RESEARCH PAPER NO. 1751

Is Business Cycle Volatility Costly? Evidence from Surveys of Subjective Wellbeing

Justin Wolfers

July 2002

RESEARCH PAPER SERIES

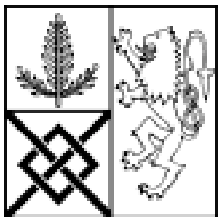

GRADUATE SCHOOL OF BUSINESS

STANFORD UNIVERSITY

Is Business Cycle Volatility Costly?
Evidence from Surveys of Subjective Wellbeing

Justin Wolfers
Stanford GSB
jwolfers@stanford.edu
www.stanford.edu/people/jwolfers

Abstract
This paper analyzes the effects of business cycle volatility on measures of subjective wellbeing, including self-reported happiness and life satisfaction. I find robust evidence that inflation, and particularly unemployment, lower perceived wellbeing. Conditional on levels of unemployment and inflation, greater macroeconomic volatility undermines wellbeing. These effects are moderate but important: eliminating unemployment volatility would raise wellbeing by an amount roughly equal to that from lowering unemployment by a quarter of a percentage point. The effects of inflation volatility on wellbeing are less easy to detect and are likely smaller.

Paper prepared for the Council on Foreign Relations / International Finance Conference on "Stabilization Policy: Why and How?", New York July 11, 2002

First Draft: June 27, 2002
This Draft: July 9, 2002

I would like to thank James Dailey, Jonathan Gardner, John McMillan, Andrew Oswald, John Simon and Betsey Stevenson for useful discussions. Ravi Pillai provided research assistance. Naturally all remaining errors are my own.

1. Introduction: The Happiness Cost of Business Cycles

The focus of macroeconomists on business cycle analysis presumably reflects a belief that macroeconomic fluctuations entail important consequences for the well-being of society. However formal attempts to quantify these costs suggest that they may be small. Lucas (1987) argues that business cycles present mainly a source of aggregate consumption risk that turns out to be of only second-order importance. He estimates the cost of this risk to be equivalent to less than one-tenth of one percent of consumption.

Yet both the general public and many economists continue to argue that smoothing the business cycle is an important objective. For instance, Shiller (1997, p.36) reports that 80% of Americans agree with the statement that preventing recessions is "as important as preventing drug abuse or preventing deterioration in the quality of our schools". Of these, 83% endorse the importance of counter-cyclical policy even if "the method of preventing economic recessions had an absolutely equal impact on economic booms… preventing really good times just as much as it prevented really bad times." Professional economists surveyed by Shiller were as likely to agree with these two statements as the public. (The corresponding proportions are 76% and 84%.)

This paper treads a path between the approaches of Shiller and Lucas. That is, I do not simply ask people whether they believe business cycles are costly. Equally, I do not assume a preference structure and make inferences about the cost of business cycles. Rather, I analyze data on subjective wellbeing, and compare the evolution of aggregate levels of happiness with macroeconomic conditions.

Each of these approaches share a common theme: they explore the extent to which macroeconomic conditions interact with preferences—standard preferences, expressed preferences or experienced utility—to induce an aversion to business cycle volatility.

In turn, this aversion to volatility can arise from several sources, although they all share a dependence on Jensen's inequality: convexity in either preferences or economic structure drives a wedge between average wellbeing in a volatile macroeconomy and wellbeing experienced under average macroeconomic conditions. As such, volatility undermines wellbeing.

The macroeconomic data that I analyze is familiar to most readers, while the use of happiness data in macroeconomics is somewhat novel. Thus, I start by introducing these data, and outlining the existing literature in section two. Section three establishes that both unemployment and inflation lower wellbeing. Section four further explores the effect of the business cycle on feelings in two national datasets. Section five is the heart of the analysis,

asking whether the data can speak to not only a preference for low unemployment and inflation, but also low volatility. Section six concludes.

To preview my findings, I find compelling evidence that unemployment and inflation both undermine wellbeing. People are much more averse to unemployment than inflation, and even with a moderate sacrifice ratio, it is a close call as to whether disinflationary policies have raised or lowered wellbeing. I find evidence that volatility in the real economy has important consequences. Current levels of aggregate unemployment volatility probably reduce wellbeing by an amount equivalent to raising average unemployment by a quarter of a percentage point. While I find some suggestive evidence that inflation volatility is costly, the evidence on this point is less clear.

2. Background: Subjective Well-being in Macroeconomics

Analysis of subjective well-being remains somewhat on the periphery of modern economics, presumably reflecting the deep grounding of our discipline in revealed preference. Nonetheless, economic analysis of feelings may still be useful for at least three reasons. First, feelings may enter directly into the utility function, and as such are an object of direct interest. Second, feelings may help predict behavior. And third, it may be that asking people how they feel is a more direct way to do welfare economics than to make inferences based on their consumption patterns.

Data on feelings of wellbeing are collected from several large-scale surveys that ask questions including:

- "Taking all things together, how would you say things are these days—would you say you're very happy, fairly happy, or not too happy these days?" [*Eurobarometer 1975-1986*]
- "On the whole, are you very satisfied, fairly satisfied, not very satisfied or not at all satisfied with the life you lead?" [*Eurobarometer 1973-1998*]
- "Taken all together, how would you say things are these days would you say that you are very happy, pretty happy, or not too happy?" [*General Social Survey: 1972-1998*]
- "Have you recently been feeling reasonably happy, all things considered?" [*British Household Panel Study: 1991-2000*]

While life satisfaction and happiness are somewhat different concepts, responses are highly correlated and hence these concepts are typically joined under the broader rubric of "subjective wellbeing". A host of validation studies have suggested that these questions reveal

something meaningful and can indeed be interpreted as reflecting levels of wellbeing. For instance, those who report being happier are typically rated by others as happier, and tend to smile and laugh more. Self-reported happiness also correlates with both heart rate and electrical activity in the brain. Measures of subjective wellbeing have relatively high test-retest correlations, similar microeconometric structure across time and space, and are highly correlated with related questions. Diener (1984) provides a useful review of psychometric analysis of the validity of these data.

In his seminal 1974 paper, Richard Easterlin, asked whether economic growth led to rising happiness.[1] Analyzing time series data in both the US and Japan (where average incomes grew by a factor of two and six, respectively, since the war), he finds that raising the income of all does not raise the happiness of all. Updates to this work in 1995 and 1996 confirm these findings. This is to be contrasted with the finding that within a society, those who are richer tend to report being happier. Easterlin reconciles these two results by reference to Duesenberry's (1949) relative income hypothesis: what matters is keeping up with the Joneses, and economic growth only makes you rich at the same rates as it does Jones. This leads to not only a different concept of welfare (Boskin and Sheshinski 1978), but also presents a challenge to the efficacy of standard macroeconomic policy.

However, the largest cross-national comparisons of happiness data—by Inglehart and Klingemann (2000) and Veenhoven (1997)—find that people in poor countries have much lower levels of self-reported happiness. These authors argue that wellbeing is highly responsive to the satisfaction of basic needs, but almost invariant to income at higher levels of development.

Finally, economic and political institutions are also related to happiness, although it is an open question as to whether these effects are independent of macroeconomic conditions, or mediated by them. Nonetheless, under this view, establishing the right institutional frameworks is the key to higher levels of happiness.

This paper takes a closer look at links between happiness and business-cycle variation in unemployment and inflation.

3. Business Cycles and Happiness

The most directly relevant analysis of happiness and the business cycle comes from di Tella, MacCulloch and Oswald (2001). Analyzing a country-year panel, they find that life

[1] Bruno and Frey (2002) provide a useful summary of recent happiness research in economics.

satisfaction declines with unemployment and inflation, controlling for country and year fixed effects. Their work was based on an unbalanced panel of Eurobarometer survey data covering twelve European nations from 1975-1991. My analysis uses the same survey, but updates it to cover sixteen countries running from 1973-98 (Schmitt and Scholz, 2002). Thus, I start by updating these results in Table 1.

