

When does Inflation become too high? Thresholds in the Inflation–Happiness relationship

Alec Steeley, University of Oxford¹

Abstract

This paper documents a threshold relationship between inflation and national happiness using annual data from 81 countries between 2006 and 2024. I show that inflation negatively impacts happiness but only when inflation exceeds 5%. Below this level, I find no effect of inflation on a country’s overall wellbeing. This null result is primarily driven by inflation observations in the range of 2% to 5% which appear to have no impact on happiness. This finding departs from the existing literature, which typically assumes a linear effect of inflation on happiness. I estimate a dynamic panel threshold model that controls for unobserved heterogeneity across countries as well as time-varying macroeconomic indicators. I also show that national happiness scores display high levels of persistence in the dynamic model, a fact which previous panel studies of happiness have ignored. I argue that the lag of happiness, which is correlated with contemporaneous happiness and inflation, is thus an essential control variable. The threshold relationship between inflation and happiness is robust across alternative specifications, well-being measures, control sets and estimation methods, and it can be considered causal under strong but plausible assumptions.

¹I am enormously grateful to the individuals who have provided support on this project. For a year of incredibly valuable supervision, guidance and support, I am grateful to Vanessa Berenguer-Rico. For valuable comments and suggestions, I am grateful to Steve Bond and Nathan Barnard.

1 Introduction

Over the past three decades, many central banks around the world have adopted explicit inflation targets, most commonly aiming to maintain annual inflation at around 2%. These policies reflect the widespread belief that while high inflation is harmful to the economy, inflation in the low single digits is likely benign or beneficial. Yet, this raises a crucial question for policymakers: at what point does inflation become actively harmful to welfare? Does public well-being respond to every tick of the inflation rate, as often assumed, or does inflation only begin to matter once it crosses a critical threshold of public awareness? Despite the prevalence of inflation targeting, and the requirement of policymakers to make tradeoffs between inflation and other macroeconomic variables, the fundamental shape of the relationship between inflation and public wellbeing remains empirically unexplored.

This paper fills that gap by documenting a distinct threshold in the inflation–happiness relationship. Using a dynamic panel of 81 countries from 2005-2024, I find that inflation is only detrimental to national happiness once it exceeds approximately 5%. Below this level, its effect on wellbeing is statistically insignificant. This null result is driven by inflation in the 2% to 5% range, where its movements appear to go largely unnoticed by the public. This finding sharply contradicts the linear relationship assumed in prior research and suggests that the public’s aversion to inflation is not constant but activates only when inflation becomes high enough to be salient and newsworthy.

Methodologically, this paper breaks new ground in the happiness literature by employing a dynamic panel data approach. I argue that controlling lagged happiness removes reverse causality (the effect of past happiness on inflation) which is likely to have biased previous results in the macroeconomics of happiness literature. The lagged dependent variable complicates estimation since lagged happiness by construction correlates with country fixed-effects. To address this, I follow the two-step procedure of dynamic panel threshold estimation from Kremer et al. (2013), which builds upon earlier asymptotic theory developed by Hansen (1999) and Caner and Hansen (2004). I substitute the estimated threshold into the system

GMM estimator of Arellano and Bover (1995) and Blundell and Bond (1998), which I demonstrate performs well for slope estimation. The approach I take first estimates the optimal threshold via a grid search, finding a threshold that splits the sample into “low-inflation” and “high-inflation” regimes. It then estimates the effect of inflation on happiness in each regime while controlling for country fixed-effects, happiness dynamics and a variety of country-time varying macroeconomic factors. By adopting this technique, I aim to move beyond the static linear models of prior work and provide more reliable estimates of inflation’s non-linear relationship with happiness.

To my knowledge, this is the first paper to study happiness in a dynamic panel setting, and the first to model inflation’s relationship with happiness using a threshold. I also show that the relationship found in this paper, where inflation’s coefficient moves from insignificant to negative at a threshold of 5%, is robust to an extensive array of alternative specifications and estimation procedures, and that it passes a key falsification test. A major concern in linking inflation to happiness is that other factors—for example, economic growth, unemployment, or institutional quality—could confound the relationship. Moreover, controlling for such factors improperly can introduce the problem of “bad controls,” whereby including variables that are themselves affected by inflation (like GDP growth) may mask part of inflation’s true effect (Angrist and Pischke, 2009). To tackle these issues, I implement a novel lagged-controls approach: instead of contemporaneous values, I include lagged macroeconomic controls (such as last year’s GDP per capita growth, unemployment rate, etc.) in the happiness regression. The idea is that lagged controls do much of the work of contemporaneous controls in removing confounding factors, but are not themselves outcomes of inflation since last year’s growth or unemployment cannot be caused by this year’s inflation. I also account for the fact that even lagged controls may be bad controls due to inflation forecasting, by controlling for inflation’s 1-year ahead forecast. The lagged control models, with both actual and unexpected inflation, confirm the threshold relationship between inflation and happiness with roughly 5% remaining the point such that inflation’s coefficient becomes negative when it exceeds this level.

In lagged control models, where “bad controls” are not a concern, I employ machine-learning aided variable selection (LASSO) on a set of 2000 country-time variables to avoid researcher bias in choosing controls. Inflation’s relationship with happiness remains remarkably similar through these variations, including when 22 “good controls” selected by LASSO from a pool of thousands are included in the model.

Having used LASSO in an effort to remove sources of omitted variable bias, I also address potential reverse causality in the inflation–happiness relationship. While prior studies have attempted to study the effect of inflation on happiness, none have addressed the possibility that higher levels of happiness might have an effect on inflation. By studying the effect of the lags of happiness on inflation, I show that higher levels of past happiness are associated with higher levels of inflation, though only when inflation is already sufficiently high. This result likely reflects that happier economies spend and invest more, but that this only leads to inflationary pressure when the economy is at capacity. Since this reverse causality runs in the opposite direction of my primary finding, it emphasises the need to include lagged happiness as a control variable to eliminate this bias in the contemporaneous relationship between inflation and happiness. Additionally, to the extent that there is also a contemporaneous channel of causation from happiness to inflation, the results for the effect of inflation on happiness may be understated.

To provide further support for inflation’s threshold relationship with happiness, I also propose a falsification test: I verify that if one replaces inflation with another macro variable (unemployment) in the threshold model, no analogous threshold emerges for unemployment’s effect on happiness. This helps confirm that the detected non-linearity is unique to inflation rather than a spurious feature of the data or method. Overall, the exhaustive robustness exercises—including dynamic specifications, lagged controls, alternative well-being metrics, reverse causality tests and falsification tests—lend credence to the conclusion that there is a genuine threshold in the inflation–happiness relationship and that the result is not driven by endogeneity or model mis-specification. While I do not exploit exogenous variation in

inflation, I provide strong but plausible assumptions under which certain results can be interpreted as causal.

To complement the interpretation of results, I explore the mechanism driving the 5% threshold. I first show that conventional explanations are incomplete; the relationship cannot be fully explained by short-run real income effects, nor is it merely a psychological artifact of official inflation targets, as the threshold exists even in countries without them. The evidence instead points towards a ‘fear of inflation’ effect, originally speculated by Di Tella et al. (2001). To provide evidence for this channel, I use Google Trends data to proxy public attention and concern. Public search interest for the term ‘inflation’ displays the same non-linear pattern as happiness: Google searches remain low and unrelated to inflation when the rate is below 5%, but they begin to rise significantly once that threshold is exceeded. This strongly supports the conclusion that the 5% threshold represents the point at which inflation becomes newsworthy and psychologically salient.

This paper is related to several strands of existing literature. First, it contributes to the economics of happiness literature that examines how macroeconomic variables affect well-being. Starting with Di Tella et al. (2001, 2003), researchers have found that inflation and unemployment significantly influence reported happiness. My findings reaffirm that inflation is an important determinant of well-being, but they depart from prior results by showing that the impact of inflation is not constant across all levels of inflation – a fact that earlier linear models could not capture. By uncovering this threshold, the paper complements studies like Blanchflower (2007); Blanchflower et al. (2014) and Ruprah and Luengas (2011), which examine the trade-off between inflation and unemployment and generally conclude that the cost (in happiness) of a one-point increase in unemployment significantly exceeds that of a one-point increase in inflation. This paper’s results suggest that such trade-off calculations may themselves depend on the inflation regime: when inflation is low, additional inflation has no impact on happiness meaning that policymakers may only have to concern themselves with unemployment rates in those cases. Yet, these results suggest that when inflation rises above

5%, the trade-off between inflation and unemployment becomes meaningful and policymakers should proceed with caution.

Second, the paper contributes to literature debating optimal inflation targets and the primacy of inflation as a macroeconomic policy objective. While the New Keynesian framework, notably articulated by Clarida et al. (1999), provides theoretical justification for inflation targeting by demonstrating that stable inflation anchors expectations and stabilises output, economists continue to debate the optimal target level—commonly at 0%, 2%, or 4%. A near-consensus around a 2% target has emerged in practice, justified by measurement biases, downward wage rigidity, and the need to maintain policy room above the zero lower bound (Bernanke et al., 1999; Blanchard et al., 2010). However, prominent voices such as Blanchard et al. (2010) have argued that raising inflation targets to around 4% could provide additional monetary policy flexibility during economic downturns, outweighing potential costs related to slightly higher inflation. The present paper suggests that a key policy focus should be on keeping inflation below the 5% threshold. When inflation is below this level then targeting other macroeconomic variables such as growth and unemployment, which also matter greatly for a country’s wellbeing, can take precedence.

Third, this paper intersects with literature on inflation and macroeconomic performance, particularly research on inflation thresholds in economic growth. Studies like Khan and Senhadji (2001), Bick (2010) and Kremer et al. (2013) have demonstrated that inflation’s effect on GDP growth is nonlinear – low to moderate inflation has little adverse effect and may even coincide with faster growth, whereas inflation beyond a certain point (often around 4–5% for developed economies) significantly hampers growth. Similarly, Glawe and Wagner (2024) show that inflation has an inequality increasing effect only when inflation exceeds 6%. The parallel between those findings and our happiness results is striking: it appears that inflation in the low single digits is generally benign or beneficial, whether one’s metric is output or subjective well-being, but once inflation climbs into the high single or double digits, the costs – in terms of both economic activity and happiness – become pronounced. By drawing this

parallel, the paper connects the subjective well-being perspective with traditional economic outcomes, suggesting a consistency in what constitutes “healthy” inflation.

Fourth, this research builds upon literature in econometrics concerning the estimation of unknown thresholds in regression models. Hansen (2000) showed that consistent thresholds could be estimated by performing a grid search over possible thresholds and selecting the threshold which minimises the sum of squared residuals, effectively choosing the threshold which best fits the model. Hansen also showed that valid confidence intervals around thresholds could be formed under mild assumptions. This theory for estimating thresholds was extended to static panel data (Hansen, 1999), and 2-stage-least-squares IV-estimation Caner and Hansen (2004). Kremer et al. (2013) combined these two models to create a procedure for threshold estimation in a dynamic panel model. The present paper relies heavily on the procedure of Kremer et al. (2013) for threshold estimation but deviates in its method for estimating slopes for inflation in the second-stage. Since the inflation threshold can be treated as known in the second-stage², I adopt a system GMM estimator for slope estimation which, given the properties of the data, is far more efficient than the standard IV-GMM estimator used in Kremer et al. (2013).

Finally, this research addresses methodological issues in the happiness literature. It has been shown that many of the literature’s results are contingent on assumptions about the linearity and comparability of happiness scores (Oswald, 2008; Bond and Lang, 2019). Though existing literature has provided evidence in favour of the reliability of happiness scores (Ferrer-i-Carbonell and Frijters, 2004; Kaiser and Oswald, 2022; Samuelsson et al., 2023), I still address relevant concerns. First, removing country-fixed effects addresses the possibility of different interpretations or translations of happiness questions across countries. The introduction of a lagged dependent variable further addresses comparability since it effectively controls a country’s initial level of happiness, meaning that inference about inflation is made by comparing countries with similar initial happiness scores. In terms of linearity, I

²Estimation error on the threshold is asymptotically negligible in slope estimation, see Hansen (1999)

confront the possibility that inflation’s non-linear relationship with happiness is driven by non-linearity in the reporting of happiness scores, rather than non-linearity in true happiness, see Oswald (2008). The falsification test I conduct showing that unemployment’s relationship with happiness displays no such non-linearities provides evidence against non-linear reporting. As an additional check on methodology, I compare results across an alternative measure of happiness: affect. In contrast to some studies which have found their results to be contingent on how happiness is measured (Kahneman and Deaton, 2010), I find consistency across life evaluation and affect measures.

