EVOLUTIONARY EFFICIENCY AND MEAN REVERSION IN HAPPINESS

LUIS RAYO AND GARY S. BECKER

ABSTRACT. We model happiness as an innate incentive mechanism molded by natural selection. The metaphorical objective is to motivate the individual towards seeking goals that favor genetic replication. The end result of this evolutionary process is a mean-reverting happiness function that is based on a context-dependent reference point. This reference point will incorporate all information that is valuable in terms of assessing the individual's performance, which in turn leads to a process of habit formation and peer comparisons. Based on principles of optimal incentives and statistical inference, the model delivers a unified account of several stylized facts surrounding the hedonic experience.

1. Introduction

Extensive studies have confirmed that happiness is no straightforward consequence of our objective conditions. Central to the psychological experience is the fact that these objective conditions are evaluated via a frame of reference that happens to change over time –e.g., Kahneman [2003]. From an economic standpoint, two phenomena are of particular interest. The first is the process of habituation, whereby permanent changes in our material conditions tend to create only transitory hedonic effects. This process has been documented for a number of variables, ranging from income to climate to the quality of the view from a condo –Brickman and Campbell [1971], Brickman, Coates and Janoff-Bulman [1978], Frederick and Loewenstein [1999]. The second phenomenon concerns social comparisons. Namely, the level of satisfaction derived from virtually any material outcome, and even from our physical state, is constantly influenced by what is achieved by our peers –Scitovsky [1992], Diener and Fujita [1996].[1]

Date: First draft: 5/03. This draft: 5/04. Work in progress.

The University of Chicago. Acknowledgments to be added.

[1]Happiness and unhappiness are interpreted broadly as positive or negative emotions and feelings. See Kahneman, Diener, and Schwarz [1999] for a comprehensive review of the current happiness knowledge.

These phenomena can be described by means of a subjective "reference point" that serves as a benchmark against which our current material conditions are compared. The comparison will establish weather an outcome constitutes a success or a failure, thus leading to either positive or negative feelings. Habits and peer comparisons can be interpreted precisely as changes in the reference point. In the psychology literature, the reference point has also been associated to the individual's level of expectations, which fits the observation that positive and negative surprises have a large impact over the hedonic experience. A constantly shifting reference point is also consistent with what is perhaps the most striking feature of happiness: its temporary nature.

The formation of habits and social comparisons has been documented in every known human culture. They are also present in young children. Habituation and envy, for example, are part of Brown's [1999] list of human universals. More generally, the psychological studies on happiness also deliver consistent results across countries and time. As a result, it is reasonable to assume that the basic mechanisms of happiness are innate; from which we can conjecture that these mechanisms served an evolutionary purpose in the descent of our species.

We develop a model that organizes the above happiness facts in an attempt to uncover their evolutionary rationale. The theory we propose is that the reference point serves as an optimal incentive device in a moral-hazard agency problem representing the process of natural selection. In this agency problem, natural selection serves as a principal that designs the happiness function of the individual. The metaphorical objective of this principal is to motivate the individual, via happiness, towards seeking goals that favor genetic replication –a simple example is rewarding the individual for having sex with an appropriate mate. The individual, in turn, simply seeks high levels of happiness; but in the process he also serves as an agent to genetic replication.

The solution of the optimal incentive problem faced by the principal represents the end point of the evolutionary process, where the world is populated exclusively by individuals equipped with optimal happiness functions. Throughout the analysis, we limit attention to this end point of natural selection. Frank [1987], and Samuelson and Swinkels [2002] develop additional applications of this type of agency approach –their purpose, however, is not the study of happiness.

The central assumption of the theory is the existence of a physiological constraint that imposes an upper and lower bound on the levels of happiness that can be experienced by the individual.[2] An analogy employed in the psychology literature is that happiness is like blood pressure –Wilson, Gilbert, and Centerbar [2002]. There is a healthy *range* for each, while the particular level experienced at any given moment depends on current needs. In the case of blood pressure, these needs are a function of physical activity. In the case of happiness, the current needs arise precisely from its role as a motivation device, whereby some outcomes are punished, and others rewarded. When these bounds on happiness are combined with a limit on how the function can be rescaled, happiness becomes scarce as an incentive instrument, implying that it must be sparingly employed.

The solution to the optimal incentive problem is a mean-reverting happiness function that is based on a context-dependent reference point. The unifying principle set forth is that the reference point optimally aggregates all information available towards assessing the individual's current performance. Habits and envy correspond to particular cases of this process, wherein past outcomes and peer success enter the reference point via their information content in terms of the individual's current potential.

The most distinctive feature of the agency approach is a difference-in-differences result, which implies that the process of habituation extends to peer output as well. A manifestation of this result is that permanent increases in peer output leading to envy have only temporary effects. Equivalently, this result can be interpreted as a habit that occurs at the level of the individual's relative social position, as opposed to her absolute income alone. In section 5, we contrast this result with the existing happiness evidence.

2. Core Model

We are interested in the end-point, or steady state, of an evolutionary process. During this process, the object over which natural selection operates is the shape of the individual's happiness function. The shape of this function is information contained in the individual's genes. The sole objective of the individual is to maximize her level of happiness, but in the process of pursuing this goal she will incur in actions that favor the multiplication of her genes to a greater or lesser extent. Over the long-run, due to selective pressures, only those happiness functions that best serve this ultimate purpose of genetic multiplication will dominate the population.

[2]The constraints can be viewed as either physical limits on the process generating emotions, or a physiological cost of extreme emotional states –e.g., Sapolsky [1999].

Crucially, when we speak of the evolutionary end-point we refer to the ancestral hunter-gather environment and the suitable adaptations developed back then. In the modern world, in contrast, the rate of environmental change has dwarfed the rate of evolutionary adaptation, resulting in a level of mal adaptation in many respects. Our approach will be relevant to the extent that the adaptations developed in the ancestral enforcement are still present in our innate characteristics today.

We model this evolutionary end-point in a stylized way, for which we provide a primitive foundation in the following section. Consider an individual, called the agent, that lives for only one period. The choice variable is the level of effort she exerts towards achieving happiness. Denote this effort by $e \in \mathbb{R}$. Effort stochastically determines the realization of a level of output $y \in \mathbb{R}$. This output is intended to represent a one-dimensional summary of "proximate" evolutionary goals. Namely, those tangible goals that favored the ultimate evolutionary goal of genetic replication during the ancestral environment. Common examples include wealth, health, and sex, as well as the well-being of friends and kin.

After output y is realized, the agent experiences a level of happiness $V(y) \in \mathbb{R}$, which is meant as a summary of emotions. We assume that the function V is bounded below and above. In particular, $V(y) \in [0, 1]$ for all y. These bounds represent a physiological constraint. When selecting her level of effort e, the agent's objective will be to maximize the expected value of happiness, conditional on effort, minus a shadow cost of effort:

$$\max_{e} E[V(y) \mid e] - e. \tag{1}$$

As explained in the following section, this shadow cost of effort will not represent an emotion itself. Rather, it will describe the opportunity cost of the resources employed in the pursuit of output.

