Working Paper No. 492

Are the Costs of the Business Cycle "Trivially Small"? Lucas's Calculus of Hardship and Chooser-dependent, Non–Expected Utility Preferences

by

Greg Hannsgen

Research Scholar, The Levy Economics Institute of Bard College

March 2007

The Levy Economics Institute Working Paper Collection presents research in progress by Levy Institute scholars and conference participants. The purpose of the series is to disseminate ideas to and elicit comments from academics and professionals.

The Levy Economics Institute of Bard College, founded in 1986, is a nonprofit, nonpartisan, independently funded research organization devoted to public service. Through scholarship and economic research it generates viable, effective public policy responses to important economic problems that profoundly affect the quality of life in the United States and abroad.

The Levy Economics Institute
P.O. Box 5000
Annandale-on-Hudson, NY 12504-5000
http://www.levy.org

ABSTRACT

In his presidential address to the American Economic Association, Robert Lucas claimed that the welfare costs of the business cycle in the United States equaled .05 percent of consumption. His calculation compared the utility of a representative consumer receiving actual per-capita consumption each year with that of a similar consumer receiving the expectation of consumption. To a risk-averse person, the latter path of consumption confers more utility, because it is less volatile. Applying Amartya Sen's chooser-dependent preferences to a non–expected utility case, I will counter Lucas's claim by arguing that people have different attitudes toward risk that is imposed and risk that is voluntarily taken on, and that policymakers, in carrying out public duties, must use sorts of reasoning different from those used by the optimizing consumers of neoclassical economic theory.

I.   INTRODUCTION

Imagine that you receive a paycheck every week for a different, random amount, because you work different hours each week. Since you have no bank account, you use your entire paycheck for consumption. It seems reasonable to assume that, other things being equal, you would prefer to have a constant paycheck with the same expected value each month. How much money would you be willing to sacrifice to ensure that you received this mean amount each month and feel equally well off?

Now consider five variants on this question. In all cases, your consumption follows the same random process. In all but the last, you are not allowed to save. In Case I, you have no choice over how many hours to work because your employer chooses your hours. You know that your employer does this because business varies with the weather, and she cannot afford to employ you on a regular schedule. In Case II, on the other hand, you also have no choice about your hours of work (also because business is irregular), but in this case, business varies because macroeconomic policymakers do not assign a high priority to stabilizing the economy. In Case III, your boss frequently flies into rages and cuts or increases your hours capriciously. In Case IV, you are investing a given amount of money on behalf of your daughter in college and you know she must support herself on the randomly varying income from this investment as long as she is a student, without any additional help from you. You feel the mean return on riskless investments is too low. Finally, in Case V, you still have a random income for some reason, but your income is large enough that you are able to save some money occasionally. You choose how much to save each week, but because your income and savings are limited, your consumption each week still varies randomly.

It seems likely that most people would be willing to pay to avoid random variation in their consumption in each of these five situations. Would their willingness to pay be different in the five different cases? Sen (1997) has introduced what he calls "chooser-dependent preferences" under which commodity bundle A is preferred to B if Person 1 makes the choice, while B is preferred to A if Person 2 makes the selection. For example, suppose a host offers us a choice of two pieces of cake of different sizes. We might choose the smaller, out of manners, but prefer to receive the larger, if the host were making the decision.

It is easy to imagine that preferences on the random consumption paths in the five cases might also depend on the chooser, or source of risk: nature, policymakers, an employer, a father, and the consumer. Both the preferences of the chooser and the person being chosen for might be affected by the identity of the chooser. To begin with, in some cases, some notion of a role of responsibility is involved. A worker might be more angry in Case II (where poor policy caused randomness) more than Case I (where the weather was the culprit), because people in whom power is entrusted have caused unnecessary hardship. The chooser herself in Case II might feel ashamed of her conduct because she had failed in her official duties, but not be so concerned about how she managed her personal finances. Sen labels these concerns "fiduciary." Certainly, the legal and ethical rules of fiduciary responsibilities entail many special responsibilities that do not extend to personal matters or, for that manner, to ordinary relationships. According to modern versions of the Hippocratic Oath, to take one example, doctors have duties not to treat illnesses in which they do not have adequate expertise and to put the welfare of the patient above all else.

Beyond ethical and legal arguments, there are other reasons to believe that preferences over consumption streams are chooser dependent. There is an enormous literature in psychology arguing that an important determinant of mental health and successful functioning is "locus of control"—whether an individual believes events are dictated by her own actions or by forces beyond her control (Seligman and Maier 1967). Recall our Case V, where the consumer could, to some extent, choose her own consumption path, in spite of the randomness of income, a situation that might result in a better sense of internal control than when randomness was determined by policy. Goldsmith and Veum (1996) have in fact argued that a change from internal to external locus of control is an important impact of unemployment. So, the person for whom policy is chosen, and not just the policymaker, is influenced in their welfare judgments by who makes the decision.

Robert Lucas has recently attempted to estimate the costs of the business cycle, treating it as an instance of the random-checks problem (2003, 1987). In his simplest calculation, he assumes that the effects of perfect monetary policy would be to completely eliminate all random variation in aggregate consumption. He first evaluates the total utility of a single representative consumer receiving actual average U.S. consumption over a particular time period. He then

compares this with the hypothetical utility that the representative consumer would have had if the average consumption variable in the utility function were replaced with its unconditional mean in each period. To put the costs in financial terms, he finds the hypothetical proportional boost in the random consumption path needed to make the consumer indifferent between the random path and the deterministic one. Lucas arrives at the estimate that the costs of the business cycle are equivalent in utility terms to .0005 times consumption. He then argues that other priorities, designed to boost trend (mean) growth, offer much higher benefits.