This paper proceeds by introducing the data sources and key variables in section 2, where I describe the happiness measures and present descriptive statistics of the data. Section 3 lays out the econometric methodology. Section 4 presents results about the inflation–happiness relationship. Section 5 reports an array of robustness checks and section 6 explores possible mechanisms underlying the empirical relationship. Section 7 concludes. Through this structure, the paper builds a case that the relationship between inflation and happiness is best understood as flat or weakly positive up to 5%, and negative above it.

2 Data

This section describes the construction of the primary dataset and provides descriptive statistics.

2.1 Happiness Data

Data on happiness is taken from the Gallup world poll. The Gallup world poll is a nationally representative survey which has been conducted in around 130 countries each year since 2005. The typical annual sample for each country is 1,000 people. Different countries and individuals are sampled at different points in the year. The measure of happiness used is called the Cantril ladder. This asks respondents the following question:

“Please imagine a ladder with steps numbered from 0 at the bottom to 10 at the top. The top of the ladder represents the best possible life for you, and the bottom of the ladder represents the worst possible life for you.

On which step of the ladder would you say you personally feel you stand at this time?”

The individual 0 to 10 responses have been averaged for each country in each year from 2005 to 2023, by the World Happiness Report (Helliwell et al., eds, 2024). This gives the average happiness of a representative sample of 1,000 individuals for each country-year combination. For the remainder of the paper, I use the word happiness to mean average happiness as defined in this way.

Happiness is treated as a continuous variable. This implicitly assumes that the happiness scale is linear and comparable across individuals. That is, a 7 out of 10 for individual i means the same thing as for individual j , and moving from a 5 to a 6 in reported happiness requires the same increase in latent happiness as moving from an 8 to a 9. These are strong assumptions, but the linearity assumption is now supported by a range of external evidence (Ferrer-i-Carbonell and Frijters, 2004; Oswald, 2008; Kaiser and Oswald, 2022; Samuelsson et al., 2023), and I confront the comparability assumption by employing country fixed-effects and controlling for lagged happiness. These assumptions are also revisited in the falsification test in section 5.

Figure 1 displays the time-series of happiness for 4 countries which roughly sit in the 4 quartiles of the happiness distribution. An immediate finding of this graph is that happiness varies much more across countries than within-countries. This is consistent with a large amount of recent literature that indeed finds large cross-country differences in happiness which may reflect the large differences in material living standards across the world.³ These cross-country happiness differences will become an important justification for the use of the system GMM procedure for slope estimation in later sections.

Using the fisher-type panel unit-root test proposed by Maddala and Wu (1999), I strongly

³For an overview of happiness trends across countries, see Our World in Data

reject the null hypothesis that happiness is non-stationary for all countries.⁴




Figure 1: Happiness over time in selected countries

2.2 Macroeconomic Variables

The cross-country happiness panel is merged with country-time varying macroeconomic data. Inflation data is taken from the World Bank’s Global Database on Inflation which was compiled by Ha et al. (2023). The chosen measure of inflation is headline consumer price index (CPI) inflation. Data on unemployment, growth and investment are taken from the international monetary fund (IMF) world economic indicators while data on openness (exports plus imports) and government expenditure are taken from the World Bank’s development indicators. Investment, government expenditure and openness are expressed as percentages of gross domestic product (GDP). For a full list of the variables used, their exact definitions

⁴I also conduct dickey-fuller tests on individual happiness time-series for each country, and find that the vast majority of countries reject the null hypothesis of non-stationarity despite the short time series.

and their data sources, see Appendix A. Following previous literature, I take the natural logarithm of investment, government expenditure and openness which makes the observations on these variables approximately normally distributed. For comparability with previous literature on the relationship between inflation, unemployment and happiness, I use the levels of inflation and unemployment. To avoid very large instances of inflation dictating results, I remove a small number of observations where inflation exceeds 50%⁵. Like happiness, all of the macroeconomic variables were found to be stationary using fisher-type panel unit root tests.

2.3 Final Data Set

Table 1: Summary Statistics of key variables

	Mean	SD	Min	Max	N
Happiness	6.04	0.91	3.66	7.97	1197
Positive Affect	0.68	0.10	0.35	0.88	1189
Inflation	4.03	4.31	-4.48	48.57	1197
Unemployment	8.47	5.68	0.70	35.30	1197
log(Investment)	3.14	0.24	2.19	4.22	1197
log(Government Expenditure)	2.79	0.28	1.97	3.73	1197
log(Open)	4.41	0.52	3.10	6.08	1197
Growth	1.99	4.07	-16.99	23.37	1104

After merging the happiness and macroeconomic data, the primary data set comprises an unbalanced panel of 81 countries observed between 2005 and 2023, totalling just over 1000 observations. ⁶ Table 1 presents the summary statistics of key variables.

⁵This removes around 1% of all observations.

⁶The cross-country panel contains a number of gaps. If improperly dealt with, this can lead to biased estimates on the lagged dependent variable. To resolve this concern, without losing too much data, I treat countries with gaps as separate countries. For example, if Argentina is observed between 2007 and 2014, as well as 2016 to 2023, it is treated as 2 separate groups with 8 time periods for each.

3 Methodology

This section outlines the econometric method used to model the threshold relationship between inflation and happiness. The model takes the form

$$Y_{it} = \alpha Y_{it-1} + \beta_1 \pi_{it} \mathbb{I}(\pi_{it} \leq \gamma) + \beta_2 \pi_{it} \mathbb{I}(\pi_{it} > \gamma) + \delta' x_{it} + \mu_i + \varepsilon_{it}, \quad (1)$$

where Y_{it} denotes the average happiness level of country i at time t . Y_{it-1} thus denotes the lagged dependent variable which captures happiness persistence through habit formation and other mechanisms. π_{it} denotes the inflation rate in country i at time t . γ denotes the inflation threshold which is the point at which the inflation–happiness relationship changes slope. x_{it} is a vector of country-time varying control variables. μ_i are time-invariant country fixed-effects which, among other things, capture differences in the translation or interpretation of happiness questions across countries. We let $\varepsilon_{it} \stackrel{iid}{\sim} (0, \sigma^2)$.

The parameters to be estimated are α , β_1 , β_2 , δ and γ . μ_i are eliminated by a transformation. The focus of this paper is the inflation–happiness relationship and thus the vast majority of discussion will be on the estimates of β_1 , β_2 and γ .

The estimation procedure broadly follows Hansen (1999), Caner and Hansen (2004) and Kremer et al. (2013). In short, this procedure performs a grid search over possible values of the threshold γ , and chooses the γ which best fits the model, based on minimising the sum of squared residuals. Appealing to the asymptotic theory developed by Caner and Hansen (2004), we can also form confidence intervals around the threshold estimate. After the threshold is estimated, it is plugged in to a system GMM estimator which is used to estimate the slope coefficients α , β_1 , β_2 , and δ .

3.1 Threshold Estimation

This section describes the estimation procedure for γ . The threshold estimation procedure needs to 1) transform the model to eliminate fixed-effects, 2) avoid endogeneity between

the lagged dependent variable and the transformed errors, 3) avoid serial correlation in the transformed errors, such that the asymptotic theory developed by Caner and Hansen (2004) can be applied to the threshold estimate $\hat{\gamma}$.

The first stage of the estimation procedure involves the forward orthogonal deviations transformation to eliminate fixed-effects. As demonstrated by Arellano and Bover (1995), this transformation eliminates fixed-effects and avoids serial correlation in the transformed errors. With no serial correlation, it can be argued that the transformed model satisfies the assumptions of Caner and Hansen (2004) and can be treated as a cross-section.

The forward-orthogonal deviation transformation involves subtracting the mean of all future periods, from the current period⁷

$$Y_{it}^* = \sqrt{\frac{T-t+1}{T-t+2}} \left[Y_{it-1} - \frac{1}{T-t-1} (Y_{it} + \dots + Y_{iT}) \right]. \quad (2)$$

Applying the forward-orthogonal deviations transformation to all variables gives the transformed model:

$$Y_{it}^* = \alpha Y_{it-1}^* + \beta_1 [\pi_{it} \mathbb{I}(\pi_{it} \leq \gamma)]^* + \beta_2 [\pi_{it} \mathbb{I}(\pi_{it} > \gamma)]^* + \delta' x_{it}^* + \varepsilon_{it}^*, \quad (3)$$

where inflation is multiplied by its indicator function before the transformation, such that the coefficients of the transformed model are identical to the coefficients of the untransformed model in equation 1. Transforming the model this way gives

$$\text{Var}[\varepsilon_i^*] = \begin{bmatrix} \sigma^2 & 0 & 0 & \dots & 0 \\ 0 & \sigma^2 & 0 & \dots & 0 \\ 0 & 0 & \sigma^2 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \sigma^2 \end{bmatrix}_{(T-1) \times (T-1)}, \quad \varepsilon_i^* = \begin{bmatrix} \varepsilon_{i2}^* \\ \varepsilon_{i3}^* \\ \vdots \\ \varepsilon_{iT}^* \end{bmatrix}.$$

⁷In this paper, I shift the notation by one period (see eq. 2). This makes the set of valid instruments, Y_{it-2}, \dots, Y_{i1} , consistent across forward orthogonal and first-difference transformed equations

This transformation ensures no serial correlation in the transformed errors, as required by Assumptions 1.2 and 1.3 in Caner and Hansen (2004, pp. 819). Following Caner and Hansen (2004) and Kremer et al. (2013), we now perform a version of 2-stage least squares estimation on the transformed model, equation (3). In the first stage, we regress the transformed lagged dependent variable Y_{it-1}^* on an instrument Y_{it-2} and the exogenous variables x_{it}^* . Since the threshold is determined in the second stage, the functional form of the estimated first stage regression can be either a linear regression or a threshold model, (see Caner and Hansen (2004)). Since, in our data, there is no reason to suspect a threshold in the first stage regression, and indeed no evidence of one, I use linear regression in the first stage. As a result, the first stage regression takes the form

$$Y_{it-1}^* = \theta_1 Y_{it-2} + \theta_2 \pi_{it}^* + \theta_3' x_{it}^* + u_{it}. \quad (4)$$

To preserve the sample size of the unbalanced panel, I use only the second lag of happiness, which is available for all panels, as an instrument in the two-stage threshold estimation procedure. Kremer et al. (2013) have shown in monte-carlo simulations that extending the instrument set makes little difference to the threshold estimate. The coefficients of the model are estimated by OLS and the fitted values of the lagged dependent variables are calculated as

$$\hat{Y}_{it-1}^* = \hat{\theta}_1 Y_{it-2} + \hat{\theta}_2 \pi_{it} + \hat{\theta}_3' x_{it}^*. \quad (5)$$

The fitted lagged dependent variable from the first-stage \hat{Y}_{it-1}^* is then plugged in to the second-stage regression,

$$Y_{it}^* = \alpha \hat{Y}_{it-1}^* + \beta_1 [\pi_{it} \mathbb{I}(\pi_{it} \leq \gamma)]^* + \beta_2 [\pi_{it} \mathbb{I}(\pi_{it} > \gamma)]^* + \delta' x_{it}^* + v_{it}^*. \quad (6)$$

This model is estimated for 140 possible thresholds for inflation, from -1 to 10, with 0.05 increments. The threshold is estimated by minimising the sum of squares of the residuals of

the second-stage. That is

$$\hat{\gamma} = \arg \min_{\gamma \in \Gamma} S(\gamma), \quad \text{with } S(\gamma) = \left(\sum_{i=1}^n \sum_{t=1}^T \hat{v}_{it}^*(\gamma) \right). \quad (7)$$

Γ is the set of possible thresholds.

