

# Is Business Cycle Volatility Costly? Evidence from Surveys of Subjective Well-being\*

Justin Wolfers

*Stanford GSB.*

---

## Abstract

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

## I. Introduction: The Happiness Cost of Business Cycles<sup>1</sup>

The focus of macroeconomists on business cycle analysis presumably reflects a belief that macroeconomic fluctuations entail important

---

\*I thank Larry Ball, Jon Bendor, Olivier Blanchard, James Dailey, Rob Dugger, Jonathan Gardner, Mark Gertler, John McMillan, Andrew Oswald, Adam Posen, John Simon, Benn Steil, Lars Svensson and Betsy Stevenson for useful discussions. Doug Geyser and Ravi Pillai provided outstanding research assistance. Naturally, all remaining errors are my own.

<sup>1</sup>Data and programs used in this paper are available at the author's homepage: [www.stanford.edu/people/jwolfers](http://www.stanford.edu/people/jwolfers).

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 0.1% 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 analyse data on subjective well-being and compare the evolution of aggregate levels of self-reported happiness with macroeconomic conditions.

All three 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 well-being in a volatile macroeconomy and well-being experienced under average macroeconomic conditions. As such, volatility undermines well-being.

The macroeconomic data that I analyse 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 II. Section III establishes that both unemployment and inflation lower well-being. Section IV further explores the effect of the business cycle on feelings in two national datasets. Section V is the heart of the analysis, asking whether the data can speak to not only a preference for low unemployment and inflation, but also to low volatility. Section VI concludes.

To preview my findings, I find compelling evidence that unemployment and inflation both undermine well-being. People are much more averse to unemployment than to inflation, and even with a moderate sacrifice ratio, it

is a close call as to whether disinflationary policies have raised or lowered well-being. I find evidence that volatility in the real economy has important consequences. Current levels of aggregate unemployment volatility probably reduce well-being 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.

## II. 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 behaviour; 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 well-being are collected from several large-scale surveys that ask questions including the following:

- ‘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 well-being’. A host of validation studies have suggested that these questions reveal something meaningful and can indeed be interpreted as reflecting levels of well-being. For instance, those who report being happier tend to smile and laugh more and are typically rated by others as happier. Self-reported happiness also correlates with both heart rate and electrical activity in the brain. Measures of subjective well-being have relatively high test–retest correlations, similar microeconomic 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. Frey and Stutzer (2002) provide an interesting survey of recent happiness research within economics.

In his seminal paper, Richard Easterlin (1974) asked whether economic growth led to rising happiness. Analysing time series data in both the USA 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 (Easterlin, 1995, 1996) confirm these results. 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, on average, economic growth only makes you rich at the same rates as it does the Joneses. This not only leads to 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 well-being is highly responsive to the satisfaction of basic needs but almost invariant to income at higher levels of development.

Finally, economic and political institutions also affect 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 the links between happiness and business-cycle variation in unemployment and inflation.

### **III. Business Cycles and Happiness**

The most directly relevant analysis of happiness and the business cycle comes from di Tella et al. (2001). Analysing a country-year panel, they find that life 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 to 1991. My analysis uses the same survey but updates it to cover sixteen countries running from 1973 to 1998 (Schmitt and Scholz 2002). Thus, I start by updating these results in Table 1, as described below.

The Eurobarometer happiness question (shown above) was asked only from 1975 to 1986 (excluding 1980 and 1981), while the life satisfaction question was asked every year from 1973 to 1998, except 1974 and 1996. Overall, this yields an unbalanced panel of 504,581 valid responses to the life

**Table 1: Well-being and Macroeconomic Conditions: Eurobarometer Survey**

	Macro data			Micro data
	Average satisfaction	% Very satisfied with life	Satisfaction index $\sim N(0,1)$	Ordered probit
<i>Panel A: Dependent variable: Life satisfaction<sup>a</sup></i>				
Unemployment	-2.305*** (0.331)	-1.096*** (0.185)	-3.566*** (0.537)	-3.452*** (0.495)
Inflation <sub>c,t</sub>	-0.626*** (0.230)	-0.157 (0.112)	-0.766** (0.342)	-0.728** (0.325)
Trade-off: $\beta/\pi$ (90% interval) <sup>b</sup>	3.7 (2.1 – 8.1)	7.0 (-13.1–35.9)	4.6 (2.4–12.6)	4.7 (4.1–5.8)
State and year fixed effects	✓	✓	✓	✓
Adj. R <sup>2c</sup>	0.9326	0.9237	0.9341	0.0645
Sample	274	274	274	n = 504,581 274 country-year clusters
<i>Panel B: Effect on % 'Very Satisfied' (Sample mean = 28%)</i>				
Unemployment <sub>c,t</sub>	n.a.	-1.1%	-1.2%	-1.1%
Inflation <sub>c,t</sub>		-0.2%	-0.3%	-0.2%

*Note*

\*\*\*, \*\*, \* 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 regression analysing individual-level data on reported satisfaction

<sup>a</sup>Micro data regressions use sample weights. Macro regressions use simple OLS (data constructed using weights).

<sup>b</sup>90% confidence interval on the unemployment-inflation trade-off reflects the 10th and 90th percentile estimates from 1000 bootstrap replications.

