Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness

Rafael Di Tella
Harvard Business School

Robert J. MacCulloch
ZEI, University of Bonn

and

Andrew J. Oswald[1]
University of Warwick

Revised: February 7, 2000

Modern macroeconomics textbooks rest upon the assumption of a social welfare function defined on inflation, π, and unemployment, U.[2] However, no formal evidence for the existence of such a function has been presented in the literature.[3] Although an optimal policy rule cannot be chosen unless the parameters of the presumed $W(\pi, U)$ function are known, that has not prevented its use in a large theoretical literature in macroeconomics.

This paper has two aims. The first is to show that citizens care about these two variables. We present evidence that inflation and unemployment belong in a well-being function. The second is to calculate the costs of inflation in terms of unemployment. We measure the relative size of the weights attached to these variables in social well-being. Policy implications emerge.

Economists have often puzzled over the costs of inflation. Survey evidence presented in Robert Shiller (1996) shows that, when asked how they feel about inflation, individuals report a number of unconventional costs, like exploitation, national prestige and loss of morale. Skeptics wonder. One textbook concludes: "we shall see that standard characterisations of the policy-maker's objective function put more weight on

the costs of inflation than is suggested by our understanding of the effects of inflation; in doing so, they probably reflect political realities and the heavy political costs of high inflation." (pp. 567-8, Olivier Blanchard and Stanley Fischer (1987)). Since reducing inflation is often costly, in terms of extra unemployment, some observers have argued that the industrial democracies' concern with nominal price stability is excessive -- and have urged different monetary policies.[4]

This paper proposes a new approach. It examines how survey respondents' reports of their well-being vary as levels of unemployment and inflation vary. Because the survey responses are available across time and countries, we are able to quantify how self-reported well-being alters with unemployment and inflation rates. Only a few economists have looked at patterns in subjective happiness and life-satisfaction. Richard Easterlin (1974) helped to begin the literature. Later contributions include David Morawetz et al (1977), Robert Frank (1985), Yew-Kwang Ng (1996), Ronald Inglehart (1990), Andrew Oswald (1997), and Liliana Winkelmann and Rainer Winkelmann (1998). More recently Ng (1997) discusses the measurability of happiness, and Daniel Kahneman, Peter Wakker and Rakesh Sarin (1997) provide an axiomatic defence of experienced utility, and propose applications to economics. Our paper also borders on work in the psychology literature; see for example Edward Diener (1984), David Myers (1993), and William Pavot, Diener, Randall Colvin, and Edward Sandvik (1991).

Section I describes the main data source and the estimation strategy. This relies on a regression-adjusted measure of well-being in a particular year and country – the level not explained by individual personal characteristics. This residual macroeconomic well-being measure is the paper's focus.

In Section II we show, using a panel analysis of nations, that reported well-being is strongly correlated with inflation and unemployment. It should be emphasised that people are not asked whether they dislike inflation and unemployment. Instead, individuals are asked in surveys how happy they are with life, and the paper demonstrates that -- possibly unknown to them -- their en masse answers move systematically with their nation's level of joblessness and rate of price change.[5] Section III concludes.

I. Happiness Data and Empirical Strategy

Our main data source is the Euro-Barometer Survey Series. Partly the creation of Ronald Inglehart at the University of Michigan, it records happiness and life satisfaction information on 264,710 people living in twelve European countries over the period 1975 to 1991. A cross-section sample of Europeans is interviewed each year. One question asks *"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?"*. Another elicits answers to a "life satisfaction" question. This question, included in part because the word happy translates imprecisely across languages, is worded, *"On the whole, are you very satisfied, fairly satisfied, not very satisfied or not at all satisfied with the life you lead?"*. We concentrate on the life satisfaction data because they are available for a longer period of time – from 1975 to 1991 instead of just 1975-86. Unsurprisingly, happiness and life satisfaction are correlated (the correlation coefficient is 0.56 for the available period 1975-86), so a focus on life satisfaction may be sufficient. The working-paper version of this paper, available on request, presents extra results using European happiness statistics.

We also study happiness data on 26,668 individuals from the United States General Social Survey (1972-1994). There the happiness question reads *"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?"*. The question was asked in each of 23 years. There is no life-satisfaction question for the U.S. It would be ideal if the well-being questions' wordings were identical in the European and US cases, but they are not. However, most of the paper's conclusions rest upon cross-Europe results, where the wording of questions is the same. For a data set on Great Britain, in which, unusually, both happiness and life-satisfaction answers are available from the same individuals, David Blanchflower and Oswald (2000) have shown that estimated happiness and life-satisfaction equations have almost identical structures.

