

The Effect of Affect on Economic and Strategic Decision Making*

Benjamin E. Hermalin[†]
University of California

Alice M. Isen[‡]
Cornell University

October 20, 2003

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 *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.

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. Nor is this merely an impression: There is a large psychological literature, based on experiments, that finds a relationship between *affect*—what non-psychologists might call mood, emotions, or feelings—and decision making and behavior (see Isen, 2000, for a survey). In particular, this research shows that relatively small changes 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. Economic modelling of decision making and game playing has, however, essentially ignored the role of affect. We seek to demonstrate that the addition of affect allows us to explain a wide variety of decisions and observed behaviors that are difficult to explain under the standard economic paradigm. Moreover, we show this can be done within a single, simple framework that maintains the assumption of rationality. For instance, our framework provides insights into the following (among other behaviors):

- Persistence of mood.
- Increased effort from increases in signing bonuses and other *non-contingent* rewards; and, conversely, why employers may prefer to fire workers than inflict across-the-board pay cuts during recessions.
- Attempts to demoralize opponents and build morale of allies in strategic interactions.
- The apparent paradox of people not pursuing behaviors correlated with well-being.
- Apparent cooperative play in finite games.

A common reaction by economists to the introduction of psychological insights into economics is that it means abandoning or relaxing the standard assumption of rationality. While it is true that many such attempts have had that flavor (see, *e.g.*, discussions in Lewin, 1996, Rabin, 1998, and Elster, 1998), our approach does *not*. In particular, the actors in our models are completely rational—they make their decisions to maximize the (discounted) value of their utility flow. What distinguishes our approach from the traditional model of decision making is that we assume that current positive affect or utility influences preferences going forward. As an example, experimental evidence (Isen and Levin, 1972) finds that positive affect tends to increase a person's willingness to aid others; that is, an increase in mood either increases an individual's pleasure from helping or lowers the psychic cost of helping.

In a one-shot setting, such a change in modelling assumptions would be difficult to distinguish from the more usual assumption of fixed preferences. Moreover, in a one-shot setting, why a person holds certain preferences over others

is not, generally, an interesting question in economics. On the other hand, if we consider dynamic settings, then affect becomes much more important: Affect at the beginning of a period influences preferences, which determine decisions, which modify the affective state at the end of the 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:

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

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 and their strategic decision making.

As Elster (1998) points out, reference to moods and other emotions in economics is rare.¹ 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.*, 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.² 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.³ 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

¹Laibson (2001), Loewenstein (2000), and Romer (2000) are some notable exceptions.

²See also Mellers (2000) and Mellers et al. (1999) for analyses of the role of anticipated emotions on decision making.

³MacLeod 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 solid, but preliminary, work in psychology (Yerkes and Dodson, 1908), 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.

sympathetic to the view that emotional state can affect cognitive ability,⁴ 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 hominid predecessors had less cognitive ability than we do, the case can be made that there was some advantage to “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 is what do they influence when it comes to behavior?

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.

When choice-set effects are absent, the intertemporal linkage in our model runs solely through affect. In particular, affect, \mathbf{u}_t , at time t is a sufficient statistic for predicting future affect levels. Among other implications, this means that *in our model* there is no *direct* effect of an individual’s past behaviors

⁴Ashby et al. (1999), for instance, note that the neurotransmitter dopamine is associated both with 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 high level of effort, etc.). Our paper concerns only the first effect, the impact on preferences.

(decisions) on her future behavior: If consumption paths $\{\mathbf{x}_\tau\}_{\tau=1}^t$ and $\{\mathbf{x}'_\tau\}_{\tau=1}^t$ both get the individual to affect level \mathbf{u}_t , then behavior thereafter will be the same. This, for instance, differs from the habit-formation literature (see, *e.g.*, §4.4 of von Auer, 1998, for a survey), in which present utility need *not* be a sufficient statistic for future behavior (*e.g.*, of two equally unhappy people, only one may consume heroin today because only he has consumed it in the past).⁵

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 explore how affect can affect the play of games. 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,

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

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 accord with hyperbolic discounting (*e.g.*, Laibson, 1997, in which $\omega_t = \gamma\delta^t$, $0 < \gamma < 1$); 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 $V_t(\mathbf{x}, U_t(\mathbf{x}, \mathbf{u}))$, we can express her objective as

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

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, for a survey). If we consider (i) positive affect to be a single dimensional construct; and (ii) positive affect to be synonymous with utility (or monotonically related to it); then there is no further loss of

⁵Appendix B more fully compares our model to the rational addition model and makes clearer the differences.