<sup>c</sup>Adjusted R<sup>2</sup> in columns 1–3; Pseudo R<sup>2</sup> in column 4.

satisfaction question in 274 country-years and 134,590 responses to the happiness question over 99 country-years.<sup>2</sup> Answers to the two questions are highly correlated and so, given the greater data availability, my analysis will focus on life satisfaction.

<sup>2</sup>Northern Ireland is included as part of the UK. For consistency, I drop East Germany from the sample and follow only West Germans.

An important technical issue relates to finding a sensible set of units with which to measure well-being. As only qualitative indicators exist, there are no natural happiness units that immediately correspond with the textbook ‘util’. Throughout this paper, I analyse four alternatives:

- Following di Tella et al. (2001), I simply take country-year averages of well-being, assigning a value of 1 to the lowest satisfaction category, and with each qualitative increment, add one.<sup>3</sup>
- 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 ‘fairly satisfied’ and ‘not very 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 well-being 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.<sup>4</sup> That is, this regression estimates numerical values for each qualitative response that are most likely given the sample proportions in the data and the assumption that the true underlying distribution of happiness is normally distributed.
- 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). The intuition underlying this regression is essentially similar to that above.

<sup>3</sup>Thus, for the Eurobarometer survey: 1 = ‘not at all satisfied’; 2 = ‘not very satisfied’; 3 = ‘fairly satisfied’; 4 = ‘very satisfied’.

<sup>4</sup>Thus, this cardinalization derives explicitly from an identifying assumption that well-being 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.

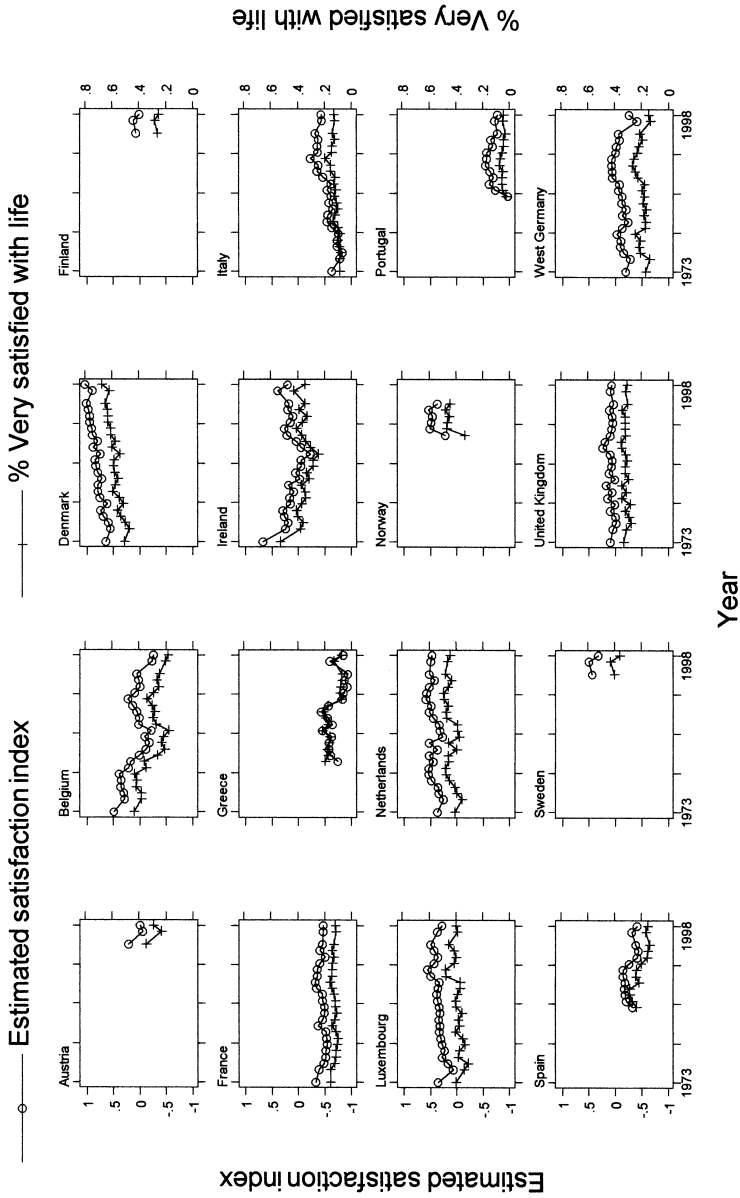


Figure 1: Evolution of life satisfaction by country

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 described above.

Table 1 analyses the relationship between measures of subjective well-being and annual measures of inflation and unemployment.<sup>5</sup> In each case, the regressions are of the form:

$$\begin{aligned} \text{Life Satisfaction}_{c,t} = & \beta \text{Unemployment}_{c,t} + \pi \text{Inflation}_{c,t} \\ & + \text{country}_c \text{ and year}_t \text{ fixed effects} + \varepsilon_{c,t} \end{aligned}$$

In Table 1, Panel A reports the main regression results. Each column shows a separate regression, and across all four columns, unemployment lowers subjective well-being. This result is robust across specifications and easily meets standard levels of statistical significance. The effect of inflation on well-being is less statistically robust but, nonetheless, these data suggest that inflation lowers well-being. 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 and 4 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 would decline by about 0.35 standard deviations. Figure 2 shows the intuition underlying these estimates.