The happiness question (shown on page 2) was asked only from 1975-86 (excluding 1980 and 1981), while the life satisfaction question was asked every year from 1973-98, except 1974 and 1996. Overall, this yields an unbalanced panel of 501,552 valid responses to the life satisfaction question in 274 country-years, and 134,590 responses to the happiness question over 99 country-years.[2] Answers to the two questions are highly correlated, and so given the greater data availability, my analysis will focus on life satisfaction.

An important technical issue relates to finding a sensible set of units with which to measure wellbeing. As only qualitative indicators exist, there are no natural happiness units that immediately correspond with the textbook "util". Throughout this paper I analyze five alternatives:

- Following di Tella et. al, I simply take country-year averages of wellbeing, assigning a value of 1 to the lowest satisfaction category, and with each qualitative increment, add one.[3]
- A less arbitrary alternative, is simply to code up the proportion of the population who are "very satisfied" in each country in each year. This has the disadvantage of effectively ignoring information from movements between "very satisfied" and "fairly satisfied", and also between "not very satisfied" and "not at all satisfied".
- A preferable alternative involves a two-stage procedure. In the first stage, I run an ordered probit regression on a full set of dummy variables for each country in each year. Under the assumption that wellbeing is an unobserved normally distributed variable within each country-year, this maximum likelihood procedure estimates the cut-points between different categorizations, thereby giving these data a more natural cardinalization.[4]

[2] Northern Ireland is included as part of the United Kingdom. For consistency, I drop East Germany from the sample, and follow only West Germans.

[3] Thus, for the Eurobarometer survey: 1 = "not at all satisfied"; 2 = "not very satisfied"; 3= "fairly satisfied"; 4 = "very satisfied".

[4] Thus, this cardinalization derives explicitly from an identifying assumption that wellbeing is normally distributed in the population. The estimates suggest that the difference between the median person who is "not at all satisfied" with their life and the median "not very" satisfied person is about two-thirds of the difference between the median in each of the "not very" and "somewhat" categories. In turn, this is about the same as the difference between the median "somewhat satisfied" and the median "very satisfied" person. Figure 2 shows this graphically.

The estimated country-year fixed effects capture shifts in the mean of the underlying distribution over time in each country. Figure 1 shows how the mean of this estimated satisfaction index varies across countries and through time. The chart also shows the proportion of the population who are very satisfied, and the two lines show close co-movement.

- An analogous but computationally more intensive procedure involves simply estimating an ordered probit regression directly on the micro data (including as controls a full set of state and year fixed effects). Standard errors are clustered at the country-year level, and hence should be comparable to those estimated on the macro data. These results are typically very similar to those derived using the two-step procedure.
- An advantage of working directly with the micro data is that a useful set of individual controls can also be added to the regression. Thus the fifth specification includes not only country and year fixed effects, but also sub-national (NUTS-1 level) region fixed-effects, and a saturated set of dummy variables for age and education, both interacted with gender.

Figure 1

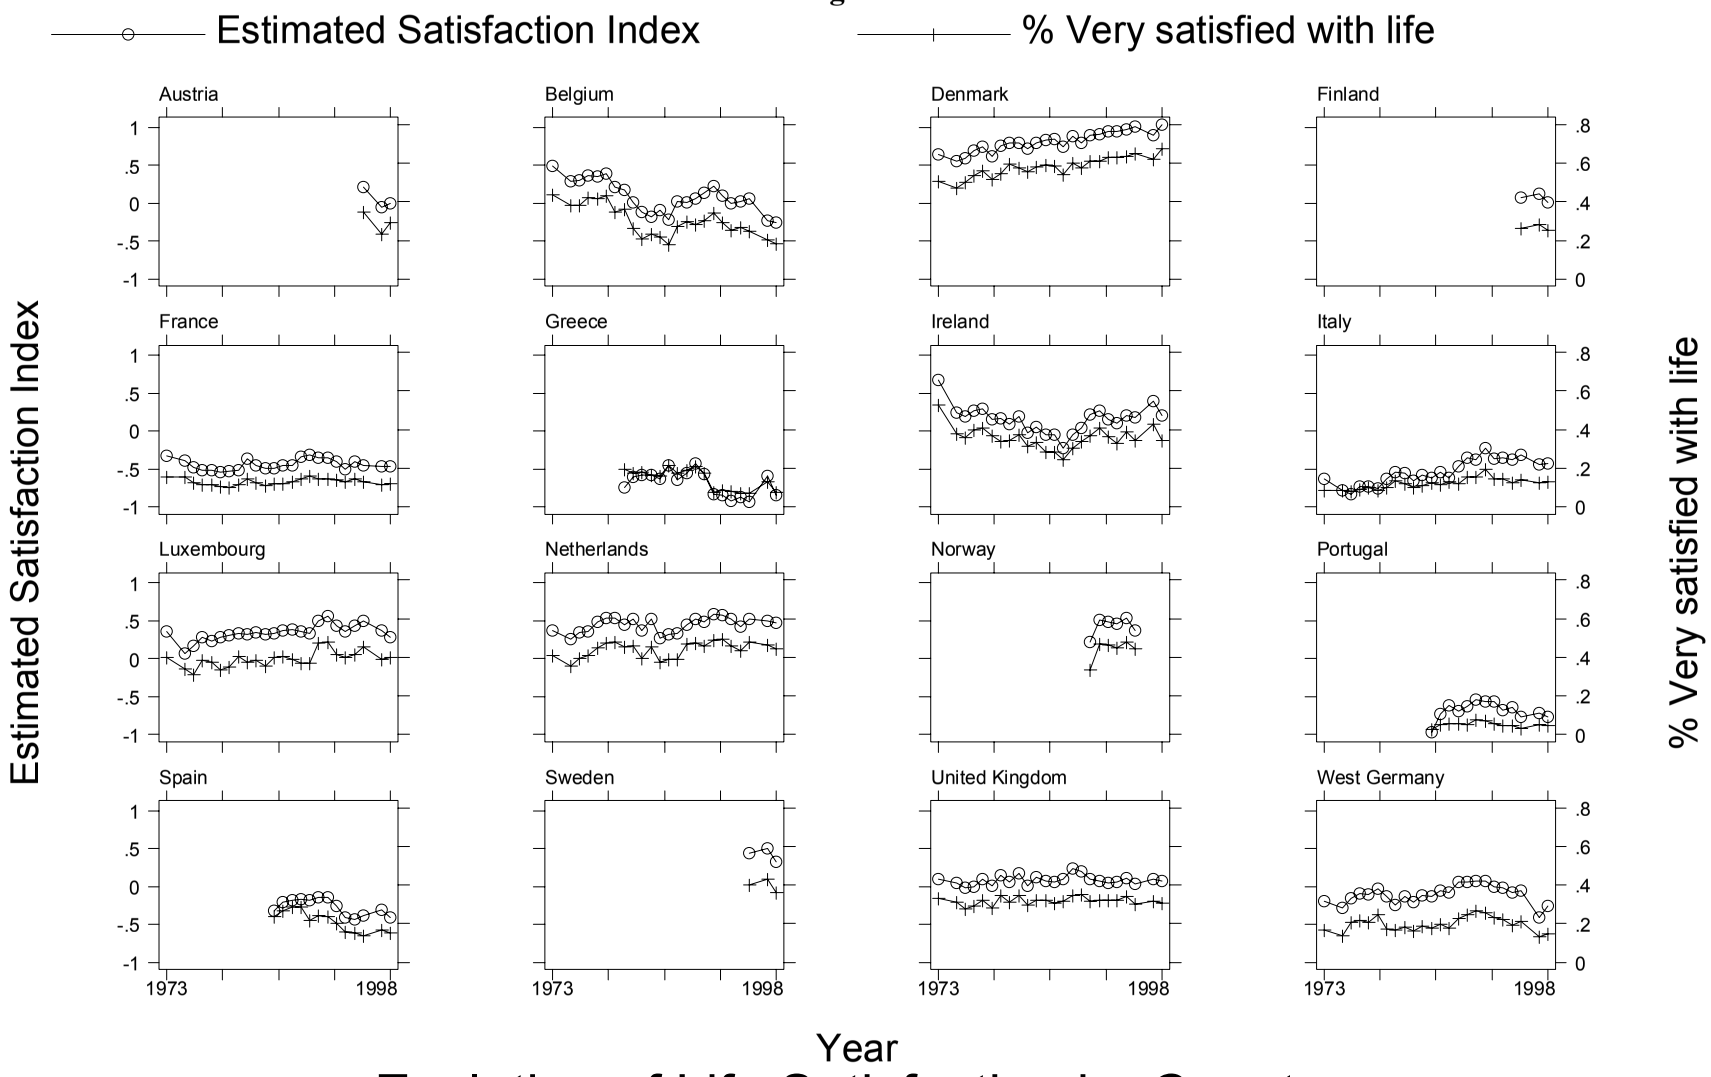

Evolution of Life Satisfaction by Country

Table 1 analyzes the relationship between measures of subjective wellbeing and annual measures of inflation and unemployment.[5] In each case, the regressions are of the form:

$$Life\ Satisfaction_{c,t} = \beta\, Unemployment_{c,t} + \pi\ Inflation_{c,t} + country\ and\ year\ fixed\ effects + \varepsilon_{c,t}$$

