A Model of the Effect of Affect on Economic Decision Making*

Benjamin E. Hermalin† University of California

Alice M. Isen‡ Cornell University

September 3, 2006

Abstract

The standard economic model of decision making assumes a decision maker's *current* emotional state has no impact on his or her decisions. Yet there is a large psychological literature that shows that current emotional state, in particular mild *positive affect*, has a significant effect on decision making, problem solving, and behavior. This paper offers a way to incorporate this insight from psychology into economic modelling. Moreover, this paper shows that this simple insight can parsimoniously explain a wide variety of behaviors.

Keywords: Affect, morale, emotion.

JEL numbers: B41, D99, C70, C73, D81.

*Financial support from the Willis H. Booth Professorship in Banking & Finance is grateful acknowledged. The authors also thank Ulrike Malmendier, Barb Mellers, Matt Rabin, Bernard Sinclair-Desgagnes, Miguel Villas-Boas, and seminar participants at Stanford, UC Berkeley, UC Davis, USC, and Vanderbilt University for helpful comments.

†Contact information. Phone: (510) 642-7575. E-mail: `hermalin@haas.berkeley.edu`. Full address: University of California • Walter A. Haas School of Business • 545 Student Services Bldg. #1900 • Berkeley, CA 94720-1900.

‡Contact information. Phone: (607) 255-4687. E-mail: `ami4@cornell.edu`. Full address: Johnson Graduate School of Management • 359 Sage Hall • Cornell University • Ithaca, NY 14853-6201.

Draft 9/3/2006 Version 2

1 Introduction

A moment's introspection will convince most people that their decisions are influenced, in part, by their mood. For instance, the decisions we make when happy are not always the same as those we make when unhappy. In addition, there is a large psychological literature that finds a relationship between *affect*—what non-psychologists might call mood, emotions, or feelings—and decision making and behavior.[1] In particular, experiments have shown that relatively small changes (induced, *e.g.,* by finding small amounts of money, receiving a small gift, or watching five minutes of a non-sexual, non-aggressive funny movie) in positive affect or happy feelings, what an economist might call utility, can markedly influence everyday thought processes and that such influence is a common occurrence (Isen, 2000). Economic modeling of decision making and game playing has, until recently, essentially ignored this role of affect.[2] We propose a model that incorporates these psychological findings. Further, we show that the model has the potential to explain a wide variety of decisions and observed behaviors that can be difficult to explain under the standard economic paradigm. Moreover, we show this can be done within a single, simple framework that maintains the basic assumption of rationality. For instance, our framework provides insights into the following (among other behaviors):

- Persistence of mood.
- Benefits of a temporary elevation of mood even if there is an eventual return to a steady state.
- Increased effort from increases in signing bonuses and other *non*-continent rewards; and, conversely, why employers may prefer to fire workers than inflict across-the-board pay cuts during recessions.
- The apparent paradox of people not pursuing behaviors correlated with well-being.

The key element to our model is the following observation: Affect at the beginning of a period influences preferences, preferences, in turn, determine decisions, which modify the affective state at the end of the

[1] See, for example, Damasio (1995) on the role of emotions on behavior and evidence from how injuries to the areas of the brain responsible for emotions affects decision making. For instance, subjects with damage to the ventromedial frontal cortex make different decisions when gambling than normal subjects (Bechara et al., 1994). Loewenstein (1996, 2000) reports on how "visceral" influences—hunger, thirst, sexual desire, moods, emotions, and physical pain—affect behavior. For a survey on experimental manipulations of positive affect and consequent effects on decision making see Isen (2000). Slovic et al. (2002) provide another survey.

[2] Some recent economic work in this area includes Loewenstein (1996, 2000), Kliger and Levy (2003), and Loewenstein and O'Donaghue (2005).

period, which then becomes the relevant affect at the beginning of the next period, and so on. In other words, if $\mathbf{u}_t$ denotes affect (possibly a multi-dimensional variable) at the *end* of period t and $\mathbf{x}_t$ denotes a vector of decisions made in period t, then we have the dynamic:

$$\mathbf{u}_t = U\left(\mathbf{x}_t, \mathbf{u}_{t-1}\right),$$

where U is a function that recognizes that period-t affect is determined, in part, by affect at the beginning of the period (*i.e.,* at $t-1$), as well as current decisions. As we will show, primarily through examples, such dynamics can explain interesting aspects of people's behavior. In particular, we can analyze how utility varies over time and, in particular, whether it tends towards a long-run steady state.

This dynamic is reminiscent of the literature on rational addiction (see, *e.g.,* Pollak, 1970, Becker and Murphy, 1988, and, for a survey, von Auer, 1998). We explore the relation between our dynamic model and this literature in depth in Appendix A. Briefly, the two approaches differ insofar, as here, the state variable is utility itself, whereas in the addiction literature it is the amount of past consumption of the addictive substance. In addiction models, consumption (today's decision) affects current utility directly and the "stock" of future past consumption, whereas in our model today's decision affects current utility only; consequently, different dynamics emerge.[3] Put somewhat differently, in our model, present utility is a sufficient statistic for predicting future utility, whereas this is not true in the rational-addiction models (in those models, *e.g.,* of two equally unhappy people, only one may consume heroin today because only he has consumed it in the past).

To give our model predictive power, we make the following assumption (see expression (4) below): If one is free to maximize utility in any given period, then her *maximized* utility in that period is greater when she starts the period in a better mood than when she starts it in a worse mood. This assumption, while compelling to us, nevertheless warrants some discussion up front. First, note we are *not* assuming that utility is taking a monotonic path; indeed, as we demonstrate, utility could always converge back towards a fixed steady-state level regardless of whether it was initially perturbed up or down; that is, along the *equilibrium* path, utility could be falling over time. We are not even assuming that current utility is an increasing

[3]Other related work on evolving preferences, such as Benhabib and Day (1981), also relate today's preferences to yesterday's choices rather than yesterday's utility or affect.

function of past utility; in our model, that property holds only if the individual is pursuing maximizing behavior. Some critics of our approach have argued that, because of reference-point effects, "current utility is a decreasing, not increasing, function of past utility." This view seems to us at odds with results showing that people seek to maintain moods, that the recollection of pleasant past experiences can boost current mood (see, *e.g.,* Wegener and Petty, 1994), and that positive affect appears able to build resources and facilitate coping, problem-solving ability, and smooth social interaction (see *e.g.,* Aspinwall, 1998; Erez and Isen, 2002; Fredrickson and Joiner, 2002; Gervey et al., in press; and Isen, 2000). It also seems to confuse serial correlation along the equilibrium path with causation: The fact that high utility at $t-1$ tends to be followed by lower utility at time t in equilibrium does not mean that high utility is *causing* lower utility; rather, it could simply reflect that, in equilibrium, it is not possible to maintain as high a level of utility, yet it could still be true that the individual in question is better off at time t than she would have been in the counter-factual world in which her utility had been lower at $t-1$. Indeed, many of the dynamics we consider have, as an equilibrium property, the feature that relatively high (low) utility in one period is followed by relatively lower (higher) utility in the next. Nevertheless, it worth noting that our approach is not dependent on the monotonicity assumption (expression (4) below), a point we develop in Appendix B.

The idea that decisions in one period can affect well-being in future periods is a well-known one in economics. The most common formulation of this is in consumption-savings models, where increasing the level of consumption today reduces possible consumption tomorrow. In much of our analysis, however, such "choice-set" effects are absent—the choice set is taken to be constant over time in our models because a time-invariant choice set makes more straightforward what the role of affect is. This is not to say that our model wouldn't apply when the choice set varies over time and we briefly consider the effect of affect on a consumption-savings model in Section 3. Among the results we identify are a tendency for affect to push decision makers to front load consumption, because affective state can serve as a store of "wealth." But at the same time, we show that decision makers may wish to postpone consumption if they anticipate consumption will be more enjoyable given future affective states. In other words, in some models, it is possible that decision makers prefer to save relatively more *on* "rainy days" (days in which their affective state is low) rather than *for* rainy days.

We present the basic model in the next section. There much of our focus on two examples that illustrate some of the potential dynamics in the model. In Section 3, we incorporate a simple consumption-savings model. In Section 4 we discuss how our model relates to some other literature in economics and is consistent with results from psychology. We conclude in Section 5.