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

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

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*

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

for all possible u . Assume, too, that,⁶

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

Then the solution to

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

is $\mathbf{x}_t = \mathbf{x}_t^*(u_{t-1})$; 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, $\mathbf{x}_t^*(u_{t-1})$ 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.

Observe that, in equilibrium, current utility is an increasing function of past utility. This distinguishes our analysis from some related work by Benhabib and Day (1981), where, in equilibrium, U can be seen as a *non-monotonic* function of u_{t-1} . Although this non-monotonicity yields interesting dynamics—including possibly chaotic dynamics—Benhabib and Day don't offer what we see as a compelling behavioral justification for their utility function.⁷

⁶Condition (4) is more general than the assumption that if $u > u'$, then $U_t(\mathbf{x}, u) > U_t(\mathbf{x}, u')$. Observe that assumption would imply (4) by revealed preference:

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

⁷They suggest that an individual is choosing the amount of leisure to enjoy each period and that the greater the level of past leisure, the more leisure the individual desires today (*e.g.*, the less the individual worked last period, the more vacation she desires today; conversely, the harder she worked last period, the less vacation she desires today).

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

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

Consider an example.

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 (Isen and Reeve, 1992; Breckler, 1993; Erez and Isen, 2002, the last, in particular, shows that positive affect increases desire for the reward effort provides). These results can, in turn, be captured by assuming

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

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, $x_t^*(u_{t-1}) = u_{t-1} \sqrt{2\beta_t}$. Equation (6) becomes

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

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-invariant 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

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

(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

$$u_t = u_0 \beta^t + \tilde{u} \frac{1 - \beta^t}{1 - \beta}. \quad (10)$$

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). 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). 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).⁸ 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

⁸Assume the support of ε_t is such that $\Pr\{u_t \leq 0\} = 0$.

period t will be

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

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

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

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

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

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

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

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: Please see Appendix A.

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.

Now consider a second 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

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

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 $x^*(u_{t-1}) = \phi(u_{t-1})$. Hence, equation (6) becomes

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

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

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

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

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

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:

- 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.⁹
- 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

⁹Barring, 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.




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

they 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.¹⁰

- 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.¹¹ 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

¹⁰The idea that people might not understand why they're unhappy (or even what would make them happy) is not implausible: If people were expert at understanding their own psychology, why would there be any market for psycho-therapists?

¹¹Since 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 $\tilde{\mathbf{x}}$, $\tilde{\mathbf{x}} \neq \mathbf{x}$. Yet this difference is not the *cause* of happiness or unhappiness, but merely a correlate.

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 neuro-transmitter 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.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

¹²This 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 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.

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 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

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

subject to

$$Y \geq \sum_{t=1}^T \delta_t x_t. \quad (18)$$

The first-order conditions are¹³

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

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

¹³In 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.

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

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

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 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, affect can lead to what appears to be impatience or lack of control, but what is really the rational stockpiling of good feeling.

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.” That is, 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).