A. Estimation Strategy

We study a regression of the form

$$LIFE\ SATISFACTION_{it} = \alpha\ INFLATION_{it} + \beta\ UNEMPLOYMENT_{it} + \varepsilon_i + \delta_t + \mu_{it}$$

where *LIFE SATISFACTION* is the average life satisfaction in country *i* in year *t* that is not explained by personal characteristics, *UNEMPLOYMENT* is the unemployment rate in country *i* in year *t*, *INFLATION* is the rate of change of consumer prices in country *i* and year *t*, ε_i is a country fixed effect, δ_t is a time effect (a year fixed effect), and μ_{it} is an error term. Life satisfaction has no natural units. It is measured here by assigning integers 1-4 to people's answers: 1 (to *"not at all satisfied"*), 2 (to *"not very satisfied"*), 3 (to *"fairly satisfied"*) and 4 (to *"very satisfied"*). We experimented with other cardinalizations; the paper's findings were unaffected. The data on unemployment and

inflation are from the Organization for Economic Cooperation and Development (OECD). Some regressions also include a country-specific time trend.

A two-step methodology is employed. In the first stage, micro-econometric OLS life satisfaction regressions are estimated for each country in the sample. The mean residual life-satisfaction is calculated for each nation in each year, which gives 150 observations in a second-stage regression. These country-by-year unexplained life-satisfaction components then become the dependent variable in a second-stage regression of the form given in the equation above. Three-year moving averages of the explanatory variables are used; the moving averages are centered on year t-1. This smooths out some of the noise evident in the data (and, we found, produces succinct estimating equations while leaving the substantive conclusions unaffected when compared to equations with many lagged and autoregressive terms).

For three reasons, issues of simultaneity are ignored. First, it might be believed that 'happiness' does not itself mold the levels of inflation and unemployment. Second, our aim is primarily to document correlations in the data. Third, it is unclear what kind of variable could serve as a persuasive instrument for macroeconomic variables in a well-being regression equation. Nevertheless, future research may have to return to this issue.

The building blocks of the analysis are thus well-being regressions for each of the countries in our sample. These are similar to emerging micro-econometric work such as that of Blanchflower and Oswald (2000) who estimate the impact of personal characteristics on happiness responses for the US and the UK.[6]

Although coefficients in our regressions do not have a ready cardinal meaning, a number of personal characteristics are positively associated with reported well-being, and

are statistically significant, in every country in our sample. These characteristics include being employed, young or old (not middle aged), and belonging to a high-income quartile. The micro-econometric structure of well-being equations is similar across nations.

Table A1 in the appendix presents a pooled micro-econometric life satisfaction regression for Europe. This is an ordinary least squares regression; we checked that an ordered probit produces the same substantive conclusions. Greater family income increases the likelihood that a respondent reports a high level of well-being. This effect of income is monotonic and is reminiscent of the utility function of standard economics. The regression evidence is also consistent with the common-sense idea that unemployment is a major economic source of human distress (on psychiatric stress data see Andrew Clark and Oswald (1997)). Our working paper reports other well-being regressions.

The main data are as follows.

B. Data Definitions

LIFE SATISFACTION: The average of the residuals from a Life Satisfaction Ordinary Least Squares regression on personal characteristics. The residuals are averaged for each country and year in the sample (Mean=-0.004; Standard deviation=0.082).

UNEMPLOYMENT: The unemployment rate (three-year moving average) from the OECD-Centre for Economic Performance data set (1992). (Mean=0.086; Standard deviation=0.037).

INFLATION: The inflation rate (three-year moving average), as measured by the rate of change in consumer prices, from the OECD-Centre for Economic Performance data set (1992). (Mean=0.081; Standard deviation=0.057).

Throughout the paper, unemployment and inflation are measured as fractions. For example, an eight percent rate of inflation is entered in our data set as 0.08, and a nine percent unemployment rate is represented as 0.09.

II. The Inflation-Unemployment Trade-Off in Well-Being Equations

Regression (1) of Table 1 studies the dependence of life satisfaction on the unemployment rate and the rate of inflation. The specification includes time and country dummies. The coefficients from regression (1) in Table 1 imply that higher unemployment and higher inflation both decrease life satisfaction.