The happiness function V will be designed by a metaphorical principal, representing the process of natural selection. The objective for this principal will be to maximize the expected level of y, via the incentives provided by V towards a high level of effort. The problem faced by the principal is mathematically equivalent to a moral-hazard agency problem, where the level of effort is hidden from the principal, and V represents an incentive payment for the agent. Letting e^* represent the effort selected in equilibrium by the agent, this problem becomes

$$\max_V \; E[y \mid e^*] \tag{2}$$

$$s.t.$$

$$e^* = \arg\max_e \; E[V(y) \mid e] - e, \tag{IC}$$

$$V(y) \in [0,1] \text{ for all } y.$$

The resulting happiness functions, which we call *efficient*, will be the output of the model.[3]

We begin by specifying a technology relating effort to output. Suppose $y = \pi(e)+z$, where π is smooth, concave, and increasing; and z is a zero-mean exogenous shock with a smooth density g that is increasing over $\mathbb{R}_-$ and decreasing over $\mathbb{R}_+$. Notice that $\pi(e)$ corresponds to $E[y \mid e]$, the principal's objective. Let $F(y \mid e)$ denote the conditional c.d.f. for output, with density $f(y \mid e)$.

We attack this incentive problem via a first-order approach (Rogerson [1985]), where the effort incentive constraint for the agent (IC) is replaced by the constraint

$$\frac{\partial}{\partial e} E[V(y) \mid e^*] - 1 \geq 0. \tag{FOC}$$

This new constraint represents the first order-condition for the agent when selecting her level of effort. The inequality is present because the constraint will only bind in the direction of less effort.[4] Once (FOC) takes the place of (IC), observe that a necessary condition for efficiency of a happiness function is that, given the equilibrium level of effort e^*, this function solves $\max_V \frac{\partial}{\partial e} E[V(y) \mid e^*]$. Otherwise, ($FOC$) could be relaxed and the level effort increased, a contradiction to the fact that V was efficient. Intuitively, this requirement means that the incentive potential of happiness has been fully exploited.

[3] In the results that follow, the bounds on happiness always bind, suggesting that the evolutionary process will presumably act on these bounds. However, provided some type of bounds remain, this process will simultaneously favor the happiness functions that best exploit the available range.

[4] This first-order approach is valid as long as $F(y \mid e)$ is concave in e, which in turn is guaranteed when π is sufficiently concave. We assume that this is the case.

For future reference, notice that $E[V(y) \mid e^*]$ can equivalently be written as $E[V(y) \mid \pi(e^*)]$, and therefore this necessary condition for efficiency can be expressed as

$$\max_V \frac{\partial}{\partial \pi} E[V(y) \mid \pi(e^*)] \tag{3}$$

$$s.t. \tag{IC}$$

$$V \in [0,1] \text{ for all } y.$$

Conversely, any happiness function that solves problem (3) will also solve the original problem (2) provided the first-order approach is indeed valid. In which case, problem (3) characterizes all efficient happiness functions.

Before solving this problem, we take a technical detour in order to motivate our setup. We do so by developing a more primitive model that serves as an equivalent dual problem delivering the same happiness functions. After doing so, we return to the above agency setup, which will provide considerable intuition for the results that follow. The reader can feel free to skip to section B.

A. Selecting a course of a action

Consider a more primitive problem where, instead of selecting a level of effort, the agent must decide upon a course of action $x \in X$. We allow X to be an arbitrary metric set, where the distance between any two courses of action x and x' is given by $d(x, x')$. Before selecting x, the agent will observe a random variable α, which summarizes the current state of nature. The combination of x and α generates a probability distribution over output y, which describes proximate evolutionary goals. This distribution is given by the conditional density $f(y \mid x, \alpha)$, which is known by the agent.[5] Suppose the evolutionary process metaphorically ranks different distributions f according to their first moment $\int yf dy = E[y \mid f]$. In other words,

[5]As an example, suppose the agent is searching for food via the recollection of fruit. Output will represent the amount of fruit she obtains, whereas the resource allocation x is the strategy followed in this pursuit. The set X is intended to represent the myriad courses of action that an individual can take in reality towards achieving a specific goal –these are meant to represent even the smallest details, from traveling in a particular direction to climbing a particular tree using a precise sequence of steps. In practice, the value of each one of these actions will depend on the state of conditions currently faced by the individual –like the presence of a fruit in a specific location. The shock α represents these conditions.

output y serves as aggregator that maps a functional space of distributions f into a single dimension, thus establishing a means to compare them.

The agent, in contrast, ranks distributions f using a different aggregator, called happiness, which is a function $V : y \to \mathbb{R}$. In particular, the agent's rank is determined by $\int V(y) f dy = E[V(y) \mid f]$.[6] After y is realized, happiness $V(y)$ will be physically experienced by the individual. As such, we assume that $V(y)$ is bounded above and below.[7] Notice that $E[V(y) \mid f]$ serves as the agent's "decision utility" –a function of actions– whereas $V(y)$ corresponds to her "experienced utility," from which decision utility is constructed. We return to this point in our concluding remarks.

In order to establish a link with the principal-agent model of the previous section, we assume the following technological structure: (i) $y = E[y \mid x, \alpha] + z$, where z corresponds to the zero-mean exogenous shock described above, realized after x is selected.[8] (ii) $E[y \mid x, \alpha]$ has a unique maximizer, denoted by $x^*(\alpha)$. And (iii) $E[y \mid x, \alpha]$ can be expressed as $E[y \mid d(x^*(\alpha), x)]$, which is smooth and decreasing in d.[9]

[6]Notice that this setup corresponds to expected utility theory, where the Bernoulli utility over outcomes if given by V.

[7]Notice also the assumption that V does not depend on the particular selection of x. In practical terms, this means that the individual is rewarded for actually finding some fruit, not from searching in any specific location. In the other extreme, the agent could be rewarded whenever she follows the exact optimal course of action $x^*(\alpha)$, regardless of the realization of y. This would be equivalent to building an automata that is hard-wired to select $x^*(\alpha)$ upon observing any state α. Important examples of this type of process include non-conscious automated responses such as our hear-beat –see also Samuelson and Swinkles [2002] on urges. But for many other decisions, this type of hard-wiring has been replaced by a happiness reward system based on a project's success –presumably due to the complexity of the environment together with the complexity of the set of feasible courses of action. In the present discussion, these are the decisions of interest.

[8]In this case, the distribution of y is derived from the distribution of z, given by $g(z)$. Namely, $f(y \mid x, \alpha) = g(y - E[y \mid x, \alpha])$.

[9]These technological assumptions are made for analytical simplicity. A weaker version of (iii) that delivers the same results is that the optimized expected output, given by $E[y \mid x^*(\alpha), \alpha]$, is independent of α. In other words, the state α determines only the optimal course of action, not the highest possible level of expected output that can be achieved. The original formulation of (iii) simplifies the exposition.

Since x only affects the distribution of y via the distance d, we can express the agent's objective as

$$\max_x \; E[V(y) \mid d(x^*(\alpha), x)] \quad \text{s.t.} \quad x \in X, \tag{4}$$

where V was previously molded by natural selection. Importantly, the expression $E[V(y) \mid d(x^*(\alpha), x)]$ is merely a mathematically simplified way of expressing the agent's objective. Underlying this expression is the fact that the agent literally ranks different courses of action x according to the distributions $f(y \mid x, \alpha)$ they generate, and the subsequent aggregation process $\int V(y)f(y \mid x, \alpha)dy$.[10]

Consider now the evolutionary problem of selecting a happiness function $V \in [0, 1]$ that maximizes $E[y \mid d]$ subject to the decision process in (4). It turns out that a large family of happiness functions will guarantee that (4) leads to the selection of $x^*(\alpha)$. Indeed, this will be true for any happiness function such that $E[V(y) \mid d]$ is decreasing in d.[11] Moreover, this property will hold for any arbitrary rescaling of V, allowing, for example, for happiness functions that are arbitrarily flat. However, this set of functions can be sharply refined by introducing a bound on the precision of the happiness instrument.