The bell curve A shows the (assumed) standard normal distribution of happiness in the population, and the estimated cut-points reflect the most likely categorization of respondents into ‘very’, ‘fairly’, ‘not very’ or ‘not at all’ satisfied with their life, given observed proportions in the data. 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. Reflecting the estimated coefficients on unemployment and inflation of 3.5 and 0.7, 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.

<sup>5</sup>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.



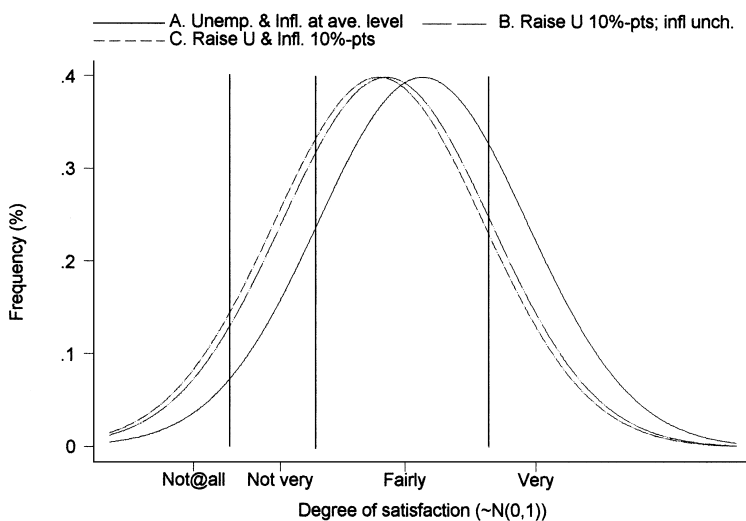


Figure 2: Distribution of happiness under 3 scenarios

Source: Author's calculation is based on column 4 of Table 1.

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 column 3. Panel B of Table 1 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 reports feeling '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 et al. (2001) 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 well-being. 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 well-being data give a new 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 sacrifice ratio, often drawn 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 above its equilibrium level by two percentage points for a year. The benefits of this temporary contraction are assumed to be permanent. 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 2%-years of higher unemployment for 16%-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 well-being. However, if the rise in unemployment following a monetary contraction is not completely transitory – as suggested by Ball (1997) or Blanchard and Wolfers (2000) – or the decline in inflation is not completely persistent, then disinflation probably lowers well-being. Indeed, the results in Table 1, when interpreted jointly with evidence of heterogeneous and persistent costs of disinflation (Ball 1994, 1997), suggest that for some countries disinflation almost certainly undermines well-being.

#### IV. Further Evidence on Well-being and Unemployment

Several other data sources shed additional light on the well-being 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, each respondent was asked to submit to a psychological test each year. This test, called the GHQ-12 (General Health Questionnaire), consists of a battery of twelve questions originally developed as a screening instrument for psychiatric illness but increasingly used as an indicator of psycho-social well-being.

Each of the GHQ-12 questions is shown in Table 2. The first column also reports the proportion of respondents who state that their well-being on that particular measure is less than usual.<sup>6</sup> 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 labour force survey and claimant counts. Unfortunately, consistent regional

---

<sup>6</sup>Responses were recoded so that higher numbers denote greater well-being. 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.

inflation indicators are not available. To test for the influence of the business cycle on well-being, I ran ordered probit regressions of well-being on the regional unemployment rate, controlling for region and year fixed effects. The second column of Table 2 shows results where the independent variable is the claimant count measure of unemployment, while the third column analyses the survey measure. The fourth and fifth columns report these estimates as marginal effects, showing the implied effect of higher unemployment on the share of people reporting that they are unwell.

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 3 turns to data from the US General Social Survey to show that state-specific economic downturns also undermine happiness in the USA. I regress happiness on two alternative measures of state labour market conditions: the CPS unemployment rate and a measure constructed as the deviation of (log) non-farm payrolls in that state from a Hodrick–Prescott filtered state trend, described in Wolfers (2002). As in the Tables 1 and 2, I show both the coefficients from an ordered probit and the corresponding marginal effects. In this case, the marginal effects shown are for the percentage who are ‘very happy’.

Beyond the effects on happiness, Table 3 also shows that economic downturns undermine public confidence in political and economic institutions. Intriguingly, the effects on faith in large companies, banks, financial institutions and even education are larger than those on political institutions such as the President and Congress.<sup>7</sup>

## V. Are People Averse to Economic Fluctuations?

Results in the previous sections suggest that both inflation and unemployment undermine well-being, that these effects are reflected in many different domains, and that unemployment undermines public confidence in economic and political institutions. In this section, I turn to examining the effects of macroeconomic *volatility* on well-being. 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 well-being. Costs of macroeconomic volatility arise either due to convexity in the structure of the economy or due to

---

<sup>7</sup> Presumably this muted anger at Washington partly reflects the fact that these are responses to state unemployment relative to the national trend (recall that the regression controls for year fixed effects).

