

DISCUSSION PAPER SERIES

DP17224

How likely is an inflation disaster?

Jens Hilscher, Alon Raviv and Ricardo Reis

FINANCIAL ECONOMICS

MONETARY ECONOMICS AND FLUCTUATIONS

CEPR

How likely is an inflation disaster?

Jens Hilscher, Alon Raviv and Ricardo Reis

Discussion Paper DP17224

Published 15 April 2022

Submitted 13 April 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Financial Economics
- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Jens Hilscher, Alon Raviv and Ricardo Reis

How likely is an inflation disaster?

Abstract

The prices of long-dated inflation swap contracts provide a much-used estimate of expected inflation at far horizons. This paper develops the methods to estimate complementary tail probabilities for persistently very high or very low inflation using the prices of inflation options. For the object of interest—inflation disasters at long horizons—we show that three adjustments to conventional measures are crucial: for real payoffs, risk, and horizon. Applying the method to the United States (US) and the Eurozone (EZ) we find that (i) the probability of US deflation in 2011-14 was not very high, (ii) the tail probability of a deflation trap in the EZ post 2015 has been high throughout in spite of varying policies meant to address this issue, as well as shocks, and (iii) there was a significant steady rise in 2021 in the risk of persistent high US inflation, and a sharp rise in 2022 in the EZ.

JEL Classification: E31, E44, G13

Keywords: Option prices, inflation derivatives, Arrow-Debreu securities

Jens Hilscher - jhilscher@ucdavis.edu
UC Davis

Alon Raviv - alon.raviv@biu.ac.il
Bar-Ilan University

Ricardo Reis - r.a.reis@lse.ac.uk
London School of Economics and CEPR

Acknowledgements

We are grateful to Daniel Albuquerque, Marina Feliciano, Rui Sousa, and Borui Zhu for excellent research assistance, and to seminar participants at Markus Academy, Johns Hopkins, LSE, PUC Chile, CEPR RPN communication series, VIMM, and the University of Glasgow. This project received funding from the European Union's Horizon 2020 research and innovation programme, INFL, under grant number No. GA: 682288. First draft: February 2022.

How likely is an inflation disaster?*

Jens Hilscher
UC Davis

Alon Raviv
Bar-Ilan University

Ricardo Reis
LSE

April 2022

Abstract

The prices of long-dated inflation swap contracts provide a much-used estimate of expected inflation at far horizons. This paper develops the methods to estimate complementary tail probabilities for persistently very high or very low inflation using the prices of inflation options. For the object of interest—inflation disasters at long horizons—we show that three adjustments to conventional measures are crucial: for real payoffs, risk, and horizon. Applying the method to the United States (US) and the Eurozone (EZ) we find that (i) the probability of US deflation in 2011-14 was not very high, (ii) the tail probability of a deflation trap in the EZ post 2015 has been high throughout in spite of varying policies meant to address this issue, as well as shocks, and (iii) there was a significant steady rise in 2021 in the risk of persistent high US inflation, and a sharp rise in 2022 in the EZ.

JEL codes: E31, E44, G13.

Keywords: option prices, inflation derivatives, Arrow-Debreu securities.

*Contacts: jhilscher@ucdavis.edu, alon.raviv@biu.ac.il, r.a.reis@lse.ac.uk. We are grateful to Daniel Albuquerque, Marina Feliciano, Rui Sousa, and Borui Zhu for excellent research assistance, and to seminar participants at Markus Academy, Johns Hopkins, LSE, PUC Chile, CEPR RPN communication series, VIMM, and the University of Glasgow. This project received funding from the European Union's Horizon 2020 research and innovation programme, INFL, under grant number No. GA: 682288. First draft: February 2022.

1 Introduction

The 5-year-5-year (5y5y) forward inflation expectations rate measures expected inflation in five years time over the following five years. It is a common indicator of whether long-run inflation expectations are well anchored (e.g., Gürkaynak, Levin and Swanson, 2010). Policymakers find it useful because it strips out current temporary fluctuations and averages over a long period of time, therefore providing focus on what an inflation-targeting central bank can achieve. Indeed, 5y5y inflation is much less volatile than headline inflation. In the United States inflation over the last twelve months reached 8.5% in March of 2022, the highest level since the early 1980s. Yet 5y5y inflation, which we plot in figure 1, rose much more modestly, consistent with this being a temporary rise and long run inflation staying anchored. A famous example of policymakers attention to this rate comes from the Eurozone (EZ). The decline in 2014 of the 5y5y measure well below 2% was used to justify the start of quantitative easing.¹

However, the 5y5y rate is an average and as such does not contain or reflect any information about the uncertainty around future inflation. The distribution of its values could be extremely tight or disperse. However, making decisions under uncertainty typically requires knowing the whole distribution of future inflation rates, not just their expected value. Especially important for making robust decisions is to know the probability of extreme inflation realizations, which we will refer to as inflation disasters. It is these large movements in actual inflation, often deviating from mean expectations, that are associated with large costs of inflation in popular discourse. Common references to the German hyperinflation of the 1920s or the Volcker disinflation of the 1980s attest to this fact. It is also these disasters that are associated with large costs in models of monetary policy that depart from certainty equivalence, and that dominate policymaking under a risk management approach.

This paper develops the methods to provide counterparts to figure 1 in the form of tail probabilities of inflation disasters. We use a data-driven approach making minimal assumptions about preferences for pricing risk or about inflation dynamics. To be specific,

¹Reporting from the August 2014 Jackson Hole meeting where the ECB justified its use of quantitative easing, the Financial Times noted: “Mr Draghi had highlighted the inflation swap rate...never before August’s Jackson Hole speech had a president of the ECB made such a clear link between its behaviour and policy action.”

the objects of interest are the following two probabilities:

$$\begin{aligned}\Phi_t^{dh} &= \text{Prob}[\pi_{T,T+H}/H > \bar{\pi} + d], \\ \Phi_t^{dl} &= \text{Prob}[\pi_{T,T+H}/H < \bar{\pi} - d].\end{aligned}$$

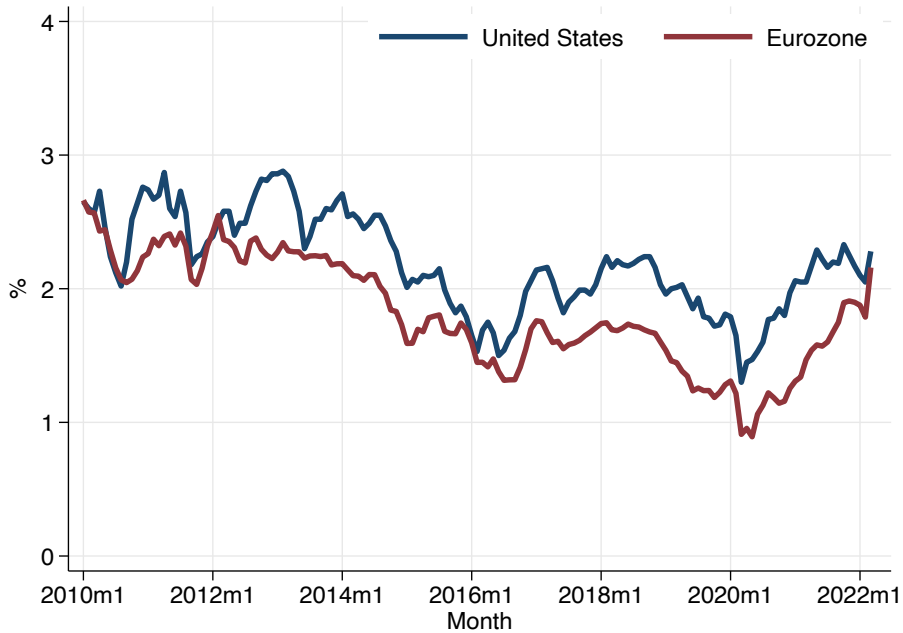
The present date is t , a future distant date is T years away, and a long horizon is denoted by H further years. Future long-term inflation is $\pi_{T,T+H}$, defined as the change in the log of the price level between the two dates in the subscript, while $\bar{\pi}$ is the inflation target, and d is the size of the disaster. These probabilities then answer the question: What is the current date t market perceived probability that inflation will be persistently above or below the $\bar{\pi}\%$ annual target between T and $T + H$? For example, what is the current probability that average inflation will be above 4% between 5 and 10 years from now?

In our empirical implementation, we provide estimates of these probabilities for the US and the EZ starting in January 2011, and setting $T = H = 5$ years, so these are 5y5y probabilities. The inflation target is $\bar{\pi} = 2\%$. For disasters we consider both high inflation and deflation, $d = 0.02$, or severe high inflation and deflation, $d = 0.03$. Note that given the 5 year horizon, high inflation is a cumulative 10 log-point deviation of inflation from target, and severely high is equal to 15 log-points, justifying the use of the word disaster. Our estimates therefore serve as a measure of the success of monetary policy at anchoring market inflation expectations, and can be used to judge the performance of different policies.

The first contribution of our paper is to provide a new method to translate the prices of traded inflation derivatives into risk-neutral and physical-measure probabilities of inflation. We show that a conventional reading of the data results in naive estimates that can provide grossly over- or under-stated estimates of the desired probabilities. These conventional readings must be adjusted in three ways. First, their units have to be adjusted so they can match Arrow-Debreu probabilities, because when inflation is high, this on the one hand raises the nominal payoff of a call option, but on the other hand lowers its real payoff. Meanwhile, when inflation is low, the payoff to a put option increases but is worth more in real terms. Probabilities based on nominal state prices are therefore too high for low inflation states since a nominal payoff of one is worth more; for high inflation, the opposite is true.² Second, options give risk-adjusted probabilities (commonly

²This point applies to other derivatives as well, so our method can be used to adjust other financial-market-based probabilities. However, for non-inflation options, this would require knowing the distribution of inflation conditional on the fundamental that the option is written on. For inflation options, that

Figure 1: Expected 5y5y inflation



referred to as risk-neutral probabilities), but since marginal utility is likely high during disasters, their prices will over-state the actual tail probabilities. Building on recent work on rare disasters we examine how inflation disasters affect the price of out-of-the-money options. We can then make an adjustment for marginal utility without having to specify the full dynamics of the stochastic discount factor that prices inflation risk. We find that periods of high inflation are indeed bad times, resulting in a large risk adjustment. In contrast, periods of deflation are associated with a much smaller drop in consumption, thus resulting in a correspondingly lower risk adjustment. Third, options pay for realizations of inflation at $\pi_{0,T}$ and $\pi_{0,T+H}$, but not for the desired forward horizon $\pi_{T,T+H}$.³ If a disaster results from the gradual unanchoring of inflation expectations, then that sluggishness requires an adjustment, otherwise both the 5-year and the 10-year probabilities will understate the probability of a 5y5y inflation disaster.

Depending on the question at hand, researchers may want to make only one or two of these adjustments. For instance, making only the real and risk (but not the horizon)

conditional distribution is a trivial point mass, making the adjustment simple.

³There are forward starting options, which we will use, but for one-year horizons ($H = 1$) as opposed to the longer horizon $H = 5$ that we would like.

adjustment, we provide estimates of average inflation being persistently d points above or below the $\bar{\pi}$ target between now and five or ten years in the future (so $T = t$ and $H = 5$ or $H = 10$). We also provide estimates of risk-neutral (or risk-adjusted) densities, so without a risk factor that would adjust for this probability reflecting elevated marginal utility in a disaster state. Of independent interest, we provide estimates of the dynamic properties of inflation, as perceived by market participants. They show a fall in stochastic volatility in the last decade, and a perception that disasters are short-lived. Finally, note that the adjustments may not all go in the same direction: for instance, the real adjustment raises the estimated probability of a high-inflation disaster, while the risk adjustment lowers it.

To judge their overall effect, we produce estimated time series of the inflation disaster probabilities Φ_t^{dh}, Φ_t^{dl} from October of 2009 (US) and January of 2011 (EZ) to March of 2022 for both.⁴ We use these measures to provide new evidence and a new perspective on three macroeconomic debates. First, we re-examine the market perceived probability of the US falling into a deflation trap in 2011-14. At the time it was judged to be very high, and justified expansionary monetary policy to fight the liquidity trap. Estimates based on our new methodology show that this probability was significantly lower than previously appreciated using the conventional measures that did not include our three adjustments. In particular the risk of short-term deflation was at times elevated, but not the risk of a deflation trap at the 5y5y horizon.

Second, we examine the probability of deflation for the Eurozone between 2015 and the present. We find that the risk of a deflation trap persisted throughout and has continued to be present until early 2022, in spite of different waves of ECB policy that tried to eliminate it. Policy since 2015 appears to have succeeded at lowering the probability of deflation in the near future, but not of a deflation trap over the long run. This justifies a mission review of the ECB that eliminates the structural features that have kept this perception of a deflation trap in the face of large policy actions. Whether the ECB's review of 2021 succeeded is too early to tell, but our estimates will allow for an assessment in a few years' time.

Third, we examine how the 5y5y probability changed in 2021 and early 2022. Policymakers throughout 2021 argued that the observed increase in inflation was temporary and that inflation expectations were anchored (Powell, 2021, Lagarde, 2021). Our estimates show a steady, significant and accelerating rise in the probability of persistent high inflation for the US since the middle of 2021. It rose above 10% by the end of 2021 and

⁴Our website provides updated probabilities at the end of every month.

approached 15% by March 2022. In the EZ there was no drift in 2021 in the market-perceived probability of a high-inflation disaster, but a remarkable fast one once we enter 2022. These numbers are relevant for policymakers navigating the resurgence of persistently high inflation, a phenomenon not seen in more than forty years. The numbers also provide hints on whether the shock to observed inflation was temporary or permanent, as well as local or global.

