Utility and Happiness

Miles Kimball and Robert Willis[1]

University of Michigan

October 30, 2006

Abstract: Psychologists have developed effective survey methods of measuring how happy people feel at a given time. The relationship between how happy a person feels and utility is an unresolved question. Existing work in Economics either ignores happiness data or assumes that felt happiness is more or less the same thing as flow utility. The approach we propose in this paper steers a middle course between the two polar views that "happiness is irrelevant to Economics" and the view that "happiness is a sufficient statistic for utility."

We argue that felt happiness is not the same thing as flow utility, but that it does have a systematic relationship to utility. In particular, we propose that happiness is the sum of two components: (1) *elation*--or short-run happiness--which depends on recent news about lifetime utility and (2) *baseline mood*--or long-run happiness--which is a subutility function much like health, entertainment, or nutrition. In principle, all of the usual techniques of price theory apply to baseline mood, but the application of those techniques is complicated by the fact that many people may not know the true household production function for baseline mood.

If this theory is on target, there are two reasons data on felt happiness is important for Economics. First, short-run happiness in response to news can give important information about preferences. Second, long-run happiness is important for economic welfare in the same way as other higher-order goods such as health, entertainment, or nutrition.

I. Introduction

In Economics, data is often a limiting factor in what research can be done. The scarcity value of data is such that it is not uncommon for one data set to be analyzed by scores of different researchers using different theoretical and statistical models. The broad area of "Cognitive Economics"[2] is defined by a willingness to consider seriously nonstandard types of data in order to increase the total set of data amenable to economic analysis. Among the important types of nonstandard data are data from the laboratory experiments of Experimental Economics, the brain-scan data of Neuroeconomics, the gene and physiological data of Biological Economics, survey measures of expectations, survey measures of preference parameters based on hypothetical choices, and self-reports on a respondent's subjective attitudes, values, beliefs, judgments and feelings.

Given the large amount of self-report data collected on "happiness"—often called "subjective well-being"—assessing the usefulness of happiness data for economic analysis is a high priority for Cognitive Economics. Jonathan Gruber and Sendhil Mullainathan (2005) conclude with these words arguing for the value of happiness data to Economics:

> Subjective well-being measures provide a possible way to directly address welfare questions. As our analysis shows, this direct approach is empirically feasible. Happiness measures may be noisy, but in our case at least, they contain sufficient signal to discern effects of moderate size policies. This is heartening because happiness data is abundant. In the U.S. the GSS is available in moderately large samples for many years. Looking beyond the U.S., the Canada data we use is not the exception but rather the rule: many countries, notably in Europe, collect cross-sections and panel data on happiness. In short, the results in this paper suggest that by using happiness data, economists may be able to directly assess the impacts of public policy on well-being.

Standard economic data are data on the constraints people face and on the choices people make subject to those constraints in real-life situations. Economics is the science of choices, and is an attempt to answer real-world questions and to deal with real-world issues. Though any one paper can only do so much, in the end, to be part of Economics, any nonstandard data should ultimately be linked back theoretically and empirically by some chain of reasoning and facts to standard data on choices "in the wild." In this paper, we attempt to forge a key link in the chain from happiness to standard economic concepts and data by presenting a theoretical model for the relationship between happiness and observed choices. We argue that while this task can be accomplished, it is a more difficult and subtle task than is often realized. In particular, many of the economists who have worked with happiness data—including Frey and Stutzer (2004b), Gruber and Mullainathan (2005),and Layard (2005)—equate happiness and utility. We argue that this equation of happiness and utility is problematic. Instead we argue for a more flexible— but still strong—two-way relationship between happiness and preferences:

1. *Preference for Happiness*: Other things being equal, people prefer to feel happy. (That is, happiness is one of the arguments of the utility function.)
2. *Happiness and News*: Temporary spikes and dips in happiness beyond a baseline level reflect recent good and bad news, where "good news" reflects a transition to a preferred situation and "bad news" reflects a transition to a less preferred situation.

---

[2] See Miles Kimball and Robert Willis (2006).

Either or both of these two connections between happiness and preferences would make happiness data important for economic analysis. First, to the extent people prefer to be happy, happiness is an appropriate subject for economic policy analysis.   Second, to the extent that spikes and dips in happiness provides a signal about what people consider good and bad news, happiness data provides important information about preferences that may not be duplicated by other available data.

II. Distinguishing Between Utility and Happiness

A. Defining Utility and Happiness. On first impression, *"utility"* and *"happiness"* seem to refer to the same concept. This goes back to the fact that the word "happiness" has two distinct meanings: (i) the greatest good for an individual and (ii) a positive feeling. To distinguish between these two meanings, we use "utility" to refer to the greatest good for an individual (as viewed by that individual) and "happiness" to refer to a positive feeling.[3] Using this narrow meaning of the word "happiness," we can say quite starkly that "Feeling happier is not necessarily a good thing, *if something more important is sacrificed in order to obtain that happiness***." This usage, which makes a logical distinction between "utility" and "happiness," is in line with the technical meanings for the words "utility" and "happiness" that economists and psychologists respectively have developed over the last century (though not with earlier usages).

The success of the Ordinalist Revolution of Lionel Robbins (1932) and of John Hicks and R. G. D. Allen (1934)—codified as "Revealed Preference" by Paul Samuelson (1938, 1947)[4]—has fixed the meaning of "*utility*" for more than a half-century of economists as a representation of an individual's preferences over alternatives.   The practice of Economics has made this concept of utility immensely valuable in thousands of applications.

In the aftermath of the Cognitive Revolution, the success of Hedonic Psychology—exemplified in the volume edited by Daniel Kahneman, Ed Diener and Norbert Schwarz (1999)—has fixed the scientific meaning of "*happiness*" within Psychology as the overall goodness or badness of an individual's felt experience at any point in time. In practice, these feelings are often gauged by questions such as "*On a scale from one to seven, where one is extremely unhappy and seven is extremely happy, how do you feel right now?*"[5]

---

[3] There is a third meaning of the word "happiness" commonly used in attempts at persuasion: the greatest good for an individual as viewed by the person using the word. We would call this "recommended utility," representing "recommended preferences."   That is, this use of the word happiness is intended to tell people what they *should* desire. We discuss this briefly in the section on utility.

[4] For more of the history of these developments, see also George Stigler (1950).

[5] For those who intend to personally use or interpret others' results based on subjective well-being data, Appendix A discusses the measurement of happiness in greater depth. Appendix A has two key points. First, a large number of different ways of measuring current experienced happiness—including self-reports—seem to give a consistent answer. Second, attempts to go beyond measuring current happiness to measure overall life-satisfaction, or overall happiness with one's entire life are fraught with problems. In particular,, because of the cognitive difficulty of answering such questions, people take shortcuts in constructing their answers.   Psychological experiments show that the measures one obtains from attempts to measure overall life-satisfaction or overall happiness with one's life are heavily influenced by happiness at the moment of the survey, survey context effects, and people's own moral and folk-psychology theories of how they *should* feel about their lives. (See Norbert Schwarz and Fritz Strack, 1999.)

Throughout this paper, we follow the convention that the technical meaning of *"utility"* is determined by the tradition in Economics, while the technical meaning of *"happiness"* is determined by the tradition in Hedonic Psychology. Thus, *utility* is a reflection of people's choices; *happiness* is a reflection of people's feelings. Once one recognizes these two concepts as distinct, *discovering the nature of the empirical relationship between utility and happiness* stands out in sharp relief as one of the central questions at the frontier between Economics and Psychology.

B. The Benthamite Tradition of Equating Utility and Happiness. One of the difficulties we face in making our viewpoint clear is that the tradition of equating "happiness" to flow utility runs deep in the history of economic thought. Indeed, Jeremy Bentham's (1781) first definition of 'utility' made the equation of utility and happiness explicit:

> "By the principle of utility is meant that principle which approves or disapproves of every action whatsoever according to the tendency it appears to have to augment or diminish the happiness of the party whose interest is in question …."

The "Revealed Preference" definition of utility—to which we resolutely adhere—is closer to Bentham's second, more inclusive, definition of utility, in the immediately following paragraph:

> "By utility is meant that property in any object, whereby it tends to produce benefit, advantage, pleasure, good, or happiness, (all this in the present case comes to the same thing) or (what comes again to the same thing) to prevent the happening of mischief, pain, evil, or unhappiness to the party whose interest is considered: if that party be the community in general, then the happiness of the community: if a particular individual, then the happiness of that individual."

Another difficulty we face in distinguishing utility and happiness is that, while "Revealed Preference" guides economic research, a more naïve Marginalism has remained very common in economic teaching. For example, "Principles of Economics" courses often teach about diminishing marginal utility by engaging students' intuitions about how happy they would feel in consuming different consumption bundles.

Let us state clearly that, throughout this paper, when we discuss utility, we do so from the perspective of Paretian Welfare Economics. Whether explicitly or implicitly, welfare questions motivate a large share of economic research; an orientation toward welfare questions is particularly important in informing our assessment of utility in cases where people are liable to mistakes. As for the focus on Pareto optimality, in our view, the use of happiness data is *not* a

Philosopher's Stone that magically solves the difficulties in comparing utility interpersonally, but happiness data—used judiciously—can give useful information about individual preferences.[7]

Any adequate theory of utility and happiness must explain why the meanings of happiness and utility seem so similar. The right nuances for explaining the semantic relationship between "happiness" and "utility" can be found in the first two definitions for "happy" in the *American Heritage Dictionary* (1976, Houghton Mifflin):

happy … 1. Characterized by luck or good fortune; prosperous. 2. Having or demonstrating pleasure or satisfaction; gratified.''

The second definition is the meaning of "happy" in Psychology. The first definition talks about prosperity, which seems closely linked to utility, but there is a hint of a stochastic element in the nature of happiness: "luck or good fortune." Our view of happiness emphasizes *recent* good luck by positing that an important component of happiness has to do with an individual's reaction to recent news about lifetime utility. We call the component of felt happiness due to recent news about lifetime utility "elation." As an example of this meaning of the word "elation" the New York Times reports the Yankees' manager Joe Torre saying "*With the danger of failing is the elation of winning. You can't get elated unless there's a danger.*" (October 11, 2006, C19.)

Although the differences are important, news about lifetime utility and lifetime utility itself are linked tightly enough that it is not surprising to find some confusion between the two meanings in the structure of the lay lexicon. In other words, if people *feel happy* whenever they receive good news about lifetime utility, it is not hard to see why they would sometimes use the word "happiness" to describe *lifetime utility* itself. Yet scientifically, we consider it crucial to have two distinct, clearly delineated concepts for revealed preference utility and happiness in the psychological sense of current feelings. Maintaining two distinct concepts—on an equal footing—in a situation where each has a certain tendency to subordinate or engulf the other, is one of the main contributions of this paper.

One way to think about the distinction between utility and happiness is that one's commitment to an Ordinalist, "revealed preference" definition of utility is confronted with an acid test when confronted with happiness data. There is a sense in which the most radical implications of the Ordinalist Revolution are apparent only in the light of data on experienced happiness.

Both felt happiness and choice-based utility are well-defined, observable concepts. Our aim is to determine the dynamic relationship between the standard psychological concept of current

---

[7] One can then make the leap from individual preferences to statements about social welfare on more or less the same terms as one could in the absence of happiness data. To the extent that happiness data give the illusion of providing a cardinal utility function, it is an illusion similar to that provided by expected utility theory—where one may sometimes need to be reminded that a monotonic transformation $f(E(U))$ of the *overall* objective function $E(U)$ leaves preferences unaltered. Just as there is no *necessary* reason why the curvature of U in the expected utility representation $E(U)$ tells us how to aggregate preferences interpersonally, there is no *necessary* reason why whatever structure is revealed in preferences *as they relate to happiness data* tells how preferences must be aggregated. At a minimum, any debate about what happiness says about social welfare must take into account the existing literature on social welfare and social choice theory.

affect—felt happiness—and the standard economic concept of lifetime utility. Establishing *any systematic relationship* between happiness and utility would provide an important bridge between Psychology and Economics, allow psychological data and theory to be used in Economics in a way that is complementary to standard economic data and theory, and enable economists to bring to bear all the tools of economic theory toward understanding happiness.

C. The Neo-Benthamites. Economists have been slower than psychologists to focus on subjective well-being data. But a growing economic literature has made use of subjective well-being data. With very few exceptions, this literature explicitly or implicitly follows the Benthamite tradition of equating utility and happiness. Richard Layard's (2005) book gives a good introduction to this literature and Bruno Frey and Alois Stutzer (2002) give a partial survey.[8] This literature lays out many provocative findings, focusing primarily on the cross-sectional and trend properties of subjective well-being. Two key motivations for the use of subjective well-being data in Economics (shared in large measure by Hedonic Psychology itself) have been (i) the desire to study the welfare implications of non-traded goods[9] (something that is especially important for older people for whom market work is a less dominant part of their lives) and (ii) the desire to study welfare implications in contexts where preferences are potentially inconsistent and to diagnose optimization mistakes.[10]

Despite this growing literature, many economists are still very skeptical of the use of subjective well-being data,[11] in large part because the theoretical status of affect--"happiness"—within economic theory is unclear.[12] A simple multiple-choice question illustrates this lack of clarity:

> What is Happiness?
>   a.   Flow utility?
>   b.   The individual's overall objective function?
>   c.   The part of the objective function that abstracts from the desire to do one's duty?
>   d.   The individual's objective function plus pleasure from memory?
>   e.   None of the above?

The answer for Neo-Benthamites is (a): current happiness is equal to flow utility; our answer is (e): none of the above.

D. Happiness, Utility and Time. Bringing the dimension of time into the discussion of utility and happiness requires a few more definitions. In Hedonic Psychology, *affect* is a useful term to refer to how happy a respondent currently feels, as opposed to judgments about his or her whole

---

[8] Easterlin (1974) deserves a lot of credit for stimulating interest by economists in happiness empirics. He followed this work up by Easterlin (1995) and Easterlin (2003). Some other recent empirical papers in economics using happiness data are John Helliwell (2002), David Blanchflower and Andrew Oswald (2004), Clark (1999), Rafael Di Tella, Alberto Alesino and Robert MacCulloch (2004), Di Tella and MacCulloch (1999), Di Tella, MacCulloch and Oswald (2001, 2003), and Wolfers (2003).

[9] See for example Frey, Simon Luechinger and Stutzer (2004) and Frey and Stutzer (2000, 2004a).

[10] See for example, Jonathan Gruber and Sendhil Mullainathan (2005) and Frey and Stutzer (2004b).

[11] See, for example, Daniel Hamermesh (forthcoming).

[12] In Appendix B, we argue that psychologists can reliably measure happiness, in the carefully defined sense of how people feel at a given time. Of course, that leaves the question of what happiness *is*. To say the same thing in a different way, some economists think happiness can't be measured well. *This is just not true.* Happiness (current affect) is one of the easiest of all subjective concepts to measure. What *is* true (that these economists are intuiting) is that once happiness is measured, we don't know what it means in terms of economic theory.

life. Throughout this paper, we use "affect" and "happiness" interchangeably.   We discuss the relationship between affect and other subjective well-being concepts in Appendix A. In economics, "lifetime utility" is a useful shorthand to refer to an individual's overall objective function—*including things the individual cares about that occur after his or her death*.   We can distinguish lifetime utility and current affect ("happiness") as follows:

- **Lifetime Utility** = The extent to which people get what they want as revealed by their *choices* in the face of varying constraints.
- **Current Affect** = How positive people's *feelings* are at a given time.

In thinking about lifetime utility, it is important to remember that people's choices clearly show that they value a wide range of goods that are not traded in markets or only partially traded in markets. Thus, the economic concept of lifetime utility is not limited to what are sometimes called "economic goods" but includes the value an individual places on *non-traded goods* such as respect, freedom, clean air, a vibrant community, being married to a particular person, and such *partially-traded goods* as time allocations--which are partially traded because people are paid for work time---and health and longevity--which are partially traded because people pay for health care. (See Kevin Murphy and Robert Topel, 2005, for an example of valuation of health and longevity.)

Lifetime utility is the standard welfare measure in economics at the individual level. It is often thought of as a discounted sum over time of "flow utility." As a counterpoint to this, Kahneman (1999), in a chapter that has been influential among psychologists who study well-being, has urged a discounted sum over time of affect (momentary experienced happiness) as the appropriate measure of overall individual welfare.[13] A *prima facie* case can be made for each of these views. Both subjective well-being and utility are based on trusting an individual's own judgment, but different judgments are trusted in each case: as a welfare measure, lifetime utility puts trust in an individual's (conscious and subconscious) judgments as reflected in choices, while the discounted sum of affect puts trust in an individual's (largely subconscious) judgments as expressed in feelings.

It would be very convenient if flow utility and affect were essentially equivalent; in that case the standard economic measure of individual welfare would match Kahneman's (1999) proposed measure of individual welfare.   One problem with this proposal is that flow utility is not a tightly-defined concept. *Lifetime* utility, anchored in revealed preference, is defined up to a monotonically increasing transformation.   By contrast, anything that adds up to a valid representation of lifetime utility can be considered flow utility. Even with the restriction that flow utility at time *t* should be measurable according to information available at time *t*, there are many candidates.

In the simple case where flow utility is taken as a function primarily of current consumption and leisure, as is common in macroeconomic applications, there are at least two serious problems in equating happiness and flow utility:

---

[13] Kahneman calls momentary affect "instant utility," but here, to avoid confusion, it is best to reserve the term "utility" for the concept of overall individual welfare in Economics.

1. The Easterlin Paradox: There is a dramatic upward trend in consumption, and in many countries a small upward trend in leisure, while felt happiness has no strong trend. (See Easterlin 1974, 1995, 2003.)
2. Hedonic Adaptation: Movements in consumption are extremely persistent, (and even movements in leisure are often moderately persistent), while happiness seems to be very strongly mean-reverting.

The phenomenon of strongly mean-reverting happiness is called hedonic adaptation.
Some of the evidence for hedonic adaptation is surveyed in Frederick and Loewenstein (1999). In response to discrete negative events with lasting practical consequences, significant hedonic adaptation over time is observed for incarceration (Zamble and Proporino, 1990; Zamble, 1992), the loss of the use of limbs, (Wortman and Silver, 1987) and for serious burns (Patterson, et al., 1993). The death of a spouse seems to have a particularly long-lasting effect on affect. But Kaprio, Koskenvuo, and Rita's (1987) finding that suicide rates the week after a spouse's death are elevated almost tenfold for women and almost seventyfold for men suggests especially low affect immediately after the loss, which then moderates to some extent. In addition, Kimball, Helen Levy, Fumio Ohtake and Yoshiro Tsutsui (2006) find direct time series evidence on the University of Michigan Surveys of Consumers for mean reversion of the average happiness of nationally representative samples after important news events.

Some of the most striking data is that on lottery winners. Less than a year after winning the lottery, Brickman, Coates and Janoff-Bulman (1978) find that winners of large lotteries displayed only slightly higher life satisfaction. Frederick and Loewenstein (1999) interpret this as evidence suggestive of substantial hedonic adaptation since it is likely that many winners of large lotteries are ecstatic immediately after winning. More recently, Gardner and Oswald (2001) look at people receiving a windfall--primarily lottery winners--in the British Household Panel Survey. They find that winning £10,000 raises affect by *six times* as much in the first year as £10,000 per year in additional income. This comparison is suggestive of income having been subject to greater hedonic adaptation than the hedonic adaptation to the relatively recent windfall.

Brickman and Campbell (1971) refer to the implications of hedonic adaptation for the trend in affect the *hedonic treadmill*. Because of the close apparent connection between the Easterlin Paradox and the phenomenon of hedonic adaptation, it seems appropriate to search for a joint explanation.

In order to maintain the equation of flow utility and measured affect in the face of the Easterlin paradox and hedonic adaptation, Neo-Benthamites often argue for strong habit formation, social comparison, or a combination of the two in the form of "external habits." (See for example, Layard, 2005.) An alternative strategy is to argue that people make systematic optimization mistakes (Layard, 2005 and Frey and Stutzer, 2004b) or that they are beset by self-control problems (for example, Gruber and Mullanaithan, 2005). We are sympathetic to the idea that such mechanisms exist, but are not persuaded that these mechanisms operate powerfully enough to justify equating happiness and flow utility. We propose a different model of the relationship between utility and happiness.

Although we make a specific proposal for the relationship between lifetime utility and happiness, we consider posing the question of this relationship as an open-ended empirical question more important than our attempt at an answer. The key point is that utility (strictly speaking *preferences*) and happiness can be measured independently making the relationship between them a standard scientific problem. Appendix A discusses the measurement of happiness, while Appendix B discusses the measurement of utility or preferences. Section III presents our hypothesis about their relationship.

III. An Integrated Theory of Preferences and Happiness

A. Preference for Happiness. This section has the limited aim of making clear (1) what our theory is and (2) that this theory is fully consistent with Ordinalism. The remainder of the paper develops and defends this theory and explores its implications. This order works best since our theory assists in important ways in its own defense. Each of the mathematical elements below points to an element of the subsequent discussion. Nevertheless, we have tried to write so that those who dislike equations can skip this section and still follow the main thread of our argument.

We follow the neoclassical tradition of allowing agents unlimited cognitive capacity to know the knowable, including the probability distributions of future uncertain events, and to calculate optimal solutions to constrained lifetime maximization problems. While we believe the relaxation of this assumption is important, we maintain it here in order to make clear that many apparent anomalies which are attributed to mistakes or inconsistencies in the happiness literature can be accounted for in our theory of rational behavior.