Table 2: Effects of Unemployment on Various Dimensions of Well-being: BHPS

Average	Estimated coefficients			
	Reported as effect sizes: Change in well-being [ $\sim N(0,1)$ ]	LFS	Claimant count	Reported as marginal effects: Change in % reporting unwell <sup>a</sup>
% Reporting unwell <sup>a</sup>	Claimant count	LFS	Claimant count	LFS
A. <i>Have you recently been able to concentrate on whatever you're doing?</i> 19.4%	-1.360 (1.390)	-0.412 (0.913)	0.4%	0.1%
B. <i>Have you recently lost much sleep over worry?</i> 18.2%	-1.673 (0.888)	-0.452 (0.622)	0.4%	0.1%
C. <i>Have you recently felt that you were playing a useful part in things?</i> 13.4%	-2.554 (1.14)	-2.563 (0.890)	0.6%	0.6%
D. <i>Have you recently felt capable of making decisions about things?</i> 9.3%	0.364 (1.119)	0.263 (0.889)	-0.1%	0.0%
E. <i>Have you recently felt constantly under strain?</i> 28.0%	-1.416 (0.839)	-1.252 (0.697)	0.5%	0.4%
F. <i>Have you recently felt you couldn't overcome your difficulties?</i> 13.9%	-1.219 (0.782)	-0.366 (0.630)	0.3%	0.1%
G. <i>Have you recently been able to enjoy your normal day-to-day activities?</i> 19.1%	0.002 (1.208)	0.417 (0.887)	0.0%	-0.1

H. Have you recently been able to face up to problems? 10.9%	-1.915 (1.361)	-0.579 (0.939)	0.4%	0.1%
I. Have you recently been feeling unhappy or depressed? 21.9%	-1.768 (0.850)	-1.805 (0.619)	0.5%	0.5%
J. Have you recently been losing confidence in yourself? 13.1%	-3.662 (0.821)	-2.153 (0.681)	0.8%	0.5%
K. Have you recently been thinking of yourself as a worthless person? 7.1%	-2.204 (0.722)	-0.871 (0.604)	0.3%	0.1%
L. Have you recently been feeling reasonably happy, all things considered? 13.0%	-1.762 (1.130)	-1.945 (0.948)	0.4%	0.4%

Notes: Each cell reports the coefficient of interest from a separate ordered probit regression, controlling for region and year fixed effects. Columns 4 and 5 report marginal effects implied by the coefficients in columns 2 and 3, respectively. 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 *xrwtgt*.

<sup>a</sup>'Unwell' equates to reporting well-being less than usual or much less than usual on a particular measure.

Table 3: Happiness and Unemployment Across US States: General Social Survey

	Average	Estimated coefficients			
		Reported as effect sizes: Change in happiness [ $\sim N(0,1)$ ]		Marginal effects: Change in % reporting very happy <sup>a</sup>	
	% Very happy or very confident	Unemp. rate	Emp. gap	Unemp. rate	Emp. gap
<i>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?</i>					
Happiness	32.5%	-1.539*** (0.519)	-0.632*** (0.218)	-0.3%	-0.1%
<i>I am going to name some institutions in this country. As far as the people running these institution 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?</i>					
Executive branch-Federal govt.	17%	-0.962 (0.725)	-0.784*** (0.301)	-0.2%	-0.2%
Congress	14%	-0.698 (0.651)	-0.666** (0.294)	-0.1%	-0.1%
US Supreme Court	34%	-0.744 (0.735)	-0.133 (0.303)	-0.3%	0.0%
Banks & financial institutions	26%	-3.084*** (0.672)	-1.276** (0.306)	-1.0%	-0.4%
Major companies	26%	-1.491** (0.677)	-0.774*** (0.298)	-0.5%	-0.3%

Organized labour	11%	-0.379 (0.630)	-0.115 (0.312)	0.1%	0.0%
Organized religion	28%	-0.601 (0.720)	-0.386 (0.320)	-0.2%	-0.1%
Education	30%	-1.445** (0.737)	-0.755** (0.320)	-0.5%	-0.3%
Press	18%	-0.352 (0.613)	-0.037 (0.285)	-0.1%	0.0%
Medicine	48%	-0.244 (0.692)	-0.205 (0.303)	0.1%	-0.1%
Television	13%	-0.601 (0.685)	-0.312 (0.315)	0.1%	0.1%
Scientific community	44%	-0.103 (0.598)	-0.034 (0.291)	0.0%	0.0%
Military	36%	-0.561 (0.660)	-0.129 (0.296)	-0.2%	0.0%

Notes: Each cell reports the coefficient of interest from a separate ordered probit regression, controlling for state and year fixed effects. Columns 4 and 5 report marginal effects implied by the coefficients in columns 2 and 3, respectively.

Sample: GSS data for 1973-98,  $n = 36,146$  valid interviews, in 773 state-year clusters for happiness;  $n = 21,245$  valid responses to all confidence questions in 667 state-year clusters.

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

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

convexity in preferences. Without this convexity, the gain to well-being in a boom offsets the loss in a recession.

Examples of convexities in economic structure include nonlinearities 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 in macroeconomic data as reduced output or employment.<sup>8</sup>

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 before one can make any inferences.

Lucas (1987) analyses the costs of volatility due to imperfect insurance in a representative agent model. He finds that aggregate consumption risk due to the business cycle undermines well-being by an amount roughly equivalent to 0.1% of the consumption of the representative agent. Essentially, he argues that aggregate consumption varies too little through the cycle to be of much concern.