In this paper, I seek to deal with several questions raised by Lucas's approach, most of them relating to the issues of chooser-dependent preferences I have described. In the next section, I review Lucas's contribution in more detail and briefly describe some of the arguments of his critics. In Section III, I attempt a theoretical analysis of the problem using a class of chooser-dependent, non–expected utility preferences that allow one to look at the problem in a very general way. This class permits consumers to be concerned with many different stochastic features of the consumption process, such as autocorrelation, that seem important in this setting. Giving several examples, I will show how such preferences could, in principle, result in a different answer to the problem. But in his book chapter and paper, Lucas is primarily concerned with measurement: he calibrates a sparsely parameterized utility function to arrive at a specific cost. In Section IV, I therefore look at the possible hazards associated with such calibration when chooser-dependency issues are involved. I link my argument to several new ideas that have been discussed recently under the rubric "behavioral welfare economics" (Bernheim 2007; Rabin and Koszegi 2007; Camerer 2007; Gul and Pesendorfer 2007). The final section concludes.

II.   REVIEW OF THE LITERATURE

In 1987, and again in his 2003 Presidential Address to the American Economic Association, Robert Lucas made the remarkable claim discussed above. Here we summarize the nature of his calculation, then describe several existing critiques and elaborations.

Figure 1. Lucas's Notion of the Costs of the Business Cycle
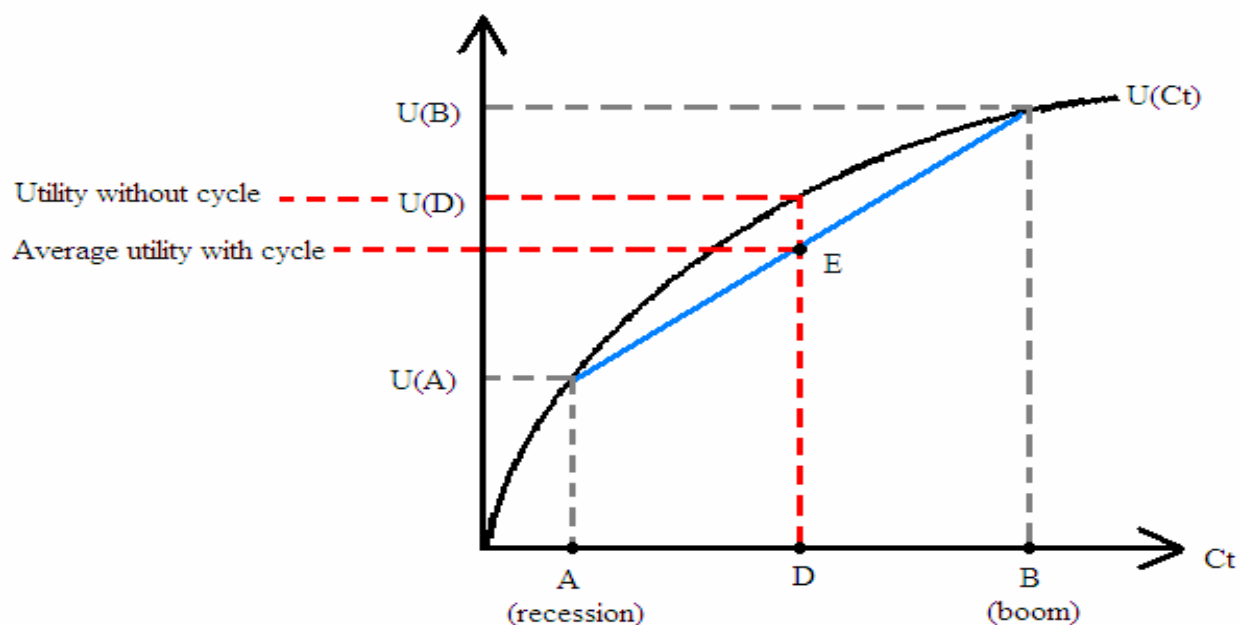
Lucas argues that the costs of the business cycle are due to the fact that utility in any given period is concave, or risk averse. In Figure 1, we show such a utility function, U(C), for total consumption in any given period. Imagine two hypothetical random consumption paths. In Path 1, consumption is at A (recession) one-half of the time and at B (boom) in the other periods. In Path 2, consumption is always at the midpoint of A and B, which is labeled D on the graph. In Path 1, average utility is given by the height of point E, the midpoint between points [A, U(A)] and [B, U(B)] (if the probability of a recession were lower than one-half, the average utility would be farther to the right on the straight line). In Path 2, on the other hand, utility is always U(D).[1] Lucas argues that consumption is generally something like Path 1, with stochastic oscillation between boom and bust. If we were to eliminate the business cycle, we would have a

---

[1] I abstract here from the secular trend in consumption, which Lucas does consider. This does not affect the principle of his analysis.

path like Path 2, which is always at its mean. Lucas seeks to measure how much people would be willing to pay, in terms of mean consumption, to move from Path 1 to Path 2. The costs are measured by discounting utility in each period until infinity and adding up the utilities.

The equation used to make the calculation is:

$$E\sum_{t=0}^{\infty} \beta^t U((1+\lambda)C_t) = \sum_{t=0}^{\infty} \beta^t U(E(C_t))$$

$0<\beta<1$
$U'(C)>0$
$U''(C)<0$

where E is the expectation taken at time 0, β is the discount factor, t is the time period, and λ is the variable for which Lucas is solving, his measure of the cost of the business cycle. Lambda gives the proportional upward shift in consumption needed to make the representative agent as satisfied as she would be if she had average consumption in every period.

To arrive at his numerical estimate of λ, Lucas uses the "constant relative risk aversion" period utility function:

$$U(C_t) = \frac{C_t^{1-\rho}}{1-\rho}$$

where ρ is a constant known as "relative risk aversion." Lucas assumes that the log of C_t follows independent and identical normal distributions. He shows that in this setup, the costs of the business cycle are

$$\lambda \approx (1/2)\rho\sigma^2$$

where σ^2 is the variance of the log of consumption. The costs depend upon ρ, because this parameter measures the concavity of the period utility function, which was important in our graphical analysis of Figure 1. Stating that macroeconomists generally use values of ρ between

one and four, Lucas uses ρ=1.[2] As mentioned before, Lucas finds that λ=.0005. Pointing out that, of course, even in an ideal world, policymakers could not reduce the variance of consumption to zero, he considers empirical evidence as to what percentage of the variance is due to demand ("nominal") shocks.