Valid confidence intervals around the threshold can be formed by constructing a likelihood ratio between the estimated threshold $\hat{\gamma}$ and other thresholds from the grid search. The likelihood ratio statistic for the null hypothesis that $\gamma = \gamma_0$ is

$$LR(\gamma_0) = N \frac{S(\gamma_0) - S(\hat{\gamma})}{S(\hat{\gamma})}.$$

where N is the total number of observations in the transformed model. The distribution of $LR(\gamma)$ under the null hypothesis is non-standard and has the inverse:

$$c(\alpha) = -2 \log(1 - \sqrt{1 - \alpha}),$$

where α is the chosen significance level. A 95% confidence interval for γ is thus the set of values of γ such that $LR(\gamma) \leq c(0.05)$.

3.2 Slope Estimation

With the threshold parameter estimated, $\hat{\gamma}$ can be plugged back in to the model, and the slope parameters β_1 and β_2 can be estimated by 2-stage-least-squares or GMM. Caner and Hansen (2004) proved that the threshold estimate is of second order asymptotic importance to the estimation of the slope parameters, and so we can treat both estimated and imposed thresholds as fixed when estimating the slope parameters. A natural choice of estimator, therefore, is the first-difference GMM estimator proposed by Arellano and Bond (1991). However, the Arellano-Bond estimator often performs poorly in finite samples when Y_{it} is highly persistent, or when the variance of the fixed-effect across panels is high (Blundell and

Bond, 1998), both of which are the case in this paper’s dataset. Therefore, I use the system GMM estimator proposed by Arellano and Bover (1995) and fully developed by Blundell and Bond (1998) to estimate the slope parameters of the model. I justify this choice of estimator further in Appendix B.

Since GMM estimation preserves sample size when using longer lags of happiness as instruments, I extend the set of lagged instruments to increase efficiency of the slope estimates. As a balance between efficiency and bias, I use half of the available lagged levels of happiness as instruments for the transformed equation ($Y_{it-2}, \dots, Y_{it-8}$ as instruments for Y_{it}^*). For panels with fewer than 8 instruments available, all available lags are used as instruments. For the levels equation, I use only the first lag as an instrument since higher order lags are redundant (Bond, 2002). Varying these lag lengths does not change the qualitative results of section 4 below. Throughout this paper, results reported as ‘System GMM’ implement two-step GMM estimation with a finite-sample correction for the two-step covariance matrix (Windmeijer, 2005). Any reported standard errors are robust to heteroskedasticity and clustering by country.

4 Results

This section presents empirical results about the relationship between inflation and happiness. In Section 4.1, I estimate a baseline dynamic panel threshold model with a small number of macroeconomic control variables. Section 4.2 addresses problems of omitted variable bias and bad controls by using lagged controls. Section 4.3 studies reverse causality in the estimates. Section 4.4 provides assumptions under which the results can be considered causal estimates.

4.1 Estimated thresholds

In column (1) and (2) of Table 2, I estimate a baseline dynamic panel threshold model of inflation and happiness. I estimate the model on the entire sample, with 5 macroeconomic

variables included as controls. In column (1), I estimate a threshold of 4.9%, with a 95% confidence interval between 4.15% and 6.05%. For the slope estimates, I find that inflation has a positive but insignificant relationship with happiness below the threshold and a negative and significant relationship with happiness above the threshold. The insignificant coefficient on below threshold inflation in Table 2 suggests that inflation goes relatively unnoticed until it reaches a sufficiently high level, roughly 5% in this case, at which point additional inflation reduces happiness. Results for other macroeconomic variables are in line with expectations: growth and government expenditure are positively associated with happiness, while unemployment is negative. Investment and openness (exports plus imports) are insignificant.

This baseline result is strikingly similar to the dynamic panel threshold model of inflation and growth which estimated a threshold of around 4% for developed countries (Kremer et al., 2013). Like its effect on happiness, inflation's effect on growth below the threshold was positive but insignificant and above the threshold was negative and significant. Since the results for the inflation–happiness relationship hold growth constant, we cannot say that growth is the mechanism by which this threshold result occurs. Instead, the similarity of these results may indicate that for both growth and happiness separately, low inflation is relatively inconsequential, while high inflation is costly.

Column (2) adds time dummies to the specification, to control for the happiness and inflation effects of any time-varying shocks which have been common to all countries, such as the financial crisis or the COVID-19 pandemic. I do not report their coefficients but most of these dummies are statistically insignificant, likely because the already included macroeconomic indicators such as unemployment and growth proxy the time-varying shocks. Regardless, the estimated relationship between inflation and happiness remains similar in this specification, with a threshold estimate of just under 5%. Inflation's coefficient below the threshold remains positive and insignificant while its coefficient above the threshold remains negative and highly significant.

In column (3), I include a control for inflation’s volatility, owing to the possibility that happiness is determined by inflation volatility rather than inflation’s level⁸. I find that while the standard deviation of inflation is negatively associated with happiness, including inflation volatility as a control variable makes little difference to the relationship between inflation levels and happiness, which continues to display a threshold at inflation rates of just under 5%.

Finally, in column (4) I explore inflation thresholds in an alternative measure of happiness, denoted as *Affect*. While the Cantril ladder measure of happiness used throughout the rest of this paper measures an individual’s overall satisfaction and happiness with their life as a whole, *affect* measures an individual’s daily emotional state.⁹

In some literature, using different measures of happiness can affect conclusions about research questions related to wellbeing. However, there are no strong reasons to think that the inflation–happiness relationship should be sensitive to the way happiness is measured. Indeed, the results in column (4) confirm that the relationship is similar when happiness is measured by life evaluation scores or by day-to-day emotions. In either case, inflation has a threshold at around 5%, with insignificance below the threshold and a negative slope above it. Due to the wide confidence interval around inflation’s threshold, we should be cautious to interpret this result too strongly, though it is consistent with our earlier estimates.¹⁰ Recalling that *Affect* is bounded between -1 and 1, the magnitude of inflation’s effect on happiness is roughly comparable to the results of column (1).

Given their consistency across specifications, we can begin to interpret the coefficients from these estimates in an economically meaningful way. Up to 5%, an increase in the inflation

⁸Inflation’s volatility is measured separately for each country and time-period as the standard deviation of the inflation rate for that country over the past 3-years

⁹To construct *affect*, I take the average scores from 3 binary response questions relating to positive emotions (smiling, enjoyment and learning) and subtract the average scores from 3 binary response questions relating to negative emotions (worry, sadness and anger). The scores for each country thus lie between 1 and -1 with positive values meaning that positive emotions outweigh negative emotions and vice versa. For a complete definition of *Affect*, see Appendix A.

¹⁰The imprecise threshold estimate likely reflects that *Affect* is constructed by combining 6 separate variables and thus is inherently more noisy. Data on the original 6 variables is unfortunately not publicly available and so I cannot analyse them separately.

rate appears to have no statistically detectable effect on a country's average happiness. However, when inflation moves above this range, this is associated with a significant decrease in happiness. For instance, an increase in the inflation rate from 5% to 15% leads to a roughly 0.167 drop in a country's average happiness, when happiness is measured on a 0 to 10 scale. The same increase in unemployment rates is associated with around twice the drop in happiness.

Table 2: Estimated Thresholds

	(1)	(2)	(3)	(4)
	Happiness	Happiness	Happiness	Affect
Y_{it-1}	0.610*** (0.0627)	0.6304*** (0.0814)	0.494*** (0.0802)	0.550*** (0.0558)
$\hat{\gamma}$	4.9 [4.15 - 6.05]	4.85 [4.6 - 5.7]	4.8 [3.45 - 5.1]	5.65 [-1 - 10]
$\pi_{it}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$	0.00682 (0.00738)	0.0068 (0.0106)	-0.000560 (0.0109)	0.000276 (0.00167)
$\pi_{it}\mathbb{I}(\pi_{it} > \hat{\gamma})$	-0.0141*** (0.00356)	-0.0165*** (0.0043)	-0.0167*** (0.00578)	-0.00194*** (0.000600)
Unemployment $_{it}$	-0.0284*** (0.00769)	-0.0263*** (0.0081)	-0.0355*** (0.0103)	-0.00429*** (0.00160)
$\log(\text{Govt Exp})_{it}$	0.448*** (0.126)	0.4108*** (0.1472)	0.625*** (0.172)	0.00875 (0.0193)
$\log(\text{Investment})_{it}$	-0.0280 (0.0968)	-0.0201 (0.0848)	-0.0328 (0.115)	0.0190 (0.0206)
$\log(\text{Open})_{it}$	-0.0140 (0.0555)	0.0143 (0.0608)	0.0246 (0.0661)	-0.0104 (0.00936)
Growth $_{it}$	0.0111*** (0.00300)	0.0138*** (0.0042)	0.00972** (0.00378)	0.00137*** (0.000504)
s.d. $(\pi)_{it}$			-0.0349*** (0.0106)	
Observations	1108	1154	891	1097
Time dummies	No	Yes	No	No

Standard errors adjusted for clustering by country in parentheses. GMM estimates are two-step GMM with corrected and robust standard errors. Government expenditure, investment and openness are expressed percentages of GDP. Estimates include a constant term which is not reported.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.2 Bad Controls and Omitted Variable Bias

Having observed a consistent pattern in the cross-country-time relationship between inflation and happiness, I now address various issues which hold back the results from having a causal interpretation.

Firstly, inflation is very likely correlated with a large number of variables which influence happiness. This motivates the use of macroeconomic control variables, such as those which were included in previous sections. However, these control variables are likely insufficient in removing all sources of omitted variable bias in the results presented so far. Additionally, including control variables in our regression introduces a new problem since many of them may be outcomes of inflation and thus “bad controls” (Angrist and Pischke, 2008). For example, high levels of inflation may cause investment to fall and this fall in investment may then cause a drop in happiness. If this is true, then holding investment constant in the happiness regression, as was done in the previous section, may lead to inflation’s negative effect on happiness being understated since lower investment is one of the channels through which high inflation leads to lower happiness. As far as I am aware, the possibility of bad controls has thus far been ignored in the empirical inflation–happiness or inflation–growth literature.

In panel data, one promising avenue for resolving the issue of bad controls involves lagging control variables. The logic behind this method is that an outcome cannot precede its cause. That is, inflation today cannot *cause* investment last year. However, this is not quite true since economic agents anticipate future inflation rates using forecasting. These expectations may influence their decisions before the inflation rates are actually observed, creating a causal relationship between future forecasted inflation and past values of other macroeconomic variables. These anticipation effects mean that lagging control variables alone does not remove bad controls.

In order to create a set of good control variables in a regression of happiness on inflation, one must eliminate these anticipation effects. To do so, I apply two novel methods. First,

I obtain data on inflation’s one-year-ahead forecasted value in each period and use this to control for inflation’s forecast. What this means, is that any variation in inflation in the regression is coming from the unanticipated components of inflation, rather than the forecasted components of inflation. I can then freely include lagged macroeconomic variables in the regression, knowing that they will not be caused by contemporaneous inflation, after controlling for inflation’s forecasts. The lagged control variables aim to proxy their contemporaneous counterparts without themselves being bad controls. As a check on this method, I also construct a variable called “unexpected inflation,” defined as the difference between actual inflation and one-year-ahead inflation forecasts. By definition, unexpected inflation in the current period is unknown in the previous period, eliminating the risk of reverse causality involving lagged controls. I thus use unexpected inflation as an explanatory regressor, knowing that this should generate similar results to the cases where inflation forecasts were controlled for.

