

The Effect of Affect on Economic and Strategic Decision Making*

Benjamin E. Hermalin[†]
University of California

Alice M. Isen[‡]
Cornell University

First Version: December 1998

This Version: July 28, 2000

Abstract

The standard economic model of decision making assumes a decision maker makes her choices to maximize her utility or happiness. Her *current* emotional state is not explicitly considered. 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 modeling. Moreover, this paper shows that this simple insight can parsimoniously explain a wide variety of behaviors.

Keywords: *Affect, morale, emotion.*

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

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—what a lay person might call happiness and what an economist might call utility—can markedly influence everyday thought processes and that such influence is a common occurrence. Economic modeling of decision making and game playing has, however, essentially ignored the role of affect. The purpose of this paper is to make amends. In particular, 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 and to do so within a single, simple framework.

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

*Financial support from NSF Grant SBR-9616675 and the Willis H. Booth Professorship in Banking & Finance is gratefully acknowledged. The authors also thank Matt Rabin, Miguel Villas-Boas, and seminar participants at UC Berkeley, UC Davis, 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.

many such attempts have had that flavor (see, e.g., discussions in Lewin, 1996; Rabin, 1998; 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. For instance, consistent with experimental evidence (Isen and Levin, 1972), 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. More generally, the happiness or utility level at the time of decision making affects preferences, which then affects the decision made.

In a one-shot setting, such a change in modeling 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 preferences are determined, in part, by affect at the beginning of the period. As we will show, primarily through examples, such dynamics can explain interesting aspects of people’s behavior and their strategic decision making.

A nice feature of this model is that, although simple, it can encompass a wide range of behaviors. As we will demonstrate, it can, for instance, explain why employers want to hire “happy” workers, workers with “good attitudes,” and why they are concerned with workers’ morale (for evidence that employers are concerned with workers’ morale see Bewley, 1999). It can reconcile rational decision making with the apparent paradox of people eschewing behaviors that correlate with happiness (e.g., socializing, becoming sober, etc.). It offers insights into why moods tend to be persistent and why, for example, pharmacological intervention can be necessary in treating the depressed. It even explains why, in common-interest situations, players try to boost each others’ morale and why, in opposing-interest situations, players try to demoralize each other. It also provides a *single* alternative explanation for behavior that has been explained by a wide variety of assumptions: for example, increased incentives from raising *fixed* wages (i.e., a result resembling efficiency wages), seemingly fair or cooperative play in *finitely* repeated games, and cooperative play *without* punishment strategies in infinitely repeated games.

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; 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. In contrast, our approach also considers how those emotions will affect decision making going forward. That

¹Laibson (1996), Loewenstein (2000), and Romer (2000) are some notable exceptions.

is, for instance, a guilty person will behave differently from a person who doesn't feel guilty (in search 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; Kaufman, 1999) have, like us, worried about the effect of emotional state on decision making. They differ from us in that they are interested in modeling the adverse consequences of emotional state on cognitive abilities.² In addition, there is evidence from psychology that positive affect can promote careful and flexible thinking, at least given minimal effort and interest in the task (see Isen, 2000, for a discussion). In contrast, our actors enjoy normal cognition, and we consider the effects of normal, everyday mild emotional states or feelings. To be sure, we are certainly sympathetic to the view that emotional state can affect cognitive ability,³ but it is worth exploring how emotional state affects behavior without departing from the rational-actor paradigm.

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; 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. Such “choice-set” effects are absent here—the choice set remains constant over time in our models.⁴ In addition, 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 (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

²MacLeod 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.

³Ashby 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 et al. (1999) 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.

⁴This isn't to suggest, however, that we can't conceive of affective state playing a role in determining the choice set, nor that our model wouldn't apply when the choice set varies over time (for whatever reason). In fact, some work already suggests that positive affect can increase the choice set (Kahn and Isen, 1993). A time-invariant choice set, however, makes more straightforward what the role of the affective state is.

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

The rest of the paper proceeds, in the next section, by presenting the basic model and analyzing two example applications. In Section 3, we explore how affect can affect the play of games. We conclude in Section 4.