Before discussing the literature spawned by Lucas's contribution, it might be useful to point out that his general approach to evaluating policies is part of a larger project, which might be called the "new welfare economics of macroeconomics." This approach uses microfounded representative-agent models to determine the effects of different policy regimes on economic variables, and then calculates the expected discounted utility associated with each regime. Textbooks by new Classical and neoclassical Keynesian authors, in fact, frequently cite the possibility of such welfare analyses as one of the chief benefits of the project of developing a microfounded macroeconomics (Sargent 2004; Obstfeld and Rogoff 1996). This paper can be seen as a critical case study of one application of this methodology. The paper focuses on one feature that microfounded macroeconomics has in common with almost all microeconomic analysis: an assumption that bundles of commodities consumed, not relations among people, are what drives all behavior in the economic sphere (Anderson 1993).

The remainder of this section will touch on some of the main points made in contributions subsequent to Lucas's book and paper. A key problem with the representative agent approach is that, in reality, the economy comprises numerous agents, whose consumption is affected differently by the business cycle. For example, while *aggregate* consumption rarely falls in a recession by more than a few percent, many unemployed individuals experience much greater falls in their consumption during a recession, and others continue to consume the same amount or more. There is a theoretical rationale for ignoring individual variation in consumption: in models of risk, it is often assumed that there are *complete markets*, which means that agents can insure against any individual-level risk, including unemployment. To see the general idea, suppose that the population is "very large," so that we can invoke the law of large numbers, and that each employee (i) has a "job" that is like a security that pays off a wage

---

[2] Strictly speaking, the period utility function is found by taking the limit of a slightly different constant relative risk aversion period utility function as ρ approaches one. This gives a log function.

$$W_{it}=a_{it}Y_t,$$

in period t, where

$P(a_{it}=u) = .05$ for all i and all t; u=constant

$P(a_{it}=e) = .95$ for all i and all t; e=constant

Y_t is national income in period t

e>u

In this model, unemployment is five percent and the unemployed receive a lower share of national income than the employed. The complete markets assumption is that an "unemployment insurance" scheme, or something like it, enables all workers to own a diversified "portfolio" of all jobs, earning a return of $W_t=((.5)u+(.95)e)Y_t$ (the average income of all jobs). In this model, what is important is the variation in Y_t (aggregate income), not in the income of individual jobs, W_{it}. Of course, creating the appropriate form of insurance is an important policy problem, from Lucas's perspective, but he sees the costs-of-the-business-cycle issue as separable from that one. Certainly, he is right in a sense: unemployment insurance is not generally seen as a responsibility of central bankers and other macro policymakers. But matters are not as clear cut as a single complete-markets model would have it. This paper will not attempt to sort out this complex and thorny issue, which has generated many articles and books of its own.[3]

Nonetheless, some work has been done within Lucas's framework to grapple with the problem of "heterogeneous agents." Krusell and Smith (1999) have estimated a model in which no insurance is available, but workers can self-insure by saving if their incomes are high enough. They calibrate their model with micro data on wealth and employment. In their paper, roughly speaking, the costs of the business cycle are the savings that could be obtained by eliminating only aggregate shocks to Y_t, not the randomness in a_{it}. They find that the steady-state costs of unemployment are 20 times as high as Lucas's estimate. Storesletten, Telmer, and Yaron (2001)

---

[3] Another "elephant in the room" is the issue of utility from work and leisure. Lucas argues that since leisure confers positive utility, his estimate is upwardly biased by the omission of leisure from the utility function (1987). Though most people would probably enjoy a vacation, work clearly has positive effects on well-being, conferring a sense of contributing to society and a feeling of accomplishment.

conduct a similar exercise with an overlapping generations model, taking into account the fact that the a_{it} may be more volatile in recessions. They estimate the costs of the business cycle at .00144 of consumption (see also Mukoyama and Şahin 2006).

A number of contributions have altered Lucas's assumption of a constant relative risk aversion utility function and a log-normally distributed consumption process. Obstfeld (1992) uses a non–expected utility function, which allows for separate parameters for risk aversion and the intertemporal elasticity of substitution (Weil 1990). He notes that with this setup, persistent shocks may have different effects from nonpersistent ones. This seems important to the business cycle issue, because autocorrelation, and perhaps nonstationarity, is one obvious feature of real-world business-cycle shocks. With a unit-root consumption process and Lucas's assumption of unitary risk aversion, the costs of the cycle amount to between .000242 and .001783, depending on the value of the elasticity of intertemporal substitution; these costs rise to between .005198 and .036957 for risk aversion of 20. We touch on papers that estimate utility function parameters later in this section.

Yellen and Akerlof (2004) have provided a broad critique. They argue that stabilization policy can probably affect the mean of consumption, not only its variance. They also stress that unemployment could have nonlinear effects on utility. For example, a week of unemployment may have greater costs during times of high aggregate unemployment, since longer spells of unemployment are more common then. Long spells have disproportionately large utility costs, for example, because people have fixed financial commitments, such as home mortgages.

Alvarez and Jermann (2004) measure costs by using observed asset prices. Essentially, they construct a composite security that yields a stream of returns "similar" to the observed process of aggregate consumption. They then determine from the prices of the underlying securities what people would be willing to pay for the composite security. They find that the costs of *all* consumption uncertainty are extremely large, but the costs of removing the business cycle are between .0008 and .0049. Alvarez and Jermann use a form of non–expected value functional similar to the utility functionals used in Section III of this paper.

Wolfers (2003) directly estimates the costs of the business cycle using surveys of subjective well-being, instead of a utility function. He looks at the effects of levels of

unemployment and inflation, but the findings most relevant to Lucas's paper show that halving the standard deviation of unemployment would increase average well-being by the same amount as reducing average unemployment by one quarter of 1 percent. Consistent with Yellen and Akerlof's argument, Wolfers finds that the volatility of unemployment has nonlinear effects on subjective well-being.

