

EVOLUTIONARY EFFICIENCY AND HAPPINESS

LUIS RAYO AND GARY S. BECKER

ABSTRACT. We model happiness as a measurement tool used to rank alternative actions. The quality of the measurement is enhanced by a happiness function that adapts to the available opportunities, a property favored by evolution. The optimal function is based on a time-varying reference point – or performance benchmark – that is updated over time in a statistically optimal way. Habits and peer comparisons arise as special cases of such 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.

1. Introduction

For long, utility was assumed to depend only on the absolute levels of our material outcomes. However, a large body of research now argues that utility, whether defined in terms of decision-making or hedonic experience, is sharply dependent on the difference between these outcomes and a time-varying reference point – examples include Markowitz [1952], Stigler and Becker [1977], Frank [1985], Constantinides [1990], Easterlin [1995], Clark and Oswald [1996], and Frederick and Loewenstein [1999]. Two pervasive phenomena in these lines are *habituation*, e.g., becoming accustomed to an expensive life-style or a physical handicap, and *peer comparisons*, e.g., caring about relative income. Both can be described by a reference point that is determined, respectively, by past outcomes and by the outcomes of peers. Moreover, these phenomena appear to be innate – they are present in young children, and have been documented in every known human culture, Brown [1999] – suggesting in turn that they served an evolutionary role in the descent of our species.¹

Date: First draft: 5/03. This draft: 4/05.

The University of Chicago. Acknowledgments to be added.

¹There is also direct neurological evidence confirming habituation to sustained stimuli. See, for example, Wright et al. [2001].

In this paper, we are interested in the hedonic dimension of utility, or happiness. We view happiness as a biological decision-making device that translates potential choices into a ranking criteria – e.g., Damasio [1994], Robson [2001a] – and seek to further uncover its evolutionary rationale. Specifically, we propose that habituation and peer comparisons can arise as special cases of the same general phenomenon, namely, a reference point that is updated over time in a statistically meaningful way. This updating enhances the individual’s ranking criteria. In addition, we argue that two related psychological phenomena can be derived from precisely the same biological role: a reference point that is influenced by the individual’s expectations of the world, and a happiness level that, although volatile, continuously reverts to its long-term mean.

The endogenous variable in our model is the individual’s happiness function. We assume that this function is implicitly designed by nature. Following the general framework of Binmore [1994], Robson [2001a], Samuelson [2004], and Samuelson and Swinkels [2004], we focus directly on the limiting outcome of evolution. This outcome is described as the solution to a metaphorical principal-agent problem, where the principal corresponds to the evolutionary process controlling the innate characteristics of the individual, or agent, who in turn serves the purpose of genetic replication. Crucially, when we speak of the evolutionary end-point, we refer to the ancestral hunter-gather environment and the suitable adaptations developed back then. Thus, our approach is relevant only to the extent that these ancestral adaptations are still present in our innate characteristics today.

The happiness function, in particular, will serve as a measurement tool used to compare alternative choices. Following Frederick and Loewenstein [1999], and Robson [2001b], our central assumption is an exogenous limit over the precision of the agent’s measurement, from which a role for adaptation is derived. For example, these authors argue that an adaptive utility is analogous to an eye that adjusts to the luminosity of the environment in order to increase its accuracy, or a voltmeter that delivers a more precise measurement when calibrated to the specific problem at hand.²

²Frederick and Loewenstein [1999] also argue that hedonic adaptation can serve a protective role against extreme emotional states. In section 6, we suggest a way in which these two approaches can be combined.

Here we take the analogy one step further. We consider an abstract choice setting where the agent must compare alternative inputs x towards the production of a random output y (consider, for example, a hunter-gatherer who is searching for fruit). Associated to each level of y , there is a real-valued hedonic utility, or happiness, $V(y)$. The agent measures the impact of x by means of the conditional expectation $E[V(y) | x]$, where $V(y)$ serves as a “lens” that can adjust to the environment. The quality of this measurement, however, is restricted by two constraints. The first is an upper and lower bound on the happiness function $V(y)$, which we interpret as a physical limit on the rewards that the brain can produce (e.g., due to its finite number of neurons). The second constraint describes a limit on the agent’s perception sensitivity: we assume that two alternatives x_1 and x_2 cannot be distinguished by the agent whenever the difference between $E[V(y) | x_1]$ and $E[V(y) | x_2]$ is smaller than some minimum threshold. When combined, these constraints provide the basis for an adaptive $V(y)$.

The fact that the agent’s choice is over inputs, while happiness depends on the level of output, leads to a parallel between the evolutionary problem and a standard problem of optimal incentives under moral hazard – where $V(y)$ corresponds to a performance reward for the agent. In both cases, the principal who designs $V(y)$ seeks to maximize the signal value of this function (due to measurement limits in the former case, and due to a cost of effort in the latter). This parallel is central to our approach, since it allows us to import a number of statistical insights from incentive theory to the study of happiness.

The optimal happiness functions we derive are one-step functions with a threshold corresponding to an endogenous reference point. (These functions can be interpreted as limiting versions of S -shaped curves – with a slope that is now entirely concentrated over the reference point.) Crucially, the reference point is positioned according to the current technological opportunities, and is updated over time so that it constantly matches the agent’s potential. In order to illustrate this adaptive mechanism, we consider an environment where the optimal reference point equals the conditional expectation of output based on all available information. In particular, this expectation exploits information contained in past levels of output and in the output of peers, from which habituation and peer comparisons are derived. Consequently, the specific functional form for these habits and peer comparisons reflects their underlying informative role.

A distinctive result is a generalized form of habituation. From a standard class of output technologies, we derive a reference point that is positively related to past levels of output as well as contemporaneous peer output (describing standard negative externalities), but negatively related to past values of peer output (which departs from usual formulations). This represents an individual who becomes habituated not only to her own success, but also to the success of her peers. For example, the negative effect of a permanent increase in peer output will fade away over time.

We begin our analysis with a static model where the basic evolutionary problem is presented. We then extend this model to a dynamic setup where adaptation is discussed.

