

Evolutionary Efficiency and Happiness

Luis Rayo

University of Chicago

Gary S. Becker

University of Chicago and Hoover Institution

We model happiness as a measurement tool used to rank alternative actions. Evolution favors a happiness function that measures the individual's success in relative terms. The optimal function is based on a time-varying reference point—or performance benchmark—that is updated over time in a statistically optimal way in order to match the individual's potential. Habits and peer comparisons arise as special cases of such an updating process. This updating also results in a volatile level of happiness that continuously reverts to its long-term mean. Throughout, we draw a parallel with a problem of optimal incentives, which allows us to apply statistical insights from agency theory to the study of happiness.

I. Introduction

For a long time, utility was assumed to depend only on the absolute level of an individual's economic conditions. However, a large body of research now shows that the relative level of these conditions also plays a central role: an individual's utility, whether defined in terms of de-

We thank Patrick Bolton, George Constantinides, Douglas Diamond, Luis Garicano, Matthew Gentzkow, Jonathan Levin, George Loewenstein, Paul Milgrom, Roger Myerson, Andrew Oswald, Philip Reny, Arthur Robson, Larry Samuelson, Hareesh Sapra, Ilya Segal, Lars Stole, Balazs Szentes, our editor Canice Prendergast, and two anonymous referees for valuable suggestions. We are also indebted to Miguel Diaz and Jose Mazoy for excellent research assistance. We thank the University of Chicago Graduate School of Business for financial support. Rayo is grateful to Alexandra Cieslik.

[*Journal of Political Economy*, 2007, vol. 115, no. 2]

© 2007 by The University of Chicago. All rights reserved. 0022-3808/2007/11502-0004\$10.00

cision making or hedonic experience, tends to be sharply influenced by his personal history and social environment. Examples include Markowitz (1952), Stigler and Becker (1977), Frank (1985), Constantinides (1990), Easterlin (1995), Clark and Oswald (1996), and Frederick and Loewenstein (1999).

In the present paper, we study hedonic utility and how it motivates choice. As a matter of terminology, we equate hedonic utility with "happiness." We are interested, specifically, in the following patterns.¹ First, the hedonic impact of sustained changes in economic conditions has a tendency to diminish over time, such as becoming accustomed to an expensive lifestyle (i.e., habit formation). Second, the level of happiness that an individual derives from his economic success is usually affected by the success of his peers (i.e., peer comparisons). Third, happiness is influenced by the individual's prior expectations concerning his own success. And fourth, while happiness is volatile, it tends to revert over time to a relatively stable long-term mean. Taken together, these features mean that the individual is mainly concerned not with his absolute level of success, but rather with the difference between his success and a benchmark that changes over time. Moreover, since these features cross cultural boundaries and age groups, it is reasonable to assume that, at a general level, they are innate.

In this paper, using economic tools, we argue that the above features can be evolutionarily advantageous in the sense of improving the individual's ability to propagate his genes. Our goal, in other words, is to provide a biological foundation for these traits.

We view happiness as a decision-making device that allows the individual to rank alternative courses of action (Damasio [1994] presents neurological evidence consistent with this approach). In particular, we study an abstract choice setting in which, in every period of his life, an individual compares alternative input choices x toward the production of a random output y . Consider, for example, a hunter-gatherer searching for fruit: x describes his foraging strategy and y his level of success. A higher y increases the individual's prospects for survival and reproduction and therefore his ability to propagate his genes. We refer to the expected value of y as the individual's fitness.

Associated with each level of y , the individual experiences a real-valued happiness level V , and he measures the desirability of alternative choices x by means of the expected value of happiness conditional on his choice: $E[V|x]$. Throughout, we assume that the happiness function V is innate. In addition, we assume that the individual's measurement $E[V|x]$ is subject to a set of physical constraints that limit its precision. These constraints are described below.

¹ See Kahneman, Diener, and Schwarz (1999) for a review of the underlying psychology.

Our plan is to show that, given our set of physical constraints, a happiness function V that is based on the individual's relative success—and exhibits the traits previously described—can lead to a more accurate ranking of choices and, therefore, better decisions from a fitness perspective.

A. *Evolutionary Approach*

Our model is based on the following stylized evolutionary approach. We begin with a class of environments that are presumably representative of the ancestral world in which we evolved. For these environments, we find the happiness functions that maximize the individual's expected level of success (i.e., his fitness) subject to our set of physical constraints. Finally, we show that these fitness-maximizing functions exhibit the desired traits. The idea behind this approach is to characterize the theoretical end point of a natural selection process in which, via trial and error, the fitness-maximizing happiness functions have come to replace all the rest. This stylized approach is also employed in different contexts by Binmore (1994), Robson (2001*b*), Samuelson (2004), and Samuelson and Swinkels (2006).

The theoretical problem of finding the fitness-maximizing happiness functions can be conveniently stated as a metaphorical principal-agent problem. As is customary in the literature, the principal represents the process of natural selection, and the agent represents an individual carrying a set of genes. In the present context, the principal designs the innate happiness function of the agent, with the goal of maximizing the propagation of the agent's genes. Importantly, the happiness function is only a means to this end: the principal does not directly care about the agent's happiness level. The agent, on the other hand, is born with the happiness function designed by the principal and, via his actions, seeks only to maximize his level of happiness. In the process, however, he inadvertently serves the principal's goal.

Crucially, this theoretical exercise refers to the ancestral environment in which humans evolved, not the modern world. In particular, when talking about fitness-maximizing happiness functions, we refer to functions that optimized genetic multiplication during hunter-gatherer times (before agriculture and animal domestication were developed). In modern times, on the other hand, we presumably share most of the innate characteristics of our hunter-gatherer ancestors. But since the technological landscape has changed so rapidly since the rise of agriculture, our happiness functions need no longer optimally promote the present multiplication of our genes.

B. Assumptions and Results

The key assumptions behind our results are two physical constraints that restrict the agent's precision when ranking alternative choices. The first constraint is a limit on the agent's perception sensitivity. We assume that two alternative choices x_1 and x_2 cannot be ranked by the agent whenever the difference between his expected happiness levels $E[V|x_1]$ and $E[V|x_2]$ is smaller than some minimum threshold. The second constraint is a bound on the range of values that V can take. In Section III, we provide a physiological motivation for these constraints based on the function of nerve cells, and we also discuss laboratory experiments on rats that are supportive of these assumptions.

As explained below, because of these imperfections, the agent will be capable of a more accurate ranking across choices when he measures the fitness associated with each choice relative to the fitness of alternative options, as opposed to measuring fitness in absolute terms. This is the starting point for all our results.

A close analogy can be drawn with a specific feature of human vision. The nerve cells in our eyes gauge the color and luminosity of an item relative to its direct surroundings, not in absolute terms. In particular, a typical light-processing cell has a center and a periphery, and its level of activation is proportional to the difference in the amount of light received across these two regions. Accordingly, the information it sends to the brain is mainly about contrast.² This property enhances the brain's ability to identify difficult items. Analogously, a happiness function that measures the merit of alternative choices directly in terms of their differences enhances the agent's capacity to single out the most efficient among them.

In previous work, Frederick and Loewenstein (1999) and Robson (2001*a*) present related arguments. Frederick and Loewenstein compare a happiness function that changes over time with an eye that, independently of the property described above, is capable of adjusting to the general luminosity of the environment (e.g., through changes in its pupil). Likewise, Robson compares happiness with a man-made measurement instrument, such as a voltmeter, that can be calibrated to fit the problem at hand. Both forms of adaptation rescale the incoming information so that it better matches the instrument's effective range of activation, therefore allowing for a more accurate discrimination across stimuli. At a general level, this notion is also present in our approach.

² A noticeable consequence is the fact that the color and luminosity of a particular item are perceived very differently depending on the color and luminosity of its direct surroundings. For a detailed exposition and examples, see Kandel, Schwartz, and Jessell (2000, 519).

Throughout our analysis, we draw a parallel between the evolutionary problem and a principal-agent problem of optimal incentives under moral hazard, where V corresponds to a performance payment for the agent. In both cases, the principal who designs V seeks to maximize the signal value of this function.³ This parallel is central to our approach since it allows us to apply statistical concepts from incentive theory to the study of happiness. Moreover, thanks to these statistical concepts, we can construct a unified theory for multiple aspects of happiness.

The optimal happiness functions we derive are based on a purely relative measure of success. In any given period, the agent's happiness depends exclusively on the difference between his output y and an endogenous reference point \hat{y} , which serves as a performance benchmark (the difference $y - \hat{y}$ is the carrier of happiness). This reference point \hat{y} is positioned according to the current opportunities faced by the agent and is updated over time in tandem with changes in these opportunities. A by-product of this updating process is that, over time, the agent's happiness will tend to its mean.

We illustrate this dynamic process in an environment in which the optimal reference point is updated every period to equal the conditional expectation of output for that period. Whenever possible, this conditional expectation will exploit information contained in past levels of output, together with any additional information contained in the output of peers. From here, we derive a rationale for habits and, simultaneously, for peer comparisons. Both these phenomena arise as special cases of the same general process: a reference point that is updated over time in a statistically meaningful way. On the basis of this notion, we derive optimal reference points with specific functional forms.

Below, we derive two types of happiness functions. The first type are step functions that deliver the maximum level of happiness whenever the agent exceeds his performance benchmark \hat{y} , and vice versa. These step functions are optimal under the two assumptions described above on limited perception and bounded happiness. The second type, derived in Appendix A, have a smoother S-shape. These functions arise when a third assumption is added, namely, that the agent has an information advantage over the principal when selecting his actions. In this case, the principal responds by spreading the slope of the happiness function more evenly across the output domain.

We begin our analysis with a static model in which the agent lives for

³ In the standard agency problem, maximizing the signal value of V is the best way to counteract the agent's temptation to shirk. In the current evolutionary problem, on the other hand, there is no equivalent temptation to shirk because the agent does not have an exogenous effort cost. In this case, maximizing the signal value of V is the best way to overcome the agent's measurement limitations.

one period. We then extend this model to a dynamic setting in which his reference point changes over time.

II. Static Model

Consider a representative agent (i.e., a hunter-gatherer) who faces an abstract one-shot project. To fix ideas, suppose that this project amounts to an opportunity to collect fruit. The agent first observes the current state of nature s , which describes the physical configuration of the world, such as the presence of fruit and dangers in specific locations. Next, he selects a course of action $x \in X$, which represents the strategy adopted, such as traveling in a certain direction or climbing a particular tree. The combination of x and s randomly determines a level of output $y \in \mathbb{R}$, namely, the amount of fruit collected.⁴ Denote the conditional probability distribution of output by $f(y|x, s)$, a function known by the agent. Once y is realized, the agent experiences a one-dimensional level of happiness given by $V(y)$ (as discussed below, V will also potentially depend on additional arguments). The agent is born with this happiness function.