For this we return to the agent's ranking of alternative courses of action. Suppose that when evaluating x via the aggregation $\int V(y)f(y \mid x, \alpha)dy$, the agent makes a small random error $\epsilon(x)$ –e.g., a random computational error. Thus, when deciding between x and x', she will compare $E[V(y) \mid x, \alpha]+\epsilon(x)$ against $E[V(y) \mid x', \alpha]+\epsilon(x')$ and select the action for which this expression is larger. Consequently, even when $E[V(y) \mid d]$ is decreasing in d, the agent will not necessarily select $x = x^*(\alpha)$. We are interested in happiness functions that minimize the evolutionary cost of this type of mistake in the limit as the errors ϵ approach zero. These functions, in turn, will coincide with the happiness functions that solve the original problem (2).[12]

[10] In other words, feelings V are literally used to evaluate alternative courses of action. Damasio [1994] presents neurological evidence supporting this view.

[11] Consider, for example, any function such that: $V(y) = 1$ whenever $y \geq \overline{y}$, where $\overline{y}$ is an arbitrary threshold, and $V(y) = 0$ otherwise. In such case, the agent's objective $E[V(y) \mid d]$ equals the probability that the shock z exceeds $\overline{y} - E[V(y) \mid d]$, and this probability is decreasing in d.

[12] For this analysis we require the technical assumption that, for all α, $x^*(\alpha)$ is an accumulation point of X. This simply means that there exist values of x arbitrarily close to $x^*(\alpha)$.

Consider the problem of minimizing the likelihood that the agent ranks an $x \neq x^*(\alpha)$ higher than $x^*(\alpha)$. Letting $\delta = d(x^*(\alpha), x)$, this problem is equivalent to that of maximizing the difference between $E[V(y) \mid d = 0]$ and $E[V(y) \mid d = \delta]$. Using a first-order Taylor expansion, we can express this difference by

$$E[V(y) \mid d = 0] - E[V(y) \mid d = \delta] \approx -\frac{\partial}{\partial d} E[V(y) \mid d = 0] \cdot \delta.$$

As the random errors ϵ become small, so will the distance between $x^*(\alpha)$ and any x selected with positive probability. Thus, the relevant values of δ will also become small. In the limit as $\delta \to 0$ the above approximation becomes exact, and the happiness functions that minimize the likelihood of a departure from $x^*(\alpha)$ correspond to those that maximize the derivative $-\frac{\partial}{\partial d} E[V(y) \mid d = 0]$.

Intuitively, the evolutionary advantage of a high derivative $-\frac{\partial}{\partial d} E[V(y) \mid d = 0]$ is that it increases the likelihood that a potential deviation away from $x^*(\alpha)$ is actually perceived by the agent as costly. In this sense, a happiness function that maximizes $-\frac{\partial}{\partial d} E[V(y) \mid d = 0]$ will also maximize the incentive to remain close to $x^*(\alpha)$.

Now observe that we can express the agent's objective as $E[V(y) \mid E[y \mid d]]$, and therefore

$$-\frac{\partial}{\partial d} E[V(y) \mid d = 0] = -\frac{\partial}{\partial d} E[y \mid d = 0] \cdot \frac{\partial}{\partial E[y \mid d]} E[V(y) \mid E[y \mid d = 0]],$$

from which it follows that the maximization of $-\frac{\partial}{\partial d} E[V(y) \mid d = 0]$ is equivalent to the maximization of $\frac{\partial}{\partial E[y|d]} E[V(y) \mid E[y \mid d = 0]]$.

Thus, after a normalization of the units in problem (2), so that $\pi(e^*) = E[y \mid d = 0]$, the maximization of $-\frac{\partial}{\partial d} E[V(y) \mid d = 0]$ will be equivalent to the maximization of $\frac{\partial}{\partial \pi} E[V(y) \mid \pi(e^*)]$, which arises from problem (2). In both cases, the objective is to maximize the marginal responsiveness of the agent's objective to changes in expected output at the point where the latter equals its highest possible level. In other words, the exogenous cost of effort in problem (2) plays the same role as the random optimization errors; both of which lead to the selection of happiness functions that are maximally responsive, in expectation, to changes in the agent's decision. The appendix below describes a further connection between problems (2) and (4) by means of a dual formulation of the latter.

B. The Happiness Reference Point

We return to the original principal-agent problem (2), and proceed to characterize the efficient happiness functions that maximize the derivative $\frac{\partial}{\partial \pi} E[V(y) \mid \pi(e^*)]$, where $\pi(e) = E[y \mid e]$, and e^* is the equilibrium level of effort. Letting $f(y \mid \pi) = g(y - \pi)$, this derivative is given by

$$\frac{\partial}{\partial \pi} \int V(y) f(y \mid \pi) dy \bigg|_{\pi = \pi(e^*)} = \int V(y) f_\pi(y \mid \pi(e^*)) dy$$
$$= \int V(y) \left\{ -g'(y - \pi(e^*)) \right\} dy.$$

Since $g'(z)$ is assumed to be positive for all $z < 0$ and negative for all $z > 0$, the term in brackets in the last integral is negative for all $y < \pi(e^*)$ and positive for all $y > \pi(e^*)$. As a result, the derivative is maximized by setting $V(y) = 0$ –the lower bound– whenever $y < \pi(e^*)$, and $V(y) = 1$ –the upper bound –whenever $y > \pi(e^*)$. The value of $V(\pi(e^*))$ is immaterial, and for concreteness we assume $V(0) = 1$. Observe that up to a zero-measure set, this function is unique.[13]

The expected value of output $\pi(e^*) = E[y \mid e^*]$ serves as a reference point. Whenever y exceeds this expectation, the agent is rewarded with the highest level of happiness, and whenever y falls short, the converse occurs. In other words, happiness is fully determined by the sign of the surprise. This function corresponds to the limiting case of the utility functions in prospect theory –Kahneman and Tversky [1979]– where the marginal impact of a surprise is always decreasing, i.e., utility is concave to the right of the reference point, and convex to the left.[14] In the limiting case, the size of the surprise is irrelevant. This shape is analytically convenient, since the entire function is summarized by its reference point.[15]

This efficient happiness function has a direct interpretation in the agency framework. Under the first-order approach to the agent's effort constraint, the purpose of V is to maximally deter a marginal deviation away from e^*. This is illustrated in figure 1. The upper panel graphs the distribution of y conditional on effort. This is a single-peaked function centered at $E[y \mid e]$. The bold line represents the distribution when the agent follows her prescribed level of effort e^*. Consider now

[13]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. Analogous rewards arise in Levin [2003] under relational contracts.

[14]For example, minor favorable surprises (such as finding some trivial amount of money) have a sizable impact on reported well-being. See Schwarz and Strack [1999].

[15]Smoother happiness functions would arise if, instead of imposing an upper and lower bound on happiness, we assumed that there is a benchmark level of happiness that is physiologically optimal, and that deviations from this level are increasingly costly. See Wilson et al. [2002] for an analogy with blood sugar, as opposed to blood pressure.