2 A Model of Positive Affect & Decision Making

Consider the following model in which positive affect or utility level influences decision making. An individual begins period t with utility $u_{t-1} \in \mathbb{R}$ determined by her past experiences. In period t , she makes decisions $\mathbf{x}_t \in \mathcal{X}$, where \mathbf{x}_t is a vector and \mathcal{X} is the time-invariant feasible set. Let her utility at the end of the period, u_t , be

$$u_t = U(\mathbf{x}_t, u_{t-1}).$$

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

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

$$\max_{\mathbf{x} \in \mathcal{X}} U(\mathbf{x}, u) \tag{1}$$

for all possible u . Assume, too, that,

$$\text{if } u > u', \text{ then } U[\mathbf{x}^*(u), u] > U[\mathbf{x}^*(u'), u'] . \tag{2}$$

Then the solution to

$$\max_{\{\mathbf{x}_t\}_{t=1}^T} \sum_{t=1}^T \delta^t U(\mathbf{x}_t, u_{t-1}) \tag{3}$$

(where $\delta > 0$ and less than one if $T = \infty$) is $\mathbf{x}_t = \mathbf{x}^(u_{t-1})$; 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 (3). Hence, $\mathbf{x}^*(u_{t-1})$ are 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 (1) 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 (2) holds). In the examples below, it is readily shown that these assumptions are met.

⁵At the level of the underlying mathematics, there is, admittedly, in all these models a means of capturing present utility as a function of past consumption (plus an initial state). At this level, then, the differences among these models come down to differences in assumptions about the functions from $\{\mathbf{x}_\tau\}_{\tau=1}^t$ to \mathbf{u}_t . A virtue of this underlying relationship is that one can, therefore, relate our analysis to the broader literature on consistent plans (see, e.g., Goldman, 1980).

Observe that the equilibrium relationship between current utility and past utility is monotonically increasing. 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.⁶

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

$$u_t = U[\mathbf{x}^*(u_{t-1}), u_{t-1}]. \quad (4)$$

Consider, now, a couple of examples.

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 et al., 1999, the last, in particular, shows that positive affect increases desire for the reward effort provides). These results can, in turn, be captured by assuming

$$U(x_t, u_{t-1}) = x_t \sqrt{2\beta} - \frac{x_t^2}{2u_{t-1}} + \bar{u}(1 - \beta), \quad (5)$$

where $\beta \in (0, 1)$ determines the marginal benefit of effort and \bar{u} is some constant. To keep the model from being pathological, assume $\bar{u} \geq 0$ 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, $x^*(u_{t-1}) = u_{t-1} \sqrt{2\beta}$. Consistent with the experimental evidence, $x^*(\cdot)$ is an increasing function. Equation (4) becomes

$$u_t = \beta u_{t-1} + \bar{u}(1 - \beta).$$

The solution to this difference equation is

$$u_t = u_0 \beta^t + \bar{u}(1 - \beta^t). \quad (6)$$

As $t \rightarrow \infty$, $u_t \rightarrow \bar{u}$; that is, \bar{u} is the long-run steady-state of this difference equation. Observe that, at any time t , u_t is the weighted average of the steady-state utility and the previous period’s utility. Hence, utility is improving if $\bar{u} > u_{t-1}$, falling if $\bar{u} < u_{t-1}$, and unchanging if $\bar{u} = u_{t-1}$. Consequently, \bar{u} is a stable fixed point of this difference equation.

Suppose that the individual in question is employed. Let the per-period benefit to her employer from x be vx , where $v > 0$. Observe that, *ceteris paribus*, the employer prefers to hire a “happier” individual and keep her happy, since the employer’s per-period revenue is

$$v\sqrt{2\beta} [u_0 \beta^{t-1} + \bar{u}(1 - \beta^{t-1})],$$

which is clearly increasing in u_0 . This corresponds to the well-known adage that “happy workers make good workers” and is consistent with the findings in Bewley (1999). If we also supposed that \bar{u} denoted a steady-state utility of the individual—e.g., something akin to “attitude”—then we see, in part, why employers seek employees with good attitudes or other attributes associated

⁶They 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).

with long-run positive affect. Relatedly, to the extent that the employer can undertake activities to raise u_0 or \bar{u} or both (e.g., pay a signing bonus, ensure a pleasant working environment, or not cut wages in a recession), he will have incentive to do so, since this yields him greater benefit.