2 A Model of Affect & Decision Making

Suppose that an individual's objective is to maximize her discounted utility flow,

$$\sum_{t=1}^{T} \omega_t V_t \left(\mathbf{x}_t, \mathbf{u}_t\right) \; ; \tag{1}$$

where $T \leq \infty$, V_t is her utility function at time t, and $\omega_t > 0$ is the weight that today she assigns utility at t. The weight ω_t could accord with traditional models of discounting (*i.e.,* $\omega_t = \delta^t$, where $\delta \in (0, 1)$ is the discount factor); or it could accord with hyperbolic discounting (*e.g.,* Laibson, 1997, in which $\omega_t = \gamma\delta^t$, $0 < \gamma < 1$ if $t > 1$ and 1 otherwise); or some other specification appropriate to the situation. When $T = \infty$, the $\{\omega_t\}_{t=1}^{\infty}$ are assumed to be such that expression (1) is finite. Observe that the utility function $V_t(\cdot, \cdot)$ takes as an argument the person's affective state, $\mathbf{u}_t \in \mathbb{R}^n$, as well as her actions, $\mathbf{x}_t \in \mathcal{X}$, where $\mathbf{x}_t$ is a vector and $\mathcal{X}$ is the time-invariant feasible set (*e.g.,* $\mathbf{x}_t$ is an allocation of the hours in a day, $\mathcal{X}$, at date t). We further assume that affect, $\mathbf{u}_t$, is determined by the function $\mathbf{u}_t = U_t(\mathbf{x}_t, \mathbf{u}_{t-1})$; that is, current mood or affect is a function of previous mood or affect, as well as decisions made. Writing $\Upsilon_t(\mathbf{x}, \mathbf{u})$ for $V_t\big(\mathbf{x}, U_t(\mathbf{x}, \mathbf{u})\big)$, we can express her objective as

$$\max_{\{\mathbf{x}_t\}_{t=1}^{T}} \sum_{t=1}^{T} \omega_t \Upsilon_t \left(\mathbf{x}_t, \mathbf{u}_{t-1}\right) \text{ subject to } \mathbf{u}_t = U_t(\mathbf{x}_t, \mathbf{u}_{t-1}) \, . \tag{2}$$

The model in (2) is fairly general and permits a wide range of analyses based on various assumptions about U_t and Υ_t. In this paper, however, we limit our attention to a specific set of assumptions based on experimental evidence about *positive* affect (see, *e.g.,* Isen, 2000, and Slovic et al., 2002, for surveys). If we consider (i) positive affect to be a single dimensional construct of which one can have more or less (*i.e.,* it is meaningful to say someone enjoys more positive affect relative to a prior state or less positive affect relative to a prior state); and (ii) we take positive affect to be synonymous with utility (or monotonically related to it); then there is no further loss of generality in treating Υ_t as the mapping $\Upsilon_t(\mathbf{x}, u_{t-1}) = u_t$. We can, therefore, rewrite her objective (2) as

$$\max_{\{\mathbf{x}_t\}_{t=1}^T} \sum_{t=1}^{T} \omega_t U_t\left(\mathbf{x}_t, u_{t-1}\right) .$$

The basic behavioral implication of this formulation is captured by the following proposition:

Proposition 1 *Assume that a solution, $\mathbf{x}_t^*(u)$, exists for the program*

$$\max_{\mathbf{x}\in\mathcal{X}} U_t\left(\mathbf{x}, u\right) \tag{3}$$

for all possible u. Assume, too, that,[4]

$$\text{if } u > u', \text{ then } U_t\left[\mathbf{x}_t^*\left(u\right), u\right] > U_t\left[\mathbf{x}_t^*\left(u'\right), u'\right] \ . \tag{4}$$

Then the solution to

$$\max_{\{\mathbf{x}_t\}_{t=1}^T} \sum_{t=1}^{T} \omega_t U_t\left(\mathbf{x}_t, u_{t-1}\right) \tag{5}$$

is $\mathbf{x}_t = \mathbf{x}_t^*\left(u_{t-1}\right)$; that is, the discounted flow of utility is maximized by making the decisions that maximize each period's utility.

Proof: Since future utility is increasing in current utility and current decisions *directly* affect current utility only, maximizing current utility period by period must maximize (5). Hence, $\mathbf{x}_t^*\left(u_{t-1}\right)$ is the optimal decisions in period t. ■

In our general analysis, we will maintain the assumptions that a person can maximize present utility (*i.e.,* the program (3) has a solution for all u) and that the greater her initial utility, the greater she can make her end-of-period utility (*i.e.,* condition (4) holds). In the examples below, it is readily shown that these assumptions are met.

As a consequence of Proposition 1, utility is defined by the difference equation:

$$u_t = U_t\left[\mathbf{x}_t^*\left(u_{t-1}\right), u_{t-1}\right]. \tag{6}$$

Consider an example.

[4] Condition (4) is more general than the assumption that if $u > u'$, then $U_t\left(\mathbf{x}, u\right) > U_t\left(\mathbf{x}, u'\right)$. Observe that assumption would imply (4) by revealed preference:

$$U_t\left[\mathbf{x}_t^*(u), u\right] \geq U_t\left[\mathbf{x}_t^*(u'), u\right] > U_t\left[\mathbf{x}_t^*(u'), u'\right].$$

Example 1 [Creativity & Cooperation]: We assume now the decision is one-dimensional. Specifically, $x_t \in \mathbb{R}_+$. Let this choice denote some measure of work effort by an individual. It could, for instance, be some measure of help provided a co-worker; it could be a measure of creativity of thought; or it could just be some measure of effort. There is experimental evidence that positive affect can increase willingness to help (Isen and Levin, 1972); enhance creativity (Isen et al., 1987); and increase intrinsic motivation (Breckler, 1993; Erez and Isen, 2002; Isen and Reeve, in press; the article by Erez and Isen, in particular, shows that positive affect increases desire for the reward effort provides). These results can, in turn, be captured by assuming

$$U_t\left(x_t, u_{t-1}\right) = x_t\sqrt{2\beta_t} - \frac{x_t^2}{2u_{t-1}} + \tilde{u}_t, \tag{7}$$

where $\beta_t \in (0,1)$ is a parameter affecting the marginal utility of effort and $\tilde{u}_t$ is some, possibly time-varying, parameter. To keep the model from being pathological, assume $\tilde{u}_t \geq 0$ and $u_0 > 0$, where u_0 is the individual's time 0 utility. Note that we've chosen to model the effect of positive affect as a reduction in the marginal cost of x; we could, however, equivalently model it as enhancing the marginal benefit of x. In this example, $x_t^*\left(u_{t-1}\right) = u_{t-1}\sqrt{2\beta_t}$. Equation (6) becomes

$$u_t = \beta_t u_{t-1} + \tilde{u}_t. \tag{8}$$

Observe that utility is improving if $\tilde{u}_t > u_{t-1}(1-\beta_t)$, falling if $\tilde{u}_t < u_{t-1}(1-\beta_t)$, and unchanging if $\tilde{u}_t = u_{t-1}(1-\beta_t)$. Consequently, if $\tilde{u}_t$ and β_t are time-*in*variant constants, $\tilde{u}$ and β, then $\tilde{u}/(1-\beta)$ is the uniquely stable fixed point of this difference equation and utility is monotonically approaching $\tilde{u}/(1-\beta)$ over time. The solution to the difference equation (8) is

$$u_t = u_0 \prod_{\tau=1}^{t} \beta_\tau + \sum_{\tau=1}^{t} \left(\prod_{\sigma=\tau+1}^{t} \beta_\sigma \right) \tilde{u}_\tau \tag{9}$$

(Elaydi, 1996, p. 3 provides a proof; note the convention $\prod_{\sigma=t+1}^{t} \beta_\sigma = 1$). If β_t and $\tilde{u}_t$ are both time-invariant constants, then expression (9) simplifies to

$$u_t = u_0\beta^t + \tilde{u}\frac{1-\beta^t}{1-\beta}\,. \tag{10}$$

Observe, in Example 1, that x_t is an increasing function of u_{t-1}. Suppose that a second party (*e.g.,* an employer) has an interest in x_t (*e.g.,* it, as assumed, represents worker effort). A strategy, then, for this second party if she wishes to affect x_t is to manipulate the decision maker's utility or mood, u_{t-1}. From (9) above, two channels seem open to this second party. She can raise $\tilde{u}_t$ or she can boost affect at a single point in time (*e.g.,* boost u_0). The former could explain why a high *fixed* wage could have important incentive effects—by boosting affect in each period it could lead to greater effort. More interesting perhaps is the latter strategy, which could speak to a number of curious phenomena: For instance, why employers find signing bonuses (common to recruiting of MBAs and sports stars) and other non-performance-contingent gifts they give new employees to be in their interests. Both strategies (an initial gift to affect u_0 or sustained higher wages to affect $\tilde{u}_t$) would also be predicted by models of "gift-exchange" incentives (see, *e.g.,* Akerlof, 1982). Affect can, thus, be seen as an alternative mechanism through which gift giving can function as an incentive scheme.