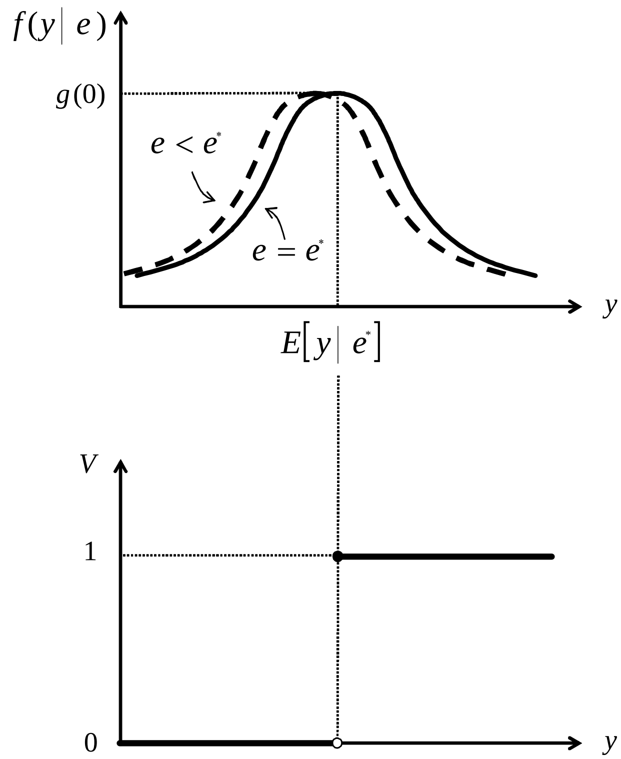

Figure 1. The First-Order Approach

a marginal deviation to $e < e^*$. This deviation will cause the distribution to shift to the left, as shown by the dashed function. As a result, the density of all values of $y < E[y \mid e^*]$ will increase, while the density for all values of $y > E[y \mid e^*]$ will decrease. As shown by the lower panel, in order to maximally deter this marginal deviation, the principal will punish all values $y < E[y \mid e^*]$ using the lowest possible level of happiness, and reward all other values with the highest possible level.[16]

3. Dynamic Environment

We now extend the core model to a dynamic setup, where the agent lives for multiple periods $t = 1, 2...$ For simplicity, we identify each period with the length of one project. At the beginning of period t, the agent selects a level of effort e_t –or a course of action x_t in the dual problem presented above. This effort determines a distribution function f_t over different levels of output y_t, which can represent

[16]This happiness function can also be interpreted as implicitly performing a Neyman-Pearson statistical test, based on y, against marginal effort deviations. The null hypothesis for this test is "$e = e^*$," an the rejection region is given by $y < E[y \mid e^*]$.

anything from successfully collecting food to finding a sexual partner to helping a friend. In a modern environment, it could also represent, for example, successfully saving for future retirement. We assume that y_t serves as a sufficient statistic for the impact of e_t over the ultimate evolutionary goal, which simply means that alternative distributions f_t can be evaluated, without loss, via their impact over y_t alone. Thus, in the metaphorical principal-agent model, we assume that the objective for the principal in period t is given by $E[y_t \mid f_t]$.

As in the core model, the agent ranks different distributions f_t via a happiness function V_t, which aggregates each distribution f_t into a single dimension. Happiness will be experienced at the end of the period, after y_t is realized; it will now be a function of y_t –the evolutionarily relevant objective– together with a vector of additional information Ω_t that is available to the agent at the time happiness is experienced. As explained below, this vector Ω_t will include information of past levels of output y_{t-s}, as well as the levels of output produced by the agent's peers, both of which will be technologically linked to y_t.

The period-t objective for the agent is given by

$$\max_{e_t} E[V_t(y_t, \Omega_t) \mid e_t] - e_t. \tag{6}$$

This formulation assumes that the agent selects e_t based on its impact over current happiness alone, without additional reference to future levels of happiness V_{t+s}. (Formally, this corresponds to a multi-self framework, where one self is alive every period.) This by no means implies that the agent is myopic in the sense of not caring about her future material well-being, it simply means that any impact that her current actions have over the future will be evaluated via current emotions based on y_t. This represents, for example, an individual that saves for her retirement precisely because it leads to a contemporaneous psychological reward. More generally, this process corresponds to a view of the human intertemporal decision process wherein current intertemporal decisions are based on current emotional rewards, not on intertemporal trade-offs of emotions.[17]

[17]Such view is closely related to the phenomena of emotional anticipation, whereby news about future material outcomes will have an immediate emotional effect. Notice, on the other hand, that an individual that cares only about current happiness will be, by construction, fully myopic terms of intertemporal trade-offs of happiness. If offered, for example, an $\epsilon > 0$ increase in current happiness, followed by a one-unit reduction in happiness for every subsequent period into the indefinite future, this agent will always accept. Why? Precisely because she will feel ϵ happier if

We proceed with several examples where the efficient happiness functions are built with habits and/or peer comparisons through the vector Ω_t. These effects will arise via the technological process behind y_t. For concreteness, y_t can be interpreted as income. For all the technologies considered, y_t will be determined by

$$y_t = E[y_t \mid e_t, \Omega_t] + z_t, \tag{7}$$

where the shock z_t is i.i.d. across time, and follows the single-peaked distribution g described above. The problem for the principal, for each period t is given by

$$\max_{V_t} \frac{\partial}{\partial e_t} E[V_t(y_t, \Omega_t) \mid e_t^*]$$
$$s.t.$$
$$V_t(y_t, \Omega_t) \in [0,1] \text{ for all } y_t \text{ and } \Omega_t,$$

which corresponds to the maximization of $E[y_t \mid e_t, \Omega_t]$ subject to the first-order condition for (6).

As in the static model, the happiness function solving this problem will have a one-step shape. The conditional expectation $E[y_t \mid e_t^*, \Omega_t]$ will serve as a reference point. Denote this reference point by $\widehat{y}_t$. After y_t is realized, the agent will experience happiness $V_t = 0$ whenever y_t falls short of $\widehat{y}_t$, and otherwise she will experience $V_t = 1$.

Example 1: A Markovian Habit

Suppose the world is populated by a single agent, and output is given by $y_t = \pi(e_t) + \theta_t$, where θ_t follows the Markovian process $\theta_t = \theta_{t-1} + z_t$. In this case, for every period t, the efficient happiness function will have a reference point $\widehat{y}_t = y_{t-1}$:

Proposition 1. *Under the above Markovian technology with a single agent, for every t, the efficient happiness function is given by*

$$V(y_t, y_{t-1}) = \begin{cases} 1 & if \ y_t \geq y_{t-1}, \\ 0 & if \ y_t < y_{t-1}. \end{cases}$$

Proof. The expected value of output for period t, conditional on $e_t = e_t^*$, is given by $\pi(e_t^*) + \theta_{t-1}$. Since output follows a single-peaked distribution around this mean, the efficient happiness function has one step and a reference point equal to $\widehat{y}_t =$

she accepts. These type of trade-offs, of course, are not offered in practice: real-world trade-offs are based on material outcomes, not emotional experience.