By analysing a representative agent model, Lucas explicitly sidesteps the issue of idiosyncratic labour market risk. While not denying the importance of idiosyncratic risk, he argues that it reflects ‘the potential or actual gain from social insurance, not from stabilization policy’. Atkeson and Phelan (1994) address idiosyncratic risk directly. They argue that the main source of idiosyncratic risk involves the risk of job loss.

However, to see why these effects may be small, consider a highly stylized model, drawing on Romer (1996, p. 415). Consider the expected utility of a worker who faces the risk that her job is eliminated next period, which would cause her utility to shift from  $U_E$  to  $U_U$ . In normal times, the job destruction rate is  $s$ , but a recession may occur with probability  $p$ , raising the destruction rate to  $s + e$ ; there is an equal likelihood of a boom that will reduce the firing rate to  $s - e$ . Thus, the worker’s expected utility is:

$$\begin{aligned} \text{Expected utility} &= p[(1 - s - e)U_E + (s + e)U_U] \\ &\quad + p[(1 - s + e)U_E + (s - e)U_U] + (1 - 2p)[(1 - s)U_U + sU_E] \\ &= sU_U + (1 - s)U_E \end{aligned}$$

<sup>8</sup>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.



The first term reflects recessions, the second term booms, and the third reflects times of calm. Importantly, note that this simplifies to an expression in which expected utility is independent of both  $p$  and  $e$  – the likelihood and magnitude of a shock to the job destruction rate, respectively. That is, dampening the business cycle has no immediate effect on average levels of well-being in this stylized example. This result reflects the fact that expected utility is linear in probabilities.

However, the risk of job loss is more highly correlated across individuals when shocks to the job destruction rate are either larger or more common. That is, the probability that both a worker and her husband lose their jobs at the same time is higher when macroeconomic volatility is greater. Consequently, the ability of husband and wife to insure each other against job loss declines with macroeconomic volatility, undermining average well-being. 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.)

Romer (1996) notes that beyond consumption volatility, cyclical variability of hours worked is much larger than that of consumption, and this may pose major welfare costs. A competitive labour market equilibrium achieves an efficient allocation. Hence, recessions may drive the marginal rate of substitution between labour and leisure below the marginal product of labour; booms potentially cause the reverse misallocation, with the marginal product below the value of leisure. These costs of volatility accrue even when imperfect competition yields an equilibrium in which employment is sub-optimally low. Figure 3, adapted from Galí et al. (2002), shows that, under imperfect competition, booms drive the economy toward the first best, while recessions drive the economy even further away from the optimum. Incrementally larger deviations from first best have an increasing welfare cost (the gap between the supply and demand curves widens), suggesting that recessions undermine well-being by more than equal-sized booms. These three authors calibrate a model to infer this cost, arguing that, for the USA, business cycle volatility yields welfare losses equivalent to about 0.8% of consumption.

Finally, on the assumption that more volatile inflation is harder to forecast, aversion to inflation volatility can be motivated by reference either to the redistributive effects of unanticipated inflation or to 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 well-being from an imposed preference structure, my point of departure is to treat these well-being costs of business cycle volatility as directly observable. Thus, I now turn to examining whether volatility in




Figure 3: Labour market

unemployment and inflation undermines well-being, conditional on average levels of unemployment and inflation.<sup>9</sup> That is, in Table 4, I regress life satisfaction against the standard deviation of both unemployment and inflation over the past eight years and also control for the mean levels of these indicators over the same period:

$$\begin{aligned}
 \text{Life Satisfaction}_{c,t} = & \beta \text{Average Unemployment}_{c,t-7..t} \\
 & + \lambda \text{Standard Deviation Unemployment}_{c,t-7..t} \\
 & + \pi \text{Average Inflation}_{c,t-7..t} \\
 & + \delta \text{Standard Deviation Inflation}_{c,t-7..t} \\
 & + \text{country}_c \text{ \& year}_t \text{ fixed effects} + \varepsilon_{c,t}
 \end{aligned}$$

<sup>9</sup>By controlling for levels of unemployment and inflation, I partial out those effects of volatility that are mediated by higher volatility causing worse macroeconomic outcomes.

**Table 4: Life Satisfaction: Levels and Volatility of Inflation and Unemployment**

	Macro data			Micro data
	Average satisfaction	% Very satisfied with life	Satisfaction index $\sim N(0,1)$	Ordered probit
Panel A: <i>Dependent variable: life satisfaction</i>				
Unemployment from $t - 7$ to $t$				
Mean	-3.109*** (0.507)	-1.666*** (0.263)	-5.141*** (0.776)	-6.029*** (0.779)
Standard deviation	-0.342 (1.007)	-0.737 (0.536)	-1.407* (1.547)	-2.411* (1.475)
Inflation from $t - 7$ to $t$				
Mean	-0.623 (0.343)	-0.343** (0.147)	-1.193*** (0.493)	-1.265*** (0.461)
Standard deviation	-0.993 (0.623)	-0.342 (0.273)	-1.610* (0.888)	-1.477** (0.736)
Adj. $R^2$	$\sqrt{0.9328}$	$\sqrt{0.9259}$	$\sqrt{0.9363}$	$\sqrt{0.0650}$
Sample	274	274	274	$n = 504,581$ 274 country-year clusters
Panel B: <i>Effect on % 'very satisfied' (Sample mean = 28%)</i>				
Mean unemployment	n.a.	-1.7%	-1.7%	-2.0%
SD unemployment		-0.7%	-0.5%	-0.8%
Mean deviation		-0.3%	-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, controlling for state and year-fixed effects.