2. The Static Model

Consider a representative agent (i.e., a hunter-gatherer) who faces an abstract one-shot project. To fix ideas, suppose 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, she 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 $y \in \mathbb{R}$ – the amount of fruit collected. Denote the conditional probability distribution of output by $f(y | x, s)$, a function known by the agent.

Beyond 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 included 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.

We are interested in the case where the agent's decision is guided by emotional rewards that are based on the realized level of output – as opposed to rewards that are based directly on x . This means that the agent is allocated full autonomy over the choice of x , rather than being prescribed a specific behavior. The idea is that because of an informational advantage that stems from her observation of s (representing a complex environment), the agent has comparative advantages when it comes to selecting the appropriate means of production – Binmore [1994, p.151]

and Robson [2001a] follow a similar approach.³ For simplicity, emotional rewards take the form of a one-dimensional level of happiness $V(y)$, which is experienced once y is realized. We assume that the agent can freely dispose of y , which allows us to focus without loss on non-decreasing happiness functions.

The agent measures the impact of alternative choices x via the expected value of happiness $E[V | x, s] \equiv \int V(y)f(y | x, s)dy$. Here the happiness function serves as a “lens” that transforms a functional space of probability distributions f into a single dimension, expected happiness, from which a decision criteria is obtained.⁴ Observe that this setup allows for a distinction between hedonic utility and decision utility. The former refers to the emotional experience, $V(y)$, whereas the latter refers to the standard notion of decision-theoretic utility (a ranking of alternative choices), given here by the state-contingent utility function $u(x, s) \equiv E[V | x, s]$.

The happiness function is implicitly designed by an evolutionary process, which we call the “principal.” When designing this function, the metaphorical objective for the principal 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 have a reproductive advantage. For concreteness, we assume that the principal seeks to maximize the expected value of y – which leads to the same results as maximizing the expected value of any other increasing function of y .⁵

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 where the principal directly selects a happiness function that maximizes her objective. (Recall that, in general, an evolutionary process where 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 coincides with a local maximum that is unique.)

³Samuelson and Swinkels [2004] study a model where part of this autonomy is subtracted in order to compensate for cognitive biases.

⁴See Damasio [1994] for a neurological foundation of emotions as a decision-making device, and Robson [2001a] for an evolutionary foundation of expected utility.

⁵The technologies we consider below fit the assumptions under which Robson [1996] shows that expected-value criteria are optimal.

2.1. Measurement Imperfections. Our theory is based on two constraints that, when combined, limit the precision of the agent’s ranking of x . The first constraint is a physical bound that limits the highest and lowest value that V can take (e.g., due to a finite brain). In particular, we assume that $V(y) \in [\underline{V}, \bar{V}]$ for all y , for some finite bounds $\underline{V}, \bar{V} \in \mathbb{R}$.

In the applications below, these bounds will bind. As a result, the principal would potentially benefit from designing an organism capable of more extreme rewards. This is fully consistent with our analysis provided that expanding these bounds also conveys a cost (e.g., additional resources devoted to augmenting the reward centers of the brain). In particular, in the presence of such a cost, the principal would benefit from using the currently available bounds in the most efficient way, since this reduces the pressure to further expand them. This is precisely the problem we solve.⁶

Provided the bounds \underline{V}, \bar{V} are finite, our results are not affected by their specific size. Thus, to simplify notation, we normalize these bounds to 0, 1.

The second constraint is a limit on the agent’s ability to measure, or perceive, small differences in her objective $E[V \mid x, s]$. We represent this imperfection using a reduced form that allows for a simple analysis. Given any pair of choices x_1 and x_2 , we assume that there is a minimum threshold $\varepsilon > 0$ such that whenever $|E[V \mid x_1, s] - E[V \mid x_2, s]| \leq \varepsilon$, these two choices cannot be ranked. Accordingly, all choices that deliver an expected reward within ε distance of the optimized value $\max_x E[V \mid x, s]$ are part of the same indifference set, denoted the “satisficing” set. We assume that the agent’s choice is randomly drawn from this set. For our purposes, it suffices to assume that her draw is monotonic in the sense that the probability assigned to any subset of the satisficing set is inversely proportional to the size of this satisficing set.⁷

⁶This argument can also be stated without using the principal-agent metaphor. Begin with a given organism who has arbitrary bounds \underline{V}, \bar{V} that are finite. Her descendants can evolve, in principle, along two different paths: (1) they may develop a brain capable of more extreme bounds, and/or (2) they may exploit the existing range of rewards $[\underline{V}, \bar{V}]$ in a more efficient way (by improving the mapping between stimuli and rewards). If, beyond a certain point, more extreme bounds are costly to grow, pressure will arise to evolve along the second path. By taking the range $[\underline{V}, \bar{V}]$ as given, we effectively study such a path.

⁷Along similar lines, Simon [1959, p.261] argues that when the utilities of two alternatives are only slightly different, the subject is likely to vacillate in his choice. As an empirical precedent, he reports an experiment where subjects are asked to rank two unequal weights: when these weights

This assumption imposes a coarseness in the agent's measurement analogous to a computer that must round any small difference to zero because it uses only finitely many digits, or an eye that cannot rank the luminosity of two sources when they are sufficiently similar. In both examples, because of the imperfection, adapting to the problem at hand improves the measurement quality: an eye uses a pupil that opens in the dark in order to maximize the relative differences in luminosity, while a computer uses a decimal point that floats. The function V will play a very similar role.

As shown below, the principal would benefit from a smaller ε . However, reducing ε is likely to require additional energy as well – e.g., in the form of repeated measurements, or more intensive optimization techniques. In this case, the principal also benefits from a happiness function that best aids an imperfect machinery, whatever the size of the underlying imperfection happens to be.

We approach the principal's problem by first solving for the optimal happiness function for any given small ε , and then characterizing the limit as $\varepsilon \rightarrow 0$. This limit serves as an analytically convenient representation of an environment where small imperfections remain.