$\pi(e_t^*) + \theta_{t-1}$. In equilibrium, θ_{t-1} can be inferred from y_{t-1} using the equality $\theta_{t-1} = y_{t-1} - \pi(e_{t-1}^*)$. As a result, the reference point can be expressed as $\widehat{y}_t = \pi(e_t^*) + y_{t-1} - \pi(e_{t-1}^*)$. It remains to show that $e_t^* = e_{t-1}^*$. The equilibrium effort can be inferred from the agent's first-order condition, which must hold with equality. In order to obtain this first-order condition, notice that the expected value of happiness is equal to the probability that output $y_t = \pi(e_t) + \theta_{t-1} + z_t$ exceeds the reference point $\widehat{y}_t = \pi(e_t^*) + \theta_{t-1}$, which in turn is equal to $prob\{z_t \geq \pi(e_t^*) - \pi(e_t)\} = \int_{\pi(e_t^*)-\pi(e_t)}^{\infty} g(z)dz$. The first-order condition for the agent therefore becomes

$$\frac{\partial}{\partial e_t} E[V_t \mid e_t^*] - 1 = g(\pi(e_t^*) - \pi(e_t^*))\pi'(e_t^*) - 1 = 0.$$

From which it follows that the equilibrium levels of effort solve, for all t, the equality $g(0)\pi'(e_t^*) = 1$, and therefore $e_t^* = e_{t-1}^*$. ∎

The argument behind the proof is the following. Using the equality $y_{t-1} = \pi(e_{t-1}^*) + \theta_{t-1}$, current output y_t can be expressed as $\pi(e_t) + \theta_{t-1} + z_t = \pi(e_t) + y_{t-1} - \pi(e_{t-1}^*) + z_t$. Moreover, the additive structure of the random shocks will imply that the equilibrium levels of effort remain constant across periods, so that $e_{t-1}^* = e_t^*$. Therefore, we can write

$$y_t = y_{t-1} + \pi(e_t) - \pi(e_t^*) + z_t. \tag{8}$$

From (7), $E[y_t \mid e_t, \Omega_t]$ corresponds to $y_{t-1} + \pi(e_t) - \pi(e_t^*)$, with $\Omega_t = y_{t-1}$. The optimal happiness reference point equals $E[y_t \mid e_t, \Omega_t]$ evaluated at $e_t = e_t^*$, which finally reduces to y_{t-1}, the best predictor of y_t.

This result can be readily interpreted within the optimal incentive framework. From (8), when assessing marginal effort deviations, the principal can restrict to the measure $y_t - y_{t-1}$, which is her best source of information regarding the agent's current behavior. In addition, notice that a marginal effort deviation will affect $y_t - y_{t-1}$ in exactly the same way as a low realization of z_t –i.e., both these effects are indistinguishable for the principal. Consequently, even though in equilibrium $y_t - y_{t-1} = z_t$, the principal must punish the agent following low realizations of z_t and vice versa. The result is a happiness function that provides the lowest payoff whenever $y_t - y_{t-1} = z_t < 0$, and vice versa, which is equivalent to setting a reference point $\widehat{y}_t = y_{t-1}$.

This example provides a stylized theory of the process of habituation, or hedonic adaptation, whereby the agent becomes accustomed to her current material standard of living –e.g., Frederick and Loewenstein [1999].[18] As in prospect theory, the carrier of happiness is the change in output, not its absolute value. In practical terms, this means that the agent will experience positive feelings if she manages to improve her standard of living, and vice versa. This type of process has been well documented, and is especially acute when it comes to income.[19]

Under this technology, the process of habituation and the formation of expectations are one and the same: the habit continuously guarantees that the reference point $\widehat{y}_t$ is equated to the expected level of y_t, so that only the random surprises have an effect over happiness. The result is a mean-reverting process in the level of happiness, where its expected value is the same every period, regardless of past levels of output. Only the sign of the random surprise z_t will determine whether current happiness turns out to lie above or below its long-term mean.

Example 2: Static Peer Comparisons

[18]Frederick and Loewenstein [1999] suggest the following functions for the process of hedonic adaptation: 1. Persistent strong hedonic states can be destructive, hedonic adaptation may help protect us from these effects. 2. It may increase our sensitivity to, and motivation to make, local changes in our objective circumstances. 3. It may prevent the continued expenditure of energy in futile attempts to change the unchangeable and redirect motivation to changes that can be made.

The conceptual addition of our theory is the notion of a statistically optimal use of the available information as a means to evaluate the individual's current performance. Under this notion, adaptation becomes one particular component of an efficient incentive mechanism, and therefore can be formally unified with the process of peer comparisons, as well as the role of expectations. Notice also that the objective of motivating the individual towards achieving local improvements in his circumstances is different from the objective of deterring local deviations away from the ideal course of action, which is the device we focus on. The latter is precisely the device that will allow us to unify habits and peer comparisons.

[19]To be sure, the purpose of our theory is not to describe the process of hedonic adaptation. Rather, our goal is to attempt a formal theory for its reason of being. By this we refer to a theory that goes beyond the following type of argument, which essentially assumes the result:

- *The hedonic intensity of a number of stimuli fades away with time because the neural process adjusts to these stimuli.* As a result, only differences in the status quo create a hedonic effect. Habits are a manifestation of this process.

Suppose the world is populated by a continuum of symmetric agents indexed by i –each one with her own principal. The output for each agent is given by $y_t^i = \pi(e_t^i) + \theta_t^i$, where $\theta_t^i = \Gamma_t + z_t^i$. The first component of the shock, Γ_t, represents an aggregate shock which is shared by all agents. The second component, z_t^i, is an individual shock that is i.i.d. across agents, and follows the distribution g. We assume that the shocks θ_t^i are independent across time, and that the population mean across the individual shocks z_t^i is equal to zero –i.e., an exact law of large numbers applies. We are interested in a symmetric steady state of the evolutionary process where every single agent is equipped with an efficient happiness function that maximizes her level of effort. Dropping the i superscript, the efficient happiness function for each individual will have a reference point $\widehat{y}_t = \overline{y}_t$, where $\overline{y}_t$ is the average output across agents:

Proposition 2. *Under the above static technology with multiple agents, for every t, and every agent, the efficient happiness function is given by*

$$V(y_t, \overline{y}_t) = \begin{cases} 1 & if\ y_t \geq \overline{y}_t, \\ 0 & if\ y_t < \overline{y}_t. \end{cases}$$

The argument is the following. In a symmetric environment where every agent follows the same level of effort e_t^*, we can write $y_t^i = \pi(e_t^*) + \Gamma_t + z_t^i$. By averaging across all individuals, this implies $\Gamma_t = \overline{y}_t - \pi(e_t^*)$. Again dropping the i superscript, the output for each agent can be expressed as

$$y_t = \overline{y}_t + \pi(e_t) - \pi(e_t^*) + z_t. \tag{9}$$

Using (7), $E[y_t \mid e_t, \Omega_t]$ now corresponds to $\overline{y}_t + \pi(e_t) - \pi(e_t^*)$, with $\Omega_t = \overline{y}_t$. Recall that the optimal happiness reference point equals $E[y_t \mid e_t, \Omega_t]$ evaluated at $e_t = e_t^*$, which in this case reduces to $\overline{y}_t$.

In terms of optimal incentives, we can see from (9) that the most informative variable for the purpose of assessing effort becomes the difference $y_t - \overline{y}_t$, which in equilibrium is equal to the shock z_t. Consequently, happiness will be a function of this surprise alone, and will follow the same mean-reverting process described in the previous example. A subtle difference is the fact that $\overline{y}_t$ does not correspond to the agent's ex-ante level of output expectations. In fact, when e_t is selected, $\overline{y}_t$ has not yet been realized. Rather, $\overline{y}_t$ reflects the agent's ex-post output potential as measured by the conditional expectation $E[y_t \mid e_t, \overline{y}_t]$.

The carrier of happiness is now relative income $y_t - \overline{y}_t$. This comparison can be interpreted as a negative externality caused by peers. However, it is not an externality in the conventional sense. The reason why $\overline{y}_t$ enters the happiness function is because it filters out the aggregate shock Γ_t, and thus increases the statistical power of the happiness incentive device. As a result, peer effects are not intrinsically related to the absolute level of peer output. Rather, they arise because of the technological interdependence that exists across peers. The resulting incentive device is analogous to a relative performance scheme inside a firm. By tightening the connection between effort and reward, it's effect is to magnify the cost of withdrawing effort.[20]

Example 3: A Markovian Habit and Dynamic Peer Comparisons

The most distinctive implication of the agency approach arises when habits and peer comparisons are combined. We begin with a simple output process that combines the technologies from the two previous examples. The output for each agent is given by $y_t = \pi(e_t) + \theta_t$, where θ_t follows the Markovian process $\theta_t = \theta_{t-1} + \Gamma_t + z_t$. The difference with example 1 is that the innovation $\theta_t - \theta_{t-1}$ is now determined by both an aggregate and an idiosyncratic shock; and the difference with example 2 is that these shocks are now fully persistent.