The origin of V is an evolutionary trial-and-error process, which we call the principal. The metaphorical goal for the principal when creating V is to promote the production of y , which is simply another way to say that in a population of individuals endowed with a diversity of happiness functions, those producing higher levels of y will have a reproductive advantage. For concreteness, we assume that the principal seeks to maximize the expected value of y , which we call *fitness*. (In our model, this maximization leads to the same results as maximizing the expected value of any other increasing function of y .)

Rather than studying the evolutionary trial-and-error dynamics, we are interested in describing the limiting outcome once sufficient experimentation and selection have taken place, while holding the environment fixed. We represent this limiting outcome by means of an optimization problem in which the principal directly selects a happiness function that maximizes her objective. (In general, an evolutionary process in which genetic traits are passed on to offspring with small random variations might converge to a local maximum that is not globally optimal. However, for the technologies considered below, the global optimum will coincide with a local maximum that is unique.)

Throughout, we assume that V is a function of the agent's output y

⁴Beyond this example, output y is meant to summarize the achievement of proximate evolutionary goals, namely, those tangible goals that favored the ultimate evolutionary goal of genetic replication during the ancestral environment. Examples presumably include wealth, health, and sex, as well as the well-being of friends and kin. Accordingly, the decision variable x represents the actions taken in pursuit of these goals.

rather than his choice of input x . This means that the principal effectively delegates the choice of x to the agent, who has the advantage of knowing the current environmental details s . Binmore (1994, 151) and Robson (2001*b*) follow a similar approach. On the other hand, notice that $V(y)$ could also potentially adjust to the state s . After analyzing our basic model, we consider this case.

The agent measures the desirability of alternative choices x , given the current state s , via his calculation of the expected value of happiness: $E[V|x, s] = \int V(y)f(y|x, s)dy$. This calculation leads to a state-contingent utility function $u(x, s) \equiv E[V|x, s]$, built by the agent himself, which delivers a decision criterion. In what follows, we assume that the agent can freely dispose of y . This allows us to focus without loss on nondecreasing happiness functions.

III. Measurement Imperfections

In the absence of any further constraints, the evolutionary problem has a simple solution: the principal can select the identity function $V(y) = y$. Under this function, the principal and agent's objectives become perfectly aligned since they both seek to maximize the expected value of y . However, once our measurement constraints are introduced, this function will no longer be optimal.

A. Limited Perception

Our first constraint imposes a limit on the agent's ability to perceive small differences in his objective $E[V|x, s]$. Given any pair of choices x_1 and x_2 , we assume that there exists a minimum threshold $\epsilon > 0$ such that whenever the difference between $E[V|x_1, s]$ and $E[V|x_2, s]$ is smaller than or equal to ϵ , the agent fails to perceive a difference between these two choices. Below, we treat ϵ as exogenous, but our results hold for all small values of ϵ . Formally, we describe this imperfection by means of a *satisficing set*.

DEFINITION 1. For any given happiness function V , perception threshold $\epsilon > 0$, and state s , the agent's *satisficing set* $\mathcal{S}(V, \epsilon, s)$ is the set of all actions x that deliver an expected happiness $E[V|x, s]$ that is within ϵ distance of the optimized value $\max_{x'} E[V|x', s]$, namely,

$$\mathcal{S}(V, \epsilon, s) \equiv \left\{ x \in X : \max_{x'} E[V|x', s] - E[V|x, s] \leq \epsilon \right\}.$$

CONDITION 1 (Limited perception). We assume that the agent randomly draws his choice from $\mathcal{S}(V, \epsilon, s)$. Moreover, this draw has full

support and is monotonic in the sense that, given ϵ and s , the probability assigned to any subset of \mathcal{S} is inversely proportional to the size of \mathcal{S} .

This condition imposes a coarseness in the agent's perception analogous to the coarseness present in any of our five senses. Our senses, like any other type of signaling activity in our nervous system, are based on the stimulation and firing of individual nerve cells. These nerve cells, however, do not have infinite sensitivity to differences in stimulation, and their firing rates are also subject to errors (see, e.g., Kandel et al. 2000, chap. 21). The overall result is that when two stimulus levels are sufficiently similar, the individual fails to tell them apart.⁵

As a preliminary exercise, it is useful to revisit the principal's problem when this first constraint alone is present. In particular, suppose that the principal uses a rescaled happiness function $V(y) = Ay$, for some constant $A > 0$. In this case, for any given state s , the agent's satisficing set is given by

$$\mathcal{S}(Ay, \epsilon, s) \equiv \left\{ x \in X : \max_{x'} E[Ay|x', s] - E[Ay|x, s] \leq \epsilon \right\},$$

where the happiness function V has been replaced by Ay . Given the linearity of the expectation operators, this set is equal to

$$\mathcal{S}\left(y, \frac{\epsilon}{A}, s\right) \equiv \left\{ x \in X : \max_{x'} E[y|x', s] - E[y|x, s] \leq \frac{\epsilon}{A} \right\}.$$

From here we learn that all the agent's potential choices (those within his satisficing set) lead to an expected value of output within ϵ/A distance of the highest achievable level $\max_{x'} E[y|x', s]$. Moreover, a larger constant A is equivalent to a smaller perception threshold ϵ . This means that the principal can virtually eliminate all choice errors by selecting an arbitrarily large constant A . Therefore, in and of itself, the limited perception constraint favors a happiness function that is arbitrarily steep.

⁵ Experiments with rats are also consistent with this type of limited sensitivity. For example, Gallistel et al. (1991) and Simmons and Gallistel (1994) offer individual rats a choice between two alternative levers that, when pressed, deliver electrical pulses to the pleasure center of their brains. By activating the associated nerve cells, these electrical pulses create a highly desirable effect (indeed, when offered the levers, rats will abandon all other activities). The trade-off faced by the rats is that only one lever can be pressed at a time (before the experiment begins, they sample both levers). Across trials, the researcher can reliably induce a given rat to spend more time pressing a specific lever by increasing the amount of electricity delivered by this lever while holding the alternative constant, provided that this increase is sufficiently large. Small increases, however, lead to essentially random responses.

B. *Bounded Happiness*

Our second constraint is a limit on the range of values that happiness V can take. We assume, in particular, that V must fall within a bounded range $[V_{\min}, V_{\max}]$, for some finite limits V_{\min} and V_{\max} . Below, we treat V_{\min} and V_{\max} as exogenous, but our results are identical for any such limits.

CONDITION 2 (Bounded happiness). For all y , $V(y) \in [V_{\min}, V_{\max}]$.

We view these bounds as a physical limit on our nervous system. As explained by Kandel et al. (2000, 21–32), the fundamental signals in our nervous system are electrical pulses called *action potentials*, which are fired by individual nerve cells. (A specific hedonic state, e.g., would result from a particular firing pattern among a group of cells.) Each one of these electrical pulses is binary (i.e., all or nothing), and they always have the same magnitude. This means that the information transmitted by a given cell is based on the frequency of these pulses, not their individual magnitudes. In addition, owing to a chemical recovery process (called the refractory period), there is a maximum rate at which these pulses can be fired. This limit imposes a physical bound on the potential activity of each cell and therefore a bound on any system composed of a finite number of cells.⁶

Once again, as a preliminary exercise, it is useful to revisit the principal's problem when only this second constraint is present. For concreteness, consider the case in which output y is distributed over a bounded domain $[y_{\min}, y_{\max}]$ of arbitrary size. In addition, consider a linear happiness function $V(y) = Ay + B$, with $A > 0$. Since output is bounded, regardless of the happiness bounds, there always exist levels of A and B such that $V(y)$ is contained within $[V_{\min}, V_{\max}]$. Moreover, given that the agent makes no mistakes, he will always choose an action that fully maximizes his expected happiness $E[Ay + B|x, s]$. But since A is positive, this is equivalent to maximizing the expected value of y . Thus the principal and agent's objectives are again fully aligned.⁷

From the above exercises, we learn that neither one of the two constraints, taken on its own, constitutes a meaningful restriction for the

⁶ The rat experiments described in the previous footnote also provide support for the bounded happiness assumption (see esp. Simmons and Gallistel 1994). The electrical pulses delivered to the pleasure center of the rat's brain cause its nerve cells to fire action potentials in direct proportion to the amount of electricity received. These action potentials are responsible for the desirable experience (nerve cells function in the same fundamental way in rats as in humans). But once the researcher increases the amount of electricity beyond a given level—by increasing current (which increases the number of active neurons) and pulse frequency (which increases the firing rate of each active neuron)—the rat's pleasure center becomes saturated with action potentials, and beyond this level, the rat no longer responds to further increases in electricity.

⁷ On the other hand, when the domain of y is unbounded, this result can still be approximated to an arbitrary degree by means of an arbitrarily small A .

principal. In either case, she can induce arbitrarily efficient decisions by merely rescaling the identity function $V(y) = y$. Notice, moreover, that these rescaled functions are still based on the absolute level of output.

The principal's problem becomes very different, however, when the two constraints are combined. In particular, recall that the limited perception constraint favors an arbitrarily steep happiness function (with an arbitrarily large A), but such a function would exceed the bounds imposed by the second constraint regardless of the size of these bounds. As shown below, this trade-off will lead to happiness functions that are steep only where it matters the most. These functions, in particular, will concentrate their slope over a performance benchmark, or reference point, that targets the agent's output potential. Moreover, the use of this benchmark means that output will effectively be measured in relative terms. We derive specific shapes for these functions below: they will be either step functions (in the main text) or **S**-shaped curves (in App. A).

In general, the principal would benefit from both a larger happiness range $[V_{\min}, V_{\max}]$ and a smaller perception threshold ϵ . This means that the natural selection process would presumably operate over these dimensions as well. However, as mentioned above, our results are not affected by the specific values of $[V_{\min}, V_{\max}]$ and ϵ , provided that the range $[V_{\min}, V_{\max}]$ remains finite and ϵ remains small but strictly positive.

IV. Output Technologies

In order to obtain explicit solutions for the principal's problem, we consider output technologies of the form

$$y = E[y|x, s] + z, \quad (1)$$

where z is a zero-mean exogenous shock that is realized after x is selected.

We assume that this shock z is drawn from a continuous density function with full support in \mathbb{R} that is single-peaked and strictly monotonic on either side of its mean (such as a normal).⁸ We also assume that (a) $E[y|x, s]$ is continuous in x ; (b) the choice space X is a compact and convex subset of \mathbb{R}^N ; and (c) for all states s , $\max_{x'} E[y|x', s] > \min_{x'} E[y|x', s]$, so that the agent's choice problem is nontrivial.