2.2. Output Technology. We focus on output technologies of the form

$$y = E[y \mid x, s] + z,$$

where z is an exogenous shock drawn from a continuous density function that has full support, and is strictly monotonic on either side of its mean – such as a normal. The shock is realized after x is selected. We assume that $E[y \mid x, s]$ is continuous in x , while X (the choice space) is a compact subset of \mathbb{R}^N , which guarantee that the choice space is sufficiently rich, and an optimal action always exists.

The above specification implies that the conditional density $f(y \mid x, s)$ is single-peaked at its mean $E[y \mid x, s]$, and is single-crossing in x : for all $x_1 \neq x_2$, $f(y \mid x_1, s)$ and $f(y \mid x_2, s)$ intersect for only one value of y (in this case, between their two means). In addition, $f(y \mid x, s)$ is ordered across x according to first-order stochastic dominance. These are the key distributional properties we employ.

approach each other, the frequency of a correct answer approaches 1/2. A related type of imperfect optimization is used for ε -equilibria in games – Radner [1980].

Under such technologies, the agent's choice problem can be simplified to an equivalent problem where her choice is one-dimensional. To do this, we define a real-valued index $\varphi(x, s)$ as follows:

$$\varphi(x, s) \equiv \frac{E[y | x, s] - \min_x E[y | x, s]}{\max_x E[y | x, s] - \min_x E[y | x, s]}.$$

For any given s , $\varphi(x, s)$ ranges from 0 to 1 as a function of x . We refer to φ as the “efficiency” of the agent's decision – for example, $\varphi(x, s) = 1$ means that, given s , the agent selected the optimal action x . Using this index, we can express $E[y | x, s]$ as a function of φ and s alone:

$$\begin{aligned} E[y | x, s] &= \varphi \max_x E[y | x, s] + (1 - \varphi) \min_x E[y | x, s] \\ &\equiv E[y | \varphi, s], \end{aligned}$$

where $E[y | \varphi, s]$ is increasing and continuous as a function of φ . In addition, output can be expressed as

$$y = E[y | \varphi, s] + z. \tag{1}$$

This formulation allows us to represent the agent's problem, without loss of generality, as one where she directly selects the level of φ , subject to $\varphi \in [0, 1]$, while x is sent to the background. To be sure, this simplification is for analytical purposes only: the existence of an underlying complex problem – where the agent actually compares values of x , not values of φ – remains important for the interpretation of the model.

2.3. The Optimal Happiness Function. We begin our analysis with the simplest case where $E[y | \varphi, s]$ is independent of the state s , which means that both $\max_x E[y | x, s]$ and $\min_x E[y | x, s]$ are independent of s . In this case, while s might affect the value of each particular choice x , it does not change the agent's overall output potential. As a result, the conditional density of y simplifies to $f(y | \varphi)$. We return to the general case in the dynamic model below.

Expressed in terms of $\varphi \in [0, 1]$, the agent's objective function is given by

$$E[V | \varphi] \equiv \int V(y) f(y | \varphi) dy.$$

Accordingly, her satisficing set corresponds to the set of choices φ that deliver an expected happiness within ε distance of $\max_{\varphi} E[V | \varphi]$:

$$\left\{ \varphi : E[V | \varphi] \geq \max_{\varphi} E[V | \varphi] - \varepsilon \right\}.$$

Under the technologies in (1), the densities $f(y \mid \varphi)$ have full support and are ordered across φ according to first-order stochastic dominance. Thus, for any V that is not constant (and non-decreasing), $E[V \mid \varphi]$ is increasing in φ , and therefore maximized at $\varphi = 1$.⁸

Consequently, the satisficing set becomes $\{\varphi : E[V \mid \varphi] \geq E[V \mid 1] - \varepsilon\}$. This set corresponds to an interval $\varphi \in [\varphi_{\min}(V, \varepsilon), 1]$, where the lower boundary $\varphi_{\min}(V, \varepsilon)$ is uniquely determined by the equality

$$E[V \mid \varphi_{\min}(V, \varepsilon)] = E[V \mid 1] - \varepsilon.$$

In particular, $\varphi_{\min}(V, \varepsilon)$ represents the lowest efficiency level φ that can arise in equilibrium.

From the principal's standpoint, the impact of V is fully summarized by the value of $\varphi_{\min}(V, \varepsilon)$, with a larger value being strictly preferred. Her problem can therefore be expressed as

$$\begin{aligned} \max_V \quad & \varphi_{\min}(V, \varepsilon) & (I) \\ \text{s.t.} \quad & V(y) \in [0, 1] \text{ for all } y, \end{aligned}$$

which corresponds to minimizing the set of inefficient choices $\varphi < 1$ that the agent confuses with $\varphi = 1$. Let φ^* denote the optimized value for this problem – which for any $\varepsilon > 0$, is smaller than 1.⁹

The following Lemma allows us to solve for the optimal V using a dual approach:

⁸From this statement we learn that, absent the ε error on the agent's measurement, *any* non-decreasing and non-constant V would be optimal – as they all deliver the same optimal preference ordering $E[V \mid \varphi]$, and induce the agent to select precisely $\varphi = 1$. In contrast, once the ε error is imposed, the cardinal properties of $E[V \mid \varphi]$, and therefore of V , play a crucial role.

⁹From this problem I , we learn the role of the bounds over V . If no bounds were imposed, the principal could attain an objective $\varphi_{\min}(V, \varepsilon)$ arbitrarily close to 1 by selecting a happiness function that is arbitrarily steep – effectively eliminating the errors caused by the agent's limited sensitivity. For example, set $V(y) = \lambda \cdot y$ for some constant $\lambda > 0$. Under this function, the expectation $E[V \mid \varphi]$ is given by $\lambda \cdot E[y \mid \varphi]$, and therefore φ_{\min} solves

$$E[y \mid \varphi_{\min}] = E[y \mid 1] - \frac{\varepsilon}{\lambda}.$$

In this case, a larger λ is equivalent to a smaller error ε . Thus, φ_{\min} approaches to 1 as λ approaches infinity. As a result, the principal would seek a λ that is arbitrarily large. This leads, in turn, to an arbitrarily steep $V(y)$ that exceeds any fixed bounds (no matter how large) for all $y \neq 0$. In contrast, once the bounds are imposed, the principal must choose V carefully so that it has a large slope only where it matters the most.