Worker output can also be a function of incentives. Suppose, for instance, that \bar{u} is $\gamma(w)$, where $\gamma(\cdot)$ is an increasing and concave function relating the per-period wage, w , to utility.⁷ For convenience, assume infinite employment. Assume a constant discount factor of $\delta \in (0, 1)$; i.e., assume an interest rate of $(1 - \delta)/\delta$. Then the wage will be set to maximize

$$\sum_{t=0}^{\infty} \delta^t \left[v\sqrt{2\beta} [u_0\beta^t + \gamma(w)(1 - \beta^t)] - w \right].$$

The solution is defined by

$$\gamma'(w^*) = \frac{1 - \beta\delta}{\delta(1 - \beta)v\sqrt{2\beta}}. \quad (7)$$

Note, first, that, although the employer is paying a fixed wage, it nevertheless has important incentive effects. In some ways, it is like an efficiency-wage story (see, e.g., Akerlof, 1982; Shapiro and Stiglitz, 1984), but differs in so far as it is not explicitly dependent on the threat of unemployment or the existence of an alternative employment sector. The two models could be more closely linked by assuming that $\gamma(w)$ is also a function of the unemployment rate and relative wages, similar to what Akerlof (1982) does.

Inverting the right-hand side of (7), we see the derivative of the inverse with respect to β is

$$\frac{1}{2}\delta v\sqrt{2}\frac{1 - 3\beta + \beta\delta + \beta^2\delta}{(1 - \beta\delta)^2\sqrt{\beta}}.$$

This is positive for β less than

$$\frac{1}{2\delta}3 - \delta - \sqrt{(9 - 10\delta + \delta^2)}$$

and negative for β greater than that, which means w^* is increasing in β for β less than that and is decreasing in β for β greater than that ($\gamma(\cdot)$, recall, is concave). In words, the wage is, at first, increasing in the worker's intrinsic motivation, $\sqrt{2\beta}$, and, then decreasing in it. This occurs because, when intrinsic motivation is low, utility is mostly a function of the wage. Hence, raising the wage has a big impact on utility. The impact of this, however, depends on intrinsic motivation. The greater the intrinsic motivation, the greater the return from inducing positive affect. Consequently, as intrinsic motivation rises, the marginal return to the employer from raising the wage is increasing, so he increases the wage. When, however, intrinsic motivation is high, utility is relatively insensitive to the wage. Consequently, there is less to be gained from a high wage, so the wage rate begins to fall with intrinsic motivation.

If x can be measured directly, then we could also consider paying a piece rate, s , as an incentive. Changing the model somewhat, suppose

$$U(x_t, u_{t-1}) = sx_t - \frac{x_t^2}{2u_{t-1}} + \tilde{u}.$$

Then $x^*(u_{t-1}) = su_{t-1}$ and the difference equation becomes

$$u_t = \frac{s^2 u_{t-1}}{2} + \tilde{u}.$$

⁷For example, consistent with Bewley's findings that pay cuts affected employees more than pay hikes, $\gamma(\cdot)$ could be steeply sloped for $w < W_P$ and less steeply sloped for $w > W_P$, where W_P is some prevailing or "fair" wage.

Its solution is

$$\begin{aligned} u_t &= \left(\frac{s^2}{2}\right)^t u_0 + \left[1 - \left(\frac{s^2}{2}\right)^t\right] \frac{\tilde{u}}{1 - \left(\frac{s^2}{2}\right)} \\ &= \beta(s)^t u_0 + \left[1 - \beta(s)^t\right] \frac{\tilde{u}}{1 - \beta(s)} \end{aligned}$$

hence,

$$x_t = s \left(\beta(s)^{t-1} u_0 + \left[1 - \beta(s)^{t-1}\right] \frac{\tilde{u}}{1 - \beta(s)} \right)$$

