THE EVOLUTIONARY OPTIMALITY OF DECISION AND EXPERIENCED UTILITIES*

Arthur Robson
Department of Economics
Simon Fraser University
8888 University Drive
Burnaby, B. C. V5A 1S6 Canada
robson@sfu.ca

Larry Samuelson
Department of Economics
Yale University
New Haven, CT
USA 06520-8281
Larry.Samuelson@yale.edu

March 7, 2009

Been Down So Long It Looks Like Up to Me—Richard Fariña.

Abstract. We show that evolution may optimally design agents to have different decision and experienced utilities.

Contents

* We thank the Canada Research Chair Program and the National Science Foundation (SES-0549946) for financial support, and thank Luis Rayo for useful discussions.

The Evolutionary Optimality of Decision and Experienced Utilities

1 Introduction

People who contemplate living in California routinely report that they expect to be significantly happier there, primarily on the strength of California's blissful climate. People who actually live in California are no happier than the rest of us (Schkade and Kahneman [8]). Far from being a California quirk, this phenomenon is sufficiently widespread as to prompt the conclusion that "Nothing ... will make as much difference as you think." [8, p. 345][1]

Psychologists interpret these findings by drawing a distinction between *decision utility* and *experienced utility* (e.g., Kahneman and Thaler [3]). Decision utilities are the utilities that determine (or at least describe, in a revealed-preference interpretation) our choices. For Schkade and Kahneman [8], these are the utilities people reveal when they contemplate living in California. Experienced utilities are the rewards we realize once the choices are made. For Schkade and Kahneman, these are reflected in the satisfaction reports from people living in California.

Experienced utilities are of no interest to a purely classical economist. Decision utilities suffice to describe the behavior that is the focus of economics. However, if we venture beyond the narrow confines of classical economics to consider the welfare of the individual, the difference is highly relevant. If experienced utilities do not match decision utilities, this casts doubt on the standard economists' presumption that decision utilities are an appropriate guide to well-being. If there is such a difference, should we not exhort people to work more diligently in discerning their future experienced utilities, and then use these to override their decision utilities (as Schkade and Kahneman [8] hint)? Once we have contending utilities (or contending selves, in the common parlance of behavioral economics), such questions are both inevitable and vexing.

We adopt a positive perspective in this paper, answering the following question: Why might we have both decision and experienced utilities in the first place? We take an evolutionary approach. We assume that evolution has equipped agents with utility functions designed to induce fitness-maximizing choices. An agent in our model must make choices in each of two periods that will (along with random events) determine his fitness. Moreover, these choices give rise to an intertemporal trade-off, in the sense that the optimal second-period choice depends upon the alternative chosen in the first period. The first-period choice may determine the agent's health or wealth or skill or status, for example, which may in turn affect how aggressive the agent should be in seeking second-period consumption. Evolution equips the agent with a first-period

[1] This phenomenon, often dubbed a "focusing illusion" (e.g., Loewenstein and Schkade [4]), was thrust into the spotlight by Brickman, Coates and Janoff-Bulman's [1] study of lottery winners and paraplegics, and has become the subject of a large literature, much of it in psychology. See Loewenstein and Schkade [4] for an introduction and Gilbert [2] for an entertaining popular account.

utility function providing the decision utilities shaping the first-period choice. Evolution also equips the agent with a second-period utility function determining the utilities he experiences as a result of his first-period and second-period choices. This latter function serves a dual role, being an experienced utility from the point of view of the first-period choice while also providing the relevant decision utility for the second period. We show that in general, the decision utility shaping the first-period choice does not match the resulting second-period experienced utility. Evolution systematically misleads the agent as to the future implications of his choices.

Why should evolution build an agent to do anything other than maximize fitness, without resorting to conflicting utility notions? Evolution's design problem is complicated by two crucial constraints. First, there are limits on how large and how small are the utilities evolution can give us.[2] By themselves, bounds on utility pose no obstacles. All that matters is that better alternatives get higher utilities, and we can accommodate this no matter how tight the range of possible utilities. However, our second assumption is that the agent is likely to make mistakes when utilities are too close. When alternative 1 provides only a slightly higher utility than alternative 2, the agent may mistakenly choose alternative 2. As a result, there is an evolutionary advantage to having the utility function be as steep as possible, so that the agent is dealing with large utility differences that seldom induce mistakes. This goal conflicts with the bounds on utility. Evolution's response is to make the utility function very steep in the range of decisions the agent is most likely to face, where such steepness is particularly important in avoiding mistaken decisions, and relatively flat elsewhere.[3] If this is to be effective, the steep spot of the utility function must be in the right place. In the second period, the "right place" depends on what happens in the first period, both in terms of what the agent chose and the realization of first-period uncertainty. Evolution thus has a incentive to adjust second-period or experienced utilities in response to first-period outcomes. But if this is to be done without distorting first-period decisions, the agent must not anticipate this adjustment—the experienced utilities guiding second-period decisions must not match the decision utilities shaping first-period decisions.

Section 2 presents the argument in a model that effectively isolates the key features of our analysis, at the cost of a number of simplifications. Section 3 shows how the model can be extended to more realistic situations and discusses its implications.

[2]Notice that in taking this position, we are explicitly adopting a view of utility maximization as a neurological process by which we make choices, rather that simply a description of consistent choices. This is the point of view of much of the current literature in behavioral economics. In particular, our view is that utilities are induced by chemical processes within our brains that are subject to physical constraints.

[3]Robson [7] argues that that utility bounds and limited discrimination between utilities will induce evolution to strategically position the steep part of the utility function. (See also Netzer [5].) Rayo and Becker [6] develop this idea in a model more closely related to that here.