The efficient reference point will not be the mere sum of y_{t-1} and $\overline{y}_t$. 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 $\overline{y}_t$. This would imply comparing an innovation against an absolute level. Rather, the carrier of happiness becomes the difference in differences

$$\Delta y_t - \Delta \overline{y}_t. \tag{10}$$

[20]Robson [2002] and Samuelson [2003] develop alternative evolutionary arguments in order to account for relative concerns. Samuelson shows that relative concerns leading to imitation can be used as a substitute for cognition when the latter cannot be fully trusted. Robson, on the other hand, points out that if utility, a function of fitness, can only attain a limited number of values, then the highest value will be reserved for the level of fitness associated to the evolutionarily optimal choice. Moreover, if this fitness level changes across time, as signaled by peers, then it will be desirable to adjust the utility scale. Both authors abstract from the more general study of happiness –e.g., mean reversion, surprises, and habits.

This result is obtained by noticing that, in a symmetric equilibrium where efforts are constant across time, output can be expressed as

$$\begin{aligned} y_t &= \pi(e_t) + \theta_{t-1} + \Gamma_t + z_t \\ &= y_{t-1} + \Delta\overline{y}_t + \pi(e_t) - \pi(e_t^*) + z_t. \end{aligned}$$

The key step is that Γ_t equals the change in average output $\Delta\overline{y}_t$. Again using expression (7), $E[y_t \mid e_t, \Omega_t]$ now corresponds to $y_{t-1} + \Delta\overline{y}_t + \pi(e_t) - \pi(e_t^*)$, with $\Omega_t = \langle y_{t-1}, \Delta\overline{y}_t \rangle$. The efficient reference point therefore becomes $\widehat{y}_t = E[y_t \mid e_t^*, \Omega_t] = y_{t-1} + \Delta\overline{y}_t$.

Proposition 3. *Under the above Markovian technology with multiple agents, for every t, and every agent, the efficient happiness function is given by*

$$V(y_t, y_{t-1}, \Delta\overline{y}_t) = \begin{cases} 1 & if \;\; y_t \geq y_{t-1} + \Delta\overline{y}_t, \\ 0 & if \;\; y_t < y_{t-1} + \Delta\overline{y}_t. \end{cases}$$

As in the previous examples, the carrier of happiness, $\Delta y_t - \Delta\overline{y}_t$, will equal the idiosyncratic shock. Happiness will therefore preserve its mean-reverting feature and its exclusive dependence of the sign of the random surprise. The novel implication of this example is the fact that average peer output is lagged with a negative sign.[21]

This difference-in-differences result can be interpreted in two alternative ways. The first is that the process of habituation also occurs at the level of peer output. Consider a sudden and permanent increase in $\overline{y}_t$. This increase will initially shift the reference point to the right, with a likely decrease in happiness, which can be interpreted as envy. But after one period, $\overline{y}_t$ will enter the reference point with a negative sign, shifting it back to its original level so that envy is gone. After one period, the agent has successfully coped.

The second interpretation is that the agent cares about her relative social position $y_t - \overline{y}_t$ as in example 2; but, in addition, she experiences a habit in terms of this social position. The habit means that her social position enters as a lagged term $y_{t-1} - \overline{y}_{t-1}$ in the reference point. As a result, the carrier of happiness becomes the change in

[21] The key is that the reference point must be informationally efficient. Thus, whenever there is a habit, there must exist a positive intertemporal correlation in the level of output. But, in a symmetric environment, this means that there must also exist a positive intertemporal correlation in the level of aggregate output. This in turn implies that $\overline{y}_t$ must contain redundant information regarding shocks that occurred in the past –in the above example notice that $\overline{y}_t = \Gamma_t + \overline{y}_{t-1}$, and only Γ_t is valuable information. In order to eliminate this redundant information, aggregate output is lagged.

social position, a second difference as well. Notice that both interpretations share the feature that the process of adaptation extends beyond the agent's own material conditions in a way that encompasses the conditions of her peers.

4. Multi-Period Habits

We conclude the model by presenting a more general result that encompasses the above propositions. Suppose output for each agent is given by $y_t = \pi(e_t) + \theta_t$, where θ_t follows the more general linear process $\theta_t = \sum_{s\geq 1} \gamma_s \theta_{t-s} + \Gamma_t + z_t$, for some coefficients γ_s. The summation allows for rich intertemporal correlations; while $\Gamma_t + z_t$ represents the period t innovation, with the assumption that z_t is i.i.d. across individuals with density g, and Γ_t is an aggregate shock common to all individuals. No restrictions are imposed over the distribution of Γ_t. This recursive process accommodates any invertible ARMA process in which both the random innovations $\Gamma_t + z_t$, and the levels of output y_t, are lagged for an arbitrary number of periods.

Under this linear structure, the model remains highly tractable. Consider a symmetric environment where every agent is equipped with the same happiness functions. Combined with the linear structure, this implies that the equilibrium efforts are constant across periods and agents. Consequently, output for period t, as a function of e_t, can be expressed as

$$\begin{aligned} y_t &= \pi(e_t) + \sum_{s\geq 1} \gamma_s \theta_{t-s} + \Gamma_t + z_t \\ &= \pi(e_t) + \sum_{s\geq 1} \gamma_s [y_{t-s} - \pi(e^*_{t-s})] + \left\{ \overline{y}_t - \pi(e^*_t) - \sum_{s\geq 1} \gamma_s [\overline{y}_{t-s} - \pi(e^*_{t-s})] \right\} + z_t \\ &= \overline{y}_t + \sum_{s\geq 1} \gamma_s [y_{t-s} - \overline{y}_{t-s}] + \pi(e_t) - \pi(e^*_t) + z_t, \end{aligned}$$

where the term in curly brackets corresponds to Γ_t. Following the representation in (7), it follows that $E[y_t \mid e_t, \Omega_t] = \overline{y}_t + \sum_{s\geq 1} \gamma_s [y_{t-s} - \overline{y}_{t-s}] + \pi(e_t) - \pi(e^*_t)$, with $\Omega_t = \langle y_{t-1}, y_{t-2}, ...; \overline{y}_t, \overline{y}_{t-1}, ... \rangle$. The efficient reference point therefore becomes $\widehat{y}_t = E[y_t \mid e^*_t, \Omega_t] = \overline{y}_t + \sum_{s\geq 1} \gamma_s [y_{t-s} - \overline{y}_{t-s}]$.