Let $f(y|x, s)$ denote the conditional probability distributions for y associated with these technologies. For any given state s , these distributions have the following properties. First, they are single-peaked at

⁸ Our results do not depend on the fact that the output domain is unbounded. This assumption is made merely for expositional convenience. In particular, given the agent's limited perception, the happiness bounds would bind regardless of the size of the output domain.

their mean $E[y|x, s]$. Second, they are single-crossing in x : for all x_1 and x_2 such that $E[y|x_1, s] \neq E[y|x_2, s]$, the distributions $f(y|x_1, s)$ and $f(y|x_2, s)$ intersect for only one value of y (between their two means). And third, they are ordered across x according to first-order stochastic dominance. These are the key distributional properties we employ.

V. An Efficiency Index

It is useful to summarize the quality of the agent's decision x using a one-dimensional index φ , which we call the *efficiency* of his decision. We define this efficiency index as follows:

$$\varphi(x, s) \equiv \frac{E[y|x, s] - \min_{x'} E[y|x', s]}{\max_{x'} E[y|x', s] - \min_{x'} E[y|x', s]}.$$

Notice that $\varphi(x, s)$ is affected by the agent's choice x only through the first term in the numerator: $E[y|x, s]$. In particular, for any given state s , this term ranges from $\min_{x'} E[y|x', s]$ to $\max_{x'} E[y|x', s]$, depending on the quality of the agent's decision. Accordingly, for any given s , the efficiency index φ ranges from zero to one.

From the definition of φ , expected output becomes

$$E[y|x, s] = \varphi \cdot \max_{x'} E[y|x', s] + (1 - \varphi) \cdot \min_{x'} E[y|x', s], \quad (2)$$

where the arguments of φ have been omitted for clarity. From this expression, we learn that the agent's choice x affects output only through φ .

A one-dimensional choice problem.—In what follows, for analytical purposes, we view the agent as directly choosing an efficiency level φ from the interval $[0, 1]$, as opposed to choosing a multidimensional action x from \mathbb{R}^N . This step is valid because the choice of x affects output only through its impact on φ . As a result, we can directly focus on φ while sending x to the background.

Expressed as a function of φ , the output technology in (1) becomes

$$y = E[y|\varphi, s] + z, \quad (3)$$

where the expectation $E[y|\varphi, s]$ corresponds to the right-hand side of (2):

$$E[y|\varphi, s] = \varphi \cdot \max_{x'} E[y|x', s] + (1 - \varphi) \cdot \min_{x'} E[y|x', s]. \quad (4)$$

For all states s , this expectation is continuous and increasing in φ . This means that higher values of φ are strictly preferred by the principal.

Notice that the state s can potentially play two different roles. First, it can affect expected output directly, independently of the efficiency

of the agent's decision φ . This occurs through its impact over the terms $\max_{x'} E[y|x', s]$ and $\min_{x'} E[y|x', s]$. Through these terms, s influences the agent's overall output potential. For example, s can capture the season of the year in which the hunter-gatherer searches for fruit, with a more favorable season leading to a higher value for $\max_{x'} E[y|x', s]$.

Second, since φ is originally a function of both x and the state s , the latter can influence what the optimal choice of x happens to be. For example, in addition to representing the season of the year, s can indicate which particular trees are likely to have fruit. However, once the choice problem is expressed in terms of φ , this second role is sent to the background.

VI. Optimal Happiness Functions

We begin our analysis with the simplest case in which, conditional on φ , the expectation $E[y|\varphi, s]$ is independent of s . Formally, the terms $\max_{x'} E[y|x', s]$ and $\min_{x'} E[y|x', s]$ in equation (4) are independent of s . (In Sec. VII below, we consider the more general case.) Output now simplifies to

$$y = E[y|\varphi] + z. \quad (5)$$

The associated densities $f(y|\varphi)$ have the following properties (inherited from the original densities $f(y|x, s)$): they are single-peaked at their mean $E[y|\varphi]$, single-crossing in φ , and increasing in φ in the first-order stochastic sense.

A. The Agent's Satisficing Set

In terms of his efficiency choice $\varphi \in [0, 1]$, the agent's satisficing set is given by

$$\tilde{S}(V, \epsilon) \equiv \left\{ \varphi \in [0, 1] : \max_{\varphi'} E[V|\varphi'] - E[V|\varphi] \leq \epsilon \right\}. \quad (6)$$

As before, when selecting φ , the agent will come only within ϵ distance of maximizing his expected happiness. Remark 1 helps us characterize this satisficing set.

REMARK 1. Under the technologies in (5), for any nondecreasing happiness function V , the expected value of happiness $E[V|\varphi]$ is nondecreasing in φ and therefore is maximized when $\varphi = 1$.

Proof. Recall that for any pair of decisions $\varphi_1 > \varphi_2$, the conditional distribution $f(y|\varphi_1)$ dominates $f(y|\varphi_2)$ in the first-order stochastic sense. As a result, the expected value of any nondecreasing function $V(y)$ is weakly higher under φ_1 than under φ_2 . QED

From this remark, it follows that the agent's satisficing set is an interval.

LEMMA 1. Under the technologies in (5), for any nondecreasing happiness function V , the agent's satisficing set $\tilde{\mathcal{S}}(V, \epsilon)$ is given by the interval $[\varphi_{\min}(V, \epsilon), 1]$, where the lower limit $\varphi_{\min}(V, \epsilon)$ is given by the smallest value of $\varphi \in [0, 1]$ (possibly zero) such that $E[V|\varphi] \geq E[V|1] - \epsilon$.

Proof. From remark 1, we know that $\max_{\varphi'} E[V|\varphi'] = E[V|1]$. From this fact and the definition of $\tilde{\mathcal{S}}(V, \epsilon)$ in (6), it follows that a choice $\varphi \in [0, 1]$ belongs to the satisficing set if and only if

$$E[V|\varphi] \geq E[V|1] - \epsilon.$$

Since $E[V|\varphi]$ is nondecreasing in φ (again from remark 1), this inequality implies that the satisficing set is an interval with an upper limit equal to one. Moreover, the lower limit of the satisficing set corresponds to the smallest value of $\varphi \in [0, 1]$ satisfying this inequality, as stated in the lemma. QED

B. The Agent's Choice

According to condition 1 (limited perception), the agent will randomly draw his choice φ from the interval $[\varphi_{\min}(V, \epsilon), 1]$. Let $E[\varphi|V, \epsilon]$ denote the expected value of this choice.

REMARK 2. Under condition 1 and the technologies in (5), the expectation $E[\varphi|V, \epsilon]$ is increasing in $\varphi_{\min}(V, \epsilon)$. In other words, the expected efficiency of the agent's choice increases when his satisficing set shrinks.

Proof. Condition 1 implies that the agent's random draw of φ is monotonic: the smaller his satisficing set $[\varphi_{\min}(V, \epsilon), 1]$, the larger the probability assigned to any subset of this interval. As a result, the expected value of φ must increase as $[\varphi_{\min}(V, \epsilon), 1]$ shrinks. QED

C. The Principal's Problem

The principal's objective is to maximize the expected output $E[y|\varphi]$ considering that the agent's efficiency choice φ is randomly drawn from $[\varphi_{\min}(V, \epsilon), 1]$. But since $E[y|\varphi]$ is linear in φ (from eq. [4]), maximizing $E[y|\varphi]$ is equivalent to maximizing the expected efficiency level chosen by the agent, $E[\varphi|V, \epsilon]$. Accordingly, the principal's problem becomes

$$\max_{\varphi} E[\varphi|V, \epsilon]$$

subject to $V(y) \in [V_{\min}, V_{\max}]$ for all y .

Moreover, from remark 2, maximizing $E[\varphi|V, \epsilon]$ is equivalent to maximizing the lower limit of the agent's satisficing set $\varphi_{\min}(V, \epsilon)$. This observation allows us to restate the principal's problem more compactly:

$$\max_V \varphi_{\min}(V, \epsilon) \quad (\text{P1})$$

subject to $V(y) \in [V_{\min}, V_{\max}]$ for all y .

Intuitively, this problem corresponds to excluding the maximum number of inefficient choices $\varphi < 1$ from the agent's satisficing set.

DEFINITION 2. Let φ^* denote the optimized value for problem (P1).

For any positive ϵ , this value φ^* is less than one. In addition, provided that ϵ is small, φ^* is larger than zero (i.e., at least some inefficient choices can be excluded from the satisficing set). This is the case we focus on.

The following lemma allows us to solve for the optimal V using a dual approach.

LEMMA 2. Suppose that V^* is a solution to problem (P1), namely, $\varphi_{\min}(V^*, \epsilon) = \varphi^*$. In addition, suppose that $\varphi^* > 0$. Then V^* must also solve

$$\max_V E[V|1] - E[V|\varphi^*] \quad (\text{P2})$$

subject to $V(y) \in [V_{\min}, V_{\max}]$ for all y .

Proof. Suppose not. Then there must exist a $V \neq V^*$ (satisfying the constraint) such that $E[V|1] - E[V|\varphi^*] > E[V^*|1] - E[V^*|\varphi^*]$. Moreover, given that $\varphi^* > 0$, we must have $E[V^*|1] - E[V^*|\varphi^*] = \epsilon$ (since this equality defines the lower bound of the agent's satisficing set). When these expressions are combined, it follows that $E[V|1] - E[V|\varphi^*] > \epsilon$. But this implies that $\varphi_{\min}(V, \epsilon) > \varphi^*$, a contradiction to the fact that V^* solves problem (P1). QED

This lemma is central to our analysis. It tells us that any optimal V^* specializes in maximizing the *difference* in expected happiness between choosing the ideal $\varphi = 1$ and choosing the less efficient $\varphi^* < 1$. By maximizing this difference, the principal prevents the agent from choosing an even less efficient level $\varphi < \varphi^*$. All other inefficient choices $\varphi \in [\varphi^*, 1)$, by definition, can never be excluded from the agent's satisficing set.

Using this lemma, we can characterize the optimal V .

PROPOSITION 1. The principal's problem (P1) is solved by a step happiness function V^* such that

$$V^*(y) = \begin{cases} V_{\max} & \text{for all } y \geq \hat{y} \\ V_{\min} & \text{for all } y < \hat{y} \end{cases}$$

where the threshold \hat{y} is uniquely determined by the equality

$f(\hat{y}|1) = f(\hat{y}|\varphi^*)$. Moreover, this solution is unique up to a subset of measure zero.