For the 5 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 well-being, while the second row shows that volatile unemployment also lowers well-being. Reflecting statistical imprecision, the results on unemployment volatility are at best suggestive: the estimate is only statistically significant in the last regression, but the point estimate is relatively large and negative in all four specifications.

The standard deviation of unemployment over an 8-year period averaged around 1.5 percentage points in my sample. Thus, the results in the last two columns suggest that halving the standard deviation of unemployment shifts the well-being distribution to the right by around  $\frac{3}{4}\% * 2 \approx 0.015$  standard

deviations, increasing the proportion of the population who are very satisfied by around 0.5 percentage points. By comparison, the gains from reducing the average level of unemployment by one percentage point are much larger – shifting the well-being distribution to the right by  $1\% * 6 \approx 0.06$  standard deviations and raising the percentage very happy by 2 percentage points. Combining these findings suggests that halving unemployment volatility yields gains in well-being approximately equal to those from reducing the unemployment rate by one-quarter of a percentage point.

The inflation results are also interesting. As in Table 1, higher average levels of inflation lower well-being. Moreover, the fourth row suggests that there may be costs to inflation volatility, with two of the four estimates statistically significant. There is insufficient precision in these reduced-form estimates to say much about the relative importance of stabilizing inflation versus unemployment.

Rather than make indirect inferences about the shape of the well-being function, Table 5 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 labour supply and labour demand curves, as shown in Figure 3.)<sup>10</sup>

All four specifications suggest that aggregate well-being is a nonlinear function of unemployment and that this nonlinearity is statistically significant; the results for inflation are less clear. Figure 4 graphs these results, showing how well-being declines with unemployment and inflation, relative to a zero unemployment/inflation baseline.

Figure 4 provides clear evidence that increasing levels of unemployment do increasing harm to well-being. The effects of inflation on well-being are close to linear, and in three of the four 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 quadruple 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%.

---

<sup>10</sup>Woodford (2001) provides detailed micro-foundations for a welfare function that is quadratic in inflation and in the output gap.

**Table 5: Are Preferences Convex?**

	Macro data			Micro data
	Average satisfaction	% Very satisfied with life	Satisfaction index $\sim N(0,1)$	Ordered probit
<i>Dependent variable: Life Satisfaction</i>				
Unemployment <sub>c</sub>	-0.461 (0.821)	-0.164 (0.488)	-0.420 (1.383)	-0.844 (1.184)
Unemployment <sup>2</sup> <sub>c,t</sub>	-9.488** (3.834)	-4.865** (2.354)	-16.252** (6.622)	-13.569** (5.378)
Inflation <sub>c,t</sub>	-1.058** (0.462)	-0.675*** (0.233)	-1.820** (0.655)	-1.773*** (0.661)
Inflation <sup>2</sup> <sub>c,t</sub>	-1.506 (1.866)	1.974** (0.843)	3.811 (2.747)	4.068 (2.646)
Adj.R <sup>2</sup>	0.9341	0.9262	0.9362	0.0647
Sample	274	274	274	n = 504,581 274 country-year clusters
<i>Marginal Effects on % 'Very Satisfied' (Sample mean = 28%)</i>				
Unemployment <sub>c,t</sub>	<i>(Holding inflation constant at sample mean of 6.5%)</i>			
At 0%	n.a.	-0.2%	-0.1%	-0.2%
At 7 1/2%		-0.5%	-0.5%	-0.6%
At 15%		-0.9%	-1.1%	-1.1%
Inflation <sub>c,t</sub>	<i>(Holding unemployment constant at sample mean of 7.7%)</i>			
At 0%		-0.7%	-0.6%	-0.5%
At 7 1/2%		-0.5%	-0.5%	-0.5%
At 15%		-0.4%	-0.4%	-0.4%

Each column shows results from a separate regression. See notes to Table 1.

This simple exercise can be taken somewhat more seriously so as to illustrate the welfare cost of business cycles. The quadratic preferences over inflation and unemployment shown above can be rewritten as preferences over the mean and variance of unemployment and inflation, respectively. To see this, note that Table 5 suggests that expected well-being takes the form:

$$\begin{aligned}
 \text{Expected Well-being} &= E[au^2 + bU + c\pi^2 + d\pi + e] \\
 &= a[E(U^2) - E(U)^2] + [aE(U)^2 + bE(U)] \\
 &+ c[E(\pi^2) - E(\pi)^2] + [cE(\pi)^2 + dE(\pi)] + e \\
 &= a\text{Var}(U) + (a\bar{U} + b)\bar{U} + c\text{Var}(\pi) + (c\bar{\pi} + d)\bar{\pi} + e
 \end{aligned}$$