The ability to incorporate time-varying parameters (*e.g.,* β_t and $\tilde{u}_t$ in Example 1) allows the model to capture a range of dynamics. For instance, the dynamics could be consistent with aspiration level theory (see Frey and Stutzer, 2002 for a short survey; but see new suggested modifications to the "hedonic treadmill" position by Diener, 2006). Suppose that the utility from income is $I\sqrt{2\beta} + \tilde{u}_t$, where I is income and $\beta \in (0, 1)$. Suppose, as assumed by aspiration level theory, that the level of income necessary to maintain a given level of utility is increasing over time; which is equivalent to assuming $\tilde{u}_t$ decreases over time. If the financial marginal return to effort, x_t, is \$1 (so $I_t = x_t$), then expression (7) represents a model in line with aspiration level theory. In particular if $u_0 < \tilde{u}_1/(1-\beta)$, then *equilibrium* utility will, at first, be increasing over time. This, in turn, will correspond to a period of time in which the decision maker works harder. At some point, however, if $\tilde{u}_t$ keeps falling, then she will become sufficiently demoralized by failing to meet her aspirations and her effort and utility will start to fall.

To be sure, the model just sketched is not the only way, or necessarily even the best way, to model aspiration level theory. In fact, some of the empirical findings that have led to aspiration level and adaptation level theory are consistent with a model in which both β_t and $\tilde{u}_t$ are time invariant (see Frey and Stutzer for a short review of the empirical evidence). A common finding from surveys of people's happiness over time is that their happiness stays fairly constant over time and that "additional material goods and services initially provide extra pleasure, but [the effect] is usually only transitory" (Frey and Stutzer, p. 414). From expression (10), it is clear that extra pleasure—that is, a boost in u_0—can have only a transitory effect, with utility eventually drifting back to the steady state level of $\tilde{u}/(1-\beta)$.

Random shocks can be readily handled by this framework. For example, for many people, sunny days boost mood, while grey days lower it (see Cunningham, 1979; Hirshleifer and Shumway, 2003 show this mood effect positively affects stock returns through its impact on trading behavior). Let $\tilde{u}_t = \bar{u} + \varepsilon_t$, where ε_t is a mean-zero random effect (*e.g.,* a measure of the deviation in sunshine during day t from the mean level of sunshine).[5] In this example, the logic of Proposition 1 still applies, so that the decision maker continues to do best by setting $x_t = u_{t-1}\sqrt{2\beta_t}$. From expression (9), utility in period t will be

$$u_t = u_0 \prod_{\tau=1}^{t} \beta_\sigma + \bar{u} \sum_{\tau=1}^{t} \prod_{\sigma=\tau+1}^{t} \beta_\sigma + \sum_{\tau=1}^{t} \left(\prod_{\sigma=\tau+1}^{t} \beta_\sigma \right) \varepsilon_\tau \, . \tag{11}$$

Given that the ε_t have zero expectation, this example shows that a deterministic model of affect change over time can be an unbiased predictor of affect change with random effects. This is not, however, a general result: As an alternative model of random shocks, suppose the β_t are distributed uniformly on $(0,1)$—for instance, the financial return to effort is now stochastic. So that the logic of Proposition 1 can be used, assume that the decision maker learns the realization of β_t *before* choosing x_t. Utility in period t is again given by expression (11), but a deterministic model (*i.e.,* with β set equal to $\mathbb{E}\{\beta_t\} = 1/2$) will only be an unbiased predictor if the β_t are uncorrelated with each other over time and if the β_t are uncorrelated with the ε_t.

As the previous paragraph suggests, some care must be taken in extending Proposition 1 when utility is subject to random shocks. To investigate further, let $\tilde{U}_t(\mathbf{x}_t, u_{t-1})$ denote the random variable that is utility at time t conditional on actions $\mathbf{x}_t$ and incoming utility level u_{t-1}. Let $U_t(\mathbf{x}_t, u_{t-1})$ denote a specific realization of this random variable.

Proposition 2 *Suppose (i) that uncertainty at time t is resolved* before *the decision maker chooses $\mathbf{x}_t$; that is, when choosing $\mathbf{x}_t$, the decision maker knows the realized function is $U_t(\cdot, u_{t-1})$. If a solution $\mathbf{x}_t^*(u)$ exists for the program* (3) *for all possible realizations $U_t(\cdot, u)$ and if condition* (4) *holds for each possible realization*

[5] Assume the support of ε_t is such that $\Pr\{u_t \leq 0\} = 0$.

$U_t(\cdot, u)$, then the solution to

$$\max_{\{\mathbf{x}_t\}_{t=1}^T} \mathbb{E}\left\{\sum_{t=1}^{T} \omega_t \tilde{U}_t(\mathbf{x}_t, u_{t-1})\right\} \tag{12}$$

is $\mathbf{x}_t = \mathbf{x}_t^*(u_{t-1})$; that is, the expected discounted flow of utility is maximized by making the decisions that maximize each period's utility.

Alternatively, suppose (ii) that uncertainty at time t is resolved after *the decision maker chooses $\mathbf{x}_t$. If, for all possible u, there exists a $\hat{\mathbf{x}}_t^*(u) \in \mathcal{X}$ such that*

$$\tilde{U}_t\left(\hat{\mathbf{x}}_t^*(u), u\right) \underset{\text{FSD}}{\geq} \tilde{U}_t\left(\mathbf{x}_t, u\right) \tag{13}$$

for all $\mathbf{x}_t \in \mathcal{X}$, where $\underset{\text{FSD}}{\geq}$ denotes ordering by first-order stochastic dominance and if

$$u > u' \text{ implies } \tilde{U}_t\left(\hat{\mathbf{x}}_t^*(u), u\right) \underset{\text{FSD}}{\geq} \tilde{U}_t\left(\hat{\mathbf{x}}_t^*(u'), u'\right), \tag{14}$$

then the solution to program (12) *is $\mathbf{x}_t = \hat{\mathbf{x}}_t^*(u_{t-1})$; that is, the expected discounted flow of utility is maximized by making the decisions that maximize each period's expected utility.*

Proof: Part (i) follows the same logic as the proof of Proposition 1.

Condition (14) asserts that the decision maker's expected utility in the future is greater the greater it is today; that is, future expected utility is an increasing function of utility today. Because it is an increasing function of utility today, then a first-order stochastic shift in utility today raises both expected utility today and expected utility in the future. ■

Proposition 2 shows that people should seek to boost their mood even if subject to random shocks and a tendency to drift to a fixed steady state.

If utility shocks are additive, then their timing *vis-à-vis* decision making is unimportant:

Corollary 1 *Suppose $\tilde{U}_t(\mathbf{x}, u) = \bar{U}_t(\mathbf{x}, u) + \varepsilon_t$, where $\bar{U}_t(\mathbf{x}, u)$ is non-stochastic and ε_t is stochastic. Suppose too that $\max_{\mathbf{x} \in \mathcal{X}} \bar{U}_t(\mathbf{x}, u)$ has a solution, $\mathbf{x}_t^*(u)$, for all possible u. Finally, suppose that*

$$u > u' \text{ implies } \bar{U}_t(\mathbf{x}_t^*(u), u) > \bar{U}_t(\mathbf{x}_t^*(u'), u').$$

Then, regardless of whether ε_t is realized prior or after the decision maker's choice of $\mathbf{x}$, the decision maker maximizes her expected discounted utility by choosing $\mathbf{x}_t = \mathbf{x}_t^*(u)$ in each period t.