The paper is related to four branches of the literature. First, it focuses on tail outcomes for inflation disasters, in common with the literature on inflation at risk (Kilian and Manganello, 2007, Banerjee et al., 2020, Andrade, Ghysels and Idier, 2012, Lopez-Salido and Loria, 2020). However, we focus on market perceptions of this risk as measured by option prices, rather than on empirical distributions based on realized inflation. Because the possibility of extreme and persistent inflation events is constantly traded, they provide many more observations on the likelihood of inflation disasters that are region specific. Estimates based on empirical distributions have to pool across many countries and long periods of time with different inflation regimes. Second, while a few other papers look at expectations of disasters in surveys (Reis, 2022, Ryngaert, 2022), we take the perspective of financial markets. Very few surveys ask respondents about tail probabilities of distant-horizon inflation, and the few that do (the Survey of Professional Forecasters for the United States) move very little over time. Time series of dispersion in surveys about long-horizon inflation are more useful, but disagreement, which many surveys capture, and uncertainty, which we measure, are not the same (Reis, 2020, Coibion et al., 2021). Third, we use inflation options data to focus on tail phenomena in common with Hilscher, Raviv and Reis (2022), Mertens and Williams (2021), Kitsul and Wright (2013), Fleckenstein, Longstaff and Lustig (2017). Yet we focus on long-term forward horizons, which raises unique challenges in using the data. For the one-year ahead questions raised in some of those papers, our risk and horizon adjustments are quantitatively less important.⁵ Fourth, we draw on the literature on equity disasters to adjust for risk (Barro, 2006, Gabaix, 2012, Barro and Liao, 2021), but we focus on inflation disasters. Our method can also be used to construct probabilities for the S&P500 or for currency changes, subject to knowing the probability that large changes in those prices coincide with high or low inflation. But, options on equities or currencies almost always have short horizons, between one week and one year, for which the adjustments are less quantitatively important.

⁵Gimeno and Ibanez (2018) is closer to us in goal but imposes very restrictive assumptions in its methods.

All market-based measures of expectations have in common that they may be influenced by noise from liquidity fluctuations as well as the opinions or hedging motives of some large traders. Our measures are not immune to this concern. However, note that because of our focus on changes in the 5y5y probability, it would take systematic changes over time in the differential liquidity of 5-year and 10-year options to bias our estimates. Still, in our empirical analysis of the recent history of US and EZ inflation, we focus on trends across many months, rather than month-to-month fluctuations, since these are likely to be more robust to liquidity changes. Finally, note that since the inflation options are actively used to hedge positions in inflation swaps, liquidity concerns would likely be common between our estimates and those in figure 1. Therefore, they would not affect our goal of providing a tail probability counterpart to the 5y5y expected inflation. Insofar as policymakers and academics have found the estimates in figure 1 reliable, they should find ours useful as well.

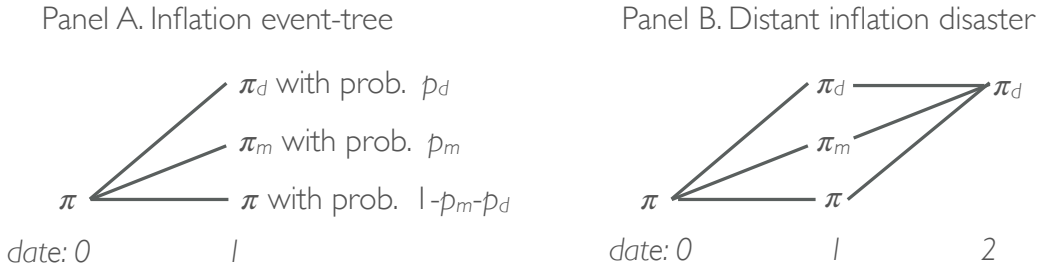
The paper is organized as follows. Section 2 lays out our approach in a simple two-period setup. Section 3 presents the general model and formally defines the probabilities of interest. Section 4 makes the real adjustment and presents time series of densities for risk-neutral probabilities, while section 5 presents our model of disasters to adjust for risk. Section 6 presents and estimates a model of inflation dynamics to adjust for correlation and build forward probabilities (horizon adjustment). Section 7 then presents our estimates to answer the three applied questions stated above. Section 8 concludes.

2 The intuition of the method in a simple setup

Consider a simple event-tree world where at present, in period 0, inflation is at its normal target level $\bar{\pi}$, at which it stays with a very high probability. However, in the next period, 1, inflation can instead rise to the moderately high level π_m , or jump to a disaster-high level π_d , according to the respective probabilities of the Markov chain p_m and p_d . We focus on a high-inflation disaster for expositional simplicity, but the arguments would be the same for a low-inflation disaster. The left panel A of figure 2 illustrates these outcomes.

The probability we are after, for now, is p_d . We have data on options that pay one nominal unit if π_d is realized in period 1, and zero otherwise. If the price of this option that pays in the disaster state at period 1 is $a_d(1)$, then by arbitrage, this price should be equal to $a_d(1) = p_d m_d \exp(-\pi_d)$. With the disaster probability, the option pays \$1, which in real terms requires an inflation adjustment $\exp(-\pi_d)$, and is discounted by the real

Figure 2: Inflation paths



stochastic discount factor m_d , reflecting the marginal utility of the future payoff.

The conventional approach in asset pricing to construct risk-neutral probabilities is to calculate: $n_d(1) = a_d(1) \exp(i(1))$. The i_1 is the nominal interest rate between the two periods, which, based on no arbitrage, must be the return from buying the three possible options.⁶ But then, it follows that $n_d(1) \geq 0$ and that $\bar{n}(1) + n_m(1) + n_d(1) = 1$. Therefore, $n_d(1)$ can be interpreted as a probability. But what does this measure?

2.1 First adjustment: risk-neutral (inflation-adjusted) probabilities

Arrow-Debreu securities instead pay one unit of consumption (not \$1) in each future state. Therefore, the price of the disaster security is: $b_d(1) = p_d m_d$. Letting $r(1)$ denote the real interest rate, the associated Arrow-Debreu probability is then $q_d(1) = b_d(1) \exp(r(1))$. It has the interpretation that, if the agents are risk-neutral, then because $m_d \exp(r(1)) = 1$, we have that $q_d(1) = p_d$, the desired object. For this reason, $q_d(1)$ is called a real risk-neutral probability.

It follows right away that:

$$q_d(1) = n_d(1) \exp(r(1) + \pi_d - i(1)) \approx n_d(1) \exp(\pi_d - \bar{\pi}) = n_d(1) \exp(d). \quad (1)$$

The approximation comes from the starting assumption that monetary policy is very credible, so that break-even inflation—the gap between the nominal and the real interest rates—is equal to the normal target inflation level. The second equality comes from recalling that the gap between disaster inflation and target inflation is what we earlier

⁶So, no arbitrage dictates: $\bar{a}(1) + a_m(1) + a_d(1) = \exp(-i(1))$.

called the disaster d . The conventionally-measured probabilities from options $n_d(1)$ must therefore be adjusted by the disaster size.

If we are calculating probabilities near the inflation target (so d is close to 0), as sometimes is done by central banks, then this adjustment factor is negligible. Likewise, even for $d = 0.03$, if the horizon is short, then the adjustment factor is quantitatively not that significant. However, if we are looking at disasters over long horizon, say 10 years, then the adjustment factor is $\exp(10 \times 0.03) = 1.35$. Reporting $n_d(1)$ from the price of a well out-of-the-money long-dated inflation option significantly underestimates the risk-neutral probability $q_d(1)$.

The intuition for this adjustment factor is simple: when the disaster happens, and the option pays, its \$1 is now worth less in real terms. Economic agents understand this, and so pay less for this option than if they were suffering from money illusion. Reading the low price, a researcher would be misled to thinking that the agents are placing a low likelihood on this event happening.

2.2 Second adjustment: actual (or physical measure) probabilities

The next adjustment is the more familiar one to financial economists: only with an estimate of the stochastic discount factor in the inflation disaster state, m_d , can we go from the risk-neutral to the actual probabilities. Importantly, to answer the question in this paper, one does not need a full model of risk. In the simple example in this section, one does not need m_n and m_m , for instance. Only the risk that is correlated with inflation in disaster times is relevant.

Moreover, it is likely that only the risk-adjustment in the disaster state will be large. Normal times will, on average, have normal marginal utility growth. This implies that our focus on the adjustment in the disaster state also represents the largest adjustment for the three states in this simple setup.

Imagine then that the main source of variation in the stochastic discount factor is whether there is a consumption disaster or not. So, $m_d(\cdot)$ is a function of consumption, which can either be normal or in a disaster. Conditional on an inflation disaster, \tilde{p} is the conditional probability that there is a consumption disaster as well. Because a consumption disaster is a time of elevated marginal utility, then the ratio of $m_d(\cdot)$ when there is a consumption disaster, to the marginal utility when there is none, call it \tilde{m} , is well above 1. Finally, continuing with the approximation that disasters are small-probability events, so that the marginal utility without a disaster is approximately equal to the expected one,

the risk-neutral probability is:

$$q(1) = ((\tilde{m} - 1) \tilde{p} + 1) p_d. \quad (2)$$

Since $\tilde{m} > 1$, the risk-neutral probability will over-state the probability of an inflation disaster, because inflation disasters are on average states of the world with high marginal utility. The rare disasters literature has argued that \tilde{m} can be quite large, and that the probability of a consumption disaster is not so small, a fact that is important when pricing equities that fall sharply when there is a consumption disaster. For inflation though, the picture is a bit different. First, the relevant probability is \tilde{p} : that conditional on an inflation disaster, there is a consumption disaster. This is well below one. There are many times, especially outside the United States, where inflation has been reasonably high or low without any sharp fall in economic activity. Second, \tilde{m} , though larger than 1, is on average lower than in a consumption disaster state, because historically, there are several episodes where an inflation disaster came with only a mild recession. Therefore, the adjustment for risk is not as dramatic as the one in the literature on the equity premium.

What the formula shows is that in order to calculate the adjustment factor the two relevant quantities to measure are \tilde{m} and \tilde{p} . Since there is already a well-established literature measuring them for consumption disasters, and since we have corresponding data on inflation, combining the two provides a path forward to identify the two parameters.

2.3 Third adjustment: forward (horizon-adjusted) probabilities

Imagine now that there is an extra period, and that the goal is to measure the probability of having a disaster in period 2. Then, as the right panel of figure 2 illustrates, the probability of having such a disaster, from the perspective of the present, is: $p_m p_{md} + p_d p_{dd} + (1 - p_m - p_d) p_d$. The price of an Arrow-Debreu security that paid one unit of consumption in period 2 if there is an inflation disaster would provide an estimate of the risk neutral probability, $q_d(2)$.

However, we do not have the option prices that match this security. Say that there are options that pay if there has been a disaster that lasted two periods. Those would provide an estimate of $p_d p_{dd}$, clearly understating the desired probability. Especially when inflation is elevated because of transitory factors, this understatement can be substantial.

Likewise, say that there are options that pay if there is a disaster in the first period, and so provide an estimate of p_d . Since inflation often moves sluggishly, $p_m p_{md} / p_d$ is quite

large: starting from a normal level, it is considerably more likely for inflation to build up to a moderately high level and then be in the disaster state in period 2, than for it to get there straight away in period 1. This makes this second probability also under-state the desired probability $q_d(2)$.⁷

Therefore, looking at either short-dated, or long-dated cumulative options, may under-state the forward probability. With two pieces of information, and three needed transition probabilities to calculate $q_d(2)$, clearly one needs one more piece of information. Intuitively, one needs some information on the extent to which inflation is sluggish. Fortunately, there are traded forward starting options, but for sub-periods of our hypothetical period 2. Namely, in the actual data, there are forward contracts for annual inflation within our 5-year desired periods. While these options are not always the most liquid, and so the inflation distributions data have to be constructed carefully, they provide the missing piece of data because their higher frequency provides an estimate of the sluggishness of inflation to enter a disaster state.

3 The theoretical result

This section derives the key theoretical result on the three adjustment factors to go from option prices to get the probability of inflation disasters.

Uncertainty about inflation: Every date, there is a state of the world s drawn from a countable set S with a probability distribution $\hat{p}(s)$, so that $\hat{p}(s) \geq 0$ for all s and $\sum_{s \in S} \hat{p}(s) = 1$. Inflation is a random variable, so it is a function of the state s and has an associated probability distribution $p(\pi)$. This is given by the standard formula: $p(\pi) = \sum_{s: \pi(s)=\pi} \hat{p}(s)$ which is calculated over the set of all possible values of inflation Π . The cardinality of Π may be lower than that of S because there may be some states s' and s'' such that $\pi(s') = \pi(s'')$. This paper is about the probability of inflation disasters alone, not about disasters more generally. Therefore, the goal is to recover $p(\cdot)$, not $\hat{p}(\cdot)$, so that we recover the probability of an *inflation* disaster. That probability may well average over states of the world where there are non-inflation disasters and others.

Inflation securities and inflation risk: The price in consumption units of an Arrow-Debreu security that pays one unit of consumption only if state s is realized is $\hat{b}(s) = \hat{p}(s)\hat{m}(s)$, where $\hat{m}(s)$ is the discounted marginal utility in that state relative to today. This

⁷The precise condition for it to understate it is: $p_m p_{md} / p_d > p_m + p_d - p_{dd}$.

is because the consumer in an Arrow-Debreu world must be indifferent between consuming one unit today, or buying $1/\hat{b}(s)$ securities that with probability $p(s)$ pay $m(s)$ utility units relative to today, so Arrow-Debreu prices twist probabilities by the marginal utility of consumption.

Assuming a full set of Arrow-Debreu securities, i.e. complete markets, is a strong data requirement. However, consider a related set of inflation securities that pay off one unit of the consumption good if inflation is π at that future date. We assume throughout that there is no arbitrage in trading inflation risk. Inflation is an aggregate variable, on which there is little inside information by any particular investor, and which is monitored by some of the largest passive investors, as well as by many speculators. By no-arbitrage, it must be that their price is $b(\pi) = \sum_{s:\pi(s)=\pi} \hat{b}(s)$. But then, it follows that:

$$b(\pi) = p(\pi)m(\pi), \quad (3)$$

where $m(\pi) = \sum_{s:\pi(s)=\pi} \hat{p}(s)\hat{m}(s)/p(\pi)$: the average marginal utility across all the states of the world where inflation is the same. The average arises because there may be states with the same level of inflation but different marginal utility: s' and s'' such that $\pi(s') = \pi(s'')$ but for which $m(s') \neq m(s'')$. As a result, $m(\pi)$ will vary *only* with inflation, or carry inflation risk, while averaging across all other sources of risk in the economy. This pattern is present in the data – over the last twenty years, the US economy has gone through booms and busts, but inflation has been approximately unchanged.