Proof. The objective in the dual problem (P2) from lemma 2 is equal to

$$\int V(y)[f(y|1) - f(y|\varphi^*)]dy.$$

This integral is maximized by setting $V(y) = V_{\max}$ for every y such that $f(y|1) \geq f(y|\varphi^*)$ and $V(y) = V_{\min}$ for every y such that $f(y|1) < f(y|\varphi^*)$. Moreover, from the single-crossing of f , we have $f(y|1) > f(y|\varphi^*)$ for all $y > \hat{y}$ and $f(y|1) < f(y|\varphi^*)$ for all $y < \hat{y}$. Finally, since $f(y|1) \neq f(y|\varphi^*)$ almost everywhere, this solution is unique up to a measure zero subset and therefore solves problem (P1) as well. QED

This result can be derived graphically. Figure 1a plots two conditional densities $f(y|\varphi)$. The bold curve represents $f(y|1)$, the most desirable density for the principal, and the dashed curve represents $f(y|\varphi^*)$. Recall that the dual objective in lemma 2 is to maximize the difference in expected happiness between these two alternatives. As depicted in figure 1b, this is achieved by delivering the maximum possible happiness V_{\max} for all values of y that are more likely to arise under $\varphi = 1$ than under $\varphi = \varphi^*$ (i.e., all values such that $f(y|1) > f(y|\varphi^*)$), and vice versa. As a result, we obtain a step function with a threshold \hat{y} that lies where the two densities intersect, namely, between their two peaks $E[y|\varphi^*]$ and $E[y|1]$.


Intuitively, recall that the limited perception condition favors an arbitrarily steep happiness function that minimizes the agent's mistakes. However, because of the happiness bounds, this function cannot be steep over the entire domain. In order to cope with this trade-off, the principal chooses a function that is infinitely steep only at a critical point (the reference point), while keeping to the bounds at all other points.⁹

A key property of this function is that it depends exclusively on relative output $y - \hat{y}$. Thanks to this feature, V^* specializes in communicating the merit of alternative choices φ relative to each other, with no reference to the absolute level of fitness $E[y|\varphi]$ associated with each one.

D. Optimal Incentives

A statistical parallel can be drawn between the principal's dual problem (P2)—leading to the step function above—and a problem of optimal

⁹ In Sec. VII and App. A, we discuss the case of smoother S-shaped curves.

FIG. 1.—The optimal V

incentives (e.g., Holmstrom 1979; Levin 2003). In particular, we can interpret V as a performance bonus, φ as an effort variable, and the bounds for V as a two-sided limited-liability constraint. Under this interpretation, the one-step bonus above maximally punishes the agent following a deviation to φ^* . Accordingly, we can view this bonus as implicitly testing the null “ $\varphi = 1$ ” against the alternative “ $\varphi = \varphi^*$ ” on the basis of the realization of y . Under this test, the null is precisely rejected when the likelihood ratio $f(y|\varphi^*)/f(y|1)$ exceeds one, which constitutes an optimal statistical test. In this way, the principal maximizes the signal value of V . This statistical theme will be present throughout.

E. The Limit When $\epsilon \rightarrow 0$

For expositional purposes, it is convenient to work with the limit as ϵ converges to zero. Even though we treat ϵ as exogenous, the optimal happiness functions that arise in this limit serve as an approximation for the happiness functions that are optimal for small values of ϵ .

PROPOSITION 2. For any given $\epsilon > 0$, let $V^*(\epsilon)$ denote the optimal happiness function characterized by proposition 1. As ϵ converges to zero, $V^*(\epsilon)$ converges pointwise to the step function

$$V(y) = \begin{cases} V_{\max} & \text{for all } y \geq E[y|1] \\ V_{\min} & \text{for all } y < E[y|1]. \end{cases}$$

Proof. See Appendix B.

The intuition behind this result is as follows. As $\epsilon \rightarrow 0$, the agent's mistakes become arbitrarily small. Consequently, the optimized lower bound for the satisficing set φ^* converges to one, and therefore, $E[y|\varphi^*]$ converges to $E[y|1]$. Graphically, this means that the lower dashed density in figure 1a converges to the higher bold one. But this implies that the happiness threshold \hat{y} , which lies between the two peaks of these densities, $E[y|\varphi^*]$ and $E[y|1]$, must converge to $E[y|1]$, as claimed by the proposition.

Up to a measure zero subset, this limiting function uniquely maximizes the derivative

$$\frac{\partial}{\partial \varphi} E[V|\varphi]_{\varphi=1} = \int V(y) f_{\varphi}(y|\varphi)_{\varphi=1} dy,$$

which represents the limiting version of the objective in the dual problem (P2). Maximizing this derivative guarantees that marginal deviations away from $\varphi = 1$ have a maximal impact over the agent's objective, thus improving his ability to discriminate. In terms of optimal incentives, this problem corresponds to a first-order approach in which all incentive power is focused over small effort deviations (e.g., Rogerson 1985; Levin 2003). As before, the optimal happiness function communicates differences across choices, not their absolute fitness levels.

VII. Extended Technologies

Here we return to the more general technologies in (3) in which expected output $E[y|\varphi, s]$ can depend directly on s . From equation (4), this means that the extreme values $\max_{x'} E[y|x', s]$ and $\min_{x'} E[y|x', s]$ vary with s .

For the time being, we assume that these extreme values are a function

of a subset of s , denoted Ω . This variable Ω represents, for example, current weather conditions. Accordingly, output becomes

$$y = E[y|\varphi, \Omega] + z, \quad (7)$$

with associated densities $f(y|\varphi, \Omega)$. In addition, we assume that the principal can use Ω as an argument of the agent's happiness function, so that happiness becomes $V(y, \Omega)$.

In this case, since Ω enters both the technology and the happiness function, it effectively becomes a parameter that indexes all our former results. In fact, save for this indexation, the optimal happiness functions are identical to those derived in the previous section. Specifically, in the limit for small ϵ , the optimal function is given by

$$V(y, \Omega) = \begin{cases} V_{\max} & \text{for all } y \geq \hat{y}(\Omega) \\ V_{\min} & \text{for all } y < \hat{y}(\Omega), \end{cases}$$

where the threshold $\hat{y}(\Omega)$, now a function of Ω , satisfies

$$\hat{y}(\Omega) = E[y|\varphi, \Omega]_{\varphi=1}.$$

As before, this threshold corresponds to the peak of the output density $f(y|\varphi, \Omega)$ conditional on the optimal choice $\varphi = 1$.

The difference with the basic model is that the agent's reference point can now explicitly adapt to the environment. For instance, in a day with good weather in which output is likely to be high (as measured by Ω), the agent's reference point will automatically shift to the right. We derive specific examples of this adaptation in the dynamic version of the model in Sections VIII and IX.

S-shaped happiness functions.—Before we proceed, it is worth emphasizing that all the happiness functions derived so far have an extreme step shape. These functions, in particular, can be interpreted as limiting versions of **S**-shaped curves with a slope that is entirely concentrated over the threshold \hat{y} . The reason for this shape is that the principal can fine-tune the happiness functions using all information Ω that is relevant for output. This allows her to precisely target incentives where they matter the most.

In order to obtain functions with a smoother, more realistic, **S**-shape, we consider an extension of the model in Appendix A. In this extension, output is affected by an additional state variable ω , observed by the agent before choosing his action, that is too detailed to be encoded in V . As a result, the agent effectively has an information advantage over the principal. In response, the principal designs **S**-shaped happiness functions with slopes that are more evenly spread over the output domain. In this way, she provides incentives that are targeted more broadly.

These **S**-shaped functions share the same key feature with the step

functions considered so far: they both depend exclusively on the difference between output y and a performance benchmark \hat{y} that adapts over time in tandem with the agent's output potential. This feature is the main focus of our paper. In what follows, with the exception of Appendix A, we focus on step functions in order to make the analysis more transparent.

VIII. Dynamic Model

Here we extend the model to a dynamic setup in which the agent lives for multiple periods $t = 1, 2, \dots$. We equate every period with one separate project. At the beginning of period t , the agent observes a state s_t and selects an action $x_t \in X$. Output is then given by

$$y_t = E[y_t | x_t, s_t] + z_t$$

which satisfies the same properties as above, with z_t independently and identically distributed across time.

Once again, we can reformulate the agent's choice problem so that, every period, he chooses an efficiency level $\varphi_t \in [0, 1]$ instead of a multidimensional action x_t . In this case, output simplifies to

$$y_t = E[y_t | \varphi_t, \Omega_t] + z_t \tag{8}$$

where Ω_t represents the relevant subset of the state s_t that affects output directly. After output is realized, the agent experiences a happiness level $V_t \in [V_{\min}, V_{\max}]$, which we allow to depend on both y_t and Ω_t .

The agent's objective for period t is to maximize the expected value of V_t as opposed to some expected discounted value of future happiness levels. In other words, everything the agent cares about, present and future, is already reflected in his present hedonic utility. This model captures forward-looking behavior by interpreting a given project as being forward-looking itself and motivated by current happiness. Consider, for example, a hunter-gatherer who eats in excess of his current needs in order to accumulate fat, or who helps a friend, precisely because it makes him happy today (a modern counterpart would be an individual who invests in his retirement funds because it increases his present happiness level).¹⁰

Notice that every period of this dynamic model is identical to the

¹⁰ In subsection B below, we also briefly discuss the case in which the agent cares about future happiness above and beyond V_t .

static model studied so far. It follows that, in the limit for small ϵ , the optimal happiness function for any given period t is given by

$$V(y_t, \Omega_t) = \begin{cases} V_{\max} & \text{for all } y_t \geq \hat{y}(\Omega_t) \\ V_{\min} & \text{for all } y_t < \hat{y}(\Omega_t), \end{cases}$$

where the optimal threshold $\hat{y}(\Omega_t)$ satisfies

$$\hat{y}(\Omega_t) = E[y_t | \varphi_t, \Omega_t]_{\varphi_t=1}. \quad (9)$$

As before, these limiting functions serve as an approximation of the happiness functions that are optimal for small values of ϵ , all of which have thresholds close to (9). In what follows, we denote $\hat{y}(\Omega_t)$ more compactly by \hat{y}_t . In addition, recall that in the limit when $\epsilon \rightarrow 0$, the agent's equilibrium action becomes $\varphi_t = 1$. This action also serves as an approximation for the actions selected for small ϵ .

The dynamic properties of this model arise from the fact that \hat{y}_t will be optimally updated over time in order to track changes in the underlying state of technology Ω_t . This updating is governed by the conditional expectation operator in (9), which means that the information contained in Ω_t will be used in a statistically meaningful way. In what follows, we use this fact to derive specific examples.

A. Habit Formation

We proceed with two simple examples in which \hat{y}_t incorporates a habit due to an Ω_t that is correlated across time. Possible causes for this correlation include shocks that determine the agent's intrinsic ability and shocks to the environment, both of which can potentially persist over time.

EXAMPLE 1: A Markovian habit. Suppose that output is given by $y_t = \varphi_t + \theta_t$, where θ_t is a random shock that follows the Markovian process $\theta_t = \theta_{t-1} + z_t$. Equivalently, output can be expressed as $y_t = \varphi_t + \theta_{t-1} + z_t$. This technology satisfies equation (8) with $\Omega_t = \theta_{t-1}$, which is correlated across time. Notice that θ_{t-1} can be inferred from the lagged equality $y_{t-1} = \varphi_{t-1} + \theta_{t-1}$. As a result, output becomes $y_t = \varphi_t + (y_{t-1} - \varphi_{t-1}) + z_t$. In equilibrium, once $\varphi_t = \varphi_{t-1} = 1$, this equation reduces to $y_t = y_{t-1} + z_t$, from which it follows that y_{t-1} (the best predictor of y_t) becomes the optimal reference point:

$$\hat{y}_t = E[y_t | \varphi_t, \Omega_t]_{\varphi_t=1} = y_{t-1}.$$

In this case, the agent will experience a high level of happiness if and only if his current output exceeds what he achieved one period ago.

