available ($z < 0$), in contrast to simply increasing the utility of steak across the board (increasing y), might reflect a situation in which evolution finds it beneficial to grant substantial influence to the agent's beliefs about the consequences of production, while allowing less influence to his beliefs about consumption.²⁴

An analogue to **Corollary 1** is:

²⁴Similarly, the enjoyment of sex is relevant for both heat of the moment decisions and for decisions

COROLLARY 2. *For any signal-processing rule ϕ , the agent's fitness is maximized without recourse to choice-set dependence ($z = 0$) if s_{p_1} and s_{p_2} are identically distributed²⁵ and*

$$E\{p_1 \mid s_{p_1} = t\} = E\{p_2 \mid s_{p_2} = t\} \quad (10)$$

for all signals t .

As with **Corollary 1**, when (10) fails, choice-set dependence is generically nonoptimal. We can thus expect choice-set dependence to be the norm rather than the exception for imperfect information processors.

As is the case with attaching utilities to actions, this result tells us that evolution will generally find it useful to exploit choice set dependence. Anyone who has ever said, "Let's put these cashews away before we spoil our dinner," has practical experience with choice-set dependence. Best of all is to be without the temptation of their presence. Once they are there, eating is the preferred choice. Worst of all is looking at the cashews without indulging.^{26,27}

Which alternatives are salient in any given context is again the result of evolution. As it turns out, a sizzling steak is salient while a steak in the grocer's freezer is not. Cashews on the table are salient; those in the pantry are less so. What is salient reflects both the technological constraints faced by evolution and the incremental value of tailoring utility to specific contexts.²⁸

It is again important to stress that there is nothing surprising about making *choices* that depend upon the set of alternatives. One may well eat vegetables when there is no alternative while eating fatty foods when they are also available. But to achieve such behavior, evolution could simply endow us with utilities for the tastes of vegetables and fat, and let us choose vegetables as the situation dictates. With choice-set dependence, we have one utility for vegetables when only they are available, and another when fatty foods are also available.

about whether to court a particular potential mate or invest in things much more indirectly linked to sexual opportunities, such as social status. The tendency to get carried away by the heat of the moment is perhaps again an instance of choice-set dependence, engineered to give relatively free reign to our beliefs about social interactions while paying less attention to our less informative beliefs about the value of a given sexual opportunity.

²⁵It is implicit in our structure that q is distributed identically in Situations 1 and 2.

²⁶Thaler (1994, p. xv) reports a similar incident, explaining it with much the same preferences.

²⁷A parent on a family outing might try to avoid the sight of ice cream parlors, but agree to get ice cream if one is inadvertently stumbled across. The preferred outcome is "no ice cream, happy kids," next is "ice cream," and worst of all is "no ice-cream, unhappy kids." Avoiding ice cream parlors can make sense even if the parent plans to refuse ice cream when sighted, as it avoids the hassle of resisting the lobbying effort from the three-foot-tall contingent. Substitute "medulla oblongata" for "kids," and one has a decent model of choice-set dependence.

²⁸For example, it may well be of evolutionary value to tailor the utility of a steak to whether the agent is thinking about the effect of the hunting trip on his social standing, since the agent may have a better understanding of social structure than of the risks involved in hunting. It is also likely that evolution has found no convenient way to build this into utility. Conversely, it seems plausible that evolution could build an agent who favors vegetables on warm days and meat on cold days. Whether this dependence is adaptive, and hence is likely to enter our utility, depends on whether the agent is more likely to misjudge the value of vegetables when it is cold and meat when it is warm.

4.2 Patterns of choice-set dependence

Choice-set dependence is a device for coping with imperfect information processing. We would then expect choice-set dependence to be stronger when the agent's information is a poorer guide to the true state of the world. There are three reasons why this intuition is incomplete. First, the relevant consideration in judging the quality of the agent's information is not just how much the agent knows about probability p or q , but rather how much he knows relative to the underlying stochastic processes from which these probabilities are drawn.²⁹ Second, we must again remember that our intuition is typically about average values, while optimal utilities are determined by marginal considerations. **Example 1** has provided an indication of how average and marginal evaluations can push in different directions. Third, utilities reflect not simply the agent's information, but expected values conditional on the agent's information. Just as (3) is not an implication of the monotone likelihood ratio property (cf. **footnote 17**), this relationship can depend upon fine details of the information structure.