Risk-neutral probabilities: Consider an alternative security that pays one unit of consumption, no matter what the state of the world is. The inverse of the price of this security is e^r , where r is the net real interest rate. Since this security has an identical payoff as buying one inflation security for each possible value of inflation, it follows that by no-arbitrage: $e^{-r} = \sum_{\pi} b(\pi) = \sum_{\pi} p(\pi)m(\pi)$. Therefore, as is standard, e^{-r} is the expected SDF or marginal utility of consumption growth. Because prices are non-negative, then we can define $q(\pi) = b(\pi)e^r$. It is non-negative and adds up to 1. It is the risk-neutral probability of this inflation rate.

The securities described so far do not exist and so their prices cannot be easily observed in the data. A different security, that matches what is traded in financial markets, pays not one unit of consumption, but rather one nominal unit at the future state-date. Again, by no-arbitrage, its price is $a(\pi) = b(\pi)e^{-\pi}$. If inflation is high, this is lower than that of $b(\pi)$, because the nominal unit delivered by this security is worth less in real

terms than that of the inflation security. The net nominal interest rate i is likewise defined as the inverse of the price of a security that delivers one nominal unit for sure next period $e^{-i} = \sum_{\pi} b(\pi)e^{-\pi}$.

Combining these two, one can define the “nominal risk-neutral probability” $n(\pi) = b(\pi)e^{i-\pi}$, which is itself non-negative and adds up to 1. Finally, let $\pi^e = i - r$, be expected inflation. It immediately follows that:

$$q(\pi) = n(\pi)e^{\pi-\pi^e}. \quad (4)$$

Time and horizons: Date 0 is the present, at which all probabilities will be calculated conditional on what is known now, while t denotes a future date. The joint risk-neutral probability density of inflation over the first T periods and over the remaining H periods is $q(\pi_{0,T}, \pi_{T,T+H})$. From the definition of marginal and conditional distributions: $q(\pi_{T,T+H}) = \sum_{\pi_{0,T}} q(\pi_{0,T}, \pi_{T,T+H})$ and $q(\pi_{T,T+H}|\pi_{0,T}) = q(\pi_{0,T}, \pi_{T,T+H})/q(\pi_{0,T})$. Finally, because of the definition of inflation, $\pi_{T,T+H}H = \pi_{T,T+1} + \pi_{T+1,T+2} + \dots + \pi_{T+H-1,T+H}$, and there is a joint distribution of $q(\pi_{T,T+1}, \pi_{T+1,T+2}, \dots, \pi_{T+H-1,T+H})$. Combining all of these probabilities and re-arranging, $q(\pi_{T,T+H})$ equals:

$$q(\pi_{0,T+H}) \sum_{\pi_{0,T}} \left[\left(\sum_{\pi_{T,T+1} + \dots + \pi_{T+H-1,T+H} = \pi_{T,T+H}} q(\pi_{T,T+1}, \dots, \pi_{T+H-1,T+H}|\pi_{0,T}) \right) \frac{q(\pi_{0,T})}{q(\pi_{0,T+H})} \right] \quad (5)$$

The expression within the round brackets takes into account the persistence of inflation across periods within the interval of time $(T, T + H)$. Multiplying it, and so within the square brackets, is the sluggishness of inflation, which requires adjusting for the relative probability that given a longer horizon $T > t$ inflation may build up.

Final result: Combining all the steps, we get the following result:

Proposition 1. *The probabilities of high and low inflation disasters are, respectively: $\Phi_t^{dh} =$*

$\sum_{\pi_{T,T+H} > H(\bar{\pi}+d)} p_t(\pi_{T,T+H})$ and $\Phi_t^{dl} = \sum_{\pi_{T,T+H} < H(\bar{\pi}-d)} p_t(\pi_{T,T+H})$ where:

$$\begin{aligned}
p_t(\pi_{T,T+H}) &= \underbrace{n_t(\pi_{T,T+H})}_{\text{Options Data}} \\
&\times \underbrace{\left(e^{(\pi_{T,T+H} - \pi_{T,T+H}^e)H} \right)}_{\text{Real Factor}} \\
&\times \underbrace{\left(e^{-r_{T,T+H}H} m(\pi_{T,T+H}) \right)}_{\text{Risk Factor}} \\
&\times \underbrace{\sum_{\pi_{0,T}} \left[\left(\sum_{\dots=\pi_{T,T+H}} q(\pi_{T,T+1}, \dots, \pi_{T+H-1,H} | \pi_{0,T}) \right) \frac{q(\pi_{0,T})}{q(\pi_{0,T+H})} \right]}_{\text{Horizon Factor}} \quad (6)
\end{aligned}$$

The proposition characterizes the adjustment factors in some generality. Each of the next three sections discusses how to implement them, and the data that we use.⁸

4 Risk-neutral probabilities: real adjustment

When data on inflation options is used to report probabilities, typically what is shown is $n_t(\pi_{T,T+H})$, what we have called the nominal risk neutral probability. Here we discuss their meaning, and why a real adjustment factor should always be used when measuring the probability of inflation.

4.1 Data

There is an active market for US and EZ inflation options. The same players that buy and sell nominal and inflation-indexed government bonds, or that trade in the inflation swap markets, will potentially also be present in these option markets to hedge some of their positions from the other markets. Therefore, even though trading volumes will

⁸An open question, common to all market prices, is whether or not there should be a fourth adjustment factor that captures the effect of potentially time-varying illiquidity of the option contracts. Not only is it difficult to measure such an effect, the direction of its impact is not obvious: since all option prices move together with the real risk-free rate, our measures would be distorted only by movements in the differential liquidity of options with strike prices that are nearer or more out of the money. It would take a new paper to investigate appropriate adjustments for liquidity in these options and to analyze whether or not any potential adjustment are economically large. We leave this task for future work.

differ, these data are as good (or as poor) as that behind figure 1, which is used very frequently.⁹

Price data exist for both call and put options for average inflation between the present and up to 15 years away for strike prices between -2% to 6% with 0.5% jumps. The typical call security with a strike price k pays at the future date the difference between the gross inflation rate (e^π) until that date and the strike price k if this is positive, or zero otherwise. The price of that option today is $a(k)$.

We use US data from October of 2009 to March 2022 from Bloomberg, while for the EZ the start date of the sample is January of 2011. While option prices are available daily, sometimes the data quality is low. Since options are not traded simultaneously, option pricing functions can violate basic options properties. Therefore to be conservative in its use, we construct data at the monthly frequency by choosing the trading days where the market prices of the options do not violate the properties of option prices, as summarized in the appendix. We focus on horizons of 5 years, and 10 years, which are also two of the more liquid markets for these securities. The appendix describes how we construct the data.¹⁰

4.2 Recovering nominal probabilities

The no-arbitrage pricing condition for the traded securities are:¹¹

$$a(k) = \sum_{\pi} \left(p(\pi) m(\pi) \max \left\{ \frac{e^\pi - k}{e^\pi}, 0 \right\} \right). \quad (7)$$

Following the seminal contribution of Breeden and Litzenberger (1978), it is convenient to approximate this by assuming a continuum of inflation states. Further using the definition of the Arrow-Debreu prices in equation (3):

$$a(k) = \int_k^\infty \left(\frac{e^\pi - k}{e^\pi} \right) b(\pi) d\pi. \quad (8)$$

⁹Baumann et al. (2021), Feldman et al. (2015) describe the use of these options data at the ECB and the Fed

¹⁰These options are traded over the counter, so a valid concern is whether inflation disasters are also times when there is a higher likelihood that the sellers of the options default on their contracts. If so, this would show up in the price of other options sold by the same intermediaries. While this might have been a concern at the start of our sample, there is no indications that it is significant for most of the period that we cover.

¹¹Note that the payoff of these securities only depends on inflation, not on the entire set of states $\tilde{p}(s)$.

Recalling the definition of $n(\pi) = b(\pi)e^{i-\pi}$, we can re-write the no-arbitrage condition as:

$$e^i a(k) = \int_k^\infty (e^\pi - k)n(\pi)d\pi. \quad (9)$$

Differentiating this expression with respect to k and using the definition of a distribution function $N(\pi)$ gives a simple formula to build this distribution:

$$N(\log(k)) = 1 + Ia'(k). \quad (10)$$

The right-hand side can be measured for different strike prices: it is how sensitive the price of the option is to the strike price. Since these strike prices are themselves inflation measures, one can easily build the whole distribution for different $k = \log(\pi)$, which is what is usually reported in the financial media.

However, from equation (3) and the definition of $n(\cdot)$, $n(\pi) = p(\pi)$ only if $m(\pi)e^{i-\pi} = 1$ that is for these conventionally-calculated probabilities to match the actual physical probabilities, it must be that there is not only risk neutrality ($m(\pi)e^r = 1$), but also that $\pi = \pi^e$ for every realization of π . But, this is only the case if there is no uncertainty about inflation. In that world, these probabilities carry little, if any, useful information; all probabilities, including the disaster probabilities in the proposition, are either trivially 1 or 0. The reason is that, even if investors are risk neutral, they still care about receiving a payoff in a high-inflation state that has lower real value. To conclude, the conventional nominal probabilities that are usually reported have no useful counterfactual in which they match the actual (physical measure) probabilities. From the perspective of theory, they are not an accurate proxy for the actual probabilities.

4.3 Risk-neutral probabilities

Instead, from the definition of $q(\pi) = b(\pi)e^r$ and equation (3), we have that $q(\pi) = p(\pi)$ as long as $m(\pi)e^r = 1$. This is the case if people are neutral with respect to inflation risk, or if the classical dichotomy holds, so inflation is uncorrelated with marginal utility. For a long horizon, as is our focus, this is not a terrible assumption, as it corresponds to believing in a long-run vertical Phillips curve, which is true in the majority of models used for monetary policy. One could obtain them from the data on $n(\cdot)$ using equation

(4). Alternatively, going back to equation (8), take derivatives with respect to k to give:

$$e^r a'(k) = - \int_k^\infty e^{-\pi} q(\pi) d\pi. \quad (11)$$

Taking another round of derivatives with respect to k gives:

$$q(\log k) = e^r k a''(k). \quad (12)$$

Since the right-hand side can be measured, this provides a way to build the Arrow-Debreu prices directly from the option prices.

Note that, given our long-run setting, we are assuming that the asset that pays off one unit of consumption in each state of the world is the risk-free asset. For short horizons, it is common to measure risk-free returns using a nominal interest rate, like the monthly T-Bill rate. Since, in the data, inflation is small and not that volatile over a few months, the nominal and real risk-free rates are similar. But, over longer horizons and for states of the world that are grouped by inflation disasters, the risk-free asset is not a nominal discount bond that pays off one unit of nominal currency in all states of the world, but instead a security that pays off one unit of consumption for sure.

4.4 Densities

Figure 3 shows the risk-neutral densities for the United States inflation for a ten-year horizon. On the left panel are the January distributions every 3 years between 2011 and 2020. They show the success of the Federal Reserve at re-anchoring inflation after the uncertainty that came with the great financial crises. In 2011, the distribution was spread out, with significant mass both on the right and left tails. As time went by, the standard deviation fell significantly. By the start of 2020, pre-pandemic, the distribution was tightly concentrated around the inflation target, with 92% of its mass between 1% and 3.5% (whereas it was 45% in 2011).

The right panel shows the evolution of the distribution in 2021 and 2022. Between January of 2020 and 2021, the distribution barely moved. There was a significant increase in the left tail during March and April of 2020 (not shown in the figure), but by December, market inflation expectations were as well anchored as they were pre-pandemic. However, from the second half of 2021 onwards the shift to the right of the median, together with the fattening of the right tail is very noticeable.

Figure 3: Risk-neutral distributions of US inflation, 10-year horizon

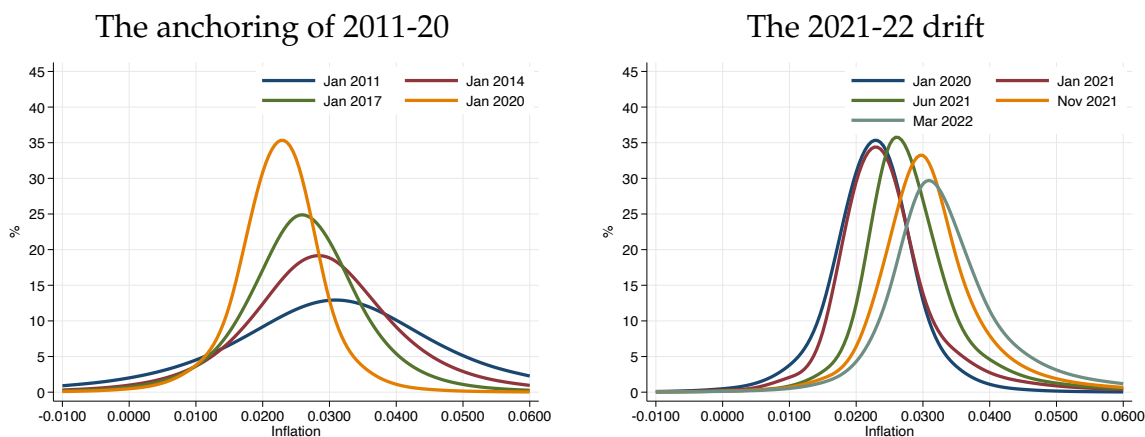
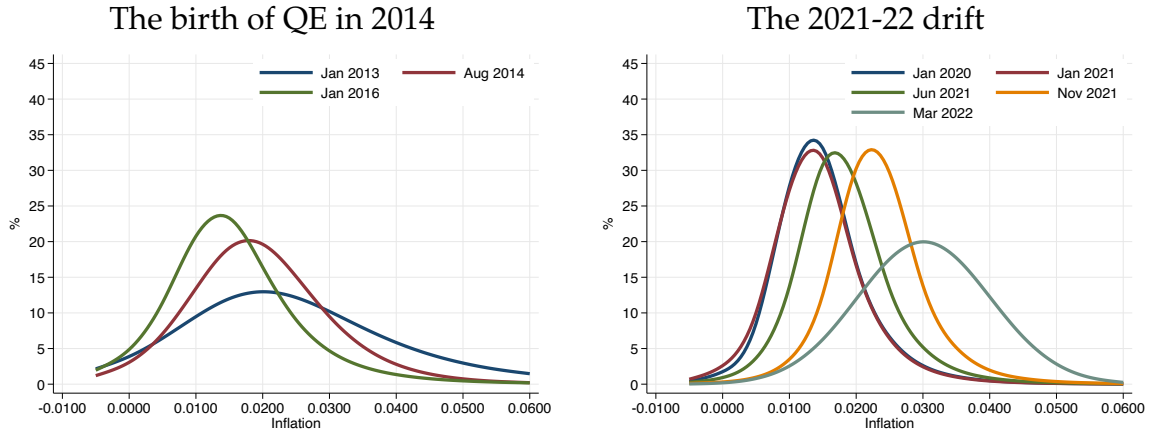


Figure 4 in its left panel shows the risk-neutral distributions between 2013 and 2016 for the Eurozone. As discussed in the introduction, the fall in the mean of the distribution between 2013 and 2014, reported in figure 1, motivated the ECB to start its program of quantitative easing. Looking at the whole distribution, perhaps there was less reason for concern at first, as the fall in the mean between January of 2013 and August 2014 was mostly driven by a decline in the far right-tail mass: the risk-adjusted 10-year horizon probability of inflation begin above 3% fell from 24% to 11%. The median was only slightly below 2%, with an accumulation of mass between 1% and 2%. However, as the figure shows, by the start of 2016, most of the mass (80% of it) was now well below 2%, with a considerable probability mass (41%) below 1%, justifying the fear that inflation expectations were anchored well below 2%.