2 Decision and Experienced Utility

2.1 The Evolutionary Environment

There are two periods. The agent makes a choice x_1 in the first period and x_2 in the second. These choices would be multidimensional in a more realistic model, but here are taken for simplicity to be elements of $[0, 1]$. Whenever it is helpful in conveying intuition, we temporarily adopt particular interpretations of x_1 and x_2, such as levels of first-period and second-period consumption or income, or as a decision to move to California (or not) and a subsequent decision of how much time to spend surfing (whether in California or Iowa), or as an investment in status and a subsequent decision of which mate to pursue.

The agent's fitness is determined by his choices x_1 and x_2 as well as the realization s_2 of a random variable $\tilde{s}_2$, reflecting environmental shocks in the second period that may perturb the link between the agent's actions and his fitness. (We add a first-period shock in Section 3.) The agent's health may depend no only on effort he invests in procuring food, but also on random events affecting the productivity of these efforts. We assume that $\tilde{s}_2$ takes on values on an interval $[-S, S]$, and has mean zero and a strictly positive, symmetric, quasiconcave, differentiable density function, with zero derivative only at 0. We let g_2 be the density of s_2.

The agent's choice must be made in ignorance of the realization of $\tilde{s}_2$. The second-period choice of x_2 is made after observing the realization s_1 of $\tilde{s}_1$ (and recalling the first-period choice of x_1), but without observing the realization of $\tilde{s}_2$.

In the absence of any constraints, evolution's task of designing an agent to make fitness-maximizing choices would be trivial (and our assumptions about the timing of the agent's choices and the realizations of the random shocks would be irrelevant). The agent's fitness-maximization problem has a maximizer (x_1^*, x_2^*). Evolution could then simply "hard-wire" agents to make this optimal decision decision.

The point of departure for our analysis is the assumption that evolution *cannot* simply hard-wire us to choose (x_1^*, x_2^*), as trivial as this sounds in the context of this model. Our interpretation here is that what it means to choose a particular value x_1 or x_2 changes with the context in which the decision is made. The first-period choice may consist of an investment in status that sometimes involves hiding food or constructing a shelter, and other times acquiring education, sometimes involves cultivating social relationships with others, and other times driving others away. Similarly, the second-period choice may involve consumption that sometimes calls for stalking food through the forest and other times sitting at a desk. Moreover, the relevant context fluctuates too rapidly for evolution to adapt. The dominant form of investment can change from clearing fields to learning *C++* too quickly for mutation and selection to keep pace. As a result, evolution must recognize that the agent will frequently face problems that are entirely novel from an evolutionary perspective.[4]

[4]Rayo and Becker [6] must similarly address the question of why evolution cannot hard-

To capture this constraint, we assume that the agent's fitness is given by

$$z_1 + z_2$$

where z_1 and z_2 determined by the agents actions and the random shocks via the functions $\tilde{z}_1$ and $\tilde{z}_2$. We interpret z_1 can z_1 as intermediate goals that evolution can identify and use as input in shaping the agent's behavior, associating z_1 with the first period and z_2 with the second. The agents choices, along with the random shocks s_1 and s_2, determine z_1 and z_2. For example, x_1 may reflect an investment in skills, and z_1 the resulting skill level. Alternatively, x_1 may reflect actions taken in pursuit of status, and z_1 the resulting status. In the second period, x_2 may be a choice of foraging strategy and z_2 may indicate how well nourished is the agent. The key distinction is that, while evolution cannot attach utilities to x_1 and x_2, she can to z_1 and z_2. Times have changed too quickly for evolution to attach utility to passing through the drive-through coffee line in the morning, but she can reward the resulting feeling of alertness. Evolution cannot tell whether status is acquired by amassing physical prowess or wealth, but can reward the acquisition of status. The linear form of this fitness assumption significantly simplifies the analysis. It would cost only extra notation to incorporate a discount factor. The technology converting the agent's choices into evolutionary goals is given by

$$\begin{aligned} z_1 &= \tilde{z}_1(x_1) && (1) \\ z_2 &= \tilde{z}_2(z_1, x_2) + s_2. && (2) \end{aligned}$$

Notice that what matters for the second-period outcome is not the means by which the first-period evolutionary goal z_1 is achieved, but the goal itself. In many cases, this seems quite natural. It matters how skilled or well-nourished or prominent is the agent, not how he got that way. We assume that $\tilde{z}_2$ is increasing in z_1, so that the first-period investment has salutary effects for second-period fitness. To keep the analysis simple, we assume that $\tilde{z}_2$ is strictly concave in x_2 (for any z_1). This ensures the existence of a unique maximizer $x_2^*(z_1)$ of $\tilde{z}_2$ for any value z_1 which, for ease of interpretation, we take to be interior. Similarly, in the first period we assume that $\tilde{z}_1(x_1) + \tilde{z}_2(\tilde{z}_1(x_1), x_2^*(z_1))$ is strictly concave in x_1, giving a unique first-period maximize maximizer that we take to be unique.

2.2 Utility Functions

Evolution can give the agent a first-period utility function $V(z_1 + z_2)$, as well as a second-period utility function $V(z_2|z_2)$ regulating the choice of x_2. Through

wire agents to make optimal choices. They assume that the evolutionarily optimal action depends upon an environmental state, and that there are so many possible values of this state that it is prohibitively expensive for evolution to hard-wire the agent to condition actions on every value. Our assumption that the state is entirely novel simply takes this formulation to a corner solution. Rayo and Becker explicitly include the state environmental within their model, while we sweep it into the background, simply assuming that evolution cannot dictate optimal choices. These accounts differ primarily in emphasis.

the technology given by (1)–(2), these implicitly become utility functions on x_1, x_2 and s_2.[5] Notice that we write the second-period utility function as $V_2(z_2|z_1)$. In the second period, the agent's task is to choose x_2, and this choice of x_2 and the realization s_2 then determine z_1 and hence the agent's utility. Our notation emphasizes that this second-period utility function may depend on the agent's previous choice of x_1 and realization of s_1, through their determination of z_1. An agent rendered well-nourished and healthy by his first-period outcome may have a different utility function than one famished and week. An agent who has climbed to the top of the status order in the first period may have a different utility function than a social outcast.