The employer's time-0 problem is, thus,

$$\begin{aligned} & \max_s \sum_{t=0}^{\infty} \delta^t (v - s) s \left(\beta(s)^t u_0 + \left[1 - \beta(s)^t\right] \frac{\tilde{u}}{1 - \beta(s)} \right) \\ &= \max_s \frac{s(v - s)}{1 - \delta} \left[\frac{(1 - \delta)u_0 + \delta\tilde{u}}{1 - \beta(s)\delta} \right] \\ &= \max_s \frac{s(v - s)}{1 - \delta} H(s). \end{aligned}$$

Suppose we had “standard preferences” that did *not* depend on mood (i.e., $\partial U_t / \partial u_{t-1} \equiv 0$). This is equivalent to assuming $\delta = 0$, which implies $s^* = v/2$.⁸ When preferences *do* depend on mood (i.e., $\delta > 0$), $s^* > v/2$ because $H(s) > 0$ and

$$H'(s) = \frac{(1 - \delta)u_0 + \delta\tilde{u}}{(1 - \beta(s)\delta)^2} \beta'(s) > 0.$$

Hence, in our model we should see stronger piece rates than the standard model: It pays to invest in a “happy” worker, which means raising the piece rate. Observe, as well, that s^* is independent of u_0 ;⁹ that is, the piece rate does *not* depend on the level of initial affect—yet affect is still important for determining the piece rate through its dynamic effect.

As a second example,

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

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

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 (4) becomes

$$u_t = \frac{\phi(u_{t-1})^2}{2}. \tag{8}$$

⁸To ensure that $\beta(s) < 1$, it's necessary that $v < 2\sqrt{2}$.

⁹Note that we need to assume that $u_0 > 0$ and $4 > 2\delta v^2$ for the model to be consistent.

Unless $\phi(u) \propto \sqrt{u}$, this is a nonlinear difference equation. The second derivative of the right-hand side function 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 function crosses the 45°-line three times (see Figure 1). This would be true, for instance, if

$$\phi(u) = \frac{\beta u}{u + 1}, \tag{9}$$

where $\beta > \sqrt{8}$. 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 (8) 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 about this model:

- 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.¹⁰
- 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 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.¹¹
- 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.¹² 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.

¹⁰This could, however, be interrupted by actions of others. For instance, harmful acts, even neglect, by others could be a shock to the dynamic system. We briefly consider such shocks below and strategic interactions with others in Section 3.

¹¹The 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?

¹²Since 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.




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

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. 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.¹³

Returning to the general formulation, define $U^*(u) = U[\mathbf{x}^*(u), u]$. Assume that $u_0 \in \mathcal{U}$ and $U^* : \mathcal{U} \rightarrow \mathcal{U}$, where \mathcal{U} is an interval in \mathbb{R} . Let $I_{\mathcal{U}}$ be the greatest lower bound on \mathcal{U} and $S_{\mathcal{U}}$ be the

¹³This is not inconsistent 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 experience is 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.

least upper bound on \mathcal{U} ; that is,

$$I_{\mathcal{U}} = \inf \mathcal{U} \text{ and } S_{\mathcal{U}} = \sup \mathcal{U}.$$

Observe, then, that

Proposition 2 *Assume that $U^* : \mathcal{U} \rightarrow \mathcal{U}$ is continuous, that $u_0 \in \mathcal{U}$, and that \mathcal{U} is an interval in \mathbb{R} . Then, in equilibrium, change in utility is (weakly) monotonic over time. Specifically, it is either non-decreasing or non-increasing (i.e., $(u_{t+1} - u_t)(u_t - u_{t-1}) \geq 0$ for all t).*