The right panel of figure 4 provides the EZ counterpart to the US distribution shown in the right panel of figure 3. Qualitatively, the evolution during 2020 and 2021 was similar across the two areas. However, the starting point had the EZ further to the left than the US. Also, quantitatively, the shift for the US is larger and the right tail becomes thicker. Noticeably, there is a very large change in the first three months of 2022 in the EZ, much larger than the change in the US.

Figure 4: Risk-neutral distributions of EZ inflation, 10-year horizon



5 Actual probabilities: risk adjustment

If the classical dichotomy held, then at the long horizons that we consider, inflation would be uncorrelated with marginal utility. Therefore, $m(\pi)$ would be a constant, equal to the inverse of the real interest rate, and the risk-neutral probabilities would be equal to the actual probabilities.¹² However, it seems likely that inflation disasters are times where marginal utility is high. Deflation and high inflation sometimes, even if not always, come at the same time as economic recessions. If so, at the tail of the distribution, $m(\pi)$ is high, in which case risk-neutral probabilities will over-state the actual probabilities of disasters, because these events are particularly costly to investors. Relative to a full model of risk, however, we only need a model to price inflation risk at the tails.

5.1 A model of risk in inflation disasters

We assume that inflation is the sum of two parts: one capturing the ups and downs of the price level during normal times, and the other capturing its sharp jumps during disasters.

¹²Note that people may still be arbitrarily risk averse: the stochastic discount factor may still be volatile, so that $\tilde{m}(s)$ is a non-degenerate function and there is plenty of risk in the economy. But, all of it would be orthogonal to inflation, $m(\pi)$ is constant.

Either component may be correlated with consumption, so their joint dynamics are:

$$\pi_{t+\Delta} = \bar{\pi} + u_{t+\Delta}^\pi + \varepsilon_{t+\Delta} + d_{t+\Delta}^h - d_{t+\Delta}^l, \quad (13)$$

$$\log(c_{t+\Delta}) = \log(c_t) + g + u_{t+\Delta}^c + \beta_0 \varepsilon_{t+\Delta} - \beta^h d_{t+\Delta}^h - \beta^l d_{t+\Delta}^l. \quad (14)$$

where Δ is a time period, consumption c_t is expected to grow at rate g , and $u_{t+\Delta}^\pi$ and $u_{t+\Delta}^c$ are shocks to inflation and consumption, respectively, that are independent of each other. The two series co-move in normal times due to the random component ε_t , (which may be driven by multiple shocks, and may be correlated over time), where the scalar β_0 measures that co-movement that is crucial to measure inflation risk premia during normal times.

Our focus is on d_t^h, d_t^l , which are two independent common disasters between inflation and output. Disasters to inflation that do not affect consumption do not produce a risk adjustment, and consumption disasters that do not come with high or low inflation do not trigger the options. They are included in $u_{t+\Delta}^\pi$ and $u_{t+\Delta}^c$, respectively. The coefficient β^h measures the size of the consumption drop when there is a high inflation disaster; the coefficient β^l the size of the drop following a deflation disaster.

To model these disasters, we follow and modify the approaches of Gabaix (2012) and Barro and Liao (2021). With probability p^h , the high-inflation disaster d_t^h is non-zero. Defining the inverse fall in consumption in a disaster by $z^h = 1/(1 - \beta^h d)$, we assume that the size of the disaster z^h follows a Pareto distribution:

$$F(z^h) = 1 - \left(\frac{z^h}{z_0^h}\right)^{-\alpha^h} \quad \text{with } z^h \geq z_0^h > 1, \alpha^h > 0. \quad (15)$$

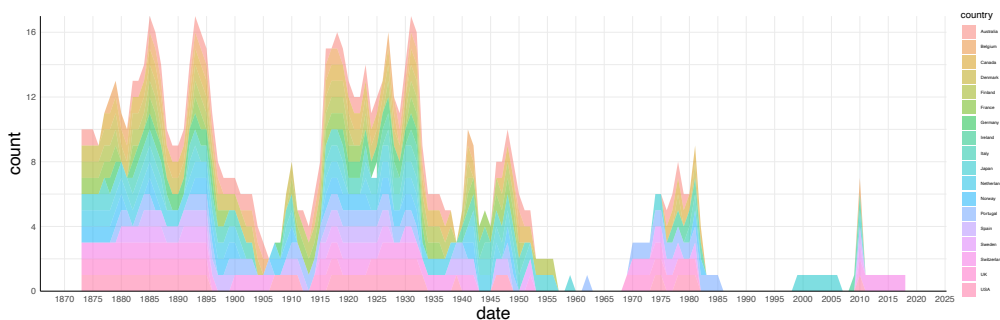
The Pareto distribution has two parameters. First, z_0^h is the minimum size of the jumps, so the higher it is, the more average consumption falls during inflation disasters. Second, the exponent α^h captures how quickly the tail of the distribution thins out, so the lower it is, the more likely is a very large consumption disaster. The same applies to (z^d, z_0^d, α^d) .

5.2 Estimating the Pareto distribution

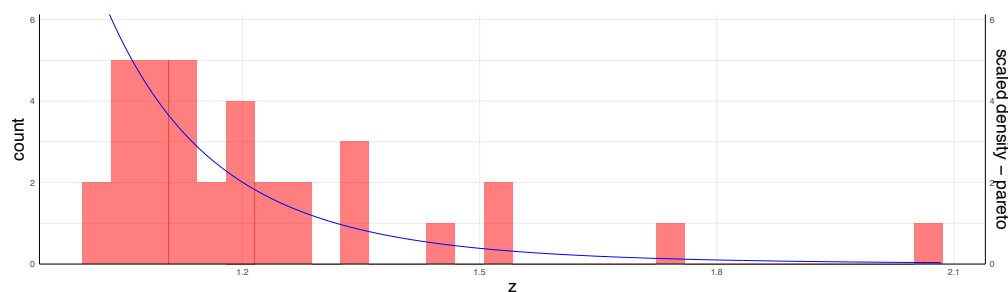
We combine data on annual output from Barro (2006) (using real GDP per capita, as he did) with data on inflation from Jordà, Schularick and Taylor (2016) between 1875 and 2015. The dataset covers 18 advanced economies, listed in the appendix. Starting with

Figure 5: Distribution of joint inflation-output disasters

(a) 150 years of inflation disasters



(b) Inverse fall in GDP (z) during an inflation and consumption disaster



one country's inflation series, we date sequential peaks and troughs by looking for local maxima and minima in 5-year rolling windows. Then, we compare the average value of inflation in a 5-year window centered around the peak (or trough), with the target level, which is taken as the trend from a band-pass filter that isolates fluctuations of frequency lower than 20 years. The appendix discusses alternatively taking 5-year fixed windows from the start of the sample, and calculating target inflation as the winsorized mean over the full sample (excluding the top and bottom quartile of observations) or instead as a moving average over the past 20 years only. Across the 6 possible methods that result from combining these, the results are quite similar.

The top panel of figure 5 shows the identified disasters across the sample. The results accord with the economic history of the time: many deflation disasters across the world in the last quarter of the 19th century and again in the 1930s, as well as three waves of high inflation disasters, after each of the World Wars and in the 1970s. Pooling all, the unconditional probability of an inflation disaster (10 log points above or below target for 5 years) is 12.9%. The key parameter for the risk adjustment though is the probability of a consumption disaster conditional on an inflation disaster. Matching the disaster 5-year

interval dates for inflation from the figure, with the dates for consumption disasters from Barro (2006), they overlap for 20.3% of the cases. Therefore, $\tilde{p} = 0.203$. Separating high and low inflation disasters, then $\tilde{p}^h = 0.374$ and $\tilde{p}^d = 0.084$.

The bottom panel plots the histograms of the observations of annual output growth for the years of joint inflation and consumption disasters, together with a simple kernel density estimate. In blue are the results from fitting the Pareto distribution in equation (15), but pooled for both high and low inflation disasters (so imposing that z^h and z^l have the same distribution). The resulting estimates are $\alpha = 6.38$ and $z_0 = 1.03$.¹³ Separating high and low inflation disasters, the estimates are: $\alpha^h = 5.45$, $z_0^h = 1.03$ and $\alpha^d = 15.18$, $z_0^d = 1.06$. That is, deflation disasters more rarely come with consumption disasters, and when they do, the falls in consumption are on average higher but with significantly thinner tails.

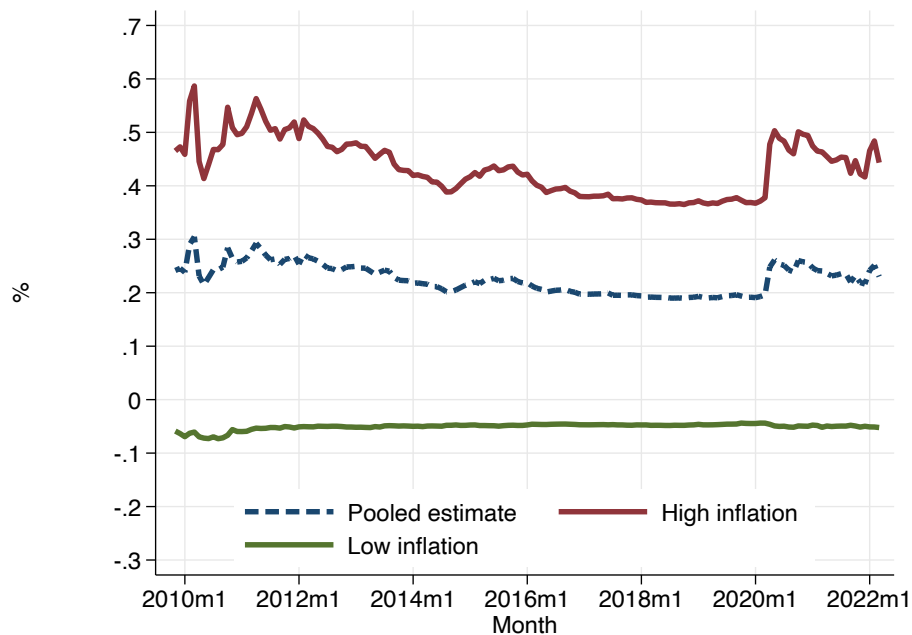
5.3 Estimating risk adjustments and risk premia

Following Gabaix (2012), Barro and Liao (2021), we then use an Epstein-Zin model for marginal utility, with a relative risk aversion coefficient of 3. Using the pooled estimates, and assuming that all of the parameters are constant over time, the risk adjustment factor is 0.82. That is, an estimate of the risk-neutral probability of $q_t(\pi_{T,T+h})$ should be multiplied by 0.82 to obtain the actual probability. Separating high and low inflation disasters, the adjustment factor is higher for the former than the latter: 0.65 versus 0.96. In fact, these results indicate that episodes of deflation, at least historically, have not been particularly risky. The reason is that across the sample there are many instances of deflation during which aggregate consumption stayed on trend.

Figure 6 plots the resulting estimates of the 10-year US inflation risk premia that come from this procedure. The risk premia are defined as $q(\pi + rp) = p(\pi)$, that is the increase in inflation to equate risk-adjusted and actual probabilities, so they are positive for high inflation and negative for deflation. The figure shows the three cases discussed so far: the pooled premium from the constant risk adjustment across both types of disasters, and the premia from treating high and low inflation disasters differently. As others before us, using quite different methods, we find only moderately high inflation risk premia for the pooled estimates: it averages to 0.23%, fluctuating between 0.19 and 0.3%. Fleckenstein, Longstaff and Lustig (2017, 2016) estimate inflation risk premia by taking the difference

¹³For comparison, Barro and Liao (2021) report α 's in the range of 6 to 8, and set $z_0 = 1.03$.

Figure 6: US inflation risk premia



between subjective expectations from analysts forecasts, and market expectations from inflation swap rates: they find risk premia in the range of 0.2-0.25%. The FRB Cleveland reports the estimated inflation risk premium from the affine term structure model of Haubrich, Pennacchi and Ritchken (2011): they average to 0.39% during our common sample.

However, the pooled estimates hide a significant difference between high inflation and deflation episodes. The average risk premium for high-inflation disasters is significantly higher at 0.61%. In contrast, the risk premium for deflation is significantly lower. Intuitively, in the data, there are many instances of deflation without a consumption disaster, so the risk-adjustment is smaller.

6 Forward probabilities: horizon adjustment

The final adjustment is to go from probabilities on cumulative inflation between the present t and a far-away date, $T + H$, to probabilities over a forward period that is far ahead in time, so between T and $T + H$. Obtaining forward expectations of inflation is

easy (as we did in figure 1), since the linearity of the expectations operator and of inflation as the difference in logs implies that $\mathbb{E}^q(\pi_{T,T+H}) = \mathbb{E}^q(\pi_{0,T+H}) - \mathbb{E}^q(\pi_{0,T})$. But, as proposition 1 shows, to estimate probabilities, it is impossible to solely use the two distributions for cumulative inflation. It requires new data, as well as a model for the time-series sluggishness and persistence of inflation as it is perceived by markets.