The effects of unemployment and inflation, which in column (1) of Table 1 have coefficients –2.8 and –1.2 respectively, are significantly different from zero at conventional levels of statistical significance. It is necessary to be clear about the units of measurement in Table 1. The numbers –2.8 and –1.2 represent the effect upon wellbeing (as cardinalized) of a one percentage point change in each of the two independent variables. As an example, consider the impact of an increase in the rate of unemployment from the mean of nine percent by one percentage point to ten percent. According to our estimate, this single-point rise in unemployment from 0.09 to 0.10 diminishes life satisfaction by 0.028 units. The number 0.028 is the product of 0.01 and 2.8. Consider instead an increase in the inflation rate from the mean of eight percent by

one percentage point to nine percent. This single-point rise in inflation from 0.09 to 0.10 leads to a 0.012 reduction in units of life satisfaction. The number 0.012 is the product of 0.01 and 1.2.

Given that the inflation and unemployment data are in fractions, these effects of unemployment and inflation are not small. Consider the consequences of a rise in unemployment of 0.04 (namely, four percentage points of joblessness, which is equal to the standard deviation in the sample). This produces a decline in well-being of 0.04 times -2.8, which is -0.11. In our cardinalization, people's levels of satisfaction are coded in four categories from 1 (not at all satisfied) up to 4 (very satisfied). Hence a movement of -0.11 is not a trivial event for a society. It is equivalent to shifting 11 percent of the population downwards from one life-satisfaction category to another. An alternative way to make the same point is to note that 0.11 slightly exceeds the standard deviation of life satisfaction in our panel of countries.

The implicit utility-constant trade-off between inflation and unemployment can now be calculated. We make the assumption that, over the relevant range, utility is linear (so that the margin is equal to the average). As in conventional economic theory, what is done in the paper is to measure the slope of indifference curves. This leads to a measure of the marginal rate of substitution between inflation and unemployment. It is useful to explain what such correlations are likely to mean within a conventional natural-rate-of-unemployment analytical framework. The estimation describes preferences themselves. Standard economic models suggest, of course, that there is no downward-sloping Phillips Curve in the long run. Knowledge of iso-utility contours is then of use to policy-makers primarily because it informs the choice of an optimal dis-inflationary path. Our

estimates, and more broadly this kind of methodology, can be viewed as aiding central bankers concerned with the choice of policy trajectories.

Regression (2) in Table 1 shows that unemployment and inflation enter strongly even if country-specific time trends are introduced into the equations. The coefficients on the two variables are negative and significantly different from zero at normal confidence levels. They are now more similar than in the first column of Table 1. However, equality of the two coefficients, in regression 2, can still be rejected statistically. Life satisfaction is therefore not captured exactly by a simple linear misery function defined on the sum of inflation and unemployment, $W=W(\pi+U)$. Unemployment has a larger weight.

Regressions (3) and (4) in Table 1 divide the sample into two time periods: before 1984 and after 1983. The coefficients keep their signs, although, as is to be expected, they are not now as well-defined. Degrees of freedom here are a source of potential concern; but this approach is primarily designed as a check on robustness. Column (5) adds into the equation a squared term in inflation -- to test if inflation is particularly bad at high levels -- but again the key result is left unaffected. If an additional squared term in unemployment is entered, its effect is negligible.

Table 2 presents further tests of the relationship between inflation, unemployment and well-being. Regression (6) in Table 2 controls for a lagged dependent variable. It finds that there is a little autoregression, with a lagged dependent variable coefficient of 0.3, but that life satisfaction data continue to be correlated with macroeconomic variables.

Regression (7) in Table 2 tests whether well-being depends on changes in the two macroeconomic variables. We use the growth in inflation (or unemployment) from one year to the next. There is some evidence that these changes matter. Both enter with the expected negative sign. Regression (8) in Table 2 shows that the inclusion of a lagged dependent variable reinforces these findings. Nevertheless, the underlying ideas remain the same.

It could be argued that the above calculations underestimate the cost of unemployment. The reason is that the first-stage regressions have already controlled for the personal cost of being unemployed. Somehow a way has to be found to measure the two unpleasant consequences of a rise in unemployment: some people lose their jobs while at the same time everyone in the economy becomes more fearful.