Table 1: Wellbeing and Macroeconomic Conditions: Eurobarometer Survey

	Panel A: Dependent Variable: Life Satisfaction[(a)]				
	Macro Data			**Micro Data**	
	Average Satisfaction	**% Very Satisfied with Life**	**Satisfaction Index ~N(0,1)**	**Ordered probit**	**Ordered probit w/ controls**
Unemployment$_c$	-2.305***	-1.096%***	-3.566***	-3.459***	
	(.348)	(.197)	(.552)	(.324)	
Inflation$_{c,t}$	-0.626***	-.157	-.776**	-.729**	
	(.205)	(.116)	(.326)	(.325)	
Tradeoff: β/π	3.7	7.0	4.6	4.7	
(90% interval)[(b)]	(2.2-8.0)	(-18.3-37.9)	(2.4-14.5)	(4.1-5.8)	
Controls					
State and year fixed effects	✓	✓	✓	✓	✓
Individual controls	×	×	×	×	✓
Adj. $R^{2(c)}$	.9326	.9237	.9341	.0645	
Sample	274	274	274	*n=501,552 274 country-year clusters*	

	Panel B: Effect on % "Very Satisfied" (Sample mean=28%)			
Unemployment$_{c,t}$	n.a.	-1.1%	-1.2%	-1.2%
Inflation$_{c,t}$		-0.2%	-0.3%	-0.2%

Notes

***, **, * denote significantly different from zero at 1%, 5% and 10% levels, respectively.

(Robust standard errors in parentheses, clustered at country-year level)

Each column reports a different regression with a different dependent variable:

(1) Average levels of happiness in each country-year (1= "not at all satisfied"; 2= "not very satisfied"; 3= "fairly satisfied"; 4 = "very satisfied").

(2) Proportion of the population in a country-year who reported being "very satisfied"

(3) The dependent variable is the value of an estimated happiness index derived from an ordered probit regression of satisfaction on a full set of interacted country and year fixed effects.

(4) Ordered probit on reported satisfaction.

(5) Ordered probit on reported satisfaction. Also includes as controls, sub-national region fixed effects, and a saturated set of dummies for age, education, and their interaction with gender.

(a) Microdata regressions use sample weights; macrodata weighted to reflect underlying microdata.

(b) 90% confidence interval on the unemployment-inflation tradeoff reflects the 10th and 90th percentile estimates from 1000 bootstrap replications.

(c) Adjusted R^2 in columns 1-3; Pseudo R^2 in columns 4-5.

[5] Unemployment and inflation data are from the OECD economic outlook. While I simply report the coefficient on that year's unemployment or inflation rate, di Tella et al use instead a three-year backward-looking average.

Panel A reports the main regression results. Each column shows a separate regression, and across all five columns, unemployment lowers subjective wellbeing. This result is robust across specifications, and easily meets standard levels of statistical significance. The effect of inflation on wellbeing is less statistically robust, but nonetheless, these data suggest that inflation lowers wellbeing. Further regressions (not shown) confirm that responses to the happiness question yield a similar pattern of responses.

Given that the dependent variable is qualitative in nature, one must take care in interpreting these magnitudes. Columns 3-5 are interpretable roughly as "effect sizes": the effects of macroeconomic variables on happiness are measured by scaling by the standard deviation of the happiness distribution. Thus, if unemployment were to rise by 10 percentage points, average levels of happiness decline by about 0.35 standard deviations. Figure 2 shows this graphically.

Figure 2

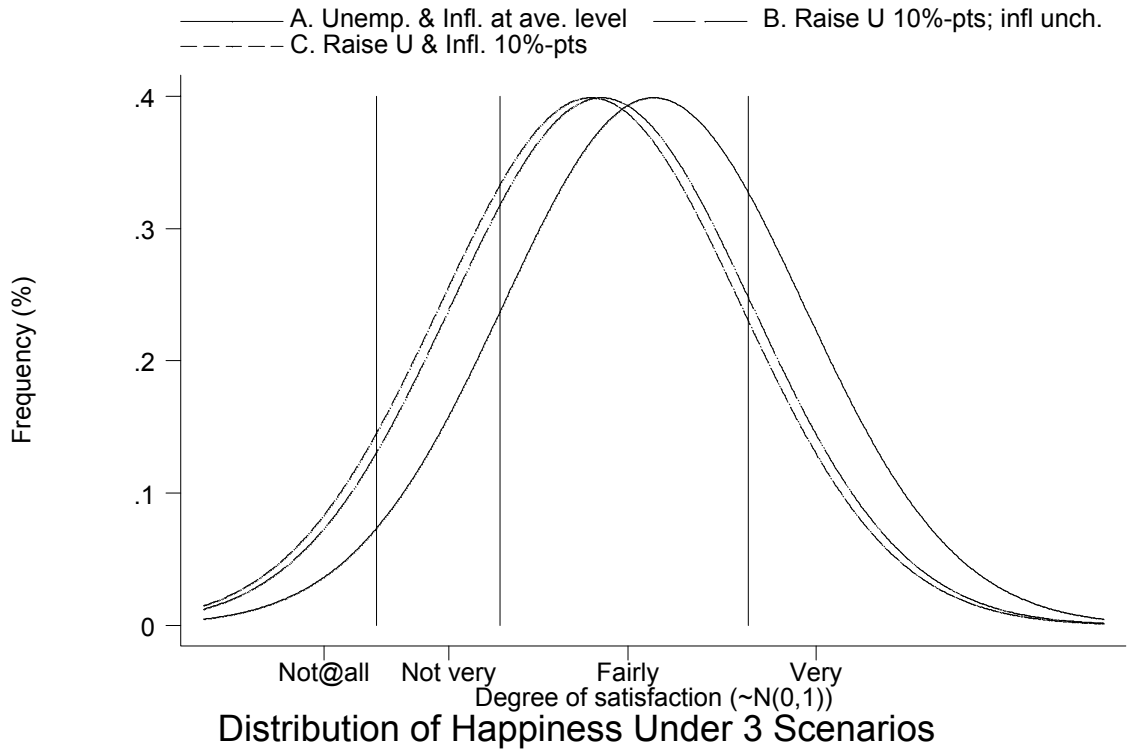

Source: Authors calculations based on results in column 4 of table 1.

The bell curve A shows the (assumed) standard normal distribution of happiness in the population, and the estimated cut-points show which respondents classify themselves as "very", "fairly", "not very" or "not at all" satisfied with their life. Unemployment shifts this distribution leftward, and curve B shows the distribution when unemployment is 10 percentage points higher.

Curve C shows this distribution if both unemployment and inflation are 10 percentage points higher. These curves are shifted left 0.35 and a further 0.07 standard deviations, respectively. To see that these are large effects, note that the median person in scenario B is as happy as the person at the 36th percentile in scenario A. Similarly, the median person in scenario C is as happy as the person at the 33rd percentile in the happiness distribution in scenario A.