6.1 Data on forward starting options

There exist markets for forward-dated options at date t that will pay out depending on the realizations of inflation in $\pi_{T,T+1}$. That is, these options are for inflation in one given year, not on the average over a longer period $H > 1$ as we would like. These data were used to estimate general stochastic processes for inflation in Hilscher, Raviv and Reis (2022), and are described there in detail, as well as in the appendix. In this paper, we use the data for one-year ahead inflation covering the one-year periods starting in 5 to 9 years; this adds 5 additional distributions that we can use.

The markets in which these trade are not as liquid, so we want to be conservative in using them. We find that all five of these one-year distributions are quite similar for almost all of our data. This indicates that a low-order Markov process with not too much persistence is an adequate model since, after 5 years, the marginal risk-adjusted distribution of inflation seems to have settled at its ergodic state. Therefore, and to allow for the possibility of data concerns, we take the average of these 5 annual distributions and use that alone for estimation, making our approach more robust to the presence of measurement noise. Using the adjustments discussed in section 4, this provides an estimate of $q(\pi_{5,6})$.

Strike prices for inflation options come in jumps of 0.5%. Correspondingly, we consider distributions for inflation in 8 bins: $\pi(i) = \{\leq -1, (-1, 0], (0, 1], (1, 2], (2, 3], (3, 4], (4, 5], > 5\}$. Our data to estimate the dynamics of inflation consist then of 21 numbers per month, for the distributions of $q(\pi_{0,5})$, $q(\pi_{0,10})$, and $q(\pi_{5,6})$. These refer to risk-neutral inflation so the model of dynamics is for risk-neutral inflation as well.

6.2 A model of inflation persistence

We start from the model of inflation dynamics laid out in the previous section, repeated here for convenience:

$$\pi_{t+\Delta} = \bar{\pi} + \varepsilon_{t+\Delta} + d_{t+\Delta}^h - d_{t+\Delta}^l. \quad (16)$$

In the previous section, we did not have to specify assumptions on the normal component of inflation ε_t . Now, we do. We make two major assumptions (and a series of minor ones). First, that the variance of ε_t is small relative to the size of the disaster jumps, so that inflation enters the disaster range only as a result of a disaster, or if inflation in the previous year was just below disaster levels. Second, that if Δ was infinitesimally small, then ε_t would approximately follow a mean-reverting Ito process with continuous sample paths in time. The result of these two assumptions is that inflation in bins follows a first-order Markov process with a particular set of restrictions on the Markov transition matrix.¹⁴

To see this, consider a discrete approximation of this process as a Markov chain where inflation can be in one of 8 states corresponding to the bins in the data. The Markov transition matrix \mathbf{P} is then 8×8 , where as usual elements in each row add up to 1. We consider the following specific model:

$$\mathbf{P} = \begin{bmatrix} 1 - 5p_l & p_l & p_l & p_l & p_l & p_l & 0 & 0 \\ p_{dl} + p_{nn} & p_{ml} & p_{mr} & 0 & 0 & 0 & 0 & 0 \\ p_{dl} & p_{nn} & p_m & p_{mr} & 0 & 0 & 0 & p_{dh} \\ p_{dl} & 0 & p_{nn} & p_n & p_{nn} & 0 & 0 & p_{dh} \\ p_{dl} & 0 & 0 & p_{nn} & p_n & p_{nn} & 0 & p_{dh} \\ p_{dl} & 0 & 0 & 0 & p_{mr} & p_m & p_{nn} & p_{dh} \\ 0 & 0 & 0 & 0 & 0 & p_{mr} & p_{mh} & p_{dh} + p_{nn} \\ 0 & 0 & p_h & p_h & p_h & p_h & p_h & 1 - 5p_h \end{bmatrix}. \quad (17)$$

Starting with the low-inflation disaster state in the first row, the economy exits with probability $5p_l$, which should be close to 1 to match the Poisson-Pareto assumption on disasters. When the disaster disappears, the economy will return to any one of the normal (non-disaster) values, though not to the state opposite and closest to the other disaster. We assume that they are equally likely reflecting the first-order Markov assumption that where it was before the disaster would not affect where it ends up now. Symmetrically, the same arguments explain the 8th row referring to the high-inflation disaster.

Turning to when inflation is close to 2%, in the third and fourth row, it may move up or down according to its normal process symmetrically with probability p_{nn} . This captures the normal inflation dynamics. Inflation may be hit by the high-inflation disaster with

¹⁴Mertens and Williams (2021) compute forward distributions under the much stronger assumption that inflation follows a Gaussian random walk.

probability p_{dh} , or with the low-inflation disaster with probability p_{dl} .

Finally, on the 2nd and 3rd (and 6th and 7th) rows, a final ingredient appears, as there is mean reversion in the normal inflation component. The probability of staying close to the target is p_n , and the probability of staying above (or below) the target is p_m .¹⁵ The probability of reverting towards target is p_{mr} , which in the data we find to be much higher than the probability of staying at that level.¹⁶

All combined, there are 6 parameters to estimate with our 21 moments: the probabilities of entering a high and low disaster p_{dh} and p_{dl} , the probabilities of exiting the disaster p_d, p_l , the probability of normal inflation moving, p_{nn} , which captures the local volatility of inflation, and the probability of elevated or low normal inflation moving back to the target, capturing mean-reversion in normal inflation p_{mr} . Given an estimate of the matrix, we can then simulate many paths to calculate the probability of inflation disasters at the forward horizon.

We estimate the six parameters using GMM. We match the three sets of moments, assigning equal weights to each set of moments and minimizing the squared deviation to the target. That is, we minimize squared differences in probabilities for each of the bins in the three sets of moments. The overall fit, which we report in the appendix, is quite good and, apart from a few isolated episodes, is reasonably stable over time.

6.3 Estimating the model

In principle, we can allow the distribution to vary over time, and so estimate \mathbf{P}_t in every month, since we have 21 moments to match every month to estimate 6 parameters. However, we instead estimate one model for the entire data set that keeps three of the parameters fixed over the whole sample, while letting three others vary across months.

In the model, local volatility is captured by the p_{nn} parameter. In the data, volatility varies substantially over time and so we need to capture this important feature of inflation dynamics. However, volatility over longer horizons is affected by both the tendency to move away from the target as well as return to it. Thus, allowing p_{nn} and the mean reversion parameter p_{mr} to both vary, leads to instability in these estimates since they are not well separately identified. To improve the precision of the other estimates and since

¹⁵Note that p_n and p_m are equal to combinations of the other parameters: $p_n = 1 - 2p_{nn} - p_{dl} - p_{dh}$. Similarly, $p_m = 1 - p_{nH} - p_{nn} - p_{mr} - p_{nL}$.

¹⁶For completeness, and again because probabilities have to add up to 1 within rows: $p_{ml} = 1 - p_{dl} - p_{nn} - p_{mr}$ and $p_{mh} = 1 - p_{dh} - p_{nn} - p_{mr}$.

p_{nn} is the clearer candidate to measure time-varying volatility, we therefore pool all the observations in the entire sample and assume that p_{mr} is constant over time. For the US, the estimate is 0.5025, while for the EZ it is 0.4720, capturing the strong evidence for mean reversion.

Similarly our pooled model assumes that the exit probabilities for disasters, p_l, p_h , are constant. For the monthly estimation, they are close to 0.2 for a large share of the sample (as expected), motivating our choice to assume that they are not time-varying. The estimates are almost exactly equal to 0.2 for the US (0.1990, 0.1998), implying that, as soon as the US enters a disaster state, it leaves again immediately. For the Eurozone, instead they are (0.1999, 0.0617); markets expect EZ high-inflation disasters to persist.

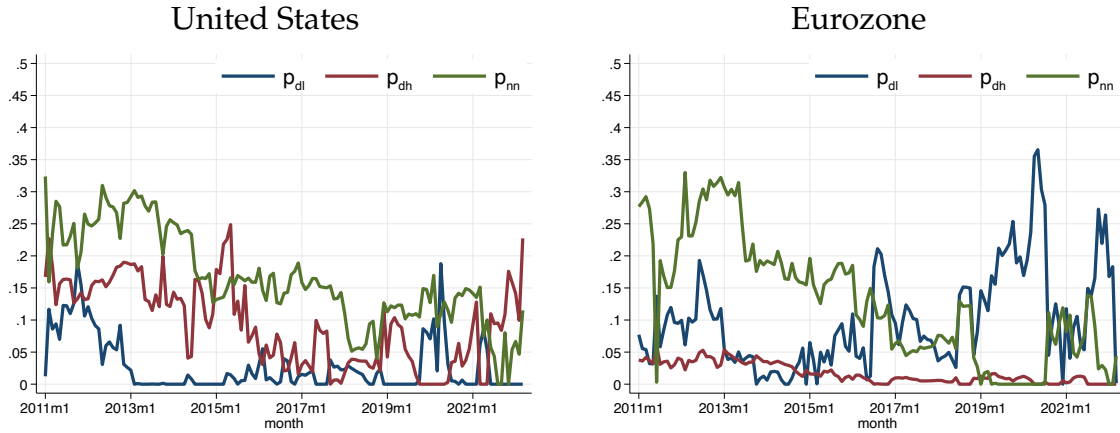
Pooling the data in this way means that we move from estimating six parameters for all of the months in our sample period to estimating three time-varying parameters plus three constant parameters all in one model. We therefore opt to reduce the frequency to quarterly and instead estimate three quarterly time-varying and three constant parameters. This approach is computationally manageable without resulting in a material loss of information. To get back to a monthly frequency for our analysis, we then re-estimate the model every month only for the three time-varying parameters, while keeping the three constant parameters at their full-sample constant estimated values.

Figure 7 shows the estimates of the three key parameters over time. In the left panel, for the US, the decline in p_{nn} since the start of the decade captures a fall in the perceived volatility of inflation. Independently of this, the probability of jumping to a low-inflation disaster was high at the start, but became quite low after mid-2012. More erratic is the pattern of probability of jumping to a high-inflation disaster. It significantly declines after 2015, but, since the start of the pandemic, it has risen significantly.

For the EZ, there is a similar decline in the stochastic volatility of inflation throughout the decade. However, the probability of a deflation disaster hitting the economy is higher than in the US throughout the sample, and varies significantly, including a significant rise in 2018-19. The probability of a high inflation disaster stays small throughout, including at the very end of the sample. In early 2022 the probability of a low inflation disaster declines markedly and there is an uptick in inflation volatility.

The appendix describes the robustness of these estimates. We estimated several other candidate models, including models with four and eight time-varying parameters, as well as one where some parameters move at an annual while others move at a monthly frequency. These other models are discussed in the appendix along with a more detailed

Figure 7: Inflation dynamics: model parameter estimates



explanation as to why we did not end up using them for our main results.

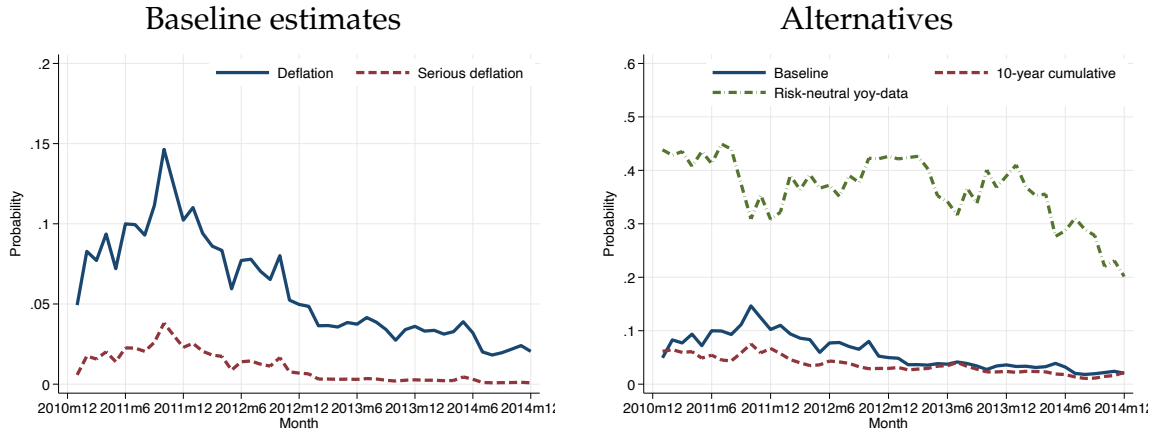
7 Estimates: inflation and deflation disasters in US and EZ

This section applies the tools we developed to measure the probability of inflation disasters for the US and the EZ, and re-examines three macroeconomic questions during this time. Throughout, our baseline estimates use our full method, and so display Φ_t^{dh} , Φ_t^{dl} , the actual probabilities of inflation disasters starting 5 years from the date in the horizontal axis, over a further 5 years. We complement these with some variants at other horizons.

7.1 US deflation fears in 2011-14

The left panel of figure 8 shows the evolution of the US probability of deflation and serious deflation (less than -1%) in the 5y5y horizon over time. At first, following the deflation of 2009, these probabilities were high and rising. Investors were perhaps doubtful of the Federal Reserve’s ability to steer inflation back on target after the large fluctuations of the previous two years. Yet, by the end of 2012, the probability of persistent deflation had fallen below 5%, and the probability of serious deflation was close to 0, staying there for many years after. Accommodative monetary policy lasted for several years, often motivated by the desire to avoid the risk of a deflation trap. According to our estimates, this fear was not justified.