Analogous to an optimal incentive scheme, in order to best guide the agent, the principal will employ her most accurate source of information regarding φ_t . From the equality $y_t - y_{t-1} = \varphi_t - \varphi_{t-1} + z_t$, we learn that the most accurate source is the difference $y_t - y_{t-1}$. As a result, this difference becomes the carrier of happiness. In contrast, if the reference point did not adapt, as θ_t drifts to extreme values, all decisions φ_t would appear increasingly good or increasingly bad, and thus increasingly similar.

On the other hand, observe that a reduction in φ_t will affect y_t in exactly the same way as a low realization of z_t , implying that these two variables cannot be distinguished by V_t . As a consequence, the principal must punish the agent following low realizations of z_t , and vice versa: luck will inevitably affect happiness.

In fact, in equilibrium, the sole carrier of happiness becomes the random shock z_t (i.e., $y_t - \hat{y}_t = z_t$). This implies that the expected value of happiness is the same for every period, regardless of past levels of output: the effects of luck are always short-lived. These features are shared by all the examples that follow.

In many languages, the word “happiness” is closely linked to “fortune” and “luck.” For the ancient Greeks, for example, happiness (*eudaimonia*) was ultimately determined by the will of the gods: “When viewed through mortal eyes, the world’s happenings—and so our happiness—could only appear random, a function of chance” (McMahon 2004, 7).

EXAMPLE 2: Autoregressive habits. Suppose that output is given by $y_t = \varphi_t + \theta_t$, and θ_t follows the autoregressive process $\theta_t = \sum_{k=1}^{\infty} \alpha_k \theta_{t-k} + z_t$ for arbitrary constants α_k . In this case, Ω_t is the vector $(\theta_{t-1}, \theta_{t-2}, \dots)$. Following steps similar to those in example 1, output becomes $y_t = \varphi_t + \sum_{k=1}^{\infty} \alpha_k (y_{t-k} - \varphi_{t-k}) + z_t$. In equilibrium, this equation reduces to $y_t = \sum_{k=1}^{\infty} \alpha_k y_{t-k} + (1 - \sum_{k=1}^{\infty} \alpha_k) + z_t$, from which we obtain

$$\hat{y}_t = E[y_t | \varphi_t, \Omega_t]_{\varphi_t=1} = \sum_{k=1}^{\infty} \alpha_k y_{t-k} + \left(1 - \sum_{k=1}^{\infty} \alpha_k\right).$$

The reference point is now a weighted average between past levels of output and the equilibrium efficiency level $\varphi_t = 1$. In this case, the individual could potentially adapt more gradually over time, or even adapt only to a partial extent (i.e., when the sum of the α_k coefficients is smaller than one). The specific weights guarantee that the carrier of happiness $y_t - \hat{y}_t$ employs only the new information contained in y_t .

B. Habits and Forward-Looking Behavior

In our model, the presence of habits does not deter the agent from seeking $\varphi_t = 1$. A possible interpretation is that the agent is simply unaware of these habits and therefore does not take them into account. This interpretation is consistent with a common conclusion in the psychology literature that the average individual sharply underestimates the degree to which he will habituate to an improvement in his economic conditions (see, e.g., Gilbert et al. 1998; Loewenstein and Schkade 1999).¹¹

Nonetheless, the presence of habits still opens the possibility that an agent who recognizes the existence of these habits, and manages to internalize their effect, might benefit from a deviation. Whether such a profitable deviation exists depends on how the reference point \hat{y}_t is determined outside equilibrium. So far, this issue has not been discussed.

Consider the technologies of example 2. In general, there are several alternative formulations for the reference point, all of which are equivalent in equilibrium. For example, \hat{y}_t may correspond to an exogenous function of past levels of output, namely,

$$\hat{y}_t = \sum_{k=1}^{\infty} \alpha_k y_{t-k} + \left(1 - \sum_{k=1}^{\infty} \alpha_k\right).$$

Alternatively, \hat{y}_t may equal the best predictor of y_t conditional on $\varphi_t = 1$ and all past information, namely,

$$\hat{y}_t = E[y_t | \varphi_t, \Omega_t] |_{\varphi_t=1} = \sum_{k=1}^{\infty} \alpha_k \theta_{t-k} + 1$$

(where the values of θ_{t-k} are inferred from the technological equalities $\theta_{t-k} = y_{t-k} - \varphi_{t-k}$). In this case, the agent uses her best cognitive abilities to form an output expectation conditional on $\varphi_t = 1$, only to then compare her actual success y_t against this self-imposed benchmark. When $\varphi_{t-k} = 1$ for all k , both formulations coincide. The difference arises outside equilibrium.

In the former case, a reduction in φ_t will reduce future reference points by reducing y_t , thus increasing the expected value of future happiness. As a result, the deviation might indeed be beneficial.¹² In the latter case, in contrast, the agent understands that a reduction in φ_t will not affect his future reference points because they depend only on the

¹¹ See also Burnham and Phelan (2001) and Frey and Stutzer (2004) for a discussion of biased predictions.

¹² For example, if the agent maximizes a geometrically discounted sum of future happiness levels at rate β , a marginal deviation away from $\varphi_t = 1$ will be beneficial if and only if $\sum_{k=1}^{\infty} \beta^k \alpha_k > 1$.

underlying technological shocks θ , and not on the particular realization of y_t . This case describes an agent who cannot change his future output expectations by merely reducing φ_t . Therefore, a deviation is never beneficial. In this case, maximizing a present discounted sum $\sum_{\tau=0}^{\infty} \beta^{\tau} V_{t+\tau}$ is identical to maximizing current happiness V_t alone.

IX. Multiple Agents

In order to derive peer effects, we extend the model to include multiple agents i . During each period t , the agents simultaneously select efficiency levels φ_t^i . These choices randomly determine an output level y_t^i for each agent. Let \bar{y}_t denote the average output across agents, and let $w_t^i \equiv y_t^i - \bar{y}_t$ denote relative output. Our new assumption is that agents will experience common productivity shocks (e.g., because of a shared environment). This implies that peer output will be valuable when assessing individual performance.¹³ Dropping the i superscript, we focus on technologies such that

$$w_t = E[w_t | \varphi_t, \Omega_t] + z_t. \quad (10)$$

Throughout, we assume that these technologies satisfy our previous assumptions with w_t in the place of y_t . In addition, we assume that z_t is independent across agents and that the population average for these shocks is zero (i.e., an exact law of large numbers applies). Finally, we allow V to depend on \bar{y}_t as well as on Ω_t .

From (10), notice that the output density $f(y_t | \varphi_t, \bar{y}_t, \Omega_t)$ depends on y_t and \bar{y}_t only through the agent's relative position w_t . As a result, happiness can be expressed without loss as $V(w_t, \Omega_t)$. It follows that this model is identical to the model with a single agent, with w_t replacing y_t .

Consequently, in the limit for small ϵ , the optimal V_t is a step function with $V(w_t, \Omega_t) = V_{\max}$ for all $w_t \geq \hat{w}_t$, and $V(w_t, \Omega_t) = V_{\min}$ for all $w_t < \hat{w}_t$, where the new benchmark is given by

$$\hat{w}_t = E[w_t | \varphi_t, \Omega_t] |_{\varphi_t=1}.$$

EXAMPLE 3: Static peer comparisons. Suppose that output for each agent is given by $y_t = \varphi_t + \Gamma_t + z_t$. The term Γ_t represents an aggregate shock that is shared by all agents, whereas z_t is the idiosyncratic shock from (10). Both Γ_t and z_t are realized after φ_t is selected. No restrictions

¹³ If the principal directly benefitted from relative output w_t , peer effects would immediately arise (see, e.g., Cole, Mailath, and Postlewaite [1992] for the potential benefits conveyed by w_t). Here we show how these effects can extend beyond any direct advantage of achieving a high w_t .

over the distribution of Γ_t are imposed. In equilibrium, by averaging across agents we obtain $\bar{y}_t = 1 + \Gamma_t$. Therefore, $y_t - \bar{y}_t = w_t = (\varphi_t - 1) + z_t$, which satisfies (10) (in this case, Ω_t is redundant). The optimal reference point for w_t becomes $\hat{w}_t = E[w_t | \varphi_t] |_{\varphi_t=1} = 0$. Accordingly, the reference point for y_t is given by $\hat{y}_t = \bar{y}_t$.

The carrier of happiness now becomes the agent's relative success $y_t - \bar{y}_t$. The reason why \bar{y}_t enters the happiness function is that it filters out the aggregate shock Γ_t , and therefore it increases the statistical power of the measurement device. The resulting happiness function is analogous to a relative performance scheme inside a firm. By tightening the connection between effort and reward, the effect of this scheme is to magnify the cost of withdrawing effort (e.g., Lazear and Rosen 1981; Green and Stokey 1983).

In related work, Samuelson (2004) proposes an additional biological role for peer influences. His work differs from the present approach in terms of both the particular phenomenon under investigation and the mechanism leading to the results. Samuelson is interested in peer imitation (e.g., why we follow the crowd in our consumption decisions), as opposed to the hedonic effects of relative success (e.g., why we derive more happiness from the size of our office, our salary, or our health, when it compares favorably to that of our peers). In particular, he argues that by introducing peer consumption in the agent's utility function, evolution can bias the agent toward imitating his peers as opposed to making an independent decision. In this way, the agent relies less on his own imperfect cognitive skills.

Habits and peers.—A distinctive implication of our model arises when habits and peer comparisons are combined. We begin with a simple example that joins examples 1 and 3.