We assume that the agent cares directly only about the joint stochastic process of a (potentially large) vector S of state variables, a vector C of control variables, happiness H and a vector B of other outputs of household production functions (such as health).[14] We can abbreviate the complete vector of ultimate goods as $Z=(S, C, H, B)$. The agent always knows the *current* value of the vector of ultimate goods Z. Moreover, to maintain conceptual clarity for the ordinalist concept of individual welfare in our model despite the full set of other complications we introduce, we assume that the agent's preferences over the ultimate goods obey the axioms for intertemporal expected utility maximization with a complete and consistent set of subjective probabilities and with dynamic consistency. Thus, at every time t, preferences can be represented by

$$V_t = E_t \Omega(Z_0,...,Z_{t-1}, Z_t, Z_{t+1},...),$$

where Ω is a time-invariant function. We take time 0 to be a point early enough in the agent's existence (for example, during the agent's time in the womb) that the time 0 vector of state variables K_0 is a sufficient statistic for the agent's information set at time 0. As noted above, we will refer to the objective function V_t as *lifetime utility* at time t (discounted to time 0), even when the agent's objective function includes events encoded in distant future Z's that will happen after the time of the agent's death. Also, note that "lifetime utility" always refers to the *expected value* of the von Neumann-Morgenstern utility index at a particular time.

---

[14] The letter B is in honor of Gary Becker, who extended the concept of household production to more abstract goods.

Informally, we stated our Preference for Happiness Axiom as *other things being equal, people prefer to feel happy.* Given a subjective intertemporal expected utility representation of preferences, it is easy to formalize this as the intertemporal expected utility index Ω being at least weakly increasing in happiness at all dates.[15]

> **The Preference for Happiness Axiom:** *There exists a subjective intertemporal expected utility index Ω representing the agent's preferences that is weakly nondecreasing in H_t for any date t, holding fixed S_t, C_t, and B_t and the entire vector of ultimate goods Z_τ at all other dates τ.*

Note that only the existence of such an Ω is required. To the extent that it psychologically impossible for happiness to vary independently of *S, C,* and *B,* then over the relevant domain there may be equivalent representations of the same preferences that are decreasing in *H.* For example, suppose there were only one period and that $H= \Phi(S_t, C_t, H_t)$. Then $\Omega(S_t, C_t, H_t)$ with partial derivative $\Omega_H>0$ would be equivalent to the alternative function Ω^* defined by $\Omega^*(S_t, C_t, H_t, B_t) \equiv (S_t, C_t, 2\Phi(S_t, C_t, B_t)-H_t, B_t)$, which has the partial derivative $\Omega^*_H < 0$.

Because the degree to which the Preference for Happiness Axiom bites depends on the extent to which happiness can vary independently of *S, C,* and *H,* it is helpful to have a notation for the (possibly empty) vectors of state and control variables that might affect happiness beyond the effects of *S, C,* and *B* and the effects of the history of lifetime utility on happiness that we discuss below. Let *J* be a vector of state variables and *Q* a vector of control variables that might affect happiness *H* but are not in the vector of ultimate goods. We include the entire past history of *S* and *C* in *J* in case this history affects happiness. We also include in *J* and *Q* all the state and control variables that might affect the outputs of the other household production functions represented by *B.* That is, we assume that *S, C, J,* and *Q* jointly determine *B.*

One of the most important aspects of *J* is that many of its elements may not be observable by the agent. It is one thing to know how happy one is; it is another thing to know *why* one is happy. It makes sense to think of preferences as being over things that are or will be observable, but it is possible that one will never know *why* one's life experiences went as they did. Thus, our theory opens up the possibility of exploring "folk theories" of happiness, "scientifically correct" theories that are recommended by experts and the connection between the two.

We think of *Q* as observable, since it stretches the meaning of a "control" variable if one does not know the setting on the dial of one's controls.

B. The Determination of Happiness. To motivate the model of happiness determination below, let us begin with the observation that—although the relationship between circumstances and happiness is weak in the long run—all the evidence suggests that subjective well-being responds in an intuitive and important way to *news* about objective circumstances. At the most trivial level, subjective well-being rises significantly after experimental subjects find a dime and falls significantly after experimental subjects are given negative test results (e.g., Schwarz, 1987).

---

[15] By this assumption, we are leaving aside the possibility of manic happiness so high that it would be unpleasant to feel.

People's happiness rises immediately after they find out they got tenure, and falls immediately after they find out they did not get tenure. People's happiness rises immediately upon discovering a new love, a new item they plan to buy or a new argument for their paper. People's happiness falls immediately after they lose a loved one, lose a limb or lose a job. The rare exceptions are when losing a job, say, is actually good news perhaps because of generous severance pay coupled with the fact that one was intending to leave anyway. The theoretical outline we propose builds on these observations by positing that a major component of affect depends directly on *news* about objective life circumstances that has arrived over the last few months rather than on the *level* of circumstances.

As noted above, we call the component of happiness due to recent news about lifetime utility *elation*. *Dismay* is the algebraic opposite of elation: *dismay = -elation*. If expectations are rational, standard results about rational expectations imply that *news*—dynamic revisions to rational expectations—will be zero-mean and unpredictable. As a result, *elation*—which is a function of recent news—will be strongly mean reverting. Intuitively, news doesn't stay news for very long. At the psychological level, the initial burst of elation dissipates once the full import of news is emotionally and cognitively processed.

Desmond Morris, at the outset of his wonderful little book *The Nature of Happiness,* writes:

> "The true nature of happiness is frequently misunderstood. It is often confused with contentment, satisfaction or peace of mind. The best way to explain the difference is to describe contentment as the mood when life is good, while happiness is the sensation we experience when life suddenly gets better. At the very moment when something wonderful happens to us, there is a surge of emotion, a sensation of intense pleasure, an explosion of sheer delight—and this is the moment when we are truly happy. Sadly, it does not last very long. Intense happiness is a transient, fleeting sensation. We may continue to feel good for quite a while, but the joyful elation is quickly lost."

Morris's description of "happiness" emphasizes *elation*—the word Morris also uses to describe this type of happiness. The "contentment" he refers to is close to our concept of *baseline mood*, which unlike Morris, we also consider a fully legitimate component of happiness, since both the contentment when life is good and the joy when life gets better are likely to affect measured subjective well-being.

We model the effect of news on happiness through the history of (expected) lifetime utility v_t. Lifetime utility V_t depends on the information at time t about the future, so news is reflected in changes in V. We propose the following three axioms relevant to news and happiness:

The News and Happiness Axioms:

1. Happiness H at time t is a function of the other ultimate goods, *S, C,* and *B*, the additional state and control variable vectors *J* and *Q,* and the realized history of lifetime utility V through time *t*. That is, $H_t = \phi_t(S_t, C_t, B_t, J_t, Q_t, V_t, V_{t-1}, ..., V_0)$.
2. Holding fixed *S, C, B, J, Q* and the *past* history of realized lifetime utility V through time *t*-1, happiness at time *t* is increasing in current lifetime utility V_t. That is, if $V_t' > V_t$, then $\phi_t(S_t, C_t, B_t, J_t, Q_t, V_t', V_{t-1}, ..., V_0) > \phi_t(S_t, C_t, B_t, J_t, Q_t, V_t, V_{t-1}, ..., V_0)$.

3. Holding fixed *S, C, B, J, Q*, and *current* lifetime utility V_t, happiness at time *t* is decreasing in previous realized values of lifetime utility. That is, for any integer $j>0$, if $V_{t-j}'>V_{t-j}$, then

$$\phi_t(S_t,C_t,B_t,J_t,Q_t,V_t,V_{t-1},...,V_{t-j}',...,V_0) < \phi_t(S_t,C_t,B_t,J_t,Q_t,V_t,V_{t-1},...,V_{t-j},...,V_0).$$

Remarks: These axioms are all ordinal. They would not be changed in meaning by monotonically increasing transformations of V and Φ. To the extent the history of lifetime utility matters, it is only the history of which indifference curves for (expected) lifetime utility the agent has been on. Furthermore, the axioms only depend on being able to distinguish more or less happiness; they do not depend on the exact scale used for measuring happiness.

Since higher current lifetime utility raises happiness, which in turn raises lifetime utility, there is a multiplier, and an additional technical assumption is needed to insure that lifetime utility is well defined, given the utility index Ω. It is an ordinal version of the assumptions for a contraction mapping.[16] Let $\mathcal{V}$ be the entire time series of V, $\Delta\mathcal{V}$ be a perfectly foreseen vector increment to this time series, let $\mathcal{H}$ be the entire time series of H, let script versions of S, C, B, J, and Q represent the corresponding time series and let Φ be the vector valued function corresponding to the scalar function ϕ. Then $\mathcal{V}$ is the solution to the fixed point problem $V_t = E_t\,\Omega(\mathcal{S}, \mathcal{C}, \mathcal{B}, \Phi(\mathcal{S}, \mathcal{C}, \mathcal{B}, \mathcal{J}, \mathcal{Q}, \mathcal{V}))$, where V_t is the t element of $\mathcal{V}$.

4. (Technical Assumption) Holding $\mathcal{J}$, and $\mathcal{Q}$ fixed, for any nonzero $\Delta\mathcal{V}$, for all t, and all states of nature, $E_t\,\Omega(\mathcal{H}, \mathcal{X}, \mathcal{B}, \Phi(\mathcal{H}, \mathcal{X}, \mathcal{B}, \mathcal{J}, \mathcal{Q}, \mathcal{V}+\Delta\mathcal{V}))$ is either equal to V_t or is strictly between V_t and $V_t+\Delta V_t$.

This technical assumption ensures that, given Ω, and the stochastic processes of the fundamentals, there is only one solution for the stochastic process of lifetime utility.

Although, by assumption, preferences over the ultimate goods *K, X, H*, and *B*, obey the axioms for subjective intertemporal expected utility, the induced preferences over these fundamentals that drive both preferences and happiness involve the agent's expectations. This means that the induced preferences over the fundamentals can exhibit the kind of dependence on a reference point familiar from Prospect Theory (Daniel Kahneman and Amos Tversky, 1979), with the previous period's expectations of lifetime utility serving as a reference point. The loss aversion that is a salient feature of Prospect Theory can be generated by strong concavity of $\Omega(\mathcal{H}, \mathcal{X}, \mathcal{B}, \Phi(\mathcal{H}, \mathcal{X}, \mathcal{B}, \mathcal{J}, \mathcal{Q}, \mathcal{V}))$ in $\mathcal{V}$.[17] In this case, loss-averse behavior would be viewed as a rational response to the nature of the household production function for happiness, not as a reflection of ultimate preferences.[18]

---

[16] Monotonic transformations in either the way lifetime utility or happiness is measured without changing the economic structure would change the functional forms of both Ω and Φ, but would leave the assumption itself intact.

[17] In the March 3, 2006 version of this paper (available at http://www-personal.umich.edu/~mkimball/) we give examples of loss aversion arising from the dependence of ultimate preferences on happiness.

[18]We argue that the persuasiveness of the von Neumann-Morgenstern axioms (or of the modernized versions of these axioms) applies primarily to ultimate preferences. Fully rational agents must take facts about their own psychologies as given. There is no reason to exclude past expectations from the production function for

C. Specializing to Additive Time Separability with Observed State Variables.

For much of our discussion, it is adequate to specialize to the case in which (i) there is an additively time-separable representation of preferences over the ultimate goods, with any intertemporal dependence represented through the state-variable vector and (ii) all state variables relevant to both preferences over ultimate goods and to the production function for happiness are observed. Let K be a concatenation of S and J, while X is a concatenation of C and Q. Then after substituting in the appropriate function of K and X for B, V_t can be written as

$$V_t = E_t \sum_{\tau=0}^{T} \beta^\tau U(K_\tau, X_\tau, H_\tau).$$

(Note that one can allow for a direct dependence of the flow utility function U on time simply by including time in the comprehensive state variable vector K.)

Define the *current value* lifetime utility v_t by

$$v_t = E_t \sum_{j=0}^{T-t} \beta^j U(K_{t+j}, X_{t+j}, H_{t+j}).$$

In addition to discounting from time t rather than time 0, the current value lifetime utility v_t omits the contributions of flow utility that have already happened and can no longer be affected by current or future actions.

In addition, define the (current value) lifetime utility innovation ι_t by

$$\iota_t = \beta^t [V_t - V_{t-1}] = v_t - E_{t-1} v_t = v_t - \beta^{-1} [v_{t-1} - U(K_{t-1}, X_{t-1}, H_{t-1})].$$

This means that we can define a function ψ giving happiness as a function of K, X and the history of ι equivalent to ϕ as a function of K, X and the history of V.

$$H = \psi_t(K_t, X_t, \iota_t, \iota_{t-1}, ..., \iota_1, V_0)$$
$$= \psi_t(K_t, X_t, \beta^t(V_t - V_{t-1}), \beta^{t-1}(V_{t-1} - V_{t-2}), ..., \beta(V_1 - V_0), V_0)$$
$$= \phi_t(K_t, X_t, V_t, V_{t-1}, ..., V_1, V_0).$$

This equivalence enables us to interpret Happiness and News Axioms 2 and 3 in terms of the lifetime utility innovations ι. Axiom 2 implies that $\dfrac{\partial H_t}{\partial \iota_t} > 0$. In addition to $\dfrac{\partial H_t}{\partial V_0} < 0$, Axiom 3 implies that $\beta^{-1} \dfrac{\partial H_t}{\partial \iota_{t-j}} > \dfrac{\partial H_t}{\partial \iota_{t-j-1}}$ for any integer j from 0 to t-1. If the discount factor β if close to 1, this is close to saying that for news of the same magnitude, recent news about future events will have a bigger effect on happiness than older news about future events.[19] On the other hand, if there is heavy discounting, with $\beta \ll 1$, then it possible for old news to have large anticipation effects as something one has known for some time would take place gets closer in time. We find

---

psychological quantities. Thus, to the extent psychological quantities such as happiness enter into preferences, past expectations cannot be excluded from the induced preferences that result when the production function for those psychological quantities is substituted into preferences.

[19] This inequality allows the possibility that distant enough lags of lifetime utility innovations could have a negative effect on happiness. Though we do not think this possibility is empirically relevant, we also do not think it should be ruled out *a priori*.

this association of anticipation effects with a high degree of impatience in preferences plausible.[20]

Finally, we argue that, other than for the application of Happiness and News Axiom 3, it is reasonable to include the initial value of lifetime utility in the comprehensive history state variable vector K. We interpret v_0 as the view of lifetime utility in the instant before birth begins, when the individual has no information about her or his life prospects other than the information that is embodied in genes and body structure at that point. Because the individual's information set is biologically limited up until birth, it is appropriate to view v_0 as an element of the state variable vector that is not subject to subtle expectational effects. In other words, we argue that John Rawls's (1971) "veil of ignorance" about one's station in society is lifted gradually throughout childhood. One is not born knowing that one's family is rich or that one's family is poor. That realization comes later—often late enough that there is a distinct memory of the moment of realization. Similarly, one may not discover whether one will be attractive to desired sexual partners until after puberty. Arrivals of news about this aspect of lifetime utility could account for some of the volatility of affect for teenagers.

After this inclusion, and the inclusion of the date t in the state vector K_t, we can write happiness as

$$H_t = \psi(K_t, X_t, \iota_t, \iota_{t-1}, ...).$$

Happiness H_t depends on the comprehensive state variable vector K_t, the current control variable vector X_t, and the history of lifetime utility innovations. This is our essential claim about the nature of happiness given an additively time-separable intertemporal expected utility function.

We wish to emphasize again that in the absence of additional structure, we consider happiness, like utility, is an ordinal concept. The ordinality of happiness means that it can also be represented by any monotonically increasing function of ψ, as long as one is consistent. In settings that have additional structure, there may be one representation that is more convenient than others. For example, whenever utility is additively separable in happiness as well as additively time-separable, it make sense for us to establish the convention (used below) of measuring happiness as equal to the additively separable happiness term that is, in any case, a monotonically increasing function of happiness.

Given $H_t = \psi(K_t, X_t, \iota_t, \iota_{t-1}, ...)$, we define baseline mood M_t as the level of happiness that would prevail in the absence of any surprises:

---

[20] The prediction that anticipation effects should be associated with high levels of impatience gains more traction if we consider a simple extension of the model to allow for dynamic inconsistency. As is common in economic models of internal conflict, imagine that several agents or selves share the same body and play an internal game to determine actions. To make the extension simple, we assume the model of happiness above applies for each agent or self. Each agent is aware and watching at every point in time, but is does not have full control of what the whole person does. This kind of story is more plausible for a dual-self model, such as Drew Fudenberg and David Levine (2004) than in the typical hyperbolic discounting model, such as in Laibson (1997). As long as each agent's happiness has a positive effect on the happiness reports the whole person makes, anticipation effects can arise when one of the agents is quite impatient and therefore focused on the immediate future.

$$M_t = \psi(K_t, X_t, 0, 0, \ldots) = M(K_t, X_t).$$

We define elation e_t as the difference suprises make to happiness:

$$e_t = d(\psi(K_t, X_t, \iota_t, \iota_{t-1}, \ldots), \psi(K_t, X_t, 0, 0, \ldots)) = e(K_t, X_t, \iota_t, \iota_{t-1}, \ldots),$$

where d is a continuous directed distance measure, equal to zero when its arguments are equal, monotonically increasing in the first argument and monotonically decreasing in the second argument. It will often be convenient to choose d as the simple difference, so that

$$e_t = \psi(K_t, X_t, \iota_t, \iota_{t-1}, \ldots) - \psi(K_t, X_t, 0, 0, \ldots),$$

but if this is the definition of elation for the benchmark happiness scale, a monotonic transformation in the representation of happiness will require shifting to a more complex directed distance measure of the transformed counterpart to ψ to maintain substantive equivalence. If elation is defined by a simple difference, then we can write

$$H_t = M(K_t, X_t) + e(K_t, X_t, \iota_t, \iota_{t-1}, \ldots).$$

If a more complex directed distance measure d is used, inverting d in its first argument yields the equation $H_t = h(M(K_t, X_t), e(K_t, X_t, \iota_t, \iota_{t-1}, \ldots))$, where the function h is monotonically increasing in both arguments, and $h(M,0)=M$.

In the following two sections, we argue that this integrated model of preferences and happiness—which for brevity we will sometimes call the "elation theory of happiness"—encodes many of the most salient empirical facts about happiness.

IV. Baseline Mood

A. The Definition of Baseline Mood. Conceptually, there are two key equations from Section III. (Both are from the simplified case of additively time-separable utility.) First,

$$v_t = E_t \sum_{j=0}^{T-t} \beta^j U(K_{t+j}, X_{t+j}, H_{t+j}),$$

where v is lifetime utility, E_t is an expectation as of time t, β is a discount factor between 0 and 1, U is the flow utility function, H is happiness, X is a vector of control variables that the agent decides on from one moment to the next, and K is a vector of state variables that might take time to alter, if they can be changed at all. That is, *a typical individual cares about happiness, but also cares about other things*. Second,

$$H_t = M(K_t, X_t) + e(K_t, X_t, \iota_t, \iota_{t-1}, \ldots).$$

In words, happiness is the sum of baseline mood M and elation e. The vector of state variables K and control variables X is extensive enough that includes all of the determinants of baseline mood. Elation depends, in addition, on recent news about lifetime utility. Indeed, this is what distinguishes elation from baseline mood. By definition, baseline mood is the part of happiness *not* due to recent news about lifetime utility. To put it another way, baseline mood is what happiness would be if the events that actually occurred in an individual's life had been predictable. This means that

1. *Any predictable aspect of happiness is part of baseline mood.* This includes any persistent aspect of happiness.

2. Any aspect of happiness that *would* be predictable if one were able to predict the relevant values of K_t and X_t is a part of baseline mood. The best example here is the effect of one's current activity on happiness. On average, one can predict that people are happier when eating dinner than when doing the dishes, even if one does not know in advance when these things will happen.[21]

Thus, baseline mood is a *moving* baseline that accounts for the straightforward effects of current activities on happiness. Our model predicts movements in happiness even under perfect foresight. Although persistent cross-sectional heterogeneity associated with genes and long-lasting personality characteristics are likely to represent a substantial share of the variance of baseline mood, daily, weekly and annual cycles and movements associated with controllable aspects of time usage (such as time spent commuting in traffic versus time spent getting needed sleep), are also likely to be important.

B. Baseline Mood as the Output of a Household Production Function. Since Gary Becker's (1965) pioneering work, much of the activity of a household outside of paid work has been reconceived as household production of goods. The dependence of baseline mood on things wholly or partially under the individual's control makes it useful to think of baseline mood as the output of a household production function. From this perspective, physical health provides a good analogy for baseline mood. Like health, baseline mood

- can be measured independently of its arguments (inputs);
- is only one argument of the flow utility function;
- depends on different things than flow utility does—or on the same things with different; weights
- has a complex household production function or subutility function.
- has a concave relationship to income.

Ultimately, it is an empirical matter what baseline mood depends on, but provisionally, we view baseline mood as depending on factors such as:
a. genes[23]
b. psychologically active drugs, such as Prozac
c. sleep
d. exercise[24]
e. eating habits
f. time spent with friends[25]
g. social rank[26]
h. the pleasantness of one's current activity.[27]

---

[21] See Kahneman, Krueger, Schkade, Schwarz and Stone (2004) on the average level of affect experienced during different activities. As one unsurprising example, people experience higher affect while eating than the affect they experience while doing housework.