Figure 8: The conquest of US deflation risk

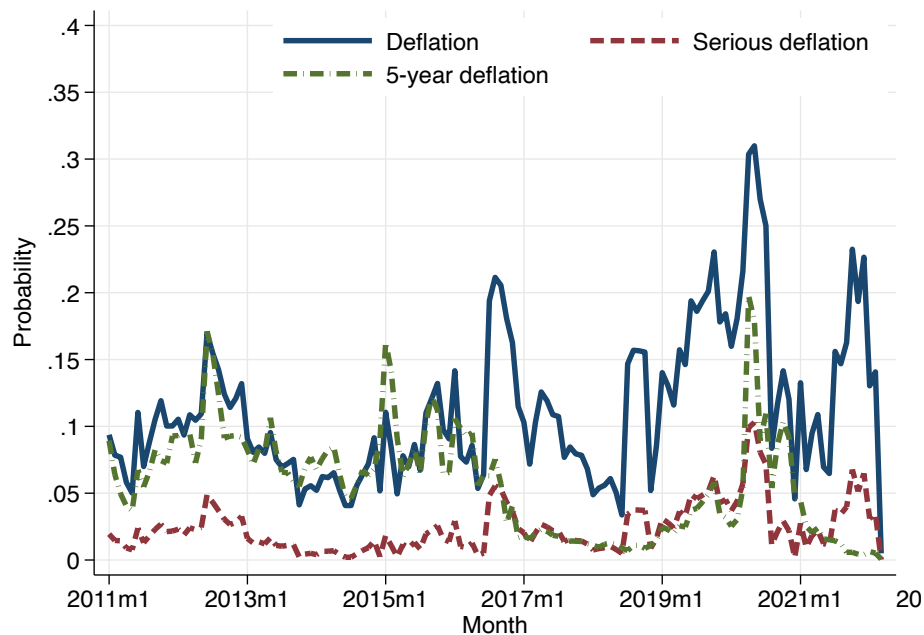


The right panel of the figure delves further into the topic by contrasting our estimate with two related ones (note that this required a wider scale). First, the figure shows also the 10-year probability (that is, doing the real and risk adjustment, but not the horizon adjustments). As expected, this was lower, as having a 10-year long deflation would be even more extreme, but not much lower.

Second, the figure shows the forward probability of deflation in a single year, averaging over the estimates from the options 5 to 9 years ahead. They measure the chance of a one-year deflation event, and do not adjust for risk, and so overstate physical measure probabilities. There is a similar downward movement in the sample period, but these numbers are much larger and more volatile. This is understandable: with current inflation over the sample period close to 2%, a one-year deflation episode is significantly more likely than deflation over a five-year period. Our Markov model estimates show both very strong mean reversion and a high probability of leaving the disaster state as soon as the economy has entered it.

Our assessment of the risk of deflation during this time is significantly smaller than the ones reported in Christensen, Lopez and Rudebusch (2015), Kitsul and Wright (2013), Fleckenstein, Longstaff and Lustig (2017). The reason is the influence of our three adjustment factors, each of which reduce the probability of deflation. First, these papers mostly focused on deflation in the near horizon and over one year, so they measured the probability of deflation, whereas our estimates are of the probability of a deflation trap, a persistent period of deflation over the long run, the event that policymakers worry most

Figure 9: Eurozone probabilities of deflation



about. Our results take into account the large horizon adjustment factor. Second, the real adjustment is significant over this large horizon, and without it the probabilities are overstated. Third, the risk premium with deflation is small, smaller than what affine models that impose a uniform risk-premium across inflation realizations would suggest.

7.2 The resilient EZ deflation-tail risk

Figure 9 shows the risk of a deflation disaster (average inflation below zero) and a serious deflation disaster (average inflation below -1%) in the Eurozone since 2011. In contrast to the United States, this risk was elevated all the way until early 2022. From 2018 until 2019, these probabilities fluctuated between high levels. The effect of the pandemic is visible, with a short-lived sharp spike in 2020Q2 (a spike is also present in the US data). More striking is that in 2021 the probability remained high, only falling in 2022. In the United States, this probability has been very close to zero since the start of 2021.

The figure shows a third series, for a deflation disaster over the next five years, so with the real and risk adjustment, but measuring the near-term rather than the 5y5y probabil-

ity, as in Boninghausen, Kidd and de Vincent-Humphreys (2018). Until 2016, this tracked the 5y5y probability. That is, the perception in markets of the probability of an inflation disaster was roughly the same over the next 5 years, or over the succeeding ones. But, from the end of 2016 onwards the probability of near deflation was low, with a short-lived spike in the middle of 2020.

What is the interpretation? The estimates suggest that the ECB's aggressive policies from 2015 onwards to contain the perception of a deflation disaster succeeded in lowering the probability of a deflation in the near term. However, the chances of falling into a deflation trap in the distant horizon did not fall, even after the increase in inflation in late 2021. If anything that probability became larger and more volatile. The market thus continues to perceive a significant structural flaw in the ECB's mandate that keeps the chance that the Eurozone will fall into a deflation trap high. Whether the recent mission review remedied this perception is too early to tell.

7.3 The pandemic and 2021 inflation fears

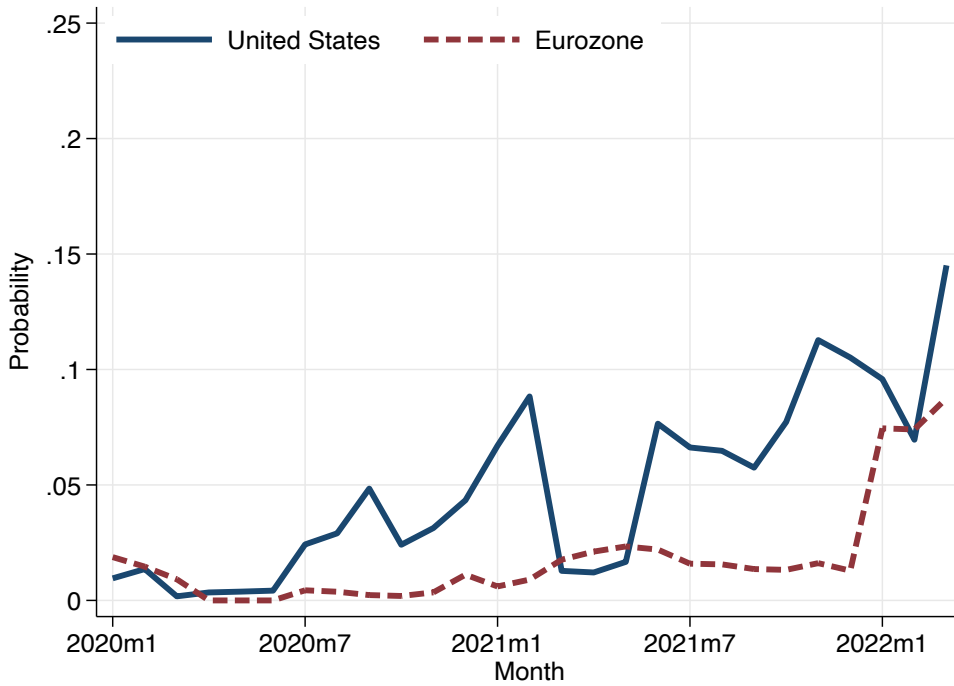
The top panel of figure 10 shows the 5-year-5-year probability of a high-inflation disaster (above 4% on average) from 2020 to early 2022 for both the US and the EZ.¹⁷ In 2020, the two were roughly similar and reasonably stable, although they started rising in the second half of the year for the US reaching 5% by the end of the year. In the second half of 2021, the two series strongly diverge, with the US probabilities steadily rising, while the EZ one stayed constant. In the first three months of 2022, the EZ probabilities sharply increased, almost catching up to the US ones.

These estimates provide a rich account of the expectations anchor during this period. In 2020 and 2021, the ECB seemed to succeed in keeping market inflation expectations anchored away from right-tail risk, even as the probability of left-tail risk remained high. Anchoring market expectations against the deflation risk seemed the bigger challenge. Yet, bad luck or bad policy led to a sharp turn in 2022, with two-sided tail risk emerging, as the probability of an inflation disaster in either direction became high. One tentative hypothesis for why the high-inflation disaster probability stayed low in the Eurozone in 2021 was that actual inflation was slower to rise there, on account of less fiscal stimulus,

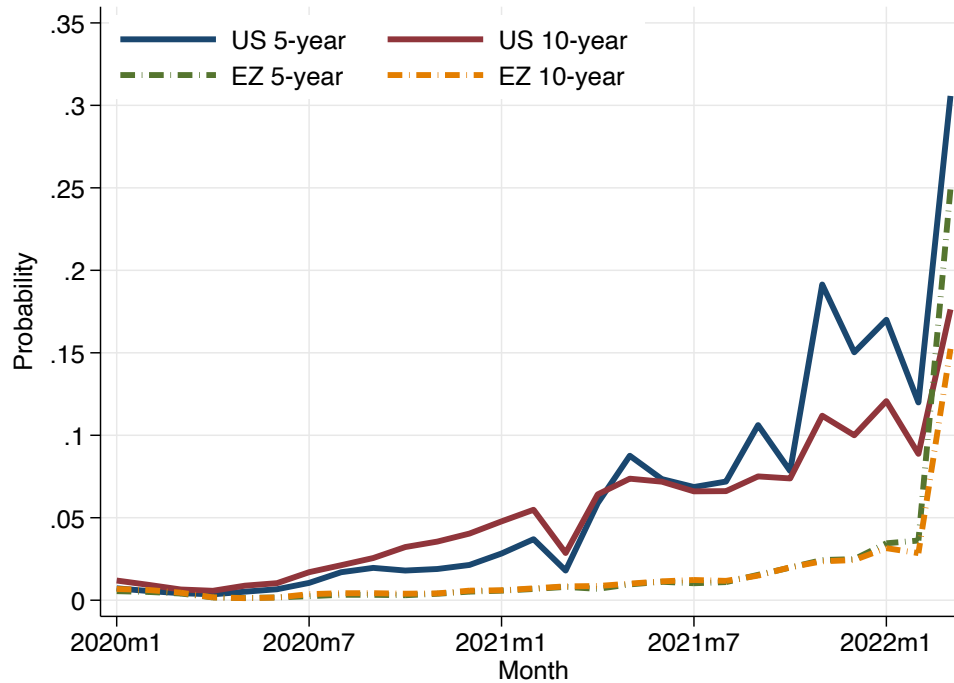
¹⁷This period was also marked by large bond purchases, of both nominal and indexed bonds by the Federal Reserve and the ECB. Insofar as these policies distort market prices, they may spillover to inflation options prices. It is much less clear whether, and in what direction, they would distort out-of-the-money options prices. A form of distortion would be violations of put-call parity conditions: in our data, these did not become more frequent during this sub-sample.

Figure 10: The risk of a high inflation disaster from 2020 to Q1 2022

(a) The 5-year-5-year (forward) horizon



(b) The 5-year and 10-year horizons



a lower starting point, and a small impact of the pandemic on labor force participation, unlike the US. Another tentative hypothesis for the large change in 2022 could be the impact of the Russian invasion of the Ukraine.

In the US, instead, the probability of a deflation risk stayed always close to zero. The market-perceived probability of high-inflation risk, however, rose steadily, especially in 2021 as inflation itself was sharply rising. By March of 2022, it reached the high value of 14%. Interestingly, this drift up of right-tail estimates is significantly more pronounced than that of the mean estimates shown in figure 1. Tail risk is drifting in a way that looking at conventional measures of expected inflation would miss. It also is in line with the general phenomenon that high inflation is often volatile inflation – in the US both are moving up, something that should and most likely does concern the monetary authority. From the perspective of economic theory, this evolution suggests that even long-horizon expectations are quite sensitive to current realizations.

The bottom panel of figure 10 digs a little deeper, by showing the 5-year and 10-year probabilities for the US and the EZ. Throughout the entire sample, before 2021, the 10-year disaster probability was always above the 5-year. As the figure shows, this changed in the US in May of 2021 and in the EZ in January of 2022, as the market perceived probability of a disaster in the next 5 years rose even more than the 5y5y horizon. The market seems to expect that some of the high inflation is temporary, but some of it will persist. Whether this is the result of over-reaction of expectations or instead a reflection of lack of credibility of monetary policy is a question left for future research.

8 Conclusion

This paper develops methods to use inflation options data to back out market-perceived probabilities for tail events in inflation. We show that producing accurate estimates requires taking into account that: (i) the options price data do not reveal the real Arrow-Debreu probabilities; (ii) the risk premium for inflation is not the same at its two tails and at the center of the distribution; and (iii) forward horizons can differ from short or long horizons because of the sluggishness of inflation. We provide some simple, but we hope robust, methods to do all of these adjustments. We show that the adjustments are quantitatively large relative to constructing probabilities using conventional methods.

In a second step, we apply our methods to data from the US and the EZ between 2009 and March of 2022. The estimates of inflation disaster probabilities lead to three

empirical findings. First, the risk of sustained future deflation in the United States in 2011-14 was quite low, in spite of policymakers' perceptions at the time. Second, there is a persistent belief in markets in the Eurozone that it may fall into a deflation disaster trap in the future. Third, in 2020-22, market perceptions of a high-inflation disaster diverged significantly between the two regions, rising significantly in the United States throughout 2021, but in Europe only sharply in 2022.

Whether in these or in other cases, probabilities of inflation disasters are informative about macroeconomic risk, monetary policy regimes, and the credibility of the central bank. They may be of particular interest in the current environment of high inflation levels, risk of stagflation, and unanchored inflation expectations (Reis, 2022). Under an inflation targeting regime, at these forward distant horizons, the success of a central bank at anchoring expectations near the target can partly be measured by whether the disaster probabilities are small. Our estimates provide inputs to future work that can use these methods and measures to diagnose the success of monetary policy regimes, and as an input in models of how economic behavior changes when agents perceive a higher chance of an economic disaster.