EXAMPLE 4: A Markovian habit and dynamic peer comparisons. Suppose that output for each agent is given by $y_t = \varphi_t + \Gamma_t + \theta_t$, where Γ_t is an arbitrary aggregate shock, and θ_t is an idiosyncratic shock that follows the Markovian process $\theta_t = \theta_{t-1} + z_t$. The difference from example 1 is the presence of Γ_t , and the difference from example 3 is the persistence of the idiosyncratic shock. Using this technology, we can write $y_t - y_{t-1} = (\varphi_t - \varphi_{t-1}) + (\Gamma_t - \Gamma_{t-1}) + z_t$. Moreover, in equilibrium, $\bar{y}_t - \bar{y}_{t-1} = \Gamma_t - \Gamma_{t-1}$. Combining these expressions, we obtain $w_t = (\varphi_t - \varphi_{t-1}) + w_{t-1} + z_t$, which satisfies (10) once $\varphi_{t-1} = 1$ (with $\Omega_t = w_{t-1} = \theta_{t-1} - \bar{\theta}_{t-1}$). As a result, the reference point in terms of w_t becomes $\hat{w}_t = w_{t-1}$. Accordingly, the reference point in terms of y_t becomes

$$\hat{y}_t = y_{t-1} + \bar{y}_t - \bar{y}_{t-1}.$$

This reference point is not the mere sum of y_{t-1} and \bar{y}_t , which would imply that the carrier of happiness $y_t - \hat{y}_t$ is the difference between the increase in output $\Delta y_t = y_t - y_{t-1}$ and the average peer output \bar{y}_t . This difference would lead to a comparison between an innovation and an absolute level, which is hard to justify on statistical grounds. Rather, the carrier of happiness becomes a double difference:

$$y_t - \hat{y}_t = \Delta y_t - \Delta \bar{y}_t.$$

In this case, while y_{t-1} and \bar{y}_t reduce current happiness, \bar{y}_{t-1} has the opposite effect. We interpret this effect as a generalized process of habituation that extends to the output of peers. Consider, for example, a sudden and permanent increase in \bar{y}_t , with y_t held constant. This increase will initially shift the reference point to the right, with a likely decrease in happiness. But after one period, \bar{y}_t will enter the reference point with a negative sign, shifting it back to its original level. As a result, the agent will have successfully coped.

Equivalently, the carrier of happiness can be expressed as the change in the agent's relative position:

$$y_t - \hat{y}_t = w_t - w_{t-1}.$$

This expression again represents a double difference: the agent compares his current output y_t with the output of his peers \bar{y}_t in order to obtain his relative social position w_t . And, on top of this, he compares w_t with the social position w_{t-1} he occupied one period ago. This leads to a twice-relative happiness function: an advance in social position increases happiness, and vice versa, but this effect is only short-lived.

We conclude with a result that encompasses all the examples above.

PROPOSITION 3. Suppose that output for each agent is given by $y_t = \varphi_t + \Gamma_t + \theta_t$, where Γ_t is an arbitrary aggregate shock, and θ_t follows the autoregression $\theta_t = \sum_{k=1}^{\infty} \alpha_k \theta_{t-k} + z_t$ for arbitrary constants α_k and z_t independently and identically distributed. Then the optimal reference point for period t is given by

$$\hat{y}_t = \sum_{k=1}^{\infty} \alpha_k y_{t-k} + \bar{y}_t - \sum_{k=1}^{\infty} \alpha_k \bar{y}_{t-k}.$$

Proof. See Appendix B.

The carrier of happiness is now a generalized double difference:

$$y_t - \hat{y}_t = \left(y_t - \sum_{k=1}^{\infty} \alpha_k y_{t-k} \right) - \left(\bar{y}_t - \sum_{k=1}^{\infty} \alpha_k \bar{y}_{t-k} \right).$$

The term $\sum_{k=1}^{\infty} \alpha_k y_{t-k}$ corresponds to a conventional habit, whereas the term $\sum_{k=1}^{\infty} \alpha_k \bar{y}_{t-k}$ again represents a habituation to peers. Regardless of the properties of the aggregate shocks (including any intertemporal

correlations), the same coefficients α_k enter both forms of habituation. The reason is that y_{t-k} is affected by the aggregate shock Γ_{t-k} , which is redundant when assessing φ_t . Subtracting \bar{y}_{t-k} from y_{t-k} filters out this shock. The implication is that lagged output and lagged peer output have the opposite effect over happiness. Consider, for example, an individual with a stable level of wealth who compares himself with a neighbor who is currently wealthier. The above formulation allows him to be happier when, for as long as he remembers, this neighbor has always been wealthy, as opposed to the case in which their relative fortunes have been recently reversed.

The carrier of happiness can also be expressed in terms of changes in social position:

$$y_t - \hat{y}_t = w_t - \sum_{k=1}^{\infty} \alpha_k w_{t-k}.$$

In this case, the individual compares his current social position w_t against a weighted sum of the positions he occupied in the past. A key property of this formulation is that a strong process of peer comparisons (where only relative position matters) can coexist with a strong process of habit formation (through coefficients α_k with a sum $\sum_{k=1}^{\infty} \alpha_k$ that is close, or even equal, to one).

In an extension to the present paper (Rayo and Becker 2007), we argue that this simultaneous presence of strong habits and strong peer comparisons can help account for several results obtained in the happiness surveys. Moreover, we argue that this feature has been absent from standard utility formulations, where habits and peer comparisons are effectively treated as substitutes for each other.

X. Concluding Remarks

We have modeled happiness as a biological measurement instrument that guides the agent's decisions. Analogous to an eye that specializes in measuring differences between neighboring objects, a happiness function that evaluates economic success in relative terms serves as a more accurate decision guide. In our model, in particular, the agent's success is evaluated against a reference point that constantly changes over time in tandem with his opportunities. This reference point integrates information that can best predict the agent's performance, including information contained in his past levels of output as well as the output of his peers. As a result, the agent is concerned not with his absolute level of success, but rather with his success relative to a benchmark that reflects his own history and social environment.

Throughout, we have suggested a statistical parallel between happi-

ness and an optimal incentive scheme that seeks to promote effort. This parallel allows us to rationalize multiple aspects of happiness using well-known concepts from incentive theory. Indeed, when viewed from an economic perspective, happiness appears to have multiple signs of statistical inference. For example, in addition to accounting for habits and peer comparisons, this incentive approach rationalizes why luck has an impact over happiness and why this impact is short-lived.

Our discussion of habits and peer comparisons is far from exhaustive. Possible extensions could address the issue of habituation patterns that differ according to the type of good involved, as well as the fact that some goods are more prone to social comparisons than others. In both cases, statistical principles may prove to play a role.

Appendix A

An Illustration of S-Shaped Functions

In the text, we considered the case in which output was determined exclusively by the agent's efficiency choice φ and a state variable Ω . Moreover, Ω was observed by the agent before choosing φ and could also be encoded in the agent's happiness function. This fact allowed the principal to perfectly target incentives via extreme step functions. Here, we extend the model by allowing the agent to observe an additional detailed signal ω , which provides further information regarding his output potential but which can no longer be encoded in V . In this case, the principal no longer foresees the full set of contingencies that can directly influence output. Anticipating this information disadvantage, the principal will design smoother happiness functions that spread incentives more evenly over the output domain.

For analytical convenience, for this exercise we assume that both state variables Ω and ω , as well as output y , are discrete, taking only integer values. At the beginning of the period, the two states are realized and observed by the agent. Conditional on Ω , the finer state ω can take $2N+1$ different values ranging from $\Omega - N$ to $\Omega + N$, for arbitrary $N \geq 0$. The realization of ω is determined by a probability distribution function $P(\omega|\Omega)$ with mean $\omega = \Omega$. We assume that $P(\omega|\Omega)$ is increasing in ω to the left of its mean and decreasing to the right.

Next, the agent selects $\varphi \in [0, 1]$, and output y is realized. We assume that, conditional on ω , output is independent of Ω . Namely, Ω influences output only indirectly via its effect on the distribution of ω . Specifically, we assume that, given ω and φ , output takes one of two values:

$$y = \begin{cases} \omega & \text{with probability } \varphi \\ \omega - 1 & \text{with probability } 1 - \varphi. \end{cases} \quad (\text{A1})$$

In this case, the agent's efficiency choice φ simply transfers probability weight from a low to a high realization of output. This binary structure is by no means critical for the results, but makes the analysis more transparent. Finally, after output is realized, the agent experiences a happiness level $V(y, \Omega)$, which can depend on the first state Ω , but not on ω . Throughout, given free disposal of output, we assume without loss that $V(y, \Omega)$ is nondecreasing in y .

A. Symmetric Information

As a preliminary step, consider the case in which the states Ω and ω are always identical to each other (formally, $N = 0$ and $P(\omega = \Omega|\Omega) = 1$), which means that information is symmetric between principal and agent. From (A1), conditional on Ω , only two values of the agent's happiness function are relevant for his decision, namely, $V(\Omega, \Omega)$ and $V(\Omega - 1, \Omega)$, depending on whether output is high ($y = \Omega$) or low ($y = \Omega - 1$). Let ΔV denote the difference between these two happiness values, which we assume is strictly positive (otherwise, the agent would be indifferent among all his decisions).

Conditional on φ and Ω , the agent's expected happiness is given by

$$\begin{aligned} E[V|\varphi, \Omega] &= \varphi \cdot V(\Omega, \Omega) + (1 - \varphi) \cdot V(\Omega - 1, \Omega) \\ &= \varphi \cdot \Delta V + V(\Omega - 1, \Omega). \end{aligned} \quad (\text{A2})$$

Since this expression is increasing in φ , and therefore is maximized when $\varphi = 1$, the agent's satisficing set is again an interval $[\varphi_{\min}(V, \epsilon), 1]$, where the lower bound $\varphi_{\min}(V, \epsilon)$ solves $E[V|\varphi_{\min}, \Omega] = E[V|1, \Omega] - \epsilon$. From (A2), this lower bound takes a simple form¹⁴

$$\varphi_{\min}(V, \epsilon) = 1 - \frac{\epsilon}{\Delta V},$$

where a larger slope ΔV is equivalent to a smaller perception error ϵ .

From this expression it follows that the principal's problem of maximizing φ_{\min} reduces to maximizing the slope ΔV (subject to the happiness bounds V_{\max} and V_{\min}). This maximization is achieved by setting

$$V(y, \Omega) = \begin{cases} V_{\max} & \text{for all } y \geq \Omega \\ V_{\min} & \text{for all } y \leq \Omega - 1, \end{cases}$$

which is a simple version of the step functions derived in the text.¹⁵

B. Private Information

Consider now the case in which, for each value of Ω , the state ω can take multiple values that are observed only by the agent (i.e., $N \geq 1$). We begin with the simple case in which ω can take one of three values (i.e., $N = 1$), given by $\Omega - 1$, Ω , and $\Omega + 1$. Suppose that the associated distribution $P(\omega|\Omega)$ is symmetric around $\omega = \Omega$. In addition, let $p \equiv P(\Omega|\Omega)$ denote the probability of the intermediate state, and normalize the happiness range $[V_{\min}, V_{\max}]$ to $[0, 1]$.

We proceed by fixing Ω and solving for the optimal function given this state. From the principal's point of view, given Ω , output y can take four values, ranging from $\Omega - 2$ to $\Omega + 1$. Denote these values, in increasing order, by y_1, \dots, y_4 .

¹⁴ This expression implicitly assumes that ϵ does not exceed ΔV . Otherwise, the lower bound φ_{\min} would become zero. However, provided that ϵ is small, this will not occur under the optimal happiness function.