[23] See Diener and Lucas (1999).

[24] See Thayer (1989), Biddle and Murtrie (1991), Steptoe, Kimbell and Basford (1996) and Argyle (1999).

[25] See Lewinsohn, Sullivan and Grosscup (1982), Reich and Zautra (1981) and Argyle (1999).

[26] See Luttmer (2004).

[27] See Kahneman, Krueger, Schkade, Schwarz and Stone (2004).

Viewing baseline mood as one of the arguments of flow utility allows the powerful language of price theory to be applied to baseline mood, just as to health. For example, Hall and Jones (2004) argue that health is a luxury good in the sense that continuing increases in per capita income will increase the budget share devoted to health-related expenditures. Similarly, one might argue that continuing increases in per capita income are likely to increase the budget share devoted to baseline-mood-related expenditures.[28]

C. Do People Know the Production Function for Baseline Mood? A key limitation on our ability to apply price theory to baseline mood is the possibility that people may not have accurate knowledge of the production function for baseline mood. People's expenditures of time and money will depend on their beliefs about the production function for baseline mood rather than the true function. Pursuing the analogy to health again, it seems reasonable that, just as people don't know the true production function for health, they may not know the true production function for baseline mood. In principle, the discovery and dissemination of facts about the determinants of baseline mood could have large positive welfare effects.[29]

One factor that could make it especially difficult for people to figure out the determinants of baseline mood is the salience of the component of happiness due to elation. Although the elation mechanism has its own functions, from the standpoint of figuring out the determinants of baseline mood, elation acts as noise.

D. Applying Price Theory to Baseline Mood.

1. A nonjudgmental view of the negative correlation between materialism and happiness.
To the extent that people do understand the determinants of baseline mood, price theory can contribute in important ways to an understanding of long-run happiness. Consider, for example, the negative correlation that has sometimes been found between "materialism" and happiness. Robert Lane (2000) gives a discussion of the mixed empirical evidence for such a negative correlation. In assessing the evidence, it is also important to be aware of the partial tautology in relating measures of unhappiness to materialism indices that contain many survey items measuring dissatisfaction and griping. Nevertheless, in order to make the logical point as clearly as possible, suppose it could be documented conclusively that materialism, in the narrow sense of valuing material goods highly, lowers happiness. Price theory suggests that as long as there is any tradeoff between happiness and material goods, those who value material goods more compared to happiness will choose a bundle with more material goods (as often found for those who are more materialistic) and less happiness. (This is called an "equalizing difference" in the labor economics literature. See for example Rosen, 1986.) The mechanics of the tradeoff could,

---

[28] The hypothesis that in the future of rich countries baseline mood will be a luxury good is inspired by Maslow (1943), who argues that once basic needs (such as physiological and safety needs) are satisfied, higher needs (such as needs for love, belonging, esteem and actualization) come to the fore. Both long-run happiness at home and long-run happiness at work might exhibit strong income effects. However, one bit of evidence running contrary to this idea that baseline mood is a luxury good is that in the Hindhu and Buddhist traditions a great deal of time and effort were often devoted to baseline-mood-raising meditation even thousands of years ago at much lower levels of per capita income than today.
[29] This view of the value of pinning down the determinants of baseline mood is consistent with the program of Positive Psychology, as described by Seligman (2002).

for example, be due to decisions such as the decision of whether to commute further to a higher paying job discussed in Section VII. Materialism lowering happiness would be similar to the effect preferences have on any choice between two distinct goods—such as when those who place an extremely high value on career success have worse physical health because they do not make time to exercise or see the doctor.

2. Explaining the Easterlin Paradox. Another important application of price theory is to the Easterlin Paradox itself. Even after accounting for the elation mechanism, since baseline mood is likely to be a normal good, there is still a version of the Easterlin Paradox that we must confront. With people much richer now, why don't they purchase more baseline mood? Trends in the externalities related to (a) social rivalry that focuses people on their relative rather than their absolute standing,[30] (b) declines in social capital, as suggested by Robert Putnam (2000) and (c) the side-effects of other people's use of the greater freedom that comes with extra income (for example, the side-effects of a higher divorce rate, with their side effects for the children of divorce) and (d) any exacerbations of internal conflicts (for example the part of the rise in obesity that results from the ability to buy more and better food) can certainly contribute toward an explanation, since most of these externalities and internal conflicts are likely to figure into happiness at least as strongly as they figure into utility. (e) Lack of knowledge of the true production function for baseline mood could also contribute in an important way toward explaining this version of the Easterlin Paradox. Moreover, (f) some of the extra resources people have are being used to lengthen life expectancy; that is, many resources are used to lengthen the duration of happiness for an individual rather than the intensity. But there may also be a price-theoretic element to the explanation. (g) *Although income has gone up, the price of baseline mood may have risen.* The most likely reason for this is if many of the inputs into baseline mood are time-intensive, such as exercise or time spent with friends. With the price of baseline mood higher, people may choose to expand their consumption of other goods rather than baseline mood. The greater people's willingness to substitute between baseline mood and other goods, the smaller the price rise necessary to explain the Easterlin Paradox. To put it plainly, although income has gone up, people still have only 24 hours in a day, and time is the main thing it takes to raise happiness. In this, happiness would be like many other labor-intensive or time-intensive goods subject to Baumol's cost disease.[32] The pervasive effects of a high dollar value of time on behavior are the central theme of Staffan Linder's (1970) wonderful book *The Harried Leisure Class*. These effects are likely to matter for the long-run level of happiness.

V. Elation, Neurobiology and Evolution

A. Evidence that Expectations Matter for Affect. One of the central predictions of the elation theory of happiness is that expectations will matter for felt happiness, since the lifetime utility

---

[30] Television may have enhanced the negative effect of social rivalry on happiness by leading people to believe the distribution of income and other advantages in society is higher than it actually is, leading people to underestimate their true social rank. See O'Guinn and Shrum (1997).

[32] See for example James Heilbrun (2003) and William Nordhaus (2006).

innovations are given by $\iota_t = v_t - E_{t-1}v_t$, and elation is a function of current and past lifetime utility innovations. The importance of expectations for happiness is indicated by the evidence surveyed in Frederick and Loewenstein (1999) that advance notice of the death of a spouse reduces the size and duration of the drop in affect after the actual death of the spouse. The following passage from Frederick and Loewenstein (1999, p. 315) is especially close to the spirit of the model here: "Even if advance notice does improve post-outcome well-being, its *overall* effect on well-being is ambiguous, since receipt of the bad news may diminish the well-being of the person between the time the notice is received and the time the event actually occurs." In the model here, it is the processing of bad news that generates a period of lower affect, whether the primary bad news occurs before the actual death of the spouse or only at the time of the actual death.

Camerer, Loewenstein and Prelec (2005, p. 28) give a good summary of some remarkable neurobiological research relevant to the role of expectations in determining affect:

> An important feature of many homeostatic systems is that they are highly attuned to changes in stimuli rather than their *levels*. A dramatic demonstration of such sensitivity to change came from single-neuron studies of monkey responding to juice rewards (see Wolfram Schultz and Anthony Dickinson 2000). These studies measured the firing of dopamine neurons in the animal's ventral striatum, which is known to play a powerful role in motivation and action. In their paradigm, a tone was sounded, and two seconds later a juice reward was squirted into the monkey's mouth. Initially, the neurons did not fire until the juice was delivered. Once the animal learned that the tone forecasted the arrival of juice two seconds later, however, the same neurons fired at the sound of the tone, but *did not* fire when the juice reward arrived. These neurons were not responding to reward, or its absence … [ellipses and all italics in original] they were responding to deviations from expectations. (They are sometimes called "prediction neurons.") When the juice was expected from the tone, but was not delivered, the neurons fired at a very low rate, as if expressing disappointment.

These results are just the tip of the iceberg in the neurobiology literature. A great deal of evidence points to machinery in the human brain that generates sophisticated short-run expectations—expectations that people are not always consciously aware of. See for example John O'Doherty et al. (2003), Jay Gottfried, O'Doherty and Raymond Dolan (2003), Ben Seymour et al. (2004), Seymour et al. (forthcoming) and O'Doherty (2005).[33]

B. The Evolutionary Significance of Elation. Though any such claim is highly speculative at this point, we are inclined toward Randolph Nesse's (2000, 2001, 2004, forthcoming) functional interpretation of affect as part of the motivational system for processing utility-relevant information. If something good happens, elation motivates the individual to think about what went right (in case there is a way to make it happen again) and how to take advantage of any new opportunities that may have arisen. If something bad happens, dismay (negative elation) motivates the individual to think about what went wrong (in case there is a way to avoid it in the future), and how to mitigate the harm of the new situation. On this view, elation and dismay are in the same genus as *curiosity,* which is part of the motivational system for processing information that is neither obviously good nor bad, but for which there may be value to finding

---

[33] More recently, Hackjin Kim, Shinsuke Shimojo and John O'Doherty (2006) find that avoiding an aversive outcome is represented in the brain in the same way as receiving a reward.

out more. Indeed, experimental inductions of elated and depressed moods have been found to change individuals' strategy of information processing across a variety of tasks (for reviews see Schwarz, 1990, 2002 and William Morris, 1999). Elated people are especially good at seeing opportunities, while dismayed people are especially good at seeing dangers.

C. The Evolutionary Significance of Hedonic Adaptation. Thinking of a temporary jump in affect occurring after utility-relevant news as functionally related to information-processing makes the functional significance of hedonic adaptation similar to the functional significance of adaptation in other aspects of perception. Frederick and Loewenstein (1999, p. 303) make this comparison explicit:

> "Adaptive processes serve two important functions. First, they *protect* organisms by reducing the internal impact of external stimuli…. Second, they *enhance perception* by heightening the signal value of changes from the baseline level…."

> "Hedonic adaptation may serve similar protective and perception-enhancing functions…. persistent strong hedonic states (for example, fear or stress) can have destructive physiological concomitants … Thus, hedonic adaptation may help to protect us from these effects."

> "Hedonic adaptation may also increase our sensitivity to, and motivation to make, local changes in our objective circumstances…."

Rayo and Becker (2005) construct a formal model that spells out the logic of Frederick and Loewenstein's (1999) claim.

D. Speculations on the Evolutionary Significance of Baseline Mood. Certain kinds of persistent situations could call for heightened sensitivity toward opportunities or toward dangers. For example, moderately high social rank or good physical health may make it safe to look more for opportunities than for dangers. Thus, it could make sense for these situations to stimulate the same machinery that is turned on by the receipt of good news.   The high variance of persistent individual differences in baseline mood suggests a frequency dependence in which there is an advantage to being a pessimist looking for dangers when most of the surrounding people are optimists who might miss dangers, while there is an advantage to being an optimist who sees opportunities if there are plenty of pessimists around to alert one to possible dangers, and few other optimists around to boldly seize opportunities.

One of the most interesting possibilities is that important aspects of the determination of baseline mood are just quirks in the affective system that have no functional significance. The mixed-strategy evolutionary equilibrium in which the fitness of moderately happy and moderately unhappy people is equal would reduce the strength of any evolutionary pressure against such quirks.

Regardless of how the "production function" for baseline mood arose, now that it is present, it makes sense to exploit it, just as Stephen Pinker (1997) argues that we exploit our sense of taste (designed, say, to motivate the search for nuts and ripe fruits) with cheesecake and our musical sense (designed, say, to help us distinguish the sounds of different kinds of objects) with symphonies and Rock and Roll.

E. Implications of the Integrated Framework for Utility and Happiness. There are three key implications of this benchmark model for the relationship between happiness and utility. First, there is a clear distinction between the psychological concept of happiness and the economic concept of flow utility. Happiness is *not* equal to either flow utility or to the overall objective function.

Second, the elation component of happiness depends primarily on unexpected *changes* in lifetime utility. For applications, the most important point about elation is that the theory here contradicts the notion that a temporary movement in happiness is unimportant because of its short duration. To the contrary, *a temporary movement in happiness may be extremely important as a signal of important utility-relevant news related to the long-term welfare of the individual.*

Third, baseline mood, while not a summary measure of flow utility, is something that people care about. As with health, the relative concern with raising baseline mood compared to raising consumption of other goods may increase along with per capita income, implying that the average share of effort and expenditures devoted to raising baseline mood may increase in the future.

Since elation depends on (mean-zero) *news* about lifetime utility, rather than on the *level* of lifetime utility, elation has no trend. Thus, utility can rise with per capita income while happiness has only the trend imparted by the growth rate of baseline mood. This guarantees that the economic concept of lifetime utility and the psychological concept of the temporal sum of happiness over time put forward by Kahneman (1999) will be numerically distinct approaches to assessing overall welfare. Distinguishing clearly between utility and happiness allows scientific questions about utility and happiness to proceed in a way that respects the insights of both Psychology and Economics without prejudging the ethical question of the proper contribution of each concept to the assessment of overall welfare–an ethical question that revolves fundamentally around the extent to which one should trust people's immediate feelings and the extent to which one should trust people's choices as indications of what most enhances their welfare. In this ethical debate, traditional Welfare Economics has implicitly staked out a position in favor of utility as the better measure of overall welfare, but the case for Kahneman's (1999) proposal deserves to be thoughtfully considered as well.[34]

Maintaining a clear distinction between happiness and flow utility also makes it possible to see where the psychological approach toward welfare assessment and the economic approach toward welfare assessment are pulling in the same direction. For example, social rank—whether appearing as an effect of other people's consumption or time use on baseline mood or on flow utility directly—will matter for both the psychological and economic measures of overall welfare. As another example, as long as baseline mood is an argument of the flow utility function, any advance in scientific understanding of determinants of baseline mood, and the dissemination of scientific knowledge about baseline mood to individuals in society will be important for both measures of overall welfare.

---

[34] The strength of Kahneman's case depends in important measure on whether, as he argues, there is no way to construct a consistent underlying set of preferences from the contradictory decisions people make, even after following the approaches discussed above in Section IV, "Measuring utility."

VI. Happiness in the Utility Function: The Easy Case

As discussed in Section III, preference for happiness, together with a dependence of happiness on recent news about lifetime utility can generate reference dependence and loss aversion despite preferences over ultimate goods that obey the axioms for intertemporal expected utility. However, there is one special case in which happiness in the utility function has more limited effects: when (1) the utility function is additively separable between happiness and other goods and (2) this additively separable function of happiness depends linearly on lifetime utility innovations. In that case, without loss of generality, one can scale happiness so that the additively separable term can be represented as linear in happiness, yielding

$$v_t = E_t \sum_{j=0}^{T-t} \beta^j U(K_{t+j}, X_{t+j}, H_{t+j})$$

$$= E_t \sum_{j=0}^{T-t} \beta^j [u(K_{t+j}, X_{t+j}) + H(K_{t+j}, X_{t+j}, \iota_{t+j}, \iota_{t+j-1}, ...)]$$

$$= E_t \sum_{j=0}^{T-t} \beta^j [u(K_{t+j}, X_{t+j}) + M(K_{t+j}, X_{t+j}) + \sum_{\ell=0}^{n} a_\ell \iota_{t+j-\ell}].$$

This special case is still rich enough that it allows us to discuss two important issues: (A) to give an explanation for why utility and happiness are often confused and (B) to show that it is sometimes permissible for research on household behavior to ignore the presence of happiness in preferences. But when happiness data are available, they can provide valuable data about preferences, even when it is *permissible* to ignore happiness. In addition, we will be able (C) to discuss whether manipulating one's information structure can add to utility and (D) to show by example that *mistakes about the determination of happiness do not always cause mistakes in one's choices.*

A. Elation Theory and the Confusion Between Utility and Happiness

Any adequate account of the relationship between utility and happiness must explain why these two concepts are often confused. Why is it that the word "happiness" has a meaning in the dictionary that is very close to what we are calling "lifetime utility" as well as the meaning referring to positive feelings that we are using? To answer this question, it is useful to compare maximizing lifetime utility with Kahneman's (1999) proposal of maximizing the true mathematical expectation of the present discounted value of happiness[35] in the context of the theory presented above.

1. Maximizing the present discounted value of happiness versus maximizing lifetime utility.
To the extent that baseline mood is different from flow utility and to some extent controllable, maximizing the expected present discounted value of happiness as Kahneman (1999) recommends will be different on that account from maximizing lifetime utility. But what about maximizing the expected present discounted value of happiness when baseline mood is beyond

---

[35] The extension of Kahneman's proposal to the true mathematical expectation in uncertain situations is not explicit in Kahneman (1999), but it seems a reasonable interpretation.

the individual's control? In that case only elation will matter in maximizing the presented discounted value of happiness. Proposition 1 addresses this case:

Proposition 1: Given (i) rational expectations, (ii) perfect memory, (iii) happiness that is the sum of baseline mood and elation, (iv) baseline mood that is exogenous to the individual, and (v) elation that is a positive linear combination of lifetime utility innovations, as of time t, maximizing the expected present discounted value of happiness is equivalent to maximizing lifetime utility.

Proof: Let elation e_t be given by

$$e_t = \sum_{\ell=0}^{n} a_\ell \iota_{t-\ell}.$$

Then the expected present discounted value of happiness is

$$E_t\left\{\sum_{j=0}^{T-t}\beta^j A_{t+j}\right\} = E_t\left\{\sum_{j=0}^{T-t}\beta^j M_{t+j} +\sum_{j=0}^{T-t}\beta^j e_{t+j}\right\} = E_t\left\{\sum_{j=0}^{T-t}\beta^j M_{t+j} + \sum_{j=-n}^{T-t} b_{j,t}\iota_{t+j}\right\},$$

where

$$b_{j,t} = \sum_{\ell=-j}^{n}\beta^{j+\ell}a_\ell$$

(as long as time t is at least n periods away from death, and somewhat less if t is less than n periods from death). Using the definition of lifetime utility innovations, perfect memory and the fact that the expectation of lifetime utility innovations conditional on previous information is zero, one can simplify the expected present discounted value of happiness further, to

$$E_t\left\{\sum_{j=0}^{T-t}\beta^j A_{t+j}\right\} = \sum_{j=0}^{T-t}\beta^j E_t M_{t+j} + b_{0,t}(v_t - E_{t-1}v_t) + \sum_{j=-n}^{-1} b_{j,t}\iota_{t+j}.$$

Given the exogeneity of baseline mood M and the perspective of time t, everything in this expression is fixed except for $b_{0,t}\, v_t$. Thus, maximizing the expected present discounted value of happiness is equivalent to maximizing $b_{0,t}\, v_t$, which in turn is equivalent to maximizing v_t.[36]

2. Maximizing current happiness. Note that under the assumptions of Proposition 1, maximizing *current* happiness alone is also equivalent to maximizing lifetime utility, since

$$A_t = M_t + e_t = M_t + \sum_{\ell=0}^{n} a_\ell \iota_{t-\ell} = M_t + a_0 v_t - a_0 E_{t-1}v_t + \sum_{\ell=1}^{n} a_\ell \iota_{t-\ell}.$$

---

[36] Note that only exogeneity of the conditional *mean* of baseline mood is needed for this result. An ability to control the variance of baseline mood, with no effect on the mean, would still leave elation totally dominant in the expected present discounted value of happiness.

Given the assumed exogeneity of baseline mood M_t, the only thing that is not fixed in this expression as of time t is the term $a_0 v_t$, so one does the same thing to maximize current happiness as to maximize lifetime utility. The reason a present discounted value of happiness is not required is that elation is already forward-looking.[37]

3. Why utility and happiness are often confused. Psychological evidence is accumulating that baseline mood can in fact, be modified deliberately—and in ways that go beyond the zero-sum game of acquiring social rank. But a lack of understanding of the determinants of baseline mood can make baseline mood seem exogenous. As noted above, one reason for this lack of understanding may be that a large fraction of the time-series variance of happiness may be accounted for by elation and dismay. To the extent that elation and dismay dominate people's perception of happiness, Proposition 1 indicates why people might think that utility and happiness are essentially the same thing.

It is when people do begin to recognize that baseline mood might be controllable that the distinction between utility and happiness becomes crucial. Understanding the ways in which baseline mood is controllable clearly matters for optimization. Understanding the distinction between utility and happiness is becoming important precisely because we are beginning to see a wider variety of ways to raise *utility* by raising *happiness* rather than being limited to raising *happiness* (temporarily) by raising *utility*.

B. Elation Independence. If happiness is additively separable in the utility function, and elation is a linear function of lifetime utility innovations, then preferences do not depend on elation. That is, given a utility function with additively separable happiness, if elation is linear in lifetime utility innovations, preferences over fundamentals will be the same as if elation were always zero. Proposition 2 addresses the essential point. Here, think of the initial utility function U as what one gets when happiness is additively separable *and* not sensitive to lifetime utility innovations:

$$U(K_t, X_t) = u(K_t, X_t) + M(K_t, X_t).$$

(Note that we are following the convention established above of measuring happiness by the additively separable happiness term in the period utility function, which in any case is a monotonically increasing function of happiness.)

Proposition 2: Given rational expectations, adding to the flow utility function a linear function of lifetime utility innovations (with positive coefficients summing to less than one) has no effect on the preferences represented by the utility function.