Lemma 1. *Suppose V^* is a solution to problem I (namely, $\varphi_{\min}(V^*, \varepsilon) = \varphi^*$). Then, V^* must also solve*

$$\begin{aligned} \max_V E[V | 1] - E[V | \varphi^*] & \quad (II) \\ \text{s.t. } V(y) \in [0, 1] & \text{ for all } y. \end{aligned}$$

Proof. Suppose not. Then there must exist a $V \neq V^*$ (satisfying the constraint) such that $E[V | 1] - E[V | \varphi^*] > E[V^* | 1] - E[V^* | \varphi^*] \equiv \varepsilon$. But this implies that $\varphi_{\min}(V, \varepsilon) > \varphi^*$, a contradiction. \square

In other words, in order for V^* to be optimal, there cannot exist an alternative V that leads to a difference between $E[V | 1]$ and $E[V | \varphi_{\min}(V^*, \varepsilon)]$ larger than ε , since this would deliver a boundary $\varphi_{\min}(V, \varepsilon)$ larger than $\varphi_{\min}(V^*, \varepsilon)$.


Proposition 1. *Problem I is solved by a one-step happiness function V^* such that*

$$V^*(y) = \begin{cases} 1 & \text{for all } y \geq \hat{y}, \\ 0 & \text{for all } y < \hat{y}, \end{cases}$$

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 measure-zero subset.

Proof. The objective in problem II is equal to $\int V(y) [f(y | 1) - f(y | \varphi^*)] dy$. This integral is maximized by setting $V(y) = 1$ for every y such that $f(y | 1) \geq f(y | \varphi^*)$, and $V(y) = 0$ 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 zero-measure subset, and therefore solves problem I as well. \square

This result can be illustrated graphically. The upper panel of Figure 1 graphs the conditional density $f(y | \varphi)$. The bold curve represents $f(y | 1)$, the most desirable function for the principal, while the dashed curve represents $f(y | \varphi^*)$ – where φ^* , by definition, always belongs to the satisficing set. The dual objective is to maximize the difference in expected happiness under these two alternatives, which is the only way to exclude every $\varphi < \varphi^*$ from the satisficing set. As depicted in the lower panel, this is achieved by a V that maximally rewards all values of y for which $f(y | 1) > f(y | \varphi^*)$, and vice versa: under any other V , the two distributions would appear more similar for the agent. The threshold \hat{y} lies where the two densities intersect. Under the technologies in (1), this occurs between the peaks $E[y | \varphi^*]$ and $E[y | 1]$.

FIGURE 1. The Optimal V

A statistical parallel can be drawn with a problem of optimal incentives – e.g., Holmstrom [1979], Levin [2003]. Interpret V as a performance bonus, φ as a costly 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 “ $\varphi = \varphi^*$.” In this case, the null is rejected whenever the likelihood ratio $\frac{f(y|\varphi^*)}{f(y|1)}$ exceeds one.

2.4. The Limit When $\varepsilon \rightarrow 0$. As ε converges to zero, the lower bound of the satisficing set φ^* converges to one. Consequently, as suggested by Figure 1, the happiness threshold \hat{y} – which lies between $E[y | \varphi^*]$ and $E[y | 1]$ – converges to $E[y | 1]$. Proposition 2 follows as a result:

Proposition 2. *Given ε , let $V^*(\varepsilon)$ denote the optimal one-step function characterized by Proposition 1. Then, in the limit as ε converges to zero, $V^*(\varepsilon)$ converges*

point-wise to the one-step function

$$V(y) = \begin{cases} 1 & \text{for all } y \geq E[y \mid 1], \\ 0 & \text{for all } y < E[y \mid 1]. \end{cases}$$

Up to a measure-zero subset, this limiting function uniquely maximizes the derivative $\frac{\partial}{\partial \varphi} E[V \mid \varphi]|_{\varphi=1} \equiv \int V(y) f_{\varphi}(y \mid \varphi)|_{\varphi=1} dy$, which represents the limiting version of the objective in the dual problem *II*. Maximizing this derivative guarantees that marginal deviations away from $\varphi = 1$ have a maximal impact over the agent's objective, thus improving her ability to discriminate. The analogy in optimal incentives is a first-order approach where all incentive power is focused over small effort deviations (e.g., Rogerson [1985], Levin [2003]).

These extreme rewards arise because the agent, by construction, is risk-neutral with respect to happiness, and happiness is costless to the principal within the bounds.¹⁰ Smoother functions would arise, for example, if we assumed that there is a neutral level of happiness that is physiologically most desirable, and deviations from this level are costly. The Appendix studies this case. There, we present a parameterization that delivers a family of smooth *S*-shaped curves with differing slopes around the reference point. In the limit, as these slopes converge to infinity, the curves converge to the one-step functions above. The analytical advantage of the limiting functions is that they are fully characterized by the position of \hat{y} .

3. The Dynamic Model

We now extend the model to a dynamic setup where the agent lives for multiple periods $t = 1, 2, \dots$. We equate every period with one separate project – the simplest possible case. 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 $y_t = E[y_t \mid x_t, s_t] + z_t$, which we assume satisfies the same properties as above, with z_t i.i.d. across time. As before, we use the one-dimensional representation in (1), where $y_t = E[y_t \mid \varphi_t, s_t] + z_t$, and the

¹⁰Also responsible for the extreme rewards is the fact that the principal and agent have symmetric information concerning the output technology. This allows the principal to precisely target incentives where they matter the most. If, on the contrary, the agent had an information advantage over the principal (e.g., in the form of fine environmental knowledge that cannot be encoded in V), the principal would no longer concentrate the entire slope of V over the reference point. In particular, by spreading this slope more evenly on either side, she provides additional incentives for an agent who happens to know, *a priori*, that the specified reference point is either too difficult or too easy to exceed. A formal illustration is available from the authors.

agent directly selects the efficiency level $\varphi_t \in [0, 1]$. After y_t is realized, the agent experiences a happiness level $V_t \in [0, 1]$.