There is a way to take account of the extra first-stage cost of joblessness, namely, to work out the sum of the aggregate and personal effects of unemployment. It is best to think of it as asking what happens if unemployment in the economy rises by 1 percentage point. We can calculate from regression (2) that an increase in the unemployment rate of a percentage point (namely, 0.01) has a cost in the chosen well-being units equal to approximately 0.02 for the average citizen. This number might be viewed as capturing a 'fear of unemployment' effect for everyone. However, it is clear from our microeconomic data that the person who actually falls unemployed experiences a much larger cost. The loss from being unemployed is equal to 0.33 when measured in the same units. This number comes from the coefficient on being unemployed in a life-satisfaction micro regression, like the one in appendix Table A1, estimated with OLS to keep the units consistent.

The entire well-being cost of a 1 percentage-point increase in the unemployment rate is therefore given by the sum of two components. One component is the 0.33 multiplied by the one percent of the population who have been unlucky enough actually to become unemployed. This is 0.33 times 0.01, which is 0.0033. The second component, which is more akin to higher fear of unemployment for everyone in society, is 0.02. Combining the two, we have 0.0033+0.02=0.0233 as society's overall well-being cost of a rise in unemployment by one percentage point.

To put this differently, in column 2 of Table 1 the well-being cost of a 1 percentage-point increase in the unemployment rate equals the loss brought about by an extra 1.66 percentage-points of inflation. The reason is that 1.66=0.0233/0.014, where 0.0233 is the marginal effect of unemployment on well-being, and 0.014 is the marginal effect of inflation on well-being (where 0.014 is derived from 0.14 multiplied by 0.01). Hence 1.66 is the marginal rate of substitution between inflation and unemployment. Because this number is larger than unity, the well-known 'misery index' is not an accurate representation of the data.

A. Inflation, Unemployment and Happiness in the United States

Since there is no question on life satisfaction in the United States General Social Survey (1972-1994), it was not possible to include the US in the panel regressions. Using GSS happiness data we estimated an OLS happiness regression -- available upon request -- on personal characteristics for the U.S. and obtained the mean residuals for each year. The year-to-year changes in the "happiness residuals" were negatively correlated with the corresponding year-to-year changes in the so-called misery index. When viewed as two individual explanatory variables, the yearly changes in happiness were somewhat more

strongly associated with changes in the unemployment rate than with inflation. Necessarily, the US findings stem from a single time-series regression. The US results are consistent with, though a little less well-defined than, the European results.

III. Conclusions

The paper studies reported well-being data on a quarter of a million people across twelve European countries and the United States. We show that people appear to be happier when inflation and unemployment are low. Consistent with the standard macroeconomics textbook's assumption that there exists a social objective function $W(\pi, U)$, randomly sampled individuals mark systematically lower in well-being surveys when there is inflation or unemployment in their country. The rates of price change and joblessness affect reported satisfaction with life even after controlling for the personal characteristics of the respondents, country fixed-effects, year effects, country-specific time trends, and a lagged dependent variable. A function strongly reminiscent of the textbook $W(\pi, U)$ exists in the data.

A large literature in economics has tried to measure the losses from inflation. By examining the appropriate area under a money demand curve, Martin Bailey (1956) and Milton Friedman (1969) originally concluded that inflation has only small costs. Similarly, Fischer (1981) and Robert Lucas (1981) find the cost of inflation to be low, at 0.3 per cent and 0.5 per cent of national income, respectively, for a 10 per cent level of inflation. The numbers implied by our happiness-equation estimates seem consistent with larger welfare losses.

At the margin, unemployment depresses reported well-being more than does inflation. In a panel that controls for country fixed-effects, year effects and country-specific time trends, the estimates suggest that people would trade off a 1 percentage-point increase in the unemployment rate for a 1.7 percentage-point increase in the inflation rate. Hence, according to these findings, the famous misery index $W(\pi+U)$ under-weights the unhappiness caused by joblessness. Unemployment appears to be more costly than inflation.

Table 1: Life Satisfaction Equations for Europe 1975-91

	(1)	(2)	(3) Pre 84	(4) Post 83	(5)
Unemployment t	-2.8 (0.6)	-2.0 (0.6)	-0.4 (1.6)	-2.0 (1.1)	-2.1 (0.6)
Inflation t	-1.2 (0.3)	-1.4 (0.4)	-0.5 (0.7)	-2.0 (0.8)	-2.3 (0.9)
Inflation^2 t					3.5 (3.0)
Time Trends[1]	No	Yes	Yes	Yes	Yes
Country Dummies	Yes	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes	Yes
No. of Obs.	150	150	72	78	150
Adj. R^2	0.27	0.54	0.57	0.66	0.55

Notes: Standard errors are in parentheses. Time trends are country-specific. Three-year moving averages of the explanatory variables are used. This is a second-stage regression. It uses as a dependent variable the regression-corrected life satisfaction levels from a first-stage OLS regression of the general kind given in the Appendix.