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

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

and

$$\frac{1 + \alpha u_1}{x_2} - \lambda = 0. \quad (21)$$

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

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

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

$$[(1 + \alpha u_0)(1 + \alpha \ln(Y - x_1))] - [1 + \alpha(1 + \alpha u_0) \ln(x_1)].$$

Rearranging, that has the same sign as

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

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

4 The Impact of Positive Affect on Strategic Interactions

We’ve so far considered only individual decision makers. In this section we extend our analysis to situations in which our decision makers interact strategically.¹⁴ Although there are many potential models to explore, we will consider only one: There are two decision makers (players), indexed by superscripts. A given player’s utility at the end of the t th period is assumed to be

$$u_t^i = x_t^i + \alpha^i x_t^j - \frac{(x_t^i)^2}{2u_{t-1}^i};$$

that is, each player i ’s utility is the same as in Example 1 (with $\beta = \frac{1}{2}$ and $\tilde{u} = 0$) but with the addition of the impact of the *other* player’s (j ’s) action on

¹⁴See Saarni (1993), *e.g.*, for a survey of the role of emotions in interpersonal interactions.

her utility. For simplicity, j 's externality on i is assumed to enter linearly. If $\alpha^i > 0$, then it's a positive externality. If $\alpha^i < 0$, then it's negative. It seems reasonable—at least in many contexts—to imagine that the externality is not too large relative to the direct effect. Consequently, we limit attention to models in which $|\alpha^i| \leq 1$ for $i = 1, 2$. We also return to the assumption that the choice sets for the players are time invariant and, moreover, are \mathbb{R}_+ for each.

To begin, consider the following *sequential* play game: Player 1 chooses her x in period 1, Player 2 chooses his x in period 2, and then the game ends. Given the finite time horizon, we may ignore discounting. Hence, the sums of the two periods' utilities are

$$\begin{aligned} U^1 &= \alpha^1 x_2^2 + x_1^1 - \frac{(x_1^1)^2}{2u_0^1} \text{ and} \\ U^2 &= x_2^2 - \frac{(x_2^2)^2}{2\alpha^2 x_1^1} + \alpha^2 x_1^1, \end{aligned}$$

respectively. So that Player 2 has a well-defined strategy, assume $\alpha^2 > 0$; Player 1 provides a positive externality for Player 2. Solving for the subgame perfect equilibrium, we see that Player 2's optimal play is

$$x_2^{2*}(x_1^1) = \alpha^2 x_1^1.$$

Hence, Player 1 maximizes

$$\alpha^1 \alpha^2 x_1^1 + x_1^1 - \frac{(x_1^1)^2}{2u_0^1}.$$

Solving yields

$$x_1^{1*}(u_0^1) = (1 + \alpha^1 \alpha^2) u_0^1.$$

Some remarks: Observe that Player 1's behavior when there's a Player 2 is different than when there is no Player 2. Specifically, she does more x when $\alpha^1 > 0$ (Player 2's action provides a positive externality for her) and she does less when $\alpha^1 < 0$ (Player 2's action imposes a negative externality on her). When she does more (*i.e.*, when $\alpha^1 > 0$), it might appear that she is “internalizing” the benefit she provides Player 2. This, however, would be an incorrect inference: She is playing selfishly, but she understands that doing more of something Player 2 likes benefits her because the happier is Player 2, the more Player 2 does of an action she likes. If Player 2 couldn't act or, equivalently, if $\alpha^1 = 0$, then Player 1 would cease to seem so generous. Consequently, although it can yield similar behavior, our model is different than an altruism model in which Player 2's utility would be an argument of Player 1's utility function. Note, too, that $x_2^{2*}(\cdot)$ is an increasing function. Particularly when $\alpha^1 > 0$, the resulting behavior of more x^1 leading to more x^2 could be viewed as Player 2 reciprocating Player 1's “generosity” or Player 2's behavior being governed by a fairness norm. Again, such a view would be, here, misleading: It need not be that Player 2 produces more x^2 because he has a preference for reciprocating or being fair—in fact, he

needn't possess such preferences at all—rather, here, he produces more because Player 1's action has changed his preferences such that he finds that he enjoys more of the action.