Proof. If $u_t - u_{t-1} = 0$, then we’re at a fixed point, so $u_{t+1} - u_t = 0$, which establishes the result. So assume, instead, that $|u_t - u_{t-1}| > 0$. We’ll consider only the case $u_t - u_{t-1} > 0$. The case $u_t - u_{t-1} < 0$ is proved similarly. Since $U^*(\cdot)$ is continuous, there exists an open interval¹⁴ in \mathcal{U} containing u_{t-1} such that $U^*(u) > u$ for all u in that interval. Let (I, S) be the largest such interval. Now I is either a fixed point or it’s $I_{\mathcal{U}}$. Similarly, S is either a fixed point or it’s $S_{\mathcal{U}}$. Either way, since $U^*(\cdot)$ is increasing, $I \leq U^*(I) < U^*(S) \leq S$. Hence, $U^*[(I, S)] \subseteq (I, S)$. Consequently, u_t is also in this interval, implying $u_{t+1} = U^*(u_t) > u_t$. The result follows. ■

Given that affective state is determined by the physical processes of the brain and all physical processes are bounded, it seems reasonable, therefore, to assume that \mathcal{U} is a bounded interval in \mathbb{R} . Without any obvious loss of generality, we may assume it’s closed as well, in which case the following corollary of Proposition 2 holds:

Corollary 1 *Assume that $U^* : \mathcal{U} \rightarrow \mathcal{U}$ is continuous, that $u_0 \in \mathcal{U}$, and that \mathcal{U} is a compact interval in \mathbb{R} . Then the sequence $\{u_t\}_{t=0}^{\infty}$ either begins at a fixed point or is monotonically approaching a fixed point.*

Proof: Please see Appendix.

Taken literally, these monotonicity results (Proposition 2 and its corollary) are unrealistic. People’s moods do not move monotonically: We *can* have good days, followed by bad days, followed by good days, and so forth. A more realistic model can be achieved by assuming that mood is subject to random events outside the decision maker’s control (e.g., for many people, sunny days boost mood, while grey days lower it—see Cunningham, 1979). A complete analysis of random perturbations is beyond the scope of this paper,¹⁵ but a few points can nevertheless be made.

First, 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). 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

¹⁴Half open if $u_{t-1} = I_{\mathcal{U}}$ or $= S_{\mathcal{U}}$.

¹⁵An earlier working paper did, however, contain some exploratory analysis—see <http://haas.berkeley.edu/~hermalin/random.pdf> or contact the authors.

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—if the dynamic process is described, for instance, by equation (6), then the two sides' impacts are symmetric. 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 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.¹⁶ 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 $\bar{u} = 0$) but with the addition of the impact of the *other* player's (j 's) action on 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 $i = 1, 2$.

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

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

¹⁶See Saarni (1993), e.g., for a survey of the role of emotions in interpersonal interactions.

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

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

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. Now,

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

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

$$\underbrace{\alpha^i \left(\alpha^j x_1^i + x_1^j - \frac{(x_1^j)^2}{2u_0^j} \right)}_{x_2^{j*}} + \frac{1}{2} \left(\alpha^i x_1^j + x_1^i - \frac{(x_1^i)^2}{2u_0^i} \right) + \underbrace{\alpha^i x_1^j + x_1^i - \frac{(x_1^i)^2}{2u_0^i}}_{u_1^i}. \quad (10)$$

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.¹⁷ 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 (10), 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 $\theta(q) = q - \frac{1}{2}q^2$. 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 $\min 2 - \frac{1}{3} \alpha^i \alpha^j = \frac{2}{3}$. Hence, when j 's action provides a positive externality for i , she prefers

¹⁷Admittedly, physical fatigue also plays a role in explaining this pattern.

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).¹⁸

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 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 3 *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.

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 4 *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 (11) equals 1 if $\mu^* = 0$. Hence the right-hand side of (11) must cross the 45° line from above. Hence, since the right-hand side is increasing

¹⁸Under 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).

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, e.g., 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} .¹⁹

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

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

¹⁹It 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.




Figure 2: Example path for x_t^{i*} for friendly and antagonistic games plus for playing alone.

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 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. Figure 2 illustrates an example time path of x_t^{i*} (here, $\delta = .9$, $\alpha^i = \alpha^j = .35$, $\tilde{\alpha}^i = \tilde{\alpha}^j = -.35$, and $u_0^i = u_0^j = 10$). Observe that, in Figure 2, 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.

Finally, in a game, a given player’s utility no longer needs to follow a monotonic path—see Figures 3. Hence, consistent with earlier discussion, strategic interactions can introduce non-monotonicity into an individual’s utility time path.