Proof: The case in which ε_t is realized prior to the decision at time t follows immediately from part (i) of Proposition 2. For the case in which ε_t is realized after the decision at time t, observe that both conditions (13) and (14) are implied by the assumptions of the corollary. ■

To better appreciate the importance of the first-order stochastic order conditions consider the following example. The decision maker lives for two periods with $\omega_1 = \omega_2 = 1$ and her utility in period t is $I_t u_{t-1}^k$, with $u_0 = 1$ and $0 < k \leq 1$. In each period, she can participate in a gamble ($x_t = 1$) or decline to participate ($x_t = 0$). If she declines, she gets .9 for sure. If she participates, she gets 2 with probability 1/2 and 0 with probability 1/2. Observe the gamble has a higher expected payoff than the sure thing, but it doesn't stochastically dominate the sure thing. Consequently, while participating maximizes expected utility in each period, she will nonetheless prefer not to participate in period 1 if $k < .725$. That is, depending on the risk consequences for *future* utility, taking a gamble today that maximizes today's expected utility could be non-optimal for lifetime expected utility.

A comprehensive mathematical analysis of risk taking in this framework is a topic we leave for future research. Nonetheless, as we will discuss shortly, even a deterministic approach can provide qualitative insights about the relation between positive affect and risk taking.

For reasons similar to those that explain attitudes towards risk, positive affect can enhance self-control under some circumstances. Empirical evidence is now accumulating that people in positive affect perform a work task that needs to be done, even when there are more pleasant tasks on which they could choose to spend their time (*e.g.,* Isen and Reeve, in press), and that people in whom positive affect has been induced are better able to persevere on a task in a situation where those in the control condition do not continue working (*e.g.,* Wan and Sternthal, 2005). This empirical evidence motivates the following example, which has strikingly different dynamics than Example 1.

Example 2 [Socializing & Sobriety]: Again assume the decision is $x_t \in \mathbb{R}_+$. Let x_t denote energy or effort expended on some task. For instance, x_t could be effort at socializing with others; or energy spent keeping to a diet or staying sober; or some other effort similar to that considered in the previous example. Suppose that

$$U_t\left(x_t, u_{t-1}\right) = \phi\left(u_{t-1}\right) x_t - \frac{x_t^2}{2};$$

that is, here, utility at the beginning of a period modifies the marginal benefit of the action, x_t. Assume that $\phi(\cdot)$ is at least twice continuously differentiable, that $\phi(\cdot) \geq 0$, and that $\phi'(\cdot) > 0$. These assumptions reflect the idea that socializing is more pleasurable the happier one is, that a positive mood makes it more rewarding to keep to a diet or stay sober, or the other behavioral evidence cited in Example 1. Observe that $x^*\left(u_{t-1}\right) = \phi\left(u_{t-1}\right)$. Hence, equation (6) becomes

$$u_t = \frac{\phi\left(u_{t-1}\right)^2}{2}. \tag{15}$$

Unless $\phi(u) \propto \sqrt{u}$, this is a nonlinear difference equation. The second derivative of the right-hand side of (15), as a function of u_{t-1}, is

$$\left[\phi'\left(u_{t-1}\right)\right]^2 + \phi\left(u_{t-1}\right)\phi''\left(u_{t-1}\right).$$

From this it follows that if $\phi(\cdot)$ is strictly concave—improving initial utility has a bigger impact on future utility when initial utility is small than when it's large—then this difference equation can be convex for low values of u_{t-1} and concave for high values of u_{t-1}. In turn, this means it is possible that the right-hand side of (15), as a function of u_{t-1}, crosses the 45°-line three times (see Figure 1). This would be true, for instance, if

$$\phi(u) = \frac{\beta u}{u+1} \tag{16}$$

and $\beta > \sqrt{8}$. The three points of crossing would then be 0, $\frac{1}{4}\beta^2 - 1 - \frac{1}{4}\beta\sqrt{\beta^2 - 8}$, and $\frac{1}{4}\beta^2 - 1 + \frac{1}{4}\beta\sqrt{\beta^2 - 8}$ (*e.g.,* if $\beta = 3$, the points would be 0, $\frac{1}{2}$, and 2). Regardless of functional form, when (15) crosses the 45°-line three times, each point of crossing is a fixed point. Only the first, $\hat{u}_1$, and third, $\hat{u}_3$, however, are stable: To the left of the second, $\hat{u}_2$, the process converges toward $\hat{u}_1$ and to the right of $\hat{u}_2$, the process converges toward $\hat{u}_3$.

Some points illustrated by this example:

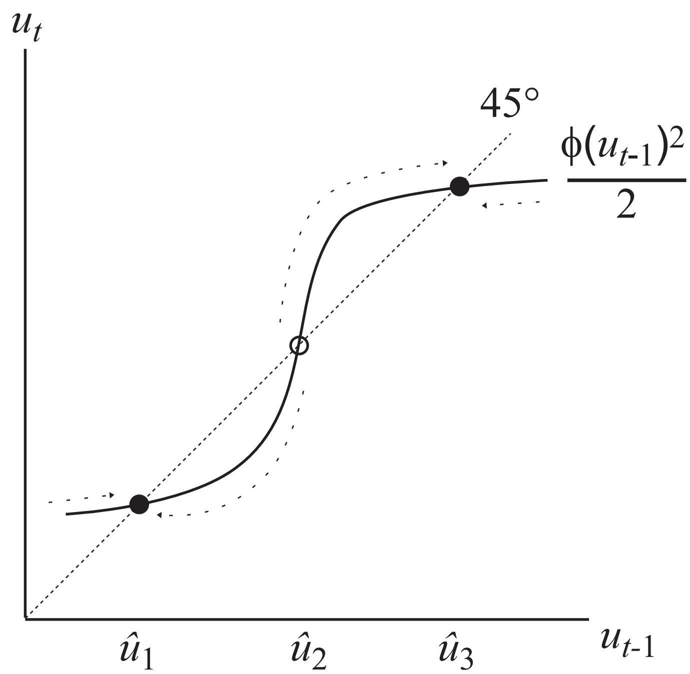

Figure 1: Possible relationship between u_{t-1} and u_t in Example 2.

- Small differences in initial utility (level of positive affect) can lead to large differences in future utility: Consider two individuals, one with initial utility $\hat{u}_2 - \varepsilon$ and one with initial utility $\hat{u}_2 + \varepsilon$ (assume $\hat{u}_3 - \hat{u}_2 > \varepsilon > \hat{u}_2 - \hat{u}_1$). The former's utility will be constantly decreasing, while the latter's will be constantly increasing. Correspondingly, there will be increasing differences in behavior.
- Differences in affective states across individuals will tend to be persistent all else being equal.[6] (This is not, however, to say that individuals' affective states don't fluctuate over time in response to exogenous conditions.)
- There will be a strong correlation between behavior and affect (*e.g.,* happy people socialize more), but the conventional causal inference will be wrong: People are *not* so much unhappy because they don't socialize or fail to keep to a diet or drink too much, rather they behave in these ways because they

[6] Barring, of course, external shocks, such as those resulting from actions of others. For instance, harmful acts, even neglect, by others could be a shock to the dynamic system.

aren't happy. That is, although behavior affects affect, affect affects behavior and it is not, therefore, always possible to modify affect in a desired way through behavior. It is even possible that unhappy people themselves confuse correlation for causation: Mistakenly declaring that they would be happier if they socialized more, kept to their diets, stayed off the bottle, etc.

- Recall that our decision makers are behaving optimally, there is *no* way for them to modify their behavior to achieve a better utility time path.[7] Hence, it would be wrong to blame the unhappiness of the "recluse" or failed dieter on his or her lack of effort or will power; and it would seem wrong, as well, to blame it on irrational behavior.

In terms of policy, this suggests that in an employment situation (*i.e.,* one in which x is a measure of effort), the firm wants to identify workers with initial utility (happiness, attitude, etc.) greater than $\hat{u}_2$ or induce such a utility initially (*e.g.,* by giving a signing bonus). Moreover, it wants to insure that positive affect is not dispelled. This could help to explain the use of layoffs over wage cuts in response to recessions. Even a short-term wage cut—a large negative shock—can have lasting consequences for morale (utility), with long-term negative consequences for the employer. That is, the employer can suffer the consequences of a wage cut long after the recession ends. Consequently, the employer could do better firing some workers in order to keep the wages and, thus, morale, of the remaining workers at a sufficient level during a recession. Bewley (1999) offers empirical evidence in support of the idea that employers' concerns about continuing worker morale lead them to fire some workers rather than cut the wages of all workers in a recession.