References

- Andrade, Philippe, E. Ghysels, and Julien Idier.** 2012. "Tails of Inflation Forecasts and Tales of Monetary Policy." Banque de France working paper 407.
- Banerjee, Ryan Niladri, Juan Contreras, Aaron Mehrotra, and Fabrizio Zampolli.** 2020. "Inflation at risk in advanced and emerging market economies." *BIS Working Papers No 883*.
- Barro, Robert J.** 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–866.
- Barro, Robert J., and Gordon Y. Liao.** 2021. "Rare disaster probability and options pricing." *Journal of Financial Economics*, 139(3): 750–769.
- Baumann, Ursel, Matthieu Darracq Paries, Thomas Westermann, Marianna Riggi, Elena Bobeica, Aidan Meyler, Benjamin Böninghausen, Friedrich Fritzer, Riccardo Trezzi, Jana Jonckheere, et al.** 2021. "Inflation expectations and their role in Eurosystem forecasting." *ECB Occasional Paper 264*.
- Boninghausen, Benjamin, Gregory Kidd, and Rupert de Vincent-Humphreys.** 2018. "Interpreting recent developments in market-based indicators of longer-term inflation expectations." *ECB Economic Bulletin, Issue 6/2018*.
- Breeden, Douglas T, and Robert H Litzenberger.** 1978. "Prices of state-contingent claims implicit in option prices." *Journal of Business*, 51(4): 621–651.
- Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch.** 2015. "Pricing Deflation Risk with US Treasury Yields." *Review of Finance*, 20(3): 1107–1152.
- Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Jane Ryngaert.** 2021. "Do You Know that I Know that You Know...? Higher-Order Beliefs in Survey Data." *Quarterly Journal of Economics*, 136(3): 1387–1446.
- Feldman, Ron, Ken Heinecke, Narayana Kocherlakota, Sam Schulhofer-Wohl, and Tom Tallarini.** 2015. "Market-based probabilities: a tool for policymakers." *Federal Reserve Bank of Minneapolis*.
- Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig.** 2017. "Deflation Risk." *Review of Financial Studies*, 30(8): 2719–2760.
- Fleckenstein, Matthias, Francis Longstaff, and Hanno Lustig.** 2016. "Inflation-Adjusted Bonds and the Inflation Risk Premium." In *Handbook of Fixed-Income Securities*, ed. Pietro Veronesi, Chapter 3. Wiley.
- Gabaix, Xavier.** 2012. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." *Quarterly Journal of Economics*, 127(2): 645–700.
- Gimeno, Ricardo, and Alfredo Ibanez.** 2018. "The eurozone (expected) inflation: An option's eyes view." *Journal of International Money and Finance*, 86: 70–92.

- Gürkaynak, Refet S., Andrew Levin, and Eric Swanson.** 2010. "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., UK, and Sweden." *Journal of the European Economic Association*, 8(6): 1208–1242.
- Haubrich, Joseph G, George Pennacchi, and Peter Ritchken.** 2011. "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps." FRB Cleveland working paper 11-07.
- Hilscher, Jens, Alon Raviv, and Ricardo Reis.** 2022. "Inflating Away the Public Debt? An Empirical Assessment." *Review of Financial Studies*, 35(3): 1553–1595.
- Jordà, Òscar, Moritz Schularick, and Alan M. Taylor.** 2016. "Macrofinancial History and the New Business Cycle Facts." In *NBER Macroeconomics Annual 2016*. Vol. 31, , ed. Martin Eichenbaum and Jonathan A. Parker. University of Chicago Press.
- Kilian, Lutz, and Simon Manganelli.** 2007. "Quantifying the Risk of Deflation." *Journal of Money, Credit and Banking*, 39(2-3): 561–590.
- Kitsul, Yuriy, and Jonathan H. Wright.** 2013. "The Economics of Options-Implied Inflation Probability Density Functions." *Journal of Financial Economics*, 110(3): 696–711.
- Lagarde, Christine.** 2021. "Monetary policy during an atypical recovery." *Speech at ECB Forum on Central Banking, Sintra, Portugal*.
- Lopez-Salido, David, and Francesca Loria.** 2020. "Inflation at Risk." FEDS working paper 2020-013.
- Mertens, Thomas, and John Williams.** 2021. "What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices." *American Economic Review*, 111(8): 2473–2505.
- Powell, Jerome H.** 2021. "Monetary Policy in the Time of COVID." *Speech at economic policy symposium of FRB Kansas City, Jackson Hole, Wyoming*.
- Reis, Ricardo.** 2020. "The People versus the Markets: A Parsimonious Model of Inflation Expectations." *CEPR discussion paper 15624*.
- Reis, Ricardo.** 2022. "Losing the Inflation Anchor." *Brookings Papers in Economic Activity*, forthcoming.
- Ryngaert, Jane.** 2022. "Inflation Disasters and Consumption." *Journal of Monetary Economics*, forthcoming.

Online Appendix to “How likely is an inflation disaster?”

Jens Hilscher
UC Davis

Alon Raviv
Bar-Ilan University

Ricardo Reis
LSE

April 2022

This appendix is split into three sections that explain: how we obtain the probability distributions for inflation from option prices; how we calculate the risk adjustment factors; and how we estimate inflation dynamics to adjust for the horizon.

A Constructing the marginal distributions of inflation

The paper uses data on two sets of distributions. First, with zero-coupon inflation caps and floors options at date t , we construct distributions of cumulative inflation from t for 5 and 10 year horizons using the formula in section 3 in the paper. Second, using year-on-year caps and floors on inflation we construct forward distributions for one-year periods starting in five to nine years. The data are from Bloomberg, for the United States (US) and the Eurozone (EZ). Our data cleaning and construction process closely follows Hilscher, Raviv and Reis (2022). Relative to their work, we use fewer maturities, have a higher frequency (monthly rather than annual), and build distributions for the EZ as well as the US.

A.1 Data pre-cleaning

Before starting the construction of the inflation distributions, we pre-process the data. The raw data includes both data errors as well as data points that are based on trades at different times of the day. This lack of simultaneity means that option prices may not pass some basic screens. We only use data if it passes the following requirements: (1) cap and

floor premia are monotonic in the strike price, (2) cap and floor premia increase monotonically with maturity, (3) butterfly spreads, which represent one way of constructing nominal Arrow-Debreu security payoffs, have positive prices, and (4) the put-call parity implied real rates are consistent across strike prices.

A.2 Implied volatility smoothing

We then transform the data and calculate Black and Scholes (1973) implied volatilities. This nonlinear transformation makes it easier to adjust for data inaccuracies and errors. Black and Scholes (1973) implied volatilities of the cap and floor contracts are smoother than the prices of the options. We therefore follow Shimko (1993) and use implied volatilities to interpolate and smooth the data. We fit the SABR model, the four-factor stochastic volatility model developed by Hagan et al. (2002) for each maturity. We search for the set of parameters that minimizes the norm of the difference between model and actual volatilities. We constrain the SABR parameters to ensure that the smoothing does not introduce any arbitrage opportunities in option prices. In this way we construct a smoothed maturity-specific implied volatility function, which we then use to convert back to option prices.

For the year-on-year data, we first extract individual caplet and floorlet prices from the market prices of caps and floors. We then use the Rubinstein (1991) transformation to price forward starting options based on their specific option tenor, which is the time between reset dates. We discount using the real interest rate which is extracted from the put-call parity relationship of the zero coupon options (Birru and Figlewski, 2012). For the individual caplet and floorlet prices we then follow the same SABR implied volatility smoothing procedure with the same constraint that smoothing cannot introduce arbitrage opportunities.

A.3 Strike prices

The zero-coupon cap and floor data for the five and ten year maturities that we are interested in has strike prices from 1% to 6% (caps) and from -2% to 3% (floors), in 0.5% increments. At times, individual data points may be missing or the range may be slightly smaller. Using our smoothing algorithm we can calculate implied prices for the missing data points and we can also extrapolate to strike price above and below the maximum and minimum strike price levels.

Starting August 10, 2021, data availability for the US drops and we only have 1% increments. For the EZ the lowest cap strike price is 1.5%.

A.4 Constructing distributions

Data quality is not constant over time. In order to construct accurate distributions we require high-quality data, that is, a combination of many observed option prices and those option prices passing the pre-screening outlined above. Each month, we choose one (or sometimes more) trading days that have the highest quality data, as close as possible to the start of the month. We ensure that spacing between observations is stable, so that we do not end up, for example, with a day at the end of February followed by a day at the beginning of March.

For the year-on-year data, for the EZ it is common that only the five, seven and ten-year maturities are available. This means that we can observe the price of the portfolio of two year-on-year caplets (or floorlets) for the one-year periods starting in five and six years and one portfolio for the following three one-year periods. For the US, we have data for the different maturities but only until June of 2018, after which available maturities also decline. We linearly interpolate the implied volatility for the missing years. Based on data for which the various maturities are available, we know that the year-on-year forward distributions from years five to nine are quite stable, supporting our interpolation technique.

A.4.1 Periods of sparse data, especially on US YOY

When constructing the distributions we use the put-call-parity-implied real interest rate for calculation of the option implied volatility. If there are sparse data, sometimes there are no overlapping observations. This happens only in the case of the YOY data for the US starting in June 2021. Before this time and for all other distributions (5Y and 10Y zero coupon), we have the necessary data. For these cases we use the Bloomberg swap rate for the relevant period. Comparing the nominal rate to the swap rate, we recover the real rate for the period.

Another data issue starting June 2021 is that there are not sufficient data to construct the 1Y distribution, which is needed for construction of the YOY distributions. In those cases the one-year implied volatility function is linearly extrapolated from the two- and three-year implied volatility functions. Given that we are using data for the six to 10-year

horizons, this adjustment has little effect.

B Model of inflation risk

This section of the appendix describes the estimation of the distribution of joint output-inflation disasters. The data on inflation comes from Jordà, Schularick and Taylor (2016), which is then merged with the output data in Barro (2006). The list of 18 covered countries is: Australia, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom, and the United States.

B.1 Identifying disasters: baseline

To classify disasters in a time-series for inflation for an individual country, we proceed in two steps. First, we identify cycles, periods between local maxima and minima. Second, we compare these to a target, or normal, inflation rate at the time. If inflation is sufficiently away from the target we call this a disaster.

On cycle identification, for the results in the main body of the paper, the local minima (maxima) of the inflation series are identified in a rolling, centered five-year window: if the midpoint is lower (higher) than all other values in the window, that point is classified as a trough (peak).

$$\text{Date } t \text{ is: } \begin{cases} \text{a peak} & , \text{ if } \pi_t > \pi_{\tilde{t}} \forall \tilde{t} \in \{t-2, t-1, t+1, t+2\} \\ \text{a trough} & , \text{ if } \pi_t < \pi_{\tilde{t}} \forall \tilde{t} \in \{t-2, t-1, t+1, t+2\} \\ \text{neither} & , \text{ otherwise} \end{cases}$$

All observations after some preceding trough/peak up to (and including) the next local extremum are classified as one cycle (method C2, see below). These cycles, often spanning several years, are the unit for evaluating whether there is a disaster. The inflation of an entire cycle $C = \{t_C, t_C + 1, \dots, t_C + T_C\}$ is the aggregation of yearly inflation within the cycle; we use the cumulative growth rate $\pi_C = \left(\prod_{t \in C} (1 + \pi_t) \right) - 1$ as aggregator.

A cycle is classified to be in a disaster state π_d if this inflation value deviates from some target by some threshold. For the baseline results, the target is given by applying a 20 year-Butterworth square-wave highpass filter on the inflation series (sub-method T3

below).

B.2 Alternatives to identifying disasters

We explored alternatives to both identifying cycles and to setting the target. Starting with the target, beyond the baseline (method T3), we also use the mean of inflation censored at the $[0.25, 0.75]$ -quantiles for each country, with the exception of the US, where we use 2%. Here the inflation target is a country-specific constant. This is method T1. Another alternative was, for each country, to compute the mean of censored inflation as above, but using the past 20 years in a rolling window, imputing for the first 19 observations the values from method T1. Here the inflation target is a country-specific constant for the first 19 years and time-moving afterwards, and we call this method T2. The threshold for deviation is chosen as the inflation target, which in the case of (T2) and (T3) is itself moving with time.

Relative to the baseline, beyond the baseline (method C2), we considered partitioning the observed time period using peaks/troughs. For each country, annual inflation and inflation target (using submethods T1-T3) are smoothed with a five-year leading window. Moving with the direction of time, if in some year inflation deviates from target, that and the next four years are classified as inflation disaster; the evaluation then continues with the year following this cycle. This procedure yields disaster cycles with a fixed length of five years, and we call it method C1.

B.3 Results under alternatives, for pooled sample

A cycle is classified as a joint inflation-and-output disaster if it has been classified as an inflation disaster, and additionally contains at least one year that has been classified as an output disaster in Barro (2006).

Overall, with two methods to partition the time period into cycles, and three methods to define an inflation target, this yields six alternative ways in total. Table 1 presents the unconditional probability of an inflation disaster, and the probability of a joint inflation-and-output disaster conditional on an inflation disaster \tilde{p} , for method $\{C1, C2\} \times \{T1, T2, T3\}$. It also shows the conditional probabilities for when the occurrence of a joint disaster is evaluated on a year-by-year basis rather than per cycle, which produces lower probabilities. Table 1 also reports estimated parameters of a Pareto fit on the (transformed) changes in output $z = 1/(1 + g)$ during joint disasters.

Table 1: Unconditional and conditional probabilities, Pareto fits

Method	C1: fixed disaster length			C2: peak/trough cycles		
	C1.T1	C1.T2	C1.T3	C2.T1	C2.T2	C2.T3
unconditional probability of an inflation disaster	21.3%	20.7%	13.2%	20.6%	21.3%	13.4%
probability of output disaster conditional on inflation disaster \tilde{p}	16.7%	18%	21.3%	16.9%	18.6%	20%
estimated z_0	1.04	1.03	1.04	1.03	1.03	1.03
estimated α	5.73	5.7	6.77	6.11	6.67	6.38

B.4 Results under alternatives, separating high and low

Table 2 presents conditional probabilities where a distinction was made between high and low inflation disasters.

C Horizon factor model

The Markov model that we use to model inflation dynamics has six parameters: symmetric movements in the middle of the distribution, p_{nm} , entering the high or low inflation disaster, p_{dh} and p_{dl} , exiting the high or low inflation disaster, p_{nH} and p_{nL} and a probability capturing mean reversion, p_{mr} . The transition matrix is reported in the main text. We chose this model because it fits well, with parameters that have clear interpretations, and it is sufficiently rich to capture the dynamics well, but not so complicated that it becomes difficult to interpret movements in the parameter estimates.

In our baseline model, the first three parameters are time-varying. This captures time varying volatility and time-varying probabilities of entering a disaster, which is the variation that this paper is interested in estimating. The other three parameters are not time-varying. The probabilities of leaving a disaster are close to constant when estimated in an unconstrained setting and the mean reversion parameter is, if left to vary freely, quite unstable due to the difficulty of identifying it relative to the local movement probability, both of which affect medium-term volatility.