While Figure 2 provides a useful metric for interpreting the results specifically in column 4, the scaling is also roughly similar for the results in columns 3 and 5.[6] Panel B reports an alternative metric that is also comparable with results in the second column: changes in the predicted share of the population who are "very satisfied". (These are marginal effects, evaluated at the mean.) On average, 28% of the population are "very satisfied" with their lives. Raising inflation by ten percentage points lowers this number by around two percentage points, while raising unemployment by ten percentage points, lowers the proportion very satisfied by around twelve percentage points.

The ratio of the coefficients on unemployment and inflation—shown in the third row—gives some sense of how the public feels about the unemployment-inflation trade-off. Whereas di Tella, Oswald and MacCulloch had found the public to be indifferent between raising unemployment for a year by one percentage point and raising inflation by 1.7 percentage points, this larger sample suggests that the inflation-unemployment trade-off is closer to five to one. That is, the public appears to be extremely averse to unemployment.

These results can be used to speculate about the effects of disinflation on wellbeing. Feldstein (1997, p.123) has claimed a widespread professional consensus that inflation has important adverse effects and that "these adverse effects justify the sacrifices in employment and output that are generally needed to reduce inflation." These wellbeing data give an alternative way to evaluate these costs and benefits.

Consider a central bank trying to decide whether to reduce inflation by one percentage point. The costs of such a policy derive from estimates of the so-called sacrifice ratio, often derived from estimates of the slope of the short-run Phillips Curve. These estimates suggest that the cost of such a policy is that unemployment must be kept below its equilibrium level by two percentage points for a year. The benefits of this temporary contraction are assumed to be

[6] This is not precisely true. With an ordered probit, the distribution of happiness *conditional on covariates* is assumed to follow a standard normal distribution. Columns 4 and 5 successively add covariates, soaking up residual variation. Consequently one might expect estimated coefficients to be slightly larger as covariates are added. The fact that these individual-level controls do not much raise R^2 leads to the conclusion that this effect is small in the present case.

permanent, and hence assuming a discount rate of 6%, permanently reducing inflation by 1 percentage point has a present value equivalent to 16-percentage point-years of lower inflation. Would the public be willing to trade-off 2 percent-years of higher unemployment for 16 percent-years of lower inflation?

The preferences over inflation and unemployment described in Table 1 suggest that this is a close call. If the effect on unemployment is temporary and the effect on inflation permanent, then disinflation probably enhances wellbeing. However, if the rise in unemployment following a monetary contraction is not completely transitory (as suggested by Ball 1997), or the decline in inflation is not completely persistent, then disinflation probably lowers wellbeing. Indeed, the results in Table 1, when interpreted jointly with evidence of heterogeneous and persistent costs of disinflation in Ball (1994, 1997) suggest that, for some countries, disinflation almost certainly undermines well-being.

4. Further Evidence on Wellbeing and Unemployment

Several other data sources shed additional light on the wellbeing costs of unemployment. In particular, the British Household Panel Survey allows one to draw a richer portrait. This survey—similar in structure to the PSID—started tracking around 10,000 Britons in 1991. Of particular note, respondents are asked to submit to the psychological tests each year. These tests, called the GHQ-12 (General Health Questionnaire), comprise a battery of twelve questions originally developed as a screening instrument for psychiatric illness, but increasingly used as an indicator of psycho-social wellbeing.

Each of the GHQ-12 questions is shown in Table 2. The first column also reports the proportion that states their wellbeing on that measure to be less than usual.[7] Typically, about 10-20% of the population exhibits each "symptom" on any specific question.

These data were matched to regional unemployment data from both the labor force survey, and claimant counts. Unfortunately consistent regional inflation indicators are not available. In order to test for the influence of the business cycle on wellbeing, I ran ordered probit regressions of wellbeing on the regional unemployment rate, including region and year fixed effects. The second column of Table 2 reports these results where the independent variable is the claimant count measure of unemployment; column 3 analyzes the survey measure. The

[7] Responses were recorded so that higher numbers denoting greater wellbeing. For questions regarding capabilities (A, C, D, G, H, L): 1 = much less; 2= less so; 3 = same as usual; 4 = more than usual. For questions regarding limitations (B, E, F, I, K): 1 = much more; 2 = rather more; 3 = no more than usual; 4 = not at all.

fourth and fifth columns report these estimates as marginal effects that indicate the effect of higher unemployment on the share of people reporting that they are unwell on a particular measure.

Table 2: Effects of Unemployment on Various Dimensions of Wellbeing: BHPS

Average	*Estimated Coefficients*			
% Reporting unwell[(a)]	**Reported as Effect Sizes: *Change in wellbeing [~N(0,1)]***		**Reported as Marginal Effects: *Change in % reporting unwell[(a)]***	
	Claimant Count	**LFS**	**Claimant Count**	**LFS**
A. Have you recently been able to concentrate on whatever you're doing?				
19.4%	-1.360	-.412	0.4%	0.1%
	(1.390)	(.913)		
B. Have you recently lost much sleep over worry (18.2%)				
18.2%	-1.673*	-.452	0.4%	0.2%
	(.888)	(.622)		
C. Have you recently felt that you were playing a useful part in things?				
13.4%	-2.554**	-2.563***	0.6%	0.6%
	(1.14)	(.890)		
D. Have you recent felt capable of making decisions about things?				
9.3%	0.364	0.263	-0.1%	-0.1%
	(1.119)	(.889)		
E. Have you recently felt constantly under strain?				
28.0%	-1.416*	-1.252*	0.5%	0.4%
	(.839)	(.697)		
F. Have you recently felt you couldn't overcome your difficulties?				
13.9%	-1.219	-.366	0.3%	0.1%
	(.782)	(.630)		
G. Have you recently been able to enjoy your normal day-to-day activities?				
19.1%	.002	.417	0.0%	-0.1%
	(1.208)	(.887)		
H. Have you recently been able to face up to problems?				
10.9%	-1.915	-.579	0.4%	0.1%
	(1.361)	(.939)		
I. Have you recently been feeling unhappy or depressed?				
21.8%	-1.768**	-1.805***	0.5%	0.7%
	(.850)	(.619)		
J. Have you recently been losing confidence in yourself?				
13.1%	-3.662***	-2.153***	1.5%	0.9%
	(.821)	(.681)		
K. Have you recently been thinking of yourself as a worthless person?				
7.1%	-2.204***	-.871	0.8%	0.3%
	(.722)	(.604)		
L. Have you recently been feeling reasonably happy, all things considered?				
13.0%	-1.762	-1.945**	0.4%	0.4%
	(1.130)	(.948)		

Sample includes the Scotland and Wales extension samples, but excludes the ECHP sub-sample, and hence includes no data for Northern Ireland. Valid responses required all twelve questions to be answered, yielding 88,695 valid interviews from 14,752 respondents over ten years, in 120 region-year clusters.
Weighted using sample weight xrwght.
(a) "Unwell" equates to reporting wellbeing less than usual or much less than usual on a particular measure.

It is worth emphasizing that these results reflect aggregate movements. Thus, the interpretation is that a rise in unemployment in a region lowers average feelings of usefulness, confidence and happiness, and raises depression and feelings of worthlessness in that region.

Finally, Table 4 turns to data from the US General Social Survey to shows that economic downturns also undermines happiness across US states. I regress happiness on two alternative measures of state labor market conditions: the CPS unemployment rate, and a measure constructed as the deviation of (log) non-farm payrolls from a Hodrick-Prescott filtered trend (described in Wolfers, 2002). As in the tables above, I show both the coefficients from an ordered probit, and the marginal effects. In this case, the marginal effects shown are for the percent "very happy".