Since Section IV discusses efforts at measurement, it is important to mention the extensive efforts to calibrate constant relative risk aversion utility functions.[4] Perhaps the highest estimates come from data on household asset holdings. The "equity-premium puzzle"—low household equity holdings, despite high average returns—is known to imply very high relative risk aversion in a standard model. Kandel and Stambaugh (1991) argue that, in order to match the first moments of asset returns, they must assume a relative risk aversion of 29. However, this estimate exceeds two previous calibrations using asset holdings. Friend and Blume estimate values of "well in excess of one and probably in excess of two" (1975). Hansen and Singleton (1982), using a generalized instrumental variable technique, find that values may well be negative.

Szpiro (1986) estimates risk aversion at between 1.2 and 1.8 by comparing total property and liability insurance in force with the total value of all property, taking into account overhead costs and profits. Barsky, Juster, Kimball, and Shapiro (1997) report ranges of relative risk aversion for respondents to the 1992 Health and Retirement Study, who were asked which of several hypothetical, large gambles on lifetime income they would accept. Sixty-five percent of all respondents chose a gamble that would imply an expected risk aversion of 15.7, though significant numbers had much lower risk aversions (p. 548). A now-large literature attempts to estimate risk aversion from the behavior of television game show contestants. For example, Hartley, Lanot, and Walker (2006) conclude from data on the show "Who Wants to Be a Millionaire?" that relative risk aversion is approximately one.

It is clear that many theoretical limitations and numerical ambiguities in the costs of

---

[4] In the constant relative risk aversion utility function, which is still very widely used, relative risk aversion also happens to be equal to the reciprocal of the intertemporal elasticity of consumption (assuming optimizing behavior). So, some might say that estimates of this parameter are also germane. However, in more general non–expected utility functions such as Weil's (1990), relative risk aversion and the intertemporal elasticity are independent parameters. Both would have effects on the costs of the business cycle, as Obstfeld (1992) points out, and the model in Section III can include such effects, but our studies of calibration procedures in Section IV will concentrate on

business cycles are well recognized. This paper attempts to deal with some of those problems. In the next section, rather than enter the debate over the proper functional form of the utility function and the characteristics of the consumption process, I use a very general non–expected utility functional. I also make no strong assumptions about how macro policy affects the consumption process, allowing for the observation of Akerlof and others that moments other than the variance could be involved. In Section IV, I deal with some of the obvious problems of using estimates of relative risk aversion like the ones discussed in the previous two paragraphs to judge the welfare effects of any policies, including macroeconomic stabilization. The entire paper concentrates on the implications of Sen's chooser-dependence idea, which, to my knowledge, has not been applied to the costs of the business cycle.

III.   A MODEL OF CHOOSER-DEPENDENT PREFERENCES AND ITS IMPLICATIONS FOR LUCAS'S CALCULATION

Sen (1997) symbolizes his notion of chooser-dependent preferences with a preference relation, defined over some choice set Q, like

$$\prec_{i,j}, \ \succ_{i,j}$$

where the symbols mean i disprefers (prefers) when j chooses.[5] To use Sen's example, discussed in the introduction, suppose A is a large piece of cake and B is a small one. Sen points out that Person 1 may prefer A when Person 2 chooses but would prefer B if she is offered a choice from a plate (to be polite). In this case

$$B \prec_{1,2} A, \ B \succ_{1,1} A$$

I wish to apply this concept to consumption streams. To keep things manageable, I will assume that the representative consumer consumes a random vector of amounts of a single good

$$(c_t) = (c_0, c_1, c_2, ..., c_T)$$

risk aversion.

11

in time periods t=0,1,2,3,....T, with

$$0 \leq c_t \leq c*$$

(it is convenient to assume that lives are finite, but T should be considered a fairly large number. Also, c* can be as large as necessary). I will abuse notation slightly here, by using the left side of the previous equation to indicate the entire consumption vector and also sometimes, as in the inequality, using c_t to refer to the value at a particular time t.

Under very general conditions, there exists a distribution function for this vector

$$F(c_0, c_1, c_2, ..., c_T) = F(c_t)$$

This allows any sort of autocorrelation, heteroskedacticity, nonstationarity, etc., in the consumption process (consumption is deterministic at time zero, so the marginal distribution for c_0 is degenerate). All distribution functions F of this type on the T-dimensional closed interval

$$S = \times_{t=0}^{T}[0, c*]$$

can be characterized by particular necessary and sufficient features. I call the set of all such functions with these features F*.

Rather than starting with an ordinal preference relationship, this model begins with utility functionals. Instead of saying utility is equal to the expected sum of the discounted values of period utility, I would like to use a utility functional—a mapping of *distributions* F onto the real numbers. A number of utility functionals defined on F* exist:

$$U_{i,j}:F* \rightarrow R$$

which is the utility for Person i, in period 0, of the consumption stream, when Person j chooses.[6]

---

[5] He also discusses preferences that depend on the choice set, but I will be dealing with that subject only indirectly.
[6] Note that in general a utility function of this type may be time inconsistent. That is, the relative utilities of consumption path A may seem preferable to B at time zero, but then I may change my mind when I get to time t+k.

This sort of utility functional is similar to a utility function defined on lotteries. This class of non–expected utility functionals defined on nonindependent multivariate distributions has been studied without chooser-dependence (Karni 1989; Safra and Segal 1993). In all cases, we assume that the same person is the actual consumer/citizen. By defining utility on *distributions*, rather than actual consumption bundles, I allow the consumer to care about many different features of the consumption process.[7]

Of course, in practice, the consumption stream enjoyed by a consumer is not chosen by any one person or political body. One way of looking at the problem is to think of a macroeconomic authority taking into account the actions of other agents—other policymakers, consumers, and firms—as a function of their own decisions. Then, the choice set of the macroeconomic decision maker is a set of joint outcomes of the strategies of all agents, given the strategies of the macroeconomic authority. In other words, the macroeconomic authority would be choosing among a set of strategies, each expected to lead to a certain consumption path, given the strategies of the other public and private agents in the economy. Then, from the perspective of the policymaker, the choice would be, in a sense, entirely her own, and the appropriate utility function would have the policymaker as j. It is also true that, even in a given situation, *perceptions* or *beliefs* about who is choosing may differ, depending on the circumstances or individual.