Theorem 1. *Under the above linear process for output with general intertemporal correlations, for every t, and every agent, the efficient happiness function is given* *by*

$$V(y_t, \Omega_t) = \begin{cases} 1 & if\ y_t \geq \overline{y}_t + \sum_{s\geq 1} \gamma_s [y_{t-s} - \overline{y}_{t-s}], \\ 0 & if\ y_t < \overline{y}_t + \sum_{s\geq 1} \gamma_s [y_{t-s} - \overline{y}_{t-s}]. \end{cases}$$

The carrier of happiness is now the generalized difference in differences

$$y_t - \sum_{s\geq 1} \gamma_s y_{t-s} - \left[\overline{y}_t - \sum_{s\geq 1} \gamma_s \overline{y}_{t-s} \right].$$

The lagged term $\sum_{s\geq 1} \gamma_s y_{t-s}$ corresponds to a habit, now a function of a more comprehensive past. This habit is subtracted in order to eliminate any component of y_t that is associated to past output as opposed to current effort e_t. Crucially, in this symmetric environment, any intertemporal correlation also impacts aggregate output. The lagged sum $\sum_{s\geq 1} \gamma_s \overline{y}_{t-s}$ is therefore subtracted from $\overline{y}_t$ as a means to enhance its informative value. This lag can again be interpreted as a process of habituation to peers. Moreover, since exactly the coefficients γ_s are employed, the general prediction is that any reference point that depends positively on past output as well as current aggregate output, must simultaneously depend negatively on past aggregate levels.

5. Habituation to Peers

Perhaps the most characteristic implication of the present agency approach to happiness is that habits are accompanied by peer effects lagged with a negative sign. Here we contrast this implication with the existing happiness evidence. We limit attention to income, a variable known to cause particularly fleeting effects. We proceed by interpreting the available evidence within a stylized formal model. In particular, suppose that the carrier of happiness at any given date is the difference between current income y_t and some time-varying reference point $\widehat{y}_t$. Suppose, in addition, that this reference point can be expressed as a linear function of past income together with the history of aggregate levels:

$$\widehat{y}_t = \sum_{s\geq 1} \alpha_s y_{t-s} + \lambda \overline{y}_t + \sum_{s\geq 1} \beta_s \overline{y}_{t-s}.$$

We impose no ex-ante assumptions over the coefficients α_s, λ, β_s. Rather, we wish to translate the available happiness evidence into simple restrictions over the values that these coefficients can take.

Consider the experiment of increasing the level of income of every individual in a given country by a constant level c for every period into the indefinite future. What is the long-term effect of this increase? Assuming this country had a sufficiently high level of income to begin with, a number of studies suggest that the change will have no impact at all –see Frederick and Loewenstein [1999] for a review. In the U.S., for example, real per-capita income increased by a factor of four between 1970 and 1990, but Myers and Diener [1996] find no change in average happiness or satisfaction. For poor countries, however, the existing studies do suggest a positive effect. Consequently, we interpret this evidence as an indication that the experiment will have a non-negative long-term effect over happiness. In terms of the coefficients of the linear model, this implies that over the long run, $\widehat{y}_t$ cannot increase by more than c, which delivers our first restriction:

$$(S1) \qquad \sum_{s\geq 1} \alpha_s + \lambda + \sum_{s\geq 1} \beta_s \leq 1.$$

Next, consider the experiment of increasing average output $\overline{y}_t$ by c units while leaving y_t intact. What is the short-term effect of this impact? It is well known that peer comparisons are a central component of income satisfaction –see Argyle [1999] for a review. These effects have been measured, for example, among co-workers, where the best predictor of income satisfaction is typically wage inequality. In some cases, individuals are even willing to receive a lower income as long as they receive more than a rival group. We interpret this effect as

$$(S2) \qquad \lambda > 0.$$

Finally, consider the experiment of increasing the income of a single individual by a constant level c for every period into the indefinite future. What is the long-term effect of this increase? Provided income is sufficiently high to begin with, the evidence again suggests that this change will have a minimal long-term effect –e.g., Easterlin [1974], Duncan [1975].[22] Following Brickman and Campbell [1971], this process has been called the "hedonic treadmill." In terms of the model, this means that over the long-run, $\widehat{y}_t$ will increase by approximately c units as well, delivering our third stylized restriction:

$$(S3) \qquad \sum_{s\geq 1} \alpha_s \approx 1.$$

[22]For example, Clark [1999] shows evidence that job satisfaction is tightly related to increases in salary, not absolute levels.

By combining $(S1) - (S3)$ we obtain

$$\sum_{s\geq 1} \beta_s \leq 1 - \sum_{s\geq 1} \alpha_s - \lambda \approx -\lambda < 0.$$

Thus, whenever λ is sufficiently different from zero, we must have $\sum_{s\geq 1} \beta_s < 0$, as predicted by the agency model.[23]

6. Concluding Remarks

Since so many decisions are motivated by emotional rewards, it has always been tempting to equate utility and happiness. These two notions, however, need not coincide. Following Kahneman [1999], we can conceptually distinguish "decision utility," which refers to utility in the standard choice-theoretical sense, from "experienced utility," which describes the hedonic experience. Our model includes a link between the two. Experienced utility corresponds to happiness, which is an emotional state that arises after the outcome of a project is realized –this outcome, for example, can represent the amount of fruit collected during a gathering project. Decision utility, on the other hand, describes a ranking among all the alternative physical courses of action in which the individual can engage during this project –for example, climbing a particular tree. These actions, per se, are not the carriers of happiness. The link between the two types of utility is the following, decision utility is constructed by computing an expected value of happiness conditional on each alternative course of action. This computation amounts to transforming a functional space of conditional distributions into a one-dimension measure: the aggregating device is happiness. Consider a modern example. Suppose an individual is offered two alternative shirts and must select one. Suppose the function of the shirt is to attract a mate in a party later that day. Following our framework, decision utility will rank the two shirts by means of ranking the more primitive actions of physically *grabbing* one shirt instead of the other. This ranking, in turn, will be determined by the expected value of happiness created by each shirt, where happiness is a function of the outcome of the project: whether or not the mate was secured.

Our objective is not to present a comprehensive description of the existing patterns and determinants of happiness. This would be a daunting task for any formal model. Rather, our intention is to set forth an agency-based mechanism that can account for some basic components of the hedonic experience. We propose, for

[23]It can even be argued that $\lambda \approx 1$, which is consistent with the evidence that, in wealthy countries, the process of economic growth has a minimal impact over happiness.

example, that habits and peer comparisons enhance the statistical power of the happiness incentive device. This can explain, for instance, why peer income can adversely affect an individual's reference point and thus lower her satisfaction. Our claim, to be sure, is not that peer income will always reduce the level of happiness. An obvious example is when our interests happen to be aligned with those of our peers –such as the case of friendship and kinship. Indeed, the evolution of humans is widely believed to have been based on intense reciprocal relationships –formally, the project of helping a friend would be modeled in the same way as the project of gathering fruit. In the model, contemporaneous peer success always hurts the individual precisely because in the examples we consider their interests happen to be unrelated –while their output technology is not.

The modern study of happiness has been closely accompanied by the study of hedonic predictions, namely, the degree to which individuals hold accurate intuitive models of the determinants of their own future tastes and feelings. The consensus among psychologists, for example, is that the average individual has profound deficiencies when it comes to these intuitive models –e.g., Gilbert et al. [1998], and Wilson et al. [2000], see also Loewenstein and Schkade [1999] for a review. Examples include underestimating the extent of hedonic adaptation as well as biased perceptions regarding the long-term determinants of happiness. Given the inherent complexities of this type of phenomena, we have opted to model the individual as a collection of multiple selves, one for each period. Each self faces one project, and responds to her happiness function alone –importantly, these projects can be forward-looking themselves, such as helping a friend, or saving for future retirement. The result is a happiness function that can be expressed exclusively in terms of material outcomes, therefore avoiding the complexities, for example, of stating current feelings as a higher order-function of future feelings, which in turn can depend on subsequent feelings, and so on. In this sense, the model is silent with respect to the formation of hedonic predictions.