We are not arguing, however, that our approach is superior to assuming altruism or norms of fairness and reciprocity. Rather we are pointing out that behavior consistent with such motives need not be due to those motives (even in one-shot or finite-horizon games). Of course, these other approaches could be complementary to our approach: We could, for instance, assume that Player 2's utility is boosted from the “warm glow” that comes from knowing that Player 1 has done something extra for him (*e.g.*, we could assume

$$\hat{U}^2 = x_2^2 - \frac{(x_2^2)^2}{\kappa_0 + \kappa_1(x_1^1 - \bar{x}^1)} + \alpha^2 x_1^1,$$

where κ_0 and κ_1 are positive constants and \bar{x}^1 is the optimal level of x were Player 1 *not* playing with Player 2—here, u_0^1).

Consider, now, a two-period game with *simultaneous* moves. The players' utilities are, thus,

$$\begin{aligned} u_2^i &= \alpha^i x_2^j + x_2^i - \frac{(x_2^i)^2}{2u_1^i} \text{ and} \\ u_1^i &= \alpha^i x_1^j + x_1^i - \frac{(x_1^i)^2}{2u_0^i}. \end{aligned}$$

Solving for the subgame perfect equilibrium, we see that

$$x_2^{i*}(u_1^i) = u_1^i.$$

Hence, player i chooses x_1^i to maximize

$$\underbrace{\alpha^i \left(\alpha^j x_1^i + x_1^j - \frac{(x_1^j)^2}{2u_0^j} \right)}_{x_2^{j*}} + \frac{1}{2} \left(\alpha^i x_1^j + x_1^i - \frac{(x_1^i)^2}{2u_0^i} \right) + \underbrace{\alpha^i x_1^j + x_1^i - \frac{(x_1^i)^2}{2u_0^i}}_{u_1^i}. \quad (23)$$

Thus,

$$x_1^{i*}(u_0^i) = \left(1 + \frac{2}{3} \alpha^i \alpha^j \right) u_0^i.$$

For future reference, note that x_t^{i*} is a linear function of i 's $t-1$ utility only; in particular, the other state variable, u_{t-1}^j , is not relevant to her decision.

Define three scenarios:

Friendly game: $\alpha^1 > 0$ and $\alpha^2 > 0$;

Antagonistic game: $\alpha^1 < 0$ and $\alpha^2 < 0$; and

Mixed game: $\alpha^i \alpha^j \leq 0$, $i \neq j$.

Observe that, relative to a situation in which player i is alone, her choice of x_1^i is greater in a friendly or antagonistic game (since $\alpha^i \alpha^j > 0$), but smaller in a mixed game (since $\alpha^i \alpha^j < 0$). The intuition is straightforward: When $\alpha^i > 0$, player i wants to make player j happier, since that will yield a greater externality in the next period. Conversely, when $\alpha^i < 0$, she wants to reduce player j 's happiness, since that will yield less of the externality in the next period. Player i boosts (reduces) player j 's happiness by *increasing* x_1^1 when her action has a positive (negative) externality on j . She boosts (reduces) his happiness by *decreasing* x_1^1 when her action has a negative (positive) externality on him. This makes sense: Consider, for example, that, in many sporting events (examples of antagonistic games), the players seem to expend more energy or play with greater intensity in the first half than the second.¹⁵ Similarly, as an example of a friendly game, the home team tries to get its fans “into the game” early on and the fans tend to cheer a lot at the beginning (*e.g.*, during player introductions). Note that if just *one* player is unaffected by the other's action, then *neither* player deviates from what he or she would have done if playing alone. That is, the apparent concern about the other player in the first period is strategic and not inherent (an insight also borne out by the fact that the last-period action is equal to the playing-alone action for both players regardless of scenario).