Evolution's task is to design these utility functions so as to induce fitness-maximizing choices. In the absence of any additional constraints (beyond the inability to write utilities directly over x_1 and x_1), evolution's utility-function design problem is trivial. She need only give the agent the utility functions

$$V_1(z_1 + z_2) \;=\; V_2(z_2|z_1) \;=\; z_1 + z_2 \;=\; \tilde{z}_1(x_1) + \tilde{z}_2(\tilde{z}_1(x_1), x_2) + s_2.$$

As straightforward as this result is, we believe it misses some important evolutionary constraints that we introduce in two steps. Our first assumption is that evolution faces limits on how large or small a utility she can induce. Our view here is that utilities must be produced by physical processes, presumably the flow of certain chemicals in the brain. The agent makes choices leading to a fitness levels z_1 and z_2 (or perhaps imagines such choices when evaluating them), and receives pleasure from the resulting cerebral chemistry. There are then bounds on just how strong (or how weak) the resulting sensations can be. Without loss, we assume that utilities must be drawn from the interval $[0, 1]$.

These constraints alone pose no difficulties. Essentially, evolution need simply recognize that utility functions are unique only up to linear transformations. In particular, in this case evolution need only endow the agent with the utility function

$$V_1(z_1 + z_2) \;=\; V_2(z_2|z_1) \;=\; A + B[z_1 + z_2],$$

where A and B are chosen (in particular, with B sufficiently small) so as to ensure that utility is drawn from the unit interval, no matter what the feasible values of x_1, x_2, s_2.[6]

[5] Notice that we are effectively acting "as if" the agent "knows" the functions $\tilde{z}_1$ and $\tilde{z}_2$. We view this not literally as a cognitive understanding, but as the agent's having effectively learned which choices of x_1 and x_2 lead to high utilities. Who among us is sure we understand the nuances o of the connection between diet and good nutrition? At the same time, who has not learned which foods they like, and which they dislike?

[6] Our assumption that the random variable $\tilde{}_2$ has bounded support allows evolution to keep utilities within $[0, 1]$ via a linear transformation. Were these supports unbounded, the corresponding transformation would have to be nonlinear. This would be immaterial in the current example, where the relevant lotteries are ordered by first-order stochastic dominance, ensuring that tradeoffs between risk and return do not arise, but this issue may reappear in more general settings. We are not troubled by assumptions that $\tilde{s}_1$ and $\tilde{s}_2$ have bounded support, on the strength of the belief that there are likely to be bounds on consumption. To put it differently, we would be unconvinced by a divergence between decision and experienced utility that depended crucially on the possibility of unbounded consumption.

We now add a second constraint to evolution's problem—there are limits to the ability of the agent to perceive differences in utility. When asked to choose between two alternatives whose utilities are very close, the agent may be more likely to choose the alternative with the higher utility, but is not certain to do so.[7] This is in keeping with our interpretation of utility as reflecting physical processes within the brain. A very slightly higher dose of a neurotransmitting chemical may not be enough to ensure the agent flawlessly chooses the high-utility alternative.[8]

To make this restriction precise, we begin with the second period and assume the individual cannot distinguish any pair of choices whose expected utility is within $\varepsilon_2 > 0$ of each other. Hence, instead of certainly choosing the maximizer $x_2^*(z_1)$ of $V_2(z_2|z_1) = V_2(\tilde{z}_2(z_1, x_2) + s_2|z_1)$ in the second period, the agent may choose any x_2 with the property that

$$E_{\tilde{s}_2} V_2(\tilde{z}_2(z_1, x_2^*(z_1)) + s_2|z_1) - E_{\tilde{s}_2} V_2(\tilde{z}_2(z_1, x_2) + s_2|z_1) \leq \varepsilon_2.$$

To keep things simple, we assume the agent chooses uniformly over the resulting satisficing set $[\underline{x}_2(z_1), \overline{x}_2(z_1)]$, where $\underline{x}_2(z_1) < x_2^*(z_1) < \overline{x}_2(z_1)$ and[9]

$$\begin{aligned} & E_{\tilde{s}_2} V_2(\tilde{z}_2(z_1, \underline{x}_2(z_1)) + s_2|z_1) && (3) \\ = \; & E_{\tilde{s}_2} V_2(\tilde{z}_2(z_1, \overline{x}_2)(z_1) + s_2|z_1) \\ = \; & E_{\tilde{s}_2} V_2(\tilde{z}_2(z_1, x_2^*(z_1)) + s_2|z_1) - \varepsilon_2. && (4) \end{aligned}$$

Evolution chooses the utility functions V_2 to maximize fitness, subject to (3)–(4), for each z_1. Notice that we will then have $\tilde{z}_2(z_1, \overline{x}_2(z_1)) = \tilde{z}_2(z_1, \underline{x}_2(z_1))$, and that maximizing fitness is equivalent to maximizing this value.

In the first period, the agent has utility function $V_1(z_1 + z_2)$. This utility is again bounded, so that $V_1 \in [0, 1]$. In addition, the individual cannot distinguish any pair of choices whose expected utility is within $\varepsilon_1 > 0$ of each other. This again leads to a random choice from a satisficing set, whose specification we delay until Section 2.4.