Table 4: Happiness and Unemployment Across US States: General Social Survey

	Average	*Estimated Coefficients*			
	% very happy or very confident	Reported as Effect Sizes: *Change in happiness [~N(0,1)]*		Marginal Effects: *Change in % reporting very happy*[(a)]	
		Unemp. Rate	Emp. Gap	Unemp. Rate	Emp. Gap
"Taken all together, how would you say things are these days would you say that you are [3] very happy, [2] pretty happy, or [1] not too happy?"					
Happiness	32.5%	1.539***	0.632***	-0.3%	-0.1%
		(.519)	(.218)		
I am going to name some institutions in this country. As far as the people running these institutions are concerned, would you say you have [3] a great deal of confidence, [2] only some confidence, or [1] hardly any confidence at all in them?					
Executive branch - Federal govt.	17%	-0.962	-0.784***	-0.2%	-0.2%
		(.717)	(.296)		
Congress	14%	-0.698	-0.666**	-0.1%	-0.1%
		(.653)	(.295)		
United States Supreme Court	33%	-0.744	0.133	-0.3%	0.0%
		(.733)	(.303)		
Banks & financial institutions	27%	-3.084***	-1.276***	-1.0%	-0.4%
		(.682)	(.305)		
Major companies	26%	-1.491**	-0.774***	-0.5%	-0.3%
		(.679)	(.296)		
Organized labor	12%	0.379	0.115	0.1%	0.0%
		(.626)	(.313)		
Organized religion	30%	-0.601	-0.386	-0.2%	-0.1%
		(.721)	(.316)		
Education	32%	-1.445**	-0.755**	-0.5%	-0.3%
		(.731)	(.317)		
Press	18%	-0.352	0.037	-0.1%	0.0%
		(.609)	(.286)		
Medicine	49%	0.244	-0.205	0.1%	-0.1%
		(.695)	(.301)		
Television	15%	0.601	0.312	0.1%	0.1%
		(.687)	(.315)		
Scientific Community	44%	0.103	0.034	0.0%	0.0%
		(.600)	(.291)		
Military	36%	-0.531	0.129	-0.2%	0.0%
		(.654)	(.299)		

Sample: GSS data for 1973-1998. *n*=36,146 valid interviews, in 773 state-year clusters for happiness; *n*=22,786 valid responses to all confidence questions in 667 state-year clusters.
Ordered probit regressions, controlling for state and year fixed effects
***, **, * denote significantly different from zero at 1%, 5% and 10% levels, respectively.
(Robust standard errors in parentheses, clustered at state-year level)

Beyond the effects on happiness, Table 4 also shows that economic downturns undermine public confidence in political and economic institutions. Intriguingly, the effects on faith in large companies, banks and financial institutions and even education are larger than those on political institutions such as the President and Congress.[8] (Note that the identification here reflects the performance of the state economy relative to the national trends picked up by the year fixed effects; consequently much of the blame may in fact lie

5. Are People Averse to Economic Fluctuations?

Results in the previous sections suggest that the both inflation and unemployment undermine wellbeing, that these effects are reflected in many different domains, and that this undermines public confidence in economic and political institutions. In this section I turn to examining the effects of macroeconomic volatility on wellbeing. The preceding results do not speak directly to this issue: it may be that the benefit of a boom exactly offsets the costs of a downturn, and hence the business cycle does not affect average levels of wellbeing. Costs of macroeconomic volatility arise either because of convexity in the structure of the economy, or convexity in preferences.[9]

Examples of convexities in economic structure include non-linearities in the short-run Phillips curve (Debelle and Laxton, 1997), investment irreversibilities (Bernanke 1983), short-run quasi-fixity in production (Ramey and Ramey 1991), and the choice of high-variance high-return projects (Black, 1987). Costs arising from convexity in the production side of the economy can be directly observed by the macroeconomist as reduced output or employment. [10]

Examples of papers focusing on convexity in preferences are discussed below. A common feature of these papers is that the welfare costs attributable to convex preferences are not immediately observable in macroeconomic data—one typically needs a well-specified welfare function in order to make any inferences.

[8] Presumably this lack of anger at Washington partly reflects the fact that these are responses to state unemployment relative to the national trend (as picked up by year fixed effects).
[9] Without this convexity, the gain to wellbeing in booms offsets the loss in a downturn.
[10] Ramey and Ramey (1995) carefully examine the cross-country evidence, and find robust evidence that macroeconomic volatility undermines growth. Specifically, a reduction in the volatility of output growth equal to one standard deviation of its cross-country variation translates into growth that is one-third of a percentage point higher in OECD countries.

Lucas (1987) analyzes the costs of volatility due to imperfect insurance, finding that aggregate consumption risk due to the business cycle undermines wellbeing by an amount roughly equivalent one-tenth of a percent of the consumption of the representative agent.

By analyzing a representative agent model, Lucas explicitly sidesteps the issue of idiosyncratic labor market risk. While not denying the important of idiosyncratic risk, he argues that this reflects "the potential or actual gain from social insurance, not from stabilization policy." Atkeson and Phelan (1994) address idiosyncratic risk directly. They argue that that the main source of idiosyncratic risk involves the risk of job loss. Figure 4 shows a highly stylized model in which recessions are periods of high rates of job loss, while booms yield offsetting low rates of job loss. (To complete the model, assume that rates of job-finding are constant through time.)

Figure 3

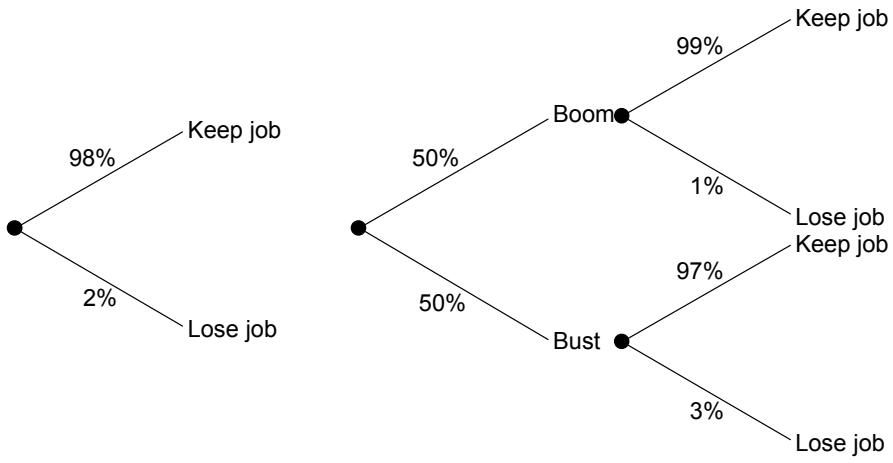

A. No Business Cycle
$E[U]=.98U_{Emp}+.02U_{Unemp}$

B. Volatile Business Cycle
$E[U]=.98U_{Emp}+.02U_{Unemp}$

The elimination of business cycles has no effect on average levels of wellbeing in this stylized example. This result reflects the fact that expected utility is linear in the probability of job loss.

However, labor market risk is more highly correlated across individuals in scenario B than scenario A. That is, the probability that both a worker and her husband lose their jobs at the

same time is higher when the firing rate is volatile. Consequently, the ability of husband and wife to insure each other against job loss declines with macroeconomic volatility, undermining average wellbeing. This same logic extends to other forms of intra-temporal pooling of idiosyncratic risk. (Formally, this arises because pooled income is not linear in the probability of job loss.)

Galí, Gertler and López-Salido(2002) propose a different set of costs of business cycle volatility, analyzing the implications of imperfect competition. These models typically involve equilibria in which employment is sub-optimally low. Booms drive the economy toward the first best, while recessions drive the economy even further away from the optimum. As shown in Figure 4, incrementally larger deviations from first best have an increasing welfare cost (the gap between the supply and demand curves widens), suggesting that recessions undermine wellbeing by more than equal-sized booms. These three authors calibrate a model to infer this cost, arguing that for the US, business cycle volatility yield welfare losses equivalent to about 0.8 percent of consumption.