¹⁵ Another case in which step functions are optimal occurs when ω can take multiple values for each value of Ω (i.e., $N \geq 1$), but neither principal nor agent observes ω up-front. In this case, as in the text, information remains symmetric and, conditional on Ω , the distributions of y remain single-crossing in φ (since the probability distributions $P(\omega|\Omega)$ are assumed to be single-peaked at their mean). Therefore, a reasoning similar to that in the text applies.




FIG. A1.—Private information

Accordingly, the happiness function can also take four values, denoted V_1, \dots, V_4 . From the agent's point of view, on the other hand, once ω is learned, output can take only two values—from (A1). For example, in the intermediate state $\omega = \Omega$, output can take only one of the intermediate values y_2 and y_3 , as shown in figure A1.

Associated with each state ω , the agent faces a different satisficing set. Denote the lower bounds of these sets by $\varphi_{\min}(\omega)$ (where the dependence on V and ϵ has been dropped for simplicity). Consider again the intermediate state $\omega = \Omega$. Under this state, the agent's expected payoff equals $\varphi V_3 + (1 - \varphi)V_2$. Accordingly, the lower bound of his satisficing set (which solves $\varphi_{\min}(\Omega)V_3 + [1 - \varphi_{\min}(\Omega)]V_2 = V_3 - \epsilon$) becomes

$$\varphi_{\min}(\Omega) = 1 - \frac{\epsilon}{V_3 - V_2}.$$

Similarly, the lower bounds for the satisficing sets in the extreme states $\Omega - 1$ and $\Omega + 1$ become $\varphi_{\min}(\Omega - 1) = 1 - \epsilon/(V_2 - V_1)$ and $\varphi_{\min}(\Omega + 1) = 1 - \epsilon/(V_4 - V_3)$.

The principal now faces a trade-off: she cannot maximize the lower bound for any given state without simultaneously reducing the lower bound for an adjacent state. For example, the lower bound for the intermediate state $\omega = \Omega$ would be maximized by setting $V_2 = 0$ and $V_3 = 1$, but this would eliminate all incentives in the two extreme states. Similarly, incentives can be improved in the extreme states by increasing V_2 and reducing V_3 , but this would reduce incentives in the intermediate state.

In order to derive an exact solution for this problem, we must assume a specific form for the randomizations followed by the agent within his satisficing sets.

Here, for simplicity, we assume that these randomizations are uniform (but very similar results are obtained under more general distributions). In this case, the principal's objective can be expressed as a weighted average of the lower bounds $\varphi_{\min}(\omega)$:¹⁶

$$\sum_{\omega=\Omega-1}^{\Omega+1} P(\omega|\Omega) \cdot \varphi_{\min}(\omega) = p \cdot \varphi_{\min}(\Omega) + \frac{1}{2}(1-p) \cdot [\varphi_{\min}(\Omega-1) + \varphi_{\min}(\Omega+1)]. \quad (\text{A3})$$

Substituting for the values of $\varphi_{\min}(\omega)$ and eliminating redundant terms, we can express this objective equivalently as

$$-\left[p \cdot \frac{1}{V_3 - V_2} + \frac{1}{2}(1-p) \cdot \left(\frac{1}{V_2 - V_1} + \frac{1}{V_4 - V_3} \right) \right].$$

Accordingly, the principal seeks to minimize the expression in brackets by selecting four values V_1, \dots, V_4 within the happiness range $[0, 1]$. This problem is solved by setting

$$(V_1, V_2, V_3, V_4) = (0, \mu, 1 - \mu, 1),$$

where

$$\mu \equiv \left[2 + \left(2 \frac{p}{1-p} \right)^{1/2} \right]^{-1}$$

is a decreasing function of p .

For example, as illustrated in figure A1, when $p = \frac{1}{3}$, we obtain $\mu = \frac{1}{3}$, and V becomes a straight line. On the other hand, when $p > \frac{1}{3}$, we obtain $\mu < \frac{1}{3}$, and the optimal V becomes S-shaped: the slope between the intermediate values y_2 and y_3 is larger than the slope between the extreme values y_1 and y_2 and between y_3 and y_4 . The reason is that a p larger than $\frac{1}{3}$ induces the principal to place more weight on the intermediate state relative to the extreme ones. In the limit when $p \rightarrow 1$, μ converges to zero, and the optimal V converges to a step function with $V_1 = V_2 = 0$ and $V_3 = V_4 = 1$, that is, the type of function that was optimal under symmetric information between principal and agent.

C. A Greater Number of States

Consider now the case in which ω can take $2N+1$ values for each value of Ω (ranging from $\Omega - N$ to $\Omega + N$) for arbitrary $N \geq 1$. As above, we fix Ω and solve for the optimal function given this state. For any given ω , output can again take one of two values: ω and $\omega - 1$. Let V_ω and $V_{\omega-1}$ denote the happiness levels

¹⁶ To see this, notice that the expected value of output given Ω (the principal's payoff) is $\sum_\omega P(\omega|\Omega) \cdot \{\omega - 1 + E[\varphi|\omega]\}$, where $E[\varphi|\omega]$ represents the expected draw of φ from the satisficing set $[\varphi_{\min}(\omega), 1]$. Moreover, under uniform randomizations, this expected draw equals $(1/2)[1 + \varphi_{\min}(\omega)]$. The objective in (A3) follows from substituting for this value and eliminating redundant terms.

associated with these values, and let $\Delta V_\omega \equiv V_\omega - V_{\omega-1}$. Accordingly, for any given ω , the agent faces a satisficing set with a lower bound $\varphi_{\min}(\omega)$ given by¹⁷

$$\varphi_{\min}(\omega) = 1 - \frac{\epsilon}{V_\omega - V_{\omega-1}} = 1 - \frac{\epsilon}{\Delta V_\omega}. \quad (\text{A4})$$

Under the assumption of uniform randomizations within the agent's satisficing sets, from the same reasoning as above, the principal's problem simplifies to maximizing the weighted average of these lower bounds $\varphi_{\min}(\omega)$. Her choice variables are now $2N + 2$ happiness levels ranging from $V_{(\Omega-N)-1}$ to $V_{\Omega+N}$, where $(\Omega - N) - 1$ is the lowest possible level of output given Ω and $\Omega + N$ is the highest. Accordingly, her problem becomes

$$\begin{aligned} & \max_{V_{(\Omega-N)-1}, \dots, V_{\Omega+N}} \sum_{\omega=\Omega-N}^{\Omega+N} P(\omega|\Omega) \cdot \varphi_{\min}(\omega) \\ & \text{subject to } V_{(\Omega-N)-1}, \dots, V_{\Omega+N} \in [V_{\min}, V_{\max}]. \end{aligned}$$

When we substitute for the values of $\varphi_{\min}(\omega)$ in the principal's objective using (A4) and eliminate redundant terms, this problem further simplifies to

$$\begin{aligned} & \min_{V_{(\Omega-N)-1}, \dots, V_{\Omega+N}} \sum_{\omega=\Omega-N}^{\Omega+N} P(\omega|\Omega) \cdot \frac{1}{\Delta V_\omega} \\ & \text{subject to } V_{(\Omega-N)-1}, \dots, V_{\Omega+N} \in [V_{\min}, V_{\max}]. \end{aligned} \quad (\text{A5})$$

The following proposition provides the solution to this problem as a function of the prior distribution $P(\omega|\Omega)$. A numerical example is provided below.

PROPOSITION 4. Problem (A5) is uniquely solved by setting

$$V_{\Omega+N} = V_{\max},$$

$$V_{(\Omega-N)-1} = V_{\min},$$

and, for all ω ,

$$\Delta V_\omega = P(\omega|\Omega)^{1/2} \cdot C,$$

where C is a constant equal to

$$(V_{\max} - V_{\min}) \cdot \left[\sum_{\omega=\Omega-N}^{\Omega+N} P(\omega|\Omega)^{1/2} \right]^{-1}.$$

Proof. See Appendix B.

From this proposition, we learn that the optimal happiness function uses the full range $[V_{\min}, V_{\max}]$, and, for any given ω , it has a slope ΔV_ω proportional to the square root of the probability of that state, $P(\omega|\Omega)^{1/2}$. In this way, the optimal function concentrates a higher slope over output ranges that are more likely given Ω .

Figure A2 presents an example with $N = 10$ and $[V_{\min}, V_{\max}] = [0, 1]$. Figure A2a graphs two distribution functions $P(\omega|\Omega)$ that are discrete versions of normal densities. Both densities are centered around an arbitrary Ω , with standard de-

¹⁷ As before, this expression implicitly assumes that ϵ does not exceed ΔV_ω . Provided that ϵ is small, this will indeed be the case under the optimal happiness function.

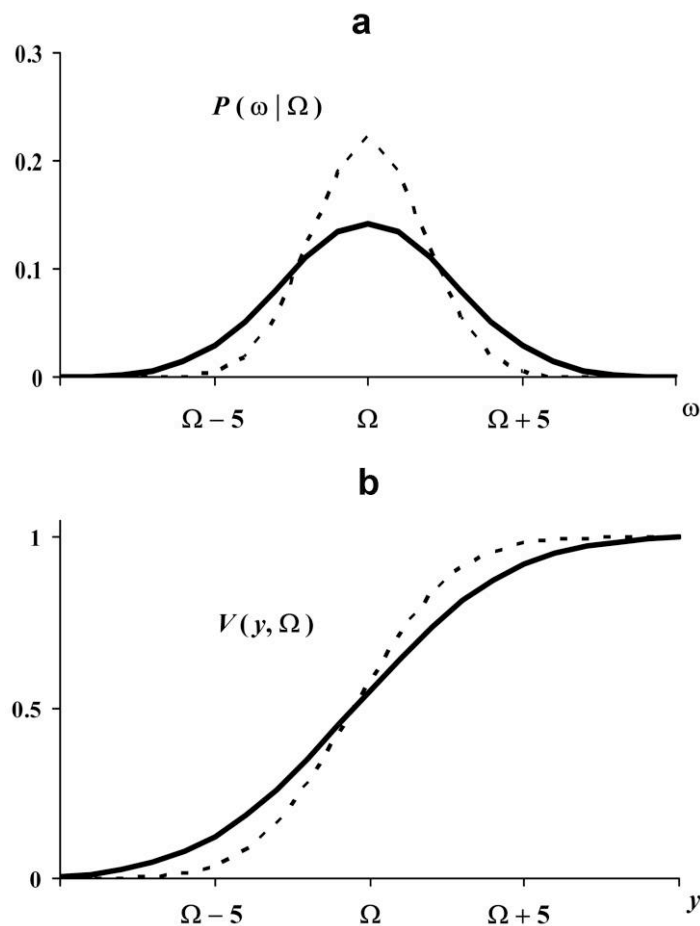


FIG. A2.—S-shaped curves

viations of 1.8 and 2.8. Figure A2*b*, on the other hand, presents the two associated optimal happiness functions derived from proposition 4. As before, a more disperse density leads to a function with a slope that is more evenly spread.