We begin with an example illustrating our basic intuition. If evolution has found it optimal to subsidize an option, then an increase in the quality of the agent's information about the option will generally lead to a smaller subsidy. Similarly, if evolution has found it optimal to discourage an option, then an increase in the quality of the agent's information will lead to less vigorous discouragement.

EXAMPLE 4. Consider two situations, as in **Figure 4**, with the option being salient in Situation 2 but not Situation 1. Let q be degenerate at $q^* = \frac{1}{2}$, and let s_q be uniformly distributed on $[0, 1]$. The agent's information about q is thus noisy, but unbiased. Let the agent's information-processing rule be $\phi(s) = s$ for all signals s .

Let the signal in Situation 2 satisfy $s_{p_2} = p_2$, and let p_2 be uniformly distributed on $[\frac{3}{8}, \frac{3}{4}]$. So, in Situation 2, the average value of the option exceeds that of rejection. If $y - z = 0$, then the expected value of acceptance conditional on indifference exceeds that of rejection, and the agent accepts the option too seldom. The optimal value of $y - z$ equals $\frac{3}{8}$.

In Situation 1, let $s_{p_1} = p_1$, with p_1 uniformly distributed on $[\frac{1}{8}, 1]$. Then the options have the same expected value in Situations 1 and 2. However, the larger variance in the value of p_1 makes it more likely that, should the agent view the option as a bad idea, this reflects an informative indication that the option is indeed not valuable. The agent's information is thus more valuable in Situation 1. Evolution still finds it optimal to subsidize accepting the option in Situation 1, but not as strongly as in Situation 2. The evolutionarily optimal utility function sets $y = \frac{1}{8}$. To preserve the desired subsidy in Situation 2, we have $z = -\frac{1}{4}$, using choice-set dependence to distinguish the two situations, with a stronger push toward accepting the option in the situation where the agent's information is relatively less useful.

²⁹Hence, the agent may have quite precise information about the value of eating and relatively noisy information about the value of hunting. But if the true evolutionary value of eating is drawn from an extremely tight distribution and the value of hunting from a more diffuse one, then the agent's informational advantage will be with respect to hunting.

A reverse pattern appears when evolution finds it optimal to discourage the option. Let us retain the specification of this example, but now assume that p_2 is distributed uniformly on $[\frac{1}{4}, \frac{5}{8}]$. Evolution now finds it optimal to set $y - z = -\frac{3}{8}$. Let p_1 be distributed uniformly on $[0, \frac{7}{8}]$. Then evolution sets $y = -\frac{1}{8}$, and hence $z = \frac{1}{4}$. Choice-set dependence is again used to exert a stronger effect on utility when the agent's information is relatively less reliable, in this case a stronger disincentive to accept the option. $\diamond$

The previous example identified the agent as having better information in Situation 1. However, the agent's signal in Situation 1 was drawn from a uniform distribution with a larger variance (and equivalent mean) than that of Situation 2, suggesting that the agent has less information in the former case. To reconcile these observations, notice that the agent's signal in each situation identifies the underlying success probability precisely. The distribution from which the signal is drawn is the evolutionarily relevant prior distribution of success probabilities. When this prior is relatively diffuse, as in Situation 1, the informational advantage conferred by observing the realized probability is relatively large.

The following example provides another illustration of the importance of the relative quality of the agent's information.

EXAMPLE 5. Once again, let there be two situations, with the option being salient in Situation 2 but not Situation 1. Let q be degenerate at $q^* = \frac{1}{6}$, and let s_q be uniformly distributed on $[0, 1]$. Let p_1 and p_2 both be uniformly distributed on $[0, 1]$. Let s_{p_2} be uniformly distributed on $[0, 1]$ but independently of p_2 , so that s_{p_2} is another useless signal. Let the agent's information-processing rule be $\phi(s) = s$ for all signals s .

The agent has no information about Situation 2, while the expected value of accepting the option is positive. The optimal utility function thus sets $y - z \geq 1$, ensuring that the agent always accepts the option.