Substituting the definitions of x_1^{i*} and x_1^{j*} into player i 's “lifetime” utility, expression (23), yields

$$U^i = \left[\alpha^i \alpha^j \left(1 + \frac{2}{3} \alpha^i \alpha^j \right) + \frac{3}{2} \theta \left(1 + \frac{2}{3} \alpha^i \alpha^j \right) \right] u_0^i + \alpha^i \left(1 + \frac{2}{3} \alpha^i \alpha^j \right) \left(2 - \frac{1}{3} \alpha^i \alpha^j \right) u_0^j,$$

where $\theta(\cdot) : q \rightarrow q - \frac{1}{2}q^2$. The sign of the last term equals the sign of α^i , since $\min 1 + \frac{2}{3}\alpha^i \alpha^j = \frac{1}{3}$ and $\min 2 - \frac{1}{3}\alpha^i \alpha^j = \frac{5}{3}$. Hence, when j 's action provides a positive externality for i , she prefers that j have high initial utility. Conversely, when j 's action imposes a negative externality on her, she prefers that j have low initial utility. This is consistent, for example, with the observation from sports that teams do better against demoralized opponents. It could also explain norms of sportsmanship and even why there are rules against directly trying to demoralize opponents (*e.g.*, against excessive celebration after a touchdown).¹⁶

Finally, consider an infinite-horizon version of this simultaneous-move game. Let δ be the common discount factor. Once we consider an infinite horizon, a myriad of equilibria arise because of the players' abilities to reward cooperative

¹⁵Admittedly, physical fatigue also plays a role in explaining this pattern.

¹⁶Under current NCAA rules for American football, a team guilty of excessive celebrating is cited for unsportsmanlike play and penalized 15 yards on the next play. Note that the motivation for this penalty is unlikely to be (solely) the fact that celebrating delays play: There already exists a delay-of-game penalty (a five-yard penalty).

play and punish uncooperative play. Since such equilibria would also emerge in a game with externalities but without moods affecting behavior, we won't explore such equilibria here. Instead, we will focus on a Markov equilibrium of the game. In a Markov equilibrium, play at period t can be conditioned only on the state variables, in this case u_{t-1}^1 and u_{t-1}^2 . Moreover, if $(u_{t-1}^1, u_{t-1}^2) = (u_{\tau-1}^1, u_{\tau-1}^2)$ for any t and τ , then the equilibrium strategies of the players going forward from either t or τ must be the same. Note that Markov equilibria are also subgame perfect. In the Markov equilibrium of the game in which past utility didn't matter (e.g., one in which $u_t^i = \alpha^i x_t^j + x_t^i - (x_t^i)^2/2$), a player would play the same x each period and that x would maximize his utility without regard for his opponent's utility. This will *not* be true in a Markov equilibrium of the game where past utility *does* matter.

Proposition 4 *There exists a Markov equilibrium of the infinite-horizon game in which*

$$x_t^{i*} = \mu^* u_{t-1}^i,$$

where

$$\begin{aligned} \mu^* &= 1 - \frac{2^{\frac{1}{3}}Y}{3\delta Z} + \frac{Z}{3\delta \times 2^{\frac{1}{3}}}, \\ Y &= 6\delta - 3\delta^2 - 6\alpha^i \alpha^j \delta^2, \text{ and} \\ Z &= \left(54\alpha^i \alpha^j \delta^3 + \sqrt{4Y^3 + 2916 (\alpha^i \alpha^j)^2 \delta^6} \right)^{\frac{1}{3}}. \end{aligned}$$

Proof: Please see Appendix A.

Observe that if $\alpha^i = 0$ or $\alpha^j = 0$, then the equilibrium response constant, μ^* , equals one—the same value it would take if each player were playing in isolation (see Example 1). Consequently, as in the two-period game, if just one player is unaffected by the other's action, then *neither* player deviates from what he or she would have done if playing alone. It follows that what distinguishes this equilibrium from isolated play is $\alpha^i \alpha^j \neq 0$. The next proposition addresses how:

Proposition 5 *The equilibrium response constant, μ^* , is increasing in $\alpha^i \alpha^j$. Hence, if the game is friendly or antagonistic ($\alpha^i \alpha^j > 0$), then both players do more of the action conditional on their utility than they would if they played alone. Conversely, if the game is mixed ($\alpha^i \alpha^j < 0$), then both players do less of the action conditional on their utility than they would if they played alone.*