In summary, employing both lagged macroeconomic controls and controlling (or removing) inflation forecasts alleviates concerns about bad controls. At the same time, the lagged controls still achieve the intended purpose of the contemporaneous controls in removing confounding factors from the estimated inflation–happiness relationship. With these insights, this section proceeds to estimate dynamic panel threshold models, incorporating either inflation or unexpected inflation as explanatory variables along with lagged macroeconomic controls. I first simply lag the 5 macroeconomic control variables used in previous sections, to check that the previous results hold when bad controls are removed. Subsequently, I exploit the fact that lagging control variables overcomes the problem of bad controls, and use LASSO to select a much larger set of control variables. While even the regressions with a very large number of good controls do not provide clean causal estimates of inflation’s effect on happiness, I address a number of relevant concerns to the methodology and provide identification assumptions in section 4.4

I first test whether the conclusions from section 4.1 hold when control variables are lagged

Table 3: Lagged Controls

	(1)	(2)	(3)	(4)
	Happiness	Happiness	Happiness	Happiness
Y_{it-1}	0.684*** (0.0624)	0.662*** (0.0705)	0.416*** (0.0515)	0.425*** (0.0520)
$\hat{\gamma}$	5.05 [4.65 - 5.65]	5.35 [-1 - 10]	4.75 [0.5 - 9.7]	4.7 [-1 - 10]
$\pi_{it}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$	0.0146 (0.0105)		0.0143* (0.00759)	
$\pi_{it}\mathbb{I}(\pi_{it} > \hat{\gamma})$	-0.0109*** (0.00408)		-0.00869** (0.00361)	
$\pi_{it}^{Unexpected}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$		0.0166* (0.00897)		0.00603 (0.00507)
$\pi_{it}^{Unexpected}\mathbb{I}(\pi_{it} > \hat{\gamma})$		-0.0125** (0.00515)		-0.0103*** (0.00352)
$\pi_{it}^{Forecast}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$	-0.0264*** (0.00946)		-0.00787 (0.00713)	
$\pi_{it}^{Forecast}\mathbb{I}(\pi_{it} > \hat{\gamma})$	-0.0121** (0.00603)		0.00285 (0.00572)	
Observations	1025	1025	1656	1656
Variable Selection	Imposed	Imposed	LASSO-CV	LASSO-CV
No. of Lagged Controls	5	5	22	22

Corrected and robust standard errors from two-step GMM in parentheses. 95% confidence intervals in square brackets. Estimates include a constant term which is not reported. Columns 1 and 2 include controls for the 1st lag of Unemployment, Govt Expenditure, Investment, Openness and Growth. Columns 3 and 4 include 22 lagged control variables selected by LASSO cross-validation prior to threshold and slope estimation

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

by one period. Columns (1) and (2) in Table 3 study regressions of happiness on inflation, with one-period lags of the macroeconomic control variables from previous regressions. In column (1), I control for inflation's forecasts, both above and below a threshold. In column (2), I instead use unexpected inflation as the explanatory regressor. The counterpart for these two models, where controls were not lagged and inflation's forecast was not controlled or removed, is column (1) in Table 2. Lagging the control variables leaves results relatively unchanged. We continue to see the inflation–happiness relationship switch from positive and insignificant to negative and significant, and the estimated threshold is still around 5%. The relationship between unexpected inflation and happiness in column (2) also displays this pattern. These results indicate that while the 5 macroeconomic control variables included in previous sections may have been bad controls, that fact does not matter much for the results. After correcting for the problem of bad controls, the conclusion that inflation reduces happiness when it exceeds 5% remains. The similarity between these results and the results with contemporaneous controls provides evidence that the lagged controls are indeed proxying for the contemporaneous controls. Indeed, the coefficients on the lagged controls (not reported for brevity) are similar to the coefficients on their contemporaneous counterparts.

4.2.1 Lagged controls selected by LASSO

Having demonstrated that *bad* controls can be proxied by *good* lagged controls and by accounting for inflation forecasts, I now use this method to control many more variables and address the problem of omitted variable bias. One implication of the model with lagged controls and corrections for anticipation effects, is that a large number of controls can now be included, to remove as many sources of omitted variable bias as possible, without any of the control variables being bad controls. I thus study a large pool of over 2000 country-time varying control variables, constructed from the World Bank's development indicators. I combine these with the global leaders database which comprises a large number of political variables to cover characteristics which describe a country's political institutions. Taken

together, these two datasets should include or proxy almost any observable cross-country variable.

To select from the pool of control variables in a non-arbitrary way, and without over-fitting, I use the least absolute shrinkage and selection operator (LASSO). LASSO has not yet been developed for system GMM estimation of dynamic panel models and so I cannot use LASSO for inference on the inflation–happiness relationship; rather I use it solely as a method to select relevant control variables. I select control variables by solving the LASSO algorithm for a pooled OLS dynamic panel regression and then identifying any lagged control variables with non-zero estimated coefficients. I then plug the selected control variables into the dynamic panel threshold model which is estimated by system GMM as in previous sections. For a detailed explanation of the LASSO procedure, see Appendix C.

Column (3) in Table 3 estimates the inflation–happiness relationship using the dynamic panel threshold procedure with 22 lagged control variables selected by LASSO cross validation (LASSO-CV). The estimated relationship remains astonishingly similar to the one estimated elsewhere in this paper. The coefficient on above threshold inflation remains negative and significant. The coefficient on below threshold inflation is positive and just significant at the 10% level. This result could be evidence that moving from very low or negative rates of inflation towards more moderate inflation of around 5% weakly makes people happier. Then again, this result is not consistent across the paper and so it should be interpreted with caution. The estimated threshold of 4.75% in column (3) is in line with results of previous sections, though the confidence interval is wider than in many other estimates. As a check, column (4) estimates the dynamic panel threshold using unexpected inflation as an explanatory regressor, instead of inflation, generating another similar set of results.

The finding of the lagged controls exercise is thus that the estimated relationship between inflation and happiness is robust to a large set of plausibly important control variables. While it is surprising that our results about the inflation–happiness relationship do not change after including many control variables, this could be explained by the fact that the country

fixed-effects and lagged dependent variable already account for any endogeneity which might be a threat to inference. After accounting for country fixed-effects and the lagged dependent variable, inflation's relationship with happiness is the same in models with no additional control variables, and models with 22 control variables selected from a pool of thousands.

4.3 Reverse Causality

Having addressed concerns of omitted variables bias and bad controls, I now turn to reverse causality. It is quite possible that happier populations consume and invest more, thereby driving up inflation. This threatens to bias our existing estimates, which cannot separate out the simultaneous effects of inflation on happiness and happiness on inflation.

In this section, I argue that this concern does not undermine our main conclusions and in fact supports them. First, the dynamic model already provides a significant control for this channel. The inclusion of the lagged dependent variable Y_{it-1} means that the regression estimates the effect of inflation on happiness while holding the previous year's happiness constant. Given that any effect of public mood on aggregate demand and inflation would likely operate with a lag, controlling for past happiness effectively removes the primary pathway for reverse causality. While news of high inflation can impact well-being almost immediately, it likely takes time for sustained consumer optimism to translate into inflationary pressure. Second, to the extent that any contemporaneous reverse causality remains, the empirical evidence in this section shows it runs in the opposite direction of our main finding. As will be demonstrated below, higher happiness appears to cause higher inflation, not lower inflation. This creates a countervailing effect that would only bias our main coefficient on inflation towards zero. Therefore, if the estimates presented in this paper are biased by reverse causality, then this only means that they are understated.

To empirically test the direction of reverse causality, I test whether lags of happiness predict future values of inflation. If the threshold relationship were explained by happiness causing inflation, we should find that higher past happiness reduces inflation when it is above

Table 4: Reverse Causality

	(1) Inflation	(2) Happiness
Happiness _{it-1}		0.603*** (0.0528)
π_{it-1}	0.125 (0.0773)	
π_{it-2}	-0.0483 (0.0479)	
$\hat{\gamma}$	7.95 [7.7 - 8.1]	5.65 [4.55 - 6.7]
$\pi_{it-1}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$		-0.00320 (0.00561)
$\pi_{it-1}\mathbb{I}(\pi_{it} > \hat{\gamma})$		-0.0158*** (0.00364)
Happiness _{it-1} $\mathbb{I}(\pi_{it} \leq \hat{\gamma})$	-0.370 (0.237)	
Happiness _{it-1} $\mathbb{I}(\pi_{it} > \hat{\gamma})$	0.674** (0.274)	
$\pi_{it}^{Forecast}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$		0.266** (0.134)
$\pi_{it}^{Forecast}\mathbb{I}(\pi_{it} > \hat{\gamma})$		0.448** (0.201)
Observations	1015	1108
Additional Macro Controls	5	5

Corrected and robust standard errors from two-step GMM in parentheses. 95% confidence intervals in square brackets. Estimates include a constant term which is not reported. Controls for Unemployment, Govt Expenditure, Investment, Openness and Growth are included in both estimates but not reported.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

5%. In Table 4, I find the opposite. In column (1), I regress inflation on lags of itself¹¹, as well as the 5 macroeconomic controls used in the baseline specification in section 4.1, and lags of happiness that vary across inflation regimes. To be sure that the past values of happiness are not predicting future values of inflation simply due to anticipation effects, I control for inflation forecasts, in line with section 4.2.

The results in column (1) support the opposite relationship to that which has been found in the rest of the paper. I find that higher levels of lagged happiness actually predict higher levels of inflation. I also find evidence that this result only holds when inflation is high. This has a simple theoretical explanation: when individuals in an economy are happier, they spend more, increasing aggregate demand. However, this only puts pressure on inflation in high inflation regimes, because this is when the economy is at full capacity and hence when demand pressure leads to inflation. In low-inflation environments with spare capacity, this demand does not translate into price increases. Thus, far from explaining our main result, the empirical findings of this section confirm that any remaining reverse causality which is not controlled by the lagged dependent variable, would cause our estimates to be understated.

The effect of lagged happiness on inflation, while statistically significant, is small in size. A 1 point increase in past happiness measured on a 0-10 scale is associated with just a 0.674 increase in inflation. This is a large increase in happiness required for a comparatively minor change in inflation. The opposing negative effect of inflation on happiness unsurprisingly dominates, especially when the lagged dependent variable which conditions out much of the reverse causality is controlled.

To confirm that the primary causal channel runs from inflation to happiness, I study regressions of happiness on lags of inflation, rather than contemporaneous inflation. These can be found in column (2) of Table 4. As expected, lagged inflation has a very similar relationship with happiness to the one uncovered in the rest of this paper: it is negative and significant, but only when it exceeds a threshold of around 5% .

¹¹I found that it was necessary to regress inflation on its first and second lag in order to remove serial correlation from the estimated model and pass diagnostic tests.

These empirical findings on reverse causality are consistent with the conclusion of the rest of this paper. They suggest that the contemporaneous correlation between inflation and happiness disguises two opposing relationships: first, a strong negative effect of inflation on happiness which only appears when inflation exceeds 5%; second, a small positive effect of happiness on inflation which only appears when inflation exceeds around 8%. By controlling for the lag of happiness, the estimates of this paper remove much of the reverse causality from happiness to inflation. As this section has shown, any remaining reverse causality has the effect of understating the true negative effect of inflation on happiness.

4.4 Identification Assumptions

Given that the results estimated in this paper do not exploit exogenous variation in inflation, we should be cautious to give them a causal interpretation. However, we can still outline the precise assumptions under which a causal interpretation of these results could be made, and I will argue that these assumptions are plausible. When using LASSO to select from a set of lagged controls, the coefficient on inflation has a causal interpretation under the following assumptions:

1. Inflation does not cause lagged values of other variables, other than through inflation forecasting (*The no bad controls assumption*)
2. All relevant sources of omitted variable bias in the inflation–happiness regression can be found in the set of LASSO-selected control variables (*The selection on observables assumption*)
3. There is no reverse causality from happiness to inflation
4. There is no measurement error in inflation

Section 4.2 provided justification for the no bad controls assumption, and common sense suggests it holds after accounting for forecasting. The no measurement error assumption is

likewise not a concern since this paper uses World Bank inflation data, and any measurement error would only serve to attenuate any reported coefficients. The reverse causality assumption was addressed in section 4.3. By far the strongest assumption, then, is that of selection on observables. However, we can defend this assumption on a few grounds. The first is that the set of control variables in this paper is so large that it either directly captures or proxies every variable that could possibly influence both inflation and happiness. We might think that as our set of control variables gets larger and larger with no duplicates, the potential for omitted variable bias in a regression which selects from these variables becomes smaller and smaller. As a result, we can think of our particular set of thousands of control variables as approaching this limit.