The class of non–expected utility functions just introduced is helpful because important theoretical objections and empirical anomalies have arisen in connection with the standard expected, discounted utility (DU), infinite horizon model, in which utility is given by

$$U_{DU}[F] = E \sum_{t=0}^{\infty} \beta^t U(c_t)$$

where U(c) is a standard period utility function, such as the constant relative risk aversion form used by Lucas.

There is no way around this problem in welfare analysis, except to evaluate utility from a given time period (in fact, the plausibility of time-inconsistent preferences has been one of the major critiques of the DU model).

[7] It turns out that under fairly weak assumptions about the "acquisitiveness" of the underlying preference relation, and with some smoothness conditions, one can say a great deal about the nature of the functional U. Specifically, it

Three key features of the expected DU model have been called into question. The first is the use, even in a deterministic setting, of a discounted sum of the values of a single period utility function (Frederick, Loewenstein, and O'Donoghue 2002). Second, in the expected DU model, utility is expressed as the expected value of a function of the random vector or sequence, a feature that brings problems of its own (Machina 1987). Third, in this model, dislike of risk is modeled solely by concavity of the period utility function, as measured by relative risk aversion (Rabin and Thaler 2001).

Numerous exceptions exist to all three of these features, but I will just give a flavor of them. The first feature is inconsistent with evidence that people may discount $100 deferred from today until one week from today more than they discount $100 deferred from 52 weeks from now until 53 weeks from now. Also, it seems that many people dislike roller-coaster utility paths or prefer paths that rise over time. The second feature of the expected DU model (the expected-utility aspect) is subject to numerous classical and behavioral paradoxes, such as the Allais paradox. The risk aversion aspect of the theory (which is not completely separate from the expected utility issue) is subject to the point that while most people certainly dislike risk in some sense and seek to avoid it, their behavior is not well explained by "risk aversion" (concavity of the utility function). Matthew Rabin (2000) has shown that with a concave utility function in money, a mild unwillingness to take small risks (say, turning down a bet offering a 50 percent chance of losing $100 and a 50 percent chance of gaining $110) implies turning down larger bets on very favorable terms (a 50 percent chance of losing $1,000 and a 50 percent chance of winning $100 billion). This statement holds regardless of how wealthy the gambler is to begin with. Concern about risk may be better modeled using ideas forwarded by psychologists and behavioral economists, such as disutility from dreading the future, myopic neglect of catastrophic risk, and loss aversion (see, for example, Loewenstein, Hsee, Weber, and Welch 2001; Krantz and Kunreuther 2006).

As an example of the flexibility of the non–expected utility preferences used here, we could have

---

can be written as a functional of the marginal distributions (Safra and Segal 1993).

$$U_{1,1}[F] = -\sum_{t=0}^{T} \theta_t F_t(\pi_t)$$

where the θs are positive weights, and

$$F_t(c_t) = F(c^*, c^*, c^*, \ldots, c_t, \ldots, c^*)$$

is the marginal distribution of consumption in a particular period t. This utility functional is the negative of a weighted sum of the probabilities that income will be below π_t, say, the poverty level, in each period t. Of course, by this measure, the costs of the business cycle in an industrialized country are probably zero, because average consumption rarely, if ever, falls below the poverty line. But the π's could be living standards of some type, possibly rising over time. If the chooser were the government, this might be a good utility function; perhaps many individual citizens would act differently, gambling everything to maximize the possibility of becoming wealthy.

Central banks are interested in preventing recessions and in promoting growth—responsibilities that do not agree perfectly with individual objectives. As an example, consider the following central bank or welfare agency objective function

$$U[F, c'_0] = \int_A 1 \, dF(c_t)$$

where

$$A = \{c_t : c'_0 \le c_1 \le c_2 \le \ldots c_{T-1} \le c_T \le c^*\}$$

This function is gives the probability that c never goes down, conditional on the known initial level of consumption, c'_0.

We have noted that psychological considerations might make risk more costly when it is imposed by another person. If the utility functional is smooth, it turns out that risk aversion can be rigorously defined (Karni 1989). A simpler approach is a chooser-dependent functional in

which risk—measured by the variances, as in Lucas's case—is more negatively weighted relative to averages when one person is choosing for another, and the autocovariances are negatively weighted. Also, the discount rates can be different for each period, allowing for hyperbolic discounting, etc.

$$U_{i,j}[F] = \Phi_1^{\delta} E \sum_{t=0}^{T} D_1(t) C_t - \Phi_2^{\delta} \left[\sum_{t=1}^{T} D_2(t) \sigma_t^2 \right] - \Phi_3 \left[\sum_{t=2}^{T} D_3(t) \rho_t^2 \right]$$

where $\Phi_1 > 1$, $0 < \Phi_2 < 1$, $0 < \Phi_3 < 1$; $\delta = 1$ when $i = j$ and $\delta = 0$ otherwise; the D(t)s are discounting functions, which take on positive, real values for integers between 0 and T; σ_t^2 are the variances of the (possibly nonstationary) consumptions; and ρ_t^2 are their first-order autocovariances. Utility functionals can also be devised in which the chooser does not know the true distribution.[8]

It is not the purpose of this section to choose a particular utility function, only to suggest the variety that might exist, once chooser-dependency is taken into account. Also, some of the examples show that utility functions for the choices of consumers do not always make sense for policymakers, and vice-versa. In the rest of this section, in which the costs of the business cycle are analyzed, no particular functional is assumed.