Table 2: Checks on Life Satisfaction Equations for Europe 1975-91

	(6)	(7)	(8)
Life Satisfaction t-1	0.3 (0.1)		0.2 (0.1)
Unemployment t	-1.7 (0.7)	-2.1 (0.6)	-1.8 (0.7)
Inflation t	-0.7 (0.5)	-1.4 (0.4)	-0.8 (0.5)
ΔUnemployment t		-1.0 (0.9)	-0.1 (0.9)
ΔInflation t		-0.7 (0.4)	-0.5 (0.4)
Time Trends[1]	Yes	Yes	Yes
Country Dummies	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes
No. of Obs.	140	150	140
Adj. R^2	0.56	0.55	0.56

Notes: Standard errors are in parentheses. Time trends are country-specific. Three-year moving averages of the explanatory variables are used. This is a second-stage regression. It uses as a dependent variable the regression-corrected life satisfaction levels from a first-stage OLS regression of the general kind given in the Appendix.

Appendix

Table A1: OLS Life Satisfaction Micro-Equation for Europe 1975-91.

Dep Var: Reported Life Satisfaction		**Coefficient**	**Standard Error**
Unemployed		-0.33	7e-3
Self employed		0.04	5e-3
Male		-0.04	3e-3
Age		-0.02	1e-3
Age Squared		2e-4	6e-6
Education to age: 15-18 years		0.03	4e-3
≥ 19 years		0.06	4e-3
Marital Status: Married		0.08	4e-3
Divorced		-0.18	0.01
Separated		-0.23	0.01
Widowed		-0.10	0.01
Num. of children between 8 & 15 yrs:	1	-0.02	4e-3
	2	-0.03	0.01
	3	-0.06	0.01
Income Quartiles : Second		0.12	4e-3
Third		0.20	4e-3
Fourth (highest)		0.30	5e-3
Retired		0.05	6e-3
In school		0.04	7e-3
At home		0.03	5e-3

Notes: [1] Number of Observations=264,710. Adj. R^2=0.17. [2] The regression includes country and year dummies from 1975 to 1991. The country dummies (standard errors) are: Belgium 0.315 (0.006), Netherlands 0.540 (0.006), Germany 0.242 (0.006), Italy -0.087 (0.006), Luxembourg 0.469 (0.009), Denmark 0.694 (0.006), Ireland 0.356 (0.007), Britain 0.328 (0.006), Portugal –0.171 (0.008), Greece –0.146 (0.007) and Spain 0.124 (0.008). The base country is France. **[**3] The exact question is: *"On the whole, are you very satisfied, fairly satisfied, not very satisfied, or not at all satisfied with the life you lead?"* Answers were coded as follows: 1 to *"not at all satisfied"*, 2 to *"not very satisfied"*, 3 to *"fairly satisfied"* and 4 to *"very satisfied"*. Micro-econometric life-satisfaction equations are used as a first stage in the paper's analysis.

References

Bailey, Martin. "The Welfare Cost of Inflationary Finance." *Journal of Political Economy,* February 1956 (1), 64, pp. 64-93.

Barro, Robert J. and Gordon, David. "A Positive Theory of Monetary Policy in a Natural Rate Model." *Journal of Political Economy*, August 1983, 91(3), pp. 589-610.

Blanchard, Olivier and Fischer, Stanley. *Lectures on Macroeconomics*. Cambridge, MA: MIT Press, 1989.

Blanchflower, David and Oswald, Andrew J. "Well-being Over Time in Britain and the USA.", National Bureau of Economic Research Paper, forthcoming, February 2000.

Burda, Michael and Wyplosz, Charles. *Macroeconomics: A European Text*. Oxford: Oxford University Press, 1993.

Clark, Andrew and Oswald, Andrew J. "Unhappiness and Unemployment.", *Economic Journal*, September 1994, 104 (5), pp. 648-659.

Diener, Edward. "Subjective Well-Being." *Psychological Bulletin*, 1984, 93(3), pp.542-575.