A complication regarding selection on observables, however, is that the entire set of control variables chosen in our regression were lagged by one period. This was done in order to ensure that none of them were bad controls. The result of lagging controls is that any contemporaneous shocks which influence both happiness and inflation simultaneously will not be controlled for and only their lagged counterparts will be. This is a major concern for inference. While many contemporaneous variables will be persistent and proxied by their first lag, many will also contain random shocks in each period. To the extent that these random shocks are correlated with inflation and happiness, then selection on observables is violated. However, this concern is weaker, in the case of inflation, if we think that inflation responds to variables with a lag. If inflation responds to contemporaneous shocks with a lag of one year, then it would be appropriate to control for only lagged variables, since their contemporaneous counterparts would not be omitted variables with respect to contemporaneous inflation. Thus, while this is by the far the strongest assumption that needs to be made, it is not impossible that a very large set of lagged controls can entirely address the selection on observables requirement for causality.

5 Robustness Checks

Having estimated a set of baseline results and defended them against concerns of omitted variable bias and reverse causality, I now conduct additional robustness checks. These will test the robustness of the conclusion that the inflation–happiness relationship moves from insignificant (or weakly positive) to negative at a threshold of around 5%.

5.1 Falsification tests

As a first check on robustness, I study two falsification tests for the threshold relationship found between inflation and happiness in previous sections. An implicit claim of this study is that the empirical threshold relationship seen for inflation and happiness appears because inflation uniquely is a variable where we would expect such a relationship. In contrast to other variables, inflation’s optimal rate is generally thought to be somewhere in the middle of the set of possible rates, hence why inflation is targeted rather than maximised or minimised. If this reasoning explains the inflation–happiness threshold relationship, then we should not expect to see notable thresholds in other variables, and we should certainly not expect to see a slope which only becomes significant above a certain level. If we were to empirically discover thresholds in other variables besides inflation, we might worry that the inflation threshold is generated by some unobserved threshold process, rather than the true relationship between inflation and happiness.

To check this, I study the unemployment–happiness relationship which was found to be negative, significant and relatively stable in all of the regressions we have studied so far. I study two subtly different falsification tests in this section. First, I study the unemployment–happiness relationship with thresholds in unemployment. This test checks whether a threshold appears in unemployment when unemployment gets higher. Second, I study the unemployment–happiness relationship with thresholds in inflation. This test checks whether when inflation gets higher, the unemployment–happiness relationship changes.

5.1.1 Unemployment thresholds in unemployment

Unlike inflation, we do not have reasons to suspect that additional unemployment is only harmful when unemployment exceeds a particular threshold, or that the unemployment-happiness relationship changes sign from positive to negative. Common sense suggests that higher and higher unemployment should make people less and less happy and that, all else equal, the optimal state is for unemployment to be as low as possible.¹²

Column (1) in Table 8 in Appendix D looks for a possible change in the slope of the unemployment-happiness relationship as unemployment rates vary.¹³ I find no evidence of such a change. This suggests that inflation's relationship with happiness is uniquely non-linear and that the existence of thresholds is not a necessary feature of the estimation procedure. The unique non-linearity of inflation also provides evidence against a non-linear reporting function (i.e. a non-linear relationship between true happiness and reported happiness scores), which has been suggested as a possible way of generating spurious non-linear results in subjective wellbeing data (Oswald, 2008).

Another way to demonstrate the uniqueness of the inflation threshold is by comparing the likelihood ratio plots for inflation and unemployment thresholds. Figure 2 plots the likelihood ratio function $S(\gamma)$ for different values of γ . The left hand panel plots the likelihood ratio function for the estimated inflation threshold from the baseline model (column 1 in Table 2). The right hand panel plots the likelihood ratio function for the estimated unemployment threshold from column (1) in Table 8. The likelihood ratio function in both cases compares the model fit of the estimated threshold with the fit of all the other thresholds which were tried in the grid search. The 95% confidence intervals for each threshold estimate are the set of thresholds such that the likelihood ratio function lies below the grey dotted line (see section 3 for details).

¹²A caveat to this is that extremely low unemployment rates can generate inflationary pressure which is bad for happiness. This is not a concern in our results since inflation is a regressor and is thus held constant.

¹³The grid search for possible thresholds includes unemployment rates of between 2% and 20%, at intervals of 0.05.

In the left-hand panel, for the baseline dynamic panel threshold model of inflation, we see a clear V-shape in the likelihood ratio function. This shape suggests that the estimated threshold (4.9%) fits the model significantly better than other thresholds which were tried and that 4.9% is clearly not an arbitrary point at which to split the data. In contrast, the likelihood ratio function for unemployment thresholds is flat across most of the grid of thresholds which were tried. The estimated threshold (18.6%) appears to be a point which splits the data no better than any other arbitrary threshold which could have been otherwise chosen.




Figure 2: Likelihood ratio plots for inflation (left) and unemployment (right) thresholds

5.1.2 Unemployment thresholds in inflation

Having studied unemployment thresholds under different unemployment regimes, I now study whether the slope of unemployment with happiness varies under different inflation regimes. The rationale for this test is that, if we observe other variables besides inflation changing slope or sign on each side of inflation thresholds, then there could be a large range of unobservable variables that also change as inflation changes its regime. For example, it is plausible that on each side of the inflation threshold, a range of unobservable variables change to reflect the fact that the economy has moved from a downturn to an expansion.

¹³95% confidence intervals are values below the grey dotted line which lies at $c(0.05)$

Column (2) in Table 8 studies the unemployment–happiness relationship, with inflation as a threshold variable. In this specification, I continue to see no sign of thresholds or sign changes in the unemployment–happiness relationship. The estimated inflation threshold for the unemployment–happiness relationship is quite near to the 4.9% threshold estimated for the inflation–happiness relationship. However, the confidence interval is very large and the slope coefficients clearly indicate that this threshold is not due to a sign change or a move from insignificance to significance in the relationship between unemployment and happiness. Overall, these results strongly suggest that, unlike inflation, the relationship between unemployment and happiness is linear and does not vary depending on the inflation rate.

5.2 Estimation procedures

While the procedure used in this paper is a first-best way of estimating thresholds in a dynamic panel setting, it is useful to check whether the results are consistent with simpler second-best procedures. In this section, I test the robustness of the 4.9% threshold estimate from the baseline model to two alternative estimation procedures.

I first estimate a saturated model, where inflation’s slope is allowed to vary for each 2% interval between 0% and 20%. This model gives an indication of how inflation’s slope varies across the set of inflation rates. Since inflation’s slope is allowed to vary at many different intervals, this procedure should capture a threshold relationship if it exists, while also capturing more complicated non-linearities that the threshold model might have missed.

The second check estimates the within-groups threshold estimation procedure of (Hansen, 1999). Although this methodology is known to be suboptimal since within-groups estimation introduces bias in the the lagged dependent variable (Nickell, 1981), I use it to confirm that my central findings are not contingent on the removal of the Nickell bias. The procedure for within-groups threshold estimation involves three steps. First, variables are transformed using the within-groups transformation. Second, the model is estimated over the usual grid

of thresholds and the threshold which produces the model with the lowest sum of squared residuals is chosen. Third, this optimal threshold is treated as fixed and slope parameters are estimated, using a within-groups fixed-effects regression (with no time dummies).

The results from both the saturated and within-groups estimation procedures are presented in Table 9. Both sets of estimates confirm this paper's baseline results. In the saturated model, Inflation's coefficient is positive and insignificant for the below 0%, 0% to 2%, and 2% to 4% inflation bands, and then negative for all inflation intervals above 4%. This provides strong support for the baseline threshold model which suggests that inflation's relationship with happiness becomes negative at 4.9%. The negative inflation coefficients do not vary much once they become negative, providing support for the idea that inflation has a single threshold in its relationship with happiness and that this threshold is due to a move from a weakly positive to a negative relationship.

The results for the lower inflation bands provide insight into what is driving the insignificance of inflation below the threshold. For the below 0% and 0% to 2% bands, inflation's coefficient is quite large, though it is just insignificant due to its high variance. In contrast, the 2% to 4% coefficient is very small and not close to being significant. This suggests that the insignificant below threshold coefficient is driven by the fact that inflation rates in the low positive range (2% to 4%) have no effect on happiness. This is despite the fact that inflation very frequently appears in this range in the sample. The positive coefficients on below 0% and 0% to 2% inflation suggest that moving away from the extremes of ultra-low inflation and deflation are weakly positive for happiness.

Results from the within-groups threshold estimation procedure also closely mirror those from the baseline model. The within-groups approach estimates the inflation threshold at 4.85%, with a confidence interval ranging from 4.4% to 5.65%. This aligns well with the baseline threshold estimate of 4.9%. The slope coefficients from the within-groups model exhibit the same sign as the baseline model, moving from positive to negative. The below threshold coefficient is statistically significant under this specification, again suggesting that

inflation's relationship with happiness could be weakly positive below the threshold. Although these within-groups estimates are known to be subject to bias, their consistency with the baseline findings suggests that the key conclusions from the baseline model do not depend on the specific estimation procedure chosen.

5.3 Heterogeneity

Having estimated the dynamic panel threshold models on the whole sample, it is important to check whether the results vary across the sample. Heterogeneity is pertinent in this study since it is plausible that the effect of inflation on happiness varies across the wide range of countries and time periods included in the data which is studied.

To check for possible sources of heterogeneity, I estimate the model on 3 pairs of samples which all roughly split the dataset in half. I first split the sample based on whether a country meets the World Bank's definition of a high-income country.¹⁴ The purpose of this is to test whether richer countries respond differently to inflation when compared to poorer countries. It should be noted that the countries in my sample which do not fit the high-income definition are almost entirely middle-income countries, since happiness data on poorer countries is generally unavailable. I secondly split the sample across time, studying observations from before and after the year 2014 which falls in the middle of the time period studied. Finally, I split the sample based on whether a country is from Europe or elsewhere in the world. This demarcation tests for geographical heterogeneity between the relatively homogenous European continent and the remaining countries in the sample which come from 5 different continents.

Results on heterogeneity can be found in Table 10. In all three sub-sample comparisons, I find little evidence of heterogeneity across the sample. Whether comparing by income-group, across-time or across-continent, the effect of inflation on happiness appears relatively stable. That is to say, over the last 20 years, in a range of high and middle-income countries around

¹⁴A list of high-income countries under this definition is available at The World Bank

the world, the effect of inflation on happiness is negative only when inflation exceeds 5%.

One caution with these results is that the confidence intervals on the thresholds for inflation are wider under many of the sub-sample specifications. This is to be expected since the precision of the threshold estimates is directly proportional to the sample size (see section 3) which is halved in all of these estimates. Moreover, given the consistency of the threshold estimates across a large range of specifications, we can be relatively confident in the interpretation of a roughly 5% threshold in the inflation–happiness relationship.

6 Mechanism

This section studies the mechanism underlying the inflation–happiness relationship. It explores why inflation negatively impacts people’s happiness and why it does so only when inflation is high.

6.1 Real income and inequality effects

One way in which inflation could impact happiness is through inflation’s influence on other macroeconomic variables. Some of these have already been controlled in previous regressions in this paper, meaning that we can rule out the possibility that inflation affects happiness solely through inflation’s impact on growth, unemployment or other controls included in our previous equations. This section explores two other mechanisms which have not yet been ruled out, namely inflation’s direct real income effects and inflation’s impact on inequality.

Inflation’s implied impact on real income is likely to reduce happiness. When prices go up, the purchasing power of a given amount of income falls, and this is likely to reduce consumption and overall happiness as a consequence. If we expect nominal incomes to catch-up with prices in the long-run, then the real income explanation would suggest that inflation’s impact on happiness is transitory.