Proof: Using an asterisk to represent the modified flow utility and lifetime utility functions, let

$$U^*(K_t, X_t, \iota_t, \iota_{t-1}, \ldots) = U(K_t, X_t) + \sum_{\ell=0}^{n} a_\ell \iota^*_{t-\ell},$$

---

[37] In an analogy to exotic financial securities due to George Akerlof when he first heard our definition of elation, elation provides a kind of *tranche* of current and future effects on flow utility.

where Happiness and News Axiom 4 requires that $\sum_{\ell=0}^{n} a_\ell < 1$. Note that the relevant lifetime utility innovations will be those for the modified lifetime utility function. Modified lifetime utility is then

$$v *_t = v_t + E_t \sum_{j=-n}^{T-t} b_{j,t} \iota *_{t+j},$$

where, as above, $b_{j,t} = \sum_{\ell=-j}^{n} \beta^{j+\ell} a_\ell$. The essential structure here is that modified lifetime utility $v*_t$ is equal to the original lifetime utility v_t plus the expected value of a linear combination of the modified lifetime utility innovations with positive coefficients running from n periods back, up to the lifetime utility innovation in the agent's last period. Because lifetime utility innovations have mean zero conditional on previous information, one can simplify this further to

$$v *_t = v_t + b_{0,t} \iota *_t + \sum_{j=-n}^{-1} b_{j,t} \iota *_{t+j} = v_t + b_{0,t}(v *_t - E_{t-1} v *_t) + \sum_{j=-n}^{-1} b_{j,t} \iota *_{t+j}.$$

The condition that $\sum_{\ell=0}^{n} a_\ell < 1$ guarantees that $b_{0,t}<1$. Therefore, one can solve for $v*_t$:

$$v *_t = \frac{v_t - b_{0,t} E_{t-1} v *_t + \sum_{j=-n}^{-1} b_{j,t} \iota *_{t+j}}{1 - b_{0,t}}.$$

Because $-b_{0,t} E_{t-1} v_t + \sum_{j=-n}^{-1} b_{j,t} \iota *_{t+j}$ is fixed as of time t, as a representation of preferences over choices at time t, $v *_t$ is equivalent to $v_t/(1-b_{0,t})$, which in turn is equivalent to v_t itself.

To recap the proof, when a linear combination of lifetime utility innovations is added to the lifetime utility function, (1) the future lifetime utility innovations do not affect decisions because their expectation is zero, (2) the past lifetime utility innovations do not affect decisions because they are predetermined and (3) the current lifetime utility innovation does not affect decisions because, to the extent it is not predetermined, it is perfectly correlated with the original lifetime utility function.

The reason we think elation-independent preferences are important is that they provide an example in which standard economic approaches that ignore happiness will obtain results consistent with those of an approach that does incorporate happiness. In this special case, the only objection to standard methods is that these standard approaches would be ignoring an important source of useful data.

C. Manipulating the Timing of News and Manipulating Expectations. In the proof of Proposition 2, have we tacitly assumed a fixed information structure? Does adding elation that is

a positive linear combination of lifetime utility innovations to the utility function affect preferences over information structures, even when both information structures would lead to the same decisions over other variables? For example, could it make people want to delay when they hear news in order to manipulate their own feelings? The answer is no. Because rational expectations take into account the information structure, there is no way to game the system with any rule set up in advance. Suppose for example, that you told your friend to tell you good news right away, but to withhold bad news. The Bayesian inference in rational expectations would cancel out any effect on the expected lifetime utility innovation, though it would certainly affect the *ex post* distribution of lifetime utility innovations. Formally, v^*_t can be expressed as a linear function of v_t and *past* expectations about lifetime utility. Choosing a different information structure *now* can only affect current and future expectations about lifetime utility. That includes choosing an information structure when your friend says "I know what happened, do you want me to tell you or not?" since any revelation is still in the future, if in the near future.

Of course, even a mean-zero effect on the distribution of lifetime utility innovations will affect lifetime utility when added elation is a nonlinear function of lifetime utility innovations or elation enters the utility function nonlinearly, as we discuss below. Also, imperfect memory of past expectations may provide an opening for gaming the system by trying to reduce one's remembered past expectations. This may be particularly relevant for the memories of past expectations parents transmit to a child about the child's prospects: the gap between parent and child can be one source of imperfect memory in a dynasty. More generally, an attitude of gratitude (whose value is not diminished by the triteness of the phrase) can serve the same purpose as manipulable memory. It often involves substituting comparisons with others in a worse situation for comparisons with one's own remembered past expectations or one's own deductions of what one ought to have expected in the past.

Given perfect memory, but irrational expectations, it may be harder to beneficially manipulate expectations than one might at first think, since then lowering one's expectations adds to flow utility in the future, but subtracts from flow utility now.[38] It is when one can manage high expectations now, but remember them in the future as if they were low expectations that there is a real opening for beneficial manipulation of beliefs.

D. Do Mistakes about the Rate of Hedonic Adaptation Matter? We argued above that because utility and happiness are distinct, the psychological phenomenon of hedonic adaptation does not have any necessary implications for the shape of the utility function. In particular, if flow utility depends only on baseline mood and not on elation, as in the model above, the determination of elation, including the rate of hedonic adaptation, has no effect on the lifetime utility function. Thus, when elation is not an argument of the utility function, misprediction of hedonic adaptation causes no material harm to utility maximization, contrary to the claims of Schkade and Kahneman (1998).

---

[38] Nevertheless, there is evidence people do some of this kind of manipulation of expectations. Nisan (1972) finds that study participants taking an immediate test were less confident than those taking a test in 4 weeks. Similarly, Shepperd, Ouellette, & Fernandez (1996) find that college seniors were more muted in estimated first-job salaries than sophomores and juniors. (See also Shepperd, Findley-Klein, Kwavnick, Walker and Perez (2000).) In each case, confidence was reduced when proximity to performance outcomes was more immediate. We are grateful to Norbert Schwarz for cluing us in to this evidence.

To pursue the question further, consider how much harm there is to mistakes about the rate of hedonic adaptation in the context of the model of subsection A, with a positive linear combination of lifetime utility innovations added to the flow utility function. Mistakes about the rate of hedonic adaptation are mistakes about the true values of the coefficients a_ℓ. Since the modified lifetime utility function is equivalent to the original utility function regardless of the values of the coefficients a_ℓ (as long as they are positive and add to less than one), mistakes about the rate of hedonic adaptation will not distort decisions at all and so will be costless!

In the light of the lack of harm to optimization from misperception of the rate of hedonic adaptation in this benchmark case, any serious claim of quantitatively significant harm to optimization from misperception of the rate of hedonic adaptation would require careful modeling. For example, when elation is a nonlinear function of lifetime utility innovations, or flow utility is a nonlinear function of elation, there is likely to be at least some harm from misperception of the rate of hedonic adaptation, but it is not clear how large this harm would be. In the case of imperfect memory, misperception of the rate of hedonic adaptation might cause one to exert too much or too little effort toward manipulating one's memories, but whether this results in a serious reduction in lifetime utility depends on how great the scope is for manipulation of memory.

One of the most important effects of underestimating the rate of hedonic adaptation is that it will cause an overestimation of the unconditional variance of elation, since the effects of unforeseen increases or decreases in lifetime utility seem like they will be long-lasting. An overestimation of the unconditional variance of elation should, in turn, cause an individual to overestimate the fraction of the variance of happiness due to elation and underestimate the fraction of the variance of happiness due to baseline mood. As shown above, this overestimation of the persistence of elation does not necessarily interfere with maximizing lifetime utility, but it *would* tend to push Kahneman's suggested alternative of maximizing the expected present discounted value of happiness in the direction of maximizing lifetime utility. Since elation embodies movements in lifetime utility, anything that exaggerates the importance of elation in happiness is likely to make maximizing happiness more like maximizing utility, as indicated by the extreme case of Section VI-A, where elation is the only controllable component of happiness.

By implication, this section using elation-independent preferences also hints at some of the additional phenomena that arise when happiness is nonseparable or is nonlinear in lifetime utility innovations. Most important is the fact that even additively separable happiness can give rise to loss aversion if happiness in strongly concave in lifetime utility innovations. Next in importance is the fact that the assumption that preferences are over happiness rather than over baseline mood and elation separately means that reactions in household behavior to lifetime utility can indicate the particular type of nonseparability of happiness in the utility function. Thus, even if one is

preimarily concerned with long-run happiness, short run reactions of happiness to news may help one to identify the shape of preferences for long-run happiness as well.

---

[52] Note that with a finite horizon, the two formally similar versions of the utility function would no longer represent

exactly the same preferences. The lifetime utility function $v_t = E_t \sum_{j=0}^{T-j} \beta^t [f(C_{t+j}) - f(C_{t+j-1})]$ would imply a

greater tendency to consume in the period immediately before death than $v_t = E_t \sum_{j=0}^{T-j} \beta^t f(C_{t+j})$. However, the

lifetime utility function $v_t = E_t \sum_{j=0}^{T-j} \beta^t [f(C_{t+j}) - f(C_{t+j-1})]$ would still be equivalent to the lifetime utility

function $v_t = E_t \left\{ \left[\sum_{j=0}^{T-j-1} \beta^t f(C_{t+j}) \right] + \frac{\beta^{T-t}}{1-\beta} f(C_T) \right\}$. This equivalent form with "flow utility" depending

only on current consumption might easily be more convenient, despite the odd-looking coefficient on $f(C_T)$.

VII. Evidence that Happiness is Not the Same Thing as Flow Utility

Having illustrated some of the basic aspects of our theory, it is appropriate to directly address some of the key controversies about utility and happiness. Let us say at the outset of this discussion that it would be convenient for many reasons if happiness in the narrow sense of current affect were proportional to flow utility. Not only would this make a welfare measure based on a discounted sum of current affect equal to lifetime utility, but such a simple measure of flow utility would make utility empirics more like production empirics, where one actually gets to see output directly. However, as we mentioned briefly in Section II, there are two reasons why it is very difficult to maintain that happiness and flow utility are even close to the same thing: the Easterlin Paradox and Hedonic Adaptation.

A. The Easterlin Paradox. As Easterlin (1974, 1995, 2003) observes, real per capita GDP and real consumption expenditure in the United States has risen dramatically in the last fifty years, but the percentage of people saying they are "very happy" has been falling slightly. The story is even more dramatic in Japan, where the percentage rise in per capita GDP is even more rapid, but the graph of subjective well-being is essentially flat. In other words, in developed societies, profound increases in average real income and in the objective standard of living over the last 50 years have not been associated with increases in the average happiness of their citizens. As argued in Appendix B, "Measuring Utility,", it is not just a matter of money not buying happiness, since there are many other positive trends. In short, there is strong evidence that utility has gone up, but happiness has not.

B. Hedonic Adaptation. **As noted in Section II,** in addition to the difference in the trend behavior of utility and affect, the shorter-run dynamic behavior of affect is also quite different from the dynamic behavior of flow utility as normally modeled. As an empirical matter, the movements in happiness in response to a number of important categories of changes in circumstances are subject to *hedonic adaptation*--regression of affect toward its previous level.

In recent work, Kahneman, Krueger, Schkade, Schwarz and Stone, 2006, seem to have discovered another important fact about hedonic adaptation: measures of *current* happiness such as data from experience sampling or from the day reconstruction method (pioneered by Kahneman, Krueger, Schkade, Schwarz and Stone, 2004) show even stronger hedonic adaptation (mean reversion) than life-satisfaction or global happiness measures. As discussed in Appendix A, "Measuring Happiness," because life satisfaction and global happiness evaluations incorporate an element of autobiography *and* people's ideas about how they "should" feel, they will tend to show more permanent effects of events such as unemployment, as Lucas, Clark, Georgellis and Diener (2004) find.

What is even more serious, the likely influence of people's folk theories of how they "should" feel on life satisfaction and global happiness evaluations may account for some of the modest relationships with income that these measures show. Also, some of the variance in income comes from recent enough income innovations that the dependence of happiness on news we argue for below could account for some of the remainder of the correlation observed between

income and life satisfaction. (Recall that life satisfaction and global happiness measures are significantly influenced by current affect.)

Given these hints, we predict that future research focusing on current happiness data as opposed to life satisfaction and global happiness evaluations will deepen the Easterlin Paradox and raise estimates of the extent of hedonic adaptation.

C. Hedonic Adaptation vs. Habit Formation. Note that hedonic adaptation is not the same thing as habit formation. Hedonic adaptation is a statement about happiness, as measured by psychologists. Habit formation is a statement about utility, as measured by revealed preference. For example, habit formation often refers to a tendency to do something more if you have done it in the past—an effect of past consumption on marginal utility. Of course, *if* happiness were proportional to flow utility *then* hedonic adaptation and habit formation would be tightly linked. This could be empirically problematic, because data on hedonic adaptation might then imply extremely strong habit formation. For example, suppose utility was of the form made popular by George Constantinides (1990), which can be represented as

$$v_t = E_t \sum_{j=0}^{\infty} \beta^j U(C_{t+j} - \theta H_{t+j}),$$

where v_t is lifetime utility, β is the discount factor for flow utility U, current consumption is denoted C, the "habit" H is a weighted average of past consumption levels, and θ is a parameter between zero and one. Given this form of the utility function, *if* happiness were proportional to flow utility, then evidence of complete hedonic adaptation would only be consistent with $\theta = 1$. For comparison, Joseph Lupton (2002) estimates $\theta \approx .75$ when estimating based on data for life-cycle portfolio choices and a value of θ close to zero when looking at consumption choices. The reason consumption data does not support a high value of θ is that, unless the lags in the habit H are quite long, a high value of θ implies there should be a strong autocorrelation for consumption growth rates that is absent in the data. If happiness were proportional to flow utility, matching the speed of hedonic adaptation would require a fast-moving habit, for which consumption data point to a value of θ near zero. Moreover, even the higher value of $\theta \approx .75$ would not match the observed extent of hedonic adaptation.

A more subtle discussion is required if—to match happiness data—someone suggests a type of habit formation that would not show up in empirical data *other* than happiness data. Suppose that everyone agreed, based on empirical results, that current affect A_t was given by $A_t = f(C_t) - f(C_{t-1})$. To make things even better, suppose that lifetime utility v_t could be represented as

$$v_t = E_t \sum_{j=0}^{\infty} \beta^j [f(C_{t+j}) - f(C_{t+j-1})].$$

One could then claim that affect was equal to flow utility, where flow utility was given by

$$U(C_t, C_{t-1}) = f(C_t) - f(C_{t-1}),$$

with lifetime utility v_t being given by

$$v_t = E_t \sum_{j=0}^{\infty} \beta^j U(C_{t+j}, C_{t+j-1}).$$

But a bit of algebra shows that

$$v_t = -f(C_{t-1}) + (1 - \beta) E_t \sum_{j=0}^{\infty} \beta^j f(C_{t+j}).$$

Since C_{t-1} is already fixed at time t, and the multiplicative factor $(1-\beta)$ does not affect preferences, this utility function represents the same preferences over choices from time t on as the lifetime utility function

$$v_t = E_t \sum_{j=0}^{\infty} \beta^j f(C_{t+j}).$$

There are enough degrees of freedom in this example to force flow utility to be equal to affect. We argue, however, that even in this case it would be clearer and more evocative of the existing economic literature to represent the lifetime utility function in the equivalent, but simpler, more convenient, and more familiar, form $v_t = E_t \sum_{j=0}^{\infty} \beta^j f(C_{t+j})$, where $f(C_t)$ is thought of as the flow utility function $U(C_t)$. The complexity in affect can then be represented in the *relationship* between flow utility and affect. In particular, the stipulated equation $A_t = f(C_t) - f(C_{t-1})$ can then be described by saying that "affect is equal to the first difference of flow utility."[52]

While the two flow utility functions $U_t=f(C_t)$ and $U_t=f(C_t)-f(C_{t-1})$ are equivalent in the preferences they represent over choices at time t and beyond, could the difference between them bear on the hypothetical choice between the consumption bundle now and the consumption bundle fifty years ago we use in Appendix B to argue that utility has been increasing? One answer is to point out that the individual and social choices we really face are those of the next fifty years, not of the past fifty years. Looking toward the future, we have the habits that we have from the past, and must take those as given. From this point of view, the two utility functions are fully equivalent.

Another answer is to carefully pose the hypothetical choice between different comprehensive consumption bundles in a way that takes into account all relevant habit formation. For example, imagine that one were forced to put one's newborn child up for adoption in one of two worlds, where one world has the comprehensive consumption bundle of fifty years ago, while the other world has the comprehensive consumption bundle we have now. Alternatively, assuming that per capita GDP and other objective circumstances improve as much in the next fifty years as they have in the last fifty years, would you rather put your newborn child up for adoption in the world that has the comprehensive consumption bundle we have now or the world that has the

comprehensive consumption bundle of fifty years from now? Because it is hard to imagine the future in detail (even after conditioning on the values of some key statistics, as here) this is a more difficult question, but an important one.[53]

The closely related choice of which society one would wish to be born into is a crucial tool in John Rawls's (1971) extremely influential book of political philosophy *A Theory of Justice* (anticipated by Rawls 1951, 1958). In the Economics literature, choices between societies are also a crucial tool in John Harsanyi's (1953, 1955) theory of social welfare (discussed ably by Pattanaik, 1968.) Of course, these are very difficult choices to make. Nevertheless, revealed preference gives some guidance here, while a simple model of 100% hedonic adaptation would guarantee that happiness data could give *no* guidance for such choices.

E. Local and Global Marginal Thinking vs. Focusing Illusion. Just as the distinction between utility and happiness breaks any tight link between hedonic adaptation and habit formation, the distinction between utility and happiness should make one cautious in using happiness data to assert that people are making systematic optimization mistakes.[54]

David A. Schkade and Daniel Kahneman (1998) consider a thought experiment familiar to readers of David Lodge's (1978) comic novel, *Changing Places*, in which a professor from the gray English industrial city of Birmingham has the opportunity to spend a sabbatical year at Euphoric State on the shores of San Francisco Bay while a California professor takes his place as a visitor at Birmingham. Schkade and Kahneman (1998) study two groups of students, one residing in a gray Midwestern climate and the other in the brilliant sunshine of California. When surveyed, students in both locations have the same distribution of subjective well being. Both Midwestern and California students also predict that either they themselves or a student like them would be more satisfied with specific aspects of California including climate, outdoor activities, social life and cultural opportunities. Schkade and Kahneman explain their results in terms of a *focusing illusion*:

> When a judgment about an entire object or category is made with attention focused on a subset of that category, a focusing illusion is likely to occur, whereby the attended subset is overweighted relative to the unattended subset. In particular, when attention is drawn to the possibility of change in any significant aspect of life, the perceived effect of this change on well-being is likely to be exaggerated. (p. 340.)

While they do not conduct such an experiment, it appears that Schkade and Kahneman believe that a person who actually moved to California would not experience a permanent increase in measured happiness or satisfaction. This would seem logical in light of the equality of overall life satisfaction they observe among Midwestern and California students. Moreover, they cite

---

[53] Note that while there is good reason to hope that utility will be higher in the future, it is not clear that the Easterlin Paradox will continue into the future. It is possible that average long-run happiness will be significantly higher in the future.

[54] In the absence of an adequate theory of the relationship between utility and happiness, it is best to be cautious about asserting that people are making systematic optimization mistakes even when it is clear that people *are* making mistakes in predicting the dynamics of happiness. We return to this issue after presenting our theory of the relationship.

other instances involving paraplegics, lottery winners and widowed spouses in which the positive or negative effect of these events on measured happiness is transient.

What do these results imply? Tacitly assuming that happiness can be set equal to flow utility, Schkade and Kahneman suggest that people mispredict utility for two reasons. First, because of the focusing illusion they overemphasize the importance of a particular aspect of life in California—say, climate—among the determinants of overall satisfaction. Second, and perhaps for the same reason, people fail to predict that their mood will adapt to local circumstances within a relatively short period of time. This appears to be consistent with Schkade and Kahneman's interpretation when they write, "At the individual level, the focusing illusion may lead to unnecessary initiatives. For example, it is not unlikely that some people might actually move to California in the mistaken belief that this will make them happier. (p. 345)

The theory we advance in this paper would predict the same pattern of survey results about happiness, both cross-sectionally and longitudinally, but the interpretation of the results would be different. When the prospect of relocating, say, from Ann Arbor to Berkeley arises, conventional economic theory suggests that an individual needs to consider global utility maximization by comparing the (ordinal) heights of two utility mountains, one corresponding to attainable levels of utility in Ann Arbor and the other to attainable levels of utility in Berkeley. The heights of these mountains depend on location-specific nontraded goods such as climate, topography and culture but also on variations in location-specific traded and partially traded goods that would be available to the person given wages, prices, employment opportunities, family and friends, leisure possibilities and so on. In conventional economic theory, an individual would make a migration decision by comparing the heights of the two utility mountains. Once in Berkeley, the utility mountain in Ann Arbor becomes irrelevant and an individual's decisions are concentrated on finding allocations of income and time to alternative bundles of traded, partially traded and non-traded goods that place her as close as possible to the summit of the local utility mountain.

We suggest that focusing is best understood not as an illusion, but rather as a mental act that plays a familiar role in economic theory. Conventional economic theory suggests that a consumer chooses an allocation that maximizes his utility subject to a budget constraint, a time constraint and other relevant constraints such as distance from family and friends. In finding this optimum, the consumer compares the marginal utility gained from a good with the utility value of its marginal cost in dollars, time, or social interaction. A mental calculation of marginal utility—a partial derivative—requires focusing because it asks how much utility would change holding everything else constant. The empirical evidence of focusing described by Schkade and Kahneman suggests that people are readily able to think about the positive or negative impact of a particular event or state, holding other aspects of life constant.