Now suppose that the agent's signal about the success probability p_1 either equals p_1 or provides no useful information. Hence, with probability $\gamma \in [0, 1]$ the signal s_{p_1} is given by p_1 and with probability $1 - \gamma$, s_{p_1} is distributed uniformly on $[0, 1]$ and is independent of p_1 . Then the optimal utility function has $y > 0$, pushing the agent toward accepting the option in Situation 1. However, as γ increases and hence the agent's information in Situation 1 gets better, the optimal value of y declines to $\frac{2}{3}$ (a value calculated in the [Appendix](#)), with z declining similarly to preserve incentives in Situation 2. As the agent's information becomes increasingly precise in Situation 1, evolution optimally gives more weight to the agent's information by reducing the utility attached to the accept option, relying upon choice-set dependence to preserve the incentives in the face of the agent's noisy information in Situation 2. $\diamond$

In Examples 4 and 5, the utilities attached to actions are smaller when the agent has better information. While we find this pattern intuitive, the result is not universal. The following example illustrates.

EXAMPLE 6. We again consider two situations, with the option salient in Situation 2. Let the actual value of q be fixed at $\frac{1}{4}$, with the agent receiving a signal s_q that is uniformly




FIGURE 5. A choice between two situations, with accepting the option being salient in the second situation, but not the first.

distributed on the interval $[0, \frac{1}{2}]$ (and hence that is unbiased). Let $s_{p_1} = p_1$ and $s_{p_2} = p_2$, with p_1 distributed on $[0, 1]$ with density $-\ln p_1$, and with p_2 uniformly distributed on $[0, \frac{1}{2}]$. It is straightforward that evolution optimally sets $y - z = 0$ and (from [Example 3](#)) $y < 0$. In this case, the agent's information is relatively better in Situation 1, while evolution optimally chooses (with the help of choice-set dependence) to attach utilities to actions that have no effect in Situation 2 while pushing the agent against accepting the option in Situation 1. $\diamond$

In this example, the agent makes good (indeed, optimal, though this is not critical to the result) decisions in Situation 2, where his information is relatively less valuable, while still making systematic errors in Situation 1 in spite of having relatively good information. Utilities attached to actions thus play no role in Situation 2, while they are useful in Situation 1.

Self control and commitment Choice-set dependence gives rise to internal conflict and self-control, with the corresponding incentives to make behavioral commitments. For example, suppose the environment is as shown in [Figure 5](#). In this setting, the agent begins by choosing between the two situations depicted in [Figure 4](#), with accepting the option being salient in Situation 2 but not Situation 1. Let us think of accepting the option as choosing an unhealthy but gratifying meal, while rejecting it corresponds to a diet meal. Situation 1 corresponds to a lonely meal at home, with a refrigerator full of diet dinners and a steak in the freezer. Situation 2 corresponds to a steakhouse with a supplementary dieter's menu. Suppose, as in [Example 3](#), that the optimal utility function entails $z < 0$, so that steak is subsidized when it is salient. Then the agent may prefer Situation 1 even if there is some cost in choosing Situation 1, in order to ensure that he rejects the steak.

Choice-set dependence has implications for self control beyond those of present bias. First, difficulties with self control can arise without intertemporal choice. One can strictly prefer junk food that is hidden to that which is exposed, knowing that one will find it painful to resist the latter, all within a span of time too short for nonstandard

discounting to lie behind the results.³⁰ More importantly, because our utility for one choice can be reduced by the salient presence of another, it may be valuable to preclude temptations that one *knows* one will resist. Someone who is certain she will stick to a diet may still go to some lengths not to be tempted by rich food.³¹

Choice-set dependence and temptation In Examples 4 and 5, choice-set dependence appears in the form of a utility penalty ($z < 0$) for not accepting an attractive option ($y > 0$). This is the force at work when we find it hard to resist a waiting (and hence salient) steak. One interpretation of such a utility penalty is as temptation.

Introspection suggests that choice-set dependence often takes the form of temptation: when a consumption opportunity is salient, forgoing the opportunity leaves us feeling worse than never having been offered the alternative in the first place. We expect choice-set dependence to often take the form of penalizing us for neglecting attractive options. This will be the case if “evolution” is especially willing to over-ride the agent when an option is salient. There are two reasons why we expect this to be the case. First, we expect the agent’s information to be relatively better about production than consumption decisions, while expecting the immediacy of the latter to render them more often salient. Second, what we find salient is also shaped by evolution. As a result, we can in general expect that our information will be systematically better about options that are not salient precisely because it is where our information is poorest that there is the greatest evolutionary value in providing direct feedback in the form of utility. The fact that we have utility for sex thus suggests that we are better at figuring out how to get sex than at calculating the benefits thereof.