In other contexts, a policy prescription for the recluse or the failed dieter might be for *outside* intervention to directly improve utility rather than to focus on the deficient behavior. For instance, pharmacological or other intervention might be beneficial by directly enhancing mood (*e.g.,* by affecting the amount of a neurotransmitter like dopamine or serotonin). In extreme cases physicians may prescribe a mood-elevating drug, and once u_t gets above $\hat{u}_2$, the pharmacological intervention could be discontinued.[8]

[7] Since we're considering a single-dimensional choice set of actions, we're abstracting from the possibility of actions on other dimensions (*e.g.,* going to an enjoyable film or giving oneself a treat to self-induce positive affect). But the point carries over to a vector of activities. That is, "happy" people could tend to choose the vector $\mathbf{x}$, while "unhappy" people tend to choose the vector $\hat{\mathbf{x}}$, $\hat{\mathbf{x}} \neq \mathbf{x}$. Yet this difference is not the *cause* of happiness or unhappiness, but merely a correlate.

[8] This is roughly consistent with actual medical practice. Informal discussions with physicians indicate that "accepted practice" for first-time treatment with selective serotonin reuptake inhibitors (SSRIs) is to put someone on them for 6 to 12

In the context of Example 2, we see that downside shocks and upside shocks can be quite different. Depending on where an individual starts, a downside shock can cause her to eventually approach $\hat{u}_1$ in Figure 1, while an upside shock can cause her to eventually approach $\hat{u}_3$. For instance, consider decision making over gambles. There is evidence that individuals in whom positive affect has been induced behave in a more risk-averse fashion than a control group (*i.e.,* those in a "neutral" affective state; see Isen et al., 1988, for details). This is not surprising if the dynamic process resembles the one in Figure 1: An individual above the unstable fixed point $\hat{u}_2$ faces a dire downside risk—she could get switched from trending up toward $\hat{u}_3$ to trending down toward $\hat{u}_1$—versus a modest upside potential. In contrast, an individual below $\hat{u}_2$ faces a sizeable upside potential—switching from trending down toward $\hat{u}_1$ to trending up toward $\hat{u}_3$—versus a modest downside risk. Alternatively, under the "happy workers are better workers" idea, a firm might set wages so that workers' well-being is trending toward $\hat{u}_3$. This would, then, help to explain Bewley's findings that employers perceive wage increases as having a modest positive effect on morale, while they perceive wage cuts as disastrous: A wage increase of a given amount will raise long-run utility a bit above $\hat{u}_3$, while a similar size wage cut could plunge workers toward $\hat{u}_1$. We are not, however, claiming that *all* the dynamic processes considered here will necessarily exhibit this asymmetry between upside and downside. Rather our point is that a model of decision making that is sensitive to the role of affect can provide new insights into decision making under uncertainty and asymmetric reactions to policy changes (*e.g.,* wage increases versus wage cuts).

3 A Model of Savings

In the analysis above, the choice set each period, $\mathcal{X}$, was time-invariant. In particular, an action at time t could not affect the choice set at a later date $t + \tau$. In this section, we briefly depart from that assumption by considering a standard consumption-savings model.

To keep matters straightforward, assume the decision maker is initially endowed with an amount $Y \in \mathbb{R}_+$ of a good, which she consumes over a finite number of periods, T. Let $x_t \in \mathbb{R}_+$ denote consumption in period

months and, then, wean him or her off them. For many patients this is sufficient (*i.e.,* u_t is now greater than $\hat{u}_2$), and future medication is not necessary. Other patients cycle on and off them, suggesting that their brain chemistry or life experiences are such that they are periodically and randomly thrown well below $\hat{u}_2$, necessitating intervention to escape. Of course, intervention by others (*e.g.,* taking the person to an enjoyable film or giving her a treat) can also be effective in many cases.

t. For convenience, assume a constant decay rate on the unconsumed portion of the good that equals the decision maker's constant discount rate (*i.e.,* $\omega_t = \delta$ and the interest rate, r, is such that $\delta = 1/(1+r)$). As usual, $0 < \delta \leq 1$.

The decision maker chooses $\{x_t\}_{t=1}^T$ to maximize

$$\sum_{t=1}^{T} \delta^t U_t\left(x_t, u_{t-1}\right) \tag{17}$$

subject to

$$Y \geq \sum_{t=1}^{T} \delta^t x_t\,. \tag{18}$$

The first-order conditions are[9]

$$\left(\delta^t + \sum_{\tau=t+1}^{T} \delta^\tau \prod_{\sigma=t+1}^{\tau} \frac{\partial U_\sigma(x_\sigma^*, u_{\sigma-1})}{\partial u}\right) \frac{\partial U_t(x_t^*, u_{t-1})}{\partial x} - \lambda\delta^t = 0\,, \tag{19}$$

where λ is the Lagrange multiplier on (18) (note the convention $\sum_{T+1}^{T} = 0$). If, as in the analysis of the previous section, the choice set were invariant across periods—*i.e.,* there were no binding budget constraint—then $\lambda = 0$ and (19) re-establishes Proposition 1: Maximizing the discounted flow of utility would again be solved by maximizing each period's utility. If affect didn't have an impact on preferences—*i.e.,* if $\partial U/\partial u \equiv 0$—and if the per-period utility function, U_t, were constant across periods, then (19) reduces to the familiar result of equal consumption per period when the financial discount factor (decay rate) and the personal discount factor coincide.

Define

$$F_t = \sum_{\tau=t+1}^{T} \delta^{\tau-t} \prod_{\sigma=t+1}^{\tau} \frac{\partial U_\sigma(x_\sigma^*, u_{\sigma-1})}{\partial u}\,.$$

Then (19) can be reexpressed as

$$\big(1 + F_t\big)\frac{\partial U_t}{\partial x} - \lambda = 0\,,$$

where arguments have been suppressed for convenience. Observe that $\partial U_t/\partial x$ is greater the smaller is F_t and *vice versa*. If, as is typical of consumption models, we assume diminishing marginal utility of consumption (*i.e.,* $\partial^2 U_t/\partial x^2 < 0$), then x_t is greater, the larger the impact u_t has on future utility. This result can

[9] In what follows, we make the standard assumption that the constraints $x_t \geq 0$ are not binding. Allowing this assumption to be relaxed is not particularly interesting in this analysis.

be interpreted as follows: Increasing positive affect today boosts positive affect tomorrow; hence, one can invest for the future financially (*i.e.,* by consuming less today) or in terms of affect by raising positive affect today (*i.e.,* by consuming *more* today). In other words, affect is an alternative means of "storing wealth." Not surprisingly, when the financial returns are such that one would smooth consumption across time, the "affective returns" to consumption today will bias the decision maker toward earlier consumption over later consumption. In essence, positive affect can lead to what appears to be impatience or lack of control, but what is really the rational stockpiling of good feeling.[10]

The above argument, however, focuses on just one potential effect. There is also the effect of affect on the magnitude of $\partial U_t/\partial x$; that is, the cross-partial derivative, $\partial^2 U_t/(\partial x \partial u)$, is also important. If increases in initial affect, u_{t-1}, raise the marginal utility of consumption in that period, then the rational decision maker will wish to postpone consumption to those periods in which u_{t-1} can be expected to be large, since it is in those periods that she gets the biggest "bang for her buck." All else equal, if being happy raises the marginal enjoyment of consumption, then, essentially, decision makers should not save *for* rainy days, but rather save *on* rainy days (*i.e.,* consume less until she is in a better mood).[11]

These effects are most effectively illustrated by an example rather than a general model. To that end, suppose that $U(x,u) = (1+\alpha u)\ln(x)$, $\alpha \geq 0$, $\delta = 1$, and $T = 2$. Assume too that $1 + \alpha u_0 > 0$. Finally, assume that Y is large enough that $1 + \alpha u_1 > 0$ if the decision maker plays optimally. Note that $\alpha = 0$ is a standard, affectless, consumption model. From expression (19), the first-order conditions are

$$\frac{1+\alpha u_0}{x_1}\left(1 + \alpha \ln(x_2)\right) - \lambda = 0\,, \tag{20}$$

and

$$\frac{1+\alpha u_1}{x_2} - \lambda = 0\,. \tag{21}$$

[10]Note that there is an important difference between positive affect's *appearing* to lead to a lack of control and it actually leading to a lack of control. Recall actors in our model are behaving rationally. Moreover, as noted earlier, empirical evidence suggests that positive affect does not lead to a loss of self-control, dangerous risk-taking, impulsive behavior, or failure to consider important negative information when there is good reason to pay attention to it. There is convincing evidence that positive affect does not foster such behavior (*e.g.,* Aspinwall, 1998; Isen et al., 1988; and Trope and Pomerantz, 1998), even though, all else equal, people in positive affect will prefer to engage in and enjoy something pleasant, and do enjoy it more than do people in a neutral feeling state (*e.g.,* Erez and Isen, 2002; Isen and Reeve, in press).