2.3 The Second Period

Consider now the solution to Nature's optimization problem in the second period. The agent's fitness will be higher the smaller is the satisficing set

[7] We might also introduce some friction into the agent's learning of the functions $\tilde{z}_1$ and $\tilde{z}_2$ transforming x_1 and x_2 into fitnesses. We put aside such frictions here to concentrate on the optimal shape of the utility function.

[8] Very small utility differences pose no problem for classical economic theory, where differences in utility indicate that one alternative is preferred to another, with a small difference serving just as well as a large one. However, it is a problem when utilities are induced via physical processes. The psychology literature is filled with studies documenting the inability of our sense to reliably distinguish between small differences If the difference between two chemical flows is arbitrarily small, we cannot be certain that the agent will invariably choose the larger.

[9] We could work with a more general assumption about how the choice from the satisficing set is made, but must preclude the possibility that an attempt by evolution to improve the agent's decisions by increasing $\underline{x}_2(z_1)$ and decreasing $\overline{x}_2(z_1)$ is thwarted by agent's pushing more and more of her choice probability toward these boundaries.

$[\underline{x}_2(z_1), \overline{x}_2(z_1)]$, and hence the larger are the fitnesses $\tilde{z}_2(z_1, \overline{x}_2(z_1)) = \tilde{z}_2(z_1, \underline{x}_2(z_1))$.

Consider the task of maximizing $\tilde{z}_2(z_1, \overline{x}_2(z_1)) = \tilde{z}_2(z_1, \underline{x}_2(z_1))$, subject to the constraints given by (3)–(4). A change of variable allows us to rewrite the constraints in (3)–(4) as

$$
\begin{aligned}
& \int V_2(\tilde{z}_2(z_1, x_2^*(z_1)) + s_2|z_1)g_2(s_2)ds_2 - \int V_2(\tilde{z}_2(z_1, \underline{x}_2(z_1)) + s_2|z_1)g_2(s_2)ds_2 && (5)\\
= & \int V_2(\tilde{z}_2(z_1, x_2^*(z_1)) + s_2|z_1)g_2(s_2)ds_2 - \int V_2(\tilde{z}_2(z_1, \overline{x}_2(z_1)) + s_2|z_1)g_2(s_2)ds_2 && (6)\\
= & \int V_2(z_2|z_1)[g_2(z_2 - \tilde{z}_2(z_1, x_2^*(z_1))) - g_2(z_2 - \tilde{z}_2(z_1, \underline{x}_2(z_1)))]dz_2 && (7)\\
= & \int V_2(z_2|z_1)[g_2(z_2 - \tilde{z}_2(z_1, x_2^*(z_1))) - g_2(z_2 - \tilde{z}_2(z_1, \overline{x}_2(z_1)))]dz_2 && (8)\\
= & \ \varepsilon_2.
\end{aligned}
$$

The difference given by (5) is decreasing in $\tilde{z}_2(z_1, \underline{x}_2(z_1))$, as is (6) decreasing in $\tilde{z}_2(z_1, \overline{x}_2(z_1))$. As a result, the optimal utility function must maximize (7) (subject to maintaining the value of (8)) and maximize (8) (subject to maintaining the value of 7)), for the optimal $\tilde{z}_2(z_1, \underline{x}_2(z_1))$ and $\tilde{z}_2(z_1, \overline{x}_2(z_1))$.[10] However, (7) and (8) are both increased by setting the utility $V_2(z_2|z_1)$ as small as possible when $g(z_2 - \tilde{z}_2(z_1, x_2^*(z_1))) - g(z_2 - \tilde{z}_2(z_1, \underline{x}_2(z_1))) = g(z_2 - \tilde{z}_2(z_1, x_2^*(z_1))) - g(z_2 - \tilde{z}_2(z_1, \overline{x}_2(z_1))) < 0$; and by setting the utility $V_2(z_2|z_1)$ as large as possible when this inequality is reversed.

There will then exists a threshold $\hat{z}_2(z_1)$, satisfying

$$
\begin{aligned}
& g(\hat{z}_2(z_1) - \tilde{z}_2(z_1, x_2^*(z_1))) - g(\hat{z}_2(z_1) - \tilde{z}_2(z_1, \underline{x}_2(z_1)))\\
= \ & g(\hat{z}_2(z_1) - \tilde{z}_2(z_1, x_2^*(z_1))) - g(\hat{z}_2(z_1) - \tilde{z}_2(z_1, \overline{x}_2(z_1)))\\
= \ & 0
\end{aligned}
$$

such that[11]

$$
\begin{aligned}
V_2(z_2|z_1) &= 0, \text{ for all } z_2 < \hat{z}_2(z_1)\\
V_2(z_2|z_1) &= 1, \text{ for all } z_2 > \hat{z}_2(z_1).
\end{aligned}
$$

Moreover, in the limit as $\varepsilon_2 \to 0$, we have

$$\hat{z}_2(z_1) = \tilde{z}_2(z_1, x_2^*(z_1)).$$

[10] Otherwise, one could (for example) find a different utility function that gives a larger value for (7) while preserving (8) and holding $\tilde{z}_2(z_1, \overline{x}_2(z_1)) = \tilde{z}_2(z_1, \underline{x}_2(z_1))$ constant at their candidate optimum values, introducing slack into the constraint given by (5). But then the value of $\tilde{z}_2(z_1, \underline{x}_2(z_1))$ that induces equality in (5) would have to be higher than the candidate optimum value, a contradiction.

[11] Our unimodality assumption on the distribution g ensures, for example, that $g(z_2 - \tilde{z}_2(z_1, x_2^*(z_1))) - g(z_2 - \tilde{z}_2(z_1, \underline{x}_2(z_1)))$ is negative for small z and positive for large z, with a unique zero identified by $\hat{z}_2(z_1)$.