In contrast to the special case studied in the static model, we now allow the extreme values $\max_{x_t} E[y_t | x_t, s_t]$ and $\min_{x_t} E[y_t | x_t, s_t]$ to vary with s_t , so that $E[y_t | \varphi_t, s_t]$ also varies with s_t . In particular, we assume that these extreme values now depend on a subset of s_t , denoted Ω_t . Accordingly, output becomes

$$y_t = E[y_t | \varphi_t, \Omega_t] + z_t. \quad (2)$$

The new variable Ω_t may represent, for example, current weather conditions. We assume that Ω_t can be encoded, together with y_t , into the happiness function (which means that Ω_t must be relatively simple). As a result, happiness becomes $V_t(y_t, \Omega_t)$. This dependence on Ω_t implies that happiness can now adapt to changes in output potential.

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. In other words, everything the agent cares about, present and future, is reflected in present emotions. This model captures forward-looking behavior by interpreting a given project as being forward looking itself, and rewarded by current happiness. Consider, for example, a hunter-gatherer who eats in excess of current needs in order to accumulate fat, or helps a friend, precisely because it makes her feel happy today (a modern counterpart might be an individual who invests in her retirement funds for precisely the same reason). Below, we also discuss the case where the agent internalizes future happiness above and beyond V_t .

Since the agent faces a separate project every period, any such period is identical to the static model above – save for the presence of Ω_t . Moreover, Ω_t simply enters as a technological constant in all the previous analysis. As a result, from Proposition 2, the optimal limiting function V_t (as $\varepsilon \rightarrow 0$) is a one-step function with threshold $\hat{y}_t = E[y_t | \varphi_t, \Omega_t]_{\varphi_t=1}$ – namely, the peak of the density $f(y_t | \varphi_t, \Omega_t)_{\varphi_t=1}$. In other words, the impact of Ω_t occurs via a threshold \hat{y}_t that is updated in a statistically optimal way.

3.1. Habituation. We proceed with two simple examples where \hat{y}_t incorporates a habit due to an Ω_t that is correlated across time. Possible causes for this correlation include environmental shocks and an intrinsic talent that persists over time.

Example 1: A Markovian Habit. *Suppose 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 (2) 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:*

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

In this case, the agent experiences a high level of happiness if and only if her current output exceeds what she achieved one period ago. This result follows from an optimal statistical inference. In order to best guide the agent, the principal employs her most accurate source of information regarding φ_t (analogous to an optimal incentive scheme). 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, all decisions φ_t would appear increasingly good or increasingly bad – and thus increasingly similar – as soon as θ_t drifted to extreme values. Thus, the agent would lose ability to discriminate.

On the other hand, observe that a reduction in φ_t affects 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 . Consequently, the principal must punish the agent following low realizations of z_t , and vice versa: happiness is inevitably affected by chance.

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

Indeed, in many languages, the word “happiness” is closely linked to “fortune” and “luck.” For the ancient Greeks, 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, p.7).

Example 2: Auto-Regressive Habits. Suppose output is given by $y_t = \varphi_t + \theta_t$, and θ_t follows the auto-regressive process $\theta_t = \sum \alpha_s \theta_{t-s} + z_t$ for arbitrary constants α_s , and $s \geq 1$. In this case, Ω_t is the vector $(\theta_{t-1}, \theta_{t-2}, \dots)$. Following similar steps to those in Example 1, output becomes $y_t = \varphi_t + \sum \alpha_s (y_{t-s} - \varphi_{t-s}) + z_t$. In equilibrium, this equation reduces to $y_t = \sum \alpha_s y_{t-s} + (1 - \sum \alpha_s) \varphi_t + z_t$, from which we obtain

$$\hat{y}_t = E[y_t \mid \varphi_t, \Omega_t] \Big|_{\varphi_t=1} = \sum \alpha_s y_{t-s} + \left(1 - \sum \alpha_s\right) \varphi_t.$$

The reference point is now a weighted average between past levels of output and the equilibrium efficiency level $\varphi_t = 1$ (e.g., allowing habituation to occur at a slower rate than in Example 1). The specific weights guarantee that the carrier of happiness $y_t - \hat{y}_t$ employs only the new information contained in y_t .

3.2. Habits and Forward-Looking Behavior. Notice that the presence of habits does not deter the agent from selecting $\varphi_t = 1$. 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 finding in the psychology literature that individuals tend to underestimate the degree to which they adapt to changing circumstances – Gilbert et al. [1998], Loewenstein and Schkade [1999].¹¹

This opens the possibility that a rational 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 will depend 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. In one extreme, \hat{y}_t may correspond to an exogenous function of past levels of output, namely, $\hat{y}_t = \sum \alpha_s y_{t-s} + (1 - \sum \alpha_s) \varphi_t$. In the other extreme, \hat{y}_t may equal the best predictor of y_t conditional on $\varphi_t = 1$ and all past information, namely, $\hat{y}_t = E[y_t \mid \varphi_t, \Omega_t] \Big|_{\varphi_t=1} = \sum \alpha_s \theta_{t-s} + 1$ (where the values of θ_{t-s} are inferred from the technological equalities $\theta_{t-s} = y_{t-s} - \varphi_{t-s}$).¹² When $\varphi_{t-s} = 1$ for all s , both formulations coincide. The difference arises outside equilibrium:

¹¹See also Frey and Stutzer [2004] for a study of biased predictions, and Burnham and Phelan [2001] for a witty account.

¹²In this case, the agent uses her best cognitive abilities to form an expectation, only to then compare her actual success against this self-imposed benchmark. Such a procedure would add flexibility to the reference point, which can be evolutionarily advantageous in an environment where the parameters of the technological process change over time.

In the former case, a reduction in φ_t reduces future reference points by reducing y_t , thus increasing the expected value of future happiness. As a result, the deviation might indeed be beneficial.¹³ In the latter case, in contrast, the agent understands that a reduction in φ_t does not affect her future reference points because these only depend on the underlying technological shock θ_t , and not on the particular value of y_t . This case describes an agent who cannot change his future output expectations by merely reducing φ_t . Therefore, a deviation is no longer beneficial. Such an agent can be described as either maximizing a present discounted value of future happiness levels, or equivalently, as maximizing current happiness alone.