Figure 4

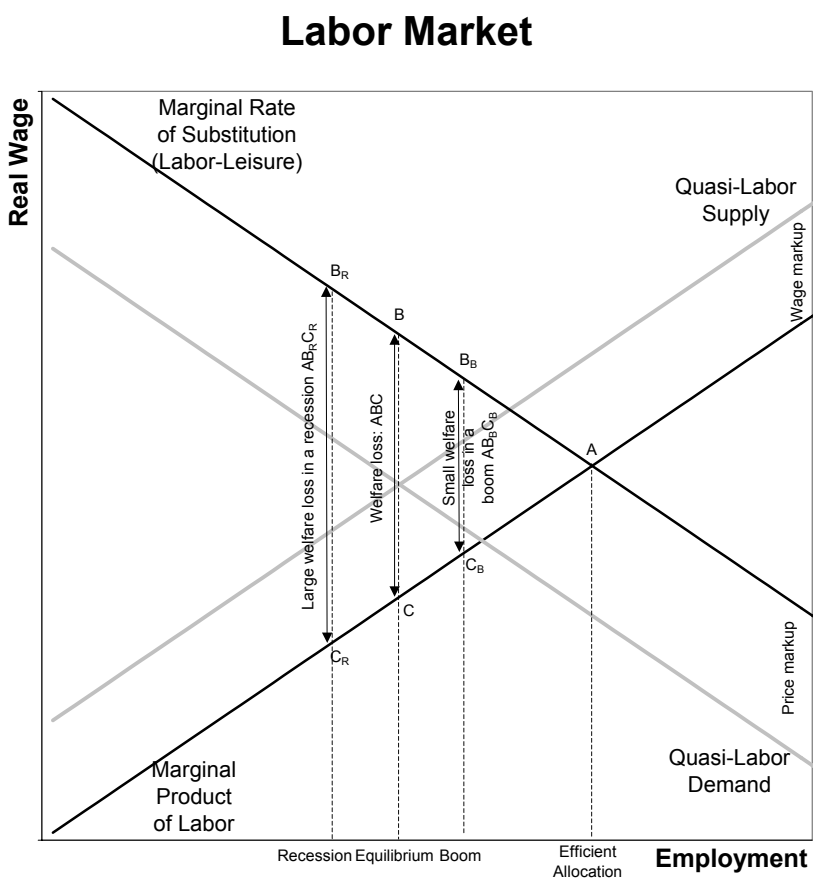

Finally, on the assumption that more volatile inflation is harder to forecast, aversion to inflation volatility can be motivated by reference to either the redistributive effects of unanticipated inflation, or the effects of relative-price distortions resulting from staggered price setting. (For more on this point, see Woodford, 2001)

While the common theme of these papers has been an attempt to infer effects on wellbeing from an imposed preference structure, my point of departure is to treat these wellbeing costs of business cycle volatility as directly observable. Thus, I now turn to examining whether volatility in unemployment and inflation undermines wellbeing, conditional on average levels of unemployment and inflation.[11] That is, I run the following regression:

$$\begin{aligned} Life\ Satisfaction_{c,t} = {} & \beta\ Average\ Unemployment_{c,t\text{-}7..t} + \lambda\ Standard\ Deviation\ Unemployment_{c,t\text{-}7..t} + \\ & \pi\ Average\ Inflation_{c,t\text{-}7..t} + \delta\ Standard\ Deviation\ Inflation_{c,t\text{-}7..t} \\ & + country\ \&\ year\ effects + \varepsilon_{c,t} \end{aligned}$$

[11] By conditioning on levels of unemployment and inflation, I partial out those effects of volatility that are mediated by higher volatility causing worse macroeconomic outcomes.

Table 5: Life Satisfaction: Levels and Volatility of Inflation and Unemployment

	Panel A: Dependent Variable: Life Satisfaction[(a)]				
	Macro Data			Micro Data	
	Average Satisfaction	% Very Satisfied with Life	Satisfaction Index ~N(0,1)	Ordered probit	Ordered probit with controls
Unemployment from *t-7* to *t*					
Mean	-3.685***	-1.858***	-6.017***	-6.040***	
	(.514)	(.276)	(.864)	(.780)	
Standard deviation	-.912	-1.041*	-2.612*	-2.419*	
	(1.048)	(.557)	(1.594)	(1.481)	
Inflation from *t-7* to *t*					
Mean	-.651	-.419***	-1.324***	-1.265***	
	(.348)	(.153)	(.508)	(.460)	
Standard deviation	-.884	-.294	-1.378*	-1.488**	
	(.559)	(.255)	(.810)	(.737)	
Controls					
State and year fixed effects	✓	✓	✓	✓	✓
Individual controls	×	×	×	×	✓
Adj. $R^{2(c)}$	.9473	.9451	.9506	.0605	
Sample	274	274	274	*n=501,552 274 country-year clusters*	
	Panel B: Effect on % "Very Satisfied" (Sample mean=28%)				
Mean unemp.	n.a.	-1.9%	-2.0%	-2.0%	
SD unemp.		-1.0%	-0.9%	-0.8%	
Mean inflation		-0.4%	-0.4%	-0.4%	
SD inflation		-0.3%	-0.5%	-0.5%	

Notes:
See notes to Table 1
Each column shows results from a separate regression with a different dependent variable.
For the 7 country-year observations with less than 8 years of inflation/unemployment data, standard deviation and mean of a smaller number of observations were used.

The first row confirms that unemployment dramatically lowers wellbeing, while the second row shows that volatile unemployment also lowers wellbeing. This coefficient is not particularly precisely estimated, but nonetheless, is statistically significant in four of five regressions.

The standard deviation of unemployment over an 8-year period averaged around 1½ percentage points in my sample. Thus, halving unemployment variability shifts the wellbeing distribution to the right by around .02 standard deviations, increasing the proportion of the population who are very satisfied by around 0.6 percentage points. By comparison, the gains from reducing unemployment by one percentage point are much larger—shifting the wellbeing distribution to the right by .06 standard deviations, and raising the percent very happy by 2 percentage points. Combining these findings suggests that halving unemployment volatility yields gains in wellbeing approximately equal to those from reducing the unemployment rate by one-third of a percentage point.

The inflation results are also interesting. As in Table 1, higher average levels of inflation lower wellbeing. Moreover, the second row suggests that there may be costs to inflation volatility, with three of the five estimates statistically significant. The standard deviation of inflation is around double that of unemployment, while the coefficient on inflation volatility is roughly half that on unemployment volatility. Take together, these point estimates suggest that the gains to halving either unemployment or inflation volatility are roughly similar.

Finally, it is worth noting that much of the literature on the evaluation of monetary policy rules starts from an objective function in which the weights on squared inflation and output deviations are equal. By contrast, my results suggest that—in terms of wellbeing—the weight on the standard deviation of unemployment is about twice that on the standard deviation of inflation. Taking account of the relative variances of inflation and unemployment, and an Okun's Law relationship in which the output gap moves two-for-one with unemployment, my estimates suggest that squared deviations of the output gap should receive about twice the weight of squared deviations of inflation.

Rather than make indirect inferences about the shape of the wellbeing function, Table 6 directly estimates the degree of convexity in preferences over inflation and unemployment. Thus, I repeat the analysis in Table 1, but also add quadratic terms in both inflation and unemployment. (A quadratic functional form is easily motivated in terms of the imperfect competition model with linear labor supply and labor demand curves, as shown in Figure 4.[12])

[12] Woodford (2001) provides detailed microfoundations for a welfare function that is quadratic in inflation and the output gap.

Table 6: Are Preferences Convex?

	Dependent Variable: Life Satisfaction				
	Macro Data			Micro Data	
	Average Satisfaction	% Very Satisfied with Life	Satisfaction Index ~N(0,1)	Ordered probit	Ordered probit with controls
Unemployment$_{c,}$	-.461	-.164	-.696	-.866	
	(.816)	(.459)	(1.301)	(1.184)	
Unemployment$_{c,t}$-squared	-9.488**	-4.865**	-14.297**	-13.496**	
	(3.766)	(2.120)	(5.970)	(5.381)	
Inflation$_{c,t}$	-1.058**	-.675**	-1.782**	-1.776***	
	(.467)	(.263)	(.715)	(.662)	
Inflation$_{c,t}$-squared	1.506	1.974**	4.114	4.075	
	(1.708)	(.961)	(2.873)	(2.651).	
Adj. R^2	.9341	.9262	.9384	.0646	
Sample	274	274	274	*n=501,552 274 country-year clusters*	
	Marginal Effect on % "Very Satisfied" (Sample mean=28%)				
Unemployment$_{c,t}$					
At 0%	n.a.	-0.2%	-0.2%	-0.3%	
At 7½%		-0.5%	-0.5%	-0.6%	
At 15%		-0.9%	-0.8%	-0.8%	
Inflation$_{c,t}$					
At 0%		-0.7%	-0.6%	-0.6%	
At 7½%		-0.5%	-0.5%	-0.5%	
At 15%		-0.4%	-0.4%	-0.4%	

See Notes to Table 1.