5. EXTENSIONS

The information processing problem that is central to our story about why we have choice set dependent utility also provides insight into other features of our decisions making processes.

Presently-biased preferences Presently-biased preferences may also be a useful evolutionary response to imperfect information processing, allowing our model to unify seemingly contending explanations. The utilities directly associated with consumption typically arise contemporaneously with the consumption, while the outcomes implied by such consumption, with their attendant utilities, occur over time. The utilities associated with a production process such as hunting or gathering are likely to be spread

³⁰The advice to dieters that they not go to the grocery store when hungry is interesting in this light. Even though one may be buying food only to be eaten in subsequent days, one’s current hunger makes choices between various types of food more salient. Since the time structure of consumption in subsequent days is unaffected by today’s hunger, this behavior is not easily explained by presently-biased preferences.

³¹Notice that under choice-set dependence (but not necessarily present bias), foreclosing an option is beneficial only if it makes it less salient. In contrast, the current self of an agent with presently-biased preferences might wish to restrict the choice of a future self, even though it makes that future self unhappy, because it makes other future selves happier. Thus, when Ulysses had himself tied to the mast, one can conclude that either (a) he had presently-biased preferences, or (b) the very fact that he knew he was physically unable to approach the sirens made it less painful to be unable to do so.

out, rather than occurring primarily at the time of the decision. Suppose that evolution designs us to discount more between the current and next period than between successive periods. When evaluating a consumption decision, this has the effect of making the agent weight the utility attached to accepting the option more heavily than the utilities attached to outcomes. When evaluating production decisions, the extra discounting in the first period is relevant only to the (limited) extent that there is any utility attached to the first-period part of the production process. For all later utilities, the extra discounting affects the alternatives of accepting and rejecting the option equally. Hence, a present bias in preferences gives relatively greater weight to the agent's information about production decisions, where our expectation is that his information is relatively more valuable.

What to fear Most of us are “irrationally” afraid of snakes (“irrationally” in the sense that we remain afraid even on hearing that a particular snake is harmless or that a particular swamp is snake free), but few of us are afraid of mushrooms. Since both can be potentially fatal and both can be good eating, this is puzzling. The informational perspective developed here suggests an explanation.

Imagine that being bitten by a poisonous snake is always fatal, and that there is 1 chance in 1000 of being bitten when venturing into a swamp actually populated by poisonous snakes, with 999 chances in 1000 of noticing nothing (that is, neither seeing nor being bitten by a snake). Imagine that ingesting a poisonous mushroom is fatal 2 times in 1000, and results only in serious discomfort the rest of the time. The mushrooms are then inherently more dangerous than the snakes, assuming comparable priors concerning the likelihood that a swamp is snake-infested or a particular species of mushroom poisonous.

An agent who wants to routinely cross a swamp is either right, and the swamp is safe, or is wrong, and will eventually be unlucky and die. Even a fairly small probability that the swamp is dangerous thus makes it evolutionarily valuable to keep the agent out. On the other hand, an agent who routinely wants to eat the local mushrooms is either right and they are good mushrooms, or is wrong and may die on his first ingestion, but with much higher probability gets sick, realizes he was wrong, and stops eating them. This makes it more valuable to allow the agent to risk the mushroom. Evolution should thus make us fear not simply things that are bad for us, but rather things whose danger we may underestimate *without* discovering our error before they kill us. There is thus evolutionary value to an “irrational” fear, one that does not dissipate simply because we claim to have information that it should.

Following your heart When gut instincts and dispassionate deliberations disagree, the “rational” prescription is to follow the head rather than the heart. In our model, a strong utility push in favor of an action indicates either that the action has been a very good idea in our evolutionary past or that this is a setting in which our information has typically been unreliable. There is thus information in these preferences. The truly rational response is to ask how much weight to place on the advice they give.³²

³²We think it unlikely that in evolutionary time, we systematically thought about the information that could be rationally drawn from our emotional state. This is reflected in our assumption that ϕ does not

Consider the question of what to eat while pregnant. Selection has given us tastes favoring nutrients that are important for the fetus. So, even if we don't know why pickles and ice cream are a good idea, the rational thing to do is to update in the direction that they are beneficial.³³ Similarly, questions of when to trust have been critical in our evolutionary past. As a result, the mere fact that one is uncomfortable about a business deal is probably relevant information.