The above distinction is also relevant for policy. In the former case, for instance, a tax policy that leads to an increasing income profile over the life cycle might increase long-term happiness. In the latter case, in contrast, expectations would fully adjust to the policy, therefore eliminating the desired effect.

4. Multiple Agents

In order to derive peer effects, we extend the model to include multiple agents i . Actions for period t , denoted φ_t^i , are selected simultaneously, and they 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. The new assumption is that the agents experience common productivity shocks (e.g., due to a shared environment), implying that peer output becomes valuable when assessing individual performance.¹⁴ Dropping the i superscript, we focus on technologies such that

$$w_t = E[w_t \mid \varphi_t, \Omega_t] + z_t, \quad (3)$$

which we assume satisfies 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.

¹³For 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 $\sum \beta^s \alpha_s > 1$.

¹⁴If the principal directly benefitted from relative output w_t^i , peer effects would immediately arise –see, for example, Cole et al. [1992] for the potential benefits conveyed by w_t^i . Here we show how these effects can extend beyond any direct advantage of achieving a high w_t^i . In a complementary approach, Samuelson [2004] derives peer effects that lead to an imitation of consumption levels (even when only absolute consumption is relevant), which may be desirable in an environment where the optimal level of consumption is not fully known by the agent.

Although the principal cares only about y_t , happiness is also allowed to depend on \bar{y}_t (as well as Ω_t). From (3), the conditional density $f(y_t | \varphi_t, \bar{y}_t, \Omega_t)$ depends on y_t and \bar{y}_t only through w_t . As a result, happiness can be expressed without loss as $V_t(w_t, \Omega_t)$. It follows that this model is identical to the model with a single agent, with w_t replacing y_t . Consequently, from Proposition 2, the optimal V_t is a one-step function with $V_t(w_t, \Omega_t) = 1$ for all $w_t \geq \hat{w}_t$, and $V_t(w_t, \Omega_t) = 0$ otherwise – where $\hat{w}_t = E[w_t | \varphi_t, \Omega_t]_{\varphi_t=1}$.

Example 3: Static Peer Comparisons. *Suppose 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 (3). Both Γ_t and z_t are realized after φ_t is selected. No restrictions over the distribution of Γ_t are imposed. In this case, Ω_t will be redundant. 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 (3). The optimal reference point for w_t becomes $\hat{w}_t = E[w_t | \varphi_t]_{\varphi_t=1} = 0$. Thus, 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 income $y_t - \bar{y}_t$. The reason why \bar{y}_t enters the happiness function is because it filters out the aggregate shock Γ_t , and thus 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, its effect is to magnify the cost of withdrawing effort – e.g., Lazear and Rosen [1981], Green and Stokey [1983].

A distinctive implication of the model arises when habits and peer comparisons are combined. We begin with a simple example that combines the technologies from Examples 1 and 3:

Example 4: A Markovian Habit and Dynamic Peer Comparisons. *Suppose 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 Markovian process $\theta_t = \theta_{t-1} + z_t$. The difference with Example 1 is the presence of Γ_t , and the difference with Example 3 is the persistence of z_t . 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 (3) (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 . Such a reference point would imply that the carrier of happiness 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 would lead to a comparison between an innovation and an absolute level. Rather, the carrier of happiness is the difference in differences

$$\Delta y_t - \Delta \bar{y}_t,$$

where the role of \bar{y}_{t-1} is to filter out the lagged aggregate shock Γ_{t-1} .

Thus, while y_{t-1} and \bar{y}_t reduce current happiness, \bar{y}_{t-1} has the opposite effect. We interpret this 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 , while holding y_t constant. This increase initially shifts the reference point to the right, with a likely decrease in happiness. But after one period, \bar{y}_t enters the reference point with a negative sign, shifting it back to its original level. As a result, the agent has successfully coped. Equivalently, this agent has become habituated to her new lower social position w_t , even though her own income has not changed.¹⁵

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

Proposition 3. *Suppose 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 auto-regression $\theta_t = \sum \alpha_s \theta_{t-s} + z_t$ for arbitrary constants α_s , and z_t i.i.d. Then, the optimal reference point for period t is given by*

$$\hat{y}_t = \sum \alpha_s y_{t-s} + \bar{y}_t - \sum \alpha_s \bar{y}_{t-s}.$$

Proof. Using the above technology, we can write

$$y_t - \sum \alpha_s y_{t-s} = \varphi_t - \sum \alpha_s \varphi_{t-s} + \Gamma_t - \sum \alpha_s \Gamma_{t-s} + z_t.$$

¹⁵This may potentially describe an agent who copes with peer success by changing her reference group. To the best of our knowledge, a satisfactory formal model of reference-group formation is yet to be developed.

Therefore, in equilibrium, $\bar{y}_t - \sum \alpha_s \bar{y}_{t-s} = (1 - \sum \alpha_s) + \Gamma_t - \sum \alpha_s \Gamma_{t-s}$. Combining these two expressions, we obtain

$$w_t = \sum \alpha_s w_{t-s} + (\varphi_t - 1) - \sum \alpha_s (\varphi_{t-s} - 1) + z_t,$$

which satisfies (3). The result follows from setting $\varphi_t = \varphi_{t-s} = 1$ (so that $\hat{w}_t = \sum \alpha_s w_{t-s}$), and rearranging terms. \square

The carrier of happiness is now the generalized difference in differences

$$y_t - \sum \alpha_s y_{t-s} - \left[\bar{y}_t - \sum \alpha_s \bar{y}_{t-s} \right].$$

The term $\sum \alpha_s y_{t-s}$ corresponds to a conventional habit, whereas the presence of $\sum \alpha_s \bar{y}_{t-s}$ again results in habituation to peers. Regardless of the properties of the aggregate shocks (including any intertemporal correlations), the same coefficients enter both forms of habituation. The reason is that y_{t-s} is impacted by the aggregate shock Γ_{t-s} , which is redundant when assessing φ_t . Subtracting \bar{y}_{t-s} from y_{t-s} 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 feel better when, for as long as he can remember, this neighbor has always been wealthier – as opposed to the case where their relative fortunes have been recently reversed.