All five columns suggest that aggregate wellbeing is a nonlinear function of unemployment, and that this non-linearity is statistically significant; the results for inflation are less clear. Figure 5 shows these results graphically.

Figure 5

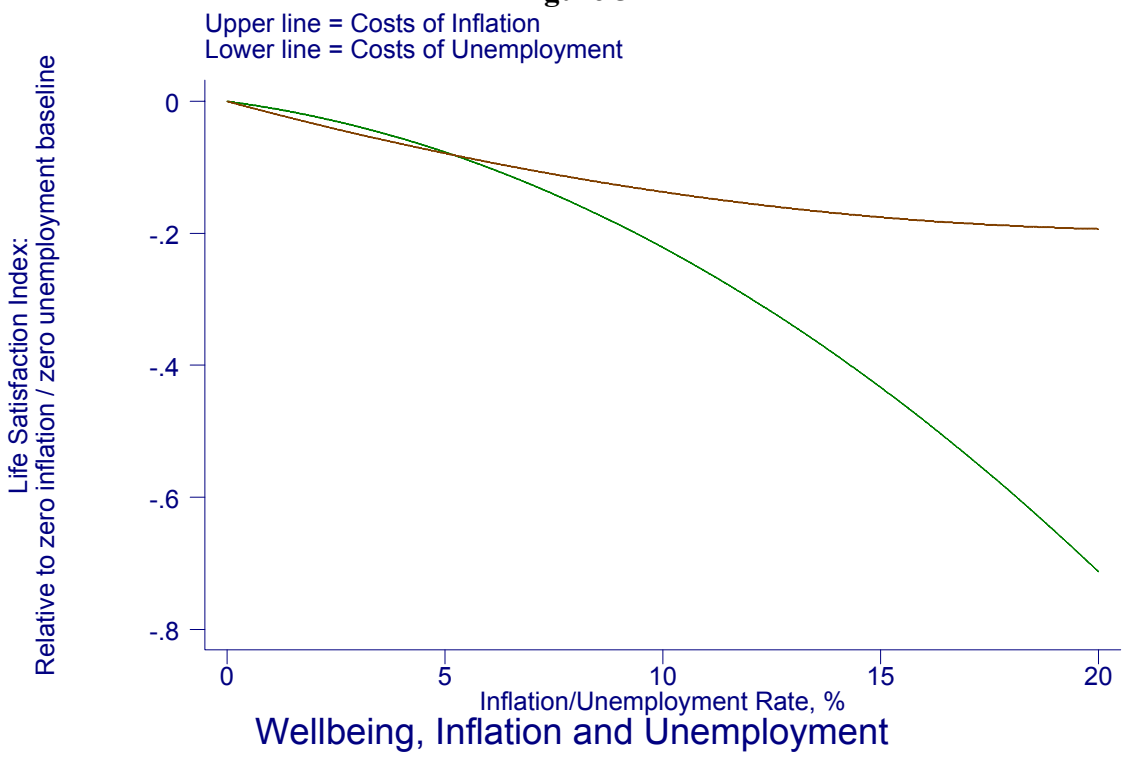

Source: Author's calculations based on results in Column 4 of Table 6.

There is clear evidence that increasing levels of unemployment do increasing harm to well-being. The effects of inflation on wellbeing are close to linear, and in four of the five columns, a null of linearity cannot be rejected.

Thus, these results suggest large gains to reducing output volatility, and little or no gain to reducing inflation volatility. For instance, these results suggest that when unemployment is 15%, the marginal effect of reducing unemployment on well-being is more than double that when unemployment is 5%. Consequently, if unemployment were to alternate between 5% and 15%, the well-being gained during booms would be less than that lost during recessions. Indeed, the average well-being in such an economy is equal to that in an economy with a constant unemployment rate of 11%.

This simple exercise can be taken somewhat more seriously in order to illustrate the welfare cost of business cycles. If one is willing to think of this aggregative function as a welfare function, $W(\cdot)$, then the equivalent variation of shifting to scenarios with greater and less unemployment volatility can be estimated.

That is, I ask: how much higher of an unemployment rate would the population be willing to tolerate in exchange for decreasing business cycle volatility? Thus, the costs of volatility associated with drawing macroeconomic variables from a distribution, $F(\cdot)$, can be calculated as:

$$W(\bar{x}) = \int W(x + equivalent\ variation) dF(x)$$

To make matters concrete, assume that each year's unemployment rate is drawn from a normal distribution with mean $\bar{x}$ and standard deviation σ^2. Table 3 presents some illustrative calculations. (Not surprisingly, estimates of compensating and equivalent variation are virtually identical.)

Table 7: Cost of Business Cycles: Measured as percentage points of unemployment

		Standard deviation of unemployment rate (percentage points)				
		.5%	**1%**	**2%**	**2.2%**	**4%**
Average unemp. rate	**5%**	-.02	-.06	-.25	-.30	-1.04
	7.7%	-.01	-.05	-.18	*-.22*	-.76
	10%	-.01	-.04	-.15	-.18	-.62
	15%	-.01	-.03	-.11	-.13	-.44

Source: Author's calculations of equivalent variation-based estimate of cost of business cycles, on the assumption that the unemployment rate is normally distributed, drawn from a distribution with standard deviation and mean shown above. Welfare function from Column 4 of Table 2.

For the 274 country-year observations in my sample, the average unemployment rate was 7.7%, and the within-country standard deviation of unemployment was 2.2%. Thus, according to estimates in Table 3, average well-being would be equally improved by either eliminating the business cycle, or lowering average unemployment rates by about a quarter of a percentage point (0.22 to be precise).

Thus, the gain in wellbeing from eliminating the business cycle, while not overwhelming, is certainly worthwhile. Further, while the aim of eliminating the business cycle is clearly unattainable, most of these gains will accrue simply from halving business cycle volatility. This can be seen by noting that the gain from eliminating business cycle volatility when the standard deviation is 4% is more than twice that from eliminating shocks drawn from an unemployment distribution with a standard deviation of 2%.

These decreasing returns to business-cycle smoothing can be seen more clearly in Figure 6.