In some settings, it is clear that the world has changed in ways to which evolution has not had time to adjust, and so the implicit advice in our preferences should be resisted. The desirability of adding 20 pounds of fat stores before winter is in this class. While we agree that starvation is undesirable, we reasonably believe that we face different odds than those that shaped our evolution. In other settings, it is clear why our utilities push us in a particular direction, but we may disagree about the desirability of the ultimate outcome. Hence, there is no real need to conclude that since sex is fun, it must be good for you.

6. CONCLUSION

There is evolutionary value in having things in the utility function not because they have direct evolutionary consequences, but rather because conditioning our utility on these features makes our choices differentially responsive to the quality of our information. The important aspects of the model from which these conclusions have emerged are the following.

- The agent has useful information about fitness. Both cultural and environmental factors can be critical to fitness and can fluctuate more rapidly than the pace of evolutionary adaptation.
- It is impossible, for reasons of complexity and (more importantly) of the difficulty of building an accurate prior into the agent, to make the agent a *perfect* information processor. Our results show that the very fact that we have utility for intermediate actions suggests that we are not perfect information processors.
- Given that we process information imperfectly, there is every reason to think that our beliefs are sometimes more and sometimes less reliable as estimates of true expected fitness.
- Evolution has compensated for this by incorporating seemingly irrelevant factors into our utility because they are correlated with the quality of our information.
- This gives rise to seemingly irrational (from the classical point of view) utility functions with their attendant self-control problems.

While self-control problems are irrational in a classical utility framework, it should be no surprise that real utility functions incorporate them.

take y or z as arguments.

³³Only recently do we know why it makes sense to avoid some foods for which pregnant women develop a distaste: they contain toxins capable of crossing the placental barrier (see, for example, Profet 1992).

APPENDIX: CALCULATIONS, **EXAMPLE 5**

Consider Situation 2 in the limiting case of $\gamma = 1$ and hence the agent having perfect information about the success probability p_2 . The agent is indifferent between accepting and rejecting the option when $\phi(s_{p_2}) + y = \phi(s_q) + z$, or $\delta_2 = s_{p_2} - s_q = z - y$. We are then interested in the difference in the expected values of accepting and rejecting, conditional on the agent being indifferent, or $E\{\Delta_2 \mid \delta_2 = z - y\}$. This is the expected value of $p_2 - q$, conditional on having drawn signals from the intersection of the line $s_{p_2} - s_q = z - y$ with the unit square, given the uniform distributions of s_{p_2} and s_q on $[0, 1]$. Given that the expected value of q is $\frac{1}{6}$, this is (for $y - z > 0$)

$$\begin{aligned} \int_0^{1-(y-z)} p_2 \frac{1}{1-(y-z)} dp_2 - \frac{1}{6} &= \left. \frac{\frac{1}{2} p_2^2}{1-(y-z)} \right|_0^{1-(y-z)} - \frac{1}{6} \\ &= \frac{1}{2}(1-(y-z)) - \frac{1}{6}. \end{aligned} \quad (11)$$

Setting this equal to zero gives

$$y - z = \frac{2}{3}.$$

In general, note from (11) that the optimal $y - z$ solves

$$\frac{1}{2}(1-(y-z)) - q^* = 0$$

or

$$y - z = 1 - 2q^*.$$

REFERENCES

- Ainslie, George W. (1992), *Picoeconomics*. Cambridge University Press, Cambridge. [119]
- Einon, Dorothy F. (1998), “How many children can one man have?” *Evolution and Human Behavior*, 19, 413–426. [120]
- Elster, Jon (1985), “Weakness of will and the free-rider problem.” *Economics and Philosophy*, 1, 231–265. [119]
- Gardner, Eliot L. and Joyce H. Lowinson (1993), “Drug craving and positive/negative hedonic brain substrates activated by addicting drugs.” *Seminars in the Neurosciences*, 5, 359–368. [119, 131]
- Gul, Faruk and Wolfgang Pesendorfer (2001), “Temptation and self-control.” *Econometrica*, 69, 1403–1436. [119]
- Harvey, Paul H., Robert D. Martin, and Tim H. Clutton-Brock (1987), “Life histories in comparative perspective.” In *Primate Societies* (Barbara B. Smuts, Dorothy L. Cheney, Robert M. Seyfarth, Richard W. Wrangham, and Thoma T. Struhsaker, eds.), 181–196, University of Chicago Press, Chicago. [125]