D. Adaptation over Time

Notice that all happiness functions derived in this appendix are fully independent of the absolute value of output y . Indeed, they can all be expressed as a function of the difference between y and the intermediate state Ω . Accordingly, the carrier of happiness becomes $y - \Omega$. As in the text, the particular benchmark $\hat{y} = \Omega$ corresponds to the ex ante expected value of output conditional on $\varphi = 1$. Consequently, as Ω changes over time, these S-shaped functions also shift in tandem with the agent's output potential.

Appendix B**Proofs***Proof of Proposition 2*

For any given $\epsilon > 0$, let $\varphi^*(\epsilon) < 1$ denote the optimized value of problem (P1). In addition, let $\hat{y}(\epsilon)$ denote the threshold for the corresponding optimal function $V^*(\epsilon)$, characterized in proposition 1. In particular, $\hat{y}(\epsilon)$ is uniquely determined by the equality

$$f(\hat{y}(\epsilon)|1) = f(\hat{y}(\epsilon)|\varphi^*(\epsilon)).$$

Since the densities $f(y|1)$ and $f(y|\varphi^*(\epsilon))$ are single-peaked at their respective means $E[y|1]$ and $E[y|\varphi^*(\epsilon)]$, the above equality means that the threshold $\hat{y}(\epsilon)$ must lie between these two means:

$$E[y|\varphi^*(\epsilon)] < \hat{y}(\epsilon) < E[y|1]. \quad (\text{B1})$$

We now claim that, as $\epsilon \rightarrow 0$, $\varphi^*(\epsilon) \rightarrow 1$. To see this, consider any feasible happiness function V_0 that is nondecreasing and nonconstant. Since the densities $f(y|\varphi)$ are stochastically increasing in φ , the expected happiness $E[V_0|\varphi]$ is increasing in φ . It follows that $\varphi_{\min}(V_0, \epsilon)$, the lower limit of the agent's satisficing set under V_0 , must converge to one as $\epsilon \rightarrow 0$. But, by definition, $\varphi^*(\epsilon) \geq \varphi_{\min}(V_0, \epsilon)$, which implies that $\varphi^*(\epsilon)$ must also converge to one.

This claim implies that, as $\epsilon \rightarrow 0$, $E[y|\varphi^*(\epsilon)] \rightarrow E[y|1]$. Therefore, from (B1), the threshold $\hat{y}(\epsilon)$ must converge to $E[y|1]$. The proposition follows as a result. QED

Proof of Proposition 3

Using the above technology, we can write

$$y_t - \sum_{k=1}^{\infty} \alpha_k y_{t-k} = \varphi_t - \sum_{k=1}^{\infty} \alpha_k \varphi_{t-k} + \Gamma_t - \sum_{k=1}^{\infty} \alpha_k \Gamma_{t-k} + z_t.$$

Thus, in equilibrium,

$$\bar{y}_t - \sum_{k=1}^{\infty} \alpha_k \bar{y}_{t-k} = \left(1 - \sum_{k=1}^{\infty} \alpha_k\right) + \Gamma_t - \sum_{k=1}^{\infty} \alpha_k \Gamma_{t-k}.$$

Combining these two expressions we obtain

$$w_t = \sum_{k=1}^{\infty} \alpha_k w_{t-k} + (\varphi_t - 1) - \sum_{k=1}^{\infty} \alpha_k (\varphi_{t-k} - 1) + z_t,$$

which satisfies (10) once $\varphi_{t-k} = 1$ for all k . The result follows from setting $\varphi_t = \varphi_{t-k} = 1$ (so that $\hat{w}_t = \sum_{k=1}^{\infty} \alpha_k w_{t-k}$) and rearranging terms. QED

Proof of Proposition 4

Consider a relaxed version of problem (A5):

$$\begin{aligned} \min_{V_{(\Omega-N)-1}, \dots, V_{\Omega+N}} \quad & \sum_{\omega=\Omega-N}^{\Omega+N} P(\omega|\Omega) \cdot \frac{1}{\Delta V_{\omega}} \\ \text{subject to} \quad & \sum_{\omega=\Omega-N}^{\Omega+N} \Delta V_{\omega} \leq V_{\max} - V_{\min}, \end{aligned}$$

where the new constraint is implied by (but does not imply) the original constraint. In this relaxed problem, the effective choice variables are the slopes ΔV_{ω} (for $\omega = \Omega - N, \dots, \Omega + N$). Expressed in terms of these new choice variables ΔV_{ω} , the relaxed problem is strictly convex and can therefore be solved using a simple Lagrangian. In particular, if we let λ denote the multiplier for the constraint, the first-order condition for each value of ΔV_{ω} is

$$-\frac{P(\omega|\Omega)}{(\Delta V_{\omega})^2} + \lambda = 0.$$

When combined with the fact that the constraint must bind (i.e., $\sum_{\omega=\Omega-N}^{\Omega+N} \Delta V_{\omega} = V_{\max} - V_{\min}$), these first-order conditions deliver a unique solution:

$$\begin{aligned} \Delta V_{\omega} &= P(\omega|\Omega)^{1/2} \cdot \lambda^{-1/2} \\ &= P(\omega|\Omega)^{1/2} \cdot (V_{\max} - V_{\min}) \cdot \left[\sum_{\omega=\Omega-N}^{\Omega+N} P(\omega|\Omega)^{1/2} \right]^{-1}, \end{aligned}$$

where the second equality follows from solving for $\lambda^{-1/2}$ by adding across all values of ΔV_{ω} and using the binding constraint. Notice that these values for ΔV_{ω} are the ones stated in the proposition.

Consider now the original problem (A5). When we set $V_{\Omega+N} = V_{\max}$ and $V_{(\Omega-N)-1} = V_{\min}$ and use the same slopes ΔV_{ω} derived in the relaxed problem, the original constraint is satisfied, and therefore a solution to (A5) is achieved. Moreover, uniqueness of this solution follows from the fact that the solution of the relaxed problem is unique and, given these unique values for ΔV_{ω} , the only way to satisfy the original constraint is by setting $V_{\Omega+N} = V_{\max}$ and $V_{(\Omega-N)-1} = V_{\min}$. QED

References

- Binmore, Kenneth G. 1994. *Game Theory and the Social Contract*. Vol. 1. *Playing Fair*. Cambridge, MA: MIT Press.
- Burnham, Terry, and Jay Phelan. 2001. *Mean Genes: From Sex to Money to Food, Taming Our Primal Instincts*. New York: Penguin.
- Clark, Andrew E., and Andrew J. Oswald. 1996. "Satisfaction and Comparison Income." *J. Public Econ.* 61 (September): 359–81.
- Cole, Harold L., George J. Mailath, and Andrew Postlewaite. 1992. "Social Norms, Savings Behavior, and Growth." *J.P.E.* 100 (December): 1092–1125.
- Constantinides, George M. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *J.P.E.* 98 (June): 519–43.

- Damasio, Antonio S. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Putnam.
- Easterlin, Richard A. 1995. "Will Raising the Incomes of All Increase the Happiness of All?" *J. Econ. Behavior and Organization* 27 (June): 35–47.
- Frank, Robert H. 1985. *Choosing the Right Pond: Human Behavior and the Quest for Status*. New York: Oxford Univ. Press.
- Frederick, Shane, and George Loewenstein. 1999. "Hedonic Adaptation." In *Well-Being: The Foundations of Hedonic Psychology*, edited by Daniel Kahneman, Ed Diener, and Norbert Schwarz. New York: Sage Found.
- Frey, Bruno S., and Alois Stutzer. 2004. "Economic Consequences of Mispredicting Utility." Working Paper no. 218 (December), Inst. Empirical Res. Econ., Univ. Zurich.
- Gallistel, Charles R., Matthew Leon, Meg Waraczynski, and Michael S. Hanau. 1991. "Effect of Current on the Maximum Possible Reward." *Behavioral Neuroscience* 105 (December): 901–12.
- Gilbert, Daniel T., Elizabeth C. Pinel, Timothy D. Wilson, Stephen J. Blumberg, and Thalia P. Wheatley. 1998. "Immune Neglect: A Source of Durability Bias in Affective Forecasting." *J. Personality and Soc. Psychology* 75 (September): 617–38.
- Green, Jerry R., and Nancy L. Stokey. 1983. "A Comparison of Tournaments and Contracts." *J.P.E.* 91 (June): 349–64.
- Holmstrom, Bengt. 1979. "Moral Hazard and Observability." *Bell J. Econ.* 10 (Spring): 74–91.
- Kahneman, Daniel, Ed Diener, and Norbert Schwarz, eds. 1999. *Well-Being: The Foundations of Hedonic Psychology*. New York: Sage Found.
- Kandel, Eric R., James H. Schwartz, and Thomas M. Jessell. 2000. *Principles of Neural Science*. 4th ed. New York: McGraw-Hill.
- Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *J.P.E.* 89 (October): 841–64.
- Levin, Jonathan. 2003. "Relational Incentive Contracts." *A.E.R.* 93 (June): 835–57.
- Loewenstein, George, and David Schkade. 1999. "Wouldn't It Be Nice? Predicting Future Feelings." In *Well-Being: The Foundations of Hedonic Psychology*, edited by Daniel Kahneman, Ed Diener, and Norbert Schwarz. New York: Sage Found.
- Markowitz, Harry M. 1952. "The Utility of Wealth." *J.P.E.* 60 (April): 151–58.
- McMahon, Darrin M. 2004. "From the Happiness of Virtue to the Virtue of Happiness: 400 B.C.–A.D. 1780." *Daedalus* 133 (Spring): 5–17.
- Rayo, Luis, and Gary S. Becker. 2007. "Habits, Peers, and Happiness: An Evolutionary Perspective." *A.E.R. Papers and Proc.* 97 (May).
- Robson, Arthur J. 2001a. "The Biological Basis of Economic Behavior." *J. Econ. Literature* 39 (March): 11–33.
- . 2001b. "Why Would Nature Give Individuals Utility Functions?" *J.P.E.* 109 (August): 900–914.
- Rogerson, William P. 1985. "The First-Order Approach to Principal-Agent Problems." *Econometrica* 53 (November): 1357–67.
- Samuelson, Larry. 2004. "Information-Based Relative Consumption Effects." *Econometrica* 72 (January): 93–118.
- Samuelson, Larry, and Jeroen M. Swinkels. 2006. "Information, Evolution and Utility." *Theoretical Econ.* 1 (March): 119–42.
- Simmons, Janine M., and Charles R. Gallistel. 1994. "Saturation of Subjective

Reward Magnitude as a Function of Current and Pulse Frequency." *Behavioral Neuroscience* 108 (February): 151–60.

Stigler, George J., and Gary S. Becker. 1977. "De Gustibus Non Est Disputandum." *A.E.R.* 67 (March): 76–90.