Within the context of the utility functions analyzed in this section, suppose that H is the distribution of the consumption path with stabilization policy and J is the distribution with no

---

[8] A nonchooser dependent version of this class of preferences, with the added wrinkle that the chooser could not perceive the true distribution of consumption, would be

$$U[F] = \int_S w(c_t) dg(F(c_t))$$

where w is a continuous, real-valued utility function of all T consumption bundles, defined on the domain S, and g is a function that preserves the necessary and sufficient conditions for a distribution. The function g distorts the distribution, resulting in an incorrect estimate of future utility. An example of such a function would be

$$g(F(c_0, c_1, ..., c_T)) = (F(c_0, c_1, ..., c_T))^{\alpha}$$

where $\alpha > 0$. For $0 < \alpha < 1$, this results in an overly optimistic view about the future and an overestimate of expected future utility. For $\alpha > 1$, the decision maker makes the reverse error. This function demonstrates that the utility function here is capable of encompassing a simplified version of "animal spirits" in which the person acts as if she "knows" an incorrect distribution (Keynes 1936).

stabilization policy. H could differ from J in any number of ways: in the spirit of Lucas's exercise, H's variance or autocovariance could be lower, or it could take a particular value with probability one. Let

$$U(\eta, F)$$

be utility when

$$(\eta c_0, \eta c_1, \eta c_3, \ldots, \eta c_T)$$

is distributed according to F. The same utility functional without the eta argument will refer to the case where

$$\eta = 1$$

Then, in Lucas's framework, the costs of the business cycle could be measured using the equation

$$U(1, H) = U(\eta, J)$$

Note that a smaller eta indictes larger business-cycle costs, unlike Lucas's lambda. Multiplying the consumption vector by eta inevitably affects the variance and other moments of a distribution, not just the mean, but this generally happens when a variable grows larger or smaller. U would still be strictly decreasing in eta, as long people are acquisitive in some sense.

We have argued that it is sometimes useful to include i, j subscripts when utility is evaluated. Earlier, we identified the utility function as the utility of i when j was the chooser. It is important to note that i may not be the consumer or citizen here, even though the arguments in the function relate to that person's consumption. Here, it is always the citizen/consumer who is doing the consuming, but the "subject" of the utility, i, may be some sort of fiduciary, such as a policymaker. Then, the function might tell us how that person should act. Sen's interest in introducing these preferences was to describe *choice* behavior, not welfare, but clearly a kind of welfare or normative judgment is involved in the issues we discuss here. Nonetheless, the

explicit motivation for some of the choice behavior studied by Sen in his article was some form of welfare consideration.

One of the two main points of this paper is that the utility functionals for policy analysis should not be the same ones used to analyze the behavior of private consumers, p. It may be true that, using the true, chooser-dependent utility functions

$$U_{p,p}(1,H) = U_{p,p}(\eta, J)$$

with eta being the same as before. However, the real subject and chooser in this problem are either the macroeconomic authorities (m) and themselves, or the citizenry and the macroeconomic authorities. Then, if people resent random variation imposed by indifferent policymakers more then variation that is self-chosen, it may be true that

$$U_{p,m}(1,H) = U_{p,m}(\eta^*, J)$$

where $\eta^* < \eta$.

Similarly, if we are looking at how the macroeconomic policymakers make their own choices, measuring the appropriate costs of the business cycle may yield

$$U_{m,m}(1,H) = U_{m,m}(\eta^{**}, J)$$

where $\eta^{**} < \eta$
if the policymakers believe they have a special duty, above and beyond anyone's personal utility, to be mindful in making their decisions of the certain kinds of welfare costs to others.[9]

---

[9] The last example raises an issue that has been discussed by critiques of neoclassical and utilitarian theories of value and rationality. There are limits to modeling the ethical obligations of decision makers with any sort of utility function. If my utility function depends on discharging my obligations *and* on my material possessions, the implication may be that I would always be willing to violate my ethics for a sufficiently large bribe. Or, similarly, some such preferences might imply citizens angry over injustice or incompetent policy could always be placated with patronage. Our preferences, however, are always on behalf of the consumer, even when they are used to

18

The work of this section has been directed at noting how, in principle, the costs of the business cycle could be chooser dependent, in the way Sen suggested. In the process, I have used a much broader set of preferences than those considered by Lucas. Lucas's actual calculation used a calibration approach common in welfare economics, to which I will turn next.

IV.   THE IMPLICATIONS OF CHOOSER-DEPENDENCY FOR QUANTITATIVE MEASUREMENTS OF BUSINESS CYCLE COSTS AND EMPIRICAL WELFARE ANALYSES

Recall that Lucas measures the costs of the business cycle by lambda, using the constant relative risk aversion preferences

$$E\sum_{t=0}^{\infty} \beta^t \frac{((1+\lambda)C_t)^{(1-\rho)}}{1-\rho} = \sum_{t=0}^{\infty} \beta^t \frac{(E(C_t))^{(1-\rho)}}{1-\rho}$$

Recall that Lucas's calculation depends crucially on rho, which is known as relative risk aversion. Lucas cites certain values for rho that macroeconomists consider plausible. As mentioned in Section II, this parameter can be, and has been, gauged in a number of ways: the behavior of game show contestants, the amount of insurance in force, survey questions on hypothetical lifetime income gambles, and so on. Of course, as mentioned in the previous section, there are numerous empirical exceptions and theoretical objections to this form of utility function. But suppose that this is the right utility function, or that one is conducting the analysis with some more plausible function that also lacks chooser-dependence.

In conducting welfare analyses, using real-world behavior to calibrate utility functions or other measures of costs and benefits is a common technique that has recently been criticized (Camerer 2007). The idea is essentially a variant of revealed preference. I will show how the

---

describe the decision making of a policymaker. So a bribe to the policymaker would not enter the utility function. Moreover, absolute prohibitions could be modeled with functions that, say, took on a higher value for any distribution F where the probability of being in poverty was zero in all periods.

idea works, then illustrate what might go wrong. The economist does an empirical study of market behavior, using the utility function that needs to be used for the welfare analysis. In our case of non–expected value preferences over finite consumption streams, we might observe a choice between two alternatives, K and L. These could be securities, insurance policies, or moves in a game.