2.4 The First Period

Attention now turns to the first period. For simplicity, we take the limit $\varepsilon_2 \to 0$ before considering the optimal first-period utility function. While not essential to the results, this allows us to avoid a host of technical issues. In particular, the satisficing set of first-period choices will then be of the form $[\underline{x}_1, \overline{x}_1]$ where

$$\begin{aligned} & E_{\tilde{s}_2} V_1(\tilde{z}_1(\underline{x}_1) + \tilde{z}_2(\tilde{z}_1(\underline{x}_1), x_2^*(\tilde{z}_1(\underline{x}_1)) + s_2) && (9) \\ = \;& E_{\tilde{s}_2} V_1(\tilde{z}_1(\overline{x}_1) + \tilde{z}_2(\tilde{z}_1(\overline{x}_1), x_2^*(\tilde{z}_1(\overline{x}_1)) + s_2) \\ = \;& E_{\tilde{s}_2} V_1(\tilde{z}_1(x_1^*) + \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*(\tilde{z}_1(x^*)) + s_2) - \varepsilon_1. && (10) \end{aligned}$$

Notice that $\tilde{z}(\underline{x}_1) = \tilde{z}(\overline{x}_1)$. In the first period, the individual randomizes uniformly over the set $[\underline{x}_1, \overline{x}_1]$. Evolution chooses the utility function $V_1(z_1 + z_2)$ to maximize fitness, subject to (9)–(10).

In the first stage, evolution thus increases fitness by increasing $\underline{x}_1$ and decreasing $\overline{x}_1$, with $\underline{x}_1 \leq x^* \leq \overline{x}_1$, subject to the constraints given by (9)–(10). As with the second period, we can execute a change of variable to rewrite the constraints as

$$\begin{aligned} & \int V_1(\tilde{z}_1(x_1^*) + \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*(\tilde{z}_1(x_1^*))) + s_2) g_2(s_2) ds_2 \\ & \qquad - \int V_1(\tilde{z}_1(\underline{x}_1) + \tilde{z}_2(\tilde{z}_1(\underline{x}_1), x_2^*(\tilde{z}_1(\underline{x}_1))) + s_2) g_2(s_2) ds_2 && (11) \\ = \;& \int V_1(\tilde{z}_1(x_1^*) + \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*(\tilde{z}_1(x_1^*))) + s_2) g_2(s_2) ds_2 \\ & \qquad - \int V_1(\tilde{z}_1(\overline{x}_1) + \tilde{z}_2(\tilde{z}_1(\overline{x}_1), x_2^*(\tilde{z}_1(\overline{x}_1))) + s_2) g_2(s_2) ds_2 && (12) \\ = \;& \int V_1(z_1 + z_2)[g_2(z_1 + z_2 - \tilde{z}_1(x_1^*) - \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*(\tilde{z}_1(x_1^*)))) \\ & \qquad - g_2(z_1 + z_2 - \tilde{z}_1(\underline{x}_1) - \tilde{z}_2(\tilde{z}_1(\underline{x}_1), x_2^*(\tilde{z}_1(\underline{x}_1))))] dz_2 && (13) \\ = \;& \int V_1(z_1 + z_2)[g_2(z_1 + z_2 - \tilde{z}_1(x_1^*) - \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*(\tilde{z}_1(x_1^*)))) \\ & \qquad - g_2(z_1 + z_2 - \tilde{z}_1(\overline{x}_1) - \tilde{z}_2(\tilde{z}_1(\overline{x}_1), x_2^*(\tilde{z}_1(\overline{x}_1))))] dz_2 && (14) \\ = \;& \varepsilon_1. \end{aligned}$$

Analogously to our argument for the second period, we note that the difference in (11) is decreasing in $\underline{x}_1$ and the difference in (12) is increasing in $\overline{x}_2$. The optimal utility function must then maximize (13) and maximize the left side of (14), each subject to maintaining the value of the other, for the optimal values of $\underline{x}_1$ and $\overline{x_1}$. This allows us to conclude that there is a value $\hat{Z}_1$ satisfying

$$\begin{aligned} & g_2(\hat{Z}_1 - \tilde{z}_1(x_1^*) - \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*(\tilde{z}_1(x_1^*)))) - g_2(\hat{Z}_1 - \tilde{z}_1(\underline{x}_1) - \tilde{z}_2(\tilde{z}_1(\underline{x}_1), x_2^*(\tilde{z}_1(\underline{x}_1)))) \\ & g_2(\hat{Z}_1 - \tilde{z}_1(x_1^*) - \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*(\tilde{z}_1(x_1^*)))) - g_2(\hat{Z}_1 - \tilde{z}_1(\overline{x}_1) - \tilde{z}_2(\tilde{z}_1(\overline{x}_1), x_2^*(\tilde{z}_1(\overline{x}_1)))) \\ = \;& \varepsilon_1 \end{aligned}$$

with the property that

$$\begin{aligned} V(z) &= 0, \text{ for all } z < \hat{Z}_1 \\ V(z) &= 1, \text{ for all } z > \hat{Z}_1. \end{aligned}$$

Moreover, in the limit as $\varepsilon_1 \to 0$, we have

$$\hat{Z}_1 = \tilde{z}_1(x_1^*) + \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*(\tilde{z}_1(x_1^*))).$$

2.5 Decision and Experienced Utility

Now let us compare the agent's decision and experienced utilities. Suppose the agent considers the possible outcome (x_1, x_2, s_2). Reinterpreting the model a bit to match the motivating psychology literature, the agent may consider moving to California (the choice of x_1), learning to surf once there (the choice of x_2), and enjoying a certain amount of sunshine (the realization of $\tilde{s}_2$). Let us assume the agent anticipates choosing x_2 optimally in the second period, so that $x_2 = x_2^*(z_1)$. We thus disregard cases in which agents do not correctly anticipate future utilities because they do not anticipate their future optimal choices. For example, the agent may understand that he will spend a great deal more time outdoors in California, but he also understand that he will not go surfing 365 days a year.