Proof: Observe that the right-hand side of equation (24) in Appendix A equals 1 if $\mu^* = 0$. Hence the right-hand side of (24) must cross the 45° line from above. Hence, because the right-hand side is increasing in $\alpha^i \alpha^j$, it follows that increasing $\alpha^i \alpha^j$ must increase the μ^* at which the right-hand side crosses the 45° line. ■

As noted, we would observe larger responses in both friendly and in antagonistic games. For the observer of a friendly game, a natural interpretation would

be that the players are exploiting infinite repetition to sustain a cooperative outcome (*i.e.*, that promotes the positive externality) or otherwise playing in some reciprocal fashion. In this case, that interpretation would, however, be wrong. Here cooperation is not a consequence of infinite repetition nor any other direct motive to reciprocate. The players appear to cooperate only because each player understands it's better to have a happy opponent than an unhappy opponent. For the observer of an antagonistic game, a natural interpretation—at least at the start of the game—is that the players are punishing each other for *not* cooperating (not doing less of the action in recognition of the negative externality). Again, this interpretation would be incorrect in this context. The players are not punishing so much as attempting to “demoralize” their opponents, since, now, it's better to have an *unhappy* opponent than a happy opponent.

The equilibrium dynamics are

$$\begin{aligned} u_t^i &= \theta(\mu^*) u_{t-1}^i + \alpha^i \mu^* u_{t-1}^j \text{ and} \\ u_t^j &= \alpha^j \mu^* u_{t-1}^i + \theta(\mu^*) u_{t-1}^j. \end{aligned}$$

Define

$$\mathbf{u}_t = \begin{pmatrix} u_t^i \\ u_t^j \end{pmatrix} \text{ and } \mathbf{M} = \begin{pmatrix} \theta(\mu^*) & \alpha^i \mu^* \\ \alpha^j \mu^* & \theta(\mu^*) \end{pmatrix},$$

so

$$\mathbf{u}_t = \mathbf{M}\mathbf{u}_{t-1}.$$

It can be shown (see Elaydi, 1996, §3.1) that

$$\mathbf{u}_t = \begin{pmatrix} \frac{1}{2}(\lambda_1^t + \lambda_2^t) & \frac{1}{2} \frac{\lambda_1^t - \lambda_2^t}{\sqrt{\alpha^i \alpha^j}} \alpha^i \\ \frac{1}{2} \frac{\lambda_1^t - \lambda_2^t}{\sqrt{\alpha^i \alpha^j}} \alpha^j & \frac{1}{2}(\lambda_1^t + \lambda_2^t) \end{pmatrix} \mathbf{u}_0,$$

where $\lambda_h = \theta(\mu^*) + (-1)^{h-1} \mu^* \sqrt{\alpha^i \alpha^j}$ are the eigenvalues of \mathbf{M} .¹⁷

Restricting attention to the friendly and antagonistic cases, we see that

$$u_t^i > \frac{1}{2}(\lambda_1^t + \lambda_2^t) u_0^i$$

in the friendly case (since $\alpha^i > 0$), but

$$\frac{1}{2}(\lambda_1^t + \lambda_2^t) u_0^i > u_t^i$$

in the antagonistic case (since $\alpha^i < 0$). Hence, the players' utilities are always greater in a friendly game than in an antagonistic game. More importantly, consider the friendly game with α^i and α^j and the “polar” opposite antagonistic

¹⁷It might seem this wouldn't apply in the mixed case, where $\alpha^i \alpha^j < 0$. But working through with the resulting complex numbers, it can be shown that it also applies in the mixed case. We will, however, restrict attention to the cases in which $\alpha^i \alpha^j > 0$. In any case, note the model is only valid as long as u_t^i is not driven below zero. In the examples considered, this property is satisfied.

game $\tilde{\alpha}^i$ and $\tilde{\alpha}^j$, where $\tilde{\alpha} = -\alpha$. In these two games, μ^* would be the same. Yet, since $x_t = \mu^* u_{t-1}$, the actions would be less in every period (except the first) in the antagonistic game than in the friendly game. An observer might be tempted to interpret this as the players in the antagonistic game internalizing the negative externality (perhaps because of repeated play). This, however, wouldn't be correct: The players do less in the antagonistic game because they want to do less having been "demoralized" by their opponents. Consequently, the level of action is initially higher in an antagonistic game than in isolated play, but then falls below it. Again, the tempting interpretation is that, after an initial "mistake," the players cooperate by doing less of the action because of the negative externality. And again, that's an incorrect interpretation in this model: The players do less, because the initial high level has succeeded in demoralizing them.