Our existing evidence does not support the real income explanation. The real income

explanation provides no basis for a threshold in the inflation–happiness relationship; it implies a negative monotonic relationship between inflation and happiness. Under the real income hypothesis, low and even negative inflation rates, which increase real incomes in the short-run, should increase wellbeing. In contrast, this paper frequently estimates a threshold relationship between inflation and happiness where inflation’s relationship is insignificant and weakly positive below 5% and negative and significant only above this threshold.

To be sure that the real income effect is not driving results, I include real income as a control in the baseline model. Results are presented in columns (1) of Table 12 in Appendix E. $\log(\text{GDP})$ is defined as the natural log of real GDP per capita in international dollars.¹⁵ Including this control does not make a meaningful difference to the baseline results; we continue to observe a sign change in the inflation–happiness relationship with a threshold at 4.9%. Interestingly, the magnitude of the below threshold coefficient increases while the magnitude of the above threshold coefficient decreases. This is consistent with the idea that increasingly high inflation is bad for happiness but not as bad when real income is held constant. Similarly, low and negative inflation, below the threshold, are much worse for happiness once you remove the fact that deflation increases real income. That being said, our conclusion about the threshold relationship between inflation and happiness is unchanged when real income is held constant.

Another macroeconomic variable commonly associated with inflation is inequality. Empirically, it has been shown by Glawe and Wagner (2024) that inflation increases inequality, but only when inflation exceeds a threshold value of 6%. This result could plausibly explain the threshold relationship between inflation and happiness observed in this paper, if inequality also makes people less happy. To test the inequality hypothesis, I include controls for a country’s Gini coefficient in column (2) of Table 12. In this specification, I find no difference in the inflation–happiness relationship after controlling for the level of inequality. Relative to

¹⁵ $\log(\text{GDP})$ is an I(1) variable in this data. Spurious regression is not a concern, however, since the number of time periods in this data is short (Baltagi, 2021). Since $\log(\text{GDP})$ is an explanatory regressor, its lack of stationarity does not affect the validity of the additional system GMM moment conditions, which depend on the stationarity of happiness.

the baseline model, I observe no change in the behaviour of the inflation coefficients or their threshold. In fact, inequality is found to have no significant relationship with happiness in the estimates. A candidate explanation for this result is that the government expenditure variable already picks up the key ways in which equality is associated with happiness, for example through better social safety nets. With this in mind, I also studied what happened to results when the government expenditure control was removed. I found that high inflation's negative effect on happiness was slightly stronger under this specification but otherwise the inflation–happiness relationship was qualitatively the same. To rule out the possibility that inflation causes changes to inequality with a lag, I also studied these models with inequality at $t + 1$ included as a control. I found no relevant difference in these estimates, confirming that inequality is not a key driving force behind the inflation–happiness relationship that we have observed in this paper.

6.2 Inflation Targets

The threshold relationship between inflation and happiness estimated in this study may also be generated by the fact that the monetary authorities of many countries explicitly target low inflation rates. It is possible that these inflation targets create a perception that inflation rates which significantly exceed target rates (which are commonly set at either 2% or 4%) are undesirable and that this perception makes people less happy when they observe inflation rates which significantly deviate from these rates. Such a hypothesis would be broadly consistent with this paper's finding that inflation negatively affects happiness when inflation reaches a rate of above 5% and that below this threshold, inflation has a weakly positive effect on happiness. This raises the question of whether low positive inflation rates are optimal to happiness even in countries which do not explicitly target these rates. Since many countries do not have explicit inflation targets, I test this hypothesis by studying the inflation–happiness relationship for countries which do not explicitly target inflation but still might have reasons to favour low positive inflation rates and dislike inflation rates which

significantly exceed the 5% threshold.¹⁶

Table 11 presents results for the inflation–happiness relationship in countries with and without inflation targets. The 5% threshold is present in both sub-samples. This suggests that low inflation rates are optimal for happiness even in countries which are not explicitly targeting them. The results uncovered in Table 11 favour the hypothesis that central banks are not altering the perception of inflation by targeting it. Instead, they suggest that regardless of whether inflation targets are in place, inflation which significantly exceeds the commonly targeted rates of 2% to 4% is damaging to a country’s happiness. The results do not, however, support the targeting of a specific inflation rate of 2% or a higher rate of 4%. Instead, they suggest that movements of inflation rates in this range go broadly unnoticed by a country’s population and do not affect their wellbeing.

It should be noted that the below threshold results cannot be used as strong evidence that all inflation rates below 5% have no impact on happiness. The present sample contains very few observations on inflation in the negative range, meaning that the null result for below threshold inflation is driven by low positive inflation rates, not negative rates. It remains possible that deflationary spirals reduce happiness. Such a conclusion will be hard to reach with present statistical evidence due to data availability. The consistent policy conclusion from these results is thus that inflation goes relatively unnoticed as it moves in the low positive range, before becoming harmful when it exceeds a threshold of roughly 5%. This result holds regardless of whether inflation is targeted.

6.3 Awareness of Inflation

Having thus far ruled out a number of alternative explanations, I now turn to the most plausible explanation for why high inflation, and only high inflation, negatively affects happiness. High inflation is disproportionately salient and newsworthy, and is often associated with panic in a way which we would expect to reduce happiness. In contrast, low inflation can get overlooked

¹⁶The list of countries which explicitly target inflation rates is based on Central Bank News (2025).

Table 5: Google Trends

	(1)	(2)	(3)
	Google Searches	Google Searches	Google Searches
Y_{it-1}	0.852*** (0.0233)	0.766*** (0.0222)	0.827*** (0.0279)
$\pi_{it}\mathbb{I}(\pi_{it} \leq 4.9)$	-0.0886 (0.141)	0.00396 (0.148)	0.0312 (0.194)
$\pi_{it}\mathbb{I}(\pi_{it} > 4.9)$	0.310*** (0.0686)	0.517*** (0.0947)	0.639*** (0.156)
$\log(\text{GDP})_{it}$		3.042*** (0.380)	
Unemployment_{it}			-0.180*** (0.0631)
$\log(\text{Govt expenditure})_{it}$			7.993*** (1.640)
$\log(\text{Investment})_{it}$			1.498 (1.623)
$\log(\text{Openness})_{it}$			0.759 (0.855)
Growth_{it}			0.484*** (0.0756)
Observations	1564	1563	1066

Corrected and robust standard errors from two-step GMM in parentheses. 95% confidence intervals in square brackets. Estimates include a constant term which is not reported.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

entirely by the public: many consumers may not even notice when their prices go up by a few percent, especially if inflation has not been in the news.

While it is impossible to read people's minds, we can use online search engine data to bring evidence to this hypothesis. To proxy public attention and fear of inflation, I use the relative number of Google searches for each country's annual average search interest in the term 'inflation', derived from monthly Google Trends data which is available for every country and year in our existing dataset. I use the Google Search data to test whether there is a stronger relationship between inflation and fear of inflation (proxied by Google Searches) when inflation is above the estimated 4.9% threshold from the baseline model.

Turning to Table 5, I study the effect of inflation on Google searches for inflation across 3 specifications. First, I include a simple dynamic model with no control variables.¹⁷ Second, I include the natural log of real GDP per capita, a measure of living standards which likely correlates with both inflation and frequency of Google searches. Finally, I use the 5 macroeconomic controls which were included in the baseline model, effectively re-estimating the baseline model with a different dependent variable. In all 3 specifications, slope coefficients are estimated by System GMM with country fixed-effects removed. I fix inflation's threshold at 4.9%: the threshold estimated in the baseline model of inflation and happiness.

Across all 3 specifications, the relationship between inflation and Google searches for inflation is small and insignificant when inflation is below the threshold, and positive and highly significant above the threshold. In other words, when inflation moves in the low positive range, but remains below 4.9%, it is not associated with a greater public perception or fear of inflation. In contrast, when inflation exceeds this threshold, higher inflation is associated with greater awareness and fear of inflation. These results are robust to alternative specifications which have not been reported, such as including year dummies, excluding the lagged dependent variable, or using time series data on global average inflation and Google searches for inflation.

¹⁷Specifications without the lagged dependent variable suffered from serious serial correlation.

Overall, the google trends data strongly supports the hypothesis that inflation goes relatively unnoticed when it moves in the low range below 5%, but becomes highly salient when it exceeds this level. This appears to be the strongest explanation for the estimated threshold relationship between inflation and happiness which has been the consistent finding of this paper.

7 Conclusion

The results of this thesis overwhelmingly provide support for a threshold relationship between inflation and happiness. Some features of this relationship, such as the threshold estimate of around 5%, as well as inflation's negative impact on happiness above this threshold were features of every model estimated. Below the threshold, inflation's impact on happiness was generally found to be insignificant, though results from saturated regressions revealed that this was driven by low positive inflation rates (in particular rates of 2% to 4%), which appear to have no effect on happiness.

Though the results do not exploit exogenous variation in inflation, they are highly consistent and survive key robustness tests. The results from the model with LASSO-selected lagged controls, which confirm the baseline results, can also be considered causal under strong but plausible assumptions. This paper thus goes far beyond typical papers in the cross-country happiness literature in defending its results which produce a set of findings that plausibly reflect the true effect of inflation on a country's happiness. The conclusion that only particularly high inflation reduces happiness is a relatively intuitive one. Most of the time, inflation is not on people's minds and this paper suggests that the reason why is rooted in the fact that inflation is typically in a range such that it has no effect on people's overall wellbeing. That said, when inflation becomes exceptionally high, making it noticeable and newsworthy, reduced happiness ensues. This study has provided the first empirical evidence of where exactly the threshold lies such that inflation becomes harmful, namely at around

5%.

A number of policy implications of this paper can thus be drawn. The study confirms that inflation targets in the low positive range are sensible choices for policymakers. However, rather than supporting a specific target, like 2%, they suggest that the key focus of policy should be on keeping inflation in the harmless range, below 5%, since inflation in this range has no negative impact on happiness. Whether countries choose to target zero inflation or low positive rates of 2% or 4% is less important than ensuring inflation doesn't significantly exceed 5%. A consequence of this finding is that monetary policy aimed at bringing inflation from, say, 3% to 2%, should be used with caution since the wellbeing effect of inflation in this range is likely neutral, and any resulting increase in unemployment is very likely to generate a large reduction in happiness.

This study should be seen primarily as uncovering a short-run relationship between inflation and happiness. It has been shown, however, that inflation's impact on happiness is not solely driven by real income effects which are cancelled out once wages catch-up to prices. Moreover, the dynamics estimated in this paper demonstrate that happiness persists over time, meaning that any change in happiness associated with inflation may last for many years, amplifying the short-run impact. Additionally, high inflation may have negative economic consequences which persist for many years, further affecting long term happiness. For that reason, one could interpret the results of this study as understating the negative long-term effect of high inflation on happiness.

A natural road forward for this research would be a study of the impact of deflation on happiness. While the present study suggests that inflation rates are not harmful to happiness unless they exceed 5%, this result is almost solely driven by positive inflation rates, due to data availability. Another important area of future work would be a causal exploration of inflation's effect on happiness, though it is notoriously difficult to find sources of exogenous variation in cross-country macro data.

Taking these empirical results as given, there will still remain questions around how they

should be interpreted. Whether a policymaker should target inflation rates which maximise happiness or rates which promote other forms of economic health is ultimately a philosophical question. Indeed, there remain important open problems surrounding precisely why inflation is bad, whether all inflation is bad, and whether any inflation is good. This study provides a speculative empirical answer to these questions, rooted in individuals' reported subjective happiness in their lives, which begins to fall when inflation exceeds 5%. While many countries already implement policies to keep inflation low, they can now do so with direct regard to the effect they will have on national happiness.