- Laibson, David (2001), "A cue-theory of consumption." *Quarterly Journal of Economics*, 116, 81–120. [119]
- LeDoux, Joseph (1996), *The Emotional Brain*. Simon and Schuster, New York. [124]
- Leutenegger, Walter (1982), "Encephalization and obstetrics in primates with particular reference to human evolution." In *Primate Brain Evolution: Methods and Concepts* (Este Armstrong and Dean Falk, eds.), 85–95, Plenum Press, New York. [125]
- Loewenstein, George (1996), "Out of control: visceral influences on behavior." *Organizational Behavior and Human Decision Processes*, 65, 272–292. [119, 122]
- Loewenstein, George and Drazen Prelec (1992), "Anomalies in intertemporal choice: evidence and an interpretation." *Quarterly Journal of Economics*, 107, 573–598. [119]
- Loewenstein, George and Richard Thaler (1989), "Anomalies: intertemporal choice." *Journal of Economic Perspectives*, 3, 181–193. [119]
- Milton, Katherine (1988), "Foraging behaviour and the evolution of primate cognition." In *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans* (Richard W. Byrne and Andrew Whiten, eds.), 285–305, Clarendon Press, Oxford. [125]
- Mischel, Walter, Yuichi Shoda, and Monica L. Rodriguez (1992), "Delay of gratification in children." In *Choice over Time* (George Loewenstein and Jon Elster, eds.), 147–166, Russell Sage, New York. [119, 131]
- O'Donoghue, Ted and Matthew Rabin (1999a), "Doing it now or later." *American Economic Review*, 89, 103–124. [119]
- O'Donoghue, Ted and Matthew Rabin (1999b), "Incentives for procrastinators." *Quarterly Journal of Economics*, 114, 769–816. [119]
- Pinker, Steven (1994), *The Language Instinct*. Harper Collins, New York. [124]
- Plutchik, Robert (1984), "Emotions: A general psychoevolutionary theory." In *Approaches to Emotion* (Klaus R. Scherer and Paul Ekman, eds.), 197–219, Erlbaum, Hillsdale, NJ. [122]
- Pollak, Robert A. (1968), "Consistent planning." *Review of Economic Studies*, 35, 201–208. [119]
- Profet, Margie (1992), "Pregnancy sickness as adaptation: a deterrent to maternal ingestion of teratogens." In *The Adapted Mind* (Jerome H. Barkow, Leda Cosmides, and John Tooby, eds.), 327–366, Oxford University Press, Oxford. [139]
- Robson, Arthur J. (2001), "Why would nature give individuals utility functions?" *Journal of Political Economy*, 109, 900–914. [120, 122]

- Rubinstein, Ariel (2003), ““Economics and Psychology”? The case of hyperbolic discounting.” *International Economic Review*, 44, 1207–1216. [119]
- Schelling, Thomas (1984), “Self-command in practice, in policy, and in a theory of rational choice.” *American Economic Review*, 74, 1–11. [119]
- Shafir, Eldar, Itamar Simonson, and Amos Tversky (1993), “Reason-based choice.” *Cognition*, 49, 11–36. [119]
- Siegel, Shepard (1979), “The role of conditioning in drug tolerance and addiction.” In *Psychopathology in Animals: Research and Clinical Implications* (J. D. Keehn, ed.), 143–168, Academic Press, New York. [119, 131]
- Strotz, Robert H. (1956), “Myopia and inconsistency in dynamic utility maximization.” *Review of Economic Studies*, 23, 165–180. [119]
- Thaler, Richard H. (1994), *Quasi-Rational Economics*. Russell Sage Foundation, New York. [133]
- Thaler, Richard H. and H. M. Shefrin (1981), “An economic theory of self-control.” *Journal of Political Economy*, 89, 392–406. [119]
- Tversky, Amos and Itamar Simonson (1993), “Context-dependent preferences.” *Management Science*, 39, 1179–1189. [119]

Submitted 2005-4-20. Final version accepted 2005-9-22. Available online 2005-12-2.