Admittedly, we've analyzed only a limited number of games in which affect could be relevant. Yet, our analysis gives some sense of the issues that arise. In particular, behavior that can be explained by altruism, fairness norms, reciprocity norms, or the exploitation of repeated play can also be explained by affect. In many ways, particularly for *finite* games, affect is more consistent with conventional models of rationality and more parsimonious: It doesn't require players' obeying norms that aren't in their immediate self interest or taking someone else's utility as an argument in their own utility functions. Rather, players simply want to do what's best for them, but they recognize that the affect they induce in their opponents will, through interaction, feedback on them (see Batson, 1991, for empirical evidence that striving to increase others' moods can be motivated by self-interest). Moreover, affect models in games build directly on behavior in single-player decision problems in a way that these other approaches don't. Finally, affect models offer an explanation within the rational-actor paradigm for such behaviors in strategic situations of trying to demoralize your rivals or cheering on your allies, phenomena for which these other approaches don't account.

Clearly, our approach can be applied to a much larger set of strategic settings. Moreover, as we've suggested above, it could serve to complement other approaches. People, for instance, may strive to be fair or reciprocate in cooperative situations because they've found, as a rule of thumb, that the positive affect it induces in others ultimately pays off for them.

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;

- 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;
- savings and consumption behavior;
- apparently cooperative play in finite-period games; and
- attempts to demoralize opponents and to build the morale of friends.

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

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

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).¹⁸ 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 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*.

¹⁸The 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.

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, that we've shown this combination offers real explanatory power with regard to real-life behavior.

Appendix A: Proofs

Proof of Proposition 2: 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. ■

Proof of Proposition 4: Suppose that player j is playing this strategy. We need to show that it is a best response for player i to do the same. As before, define $\theta(q) = q - q^2/2$. Then, for player j ,

$$u_t^j = \alpha^j x_t^i + \theta(\mu^*) u_{t-1}^j.$$

Solving this difference equation yields

$$u_t^j = \alpha^j \sum_{\tau=1}^t \theta(\mu^*)^{t-\tau} x_\tau^i + \theta(\mu^*)^t u_0^j$$

(note we're employing the convention that $\sum_{\tau=1}^0 = 0$). Player i 's utility is

$$\begin{aligned} u_t^i &= x_t^i - \frac{(x_t^i)^2}{2u_{t-1}^i} + \alpha^i \mu^* u_{t-1}^j \\ &= x_t^i - \frac{(x_t^i)^2}{2u_{t-1}^i} + \alpha^i \mu^* \alpha^j \sum_{\tau=1}^{t-1} \theta(\mu^*)^{t-1-\tau} x_\tau^i + \alpha^i \mu^* \theta(\mu^*)^{t-1} u_0^j. \end{aligned}$$

If i is playing a best response to j , then her strategy, $\{x_t^i\}_{t=1}^\infty$, maximizes

$$\begin{aligned} \sum_{t=1}^\infty \delta^t u_t^i &= \sum_{t=1}^\infty \delta^t \left(x_t^i - \frac{(x_t^i)^2}{2u_{t-1}^i} + \mu^* \alpha^i \alpha^j \sum_{\tau=1}^{t-1} \theta(\mu^*)^{t-1-\tau} x_\tau^i + \alpha^i \mu^* \theta(\mu^*)^{t-1} u_0^j \right) \\ &= \sum_{t=1}^\infty \delta^t \left(x_t^i - \frac{(x_t^i)^2}{2u_{t-1}^i} + \mu^* \alpha^i \alpha^j x_t^i \left[\sum_{\tau=1}^\infty \delta^\tau \theta(\mu^*)^{\tau-1} \right] + \alpha^i \mu^* \theta(\mu^*)^{t-1} u_0^j \right). \end{aligned}$$