If the outcome considered by the agent gives $\tilde{z}_1(x_1) + \tilde{z}_2(\tilde{z}_1(x_1), x_2) + s_2 > \tilde{z}_1(x_1^*) + \tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*) = \hat{Z}_1$, then she anticipates the maximal utility of one, since

$$\tilde{z}_1(x_1)+\tilde{z}_2(\tilde{z}_1(x_1), x_2)+s_2 \; > \; \hat{Z}_1 = \tilde{z}_1(x_1^*)+\tilde{z}_2(\tilde{z}_1(x_1^*), x_2^*) \implies V_1(z_1+z_2) \; = \; 1.$$

However, if the scenario contemplated by the agent at the same time involves a value $s_2 < 0$ (the agent is sufficiently exited about the impending move to California even though anticipating mediocre weather), then his realized experienced utility will be zero, since then

$$z_2 \; = \; \tilde{z}_2(z_1, x_2^*(z_1)) + s_2 \; < \; \tilde{z}_2(z_1, x_2^*(z_1)) \implies V_2(z_2|z_1) \; = \; 0.$$

The agent's decision utility of one thus gives way to an experienced utility of zero.

Alternatively, if the agent considers a situation where $\tilde{z}_1(x_1) + \tilde{z}_2(z_1, x_2) + s_2 > \tilde{z}_1(x_1^*) + \tilde{z}_2(z_1, x_2^*) < 0$, then this generates a decision utility level of zero. However, if, at the same time $s_2 > 0$, her experienced utility will be one.

The agent's decision and experienced utilities will thus sometimes agree, but the agent will sometimes believe he will be (maximally) happy, only to end up miserable, and sometimes he will believe at the start that he will be miserable, only to turn out happy. The agent will be mistaken about his experienced utility whenever his utility projection depends importantly on the realization of the first-period choice than second-period uncertainty (i.e., anticipating a good outcome because he is moving to a great location, despite mediocre weather;

or anticipating a bad outcome because his location is undesirable, despite good weather). The agent's decision utilities fail to take into account that once the first-period choice has been realized, his utility function will adjust to focus on the second period, bring second-period realizations to heightened prominence.

More generally, if we compare the decision utility an agent derives from an arbitrary consumption level z with the expected experienced utility that would arise for z, we have that

$$V_1(z_1 + z_2) \neq E_{\tilde{s}_2} V_2(z_2|z_1) = \Pr\{V_2(z_2|z_1) = 1\} = \Pr\{\tilde{s}_2 \geq 0\}. \qquad (15)$$

In this constrained utility-design problem, it is efficient to dissociate the V_1 function, or decision utility, from the expectation of the $V_2(z_2|z_1)$ functions, or experienced utilities. Indeed, the optimal decision utility V_1 exaggerates the consequences of the first period choice and so is steeper than is expected experienced utility. Figure **??** illustrates.

2.6 Sophisticated Agents?

Evolution here has designed the agent to be naive. Why not make the agent sophisticated? Why not simply let the agent make decisions on the basis of experienced utilities? To gain some insight into this question, consider the agent's first-period choice of z_1. Evolution induces the agent to make an appropriate choice of z_1 by designing the agent to maximize the decision utility $V(z_1 + z_2)$, and hence calling on the agent to maximize expected utility, given by[12]

$$E_{\tilde{s}_2} V(z_1 + z_2) = E_{\tilde{s}_2}[1_{z_1 \geq \tilde{z}_1(x_1^*)} + s_2] = 1_{z_1 \geq \tilde{z}_1(x_1^*)}.$$

This utility function induces fitness-maximizing choices of z_1.

Suppose that instead, evolution designed the agent to maximize the expected value of the correctly anticipated, expected experienced utility produced by the sum $z_1 + z_2$. To examine this possibility, we first note that the second-period utility function, given a value of z_1, is given by (cf. (15))[13]

$$\begin{aligned} E_{\tilde{s}_2} V_2(\tilde{z}_2(z_1, x_2)|z_1) &= \Pr[\tilde{z}_2(z_1, x_2) + s_2 > \tilde{z}_2(z_1, x_2^*(z_1))] \\ &= \Pr[s_2 > \tilde{z}_2(z_1, x_2^*(z_1)) - \tilde{z}_2(z_1, x_2)] \\ &= 1 - g_2(\tilde{z}_2(z_1, x_2^*(z_1)) - \tilde{z}_2(z_1, x_2)) \end{aligned}$$

When choosing x_1, and anticipating an optimal choice of $x_2^*(\tilde{z}_1(x)1))$ in the second period, the agent then has utility

$$E_{\tilde{s}_2} V_2(z_2|z_1) = 1 - g_2(0) = \frac{1}{2}.$$

[12]Notice that this is a strictly increasing, continuous function of z_1 on $[0, 1]$. Hence, while the underlying rewards with which evolution motivates the agent are generated by a step function taking only values 0 and 1, the induced expected utility is continuous.

[13]This is again a function that increases strictly and continuously in c.

Expected experienced utility is thus independent of z_1. Making the agent sophisticated, i.e., allowing the agents to make decisions on the basis of expected experienced utility, leaves the agent with no first-period incentives at all.

Why does making the agent sophisticated destroy incentives? The naive agent understands that a suboptimal choice of z_1 will decrease utility. Should such a suboptimal choice z_1 be made, however, the agent's second-period utility function will adjust to the first-period choice z_1 to still yield an expected experienced utility of $\frac{1}{2}$. From evolution's point of view, this adjustment plays the critical role of enhancing second-period incentives. Should the agent be sophisticated enough anticipate it, however, first period incentives evaporate, with expected utility now being independent of the first-period choice.