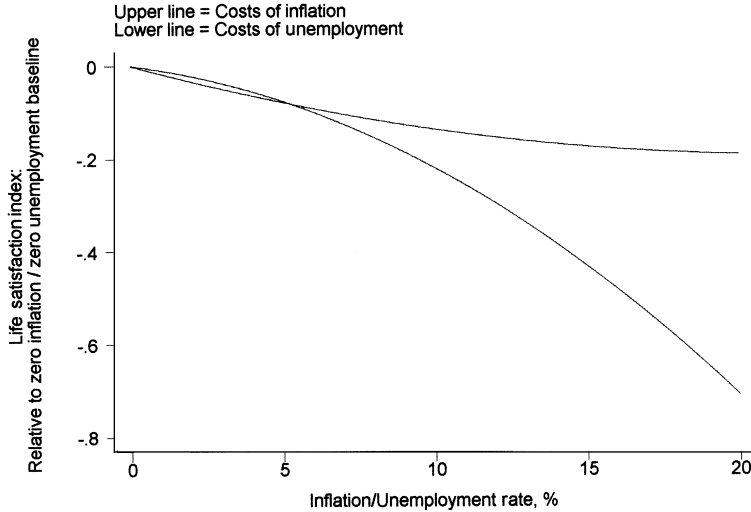


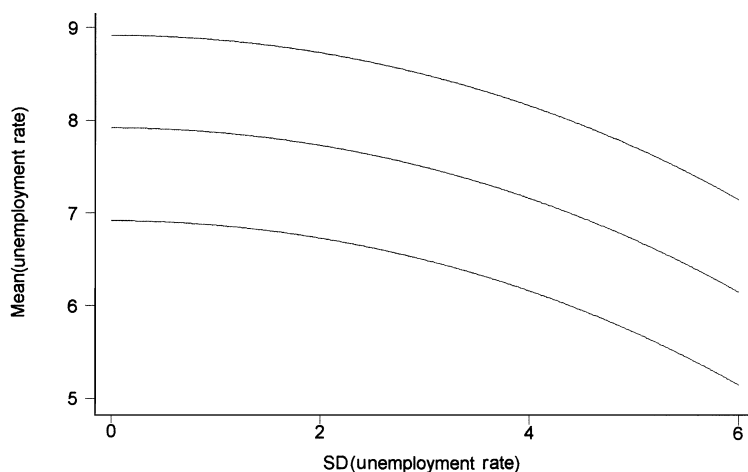
Figure 4: Well-being, inflation and unemployment

This allows me to show the results in Table 5 as indifference curves defined over the mean and standard deviation of unemployment and inflation, respectively.

Figure 5 clearly shows large gains from eliminating high levels of unemployment volatility. However, while eliminating high levels of volatility is extremely valuable, diminishing returns set in rather quickly. Table 6 shows some specific costs of business cycle volatility – essentially tracking the middle indifference curve. These welfare costs are measured in units equivalent to the rise in average unemployment that would make workers indifferent between a specific level of volatility and eliminating the business cycle.

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 well-being equivalent to those from lowering unemployment by  $0.74 - 0.22 \approx$  half a percentage point. Within my sample, the within-country standard deviation of unemployment actually averages 2.2%. It is quite clear that decreasing returns to smoothing the cycle have set in: the elimination of the remaining volatility would have an effect on well-being equivalent to lowering unemployment by only one-quarter of a percentage point (0.22, to be precise).

Note that the estimate of the well-being cost of volatility derived from the coefficients in Table 4 was somewhat larger: those numbers suggested that one needs only to halve the standard deviation of unemployment to yield increases in well-being equivalent to decreasing the unemployment rate by a quarter of a percentage point. The figure above offers a useful partial



**Figure 5:** Indifference curves: Level and volatility of unemployment  
*Notes:* Middle line passes through sample means:  $U=7.7\%$ ,  $SD(U)=2.2\%$  Based on results in Table 5, final column

**Table 6: Welfare Costs of Business Cycle Volatility**

SD (Unemployment rate)	0.5%	1%	2%	2.2%	4%	8%
Total welfare cost of volatility (%pts of higher unemployment)	0.01	0.04	0.18	0.22	0.74	3.37
$\delta(\text{Welfare cost})/\delta(\text{SD of Unemployment})$	0.04	0.09	0.18	0.20	0.38	1.02

*Source:* Author’s calculations based on results in column 4 of Table 5 (or middle indifference curve in Figure 5).

reconciliation. The fact that the larger number was derived from a reduced-form estimate is crucial: an estimated coefficient reflects only local variation in the data. Consequently, the estimated effect from Table 4 effectively extrapolates something close to the slope of a tangent to the indifference curve at the mean level of volatility, while the smaller estimate suggested here reflects the height of a secant to the  $y$ -axis. Given the curvature shown in the figure, it is clear that the slope of the tangent yields estimates of the total welfare effects of business cycle volatility that are too large. (Alternatively, compare the total welfare cost of volatility reported in the middle line of Table 6 and that implied by the slope – multiplying the top and bottom lines.)

Finally, one can also use Table 6 to speculate on advances made by macroeconomic research. Christina Romer (1986) has compiled comparable unemployment rate data for the periods before 1930 and after 1948,

concluding that the standard deviation of US unemployment may have declined by 10%. This corresponds to an effect on well-being roughly equivalent to reducing unemployment by one-twentieth of a percentage point.