For our parameterized utility function, suppose that we calculate that

$$U_{\rho_1}(K) > U_{\rho_1}(L)$$
$$U_{\rho_2}(K) < U_{\rho_2}(L)$$

where the two rhos are parameter values; they could be levels of relative risk aversion, or any other parameter in the utility function. The consumer chooses alternative K. In light of the consumer's choice, or maybe several choices, we would conclude that

$$\rho = \rho_1$$

rather than

$$\rho = \rho_2$$

We would then conduct the welfare analysis discussed earlier, using the following equation

$$U_{\rho_1}(1, H) = U_{\rho_1}(\eta, J)$$

But I argue that a problem can arise here: the use of empirically estimated parameters for welfare analysis can be misleading when the social context of the market choice used for calibration of the parameter is different from the social context of the policy decision.[10] This problem arises

---

[10] The idea that welfare analysis can be contaminated if the social conditions of choices observed by the economist (such as the degree of autonomy) are inappropriate for the welfare judgment at hand is developed in Anderson's critique of cost-benefit analysis (1993). Sen (1973) discusses how revealed preferences may not reflect welfare or

here because the true utility function must be indexed by i and j, and the studies used for calibration mentioned above almost all use observations of voluntary, private decisions, where i and j are both the consumer. There is a mismatch between the choosers in these empirical studies and the choosers in macroeconomic policy decisions.

Exactly where does the problem arise? We have arrived at the previous parameter estimate by observing the choice of K over L. So, it is true that, assuming optimizing behavior, even with the properly indexed utility functions

$$U_{\rho_1,i,j}(K) > U_{\rho_1,i,j}(L)$$
$$U_{\rho_2,i,j}(K) < U_{\rho_2,i,j}(L)$$

for the appropriate i and j. Indeed, it may be true that

$$U_{\rho_1,i,j}(1,H) = U_{\rho_1,i,j}(\eta,J) \qquad (1)$$

where eta is the same as before. But, let us say that the welfare analysis, say the estimate of the costs of the cycle, is conducted under conditions in which i' is choosing, and j' is the subject of the utility function, with

$$i \neq i'$$

or

$$j \neq j'$$

Now, getting back to the insurance study, or some other form of calibration, recall that K

---

self-interest.

was chosen over L. But, if we had done the study with the same choosers as the policy problem, that is, with the same chooser and subject as "insurance purchasers," we might have found a different result if it were true that

$$U_{\rho_1, i', j'}(K) < U_{\rho_1, i', j'}(L)$$
$$U_{\rho_2, i', j'}(K) > U_{\rho_2, i', j'}(L)$$

This implies, that for a policy choice involving i' and j', the better parameter to use is in fact ρ_2 rather than ρ_1. This paper does not address how such a calibration might be done. Conceivably, for some cases when the chooser was a fiduciary, one could observe the portfolios of trusts.

In any case, with the correct parameter value in hand, we use the following equation

$$U_{\rho_2, i', j'}(1, H) = U_{\rho_2, i', j'}(\eta, J)$$

for the business-cycle-costs problem. There is no reason why one would necessarily arrive at the same eta as in (1).

Assuming one were interested in using the constant relative risk aversion utility function, it might be that a correct, chooser-dependent version would be

$$U_{i,j}[F] = E \sum_{t=0}^{\infty} \beta^t \frac{C_t^{(1-\gamma^\delta \rho)}}{1 - \gamma^\delta \rho}$$

$0 < \gamma < 1$

where again δ=1 if i=j, and δ=0 otherwise

This function increases risk aversion for situations in which someone is choosing for another person. This would result in an increase in the measured costs of the business cycle. This may well not be the proper utility function; in fact, it is probably not. It simply exemplifies how things might change if chooser-dependence were taken into account.

Here, the focus has been on preferences that depend on the identity of the chooser. Some similar observations have recently been made about problems in welfare analyses when utility or choice depends upon various "ancillary conditions" of the type studied by behavioral economists (Rabin and Koszegi 2007; Camerer 2007; Bernheim 2007; Gul and Pesendorfer 2007). The point of their research is to point out that when welfare economists make estimates of parameters for the purpose of policy analysis, they may observe choices that are influenced by irrelevant ancillary conditions, such as the "frame" of the decision. These are often clearly "mistakes" that are recognized by people when they are pointed out. To base a welfare analysis on a parameter estimated in this way is difficult and may call for special techniques, heretofore unknown, to take into account the importance of the ancillary conditions. Rabin and Koszegi (2007) argue that preferences may be "revealed," even if they are not "implemented" in the form of actual behavior.

Clearly, this problem may apply to the measurement of risk aversion. In the case of risk aversion, people are sensitive to various contextual and ancillary conditions. Even when these anomalies are pointed out to experimental subjects, the subjects often do not agree that their behavior was a "mistake" in any sense. Take the case of insurance purchases, which has often been used to estimate relative risk aversion (see the references in Section II). Subjects often report, for example, that they purchased insurance for "peace of mind," a form of "anticipatory utility" (Krantz and Kunreuther 2006). They may be more likely to buy insurance in contexts in which anxiety is heightened, e.g., the purchase of flight insurance in an airport. They often overinsure against minor losses; this is explained by "myopic loss aversion." An example would be a reluctance to invest in stocks if portfolios are anxiously checked for losses each quarter. Also, people tend to underinsure against catastrophic events. An example would be the difficulty in selling flood insurance, especially when there has been no recent flood in a given area. Clearly, these are all reasons to doubt that parameters estimated from insurance data could be a sound basis for any sort of normative or welfare analysis.

V.   CONCLUSION

Robert Lucas (1987, 2003) has attempted to measure the costs of the business cycle by calculating the utility of an average consumer in two different situations. In the first, the consumer receives a sequence of consumption bundles equal in size to actual average consumption in the United States. In the second, the consumer receives a deterministic stream, set equal to expected consumption in each period. The costs of the business cycle are then the proportion by which the first consumption stream would have to be increased to make the consumer just as well off as with the second, deterministic stream. This paper has applied Amartya Sen's (1997) notion of chooser-dependence to this problem, arguing that the *source* of the randomness in the consumption stream probably affects how consumers and policymakers feel about it. Aside from the theoretical observation that true costs could be affected by the identity of the chooser, I argued that welfare calculations of the type used by Lucas could be distorted by faulty calibration. This could occur if parameters were calculated in a market in which Person j was the chooser, and then welfare was "measured" using that estimated parameter in a situation when a different Person, j', was the chooser or subject of the utility function.