[11]But, as previously noted, this does not imply impulsiveness or loss of self-control.

Observe that, for $\alpha > 0$, the marginal return to consumption in the first period tends to be greater, which pushes the decision maker to front load consumption to the first period. (If $\alpha = 0$, then consumption is the same across periods.) At the same time, however, u_0 also plays a role. For instance, if $u_0 = 0$, then (20) and (21) reveal that $x_1 = x_2$; the "front-loading" effect is completely offset by the "rainy-day" effect.

Proposition 3 *Suppose $U(x,u) = (1+\alpha u)\ln(x)$, $\alpha > 0$, $\delta = 1$, and there are two periods. Consumption is greater in the first period than the second if $u_0 > 0$; less in the first period than the second if $u_0 < 0$; and the same in the two periods if $u_0 = 0$.*

Proof: Substituting $Y - x_1$ for x_2, expressions (20) and (21) can be rewritten as

$$\frac{(1+\alpha u_0)\big(1+\alpha\ln(Y-x_1)\big)}{x_1} = \frac{1+\alpha(1+\alpha u_0)\ln(x_1)}{Y-x_1}\,. \tag{22}$$

Clearly, if $u_0 = 0$, then $x_1 = Y - x_1$. Moreover, if $x_1 = Y - x_1$, then (22) can hold only if $u_0 = 0$.

Consider, now, $u_0 \neq 0$. As just shown, it must therefore be that $x_1 - (Y - x_1) \neq 0$. Hence, (22) implies that the sign of $x_1 - (Y - x_1)$ is the same as the sign of

$$\big[(1+\alpha u_0)\big(1+\alpha\ln(Y-x_1)\big)\big] - \left[1+\alpha(1+\alpha u_0)\ln(x_1)\right]\,.$$

Rearranging, that has the same sign as

$$\alpha u_0 - \alpha(1+\alpha u_0)\left(\ln(x_1) - \ln(Y-x_1)\right)\,.$$

But this, in turn, means u_0 has the same sign as $x_1 - (Y - x_1)$. ■

4 Related Literature

As Elster (1998) points out, reference to moods and other emotions in economics is rare.[12] When such reference is made, it's usually to make sense of some behavior that seems inconsistent with narrow self interest. For example, honesty in situations where dishonesty would appear to have the larger expected payoff. In such cases, economists have "rationalized" honesty by appealing to the cost of guilt (see, *e.g.,*

[12] Laibson (2001), Loewenstein (2000), and Romer (2000) are some notable exceptions.

Becker, 1976, Frank, 1988, or Kandel and Lazear, 1992). Like our approach, these models consider the impact of actions *on* emotions. Decision making in these models, however, is affected by the *anticipation* of the emotional consequences only.[13] In contrast, our approach also considers how those emotions will affect decision making going forward. That is, for instance, a guilty person could behave differently from a person who doesn't feel guilty (in search of a sense of atonement, *e.g.,* the former may donate more to charity than the latter). In addition, we consider the consequences for behavior from improving mood as well as from a worsening of mood.

Two recent papers in economics, MacLeod (1996) and Kaufman (1999), have, like us, worried about the effect of emotional state on decision making. They are interested in modeling the adverse consequences of emotional state on cognitive abilities.[14] However, there is also evidence from psychology that affect, in particular positive affect, can enhance cognitive abilities, at least given minimal effort and interest in the task (see Isen, 2000, for a discussion). We are certainly sympathetic to the view that emotional state can affect cognitive ability.[15] Here, however, we ignore the possible effects of affect on cognitive ability and explore, as a baseline, how emotional state affects behavior, while remaining as close to the rational-actor paradigm as possible; that is, without postulating effects on cognitive ability.

Another strain of the economics literature focuses on "rationalizing" emotions; in particular, to explain why evolutionary forces may have produced them (see, *e.g.,* Frank, 1988, and Romer, 2000). Under the supposition, consistent with the fossil record on brain cases (see, *e.g.,* Johanson, 1996), that our homonid predecessors had less cognitive ability than we do, the case can be made that there was some advantage to

[13]See, for instance, Bell (1982) and Loomes and Sugden (1982) on the role of potential regret on decision making. Mellers (2000) and Mellers et al. (1999) also examine the role of anticipated emotions on decision making.

[14]MacLeod turns to emotions to justify his model of heuristic problem solving versus the standard optimization techniques that economists typically model decision makers as using. He argues, based on clinical observations of brain-damaged individuals reported in Damasio (1995), that people's heuristic problem-solving abilities are tied to their emotions. MacLeod does not, however, consider how *different* emotional states affect decisions, as we do. Kaufman, building on the well-known work of Yerkes and Dodson (1908) on arousal, suggests that emotional state can enhance or inhibit cognitive function: People who are completely uninterested in a problem or who are panicked over it are less able to solve it (or solve it less efficiently or effectively) than people exhibiting less extreme emotions.

[15]Ashby et al. (1999), for instance, note that the neurotransmitter dopamine is associated both with positive affect and the ability to perform cognitive tasks. Dopamine, thus, could serve as a biological explanation for a link between affect and cognition.

Relatedly, Erez and Isen (2002) find experimental evidence that positive affect increases the components of expectancy motivation: preferences (how well liked the reward is), but also perceptions (the estimated strength of the link between effort and reward, the likelihood of reward given a reasonable level of effort, etc.). Our paper concerns only the first effect, the impact on preferences.

"hardwiring" certain responses. For example, Romer notes that people (like rats) exhibit nausea aversion: If we suffer nausea—for whatever reason—within a short time after eating a particular food, we become averse to that food. For a species with limited cognitive ability, this would seem to be a good way to "learn" what foods are harmful. We, in contrast, do not seek to explain why people have emotions. We simply take the existence of emotions as given. Our question has been what do they influence when it comes to behavior?

Loewenstein (1996, 2000) and Loewenstein and O'Donaghue (2005) consider "visceral" influences on behavior. Although it doesn't fully capture the richness of these articles' analysis, these articles can be seen as studying decision making when preferences are given by $U_t(\mathbf{x}_t, \mathbf{a}_t)$, where $\mathbf{x}_t$ is the vector of decisions at time t and $\mathbf{a}_t$ is the vector of visceral factors (*e.g.,* hunger, thirst, pain, cravings, sexual desire; as well as emotions and moods). Unlike our approach, these articles do not consider how visceral states at time t could be a function of, among other factors, decisions at time $t-1$. Instead, they are more focused on how people may imperfectly forecast or even forecast with bias their future visceral states and, so, do not really know their preferences in the future. In contrast, we can be seen as making the hyper-rational assumption that people are good forecasters of future visceral states, but showing that nevertheless these visceral states continue to have important dynamic effects on decision making. As the pejorative "hyper-rational" suggests, we are not necessarily of the opinion that real people are perfect predictors of their future visceral states, but we do believe that the rationality assumption can be a useful approximation in this as in other aspects of economic modeling of decision making.[16]

5 Conclusions

In this paper, we have shown that incorporating psychological notions of affect can greatly enrich rational-actor models of decision making and strategic interaction. Although it's a modest change in our standard assumptions—yet one possessing strong empirical backing (*e.g.,* Isen, 2000)—it nevertheless gives insights into a number of behavioral phenomena:

- the persistence of mood, especially a happy mood;

[16]Indeed, as our discussion of uncertainty indicates, we don't require that our decision makers be perfect forecasters, although some aspects of the model may not carry over if they were biased forecasters.