As is often noted both by economists and non-economists, the optimization task assumed to take place in standard economic theory is daunting in the complexity of its cognitive demands on both information and calculation. To find the local optimum associated with a given utility mountain corresponding to a given location and a given time, an individual may need to consider only variations in a small number of aspects of life because many others are already settled through past decisions, trials and errors. In most day-to-day decisions, focusing on the few dimensions at issue yields a large savings in deliberation costs. The person already has a job, a spouse and

children, a home and, perhaps, the only significant decision at the moment is whether to go to a Chinese or Italian restaurant tonight. By contrast, a large decision involving changing a location, choosing a spouse or changing jobs will cause many aspects of life to change simultaneously. To find the optimum in such cases, the person needs some way to discover the highest utility mountain in a vast range of (high dimensional) mountains, each associated with a particular discrete choice. Just knowing that the next step has higher altitude is not enough.

Even in deciding about an actual move, a fully rational *homo economicus* might conduct a series of thought experiments, similar to those on Schkade and Kahneman's questionnaire concerning satisfaction with aspects of life in the Midwest or California, for each relevant aspect of life. If we assume that utility is additively separable in different aspects of the locations, then the total difference in utility from a move is

$$\Delta U = \Delta U_1 + \ldots + \Delta U_n.$$

Some of these aspects will be essentially the same in both locations, so the individual can focus on just those that are different, together with whatever combinations of aspects interact in a nonseparable way.

Daydreaming in a focused way may be a very helpful way of sorting through particular aspects of a location choice before getting on to the difficult task of making an actual location decision— which entails a summary valuation. After all, Hawaii, New Zealand or the South of France may have even better climates than California. But, among these, perhaps France and California are the best of these in culture and cuisine, on which a given person places a higher marginal utility value. But, after considering the value of these particular aspects of other mountains, it may be that the advantages of the current mountain dominate because it is close to family and friends, its properties are more certain and staying avoids the costs of moving. In long run equilibrium, migration takes place until the expected utility of individuals in the place they reside is at least as high as it is in other places. Moreover, in equilibrium, location-specific advantages such as climate will tend to generate offsetting compensating disadvantages such as high housing prices or low wages (Sherwin Rosen, 1986). It would not be surprising to find that utility is nearly equated in those locations that seem like relevant alternatives.

In this subsection, we have argued that our theory would yield the pattern of survey results reported by Schkade and Kahneman (1998), but that our interpretation of these results would be quite different from their theory of focusing illusion. This raises the question of whether there are any testable differences between the two theories. The most obvious concerns regret. If focusing creates an illusion that leads to the misprediction of utility, we would expect that, on average, people who actually moved to California would experience regret. In our theory, focusing is just an intermediate mental step in forming a summary judgment involving weighting a broad range of relevant issues and aspects of life. While the summary judgment might be erroneous from an ex post point of view, there is no reason to think that the errors are in one direction or the other—California might turn out to be even better than one imagined in a Midwestern college classroom. Similar testable differences between focusing illusion and our theory could be sought from data on regret from other sources such as new car purchases, dating behavior and many other areas of life.

F. Is Choosing Lower Long-Run Happiness Evidence of a Mistake or Evidence that Happiness and Utility are Not the Same Thing?

Equating happiness with utility is a key assumption in what has become an established theoretical consensus among happiness researchers in Economics as well as Psychology. This consensus challenges the validity of the foundations of conventional Welfare Economics which lie in revealed preference theory. In this section, we briefly describe the established consensus in the context of a specific empirical application by Frey and Stutzer (2004b) which examines the relationship between happiness and time spent commuting. Their assumption that happiness and utility are the same thing, in conjunction with the empirical relationship between happiness and commuting, leads them to conclude that individuals systematically mispredict utility. This conclusion, in turn, calls into question the key assumption of revealed preference theory: namely, that the chosen alternative yields higher utility to the consumer than those which are not chosen. Instead, in the spirit of "Subjective Well-Being is Desirable, But Not the Summum Bonum," (Diener and Scollon, 2003), we argue that Frey and Stutzer's (2004b) findings provide evidence that utility and happiness are empirically distinct, but do not bear on the validity of welfare theory based on revealed preference.

In the consensus theory, as summarized by Frey and Stutzer (2003, 2004b), reported subjective well-being is taken as a proxy measure for utility. Maintaining this very strong assumption opens up a wide range of empirical applications and allows for direct tests of conventional theory, as Frey and Stutzer (2004b) illustrate with their analysis of commuting time. Most people find commuting time unpleasant, but endure it as a necessary evil in order to work at a more interesting or better paying job while living in a nicer or cheaper location. In equilibrium, along the lines of Rosen (1986), they argue that individuals should sort themselves among locations such that the disutility of additional commuting is offset by compensating monetary or nonpecuniary benefits associated with a better job or residential location. In such an equilibrium, they argue, total utility should not be related to total commuting time.

Frey and Stutzer (2004b) test this hypothesis in a regression of happiness on commuting time using data from the German Socioeconomic Panel (GSOEP), holding a number of socioeconomic characteristics constant but leaving labor income free to vary. They find a significant negative coefficient for commuting time, contrary to the zero coefficient expected under a Rosen-esque theory. Of course, persons with higher (non-labor) wealth might have higher utility and choose both a better job or house and a shorter commute, thus creating a spurious negative correlation between happiness and commuting time. However, Frey and Stutzer find a significant (although somewhat smaller) coefficient on commuting time in an alternative specification in which permanent differences in wealth or other differences are controlled with the use of individual fixed effects. While their econometric model might be subject to other criticisms, for our purposes we provisionally accept their empirical finding of a negative relationship between happiness and commuting time. We also note that other investigators have suggested similar results for other kinds of decisions. For example, Gruber and Mullainathan (2005) find that cigarette tax increases raised the happiness of potential smokers; Schorr (1991) argues that people mismanage the balance between work and leisure, tending to overwork; and Loewenstein, Ted O'Donoghue and Matthew Rabin (2002) suggest

that misguided purchases of consumer durables such as fancy cars occur because people overestimate the future satisfaction the purchase will bring.

The theoretical explanation advanced by Frey and Stutzer (2004b) and other happiness researchers for such findings is that people systematically mispredict the future utility or, equivalently in this view, the future happiness they will obtain by taking a given action. In particular, Frey and Stutzer (2004b) hypothesize that misprediction is most severe for goods or activities with extrinsic attributes that can be purchased in the market relative to those with intrinsic attributes involving nonmarket social interactions. Thus, while a commuter may choose his home and job with the expectation that the extra money he gets from lower rent or a higher wage will offset the utility loss resulting from spending less time with his family and friends, the negative relationship between happiness and commuting time is interpreted to imply that people systematically overestimate the future relative utility of the things they obtain with the extra income.

Standard economic theory can easily accommodate unsystematic mistakes by consumers, but has a much more difficult time making sense of systematic mistakes. If people get lower net utility from long commutes, why don't they learn this and change the location of their home or job accordingly? Frey and Stutzer (2004b, p. 9) explicitly address this issue. They argue that the formation of expectations about future utility depends on reconstructions of feelings in the past. Failure to learn from mistakes results because "…remembered utility and predicted utility become similar and relatively independent of actually experienced utility." For example, they cite studies in which participants on vacation or holiday trips enjoyed the actual trip less than they had predicted, but report enjoyment levels similar to the ones predicted when they recall the experience afterward.

This is a remarkable argument. Revealed preference suggests that the people who took a vacation gained expected utility. Moreover, in recalling the trip they believe they actually received as much satisfaction as they had expected to get. Presumably, they felt no regret. Their decision appears to involve no mistake according to standard revealed preference arguments and certainly nothing in their recalled experience would cause them to be less likely to take a similar trip in the future. Despite all that, it is alleged that these people actually experienced less utility during the trip than they had expected to receive. This discrepancy is interpreted as a mistake and the failure to notice it after the fact is regarded as the reason that people do not learn from experience and correct their mistakes. Hence, misprediction of utility is common, causing people to make wrong decisions repeatedly which, against their own interest, result in lower levels of experienced utility than could be achieved by alternative decisions. The hypothesis that utility misprediction is relatively greater for actions with extrinsic aspects suggests that materialistic people will be most harmed by these mistakes.

In our view, to make a convincing argument that the individuals were making a mistake, one would need either to find evidence of regret or indications that being presented with the purported evidence of misprediction of utility caused people to want to change their decisions. In the case of commuting, we do not think people would be surprised to be told that commuting is quite unpleasant. Learning evidence that it is difficult to buy much happiness with money or that the effects of additional money on happiness are transient could have more impact on

people's decisions. In the case of trips, forgetting some of the annoyances of travel may, in fact, distort people's decisions; being reminded of these annoyances might affect their decisions to some extent. However, some of the most important benefits of a trip are precisely the memories one brings back. To the extent those memories are positive, the traveler has achieved one of the main objectives of a trip—with the forgetting of annoyances serving as a helpful aspect of the household production function for vacation memories. The incidence of regret and second-thoughts after being presented with relevant data is ultimately an empirical matter for which the quantitative size of effects is just as important as the qualitative direction of effects.

In the absence of evidence of regret or second-thoughts upon being presented with relevant data, the other possibility (which we highlight) is that utility and happiness are not the same thing. Under this alternative, the interpretation of much of the evidence cited by happiness researchers about utility misprediction and systematic mistakes in decisionmaking is simply misleading. While we present our argument using a formal model in the second half of this paper, it is useful to provide some informal intuition now for our contention that evidence from the happiness literature is not inconsistent with a conventional economic model of rational (albeit not omniscient) utility maximizing consumers. It seems quite reasonable, as Diener and Scollon (2003) argue, to assume that maximizing subjective happiness is not the only goal of many consumers because happiness competes with other values or objectives, some of which do not have positive effects on affect. Concretely, much like Becker (1965) or Lancaster (1966), we think of happiness as the outcome of one of a number of household production processes each of which combines inputs of goods, time, and social and physical environment to generate outputs of final commodities according to a household technology. For instance, in the commuting example of Frey and Stutzer (2004b), an individual may endure unpleasant commuting in part because it affords additional money or a more desirable residential location that enables him to buy nice things for his wife and children, to have his children attend a better school or to be able to contribute to a charity to relieve the suffering of others. An empirical question, mostly not addressed in the happiness literature, is whether each of these ways to use money has the same effect on subjective mood or happiness. It seems possible that they do not but, nonetheless, that the individual would be willing to sacrifice his own happiness to benefit others. If so, the negative correlation between commuting distance and happiness observed by Frey and Stutzer (2004b) is quite consistent rational with utility maximizing behavior by persons whose preferences include goals beyond narcissistic fixation on their own pleasure. *Though an altruistic motivation makes the example especially clear, the same logic applies if the objective the individual is pursuing in preference to happiness is a non-altruistic goal.*

G. Summary of the Argument that Utility and Happiness are Empirically Distinct. Here is the underlying structure of the argument that utility and happiness are empirically distinct. First, using standard utility representations, utility has a strong upward trend, while happiness has very little trend. Moreover, happiness is strongly mean-reverting even after permanent changes in circumstances, while utility is not. Second, if one is willing to use nonstandard utility representations (including the flexibility one has in choosing flow utility functions that add up to equivalent lifetime utility functions), one can say the following:

> (a) On one hand, if changes or innovations in lifetime utility were the *only* component of happiness, then maximizing happiness and maximizing lifetime utility would be

essentially equivalent; indeed, happiness could even be viewed as an exotic way of representing lifetime utility *except* that since happiness is focused on changes, it still could not represent preferences over initial levels or initial paths. To put the issue dramatically, though happiness is quite tightly linked to utility in this case, because it is focused on changes, happiness provides no representation of people's views over which society it is best to be born into.

(b) The frequent use of the concepts of utility and happiness to make social welfare statements makes it ill-advised to dismiss the representation of preferences over which society to be born into as unimportant or meaningless. Indeed, this kind of preference is closely related to important conceptions of social justice. These preferences over different comprehensive social situations do not necessarily line up with measured happiness.

(c) Any evidence for persistent, predictable effects of choice variables on happiness imples that changes or innovations in lifetime utility are *not* the only component of happiness.

(d) The fact that at times people knowingly, thoughtfully and without regret make choices that predictably lower their mood, day after day, implies that utility and happiness are empirically distinct.

All of the statements (a—d) remain true regardless of what utility representation one uses for a given set of preferences. Further discussion of these arguments must wait until we have laid out our model.

VIII. Implications of the Theory for Happiness Empirics

Even in the simple case of elation-independent preferences discussed in Section VI, the integrated framework for utility and happiness laid out in Section III has many important implications for empirical work using happiness data. One of the most basic tests of the value of our framework is whether pursuing these implications for empirical work turns out to be fruitful.

A. The Time-Series Properties of Happiness Matter. The most obvious implication of our framework is the need for more research on the time-series properties of happiness. For example,

- How fast and how extensive is hedonic adaptation for affect, as compared to the (so far) better-studied hedonic adaptation for life satisfaction and global happiness measures?
- Do the time series properties of happiness have any implications for econometric practice in research to identify the determinants of happiness?

We have work in progress along both of these lines, but the details must to left to other papers.

B. Price Theory Can be Used to Study Baseline Mood. Second, the theory of baseline mood implies that standard price theoretic tools can be applied to the low-frequency movements of happiness. For example, the dollar value people place on feeling permanently happier can be

gauged by how much they are willing to pay for psychotherapy in time and money (beyond what insurance pays for), divided by the effect of the psychotherapy on happiness. The less effective psychotherapy is at actually raising happiness, the higher the implied valuation on happiness. In the case of antidepressants, in addition to the monetary and time cost, one would have to determine how much people would be willing to pay to have an antidepressant free of side effects and add that value before dividing by the effect on happiness. Such ratios can begin to identify the marginal value of happiness.

Several other price-theoretic issues have been discussed above. Even in the context of our framework, normality of baseline mood still leads to a version of the Easterlin Paradox. It is important to construct measures of the price of happiness over time to see if an upward trend in that price can explain why people are not choosing higher baseline mood in their ever-expanding consumption bundle.

As mentioned above, one limitation in the use of these price theoretic tools is that they depend on knowing people's beliefs about the household production function for baseline mood. Would people do more things that add to happiness if they knew what they were? For example, there are some hints that, in addition to its other benefits, getting more sleep might add significantly to happiness.[55] If this is true, and people knew it, this could place a strong upper bound on the value people place on happiness (the hourly wage divided by the effect of an extra hour's sleep on happiness), but such a conclusion would only be warranted if people really knew exactly how much benefit an extra hour's sleep would have for happiness. On might obtain a more reasonable estimate of the value of happiness by conditioning on people's reports of how much they *believe* an extra hour of sleep each night would add to happiness.

Given measures of the marginal value of happiness, any evidence about the determinants of happiness should be included in cost-benefit analyses. If the marginal dollar value of happiness is high, it could motivate ever more careful empirical work to measure the strength of the effects of variables on long-run happiness. In particular, it could motivate many clever minds to look for good instruments for the possible determinants of long-run happiness.[56]

In principle, the application of price-theoretic tools to baseline mood should yield tests of the theory as well as applications. This kind of test of the theory is likely to emerge over time as the measures of the relevant concepts are refined.

C. The Elation Theory is Readily Testable. Because the theory of elation is the most highly structured aspect of our theory, it is also the most readily testable. In particular, we hope to test whether or not people's hypothetical choices between alternatives A and B always match their predictions of how happy they would be immediately after receiving the news that A had happened or that B had happened with no action on their part. Our theory predicts that people will choose the alternative that would seem like the best news to them (as indicated by their

---

[55] Norbert Schwarz, personal communication, and Kahneman, Krueger, Schkade, Schwarz and Stone (2004). The results are not definitive because of the lack of a good econometric instrument for hours of sleep that is known *a priori* not to affect happiness directly.

[56] Kerwin Charles (2002) is a good example of the kind of attention to exogeneity in happiness research that we mean.

happiness immediately after hearing). To the extent that this does not seem like a very daring prediction, it indicates that the reader has a prior belief in favor of one of the key linchpins of our theory. Nevertheless, it *is* a testable prediction. It is *not* true by definition.

D. Elation Provides Information about Preferences and Expectations. The theory of elation implies that, if it is possible to control for variation in baseline mood, the response of happiness to news will give direct information about preferences. Indeed, the elation component of happiness is essentially an "excess returns" measure for lifetime utility. Therefore, in principle, happiness data can serve as the basis for exactly the same kinds of "event studies" as those carried out using data on excess financial returns.

To implement this insight about the use of happiness data for "event studies," it would be very helpful to have a regular monthly, or even daily, time series on average aggregate affect. This would allow a test of average preferences over aggregate events. In particular, after accounting for the lagged effects of the previous months' news, the theory of elation implies that whether affect goes above baseline or below baseline indicates whether the month's news has been on average good news overall or bad news overall. Since many things happen in a month, each month's data would give information about whether a different innovation vector for the expected consumption bundle represented an increase or a decrease in lifetime utility. Over time, this would tell a great deal about average preferences for aggregate events. Information about preferences for aggregate events is particularly valuable because many of these things do not have regular markets. For example, one might want to know about the relative importance people put on geopolitics compared to economics. Any month in which there is good economic news but bad geopolitical news, or the reverse, would provide relevant information. Election returns are often read as giving this kind of information, but affect data immediately on the heels of news may give more detail. At a minimum, high-frequency happiness data serves as a kind of general-purpose poll question that can give useful insight into how people feel about any big event that comes along.

Kimball, Helen Levy, Fumio Ohtake and Yoshiro Tsutsui (2006) report a pilot study using a few months' worth of happiness data on the University of Michigan Surveys of Consumers. They find a significant dip in measured happiness both in the first week of September 2005, right after Hurricane Katrina, and in the week after the earthquake in Pakistan that occurred in October 2005. Adaptation to the hedonic effects of these national and international news events was close to complete after two weeks. The dip in happiness after Hurricane Katrina was significantly greater in the South Central region of the United States, closest to the hurricane's landfall. The size of the average U.S. reaction to the earthquake in Pakistan is almost as great as the size of the average U.S. reaction to Katrina. Although the human toll from the earthquake in Pakistan was much greater than from Katrina, this still indicates a surprisingly high degree of concern for people on the other side of the world if we are correct in our hypothesis that the size of short-run spikes in happiness indicates the magnitude of the implications of news for lifetime utility. Tsutsui, Kimball and Ohtake (2005) apply a similar event-study methodology to the hedonic reactions in a Japanese sample to the overwhelming electoral victory of Japanese Prime Minister Koizumi in October 2005.

In situations where preferences are clear, the theory of elation draws a strong link between happiness and expectations. This provides another avenue for testing the theory. Section V-A discusses some of the supporting evidence that has already been established on this score. More can be done in this area. One of the most interesting tests would be in areas where people are known to violate rational expectations or where the memory of past expectations is likely to fade. Here the test would be to see if the pattern of people's reported happiness matched the quirks in their expectation formation and memory.

Assuming that the elation theory is valid, it may have relevance for the survey measurement of preferences and expectations quite broadly. On the preference side, since the elation mechanism seems to be fairly automatic as a psychological process, it may be that it is easier and more reliable for respondents to predict their happiness after option A and after option B than it is for them to make a direct choice. On the expectations side, given the unfamiliarity of precise probabilities compared to the familiarity of happiness, it may be easier and more reliable for respondents to report happiness than for them to directly report probabilities. For example, after setting the stage by asking how happy a sample of people on one side of the political divide would feel (A) if their preferred presidential candidate won or (B) if their less preferred presidential candidate won, their average happiness in the days before the election might be an efficient way to assess their subjective probabilities of victory for their preferred candidate.

We have an example from personal experience of using elation to gather information about the strength of preferences. One of us was present when a daughter opened letters from the admissions departments of the colleges she had applied to. The evident strength of the daughter's positive affective reaction was persuasive in establishing the extra value she placed on going to her much more expensive first-choice college, as opposed to her much less expensive second-choice college. Of course, this did not indicate what the ultimate wisdom of each choice would be, but it did indicate her preferences given her beliefs about what it would be like to go to each college.

E. Sufficient Statistics. As we have argued at length, elation and dismay measure the effect of news on expected lifetime utility. Somewhat more formally, elation provides a sufficient statistic that captures the effect on the (expected) lifetime utility of current or future events that had not previously been anticipated. This interpretation of elation is similar to the more familiar idea that in analyzing lifecycle maximization problems the marginal utility of income, a scalar quantity denoted by λ_t, "serves as the sufficient statistic which captures all information from other periods that is necessary to solve the current-period maximization problem." (Blundell and MaCurdy, 1999, p.1594). In particular, λ_t measures the utility value of a dollar saved for expenditure in the future relative to a dollar spent on consumption today and also the marginal utility value of additional time spent on leisure or household production per unit of foregone earnings. Unanticipated changes in the marginal utility of income provide a signal about changes in the optimal allocation of consumption and leisure over the lifecycle.