Figure 6

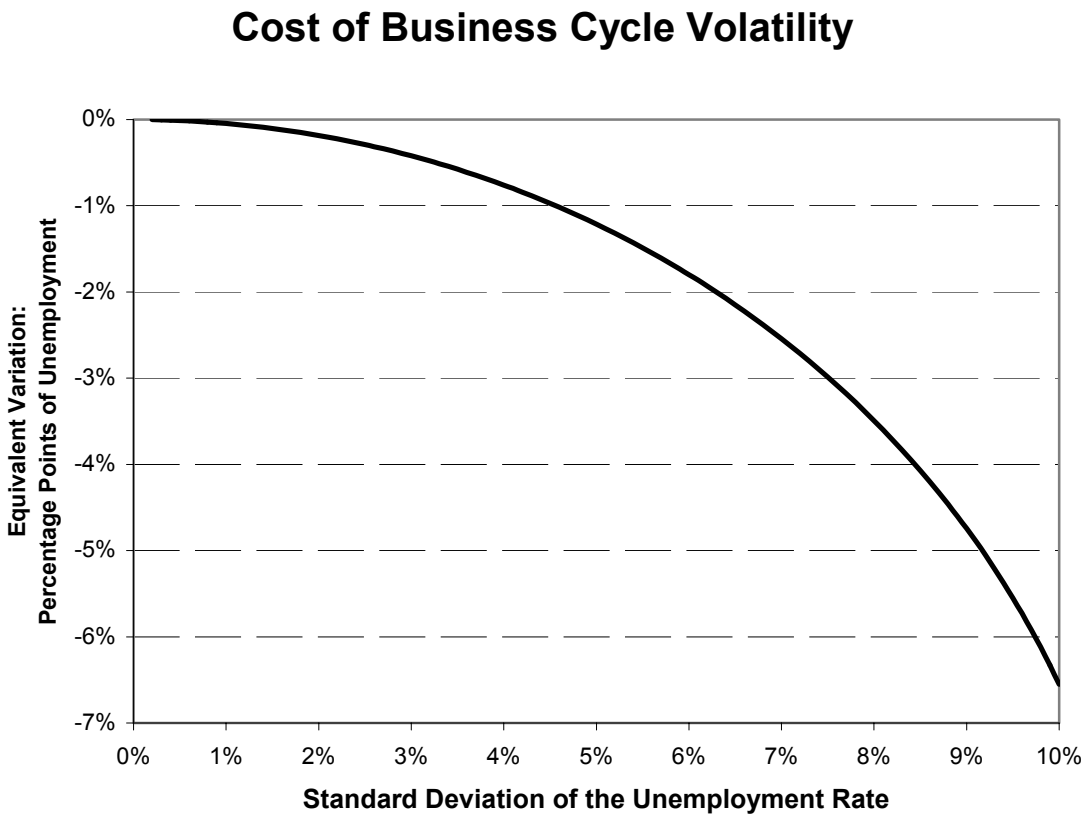

Figure 6 clearly shows large gains from eliminating high levels of unemployment volatility. For instance, if through adept macroeconomic management unemployment exhibits a standard deviation of 2.2% rather than 4%, then we can say that managing the cycle yields gains in wellbeing equivalent to lowering unemployment by half a percentage point. However, it is quite clear that decreasing returns have set in: the elimination of the remaining volatility would have an effect on wellbeing equivalent to lowering unemployment by only one-quarter of a percentage point.

Note that the estimate of the wellbeing cost of volatility in Table 5 was somewhat larger: halving the standard deviation of unemployment was estimated to increase wellbeing by an amount equivalent to decreasing the unemployment rate by around one-third of a percentage point. The figure above offers a useful reconciliation. The fact that the larger number was estimated is crucial: an estimated coefficient reflects only local variation in the data. Consequently the estimated effect from Table 5 measures something closer to the slope of a tangency line, while the smaller estimate derived in Table 7 is the slope of a ray to the origin.

Given the curvature in Figure 6, it is clear that estimate yields welfare effects of business cycle volatility that are too large.

Finally, one can also use Figure 6 to speculate as to the value of advances made by macroeconomic research. Christina Romer (1986) has compiled comparable unemployment rates series for the periods before 1930 and after 1948, concluding that there is the standard deviation of US unemployment may have declined by 10 percent. This corresponds to an effect on wellbeing roughly equivalent to reducing unemployment by one-twentieth of a percentage point.

6. Conclusion

Data on subjective wellbeing yield several important lessons for macroeconomists. Both unemployment and inflation significantly undermine aggregate measures of life satisfaction, happiness, and associated measures of psycho-social wellbeing. While feelings of depression and worthlessness rise with unemployment, feelings of usefulness, confidence, happiness and satisfaction with life fall. Not surprisingly, the public's faith in political institutions, and particularly the corporate sector also declines.

The relative size of the effects of wellbeing on inflation and unemployment suggest that disinflationary policies may have raised wellbeing in some countries, but lowered wellbeing in others. It is a close enough call that the happiness effects of such a policy surely depend on country-specific circumstances.

Whereas previous research has attempted to infer the costs of business cycle volatility by imposing utility or welfare functions, the evidence in this paper is based upon empirical observation of wellbeing. I find evidence that unemployment volatility undermines wellbeing; the evidence on inflation volatility is weaker, but suggestive. In terms of the magnitudes of these effects, it appears that there are important returns to minimizing business cycle volatility, but that diminishing returns have set in. Further efforts to tame the business cycle might increase wellbeing by an amount equivalent to reducing the unemployment rate by a quarter of a percentage point. Is this large? Compared with previous estimates, such as Lucas (1987), it is rather large. Compared with the effects of structural labor market reform (Blanchard and Wolfers, 2000), it is quite small. Compared with evidence from microeconometric happiness equations, this is an important gain—these data tend to suggest that factors like employment and marital status matter much more for wellbeing than does income.

References

Atkeson, Andrew and Christopher Phelan, "Reconsidering the Costs of Business Cycles with Incomplete Markets", in Stanley Fischer and Julio Rotemberg (eds) *NBER Macroeconomics Annual 1994* (Cambridge: MIT Press, 1994) pp.187-207.

Ball, Laurence, "What Determines the Sacrifice Ratio?" in N.G. Mankiw (ed.), *Monetary Policy*, (Chicago: University of Chicago Press, 1994) pp.155-182.

Ball, Laurence, "Disinflation and the NAIRU" in Christina Romer and David Romer (eds.), *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997.) pp. 167-185.

Bernanke, Ben S. "Irreversibility, Uncertainty, and Cyclical Investment", *Quarterly Journal of Economics*, 98(1) February 1983, 85-106.

Black, Fischer, *Business Cycles and Equilibrium*. (Cambridge: Basil Blackwell, 1987).

Blanchard, Olivier and Justin Wolfers, "The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence", *Economic Journal*, 110 March 2000.

Boskin, Michael J. and Eytan Sheshinski, "Optimal Redistributive Taxation When Individual Welfare Depends Upon Relative Income", *Quarterly Journal of Economics*, 92(4) November 1978, 589-601.

Debelle, Guy and Doug Laxton "Is the Phillips Curve Really a Curve?" Some Evidence for Canada, the United Kingdom and the United States", *IMF Staff Papers*, June 1997.

Diener, Ed. "Subjective Well-Being", *Psychological Bulletin*, 95(3) 1984, 542-575.

di Tella, Rafael; Robert J. MacCulloch and Andrew J. Oswald "Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness", *American Economic Review*, 91(1) March 2001 335-341.

Duesenberry, James, *Income, Saving and the Theory of Consumer Behavior* (Cambridge: Harvard University Press, 1949)

Easterlin, Richard, "Does Economic Growth Improve the Human Lot? Some Empirical Evidence". In P.A. David and M.W.Reder (eds) *Nations and Households in Economic Growth: Essays in Honour of Moses Abramovitz* (New York and London: Academic Press, 1974).

Easterlin, Richard, "Will Raising the Incomes of All Increase the Happiness of All?" *Journal of Economic Behavior and Organization*, 27(1) 1995, pp.35-48.

Easterlin, Richard, *Growth Triumphant: The Twenty-first Century in Historical Perspective*, (Ann Arbor, MI: University of Michigan Press, 1996).

Feldstein, Martin "The Costs and Benefits of Going from Low Inflation to Price Stability", in Christina Romer and David Romer (eds.), *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997.)

Galí, Jordi, Mark Gertler and J. David López-Salido "Markups, Gaps, and the Welfare Costs of Business Fluctuations", NBER Working Paper 8850, March 2002.

Inglehart, Ronald and Hans-Dieter Klingemann *Genes, Culture and Happiness*, (Cambridge: MIT Press, 2000)

Lucas, Robert E. *Models of Business Cycles* (Oxford: Basil Blackwell, 1987)

Ramey, Garey and Valerie A. Ramey, "Technology Commitment and the Cost of Economic Fluctuations", NBER Working Paper No. 3755, 1991.

Ramey, Garey and Valerie A. Ramey "Cross-Country Evidence on the Link Between Volatility and Growth" *American Economic Review* 85(5) December 1995, 1138-1151.

Romer, Christina, "Spurious Volatility in Historical Unemployment Data" *Journal of Political Economy* 94(1) 1986, 1-37.

Schmitt, Hermann and Evi Scholz, *Mannheim Eurobarometer Trend Files, 1970-1999* ICPSR Computer file 3384, 2002.

Veenhoven, Ruut, "Advances in the Understanding of Happiness", *Revue Quebecoise de Psychologie,* Vol 18, 1997 267-293.

Shiller, Robert J., "Why Do People Dislike Inflation?" in Christina Romer and David Romer (eds.), *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997.), pp.13-65.

Wolfers, Justin "Are Voters Rational? Evidence from Gubernatorial Elections", Stanford GSB Working Paper #1730, March 2002

Woodford, Michael "Inflation Stabilization and Welfare", NBER Working Paper 8071, January 2001.