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

Economist, The. "Stable Prices and Fast Growth: Just Say No", August 31st, 1996.

Fischer, Stanley. "Towards an Understanding of the Costs of Inflation: II.", *Carnegie-Rochester Conference Series on Public Policy*, 1981, 15, pp. 5-41.

Frank, Robert H. *Choosing the Right Pond*, New York: Oxford University Press, 1985.

Friedman, Milton. "The Optimum Quantity of Money." In *The Optimum Quantity of Money and Other Essays*, Chicago: Aldine, 1969.

Hall, Robert E. and Taylor, John. *Macroeconomics*, 5th edition, New York: Norton Press, 1997.

Inglehart, Ronald. *Culture Shift in Advanced Industrial Society*, Princeton: Princeton University Press, 1990.

Kahneman, Daniel; Wakker, Peter P.; and Sarin, Rakesh. "Back to Bentham? Explorations of Experienced Utility." *Quarterly Journal of Economics*, May 1997, 112(2), pp. 375-406.

Lucas, Robert E., Jr. "Discussion of: Stanley Fischer, 'Towards an Understanding of the Costs of Inflation: II'." *Carnegie-Rochester Conference Series on Public Policy*, 1981, 15, pp. 43-52.

Mankiw, N. Gregory. *Macroeconomics*, 3rd Edition, New York: Worth Publishers, 1997.

Morawetz, David and ten co-authors. "Income Distribution and Self-Rated Happiness: Some Empirical Evidence." *Economic Journal*, September 1977, 87 (3), pp. 511-22.

Myers, David. *The Pursuit of Happiness*, London: Aquarian, 1993.

Ng, Yew-Kwang. "Happiness Surveys: Some Comparability Issues and an Exploratory Survey Based on Just Perceivable Increments." *Social Indicators Research*, 1996, 38(1), pp. 1-27

Ng, Yew-Kwang. "A Case for Happiness, Cardinalism, and Interpersonal Comparability." *Economic Journal*, September 1997, 107(5), pp. 1848-1858.

Oswald, Andrew J. "Happiness and Economic Performance." *Economic Journal*, September 1997, 107(5), pp. 1815-1831.

Pavot, William; Diener, Edward; Colvin, C. Randall; Sandvik, Edward. "Further Validation of the Satisfaction with Life Scale: Evidence for the Cross-Method Convergence of Well-Being Measures." *Journal of Personality Assessment*, 1991, 57(1), pp. 149-161.

Shiller, Robert. "Why Do People Dislike Inflation?" In *Reducing Inflation: Motivation and Strategy*, Romer, Christine and Romer, David, eds., Chicago: University of Chicago Press, 1997.

Winkelmann, Liliana and Winkelmann, Rainer. "Why are the Unemployed so Unhappy?" *Economica*, 1998, 65(257), pp. 1-15.

---

[1] Rafael Di Tella, Harvard Business School, Morgan Hall, Soldiers Field, Boston, MA 02163, USA, Robert MacCulloch, ZEI, University of Bonn, Walter Flex Strasse 3, D-53113, Bonn, Germany, Andrew Oswald, Department of Economics, University of Warwick, Coventry CV4 7AL, England. For helpful discussions, we thank George Akerlof, Danny Blanchflower, Andrew Clark, Ben Friedman, Duncan Gallie, Sebastian Galiani, Ed Glaeser, Berndt Hayo, Daniel Kahneman, Guillermo Mondino, Steve Nickell, Julio Rotemberg, Hyun Shin, John Whalley, three referees, and seminar participants at

Oxford, Harvard Business School, and the NBER Behavioral Macro Conference in 1998. The third author is grateful to the Leverhulme Trust for research support.

[2] See, for example, Olivier Blanchard and Stanley Fischer (1989), Michael Burda and Charles Wyplosz (1993) and Robert Hall and John Taylor (1993). Early influential papers include Robert Barro and David Gordon (1983).

[3] Gregory Mankiw [1997] describes the question "How costly is inflation?" as one of the four major unsolved problems of macroeconomics.

[4] A recent contribution to this debate in the U.S. is Paul Krugman's piece "Stable Prices and Fast Growth: Just Say No", *The Economist,* August 31st, 1996.

[5] Our analysis complements the survey approach of, for example, Shiller (1996) who uses questions regarding inflation.

[6] Inglehart (1990) also documents the patterns in the micro data by looking at cross-tabulations.