References

- Angrist, Joshua D. and Jörn-Steffen Pischke**, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton, NJ: Princeton University Press, 2009.
- Arellano, Manuel and Olympia Bover**, "Another Look at the Instrumental Variable Estimation of Error-components Models," *Journal of Econometrics*, 1995, 68 (1), 29–51.
- **and Stephen Bond**, "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *The Review of Economic Studies*, 1991, 58 (2), 277–297.
- Baltagi, Badi H.**, *Econometric Analysis of Panel Data* Springer Texts in Business and Economics, 6 ed., Springer Cham, 2021.
- Bernanke, Ben S, Thomas Laubach, Frederic S Mishkin, and Adam S Posen**, "Inflation Targeting: Lessons from the International Experience," *Princeton University Press*, 1999.
- Bick, Alexander**, "Threshold effects of inflation on economic growth in developing countries," *Economics Letters*, August 2010, 108 (2), 126–129.
- Blanchard, Olivier, Giovanni Dell'Ariccia, and Paolo Mauro**, "Rethinking Macroeconomic Policy," *Journal of Money, Credit and Banking*, 2010, 42, 199–215.
- Blanchflower, David G**, "Is Unemployment More Costly Than Inflation?," Working Paper 13505, National Bureau of Economic Research October 2007.
- Blanchflower, David G., David N. F. Bell, Alberto Montagnoli, and Mirko Moro**, "The Happiness Trade-Off between Unemployment and Inflation," *Journal of Money, Credit and Banking*, 2014, 46 (S2), 117–141.
- Blundell, Richard and Stephen Bond**, "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models," *Journal of Econometrics*, 1998, 87 (1), 115–143.

- Bond, Stephen**, “Dynamic Panel Data Models: A Guide to Micro Data Methods and Practice,” *Portuguese Economic Journal*, 2002, 1, 141–162.
- Bond, T.N. and K. Lang**, “The sad truth about happiness scales,” *Journal of Political Economy*, 2019, 127 (4), 1629–1640.
- Caner, Mehmet and Bruce E. Hansen**, “Instrumental Variable Estimation of a Threshold Model,” *Econometric Theory*, 2004, 20 (5), 813–843.
- Central Bank News**, “Inflation Targets,” 2025. <https://www.centralbanknews.info/p/inflation-targets.html> Accessed: 2025-05-10.
- Clarida, Richard, Jordi Gali, and Mark Gertler**, “The science of monetary policy: a New Keynesian perspective,” *Journal of Economic Literature*, 1999, 37 (4), 1661–1707.
- Ferrer-i-Carbonell, Ada and Paul Frijters**, “How important is methodology for the estimates of the determinants of happiness?,” *The Economic Journal*, 2004, 114 (497), 641–659.
- Glawe, Lisa and Helmut Wagner**, “Inflation and inequality: new evidence from a dynamic panel threshold analysis,” *International Economics and Economic Policy*, 2024, 21, 297–309.
- Ha, Jongrim, M. Ayhan Kose, and Franziska Ohnsorge**, “One-Stop Source: A Global Database of Inflation,” *Journal of International Money and Finance*, October 2023, 137, 102896.
- Hansen, Bruce E.**, “Threshold Effects in Non-Dynamic Panels: Estimation, Testing, and Inference,” *Journal of Econometrics*, 1999, 93 (2), 345–368.
- , “Sample splitting and threshold estimation,” *Econometrica*, 2000, 68 (3), 575–603.
- Helliwell, John F., Richard Layard, Jeffrey D. Sachs, Jan-Emmanuel De Neve, Lara B. Akin, and Shun Wang**, eds, *World Happiness Report 2024*, University of Oxford: Wellbeing Research Centre, 2024.
- Kahneman, Daniel and Angus Deaton**, “High income improves evaluation of life but not emotional well-being,” *Proceedings of the National Academy of Sciences*, 2010, 107 (38), 16489–16493.
- Kaiser, Caspar and Andrew J. Oswald**, “The scientific value of numerical measures of human feelings,” *Proceedings of the National Academy of Sciences*, 2022, 119 (42).
- Khan, Mohsin S. and Abdelhak S. Senhadji**, “Threshold Effects in the Relationship Between Inflation and Growth,” *IMF Staff Papers*, 2001, 48 (1), 1–21.
- Kremer, Stephanie, Alexander Bick, and Dieter Nautz**, “Inflation and Growth: New Evidence from a Dynamic Panel Threshold Analysis,” *Empirical Economics*, 2013, 44, 861–878.

- Maddala, G. S. and Shaowen Wu**, “A Comparative Study of Unit Root Tests with Panel Data and a New Simple Test,” *Oxford Bulletin of Economics and Statistics*, 1999, *61* (S1), 631–652.
- Nickell, Stephen J.**, “Biases in Dynamic Models with Fixed Effects,” *Econometrica*, 1981, *49* (6), 1417–1426.
- Oswald, Andrew J.**, “On the Curvature of the Reporting Function from Objective Reality to Subjective Feelings,” *Economics Letters*, 2008, *100* (3), 369–372.
- Ruprah, Inder J. and Pavel Luengas**, “Monetary policy and happiness: Preferences over inflation and unemployment in Latin America,” *The Journal of Socio-Economics*, 2011, *40* (1), 59–66.
- Samuelsson, Conrad, Samuel Dupret, Michael Plant, and Caspar Kaiser**, “Can We Trust Wellbeing Surveys? A Pilot Study of Comparability, Linearity, and Neutrality,” March 2023. Happier Lives Institute.
- Tella, Rafael Di, Robert J. MacCulloch, and Andrew J. Oswald**, “Preferences over inflation and unemployment: Evidence from surveys of happiness,” *American Economic Review*, 2001, *91* (1), 335–341.
- , – , and – , “The macroeconomics of Happiness,” *Review of Economics and Statistics*, 2003, *85* (4), 809–827.
- Windmeijer, Frank**, “A finite sample correction for the variance of linear efficient two-step GMM estimators,” *Journal of Econometrics*, 2005, *126* (1), 25–51.

A Variable Definitions

Variable	Definition	Data Source
Inflation	Headline consumer price index inflation.	World Bank Global Database on Inflation
Unemployment	The number of unemployed persons as a percentage of the labour force, using the OECD harmonized unemployment rate.	IMF World Economic Indicators
Growth	Percentage change in GDP per capita, where GDP is measured in constant international dollars.	IMF World Economic Indicators
Investment	Total value of gross fixed capital formation and changes in inventories and acquisitions less disposals of valuables for a unit or sector. Expressed as a ratio of total investment in current local currency and GDP in current local currency.	IMF World Economic Indicators
Government Expenditure	General government final consumption expenditure, as a percentage of GDP.	World Bank Development Indicators
Openness	Sum of the value of exports and imports of goods and services, as a percentage of GDP.	World Bank Development Indicators
Positive Affect	Average of 3 binary response questions: <ul style="list-style-type: none"> • “Did you smile or laugh a lot yesterday?” • “Did you experience the following feelings during A LOT OF THE DAY yesterday? How about Enjoyment?” • “Did you learn or do something interesting yesterday?” 	Gallup world poll / World Happiness Report
Negative Affect	Average of 3 binary response questions: <ul style="list-style-type: none"> • “Did you experience worry during A LOT OF THE DAY yesterday?” • “Did you experience sadness during A LOT OF THE DAY yesterday?” • “Did you experience anger during A LOT OF THE DAY yesterday?” 	Gallup world poll / World Happiness Report
Affect	Positive Affect - Negative Affect	

Table 6: Variable Definitions and Data Sources

B Happiness Dynamics and Specification Tests

This Appendix studies the dynamic properties of cross-country happiness and provides specification tests. Previous literature studying happiness and subjective wellbeing in panel data has generally ignored its auto-regressive properties. This paper takes a different approach by showing that the lagged dependent variable in happiness is highly significant and is thus an important control variable. I use this section to demonstrate the presence of the Nickell Bias and the fact that system GMM estimation leads to efficiency improvements when estimating dynamic panel models of cross-country happiness.

Monte-carlo simulations have shown that the system GMM estimator performs better than the Arellano-Bond estimator in the presence of either high persistence or a high variance of the fixed-effects across panels (Blundell and Bond, 1998). The data I use in this paper is characterised by both of these properties, in particular a very high variation in happiness across countries. This should mean that the system GMM estimator is more efficient than the Arellano-Bond estimator, and thus preferable in the finite sample that this paper uses.

One way of comparing the performance of the system GMM estimator with that of the Arellano-Bond estimator is by comparing results from these estimators with the results of the pooled OLS and within-groups estimators, which are both known to be biased in dynamic panel settings with cross-country heterogeneity. Specifically, in a dynamic panel with a large N and small T , theory suggests that the coefficient on the lagged-dependent variable should be upwardly biased when estimated by pooled OLS and downwardly biased when estimated using the within-groups estimator (Nickell, 1981). Therefore, when using consistent GMM estimators, (such as the Arellano-Bond or system GMM estimators) we should expect them to generate a coefficient which lies between those of pooled OLS and within-groups. If this finding does not hold, then that may be a sign of serious inefficiency in the GMM estimates (Bond, 2002).

In Table 7, I use the above criterion to assess the performance of the system GMM estimator used in this thesis. Columns (1) to (4) present panel regressions of happiness

on its first lag, using the system GMM, Arellano-Bond, pooled OLS and within-groups estimators. For the system GMM estimator, the coefficient on the lagged dependent variable lies between that of pooled OLS and within-groups estimation as expected. The Arellano-Bond estimator, in contrast, produces a much smaller estimate on the lagged dependent variable, even smaller than within-groups, suggesting that it might be performing poorly. The efficiency improvement of system GMM thus appears to be an important consideration for our estimation procedure. For the system GMM estimator, we find that a one point increase in average happiness in a country in its previous year is associated with around a 0.6 increase in its happiness in the current year.

Columns (5) to (8) study the coefficient on the lagged dependent variable when we include the macroeconomic predictors of happiness which are studied in section 4. The additional regressors are inflation¹⁸, unemployment, growth, and the natural log of investment, government expenditure and openness as percentages of GDP. When these controls are included, we continue to find that the system GMM estimate of the coefficient on the lagged dependent variable lies between the estimates from pooled OLS (7) and within-groups estimation (8). This provides further justification for the use of the system GMM estimator.

In both the system GMM and Arellano-Bond estimates, I find evidence of first-order auto-correlation in first-differenced residuals, but no evidence of 2nd or higher order auto-correlation¹⁹. This suggests that using a single lag of happiness accounts for any serial correlation in the error term and no higher order lags of happiness are necessary. Sargan-Hansen tests fail to reject the null hypothesis, suggesting that the set of instruments is valid. A difference-in-Sargan/ difference-in-Hansen test comparing the instrument sets for the Arellano-Bond and system GMM type estimators also fails to reject the null, suggesting that the additional moment conditions used in the system GMM estimator are valid. For the

¹⁸For now, I ignore thresholds in inflation since they make no difference to the coefficient on the lagged dependent variable

¹⁹For the system GMM estimator, I use forward-orthogonal deviations transformations for coefficient estimation. The Arellano-Bond tests are conducted by first-differencing each variable and using the estimated coefficients to construct first-differenced residuals which are tested for auto-correlation.

rest of this paper, these diagnostic tests continue to pass but are not reported.

When I refer to system GMM in this paper, I mean the estimator outlined in section 3 which uses lagged levels as instruments for a forward orthogonal deviations transformed equation and lagged differences as instruments for a levels equation. An alternative system GMM estimator could replace the forward orthogonal deviations transformation with a first-difference transformation. In practice, switching between these two choices made almost no difference to any of the results presented in this paper, and so I chose the forward orthogonal deviations transformation to be consistent with the procedure used to estimate the threshold. In system GMM estimates, I allow for a non-zero constant, though I generally do not report its coefficient. Restricting the constant to zero makes no meaningful difference to any results reported in this section or elsewhere.

Though not the main focus of this paper, the finding of a highly significant lagged dependent variable is an interesting finding in its own right. It suggests that previous literature studying happiness in a cross-country panel setting has missed an important feature of happiness data, and it opens the door for a time-series literature on happiness. In this paper, the lag of happiness serves as an important control variable. Lagged happiness, which I have shown is highly correlated with contemporaneous happiness is also correlated with contemporaneous inflation, meaning that it can be a source of omitted variable bias if it is not included as a regressor when studying the inflation–happiness relationship. Direct evidence of the effect of lagged happiness on inflation is studied in section 4.3, confirming the need to control lagged happiness. Furthermore, simulation studies from Kremer et al. (2013) confirm that the dynamic panel threshold model performs better when a lagged dependent variable which exists in the data generating process for happiness is included in the estimation procedure.