It is interesting to consider the roles of these two quantities for consumer behavior. The role of the marginal utility of income is well known. Any news about changes in expected future income or price which causes the marginal utility of income to fall serve as a signal to a utility-maximizing consumer to increase current consumption and leisure by reducing saving; news that

leads to an increase in the marginal utility of income is a signal to reduce consumption and leisure and increase saving. Both elation and λ_t are derived from the (expected) lifetime utility function. It follows that news that affects probability beliefs about future incomes, prices, health or any other variables that affect preferences or constraints that causes changes in λ_t will also tend to cause elation or dismay.

Mathematically, there are two differences between elation in our theory and λ_t as sufficient statistics. First, elation is about the *total* lifetime utility, while λ_t is about the *derivative* of lifetime utility with respect to wealth. Second, it is *temporary spikes* in the level of elation that indicate a change lifetime utility, while it is *permanent changes* in the level of λ_t that indicate a change in the first derivative of lifetime utility.[57]

In principle, economists could gain insight into the effects of news by studying the dynamics of consumption and labor supply jointly with longitudinal data on subjective well-being. Observed changes in savings, consumption or labor supply respond to news that influences the marginal value of a dollar while spikes in subjective well-being reflect changes in total lifetime utility. It is easy to show that elation and changes in the marginal utility of income are not necessarily correlated in a simple way. For example, a person who receives news that he has been promoted and will be receiving a higher salary next year will feel elation—a signal that his lifetime utility has gone up—and his marginal utility of income will fall—a signal that he should allocate more of his wealth to current consumption. Today, he might choose to celebrate his promotion with dinner at a fancy restaurant. Consider a less fortunate person who has just heard a jury convict him of a long prison sentence. Obviously, this person suffers dismay. However, given the difference in the availability of fine food in prison and in town, his marginal utility of income also falls and he may also choose a fancy restaurant meal today, assuming that his appeal allows him to stay out of prison for a while. As another example, a person (with full medical insurance) who has had a successful operation on a malignant tumor will experience elation and an increase in his marginal utility of income, signaling that he should increase his saving to accommodate his longer life expectancy.

Although elation and changes in the marginal utility of income could be correlated in either direction, it is likely that most news in the economic domain reflects good or bad news about future income and wealth. That is, macroeconomic news about expansions or recessions or microeconomic news about one's own promotion or layoff tend to produce both unexpected gains or losses in lifetime utility and signals to increase or decrease current consumption. This might be consistent with psychological evidence that has been interpreted as suggesting that an up mood is a signal to move forward with bold plans, putting aside worries, while a down mood is a signal to focus on things that might go wrong and to proceed with caution. It would be interesting to examine whether the correlation between mood and consumption is usually

---

[57] A similarity between elation and λ_t is that econometric identification of both spikes in elation and changes in λ_t requires subtracting out an individual fixed effect. It may also be necessary to control for a few other factors that have predictable effects on changes in subjective well-being or behavior. For example, the real interest rate can have predictable effects on the evolution of consumption and labor supply, even in the absence of news, while time-varying determinants of baseline mood can have predictable effects on overall subjective well-being.

positive and also to test whether this correlation is reversed in the less usual situations such as those described above when mood and the marginal utility of income move together rather than in opposite directions.

IX. Implications of the Theory for Policy

A good way to discuss the policy implications of our framework is to contrast the views that we believe follow from our framework with those of Layard (2005). Layard (2005) is very bold in making policy recommendations based on happiness theory and empirics. Although he is especially bold, we consider the general tenor of his recommendations to be reasonably representative of views expressed in much of the existing happiness literature.

Layard explicitly accepts Kahneman's (1999) proposal to use the expected present discounted value of measured happiness averaged across people as the social welfare function. Besides the issues we discuss in this paper, Layard is assuming a solution to interpersonal comparability issues that we think have not been solved, but let us leave that aside, since all of our social welfare measures share that difficulty in all but the easiest applications.[58]

Many of Layard's recommendations depend only on happiness being more valuable than current public policy recognizes. The general discounting of intangibles in policy discussions makes this likely. Generating and popularizing happiness accounts in parallel to GDP accounts is a reasonable step to rectify insufficient attention to these intangibles. Taking happiness more seriously also suggests many other concrete steps, such as fighting the stigma to antidepressants and psychological treatment, and devoting more resources to mental health care, mental health research, happiness research, and public education about the determinants of long-run happiness.

Other recommendations depend on the externalities inherent in people caring about social rank. Since both revealed preference and happiness data indicate that social rank is important, these recommendations remain on the mark. Quantitatively, a revealed preference measure of the importance of social rank may be different from an affective measure, but qualitatively, the implications of social rank mattering are the same either way.

Affective data also provide a good reminder of the importance of many other externalities and public goods—an importance that can be verified by revealed preference. The sense of community matters, the strength and quality of marriages and families matter: and the

---

[58] The Ordinalist Revolution made it clear that the key philosophical issues in judging social welfare for purposes of public policy could not be avoided even if a perfect direct measure of individual welfare existed. Most notably, there is no easy escape from the difficulties surrounding interpersonal comparison. For example, should those with more refined tastes who can distinguish more minute differences in quality therefore be accorded greater weight in social choice? See Stigler (1950).

responsiveness of government matters. Again, a revealed preference measure may differ quantitatively in the importance it suggests for these externalities, but it is likely to agree qualitatively.

There are two areas where we differ with Layard. First, Layard makes many recommendations based on Kahneman's (1999) social welfare measure, where we would turn to utility based on revealed preference as the appropriate welfare measure. This leads to stark differences in recommendations about tradeoffs between economic growth and other values. For example, Layard argues that since economic growth does very little to raise happiness, while being forced to move from one city to another lowers happiness significantly, it is worth sacrificing a great deal of economic growth in order to slightly reduce the need for mobility. To us, this either tacitly assumes that feeling happy is the only thing people care about (which we dispute), or it requires forcing upon them the objective of maximizing happiness when, given the choice, they reject this objective for themselves. There are many indications that economic growth is, in fact, important to people, even if it does not raise happiness. In principle, the dollar value of happiness could be high enough to make such a sacrifice worthwhile even if happiness is not the only thing in the utility function, but it would require an extremely high value. Even given existing lay knowledge about the determinants of baseline mood, if the value people place on happiness were high enough to make this kind of sacrifice worthwhile, we should see many more people seeking psychological treatment, sleeping more, exercising more, eating better, pursuing meditative practices, and so on, than we observe. Of course, if stability of residence enters the utility function beyond its effect on happiness, its valuation could be higher, but this is not Layard's argument.[59]

It is worth being very explicit about why Kahneman's (1999) social welfare measure differs so much in its implications from standard social welfare measures in Economics. Our theory implies that the present discounted value of affect Kahneman points to is the sum of two very different components: the present discounted value of baseline mood and the present discounted value of elation. Conceptually, we view the present discounted value of baseline mood as something like the present discounted value of any other sub-utility function, such as the present discounted value of a health measure. While likely to be correlated with lifetime utility, this present discounted value of baseline mood represents only one of the things people care about. (Less importantly, this present discounted value also has in it no representation of any decreasing returns to baseline mood in the utility function.) By contrast, the present discounted value of elation is a very interesting quantity that (at least approximately) represents the cumulative innovation in lifetime utility over the interval of time covered in the discounted sum. In other words, the present discounted value of elation answers the question of how well one's life has turned out compared to what one expected at an earlier time, perhaps long ago, when the discounted sum began. If for example, one could separate out elation from baseline mood in

---

[59] Note that for the set of things that only enter utility through happiness, valuing them at their effect on happiness times the revealed preference dollar value of happiness is a very different procedure from the common valuation procedure of dividing the effect of a variable on happiness by the effect of *income* on happiness. This procedure is wrong because it assumes that income only affects utility by affecting happiness—something we know to be false for income to the extent that people thoughtfully sacrifice happiness for higher income. In order to use income as a *numeraire*, all the benefits of income on utility need to be accounted for, not just the (possibly small) fraction of the benefits of income that show up in a higher level of happiness.

measured affect and constructed present discounted sums of elation from a long panel of the adult population over 18, the average discounted sum of elation would represent how much better or worse people's lives turned out than they expected at age 18.   As a social welfare measure, this intriguing quantity has one serious problem: it does not credit as social improvement any improvement in how people's lives in a society look as of age 18. Even if the panel were extended back to five-year olds, the present discounted sum of elation from that age on would not give due weight to improvement in life prospects as life prospects appear as of age 5. In our view, many of the most valuable aspects of progress over the past few centuries, or even the past few decades, are ones that would be highly valued by five-year olds, not just in the moment, but as they think about what their lives will be like when they grow up.

Second, we consider Layard too quick to believe that people are making systematic mistakes in optimization.[60] People no doubt do make mistakes, but because happiness is not the only thing people care about, happiness data alone is seldom enough to identify optimization mistakes. The key types of evidence we would point to for identifying mistakes are regret and people changing their minds on a decision after thinking more carefully or getting better information.   Also, not all factual mistakes lead to optimization mistakes. In Section VI-D, we showed that mistakes in predicting the dynamics of affect do not necessarily lead to optimization mistakes.

X. Conclusion

Happiness research matters because—even if economic progress continues unabated over the next 50 years in the U.S. and other advanced countries—whether the citizens of these countries end up rich and happy or rich and unhappy depends on whether money *can* buy happiness *and* on whether the additional economic resources will, in fact, be used to obtain additional happiness. To the extent there is a tradeoff between happiness and other values, the increases in income and wealth that accompany economic progress are likely to make improvements in subjective well-being increasingly important for welfare compared to further improvements in other areas.[61]

In order for happiness research to fully tap into the vast accumulated human capital of the Economics profession, we consider it important to develop a theory that respects the canons of Economics as well as the findings of Psychology.   One of the most important canons of Economics is Ordinalism, or the principle of Revealed Preference.

Revealed Preference, applied to the Easterlin Paradox of nontrending happiness in the face of dramatic improvements in per capita income and many other areas of life, clearly distinguishes utility from happiness. Utility is the extent to which people achieve what they care about, as indicated by their choices; happiness is how they feel. This distinction is important. In

---

[60] In this, Layard follows Gruber and Mullainathan (2005) and Frey and Stutzer (2004b).

[61] One area where trends in happiness could have important macroeconomic effects is in the area of happiness on the job. For example, it is possible that, in the coming decades, advances in subjective well-being at work could alter people's relationship to work in a way that significantly raises the average retirement age. Happiness on the job is likely to be an increasingly important element of competitive advantage—particularly for firms that need to attract skilled workers who may place a higher dollar value on happiness.

particular, the distinction between utility and happiness leads to many insights and productive questions that would be difficult to see as long as utility and happiness are confused with one another.

Appendix A: Measuring Happiness

Psychologists have taken measurement issues in assessing emotions in general, and happiness in particular, very seriously. Randy Larsen and Barbara Fredrickson (1999) give a survey of research touching on this issue. Self-report measures of happiness and sadness (the most common type of measure) have been related to impressionistic observer ratings of happiness, highly-structured coding of facial expressions by trained observers, assessment of voice tone, electromyographic measurement of face muscle activation, measurement of skin conductance, heart-rate, blood pressure and respiration, electro-enchephalograms, positron emission tomography and functional magnetic resonance imaging of brain activity (where "…approach related positive emotions are associated with left anterior activation whereas withdrawal related negative emotions are associated with right anterior activation" Larsen and Fredrickson (1999), p. 53.) Self-report measures of happiness and sadness have also been shown to predict many types of cognition and behavior in the laboratory, including writing speed and performance speed on other tasks, judgments of probabilities, the output of free word association and word completion under time pressure, speed of judging positive and negative words versus nonwords, and the speed of the startle reflex after a loud sound. All of these experiments add up to consistent evidence that happiness is a measurable psychological state.

In general, self-report measures of emotions can be affected by social desirability and the semantic framing effects that arise cross-culturally, and lack of conscious awareness of emotions. For the most part, social desirability and semantic framing effects should be fairly constant over time within a given culture and can be dealt with empirically using fixed effects. The likelihood that people might lack conscious awareness of emotions is a subject of debate within Psychology. Some psychologists insist on conscious awareness as part of the definition of an emotion. (Larsen and Fredrickson, 1999 reports that "some would question whether an unperceived emotion is an emotion at all.") But even Tim Wilson (2002), in a book-length argument for the possibility of unconscious feelings, points out that "feelings differ from the rest of the adaptive unconscious in their potential to reach awareness" and allows that "It might even be the case that the default is for feelings to emerge into awareness, and that it takes special circumstances to prevent them from doing so." (See Wilson, 2002, p. 134.) It seems likely that the overall positive or negative aspect of feelings that we are focusing on under the label of "happiness" makes it into consciousness more reliably than the detailed reasons behind feelings or finer categorizations of emotions. Wilson (2002) goes on to discuss repression, inattention and "the obscuring of feelings by the smoke screen of people's conscious theories and confabulations." Repression and inattention seem unlikely to cause serious problems for the survey measurement of happiness. However, "the smoke screen of people's conscious theories" about happiness is a serious issue, which we address below.

Among self-report measures of happiness, the gold standard is experience sampling, in which people are signaled at random intervals to report their current happiness. Kahneman, Alan Krueger, David Schkade, Schwarz, and Arthur Stone (2004) argue that the day reconstruction method is a close second. Measuring happiness as part of a large-scale survey presents an extra issue in that the survey itself may represent a significant slice of a day. To avoid too much emphasis on the feeling states engendered by the interview process itself one can ask about

happiness over a longer, but still relatively short span of time. (Michael Robinson and Gerald Clore (2002, p. 950) looked at evidence on happiness reports with different time frames. Their evidence led them to conclude that a few week's time is the longest interval for which one can get happiness reports that are not contaminated in an important way by people's theories of how they "should" feel.) The Health and Retirement Study measures affect by the following series of questions:

> "Now think about the past week and the feelings you have experienced. Please tell me if each of the following was true for you much of the time this past week: [62]
>> a. Much of the time during the past week, you felt you were happy. (Would you say yes or no?)
>> b. (Much of the time during the past week,) you felt sad. (Would you say yes or no?)
>> c. (Much of the time during the past week,) you enjoyed life. (Would you say yes or no?)
>> d. (Much of the time during the past week,) you felt depressed. (Would you say yes or no?)"

Operationally, one can treat happiness as the latent variable behind these four yes/no questions. This series of questions on the Health and Retirement Study is a subset of the Center for Epidemiologic Studies Depression (CES-D) measure of depressive symptoms.[63] These questions illustrate what we mean when we say that the concept of happiness we are referring to is about current feelings. These questions ask about easily accessible feelings and memories of feelings in the past week. One indication of how readily respondents answer these questions is that the average amount of survey time required for all four questions put together is less than 35 seconds.

It is important to contrast current affect measures like those on the HRS with life satisfaction measures, such as those on the German Socioeconomic Panel—"On a scale from 1 to 10, how satisfied are you with your life?"—and "global happiness" questions, such as those on the World Values Survey:

> "Taking all things together, would you say you are
>> 1. Very happy
>> 2. Quite happy
>> 3. Not very happy
>> 4. Not at all happy
>> 9. Don't Know [DO NOT READ OUT]"

An extensive body of psychological research explores the cognitive processes underlying global judgments of happiness and life-satisfaction (for a review and process model see Schwarz and Strack, 1999). It converges on the following conclusions[64]:

---

[62] In the first wave respondents were instead asked "*Please tell me how often you have experienced the following feelings during the past week: all or almost all of the time, most of the time, some of the time or none or almost none of the time.*"

[63] See Steffick (2000) for a detailed description and assessment of the CES-D questions in the HRS. Besides omitting the other less relevant questions, we have reversed the order of the first two questions even after those omissions in order to give the version of the question that we would recommend for use on other surveys that do not have a more extensive CES-D battery of questions.

[64] We are particularly grateful to Norbert Schwarz for this summary of the psychological research on different subjective well-being measures.

1. Reported life-satisfaction does not reflect stable inner states of respondents. Instead, these judgments are formed on the spot and depend on which aspects of life happen to come to mind at the time of judgment, which gives rise to pronounced context effects. For example, when students are asked to report their overall life-satisfaction and their dating frequency, both correlate r = .1 when the life-satisfaction question is answered first, but r = .7 when the dating frequency question precedes the life-satisfaction question, thus bringing the domain of dating to mind (Strack, Martin, and Schwarz, 1988).
2. The use of comparison standards is similarly context dependent. People can evaluate their current lives relative to their expectations, their past situation, the situation of others, and so on, resulting in profoundly different judgments. For example, the mere presence of a handicapped other in the room is sufficient to increase global life-satisfaction (Strack, Schwarz, Chassein, Kern, and Wagner, 1990) and one's current life looks good or bad depending on which aspect of one's past was brought to mind (Strack, Schwarz, and Gschneidinger, 1985).
3. People can simplify the complex task of evaluating their life-as-a-whole by drawing on their current feelings as an indicator of their overall well-being. For example, survey respondents report higher life-satisfaction when called on sunny rather than rainy days—unless a preceding question about the weather makes them aware that their current mood may not provide diagnostic information about the overall conditions of their lives (Schwarz and Clore, 1983).[65]

This context-dependence of evaluative measures of well-being attenuates any meaningful relationship with objective circumstances of life and motivates approaches to the measurement of well-being based on people's momentary affective experience.

In comparison to global evaluations of one's life-as-a-whole, assessments of current affect pose more reasonable cognitive demands. As noted in point 3 above, experimental evidence suggests that survey responses to questions about overall life satisfaction or about global happiness with life rely heavily on the readily accessible internal information a respondent has about current affect (Schwarz and Clore, 1983).[66] Thus, how a respondent feels *right now* has a strong effect on answers to overall life-satisfaction and global happiness questions, whether we like it or not. We maintain that it is clearer to focus on current happiness directly, so that we know what we are getting, in a transparent way. Finally, to the extent that respondents are not using current affect as a shortcut to make an overall evaluation of life satisfaction or global happiness, there is a serious danger that they will report how happy or satisfied they think they *should* feel about their lives according to whatever folk theories they have about happiness and satisfaction.[67]

Note that one mechanism for manipulating answers to life-satisfaction and global happiness questions by questionnaire context or other psychological intervention is precisely through manipulating current affect. This can explain the sunny-day effect and the effect of finding a coin on global evaluations. In these experiments, there is no reason to doubt that there is a genuine effect on affect. Another mechanism is through manipulating people's sense of how happy or satisfied they *should* feel, as when people's attention is drawn to dating frequency or when they are encouraged to compare their situation to that of a handicapped other. (The

---

[65] The relationship of such context-dependence to decision-making is an important research question. For example, Hirshleifer and Shumway (2003) indicates that sunny days have a detectable effect on stock-market trading.
[66] See also Schwarz (1996, 1999) and Schwarz and Bohner (2001).
[67] For example, consider the fact reported by Lucas, Clark, Georgellis and Diener (2004), that life satisfaction is permanently dragged down by an episode of unemployment. Even if the affective sting of past unemployment has long since faded away, asking for an overall evaluation of life satisfaction invites the respondent to evaluate the past as well as the present. It is not surprising that a past episode of unemployment permanently affects one's assessment of one's autobiography.

handicapped other may also alter current affect.) Since anyone can feel bad or good at a particular moment, regardless of their overall life situation, considerations of how one *should* feel are likely to have less effect on reports of how one feels right now.

One important application of these psychological principles about overall life-satisfaction measures is to interpreting the fact reported by Lucas, Clark, Georgellis and Diener (2004), that life satisfaction is permanently dragged down by an episode of unemployment. Even if the affective sting of past unemployment has long since faded away, asking for an overall evaluation of life satisfaction invites the respondent to evaluate the past as well as the present.   It is not surprising that a past episode of unemployment permanently affects one's assessment of one's autobiography.   Their evidence does not yield a strong prediction about the persistence of the effect of a spell of unemployment on a measure of current affect.

Appendix B: Measuring Utility

For economists, a discussion of measuring utility is only a reminder. Utility is defined by revealed preference—the information gleaned from the choices people make.   Some of the accumulated wisdom from economic research is encoded in standard functional forms that are repeatedly applied and tested. The techniques of revealed preference can be applied to tradeoffs over seemingly incommensurable values, and apply even to situations involving choices over time.

For non-economists, one can say that the concept of utility relies on an individual's judgment of his or her priorities, as reflected especially in his or her actual choices when faced with a tradeoff, or, at a minimum, his or her choices in a hypothetical situation. Higher utility out of any two choices is defined by what the individual chooses (or would choose) when presented with those two choices. Thus, utility is a measure of the extent to which people get what they want, and differences in utility are predictors of behavior. This allows a deep connection between positive (descriptive) and normative (prescriptive) aspects of utility theory in Economics.

1. The Upward Trend in Utility. In view of the Easterlin (1974, 1995, 2003) Paradox of secularly nontrending happiness, an important application of the principle of revealed preference is to a hypothetical choice between the comprehensive consumption bundle (including all externalities, public goods and time use patterns) now and the comprehensive consumption bundle fifty years ago in the U.S. Real per capita income has grown dramatically over that period of time, which means that the total set of marketed consumption bundles that people can choose from has expanded.   Higher real per capita income allows people more choices, out of which they typically elect to spend in ways they could not previously afford--rejecting the available option of continuing to spend in the same way they did at the lower income level. Average work hours have trended slightly downward. Moreover, as Easterbrook (2003) points out, a large set of goods not traded or only partially traded in the market have either stayed about the same or improved over the last fifty years.   Among partially-traded goods, medical care and longevity have been improving dramatically, while household conveniences have reduced the time necessary for housework and increased the time available for genuine leisure even for many who do spend longer hours in market work. Equality between the sexes and races, while far

from complete, is much better than two generations ago; the number of democratic nations is on the rise; and even the War on Terror, which at worst could involve the nuclear destruction of a large portion of Manhattan, is an improvement over the Cold War, which at worst could have destroyed human life from the face of the earth. Finally, many of the non-traded goods that worsened for a while after 1955, such as rates of crime, teenage pregnancy and drug abuse, have turned the corner and begun trending in a favorable direction for the last two decades. In short, although many problems remain, and are the focus of nightly news reports, we argue that it would be a bad deal to trade the problems we face today for the problems of yesteryear, implying that utility is higher than fifty years ago.