4.1. Income and Happiness Surveys. A reference point that depends negatively on the lagged income of peers might also be useful when describing the happiness surveys. To illustrate this, suppose we restrict to linear reference points of the form

$$\hat{y}_t = \sum \alpha_s y_{t-s} + \lambda \bar{y}_t + \sum \beta_s \bar{y}_{t-s},$$

where y_t denotes income, perhaps measured in logs, and where the sum of the right-hand coefficients does not exceed one – a condition required if a general increase in income is to have a non-negative impact over happiness. Such a model would explain habituation to own income via positive coefficients α_s (for example, full habituation to permanent changes in y_t would require $\sum \alpha_s = 1$), and it would explain a concern for relative income via a positive λ .

A number of authors have argued that beyond the point where basic needs have been covered, an individual's absolute level of income has a minimal impact over her reported happiness. Rather, the bulk of its impact comes from the relative social position it conveys – e.g., Easterlin [1995], Oswald [1997], Frank [2004]. In

the above model, this would correspond to a λ close to 1 – so that simultaneous increases in y_t and \bar{y}_t mostly cancel each other out.¹⁶ But when combined with the initial restriction that $\sum \alpha_s + \lambda + \sum \beta_s \leq 1$, notice that a λ that is close to 1 may coexist with significant habits for own income only when the term $\sum \beta_s$ happens to be negative – as suggested by the Proposition above.

5. Conclusions

We have modeled happiness as a measurement instrument that guides the agent’s decisions. The quality of the measurement is enhanced when the happiness function adapts to the current decision environment – adaptation is thus favored by nature. In the model, this adaptation occurs through an output benchmark, or reference point, that integrates all information that can predict the agent’s performance. Whenever output is correlated across time and across agents, habits and peer comparisons arise as special cases.

Our goal has been to argue that happiness, as best observed by current empirical methods, contains the signs of statistical inference. In particular, we have suggested a statistical parallel between happiness and an optimal incentive scheme that seeks to promote effort, allowing us to import the insights of agency theory to the study of our innate psychological features. This approach rationalizes, for example, why serendipity has an impact over happiness – and why this impact is short-lived.

The examples of habits and peer comparisons that we considered are far from exhaustive. Possible extensions could address the issue of habituation patterns that differ systematically according to the type of good involved, as well as the endogenous formation of reference groups. In both cases, statistical principles may prove to play a role.

6. Appendix

Here we explore a possible variation of the model. In particular, we consider the case where there is an intermediate level of happiness V_0 that is optimal from a physiological point of view, and that deviations away from V_0 are, *ceteris paribus*, undesirable.¹⁷ Specifically, suppose that whenever $V(y)$ departs from V_0 , it causes a

¹⁶Easterlin [1995], for example, implicitly favors a $\lambda = 1$. Specifically, this extreme version is necessary to eliminate any time-series association between income and happiness when the growth rate for income varies with time.

¹⁷This setup follows a notion in the psychology literature that happiness is like blood sugar: deviations from an intermediate level are harmful for the organism, e.g., Wilson et al. [2002]. See




FIGURE 2. Costly Happiness

cost $C(V(y), V_0) \in \mathbb{R}$ to the principal (via the negative effect of this departure over the agent). Because of this cost, the bounds \underline{V}, \bar{V} may no longer bind.

Assuming an additive structure for this cost, the objective for the principal becomes

$$\max_V \left. \frac{\partial}{\partial \varphi} E[V | \varphi] \right|_{\varphi=1} - E[C | V],$$

where the first term is the limiting version of the objective in *II*, and the second term is the expected value of C in equilibrium. The principal must now trade off the signal value of a departure from V_0 against its associated cost.

We illustrate this trade-off for the case where $C(V(y), V_0)$ takes the simple form $\frac{1}{n} |V(y) - V_0|^n$, and output is given by $y = \varphi + z$, with z distributed as a standard normal. This specification allows for a solution in closed form that is parameterized by the degree of curvature n in the cost function. Ignoring the bounds \underline{V}, \bar{V} , the

also Sapolsky [1999] and Frederick and Loewenstein [1999] for a discussion of costs associated to extreme emotional states.

optimal values for $V(y)$ satisfy

$$\begin{aligned} (V(y) - V_0)^{n-1} &= \frac{f_\varphi(y | 1)}{f(y | 1)} = y - \hat{y}, \text{ for } V(y) \geq V_0, \text{ and} \\ (V_0 - V(y))^{n-1} &= -\frac{f_\varphi(y | 1)}{f(y | 1)} = \hat{y} - y, \text{ for } V(y) < V_0, \end{aligned}$$

where the equalities on the left correspond to the first-order conditions for the objective above,¹⁸ and the equalities on the right are derived from the normal density $f(y | \varphi) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}(y - \varphi)^2)$ together with the equality $\hat{y} = E[V | 1]$. The left-hand sides of these equations captures the marginal cost created by $V(y)$ (according to whether $V(y)$ is above or below V_0), while the likelihood ratios $\frac{f_\varphi(y|1)}{f(y|1)}$ capture the marginal benefit of $V(y)$ over the agent's incentives.

Figure 2 (drawn to scale) plots $V(y) - V_0$ as a function of $y - \hat{y}$, the carrier of happiness. When $n = 2$, absent any bounds, the optimal happiness function would be a 45° line. Whenever $n > 2$, the optimal function becomes S -shaped: concave to the right of the reference point, and convex to the left. Moreover, as n becomes large, the optimal function converges to a one-step function with $V(y) - V_0 = 1$ to the right of \hat{y} , and $V(y) - V_0 = -1$ to the left. After a normalization of units, this limiting case coincides with the functions derived in the text.