The intuition behind this result is straightforward. Evolution must create incentives in the first period, and naturally constructs decision utilities to penalize suboptimal choices. However, once a first-period alternative is chosen, evolution must now induce the best possible second-period choice. She accordingly adjusts the agent's utility function in response to the first-period choice, causing the optimal second-period choice to induce the same expected utility, regardless of its first-period predecessor. Suboptimal first-period choices thus lead to the same experienced utility in the second period as do optimal ones. The decision-utility penalty attached to suboptimal choices in the first period is removed in the second in order to construct better second-period incentives.

3 Discussion

3.1 Uncertainty

This section discusses the effect of having a random variable $\tilde{s}_1$ in the first period as well. The key to getting nice results is that $z_1 + z_2$ in the first period satisfy a monotone likelihood ratio property with respect to x_1.

3.2 Nonlinearity

What if, for example, the distribution of s_2 depends on x_1 (it may be sunnier in California than Indiana), or a host of other complications?

3.3 Smooth Utility Functions

Consider now how it is possible to generate a more plausible logistic-shaped utility function, with the addition of a new random shock that is known by the individual, but cannot be conditionned upon by Nature. This shock is perhaps too recently realized for utility adaptation to occur. For simplicity, we revert to a linear formulation.

There are two periods $1, 2$. In the first period, the individual must choose $b \in [0, 1]$ and, in the second, she must choose $c \in [0, 1]$. The first period choice must be made in ignorance of a real valued random variable, $\tilde{r}$, which can be decomposed into the sum of two further random variables as $\tilde{r} = \tilde{s} + \tilde{\eta}$. In the second period, the choice is made after the realization of $\tilde{s}$ is known, but still in ignorance of $\tilde{\eta}$. Both $\tilde{r}$ and $\tilde{\eta}$ are symmetric and unimodal, with mode and mean zero, for simplicity. The cdf's of $\tilde{r}$ and $\tilde{\eta}$ are F and G, respectively, with pdf's f and g. Furthermore, $\tilde{r}$ has support contained in $[-\gamma, \gamma]$ and $\tilde{\eta}$ has support contained in $[-\delta, \delta]$.

In addition, there is a further random shock $\tilde{t}$ which is identically and independently distributed across the two periods. It is assumed that $\tilde{t}$ has a distribution given by $(t_1, ..., t_N; p_1, ..., p_N)$. Each realization of this shock is known to the individual, but does not lead to utility adaptation. Perhaps the shock $\tilde{s}$, which is adapted to, is realized substantially sooner than is the second period shock $\tilde{t}$, which is not adapted to.

In the second period, the agent is endowed with a hedonic Bernoulli utility function which conditions on the realizations s of the random variable $\tilde{s}$ and on the realized choice of b. Suppose this utility is then $V_{s,b}(u)$, which is strictly increasing in fitness u. As before, $V_{s,b}(u) \in [0, 1]$. There is again a satisficing set of the form $[\hat{c}, 1]$ where

$$E_{\tilde{\eta}} V_{s,b} \left(b + 1 + s + \tilde{\eta} + t\right) - E_{\tilde{\eta}} V_{s,b} \left(b + \hat{c} + s + \tilde{\eta} + t\right) = \varepsilon_2. \qquad (16)$$

The individual randomizes uniformly over the satisficing set $[\hat{c}, 1]$.

In the first period, the agent has hedonic Bernoulli decision utility function $V(u)$, which is strictly increasing in fitness u, where $V(u) \in [0, 1]$. The satisficing set of choices will then be of the form $[\hat{b}, 1]$ where

$$E_{\tilde{\eta}} V(1 + c + s + \tilde{\eta} + t) - E_{\tilde{\eta}} V(\hat{b} + c + s + \tilde{\eta} + t) = \varepsilon_1. \qquad (17)$$

In the first period, the individual randomizes uniformly over $[\hat{b}, 1]$.

Consider now the solutions to Nature's optimization problems. For the second period, we rewrite (16) as

$$\begin{aligned} \int V_{s,b}(b + 1 + s + \eta + t) g(\eta) d\eta - \int V_{s,b}(b + \hat{c} + s + \eta + t) g(\eta) d\eta &= \qquad (18) \\ \int V_{s,b}(u) \left(g(u - b - 1 - s - t) - g(u - b - \hat{c} - s - t)\right) du &= \varepsilon_2 \end{aligned}$$

Define now $\hat{c}_n = \hat{c}(t_n)$ and $u_n(s, b)$ (temporarily abbreviated to u_n) by the requirement that

$$g(u_n - b - 1 - s - t_n) = g(u_n - b - \hat{c}_n - s - t_n).$$

It is assumed now that δ is small enough that there is at most one value of t_n giving rise to nonzero terms of the form $g(u_n - b - 1 - s - t_n)$ or $g(u_n - b - \hat{c}_n - s - t_n)$ in each range of fitness u. It is clear then that the optimal choice of $V_{s,b}$ is constant on $[u_{n-1}, u_n)$, then jumps and is again constant on $[u_n, u_{n+1})$. Suppose these values are V_n and V_{n+1}, respectively. It is also clear that optimal choice of $V_1 = 0$ and $V_{N+1} = 1$.