Lucas's estimate represents an approach to policy analysis marked by the use of the utility functions of representative agents. This approach is a relatively new one, which has accompanied the revolution over the past 35 years in the way macroeconomics is done. The issues raised in this paper may arise in other areas in the new welfare economics of macroeconomics: the measurement of the welfare costs of inflation, the welfare effects of monetary shocks, and so on. Problems with chooser-dependence are closely related to a more general set of pitfalls in welfare economics that arise when choices depend on "ancillary conditions." If these difficulties prove insurmountable, it may be that simpler criteria for policy analysis, in terms of outcomes such as unemployment or capabilities, are no less reliable than more modern techniques, and are better able to provide practical guidance.

REFERENCES

Alvarez, F. and Jermann, U. 2004. "Using asset prices to measure the cost of business cycles." Journal of Political Economy 112(6): 1224–1256.

Anderson, E. 1993. *Value and Ethics in Economics*. Cambridge, MA: Harvard University Press.

Barsky, R., Juster, F.T., Kimball, M., and Shapiro, M. 1997. "Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement study." Quarterly Journal of Economics 112(May): 537–579.

Bernheim, D. 2007. "Untitled presentation." Allied Social Sciences Association Annual Meetings, Chicago, January 5.

Camerer, C. 2007. "Untitled presentation." Allied Social Sciences Association Annual Meetings, Chicago, January 5.

Frederick, S., Loewenstein, G., and O'Donoghue, T. 2002. "Time discounting and time preference: a critical review." Journal of Economic Literature 40(June): 351–401.

Friedman, B. 1974. "Risk aversion and the consumer choice of health insurance option." Review of Economics and Statistics 56(2): 209–214.

Friend, I. and Blume, M. 1975. "The demand for risky assets." American Economic Review 65(5): 900–922.

Goldsmith, A. and Veum, J. 1996. "The psychological impact of unemployment and joblessness." Journal of Socio-Economics 25(3): 333–359.

Gul, F. and Pesendorfer, W. 2007. "Untitled Presentation." Allied Social Sciences Association Annual Meetings, Chicago, January 5.

Hansen, L. P. and Singleton, K. 1982. "Generalized instrumental variables estimation of nonlinear rational expectations models." Econometrica 50(5): 1269–1286.

Hartley, R., Lanot, G., and Walker, I. 2006. "Who really wants to be a millionaire? Estimates of risk aversion from gameshow data." Revised version. Warwick Economic Research Papers No. 747, University of Warwick.

Kandel, S. and Stambaugh, R. 1991. "Asset returns and intertemporal preferences." Journal of Monetary Economics 27(1): 39–71.

Karni, E. 1989. "Generalized expected utility analysis of multivariate risk aversion." International Economic Review 30(2): 297–305.

Keynes, J. 1936. *The General Theory of Employment, Interest, and Money*. New York: Harcourt, Brace.

Krantz, D. and Kunreuther, H. 2006. "Goals and plans in protective decision making." NBER Working Paper No. 12446.

Krusell, P. and Smith, A. 1999. "On the welfare effects of eliminating business cycles." Review of Economic Dynamics 2(1): 245–272.

Loewenstein, G., Hsee, C., Weber, E., and Welch, N. 2001. "Risk as feelings." Psychological Bulletin 127(2): 267–286.

Lucas, R. 1987. "Models of Business Cycles." Oxford: Basil Blackwell.

————. 2003. "Macroeconomic Priorities." American Economic Review 93(1): 1–14.

Machina, M. 1987. "Choice under uncertainty: problems solved and unsolved." Economic Perspectives 1(1): 121–154.

Mukoyama, T. and Şahin, A. 2006. "Costs of business cycles for unskilled workers." Journal of Monetary Economics 53(8): 2179–2193.

Obstfeld, M. 1992. "Evaluating risky consumption paths: the role of intertemporal substitution." NBER Technical Working Paper No. 120.

Obstfeld, M. and Rogoff, K. 1996. "Foundations of International Macroeconomics." Cambridge, MA: MIT Press.

Rabin, M. 2000. "Risk aversion and expected-utility theory: a calibration theorem." Econometrica 68(5): 1281–1292.

Rabin, M. and Koszegi, B. 2007. "Untitled Presentation." Allied Social Sciences Association Annual Meetings, Chicago, January 5.

Rabin, M. and Thaler, R. 2001. "Anomalies: risk aversion." Journal of Economic Perspectives 15(1): 219–232.

Safra, Z. and Segal, U. 1993. "Dominance axioms and multivariate non–expected utility preferences." International Economic Review 34(2): 321–334.

Sargent, T. 2004. *Recursive Macroeconomic Theory*. Second Edition. Cambridge, MA: MIT Press.

Seligman, M. and Maier, S. 1967. "Failure to escape traumatic shock." Journal of Experimental Psychology 74(1): 1–9.

Sen, A., 1973. "Behaviour and the concept of preference." Economica 40(159): 241–259.

————. 1997. "Maximization and the Act of Choice." Econometrica 65(4): 745–779.

Stokey, N. and Lucas, R. 1987. *Recursive Methods in Economic Dynamics.* Cambridge, MA: Harvard University Press.

Storesletten, K., Telmer, C., and Yaron, A. 2001. "The welfare cost of business cycles revisited: finite lives and cyclical variation in idiosyncratic risk." European Economic Review 45(7): 1311–1339.

Szpiro, G. 1986. "Measuring risk aversion: an alternative approach." Review of Economics and Statistics 68(1): 156–159.

Weil, P. 1990. "Non-expected utility in macroeconomics." Quarterly Journal of Economics 105(1): 29–42.

Wolf, C. and Pohlman, L. 1983. "The recovery of risk preferences from actual choices." Econometrica 51(3): 843–850.

Wolfers, J. 2003. "Is business cycle volatility costly? Evidence from surveys of subjective well-being." International Finance 6(1): 1–26.

Yellen, J. and Akerlof, G. 2004. "Stabilization policy: a reconsideration." Federal Reserve Bank of San Francisco and University of California, Berkeley.