## VI. Conclusion

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

The relative size of the effects on well-being of inflation and unemployment suggests that disinflationary policies may have raised well-being in some countries but lowered well-being 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 on empirical observation of well-being. I find evidence that unemployment volatility undermines well-being, while 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 smoothing the business cycle, but that diminishing returns have set in. Further efforts to tame the business cycle might increase well-being 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 labour market reform, though, it is quite small. Regardless, compared with evidence from microeconomic happiness equations, this is an important gain – these data tend to suggest that factors like employment and marital status matter much more for well-being than does income.

Justin Wolfers  
Political Economy Group  
Graduate School of Business  
Stanford University  
518 Memorial Way  
Stanford, CA 94305, USA  
jwolfers@stanford.edu



## References

- Atkeson, Andrew, and Christopher Phelan (1994), 'Reconsidering the Costs of Business Cycles with Incomplete Markets', in Stanley Fischer and Julio Rotemberg (eds), *NBER Macroeconomics Annual 1994*, Cambridge, MA: MIT Press, 187–207.
- Ball, Laurence (1994), 'What Determines the Sacrifice Ratio?' in N. G. Mankiw (ed.), *Monetary Policy*, Chicago: University of Chicago Press, 155–82.
- Ball, Laurence (1997), 'Disinflation and the NAIRU', in Christina Romer and David Romer (eds), *Reducing Inflation: Motivation and Strategy*, Chicago: University of Chicago Press, 167–85.
- Bernanke, Ben S. (1983), 'Irreversibility, Uncertainty, and Cyclical Investment', *Quarterly Journal of Economics*, 98(1), February, 85–106.
- Black, Fischer (1987), *Business Cycles and Equilibrium*. Cambridge: Basil Blackwell.
- Blanchard, Olivier, and Justin Wolfers (2000), 'The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence', *Economic Journal*, 110, March, 1–33.
- Boskin, Michael J., and Eytan Sheshinski (1978), 'Optimal Redistributive Taxation When Individual Welfare Depends Upon Relative Income', *Quarterly Journal of Economics*, 92(4), November, 589–601.
- Debelle, Guy, and Doug Laxton (1997), 'Is the Phillips Curve Really a Curve? Some Evidence for Canada, the United Kingdom and the United States', *IMF Staff Papers*, 44, June, 249–82.
- di Tella, Rafael; Robert J. MacCulloch and Andrew J. Oswald (2001), 'Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness', *American Economic Review*, 91(1), March, 335–41.
- Diener, Ed (1984), 'Subjective Well-Being', *Psychological Bulletin*, 95(3), 542–75.
- Duesenberry, James (1949), *Income, Saving and the Theory of Consumer Behaviour*. Cambridge, MA: Harvard University Press.
- Easterlin, Richard (1974), '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 Honor of Moses Abramovitz*, New York and London: Academic Press, 89–125.
- Easterlin, Richard (1995), 'Will Raising the Incomes of All Increase the Happiness of All?', *Journal of Economic Behavior and Organization*, 27(1), 35–48.
- Easterlin, Richard (1996), *Growth Triumphant: The Twenty-first Century in Historical Perspective*. Ann Arbor, MI: University of Michigan Press.
- Feldstein, Martin (1997), '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, 47–94.

- Frey, Bruno, and Stutzer, Alois (2002), 'What can Economists Learn from Happiness Research?' *Journal of Economic Literature*, 40(2), 402–35.
- Galí, Jordi, Mark Gertler and J. David López-Salido (2002), 'Markups, Gaps, and the Welfare Costs of Business Fluctuations', NBER Working Paper Na 8850, March.
- Inglehart, Ronald, and Hans-Dieter Klingemann (2000), *Genes, Culture and Happiness*. Cambridge, MA: MIT Press.
- Lucas, Robert E. (1987), *Models of Business Cycles*. Oxford: Basil Blackwell.
- Ramey, Garey, and Valerie A. Ramey (1991), 'Technology Commitment and the Cost of Economic Fluctuations', NBER Working Paper No. 3755.
- Ramey, Garey, and Valerie A. Ramey (1995), 'Cross-country Evidence on the Link Between Volatility and Growth', *American Economic Review*, 85(5), December, 1138–51.
- Romer, Christina (1986), 'Spurious Volatility in Historical Unemployment Data', *Journal of Political Economy*, 94(1), 1–37.
- Romer, David (1996), *Advanced Macroeconomics*. New York: McGraw Hill.
- Schmitt, Hermann, and Evi Scholz (2002), *Mannheim Eurobarometer Trend Files, 1970–1999*, ICORE News, 2, 4–5. ICPSR, Computer file 3384.
- Shiller, Robert J. (1997), 'Why Do People Dislike Inflation?' in Christina Romer and David Romer (eds.), *Reducing Inflation: Motivation and Strategy*, Chicago: University of Chicago Press, 13–65.
- Veenhoven, Ruut (1997), 'Advances in the Understanding of Happiness', *Revue Quebecoise de Psychologie*, 18, 267–93.
- Wolfers, Justin (2002), 'Are Voters Rational? Evidence from Gubernatorial Elections', Stanford GSB Working Paper No. 1730, March.
- Woodford, Michael (2001), 'Inflation Stabilization and Welfare', NBER Working Paper No. 8071, January.