The constraints in (18) become

$$G(u_n - b - \hat{c}_n - s - t_n) - G(u_n - b - 1 - s - t_n) = \frac{\varepsilon_2}{V_{n+1} - V_n}.$$

Nature's problem is to choose $\{V_n\}_{n=1}^{N+1}$ so as to maximize $\sum_{n=1}^{N} p_n \hat{c}_n$. Notice that $\hat{c}_n$ depends only on V_{n+1} and V_n and so is written $\hat{c}_n(V_{n+1}, V_n)$. The first-order conditions for Nature's problem are then

$$p_n \frac{\partial \hat{c}_n}{\partial V_n} + p_{n-1} \frac{\partial \hat{c}_{n-1}}{\partial V_n} = 0, n = 2, ..., N.$$

Using the envelope theorem, we have

$$\begin{aligned}
\frac{\partial \hat{c}_n}{\partial V_n} &= \frac{-\varepsilon_2}{g(u_n - b - \hat{c}_n - s - t_n)(V_{n+1} - V_n)^2} \\
\frac{\partial \hat{c}_{n-1}}{\partial V_n} &= \frac{\varepsilon_2}{g(u_{n-1} - b - \hat{c}_{n-1} - s - t_{n-1})(V_n - V_{n-1})^2}
\end{aligned}$$

so the first-order conditions become

$$\frac{p_n}{g(u_n - b - \hat{c}_n - s - t_n)(V_{n+1} - V_n)^2} = \frac{p_{n-1}}{g(u_{n-1} - b - \hat{c}_{n-1} - s - t_{n-1})(V_n - V_{n-1})^2},$$

for $n = 2, ..., N$.

In the limit as $\varepsilon_2 \to 0$, we have $\hat{c}_n \to 1$ and $u_n(s, b) \to b + 1 + s + t_n$. In this limit, we then have

$$\frac{V_{n+1} - V_n}{V_n - V_{n-1}} = \sqrt{\frac{p_n}{p_{n-1}}}.$$

It follows that

$$\begin{aligned}
\Delta V_n &= V_{n+1} - V_n = K\sqrt{p_n} \\
\text{and so } V_n &= \sum_{m=1}^{n-1} \Delta V_m = K \sum_{m=1}^{n-1} \sqrt{p_m} \\
\text{where } K &= \frac{1}{\sum_{n=1}^{N} \sqrt{p_n}}.
\end{aligned}$$

For simplicity, we take this limit as $\varepsilon_2 \to 0$ before considering first stage choices.

Considering now the first stage, we rewrite (17) as follows—

$$\begin{aligned}
\int V(2 + r + t) f(r) dr - \int V(\hat{b} + 1 + r + t) f(r) dr &= \\
\int V(u) \left(f(u - 2 - t) - f(u - 1 - \hat{b} - t) \right) du &= \varepsilon_1.
\end{aligned}$$

Define then $\hat{b}_n = \hat{b}(t_n)$ and u_n by the requirement that

$$f(u_n - 2 - t_n) - f(u_n - 1 - \hat{b}_n - t_n).$$

By an entirely analogous argument to that used for the second stage problem, under the assumption that γ is sufficiently small relative to the gaps between the $\{t_n\}_{n=1}^N$, it follows that the optimal choice of V is constant, with value V_n on each interval $[u_{n-1}, u_n)$, for $n = 1, ..., N$. In the limit as $\varepsilon_1 \to 0$, it follows that $\hat{b}_n \to 1$ and $u_n \to 2 + t_n$. Since it again follows that

$$V_n = \sum_{m=1}^{n-1} \Delta V_m = \frac{\sum_{m=1}^{n-1} \sqrt{p_m}}{\sum_{n=1}^{N} \sqrt{p_n}},$$

second period utility is simply first period utility translated horizontally by s.

This generalization also then preserves the key features of the original linear case, as outlined at the end of Section 2. It is optimal to dissociate decision utility $V(u)$ from expected experienced utility $E_{\tilde{s}} V_{\tilde{s},1}(u)$. Indeed, the present case highlights a feature of all three cases considered. That is, decision utility $V(u)$ is a steeper function of fitness u than is expected experienced utility $E_{\tilde{s}} V_{\tilde{s},1}(u)$. This is a consequence of the optimal adaptation of utility. It is then advantageous to exaggerate the actual consequences of the first period choice when this choice is before you. Thus, the "focussing illusion" of Schkade and Kahneman (1998) arises as evolutionarily optimal. At the same time, there is no necessary loss from this "illusion"—the only loss of evolutionary efficiency derives from a limited ability to discriminate. In the limit, as this discrimination error tends to zero and evolutionary optimality obtains, the focussing illusion remains.

3.4 More Periods

This allows us to tell richer stories.

3.5 More Dimensions

Also good for telling richer stories.

3.6 Implications

The punchline goes here.

References

[1] Philip Brickman, Dan Coates, and Ronnie Janoff-Bulman. Lottery winners and accident victims: Is happiness relative? *Journal of Personality and Social Psychology*, 32(8):917–927, 1978.

[2] Daniel Gilbert. *Stumbling on Happiness*. Vintage Books, New York, 2007.

[3] Daniel Kahneman and Richard Thaler. Anomolies: Utility maximization and experienced utility. *Journal of Econommic Perspectives*, 20(1):221–234, 2007.

[4] George Loewenstein and DAvid Schkade. Wouldn't it be nice? Predicting future feelings. In Daniel Kahneman, Ed Diener, and Norbert Schwarz, editors, *Well-Being: T Foundations of Hedonic Psychology*, pages 85–105. Russell Sage Foundation, New York, 1999.

[5] Nick Netzer. Evolution of time preferences and attitudes towards risk. *American Economic Review*, 2008. Forthcoming.

[6] Luis Rayo and Gary Becker. Evolutinary efficiency and happiness. *Journal of Political Economy*, 115(2):302–337, 2007.

[7] Arthur J. Robson. The biological basis of economic behavior. *Journal of Economic Literature*, 39(1):11–33, 2001.

[8] David A. Schkade and Daniel Kahneman. Does living in California make people happy? A focusing illusion in judgments of life satisfaction. *Psychological Science*, 9(5):340–346, 1998.