7. Appendix: A Dual Problem

We wish to describe a connection between problem (4) in section $2.A$, based on the selection of x, and the original problem (1), based on the selection of e. For this connection, we must place a restriction over the structure of X. Suppose the agent can invest resources in a continuum of projects indexed by $s \in [0, 1]$, and x is a function that maps $[0, 1]$ into $\mathbb{R}_+$. The value of $x(s)$ represents the level of resources invested in project s –e.g., time. In addition, the agent faces the budget constraint $\int x(s)ds \leq 1$, i.e., X equals the set of functions x for which this inequality holds.

We proceed by constructing a dual problem to (4). Suppose that the agent is hypothetically given control over the size of her resource budget. Denote the size of this budget by e, which determines the maximum size for $\int x(s)ds$. Denote the optimal course of action by $x^*(\alpha, e)$, which may depend on the size of the budget. Expected output is now given by $E[y \mid d, e]$, which is assumed to be increasing in e, and also decreasing in d. Finally, suppose the agent faces a marginal cost λ of expanding this budget, so that her problem becomes

$$\max_{x,e} \; E[V(y) \mid d, e] - \lambda e \quad (5)$$

$$s.t.$$

$$\int x(s)ds \leq e.$$

By selecting λ appropriately, problem (5) becomes equivalent to problem (4), i.e., when λ is equal to the Lagrange multiplier associated to the underlying constraint $e \leq 1$. Notice that λ will be positive whenever $E[V(y) \mid d, e]$ is increasing in $E[y \mid d, e]$, which we assume is the case –this condition is satisfied by all the happiness functions derived above.

Next, define $\pi(e) = E[y \mid d = 0, e]$, and replace the objective in (5) by $E[y \mid \pi(e)]$. As a result, the problem of selecting e becomes $\max_e \; E[V(y) \mid \pi(e)] - \lambda e$; which, after a normalization of λ, corresponds to problem (1). The conclusion is that problem (1) can be interpreted as a dual problem for the more primitive (4), where effort represents the total resource expenditure.

Although (5) and (1) mirror each other, the evolutionary problem based on (5) is not equivalent to the evolutionary problem based on (1) –namely, problem (2) above. The reason is that, in problem (2), the selection of V occurs while taking the shadow cost of effort as given; whereas, in problem (5), the implicit shadow cost of resource expenditures is itself a function of V. However, under the random computational errors ϵ assumed in section $2.A$, both approaches become equivalent.

References

[1] Argyle, M. [1999], "Causes and Correlates of Happiness," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 353-73.

[2] Becker, G.S. [1992], "Habits, Addictions, and Traditions," *Kyklos,* 45, pp. 327-46.

[3] Brickman, P., and D.T. Campbell [1971], "Hedonic Relativism and Planning the Good Society," in Apley, M.H. (ed.), *Adaptation-Level Theory: A Symposium*, New York: Academic Press, pp. 287-302.

[4] Brickman, P., D. Coates, and R. Janoff-Bulman [1978], "Lottery Winners and Accident Victims: Is Happiness Relative?," *Journal of Personality and Social Psychology,* 37, pp. 917-27.

[5] Brown, D.E. [1999], "Human Universals," in Wilson, R.A., and F.C. Keil, *The MIT Encyclopedia of the Cognitive Sciences*, Cambridge: MIT Press.

[6] Burnham, T., and Phelan, J. [2001], *Mean Genes: From Sex to Money to Food, Taming our Primal Instincts,* New York: Penguin Books.

[7] Clark, A.E. [1999], "Are Wages Habit-Forming? Evidence from Micro Data," *Journal of Economic Behavior and Organization*, 39, pp.179-200.

[8] Damasio, A. [1994], *Descartes' Error: Emotion, Reason, and the Human Brain*, New York: Putnam.

[9] Diener, E., and F. Fujita [1996], "Social Comparisons and Subjective Well-Being," in Buunk, B., and R. Gibbons (eds.), *Health, Coping, and Social Comparison*, Hillsdale, N.J.: Erlbaum.

[10] Duncan, O. [1975], "Does Money Buy Satisfaction?" *Social Indicators Research*, 2, pp. 267-74.

[11] Easterlin, R.A. [1974], "Does Economic Growth Improve the Human Lot? Some Empirical Evidence," in David, P.A., and M.W. Reder (eds.), *Nations and Households in Economic Growth*, New York: Academic Press, pp. 89-125.

[12] Frank, R.H. [1987], "If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?," *American Economic Review,* 77, 4, pp. 593-604.

[13] Frank, R.H. [1999], *Luxury Fever: Money and Happiness in an Era of Excess*, Princeton and Oxford: Princeton University Press.

[14] Frederick, S., and G. Loewenstein [1999], "Hedonic Adaptation," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 302-29.

[15] Frey, B.S., and A. Stutzer [2002], *Happiness and Economics*, Princeton and Oxford: Princeton University Press.

[16] 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, 3, pp. 617-38.

[17] Kahneman, D. [1999], "Objective Happiness," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 3-25.

[18] Kahneman, D. [2003], "A Psychological Perspective on Economics," *American Economic Review: Papers and Proceedings*, 93, 2, pp. 162-8.

[19] Kahneman, D., E. Diener, and N. Schwarz [1999], *Well-Being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation.

[20] Kahneman, D., and A. Tversky [1979], "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 2, pp. 263-92.

[21] Layard, R. [2003], "Happiness: Has Social Science a Clue?," Lionel Robbins Memorial Lectures 2002/03, London School of Economics.

[22] Levin, J. [2003], "Relational Incentive Contracts," *American Economic Review,* forthcoming.

[23] 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.

[24] March, J.G. [1978], "Bounded Rationality, Ambiguity, and the Engineering of Choice," *Bell Journal of Economics*, 9, 2, pp. 587-608.

[25] Myers, D.G., and E. Diener [1996], "The Pursuit of Happiness," *Scientific American* (May), pp. 54-56.

[26] Peterson, C. [1999], "Personal Control and Well-Being," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 288-301.

[27] Rabin, M. [1998], "Psychology and Economics," *Journal of Economic Literature*, 36, 1, pp. 11-46.

[28] Rodriguez-Giroñes, M.A., and A. Kacelnik [2001], "Relative Importance of Perceptual and Mnemonic Variance in Human Temporal Bisection," *Quarterly Journal of Experimental Psychology*.

[29] Robson, A.J. [2002], "Evolution and Human Nature," *Journal of Economic Perspectives*, 16, 2, pp. 89-106.

[30] Rogerson, W.P. [1985], "The First-Order Approach to Principal-Agent Problems," *Econometrica,* 53(6), pp. 1357-68.

[31] Samuelson, L., [2003], "Information-Based Relative Consumption Effects," *Econometrica*, forthcoming.

[32] Samuelson, L., and J. M. Swinkels [2002], "Information and the Evolution of the Utility Function," mimeo.

[33] Sapolsky, R. [1999], "The Physiology and Pathophysiology of Unhappiness," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 453-69.

[34] Schwarz, N., and F. Strack [1999], "Reports of Subjective Well-Being: Judgmental Processes and Their Methodological Implications," in Kahneman, D., E. Diener, and N. Schwarz [1999], pp. 61-84.

[35] Scitovsky, T. [1992], *The Joyless Economy*, Oxford: Oxford University Press.

[36] Vaughan, S.C. [2002], *Half-Full, Half-Empty: Understanding the Psychological Roots of Optimism*, New York: Harcourt.

[37] 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. 209-233.

[38] Wilson, T.D., T.P. Wheatley, J.M. Meyers, D.T. Gilbert, and D. Axsom [2000], "Focalism: A Source of Durability Bias in Affective Forecasting," *Journal of Personality and Social Psychology*, 78, 5, pp. 821-36.