Of course there are those who look back at the past with nostalgia. In part, the increasing individual freedom that comes with higher per capita income may have some undesirable side effects such as a diminished sense of community, of the sort Robert Putnam (2000) describes in his book *Bowling Alone: The Collapse and Revival of American Community*. (Trends in the divorce rate and other aspects of family structure can be seen as part of the same phenomenon.) Similarly, the rise in per capita income may have increased the availability of illegal drugs and access to a wide variety of delicious foods and drinks that create important intra-psychic conflicts. Still, how many would really want to go back to the way it used to be if they saw clearly the way it *really* used to be? It is easy to forget the legitimate and the irrational fears engendered by the Cold War, the toll of racial and other injustices on those mistreated, the enforced conformity that went along with the greater sense of community in the past, and how effective the long-available drug of alcohol is at messing up the lives of those who are prone to intra-psychic conflict. And it is easy to take for granted boons such as word processing, the ability to watch any of a huge range of movies at home, the existence of J. K. Rowling's *Harry Potter* series, and inexpensive access by means of free internet access at the public library to a huge range of fascinating scientific findings that were not known fifty years ago, let alone available at the click of a mouse.

Besides ordinary selective memory that often leads us to forget former difficulties once they have been surmounted, there is another kind of bias that helps to fuel such nostalgia: despite recent trends in historiography toward telling the stories of those at the bottom of the social ladder[68] as well as the stories of common men and women, our image of the past is still often dominated by the biographies of those near the top of the social ladder who were much better off than the average person in their time. Even when we assess the past by thinking of the experience of our own grandparents, they are far from a random sample of people in their time. The ancestors of a randomly chosen individual in the present are likely to be people who were more successful than average in number of descendants--and likely to be above average in the degree of success in their life experience more generally.[69] The experience of those who died young in the past or who never found mates is not remembered as well as the experience of those

---

[68] Our choice of the "social ladder" metaphor is influenced by the form of a question on the HRS leave-behind survey asking people to mark their perceived social rank on a printed ladder, pioneered by Michael Marmot's Whitehall II studies of British civil servants. (See Marmot, 2004.)

[69] In the future, this bias could go the other way, since the ratio of family size for high-status parents to family size for low-status parents seems to be falling over time.

who did. The travails of those in the present who die young or who never find mates is more apparent.[70]

One reason it is important to hold imagined social rank constant when assessing the past versus the present is that most people care a lot about social rank.[71] From the standpoint of revealed preference, it is not difficult to observe people making choices that sacrifice other valuable things in order to attain higher social rank. As a consequence, it is not incredible that someone might choose to be a king or queen in a bygone era rather than a middle-class person today, even if the real value of the market consumption bundle of the middle-class person today is worth much more. Individuals' positional concerns are not irrational. In many domains of life, relative standing is more crucial to obtaining desired outcomes than absolute standing (for discussions see Frank, 1985, 1999; Hirsch, 1976; Sen, 1983). Social rank yields hard-to-measure but real benefits in terms of respect and favorable treatment by others. Because, in practice, social rank is so highly correlated with income, at least in the United States, secular comparisons are useful for distinguishing concern with income from concern with social rank.[72]

The point of this extended discussion is to argue that average lifetime utility for people of a given age is higher than it was fifty years ago, even after accounting for a wide range of tradeoffs going far beyond those that are in obvious monetary terms.   Of course, the choice between the comprehensive consumption bundle of fifty years ago and the comprehensive consumption bundle now is a hypothetical choice. But every year millions of people make a choice that is similar in important respects by migrating from a poor home country to the U.S. and other rich countries. Many leave behind tight-knit communities in which they have high local social status for a foreign land where they will be at the bottom of the social status ladder and where they cannot even speak the language. Clearly, in making the enormous effort of migrating, with all the psychic costs of being uprooted from one's familiar cultural surroundings, they are choosing something that they value highly—the modern consumption bundle in the U.S. and other rich countries.[73]

2. Mistaken Choices. The greatest difficulties in measuring utility arise when people make mistaken choices or have inconsistent preferences.

Garden-variety mistakes based on a lack of knowledge of objective facts are the easiest to deal with. Consider the case of someone who chooses a particular car, thinking that it will get good

---

[70] Note that, in comparing the past to the present, it is important to abstract from people's preference for the familiar and status-quo bias more generally, which would have worked in favor of the actual experience in the past as much as in the present. A good way to abstract from the attraction of familiar idiosyncratic details of one's life is to imagine a choice between being (a) thrown into the life of a randomly chosen individual in the present and (b) being thrown into the life of a randomly chosen individual in the past. Unfortunately, there is no such helpful device to help in abstracting from the familiarity of one's entire era.

[71] For some direct evidence on the strength of preferences over social rank based on hypothetical choices, see Solnick and Hemenway (1998). For a discussion of social comparison by psychologists, see Sulls and Wills (1991).

[72] There are some instructive instances of social rank diverging from income rank even in the present-day U.S. Clergy and teachers (including professors) often have considerably higher social rank than their income rank. This relatively high social rank is important in making many people willing to sacrifice a significant amount of income to go into these fields.

[73] Note that, at $9600, the per capita GDP of Mexico—an important source of migration to the U.S.—is not far below the U.S. real per capita GDP in 1955.

gas mileage, but then discovers that it gets bad gas mileage. Once she learns this, the purchaser regrets the earlier decision and wishes that she had chosen a different car. In this case, in judging utility one needs to either use the choice the purchaser *would have made* if fully informed, *or* the choice the purchaser did make between ideas of cars with assumed characteristics.

Sometimes a mistake arises not from lack of basic facts, but from a failure of computation. For example, one of the authors has only very recently begun to adjust his book-buying habits for the high shadow-cost of available book-reading time that is generated by the large number of books already in his personal library, by thinking "What are the chances that on any future date this book will win out over all of the competition?" This is the kind of reasoning that does affect one's choices when the calculation has finally been made. Taking Paretian Welfare Economics as our touchstone, we consider it relatively uncontroversial to suggest that utility (strictly speaking, preferences) be measured according to what people would choose when not only well-informed in terms of raw information, but also when they are aware of relevant calculations and lines of reasoning.[74]

A third type of mistake is making mistakes about what one's subjective experience will be after a given choice. There is nothing disruptive of standard economic theory about the existence of experience goods, such as a new flavor of ice cream, for which preferences are known only after trying some of the good.   Marketing strategies by firms selling experience goods vary from actively providing a free taste to forcing people to buy a substantial package based on guesswork. When free samples are not provided, it is easy to make mistakes due to not fully knowing one's own tastes, even if the physical properties of the product itself are well-known.

For an expensive durable good, an optimal decision of whether to purchase the good should involve considering the time-path of one's subjective experiences with the good. It is not hard to imagine someone changing her mind about buying an expensive car upon being shown evidence that after a year people report roughly the same experience when driving an expensive car as when driving a much cheaper car—say because driving is one of those activities that becomes reasonably automatic and so fades into the background of awareness—pushed out by thoughts of where one is going to and where one has just been.   Indeed, people might not even need formal evidence for things like this; their own past experience of being less excited by a new good after the first few months could inform their later decisions. However, a considerable body of evidence reviewed in George Loewenstein and David Schkade (1999) indicates that people make serious mistakes in predicting future happiness. Though it is larger in size and scope, one can view a mistake in predicting future happiness as akin to a mistake in predicting whether one will like a particular type of ice cream. It is not clear whether people are making the right decisions or not until they are well-informed about the modification of the time-path of happiness that will actually result from a purchase.

---

[74] There is a practical problem of distinguishing between the force of a calculation or line of reasoning itself and the desire to agree with the person urging that line of reasoning. In principle there are ways to deal with that problem. For example, in presenting a hypothetical choice, it is important to even-handedly present correct calculations and lines of reasoning that favor both the pro and con side of a decision. Also, to minimize social pressure, it may be possible to present calculations and lines of reasoning by a prepared text or an interactive computer setup. Making sure the agent is able to make the decision with as much anonymity as possible may also be helpful.

Some psychologists have gone further, to maintain that the fact that reported happiness with a new car is often high in the first few months after purchase and much lower thereafter—in a way that people are bad at predicting—*necessarily* means that someone has made a mistake in purchasing it. To our way of thinking, this is going too far. The key question is whether a correct knowledge of the modification of the time-path of happiness that will actually be induced by a good would make a material difference to a decision. The issue becomes clearest if we consider the choice of a purchaser who is fully conversant with the Hedonic Psychology literature and carefully observant of the pattern of his or her own affective reactions. As long as the purchaser is aware of and thoughtfully considers the fact that happiness with a new durable is likely to fade after a time (and absent the kind of inconsistency discussed below), it seems appropriate to defer to that well-informed, thoughtful decision in judging utility, regardless of the time-path of subjective experience with the new car. Indeed, there is every reason to think that people care about many attributes of a car other than its price and the subjective experience they will have with it—such as its ability to get reliably to work and back. Even the set of all indirect effects of a purchase on happiness (for example, including the reduced likelihood of sorrow from being scolded for getting to work late) should not necessarily be dominant in the decision of whether to make a purchase. The concept of utility (or equally in this context, preferences) involves deferring to each individual's own view of how much to factor in the modification of the time-path of happiness that would result from a purchase when making the decision of whether to buy or not to buy. Since the consequences for happiness are only one aspect of a good, it would not necessarily be irrational to give those consequences only a small weight even after understanding them fully.

A fourth type of mistake is described by Barry Schwartz (2004) in his well-publicized recent book *The Paradox of Choice.* He emphasizes the mistake of trying "too hard" to optimize. The "maximizers" he identifies by an abbreviated personality test seem to optimize without regard to the costs of the time and effort devoted to deliberation about a choice. Of course, this is not true optimization in a larger sense; a fully optimal choice must take into account deliberation costs. However, this raises an important issue for the measurement of utility. We argue that the utility function for everything other than deliberation costs should be measured by the choices that would be made by exactly such an agent who disregards deliberation costs.[75] Where people's preferences are similar to one another, this concept of utility approximates the utility that can be achieved when a small jury selected from a large group of people with similar preferences pays the deliberation costs for everyone in the whole group. For example, *Consumer Reports* is a practical effort to help people approximate this level of utility.

The general point that arises from thinking about mistakes is that, even when underlying preferences are fully consistent, choices arise from the interaction of preferences, ordinary constraints *and information structures.* The relevant information structures can include both external information constraints and internal cognitive constraints. In principle, internal cognitive constraints are no harder for economic theory to deal with than external informational

---

[75] Calculation and deliberation costs should be recognized just as much as any other costs an agent faces. The difficulties in modeling "bounded rationality" problems due to the "infinite regress" problem discussed by John Conlisk (1996) among others should not blind us to the obvious fact of deliberation costs. Because we do not view the recognition of deliberation costs as a departure from "rationality" at all, we favor the more neutral term "bounded cognition" for what has traditionally been called "bounded rationality."

constraints. For example, Woodford's (2002) model of monetary nonneutrality (based on the model of rational inattention in Christopher Sims, 2002) has been criticized for relying on extremely low bit transmissions rates. One way to defend Woodford's model is to locate the low bit transmission rates inside the small portion of the typical decision-maker's psyche devoted to thinking about macroeconomics.

Rayo and Becker (2005) provide an interesting example of explicitly modeling decision-making as the outcome of an underlying utility function ("evolutionary efficiency") filtered through internal informational constraints. Assuming a limit on the total number of gradations into which values of the underlying utility function can be distinguished, they show that an optimal deployment of that limited available total precision is to make fine distinctions in the neighborhood of values where an agent will actually be operating but only gross distinctions at outlying values.[76]

Could consideration of mistakes overturn the conclusion of Appendix B, part 1, that utility is rising? It is instructive to consider again the choice hundreds of thousands of people make every year to migrate from a poor home country to a foreign rich country. Certainly some regret their decision *ex post* and wish they had never migrated, and some even pay the fixed cost of returning to their home country to a situation no richer than before. But among those who are able in the end to bring their families as well, there is no evidence of widespread regret at migrating to a rich country. There is even less regret among the grandchildren of those who migrated, who escape most of the large fixed costs of migrating.

3. Inconsistent Preferences. Inconsistent preferences are more difficult to deal with than mistakes. Just as mistakes involving one's own preferences can be modeled as an underlying utility function together with internal information acquisition, transmission and processing constraints, inconsistent preferences are now routinely modeled as an intra-psychic game between multiple agents within the same person, each having a distinct set of preferences. However, if there is more than one set of preferences operating within a single individual, normative analysis faces a version of the Social Choice problem even for evaluating individual welfare. Unlike the standard Social Choice problem, there is no reason for a presumption of equal ethical value for all the different intra-psychic agents.

In order to do normative welfare analysis even at the individual level, one must take some stand on this intra-psychic social choice problem. Our proposal in this regard would be to rely, in all ordinary cases, on the psyche's own dispute resolution system. As Camerer, Loewenstein and Prelec (2005) discuss, the brain's controlled cognitive system is often brought into action to use deliberation to resolve disputes between other systems in the brain. Thus, we consider well-

---

[76] In their paper, Rayo and Becker (2005) call the filtered version of underlying utility "happiness." However it is unclear in what way it would relate to happiness as we are using the term. In particular, like visual processing (which they use as an analogy), the filtered version of utility they discuss might operate at a very early unconscious or "automatic" stage in the sense of Colin Camerer, George Loewenstein and Drazen Prelec (2005) and so could be several cerebral processing steps prior to "happiness" in the experiential sense.

informed, thoughtful, revealed preference to be the best practical gold standard for an individual's preferences for the purposes of welfare analysis at the individual level.[77]

When multiple preferences coexist within the same individual, the whole-person utility function that is the solution to the intra-psychic social choice problem may not be the utility function that has the tightest relationship with happiness. Here, a key issue is which "selves" give happiness reports to an interviewer. A relatively straightforward case is when the problem is short-sightedness in the sense of hyperbolic discounting of the sort described by David Laibson (1997), where there is one "self" in command at each point in time, in a known sequence. Other cases could be more complex. Although we consider it a high priority for the future, modeling the relationship between reported happiness and either whole-person utility or the set of utility functions within the psyche is beyond the scope of this paper. From here on, in discussing the relationship between happiness and utility, we will assume that the individual has only one set of preferences, which are internally consistent.

4. Is Happiness in the Utility Function? The principle of revealed preference indicates that happiness is in the utility function. Hundreds of thousands of people spend thousands of dollars each on therapy that is not reimbursed by insurance in hopes of becoming happier or at least less unhappy. Millions of people endure the significant negative side effects of chemical antidepressants in order to feel happier. Self-help books and magazines featuring cover articles on happiness sell briskly. Moreover, many products that may not actually make one's psychological state significantly more positive are advertised as if they will, as described in great detail by Melinda Davis (2002). Advertising aimed at suggesting that a product will improve one's brain state would not be so prevalent if a desire for positive affect were not an element of preferences. In sum, many people want to feel happier and are willing to sacrifice other things in order to attain that psychological state.

On the other hand, it seems clear that happiness is neither the only thing in the utility function nor a sufficient statistic for all of the goods that are in the utility function. To make this clear it is sometimes useful to use the technical term "affect" for current happiness as a reminder that we are talking about feelings. People care about things other than how they feel. Most obviously, they sometimes sacrifice current affect for a later benefit; for example (to take one of people's lowest affect activities according to Kahneman, Krueger, Schkade, Schwarz and Stone, 2004) it makes sense to spend time on household chores despite the low momentary affect associated with that activity because of the later benefits of having a clean house. There might be a later increase in affect as a result, but it is not clear that the later benefits show up in affect and utility in the same way.[78] Second, people often sacrifice their own affect to benefit their children, as when one spends long hours at a grueling job to finance college educations for one's children. In

---

[77] Insisting on transitivity is one aspect of "thoughtfulness" here. Thus, in principle, in assessing preferences, we would rely on an individual's deliberative choices for an entire menu of decisions at once, with an iterative process where the individual is forced to resolve non-transitivities.

[78] This is a very interesting empirical question. In testing whether intertemporal tradeoffs in utility match intertemporal tradeoffs in happiness alone, one must address the problem that people are not good at predicting their future happiness, as pointed out by Loewenstein and Schkade (1999). It may be possible to address this problem with some combination of educating people about the likely consequences of a decision for future feelings and eliciting what their expectations about future feelings are after that education to control for any remaining mispredictions of future feelings.

standard economic models, the benefits to one's children show up in one's own utility function, but it is not clear that the benefits to others show up in one's own affect in the same proportion as in utility. Third, some people genuinely care about things that contribute to their lives but do not on average contribute to affect. For example, it would not necessarily be a bad decision to pursue excellence even if one knew that the effort would lower one's expected level of affect over a lifetime. Even striving for social rank—a dimension in which affect and utility track each other especially well—provides good examples of the divergence between affect and utility. Think of how many people have knowingly and deliberately sacrificed happiness (affect) for the sake of ambition. Some of these people would do it all over again if they had their lives to live over.

5. Persuasion about Preferences. Since most people do care about happiness at least somewhat to begin with, extolling the wonders of happiness and exhortations to value happiness more highly can often be effective tools for those who desire to persuade others to change their priorities—that is, to change their preferences and the utility function that would be needed to represent those preferences. To the extent that people *are* genuinely persuaded by such arguments, their utility function will shift to be somewhat more tightly related to happiness. The adverb "genuinely" in "genuinely persuaded" is needed to subtract out the effects of social pressure in which people are brought to outwardly assent to something they do not really agree with. From the standpoint of Pareto optimality, there is no reason to question the new, post-persuasion utility function if the persuasion is, indeed genuine. However, it seems only fair that people be made aware that it is not *illogical* to put a low valuation on happiness in one's preferences, if one so chooses in a top-down process of concretizing one's own preferences.[79]

---

[79] Of course, in a large fraction of cases of attempted persuasion about preferences, the desired preferences for the other person will be given the rhetorical label "happiness," "true happiness," "genuine happiness," or "authentic happiness," regardless of how important happiness in the narrow sense of positive affect is in those preferences. For logical clarity (which can be at variance with persuasive power), the phrase "recommended preferences" can be substituted in place of "true happiness" or similar phrases. Aristotle's use of *eudaimonea* (the Greek word for happiness) in the *Nicomachean Ethics* (fourth century B.C.E.) can be seen as an example of using "happiness" as a label for such recommended preferences. Saying this in no way diminishes the cogency of Aristotle's recommendations.

References

Argyle, Michael, 1999. "Causes and Correlates of Happiness," chapter 18 in Daniel Kahneman, Ed Diener and Norbert Schwarz eds., *Well-Being: The Foundations of Hedonic Psychology*, Russell Sage Foundation, New York.

Aristotle, fourth century B.C.E. *Nicomachean Ethics*, available in many modern editions.

Becker, Gary, 1965. "A Theory of the Allocation of Time," *Economic Journal* 75, pp. 492-517.

Benjamin, Daniel, Miles Kimball and Robert J. Willis, 2006. "Cognitive Economics," mimeo, University of Michigan.

J. Bentham (1789) Introduction to the Principles of Morals and Legislation. 1948 edition, New York: Hafner.

Biddle, S., and N. Murtrie, 1991. *Psychology of Physical Activity and Exercise,* London: Springer/Tavistock/Routledge.

Blanchflower, David G., and Andrew Oswald, 2004. "Well-Being Over Time in Britain and the USA," *Journal of Public Economics*, *88*, 1359-1386.

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

Brickman, P., D. Coates and R. Janoff-Bulman, 1978. "Lottery Winners and Accident Victims: Is Happiness Relative?" *Journal of Personality and Social Psychology*, 36 (8), pp. 917-927.

Camerer, Colin, George Loewenstein and Drazen Prelec, "Neuroeconomics: How Neuroscience Can Inform Economics," *Journal of Economic Literature* 43 (March 2005), pp. 9—64.

Charles, Kerwin K., 2002. "Is Retirement Depressing?: Labor Force Inactivity and Psychological Well-Being in Later Life" NBER working paper #9033.

Clark, A., 1999. "Are wages habit-forming? Evidence from micro data," *Journal of Economic Behaviour and Organisation, 39*, 179-200.

Conlisk, John, 1996. "Why Bounded Rationality," *Journal of Economic Literature* XXIV:2 (June), 669-700.

Constantinides, George, 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy* 98 (June), 519-43.

Davis, Melinda, 2002. *The New Culture of Desire: The Pleasure Imperative Transforming Your Business and Your Life***,** Free Press.

Di Tella, Rafael, Alberto Alesina, and Robert MacCulloch. "Happiness and Inequality: Are Europeans and Americans Different?" *Journal of Public Economics* 88, nos. 9-10 (August 2004): 2009-42.

Di Tella, Rafael, and Robert MacCulloch, 1999. "Partisan Social Happiness," *Review of Economic Studies* (forthcoming).