- the use of *non*-contingent rewards to boost worker morale; and, conversely, the avoidance of negative "rewards," such wage cuts in response to recession;
- the apparent paradox of people not pursuing behaviors correlated with well-being;
- pharmacological treatment strategies for depression; and
- savings and consumption behavior.

Moreover, we suspect that we've only scratched the surface with respect to the economic applications. We can, for instance, foresee applications to issues of morale building within organizations, promotion of products, building of customer loyalty, relationship marketing, and policy issues and social welfare, among others.

In addition, the methodology outlined here can be extended to other emotions (we've focused on positive affect—happiness—because the experimental evidence gave us a clear guide as to the nature of the relationship between affect and behavior in that context). For instance, we might imagine that guilt affects behavior; perhaps according to a dynamic similar to

$$V(x_t, g_t) = -g_t\zeta\big(-x_t\big) \text{ and } U(x_t, g_{t-1}) = g_{t-1}\xi\big(x_t\big)\,,$$

where g_t is the level of guilt at time t, $\zeta'(\cdot) > 0$, $\xi'(\cdot) > 0$, and x_t is a guilt-inducing action ($x_t < 0$ is a guilt-reducing action).[17] For example, x_t could be the amount of money the individual spends on herself in period t (a negative x_t would, then, indicate spending on others). Although simple, this model does capture the ideas that guilt reduces utility and engaging in a guilt-inducing activity increases guilt going forward, but is pleasurable today. Other emotions could be similarly incorporated into models.

Both psychological theory, including evolution-based theorizing (*e.g.,* Johnston, 1999), and empirical work (*e.g.,* Isen, 2000, or LeDoux, 1998)—to say nothing of common sense—make it clear that behavior is affected by emotions. In addition, increasingly work in neuro-science (*e.g.,* Damasio, 1995, Ashby et al., 1999, and LeDoux, 1998) is working out the links and underpinnings between feelings and behavior. Further, clinical evidence from those who've suffered certain brain injuries demonstrates a rather suggestive set of

[17] The idea that guilt is, for some, a persistent emotion is borne out by numerous anecdotes of people who devote large portions of their lives seeking to atone for their misdeeds or the misdeeds of their family or people.

correlations between behavior and emotion. Given all this, it seems that economic modeling of behavior should pay attention to emotions. Otherwise, our models will be better suited to Mr. Spock and his fellow Vulcans than to *homo sapiens.* At the same time, one of the amazing attributes of our species is the ability to employ rational thought, including planning; our view is that human affect is a part of this system. Consequently, we've sought to develop a model that integrates these elements. Building on over 20 years of psychological research on positive affect, we've shown it's possible to build a model that reflects what we know about its role in decision making while maintaining the assumption of rationality. Moreover, we believe, we've shown that this combination offers real explanatory power with regard to real-life behavior.

Appendix A: Comparison of Affect with Rational Addiction

The rational addiction literature (see, *e.g.,* Pollak, 1970, Becker and Murphy, 1988, and, for a survey, von Auer, 1998) would seem related to the model of affect presented here. This appendix provides a brief comparison.

In a rational addiction model, utility at time t is given by the formula:

$$u_t = V\left(x_t, S_{t-1}\right) , \tag{23}$$

where x_t is time-t consumption of the addictive good and S_{t-1} is some "stock" of the addictive good consumed to the beginning of period t. The stock follows the process:

$$S_t = x_t + \rho S_{t-1} , \tag{24}$$

where $0 < \rho < 1$. The key behavioral assumption is that the cross-partial derivative of V is positive: A larger stock of the addictive good raises the marginal utility of consuming it today.

Observe that instead of working in terms of per-period consumption, one can work in terms of the stock: Using (24), (23) can be rewritten as

$$u_t = V\left(S_t - \rho S_{t-1}, S_{t-1}\right) . \tag{25}$$

Given an initial stock, S_0, the rational addict chooses $\{S_t\}_{t=1}^{\infty}$ to maximize the discounted utility flow of (25) subject to the constraint that $S_t \geq \rho S_{t-1}$ (*i.e.,* negative consumption is not permitted). For a constant discount factor (*i.e.,* $\omega_t = \delta^t$), the optimal S_t solves

$$V_1\left(S_t - \rho S_{t-1}, S_{t-1}\right) - \delta\rho V_1\left(S_{t+1} - \rho S_t, S_t\right) + \delta V_2\left(S_{t+1} - \rho S_t, S_t\right) = 0 \tag{26}$$

if the constraint doesn't bind and equals ρS_{t-1} otherwise.

The affect model also has a stock variable. But in the affect model it is utility itself. That is, loosely, in the affect model, expression (23), with u_{t-1} substituted for S_{t-1}, defines *both* the flow of utility and the transformation of the stock variable (*i.e.,* it is as if it combines (23) and (24) in one expression). It is this difference—combined with the fact the affect model requires no particular sign on the cross-partial derivative of V—that distinguishes these models. It also explains why the dynamic programming problem in the affect model is so much more straightforward than in the rational addiction literature.

Appendix B: Dynamics without Monotonicity

As discussed in the text, assumption (4) served to vastly simplify the dynamics in our model. While we believe it to be a reasonable assumption, we briefly consider, to be complete, what would happen were it not to hold. Let $\mathbf{x}_t^{**}(u_{t-1})$ denote the optimal decision rule. Observe, by the envelope theorem, that a change in x_τ, $\tau < t$ will not have an impact on future utility through $\mathbf{x}_t^{**}(u_{t-1})$. Hence, the $\mathbf{x}_t^{**}(u_{t-1})$ satisfy

$$\left(\omega_t + \sum_{s=t+1}^{T} \omega_s \prod_{n=t+1}^{s} \frac{\partial U_n\big(\mathbf{x}_n^{**}(u_{n-1}), u_{n-1}\big)}{\partial u}\right) \frac{\partial U_t\big(\mathbf{x}_t^{**}(u_{t-1}), u_{t-1}\big)}{\partial x} = 0\,. \tag{27}$$

We can, thus, conclude

Proposition 4 *Assume that a solution,* $\mathbf{x}_t^*(u)$, *exists for the program* (3) *for all possible* u. *Assume further that* $U_t(\cdot, u)$ *is strictly concave for all* u *and all* t. *Finally, assume that a solution* $\mathbf{x}_t^{**}(u_{t-1})$ *exists to the decision maker's dynamic programming and that solution satisfies*

$$\left(\omega_t + \sum_{s=t+1}^{T} \omega_s \prod_{n=t+1}^{s} \frac{\partial U_n\big(\mathbf{x}_n^{**}(u_{n-1}), u_{n-1}\big)}{\partial u}\right) > 0 \tag{28}$$

along the equilibrium path. Then $\mathbf{x}_t^{**}(u_{t-1}) = \mathbf{x}_t^*(u_{t-1})$*; that is, the optimal path involves maximizing per period utility.*

Proof: If (28) holds, then the first-order condition (27) can hold if and only if

$$\frac{\partial U_t\big(\mathbf{x}_t^{**}(u_{t-1}), u_{t-1}\big)}{\partial x} = 0\,.$$

But given that $U_t(\cdot, u)$ is strictly concave, this implies $\mathbf{x}_t^{**}(u_{t-1}) = \mathbf{x}_t^*(u_{t-1})$. Given that t was arbitrary, this completes the proof. ■

For example, suppose $T = \infty$ and $\omega_t = \delta^{t-1}$, where δ is a constant discount factor. Further assume $U_t(x, u) = x - ux^2/2$ (observe this utility function does *not* satisfy expression (4)). It is readily shown that $x_t^* = 1/u_0$ if t is odd and $x_t^* = 2u_0$ if t is even. Substituting these values into (28) we have that the sign of (28) is the same as the sign of

$$1 + \delta \frac{\delta - 2u_0^2}{1 - \delta^2}$$

if t is even and

$$1 + \delta \frac{2u_0^2\delta - 1}{2u_0^2(1 - \delta^2)}$$

if t is odd. Both expressions are positive if $\max\{\frac{1}{2u_0^2}, 2u_0^2\} < 1/\delta$. This condition would hold, for instance, if $u_0 = 1$ and $\delta < 1/2$. Observe that if this condition holds, then utility oscillates period to period between u_0 and $\frac{1}{2u_0}$. Similarly, x oscillates between $1/u_0$ and $2u_0$. Such a "yoyo"-path with regard to the choice of x seems at odds with observed behavior with regard to many decisions and is one reason for questioning models that do not satisfy condition (4).