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, using affect, 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




Figure 3: u_t^i in a friendly game where $\alpha^i = .86$, $\alpha^j = .23$, $\delta = .9$, and $u_0^i = u_0^j = 10$.

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, and consequently the level of others' actions, ultimately pays off for them.

4 Conclusions

In this paper, we have shown that incorporating the psychological finding that affect influences decision making can greatly enrich rational-actor models of decision making and strategic interaction. Although a modest change in our standard assumptions—yet 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;
- changes in productivity from changes in a *fixed* wage;
- the setting of piece rates;
- the paying of signing bonuses and other *non*-contingent rewards that serve to boost worker morale;
- the apparent paradox of people not pursuing behaviors correlated with well-being;
- pharmacological treatment strategies for depression;
- decision making under uncertainty;
- 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

$$u_t = -g_{t-1}e^{-x_t} \text{ and } g_t = g_{t-1}e^{x_t}$$

where g_t is the level of guilt at time t and x_t is a guilt-inducing action ($x_t < 0$ is a guilt-reducing action).²⁰ 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; 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; 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*. At the same time, one of the amazing things about 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

Proof of Corollary 1: Since $U^*(\cdot)$ is continuous and increasing, $U^*(\mathcal{U}) = [U^*(I_U), U^*(S_U)] \subset \mathcal{U}$. From the Brouwer Fixed-Point Theorem, it follows, then, that $U^*(\cdot)$ has at least one fixed point. From the proof of Proposition 2, if $u_t \neq u_{t-1}$, then $u_{t+1} \neq u_t$. Hence, if the sequence doesn’t start at a fixed point, it is either strictly increasing or decreasing. Now, between any two consecutive fixed points a positive distance apart, we have from the proof of Proposition 2 that the sequence is either strictly increasing or strictly decreasing. Moreover, it can’t “escape” the interval, so it must be approaching one of the two fixed points. Now $U^*(I_U) \geq I_U$, hence either I_U is the smallest fixed point or there is a smallest fixed point strictly greater than I_U . Only the latter case concerns us. Let I be this smallest fixed point. Then, since $U^*(I_U) > I_U$, $U^*(u) > u$ for all $u \in [I_U, I)$.

²⁰The 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.

A sequence that starts in that interval, therefore, increases monotonically toward I . A similar argument takes case of what happens “at the top” (i.e., to the left of S_U). ■

Proof of Proposition 3: 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

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

The “principle of optimality” (see, e.g., Stokey and Lucas, 1989, §4.1) tells us that we can solve this maximization problem by ensuring that each x_t^i maximizes the discounted sum of all the utilities it *directly* affects. 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

$$\mu^* = 1 + \frac{\delta \mu^* \alpha^i \alpha^j}{1 - \delta \theta(\mu^*)}. \tag{11}$$

Tedious algebra reveals that is precisely what it equals.²¹ ■

²¹It can be shown that the equation (11) has only one real solution (the other solutions are imaginary).

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.
- 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, Alice M. Isen, and Christopher Purdy**, “The Influence of Positive-Affect on Expectancy Motivation: Integrating Affect and Cognition into Motivation Theories,” 1999. Working Paper, Johnson Graduate School of Management, Cornell University.
- Frank, Robert H.**, *Passions within Reason*, New York: Norton, 1988.
- Goldman, Steven M.**, “Consistent Plans,” *Review of Economic Studies*, 1980, XLVII, 533–537.
- 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.
- Kahn, Barbara and Alice M. Isen**, “The Influence of Positive Affect on Variety Seeking among Safe, Enjoyable Products,” *Journal of Consumer Research*, 1993, *20*, 257–270.
- 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.**, “A Cue-Theory of Consumption,” 1996. Working paper, Department of Economics, Harvard University.
- 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.
- 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 Jeannette M. Haviland, eds., *Handbook of Emotions*, New York: The Guilford Press, 1993.
- Shapiro, Carl and Joseph E. Stiglitz**, “Equilibrium Unemployment as a Worker Discipline Device,” *American Economic Review*, June 1984, *74* (3), 433–444.
- 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.