Di Tella, Rafael, Robert MacCulloch, and Andrew J. Oswald, 2003. "The Macroeconomics of Happiness." *Review of Economics and Statistics* 85 (4), 793-809.

Di Tella, Rafael, Robert MacCulloch, and Andrew Oswald, 2001. "Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness," *American Economic Review* 91 (March).

Diener, E., and Lucas, R. E., 1999. Personality and subjective well-being. In D. Kahneman, E. Diener and N. Schwarz (Eds.), *Well-Being: The Foundations of Hedonic Psychology* (pp. 213-229). New York: Russell Sage Foundation.

Diener, Ed, and Christie Scollon, 2003. "Subjective Well-Being is Desirable But Not the Summum Bonum," mimeo, University of Illinois at Urbana-Champaign.

Easterbrook, Gregg, 2003. *The Progress Paradox: How Life Gets Better While People Feel Worse*, Random House, New York.

Easterlin, Richard A., 1974 . "Does Economic Growth Improve the Human Lot?: Some Empirical Evidence," in P. A. David and W. R. Levin eds., *Nations and Household in Economic Growth*, Stanford University Press, pp. 98-125.

Easterlin, Richard A., 1995. "Will Raising the Incomes of All Increase the Happiness of All?" *Journal of Economic Behavior and Organization* 27, pp. 35-47.

Easterlin, Richard A., 2003. "Explaining Happiness," *Proceedings of the National Academy of Science*, 100 (19), pp. 11176-11183.

Frank, R. H., 1985. *Choosing the right pond*. New York: Oxford University Press

Frank, R., 1999. Luxury Fever: Why Money Fails to Satisfy in an Era of Excess, New York: N.Y. Free Press.

Frederick, Shane, and George Loewenstein, 1999. "Hedonic Adaptation," chapter 16 in Daniel Kahneman, Ed Diener and Norbert Schwarz eds., *Well-Being: The Foundations of Hedonic Psychology*, Russell Sage Foundation, New York.

Frey, Bruno S., Simon Luechinger and Alois Stutzer, 2004. "Valuing Public Goods: The Life Satisfaction Approach," Center for Economic Studies and Ifo Institute for Economic Research Working Paper #1158, University of Zurich.

Frey, Bruno S. and Stutzer, Alois, 2003, "Testing Theories of Happiness" In *Paradoxes of Happiness in Economics,* Conference held at the University of Milano-Bicocca, March 21-23, 2003.

Frey, Bruno S., and Alois Stutzer, 2000. "Happiness, Economy and Institutions," *The Economic Journal*, 110 (October), 918-938.

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

Frey, Bruno S., and Alois Stutzer, 2004a. "Happiness Research: State and Prospects," Institute for Empirical Research in Economics Working Paper #190, University of Zurich.

Frey, Bruno S., and Alois Stutzer, 2004b. "Economic Consequences of Mispredicting Utility," Institute for Empirical Research in Economics Working Paper #218, University of Zurich.

Fudenberg, Drew, and David Levine, 2004. "A Dual-Self Model of Impulse Control," mimeo, Harvard University.

Gardner, Jonathan, and Andrew Oswald, 2001. "Does Money Buy Happiness? A Longitudinal Study Using Data on Windfalls," mimeo, available at http://www2.warwick.ac.uk/fac/soc/economics/staff/faculty/oswald/

Geneakoplos, John, David Pearce and Ennio Stacchetti, 1989. "Psychological Games and Sequential Rationality," G*ames and Economic Behavior* 1, 60-79.

Gottfried, Jay A., John O'Doherty and Raymond J. Dolan, 2003. "Encoding Predictive Reward Value in Human Amygdala and Orbitofrontal Cortex," *Science* 301 (22 August), 1104-1107.

Gruber, Jonathan, and Sendhil Mullainathan, 2005. "Do Cigarette Taxes Make Smokers Happier?" *Advances in Economics Analysis & Policy*, Berkeley Electronic Press, http://www.bepress.com/bejeap.

Hamermesh, Daniel S., forthcoming. "Subjective Outcomes in Economics," *Southern Economic Journal*.

Hall, Robert E., and Charles I. Jones, 2004. "The Value of Life and the Rise in Health Spending," National Bureau of Economic Research Working Paper #10737.

Harsanyi, John C., 1953. "Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking," *Journal of Political Economy*, 61 (October), pp. 434-435.

Harsanyi, John C., 1955. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy*, 63, pp. 309-321.

Heckman, James J. and MaCurdy, Thomas E. (1980) "A Life Cycle Model of Female Labor Supply." *Review of Economic Studies* 47: 47-74.

James Heilbrun (2003). "Baumol's Cost Disease," in Ruth Towse, ed., *A Handbook of Cultural Economics*, Edward Elgar Publishing.

Helliwell, John F., 2002. "How's Life? Combining Individual and National Variables to Explain Subjective Well-Being," NBER Working Paper #9065.

Hicks, J. R. and Allen, R. G. D., "A Reconsideration of the Theory of Value", Economica, February/May, 1934, pp.52-60 and 196-219.

Hirsch, F. (1976). *Social limits to growth*. Cambridge, MA: Harvard University Press.

Hirshleifer, David, and Tyler Shumway, 2003. "Good Day Sunshine: Stock Returns and the Weather," *Journal of Finance*, 1009-1032.

Kahneman, Daniel, Ed Diener and Norbert Schwarz, 1999. *Well-Being: The Foundations of Hedonic Psychology*, edited volume, Russell Sage Foundation, New York.

Kahneman, Daniel, 1999. "Objective Happiness," chapter 1 in Daniel Kahneman, Ed Diener and Norbert Schwarz eds., *Well-Being: The Foundations of Hedonic Psychology*, Russell Sage Foundation, New York.

Kahneman, Daniel, Alan B. Krueger, David Schkade, Norbert Schwarz, and Arthur A. Stone, 2004. "Toward National Well-Being Accounts," *American Economic Review*, 94, pp. 429-434.

Kahneman, Daniel, Alan B. Krueger, David A. Schkade, Norbert Schwarz and Arthur A. Stone, 2004. "A Survey Method for Characterizing Daily Life Experience: The Day Reconstruction Method (DRM)," *Science* (December), 1776-1780.

Kahneman, Daniel, Alan B. Krueger, David A. Schkade, Norbert Schwarz and Arthur A. Stone, 2006. "Would You Be Happier if You Were Richer? A Focusing Illusion," *Science* 312 no. 5782 (June 30), pp. 1908-1910.

Kahneman, Daniel, and Amos Tversky, 1979. "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47 (2), 263-291.

Kaprio, J., M. Koskenvuo, and H. Rita, 1987. "Mortality after Bereavement: A Prospective Study of 95,647 Widowed Persons," *American Journal of Public Health* 77, pp. 283--287.

Kapteyn, A., and Wansbeek, T., 1985. "The individual welfare function: A review," *Journal of Economic Psychology, 6,* 333-363.

Kezdi, Gabor, and Robert J. Willis, 2003. "Who Becomes a Stockholder? Expectations, Subjective Uncertainty, and Asset Allocation," mimeo, University of Michigan.

Kim, H., Shimojo, S., O'Doherty, J. P., 2006. "Is Avoiding an Aversive Outcome Rewarding? Neural Substrates of Avoidance Learning in the Human Brain," *Public Library of Science (PLoS) Biology* 4(8), e233."

Kimball, Miles, Helen Levy, Fumio Ohtake and Yoshiro Tsutsui, 2006. "Unhappiness after Hurricane Katrina," NBER Working Paper #12062 (March).

Kreps, David, and Evan Porteus, 1978. "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46, 185-200.

Laibson, David, 1997. "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 62 (2), pp. 443-478.

Lane, Robert E., 2000. *The Loss of Happiness in Market Democracies*, Yale University Press, New Haven, pp. 141-158.

Larsen, Randy J., and Barbara L. Fredrickson, 1999. "Measurement Issues in Emotion Research," chapter 3 in Daniel Kahneman, Ed Diener and Norbert Schwarz eds., *Well-Being: The Foundations of Hedonic Psychology*, Russell Sage Foundation, New York.

Layard, Richard, 2005. *Happiness*. London, UK: Penguin Press. This book is based to a large extent on Layard's 2003 Lionell Robbins Memorial Lectures, London School of Economics.

Lewinsohn, P. M., J. M. Sullivan and S. J. Grosscup (1982). "Behavioral Therapy: Clinical Applications," in A. J. Rush (ed.), *Short-Term Therapies for Depression* (pp. 50-87), New York: Guilford.

Linder, Staffan, 1970. *The Harried Leisure Class*, New York, Columbia University Press.

Lodge, David. (1978) *Changing Places*. Penguin Books.

Loewenstein, George, Ted O'Donoghue and Matthew Rabin (2002), "Projection Bias in Predicting Future Utility," mimeo. Carnegie-Mellon University.

Loewenstein, George, and David Schkade, 1999. "Wouldn't It Be Nice? Predicting Future Feelings," chapter 5 in Daniel Kahneman, Ed Diener and Norbert Schwarz eds., *Well-Being: The Foundations of Hedonic Psychology*, Russell Sage Foundation, New York.

Lucas, Richard E., Andrew Clark, Yannis Georgellis, and Ed Diener, 2004. "Unemployment Alters the Set Point for Life Satisfaction," *Psychological Science* 15 (1), pp. 8--13.

Lupton, Joseph, 2002. "The Habit Liability in Life-Cycle Consumption and Portfolio Choice," mimeo, Federal Reserve Board.

Luttmer, E.F.P., 2004. "Neighbors as negatives: Relative earnings and well-being," NBER Working Paper, No. w10667.

Lykken, D., and Tellegen, A., 1996. Happiness is a stochastic phenomenon. *Psychological Science, 7*, 186-189.

Marmot, Michael, 2004. *How Our Position on the Social Gradient Affects Longevity and Health*, London: Bloomsbury.

MaCurdy, Thomas E. (1980) "An Empirical Model of Labor Supply in a Life-Cycle Setting." *Journal of Political Economy* 89: 1059-1085.

MaCurdy, Thomas (1985) "Interpreting Empirical Models of Labor Supply in An Intertemporal Framework with Uncertainty." In J.J. Heckman and B. Singer, eds. Longitudinal Analysis of Labor Market Data (Cambridge University Press, Cambridge UK).

MaCurdy, Thomas E. and Blundell, Richard (1999) "Labor Supply: A Review of Alternative Approaches." In O. Ashenfelter and D. Card., eds. Handbook of Labor Economics, Vol. 3A. Elsevier: 1559-1695.

Maslow, Abraham, 1943. "A Theory of Human Motivation" *Psychological Review, 50*, pp. 370-396.

Morris, Desmond, 2004. *The Nature of Happiness*, Little Books Ltd., London.

Morris, W. N., 1999. The mood system. In D. Kahneman, E. Diener and N. Schwarz (Eds.), *Well-Being* (pp. 169-189). New York: Russell Sage Foundation.

Murphy, Kevin, and Robert Topel, 2005. "The Value of Health and Longevity," NBER Working Paper # 11405.

Nesse, Randolph M., 2000. Is depression an adaptation? *Archives of General Psychiatry, 57*, 14-20.

Nesse, Randolph M., 2001. Motivation and melancholy: a darwinian perspective. *Nebraska Symposium on Motivation, 47*, 179-203.

Nesse, Randolph M., 2004. Natural selection and the elusiveness of happiness. *Philos Trans R Soc Lond B Biol Sci, 359*, 1333-1347.

Nesse, Randolph M., forthcoming. "Natural Selection and the Regulation of Defenses," *Evolution and Human Behavior.*

Nickerson, Carol, Norbert Schwarz, Ed Diener, and Daniel Kahneman, 2003. "Zeroing in on the Dark Side of the American Dream: A Closer Look at the Negative Consequences of the Goal for Financial Success," *Psychological Science*, 14 (November), 531-536.

Nisan, M., 1972. "Dimension of time in relation to choice behavior and achievement orientation," *Journal of Personality and Social Psychology,* 21, 175-182.

Nordhaus, William, 2006.. "Baumol's Diseases, a Macroeconomic Perspective," NBER Working Paper #12218.

O'Doherty, John P., 2005. "Reward Representations and Reward-Related Learning in the Human Brain: Insights from Neuroimaging," Wellcome Department of Imaging Neuroscience.

O'Doherty, John P., Peter Dayan, Karl Friston, Hugo Critchley and Raymond Dolan, 2003. "Temporal Difference Models and Reward-Related Learning in the Human Brain," *Neuron* 28 (April 23), 329-337.

O'Guinn, T.C. and Shrum, L. J., 1997. "The role of television in the construction of consumer reality". *Journal of Consumer Research,* 23, 278-294

Pareto, Vilfredo, 1906. Manuel of Political Economy trans. by Ann S Schwier, edited by Ann S. Schwier and Alfred N. Page (New York: Augustus M. Kelley, 1971).

Pattanaik, P. K., 1968. "Risk, Impersonality and the Social Welfare Function," *Journal of Political Economy*, 76 (December), 1152-1169.

Patterson, D. R., J. J. Everett, C. H. Bombardier, K. A. Questad, V. K. Lee, and J. A. Marvin, 1993. "Psychological Effects of Severe Burn Injuries," *Psychological Bulletin* 113, pp. 362--378.

Pinker, Stephen, 1997. *How the Mind Works,* Norton.

Putnam, Robert D., 2000. *The Collapse and Revival of American Community,* Simon & Schuster.

Rawls, John, 1951. "Outline of a Decision Procedure for Ethics," *Philosophical Review*, 60, pp. 177-197.

Rawls, John, 1958. "Justice as Fairness," *Philosophical Review*, 67, pp. 164-194.

Rawls, John, 1971. *A Theory of Justice,* revised edition 1999, Belknap Press of Harvard University Press, Cambridge, MA.

Rayo, Luis and Gary S. Becker, 2005. "Evolutionary Efficiency and Happiness," mimeo, University of Chicago.

Reich, J. W. and J. Zautra, 1981. "Life Events and Personal Causation: Some Relationships with Satisfaction and Distress," *Journal of Personality and Social Psychology*, 41, 1002-1012.

Robbins, Lionel C., 1932. *An Essay on the Nature and Significance of Economic Science*. 1984 edition, New York: New York University Press.

Robinson, Michael D. and Gerald L. Clorem 2002. "Belief and Feeling: Evidence for an Accessibility Model of Emotional Self Report," *Psychological Bulletin*, 128, 934-960.

Rosen, Sherwin (1986) "The Theory of Equalizing Differences," in O. Ashenfelter and R. Layard (eds.), *Handbook of Labor Economics*

Samuelson, Paul A., 1938. "A Note on the Pure Theory of Consumer's Behaviour," *Economica* February), 61—71.

Samuelson, Paul A., 1947. *Foundations of Economic Analysis.* Cambridge: Harvard University Press, (enlarged edition, 1983), esp. chapters V, VII, and VIII.

Schkade, D. A. and Kahneman, D. (1998). Does living in California make people happy? A focusing illusion in judgments of life satisfaction. *Psychological Science*, 9 (5), 340-529.

Schor, Juliet B. (1991) *The Overworked American: The Unexpected Decline of Leisure*. New York. Basic Books.

Schultz, Wolfram, and Anthony Dickinson, 2000. "Neuronal Coding of Prediction Errors," *Annual Review of Neuroscience*, 23: 473—500.

Schwartz, Barry, 2004. *The Paradox of Choice: Why More is Less*, New York: HarperCollins.

Schwarz, N. (1987). *Stimmung als Information: Untersuchungen zum Einfluß von Stimmungen auf die Bewertung des eigenen Lebens*. [Mood as information.] Heidelberg, Germany: Springer Verlag.

Schwarz, N. (1990). Feelings as information: Informational and motivational functions of affective states. In E. T. Higgins and R. M. Sorrentino (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2; pp. 527-561). New York, NY: Guilford Press.

Schwarz, N., 1996. *Cognition and communication: Judgmental biases, research methods and the logic of conversation.* Hillsdale, NJ: Erlbaum.

Schwarz, N., 1999. "Self-reports: How the questions shape the answers," *American Psychologist, 54*, 93-105.

Schwarz, N. (2002). Situated cognition and the wisdom of feelings: Cognitive tuning. In L. Feldman Barrett and P. Salovey (Eds.), *The wisdom in feelings* (pp.144-166). New York: Guilford.

Schwarz, N., and Bohner, G., 2001. "The construction of attitudes," In A. Tesser and N. Schwarz (Eds.), *Blackwell handbook of social psychology: Intraindividual processes* (Vol.1, pp. 436-457). Oxford, UK: Blackwell.

Schwarz, N., and Clore, G.L., 1983. Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology, 45*, 513 - 523.

Schwarz, Norbert and Fritz Strack, 1999. "Reports of Subjective Well-Being: Judgmental Processes and Their Methodological Implications," chapter 4 in Daniel Kahneman, Ed Diener and Norbert Schwarz eds., *Well-Being: The Foundations of Hedonic Psychology*, Russell Sage Foundation, New York.

Seligman, Martin E. P., 2002. *Authentic Happiness: Using the New Positive Psychology to Realize Your Potential for Lasting Fulfillment*, Free Press, New York.

Sen, Amartya, 1983. "Poor, relatively speaking," *Oxford Economic Papers, 35,* 153-169.

Seymour, Ben, John P. O'Doherty, Martin Koltzenburg, Katja Wiech, Richard Rackowiak, Karl Friston and Raymond Dolan, forthcoming. "Opponent appetitive-aversive neural processes underlie predictive learning about pain relief," *Nature Neuroscience*.

Seymour, Ben, John P. O'Doherty, Peter Dayan, Martin Koltzenburg, Anthony K. Jones, Raymond J. Dolan, Karl J. Friston, 2004. "Temporal Difference Models Describe Higher-Order Learning in Humans," *Letters to Nature* 429 (June), 664-667.

Shepperd, J. A., Findley-Klein, C., Kwavnick, K. D., Walker, D., and Perez, S., 2000. "Bracing for loss," *Journal of Personality and Social Psychology*, 78, 620-634.

Shepperd, J. A., Ouellette, J. A., and Fernandez, J. K., 1996. "Abandoning unrealistic optimism: Performance estimates and the temporal proximity of self-relevant feedback," *Journal of Personality and Social Psychology*, 70, 844-855.

Sims, Christopher, 2002. "Implications of Rational Inattention," unpublished paper, Princeton University.

Solnick, Sara J., and David Hemenway, 1998. "Is More Always Better?: A Survey of Positional Concerns," *Journal of Economic Behavior and Organization*, 37, 373--383.

Steffick, Diane E., 2000. "Documentation of Affective Functioning Measures in the Health and Retirement Study," HRS Documentation Report DR-005, available online at http://hrsonline.isr.umich.edu/docs/userg/dr-005.pdf

Steptoe, A., J. Kimball and P. Basford, 1996. "Exercise and the Experience and Appraisal of Daily Stressors: A Naturalistic Study," *Journal of Behavioral Medicine,* 21, 363—374.

Stigler, George, 1950. "Development of Utility Theory," Journal of Political Economy 58, pts 1-2 (August-October), pp. 307-327, 373-396.

Strack, F., Martin, L.L., and Schwarz, N. (1988). Priming and communication: The social determinants of information use in judgments of life-satisfaction. *European Journal of Social Psychology, 18*, 429 - 442.

Strack, F., Schwarz, N., Chassein, B., Kern, D., and Wagner, D. (1990). The salience of comparison standards and the activation of social norms: Consequences for judgments of happiness and their communication. *British Journal of Social Psychology, 29*, 303-314.

Strack, F., Schwarz, N., and Gschneidinger, E. (1985). Happiness and reminiscing: The role of time perspective, mood, and mode of thinking. *Journal of Personality and Social Psychology, 49*, 1460 -1469.

Sulls, J., and Wills, T.A. (Eds) (1991). *Social comparison: Contemporary theory and research*. Hillsdale, NJ: Erlbaum.

Thayer, R. E., 1989. *The Biopsychology of Mood and Emotion*, New York: Oxford University Press.

Tsutsui, Yoshiro, Miles Kimball and Fumio Ohtake, 2005. "Koizumi Carries the Day: Happiness and Unhappiness after the Japanese Election," Osaka University and University of Michigan.

Wilson, Timothy D., 2002. *Strangers to Ourselves: Discovering the Adaptive Unconscious*, Harvard University Press, Cambridge, MA and London, England, especially chapter 6, pp. 137-159.

Wolfers, Justin, 2003. "Is Business Cycle Volatility Costly? Evidence from Surveys of Subjective Well-Being," NBER Working Paper #9619.

Woodford, Michael, 2002. "Imperfect Common Knowledge and the Effects of Monetary Policy," in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps,* Princeton Univ. Press.

Wortman, C. and R. Silver, 1987. "Coping with Irrevocable Loss," in *Cataclysms, Crises and Catastrophes: Psychology in Action*, Master Lecture Series, v. 6, Washington, D. C., American Psychological Association, pp. 189-235.

Zamble, E., 1992. "Behavior and Adaptation in Long-Term Prison Inmates: Descriptive Longitudinal Results," *Criminal Justice and Behavior* 19, pp. 409--425.

Zamble, E., and F. Proporino, 1990. "Coping, Imprisonment, and Rehabilitation: Some Data and Their Implications," *Criminal Justice and Behavior* 17, pp. 53-70.