The “principle of optimality” (see Stokey and Lucas, 1989, §4.1) tells us that we can, thus, solve this maximization problem “period by period” (*i.e.*, treat u_{t-1} as “sunk” with respect to choosing x_t). This yields the first-order condition for x_t^i :

$$\delta^t \left(1 - \frac{x_t^i}{u_{t-1}^i} + \delta \mu^* \alpha^i \alpha^j \sum_{\tau=0}^{\infty} \delta^\tau \theta (\mu^*)^\tau \right) = 0.$$

Hence,

$$x_t^i = \left(1 + \frac{\delta \mu^* \alpha^i \alpha^j}{1 - \delta \theta (\mu^*)} \right) u_{t-1}^i.$$

Observe that the expression in the large parentheses is time-*invariant*. The proof is, therefore, complete if

$$\mu^* = 1 + \frac{\delta \mu^* \alpha^i \alpha^j}{1 - \delta \theta (\mu^*)}. \quad (24)$$

Tedious algebra reveals that is precisely what it equals.¹⁹ ■

Appendix B: 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:

$$u_t = V(x_t, S_{t-1}), \quad (25)$$

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:

$$S_t = x_t + \rho S_{t-1}, \quad (26)$$

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 (26), (25) can be rewritten as

$$u_t = V(S_t - \rho S_{t-1}, S_{t-1}). \quad (27)$$

Given an initial stock, S_0 , the rational addict chooses $\{S_t\}_{t=1}^{\infty}$ to maximize the discounted utility flow of (27) subject to the constraint that $S_t \geq \rho S_{t-1}$ (*i.e.*,

¹⁹It can be shown that the equation (24) has only one real solution (the other solutions are imaginary).

negative consumption is not permitted). For a constant discount factor (*i.e.*, $\omega_t = \delta^t$), the optimal S_t solves

$$V_1(S_t - \rho S_{t-1}, S_{t-1}) - \delta \rho V_1(S_{t+1} - \rho S_t, S_t) + \delta V_2(S_{t+1} - \rho S_t, S_t) = 0 \quad (28)$$

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 (25), 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 (25) and (26) 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.

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. Forthcoming.
- Batson, **C. Daniel**, *The Altruism Question: Toward a Social-psychological Answer*, Hillsdale, NJ: Erlbaum, 1991.
- 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.
- 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.
- 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.
- Frey, **Bruno S.** and **Alois Stutzer**, "What Can Economists Learn from Happiness Research," *Journal of Economic Literature*, June 2002, 40 (2), 402–435.

- 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 Motivation," 1992. Working paper, Cornell University.
- **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.
- 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.
- Lewin, Shira B.**, "Economics and Psychology: Lessons for Our Own Day from the Early Twentieth Century," *Journal of Economic Literature*, September 1996, *34* (3), 1293–1323.
- Loewenstein, George**, "Emotions in Economic Theory and Economic Behavior," *American Economic Review*, May 2000, *90* (2), 426–432.
- 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.
- Rabin, Matthew**, “Psychology and Economics,” *Journal of Economic Literature*, March 1998, *36*, 11–46.
- Romer, Paul M.**, “Thinking and Feeling,” *American Economic Review*, May 2000, *90* (2), 439–443.
- Saarni, Carolyn**, “Socialization of Emotion,” in Michael Lewis and Jeanette M. Haviland, eds., *Handbook of Emotions*, New York: The Guilford Press, 1993.
- Stokey, Nancy L. and Robert E. Lucas Jr.**, *Recursive Methods in Economic Dynamics*, Cambridge, MA: Harvard University Press, 1989. With Edward C. Prescott.
- von Auer, Ludwig**, *Dynamic Preferences, Choice Mechanisms, and Welfare*, Berlin: Springer-Verlag, 1998.
- 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.