REFERENCES

- [1] Binmore, K. [1994], *Game Theory and the Social Contract. Volume 1: Playing Fair*, Cambridge: MIT Press.
- [2] Brown, D.E. [1999], "Human Universals," in Wilson, R.A., and F.C. Keil, *The MIT Encyclopedia of the Cognitive Sciences*, Cambridge: MIT Press.
- [3] Burnham, T., and Phelan, J. [2001], *Mean Genes: From Sex to Money to Food, Taming our Primal Instincts*, New York: Penguin Books.
- [4] Clark, A.E., and A.J. Oswald [1996], "Satisfaction and Comparison Income," *Journal of Public Economics*, 61, pp. 359-81.

¹⁸To derive these equalities, express the objective function as

$$\frac{\partial}{\partial \varphi} E[V | \varphi] \Big|_{\varphi=1} - E[C | V] = \int V(y) f_\varphi(y | 1) dy - \int \frac{1}{n} |V(y) - V_0|^n f(y | 1) dy.$$

The corresponding first-order conditions for $V(y)$ are

$$\begin{aligned} f_\varphi(y | 1) - (V(y) - V_0)^{n-1} f(y | 1) &= 0 \text{ for } V(y) \geq V_0, \text{ and} \\ f_\varphi(y | 1) + (V_0 - V(y))^{n-1} f(y | 1) &= 0 \text{ for } V(y) < V_0. \end{aligned}$$

The desired equalities follow from dividing through by $f(y | 1)$ and rearranging terms.

- [5] Cole, H.L., G.J. Mailath, and A. Postlewaite [1992], "Social Norms, Savings Behavior, and Growth," *Journal of Political Economy*, 100, pp. 1092-125.
- [6] Constantinides, G.M. [1990], "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98, pp. 519-43.
- [7] Damasio, A. [1994], *Descartes' Error: Emotion, Reason, and the Human Brain*, New York: Putnam.
- [8] Easterlin, R.A. [1995], "Will Raising the Incomes of All Increase the Happiness of All?," *Journal of Economic Behavior and Organization*, 27, pp. 35-47.
- [9] Frank, R.H. [1985], *Choosing the Right Pond: Human Behavior and the Quest for Status*, Oxford University Press.
- [10] Frank, R.H. [2004], "Are Positional Externalities Different from Other Externalities?," mimeo, Cornell University.
- [11] Frederick, S., and G. Loewenstein [1999], "Hedonic Adaptation," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 302-29.
- [12] Frey, B.S., and A. Stutzer [2004], "Economic Consequences of Mispredicting Utility," mimeo, University of Zurich.
- [13] Green, J.R., and N.L. Stokey [1983], "A Comparison of Tournaments and Contracts," *Journal of Political Economy*, 91, pp. 349-64.
- [14] Gilbert, D.T., T.D. Wilson, E.C. Pinel, S.J. Blumberg, and T.P. Wheatley [1998], "Immune Neglect: A Source of Durability Bias in Affective Forecasting," *Journal of Personality and Social Psychology*, 75, pp. 617-38.
- [15] Holmstrom, B. [1979], "Moral Hazard and Observability," *Bell Journal of Economics*, 10, pp. 74-91.
- [16] Kahneman, D., E. Diener, and N. Schwarz [1999], *Well-Being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation.
- [17] Kahneman, D., and A. Tversky [1979], "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, pp. 263-92.
- [18] Lazear, E.P., and S. Rosen [1981], "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 89, pp. 841-64.
- [19] Levin, J. [2003], "Relational Incentive Contracts," *American Economic Review*, 93, pp. 835-57.
- [20] Loewenstein, G., and D. Schkade [1999], "Wouldn't It Be Nice? Predicting Future Feelings," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 85-105.
- [21] Markowitz, H. [1952], "The Utility of Wealth," *Journal of Political Economy*, 60, pp. 151-8.
- [22] McMahon, D.M. [2004], "From the Happiness of Virtue to the Virtue of Happiness: 400 B.C. – A.D. 1780," *Daedalus*, 133, pp. 5-17.
- [23] Oswald, A.J. [1997], "Happiness and Economic Performance," *Economic Journal*, 107, 1815-31.
- [24] Radner, R. [1980], "Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies With Long but Finite Lives," *Journal of Economic Theory*, 22, pp. 136-54.

- [25] Robson, A.J. [1996], "A Biological Basis for Expected and Non-Expected Utility," *Journal of Economic Theory*, 68, pp. 397-424.
- [26] Robson, A.J. [2001a], "Why Would Nature Give Individuals Utility Functions," *Journal of Political Economy*, 109, pp. 900-14.
- [27] Robson, A.J. [2001b], "The Biological Basis of Economic Behavior," *Journal of Economic Literature*, 39, pp. 11-33.
- [28] Rogerson, W.P. [1985], "The First-Order Approach to Principal-Agent Problems," *Econometrica*, 53, pp. 1357-68.
- [29] Samuelson, L. [2004], "Information-Based Relative Consumption Effects," *Econometrica*, 72, pp. 93-118.
- [30] Samuelson, L., and J.M. Swinkels [2004], "Information and the Evolution of the Utility Function," mimeo, University of Wisconsin-Madison.
- [31] Sapolsky, R. [1999], "The Physiology and Pathophysiology of Unhappiness," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 453-69.
- [32] Simon, H.A. [1959], "Theories of Decision-Making in Economics and Behavioral Science," *American Economic Review*, 49, pp. 253-83.
- [33] Stigler, G.J., and G.S. Becker [1977], "De Gustibus Non Est Disputandum," *American Economic Review*, 67, pp. 76-90.
- [34] Wilson, T.D., D.T. Gilbert, and D.B. Centerbar [2002], "Making Sense: The Causes of Emotional Evanescence," in Carrillo, J., and I. Brocas (eds.), *Economics and Psychology*, Oxford: Oxford University Press, pp. 209-33.
- [35] Wright, C.I., H. Fischer, P.J. Whalen, S.C. McInerney, L.M. Shin, and S.L. Rauch [2001], "Differential Prefrontal Cortex and Amygdala Habituation to Repeatedly Presented Emotional Stimuli," *NeuroReport*, 12, pp. 379-83.