References

Akerlof, George A., "Labor Contracts as Partial Gift Exchange," *Quarterly Journal of Economics*, November 1982, *97* (4), 543–569.

Ashby, F. Gregory, Alice M. Isen, and And U. Turken, "A Neuropsychological Theory of Positive Affect and its Influence on Cognition," *Psychological Review*, 1999, *106*, 529–550.

Aspinwall, Lisa G., "Rethinking the Role of Positive Affect in Self-regulation," *Motivation and Emotion*, March 1998, *22* (1), 1–32.

Bechara, Antoine, Antonio R. Damasio, Hanna Damasio, and Steven W. Anderson, "Insensitivity to Future Consequences Following Damage to Human Prefrontal Cortex," *Cognition*, 1994, *50* (1–3), 7–15.

Becker, Gary S., *The Economic Approach to Human Behavior*, Chicago: University of Chicago Press, 1976.

— **and Kevin M. Murphy**, "A Theory of Rational Addiction," *Journal of Political Economy*, August 1988, *96* (4), 675–700.

Bell, David E., "Regret in Decision Making under Uncertainty," *Operations Research*, Sep.-Oct. 1982, *30* (5), 961–981.

Benhabib, Jess and Richard H. Day, "Rational Choice and Erratic Behaviour," *Review of Economic Studies*, July 1981, *48* (3), 459–471.

Bewley, Truman F., *Why Wages Don't Fall During a Recession*, Cambridge, MA: Harvard University Press, 1999.

Breckler, Steven J., "Emotion and Attitude Change," in Michael Lewis and Jeannette M. Haviland, eds., *Handbook of Emotions*, New York: The Guilford Press, 1993.

Cunningham, Michael R., "Weath, Mood, and Helping Behavior: Quasi-Experiments in the Sunshine Samaritan," *Journal of Personality and Social Psychology*, 1979, *37*, 1947–1956.

Damasio, Antonio R., *Descartes' Error*, New York: Avon Books, 1995.

Diener, Ed, "Beyond the Hedonic Treadmill: Revising the Adaptation Theory of Well-being," *American Psychologist*, 2006, *61*, 305–314.

Elaydi, Saber N., *An Introduction to Difference Equations*, Berlin: Springer-Verlag, 1996.

Elster, Jon, "Emotions and Economic Theory," *Journal of Economic Literature*, March 1998, *36* (1), 47–74.

Erez, Amir and Alice M. Isen, "The Influence of Positive Affect on the Components of Expectancy Motivation," *Journal of Applied Psychology*, December 2002, *87* (6), 1055–1067.

Frank, Robert H., *Passions within Reason*, New York: Norton, 1988.

Fredrickson, Barbara L. and Thomas Joiner, "Positive Emotions Trigger Upward Spirals Toward Emotional Well-being," *Psychological Science*, 2002, *13*, 172–175.

Frey, Bruno S. and Alois Stutzer, "What Can Economists Learn from Happiness Research," *Journal of Economic Literature*, June 2002, *40* (2), 402–435.

Gervey, Ben, Eric R. Igou, and Yaacov Trope, "Positive Mood and Future-oriented Self-evaluation," *Motivation and Emotion*, in press.

Hirshleifer, David and Tyler Shumway, "Good Day Sunshine: Stock Returns and the Weather," *Journal of Finance*, June 2003, *58* (3), 1009–1032.

Isen, Alice M., "Positive Affect and Decision Making," in Michael Lewis and Jeannette M. Haviland, eds., *Handbook of Emotions*, 2nd ed., New York: The Guilford Press, 2000.

_ **and Johnmarshall Reeve**, "The Influence of Positive Affect on Intrinsic and Extrinisic Motivation: Facilitating Enjoyment of Play, Responsible Work Behavior, and Self-Control," *Motivation and Emotion*, in press.

_ **and Paula F. Levin**, "Effect of Feeling Good on Helping: Cookies and Kindness," *Journal of Personality and Social Psychology*, March 1972, *21* (3), 384–388.

_ **, Kimberly A. Daubman, and Gary P. Nowicki**, "Positive Affect Facilitates Creative Problem Solving," *Journal of Personality and Social Psychology*, June 1987, *52* (6), 1122–1131.

_ **, Thomas E. Nygren, and F. Gregory Ashby**, "Influence of Positive Affect on the Subjective Utility of Gains and Losses; It is Just not Worth the Risk," *Journal of Personality and Social Psychology*, 1988, *55* (5), 710–717.

Johanson, Donald, *From Lucy to Language*, New York: Simon & Schuster, 1996. With contributions by Blake Edgar.

Johnston, Victor S., *Why We Feel*, Reading, MA: Perseus Books, 1999.

Kandel, Eugene and Edward P. Lazear, "Peer Pressure and Partnerships," *Journal of Political Economy*, August 1992, *100* (4), 801–817.

Kaufman, Bruce E., "Emotional Arousal as a Source of Bounded Rationality," *Journal of Economic Behavior & Organization*, 1999, *38*, 135–144.

Kliger, Doron and Ori Levy, "Mood-induced Variation In Risk Preferences," *Journal of Economic Behavior & Organization*, 2003, *52*, 573–584.

Laibson, David I., "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, May 1997, *112* (2), 443–477.

_ , "A Cue-Theory of Consumption," *Quarterly Journal of Economics*, February 2001, *116* (1), 81–120.

LeDoux, Joseph, *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*, New York: Simon and Schuster, 1998.

Loewenstein, George, "Out of Control: Visceral Influences on Behavior," *Organizational Behavior and Human Decision Processes*, March 1996, *65* (3), 272–292.

_ , "Emotions in Economic Theory and Economic Behavior," *American Economic Review*, May 2000, *90* (2), 426–432.

_ **and Ted O'Donaghue**, "Animal Spirits: Affective and Deliberative Processes in Economic Behavior," May 2005. unpublished manuscript.

Loomes, Graham and Robert Sugden, "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty," *The Economic Journal*, December 1982, *92* (368), 805–824.

MacLeod, W. Bentley, "Decision, Contract, and Emotion: Some Economics for a Complex and Confusing World," 1996. Working paper, C.R.D.E., Université de Montréal.

Mellers, Barbara A., "Choice and the Relative Pleasure of Consequences," *Psychological Bulletin*, November 2000, *126* (6), 910–924.

_ , **Alan Schwartz, and Ilana Ritov**, "Emotion-Based Choice," *Journal of Experimental Psychology: General*, September 1999, *128* (3), 332–345.

Pollak, Robert A., "Habit Formation and Dynamic Demand Functions," *Journal of Political Economy*, July-August 1970, *78* (4, Part 1), 745–763.

Romer, Paul M., "Thinking and Feeling," *American Economic Review*, May 2000, *90* (2), 439–443.

Slovic, Paul, Melissa Finucane, Ellen Peters, and Donald G. MacGregor, "Rational Actors or Rational Fools? Implications of the Affect Heuristic for Behavioral Economics," *Journal of Socio-Economics*, 2002, *31*, 329–342.

Trope, Yaacov and Eva M. Pomerantz, "Resolving Conflicts Among Self-Evaluative Motives: Positive Experiences as a Resource for Overcoming Defensiveness," *Motivation and Emotion*, March 1998, *22* (1), 53–72.

von Auer, Ludwig, *Dynamic Preferences, Choice Mechanisms, and Welfare*, Berlin: Springer-Verlag, 1998.

Wan, Wen and Brian Sternthal, "Eliminating Depletion Effects," October 2005. Paper presented at the meeting of the Association for Consumer Research, San Antonio, TX.

Wegener, Duane T. and Richard E. Petty, "Mood Management Across Affective States: The Hedonic Contingency Hypothesis," *Journal of Personality and Social Psychology*, 1994, *66* (6), 1034–1048.

Yerkes, R.M. and J.D. Dodson, "The Relation of Strength of Stimulus to Rapidity of Habit-Formation," *Journal of Comparative Neurology and Psychology*, 1908, *18*, 459–482.