Table 7: System GMM diagnostics

	(1)	(2)	(3)	(4)
	Happiness	Happiness	Happiness	Happiness
Y_{it-1}	0.625*** (0.0500)	0.450*** (0.0751)	0.917*** (0.0104)	0.592*** (0.0251)
Macro Controls	No	No	No	No
Estimation Procedure	System GMM	Arellano-Bond	Pooled OLS	Within-Groups
Observations	1108	1108	1108	1108
m1	-5.06	-4.39		
m2	-0.77	-0.94		
Sargan-Hansen (p-value)	1.00	0.92		

	(5)	(6)	(7)	(8)
	Happiness	Happiness	Happiness	Happiness
Y_{it-1}	0.596*** (0.0570)	0.324*** (0.0785)	0.842*** (0.0157)	0.457*** (0.0275)
Macro Controls	Yes	Yes	Yes	Yes
Estimation Procedure	System GMM	Arellano-Bond	Pooled OLS	Within-Groups
Observations	1108	1108	1108	1108
m1	-5.91	-4.86		
m2	-0.78	-0.81		
Sargan-Hansen (p-value)	1.00	1.00		

Standard errors adjusted for clustering by country in parentheses. GMM estimates are two-step GMM with corrected and robust standard errors. m1 and m2 denote the test statistic from Arellano-Bond tests of first and second-order serial correlation in first-differenced residuals.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

C LASSO Estimation Procedure

The LASSO algorithm solves:

$$\arg \min_{\alpha, \beta_1, \beta_2, \delta} \sum_{i=1}^N (Y_{it} - \alpha Y_{i,t-1} - \beta_1 \pi_{it} I(\pi_{it} \leq \gamma) - \beta_2 \pi_{it} I(\pi_{it} > \gamma) - \delta' x_{it-1}) + \lambda \sum_{k=1}^K |b_k| \quad (8)$$

where Y_{it} , π_{it} and γ are defined as in section 3²⁰ and x_{it-1} is a vector of around 2000 lagged control variables. λ is a hyperparameter which governs the penalty for including an additional regression coefficient.

I select λ using 4-fold cross-validation. This method is commonly applied in the literature since it uses out-of-sample prediction as a method of testing the optimal λ . The procedure generates an optimal value of λ which minimises the cross-validation function, and substituting this into the LASSO algorithm gives 22 non-zero control variables in this case. After the set of control variables is selected, I plug them in to the dynamic panel system GMM estimation procedure, in place of the researcher-chosen controls (Unemployment, Growth etc.) and estimate the threshold and slope coefficients in the usual way.

²⁰I impose a 5% inflation threshold when performing LASSO estimation but I allow the threshold to be estimated once the control variables are selected.

D Robustness Checks

Table 8: Unemployment falsification test

	(1)	(2)
	Happiness	Happiness
Y_{it-1}	0.597*** (0.0562)	0.603*** (0.0634)
$\hat{\gamma}$	18.60 CI:[2.25 - 19.65]	5.65 CI: [-1.7 - 15.8]
$U_{it}\mathbb{I}(U_{it} \leq \hat{\gamma})$	-0.0370*** (0.00729)	
$U_{it}\mathbb{I}(U_{it} > \hat{\gamma})$	-0.0292*** (0.00672)	
$U_{it}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$		-0.0314*** (0.0108)
$U_{it}\mathbb{I}(\pi_{it} > \hat{\gamma})$		-0.0295*** (0.00761)
π_{it}	-0.0162*** (0.00372)	-0.0156*** (0.00346)
$\log(\text{Govt_expenditure}_{it})$	0.456*** (0.119)	0.460*** (0.123)
$\log(\text{Investment}_{it})$	-0.0395 (0.0991)	-0.0283 (0.0964)
$\log(\text{Open}_{it})$	-0.0194 (0.0506)	-0.0134 (0.0477)
Growth_{it}	0.0111*** (0.00296)	0.0110*** (0.00280)
Observations	1108	1108

Corrected and robust standard errors from two-step GMM in parentheses. U_{it} denotes the unemployment rate. π_{it} denotes the actual inflation rate. 95% confidence intervals in square brackets. Estimates include a constant which is not reported.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Alternative threshold estimation procedures

Saturated Model		Within-groups Model	
	Happiness		Happiness
Y_{it-1}	0.594*** (0.0633)	Y_{it-1}	0.498*** (0.0452)
$\pi_{it}\mathbb{I}(\pi_{it} < 0)$	0.0118 (0.0347)	$\hat{\gamma}$	4.85 CI:[4.4 - 5.65]
$\pi_{it}\mathbb{I}(0 \leq \pi_{it} < 2)$	0.0421 (0.0278)	$\pi_{it}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$	0.0158** (0.00761)
$\pi_{it}\mathbb{I}(2 \leq \pi_{it} < 4)$	0.000245 (0.0122)	$\pi_{it}\mathbb{I}(\pi_{it} > \hat{\gamma})$	-0.00797*** (0.00255)
$\pi_{it}\mathbb{I}(4 \leq \pi_{it} < 6)$	-0.0139 (0.00991)	Unemployment _{it}	-0.0356*** (0.00599)
$\pi_{it}\mathbb{I}(6 \leq \pi_{it} < 8)$	-0.0195*** (0.00732)	log(Govt Exp) _{it}	0.194 (0.142)
$\pi_{it}\mathbb{I}(8 \leq \pi_{it} < 10)$	-0.00887 (0.00609)	log(Investment) _{it}	-0.0335 (0.0867)
$\pi_{it}\mathbb{I}(10 \leq \pi_{it} < 12)$	-0.0120* (0.00653)	log(Open) _{it}	0.160 (0.105)
$\pi_{it}\mathbb{I}(12 \leq \pi_{it} < 14)$	-0.0232** (0.00906)	Growth _{it}	0.00668** (0.00306)
$\pi_{it}\mathbb{I}(14 \leq \pi_{it} < 16)$	-0.0155** (0.00688)	Observations	1108
$\pi_{it}\mathbb{I}(16 \leq \pi_{it} < 18)$	-0.0129** (0.00628)	Threshold Estimation	Hansen (1999)
$\pi_{it}\mathbb{I}(18 \leq \pi_{it} < 20)$	-0.0105 (0.0111)	Slope Estimation	W-G
$\pi_{it}\mathbb{I}(\pi_{it} \geq 20)$	-0.0124*** (0.00374)		
Observations	1108		
Slope Estimation	System GMM		
Controls(Not Reported)	See footnote		

Corrected and robust standard errors from two-step GMM in parentheses (left). Cluster-robust standard errors from within-groups estimation (right). The left hand column includes 5 macroeconomic controls. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Heterogeneity

	(1)	(2)	(3)	(4)	(5)	(6)
	Happiness	Happiness	Happiness	Happiness	Happiness	Happiness
Y_{it-1}	0.716*** (0.0865)	0.685*** (0.0751)	0.477*** (0.0975)	0.804*** (0.0835)	0.832*** (0.0952)	0.551*** (0.0822)
$\hat{\gamma}$	4.85 [-1 - 10]	4.9 [4.15 - 6.05]	4.8 [3.45 - 5.1]	5.65 [-1 - 10]	4.85 [1.05 - 8.5]	4.90 [4.2 - 6.70]
$\pi_{it}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$	0.0124 (0.0147)	0.00934 (0.0110)	0.0164 (0.0146)	-0.00203 (0.00902)	-0.000473 (0.0128)	0.0171 (0.0139)
$\pi_{it}\mathbb{I}(\pi_{it} > \hat{\gamma})$	-0.00935* (0.00487)	-0.00875*** (0.00340)	-0.0213*** (0.00796)	-0.00969*** (0.00338)	-0.0119*** (0.00419)	-0.00946* (0.00540)
Unemployment $_{it}$	-0.0251*** (0.00716)	-0.0134** (0.00534)	-0.0417*** (0.0104)	-0.0138* (0.00804)	-0.0109** (0.00553)	-0.0333*** (0.00802)
$\log(\text{Govt Exp})_{it}$	0.269 (0.186)	0.0703 (0.167)	0.483** (0.242)	0.209 (0.141)	0.147 (0.239)	0.500*** (0.190)
$\log(\text{Investment})_{it}$	-0.0302 (0.180)	-0.0483 (0.149)	-0.169 (0.190)	-0.0451 (0.0871)	0.147 (0.140)	-0.222 (0.145)
$\log(\text{Open})_{it}$	-0.00852 (0.0554)	-0.0411 (0.0814)	-0.0575 (0.116)	0.0208 (0.0261)	-0.00459 (0.0912)	-0.0194 (0.0917)
Growth $_{it}$	0.00737 (0.00459)	0.0147*** (0.00466)	0.0100 (0.00615)	0.0111** (0.00503)	0.00844* (0.00455)	0.0185*** (0.00434)
Observations	551	557	460	648	545	563
Sub-sample	High-income	Middle-income	< 2014	\geq 2014	Europe	Non-Europe

Standard errors adjusted for clustering by country in parentheses. GMM estimates are two-step GMM with corrected and robust standard errors. Government expenditure, investment and openness are expressed percentages of GDP. Estimates include a constant term which is not reported.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

E Mechanism

Table 11: Countries with and without inflation targets

	(1)	(2)
	Happiness	Happiness
Y_{it-1}	0.806*** (0.0630)	0.623*** (0.128)
$\hat{\gamma}$	5.15 CI:[4.4 - 5.75]	4.90 CI:[3.35 - 6.3]
$\pi_{it}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$	0.00294 (0.00995)	0.000389 (0.0117)
$\pi_{it}\mathbb{I}(\pi_{it} > \hat{\gamma})$	-0.00725** (0.00352)	-0.0169*** (0.00647)
Unemployment _{it}	-0.0132** (0.00632)	-0.0301** (0.0153)
log(Govt Expenditure _{it})	0.252** (0.114)	0.282 (0.341)
log(Investment _{it})	0.0378 (0.139)	-0.119 (0.242)
log(Open _{it})	-0.0000439 (0.0610)	0.00855 (0.0843)
Growth _{it}	0.0124*** (0.00423)	0.00935** (0.00385)
Observations	542	566
Inflation Targets?	Yes	No

Corrected and robust standard errors from two-step GMM in parentheses. 95% confidence intervals in square brackets. Estimates include a constant term which is not reported.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 12: Real income and Inequality effects

	(1)	(2)
	Happiness	Happiness
Y_{it-1}	0.597*** (0.0522)	0.551*** (0.0762)
$\hat{\gamma}$	4.90 CI:[4.65 - 5.65]	5.10 CI:[4.7 - 5.65]
$\pi_{it}\mathbb{I}(\pi_{it} \leq \hat{\gamma})$	0.0170** (0.00782)	0.00294 (0.0107)
$\pi_{it}\mathbb{I}(\pi_{it} > \hat{\gamma})$	-0.00617** (0.00289)	-0.0191*** (0.00502)
Unemployment $_{it}$	-0.0204*** (0.00635)	-0.0310*** (0.00862)
$\log(\text{Govt Expenditure})_{it}$	0.198** (0.0956)	0.641*** (0.186)
$\log(\text{Investment})_{it}$	-0.0217 (0.0907)	0.0485 (0.127)
$\log(\text{Open})_{it}$	-0.0974** (0.0447)	-0.0518 (0.0804)
Growth $_{it}$	0.0100*** (0.00296)	0.00921** (0.00405)
$\log(\text{GDP})_{it}$	0.258*** (0.0483)	
Inequality $_{it}$		0.00229 (0.00541)
Observations	1108	745

Corrected and robust standard errors from two-step GMM in parentheses. 95% confidence intervals in square brackets. Estimates include a constant term which is not reported.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$