The main model is estimated at the quarterly frequency. Instead of six parameters each period, it includes three fixed parameters and three time-varying parameters for

Table 2: With distinction between deflation and inflation: unconditional and conditional probabilities, and Pareto fits

Method	C1: fixed disaster length			C2: peak/trough cycles		
	C1.T1	C1.T2	C1.T3	C2.T1	C2.T2	C2.T3
Low-inflation disasters only						
probability of output disaster conditional on low-inflation disaster	12.1%	11.4%	14.2%	8.5%	8.5%	8.5%
estimated z_0	1.04	1.03	1.04	1.06	1.06	1.06
estimated α	10.84	8.62	11.55	15.18	15.18	15.18
High-inflation disasters only						
probability of output disaster conditional on high-inflation disaster	19.6%	22.7%	29%	22.7%	25.4%	35.6%
estimated z_0	1.07	1.03	1.05	1.03	1.03	1.03
estimated α	5.4	5.11	5.73	5.4	6.09	5.45

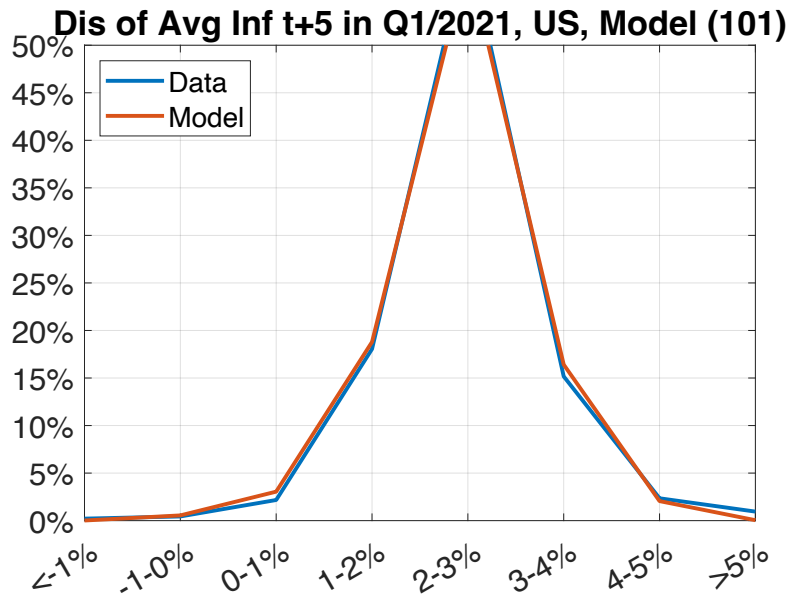
the quarterly sample. Estimating the model using monthly frequency data proved to be computationally too costly relative to the very small potential benefit of higher-frequency estimates of the time-varying parameters based on the full model. To obtain monthly estimates, we re-estimate the model separately at each month, maximizing fit only over the three time-varying parameters, while keeping fixed the three constant parameters estimated with the quarterly data. We verify that in the months in the middle of the quarter, the quarterly and monthly estimates are very close to each other.

C.1 Model fit

The model is fit by GMM. We fit three sets of moments (i) the five-year zero-coupon distribution, (ii) the ten-year zero-coupon distribution, and (iii) the average of the t+6 to t+10 year-on-year distributions. Each set of moments has eight moments associated with it. Each of the sets has an equal weight when minimizing the squared deviations of the model from the actual probability.

As an example, Figures 1, 2 and 3 present data and model distributions for the first quarter of 2021. In order to calculate the model average inflation for five and ten years we

Figure 1: 2021 Q1, 5-year cumulative distribution



need to choose a value for average inflation in the high and low inflation disaster states (below -1% and above 5%); we set these equal to -2% and 6%.

We next compare model fit of our main model (Model 101) to the fit of the alternative model (Model 1) for which all parameters vary freely over time. Figures 4 and 5 plot model fit over time as well as average model fit. What the figures report is the root mean squared error of the model and model R^2 . Figure 5 reports results from fitting a quarterly model and then, in a second step, fitting the time-varying parameters for each of the missing months, but holding fixed the time invariant parameters.

Though there is some heterogeneity over time, with a spike in the early days of the pandemic, overall fit is quite good. Importantly, though overall fit declines when moving from the flexible time-varying model (Model 1) to the more restricted model (Model 101), for R^2 this is driven primarily by poor fit early in the sample period. It is also useful to note that fit, as measured by the RMSE, move together for both models. It is therefore not the model but rather time variation in the data that leads to time variation in fit.

For the EZ the pattern is similar to the US. Figures 6 and 7 show the fit of models 1 and 101 respectively. The pattern is similar to the US data. Overall fit is comparable for both models, though again a little better for the flexible model, as expected. Time variation

Figure 2: 2021 Q1, 10-year cumulative distribution

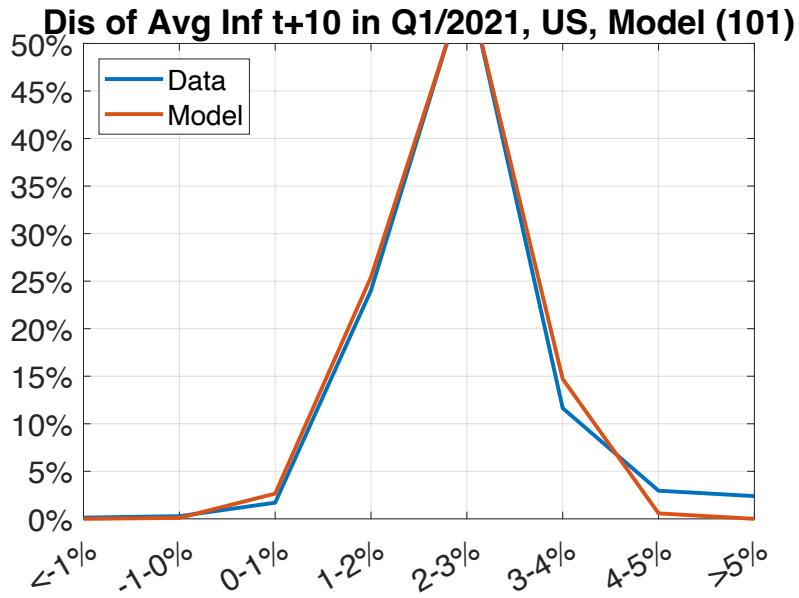


Figure 3: 2021 Q1, one-year forward distribution

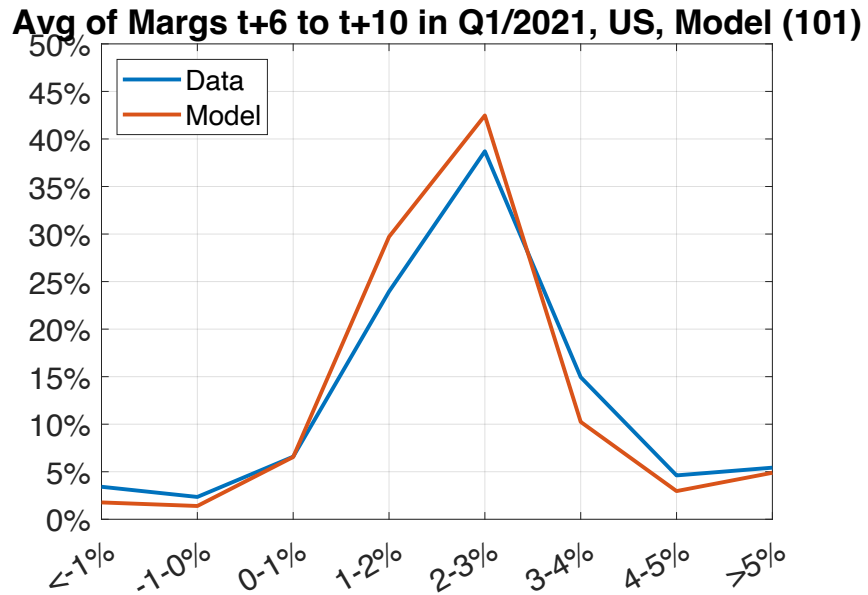


Figure 4: Model fit for US monthly and quarterly models

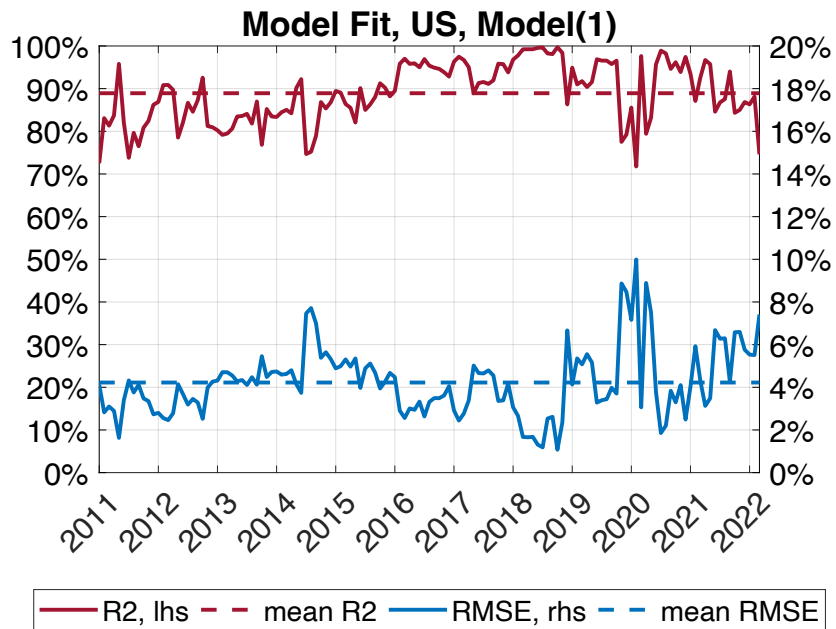


Figure 5: Model fit for US monthly and quarterly models

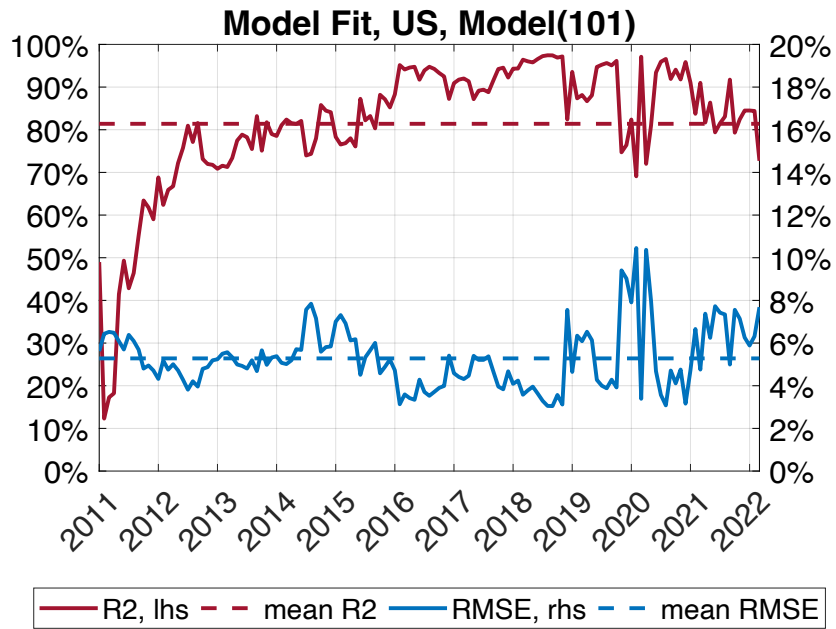
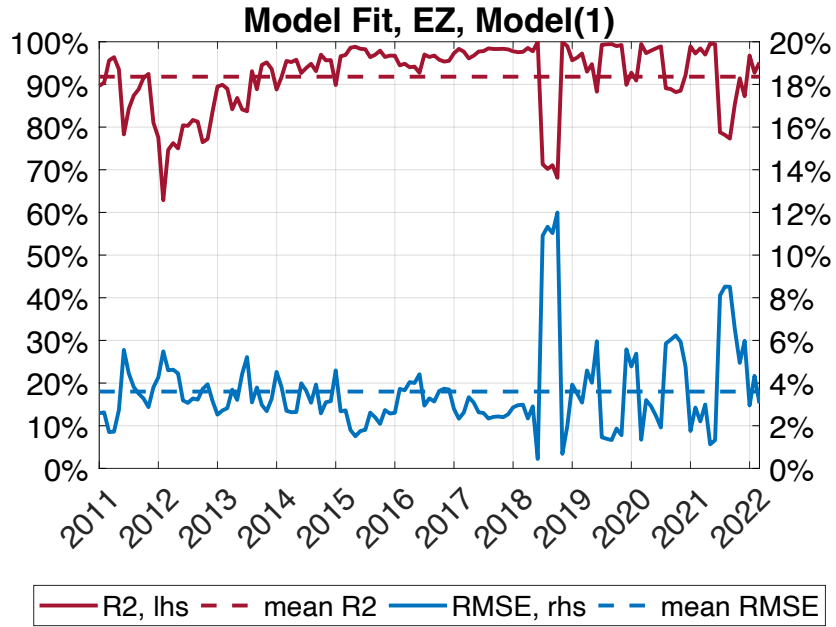


Figure 6: Model fit for EZ monthly and quarterly models



in model fit is also similar, with the exception of the early days of the pandemic, during which model 101 underperforms by a little more.

We also explore another model in which the parameters that are held fixed in model 1 are allowed to vary at the annual frequency, rather than at the monthly frequency, which is what we assume in the fully flexible model. This approach, which we refer to as model 101A, results in a substantial increase in parameters relative to model 101, our main model, and it also improves fit a bit, but it has the same feature as the monthly model, which is an inability to clearly identify long-term trends in volatility through the probability of local changes. This is because this model allows for slow-moving changes in mean reversion, which also affects long-run volatility, itself slow-moving. Figure 8 shows model fit compared to the fully flexible monthly model.

C.2 Model parameters

For completeness, we show the full set of parameters for model 1, both for the US and the EZ. The time-varying parameters are also presented in the main paper. Figures 9 and 10 show the estimated model parameters.

Figure 7: Model fit for EZ monthly and quarterly models

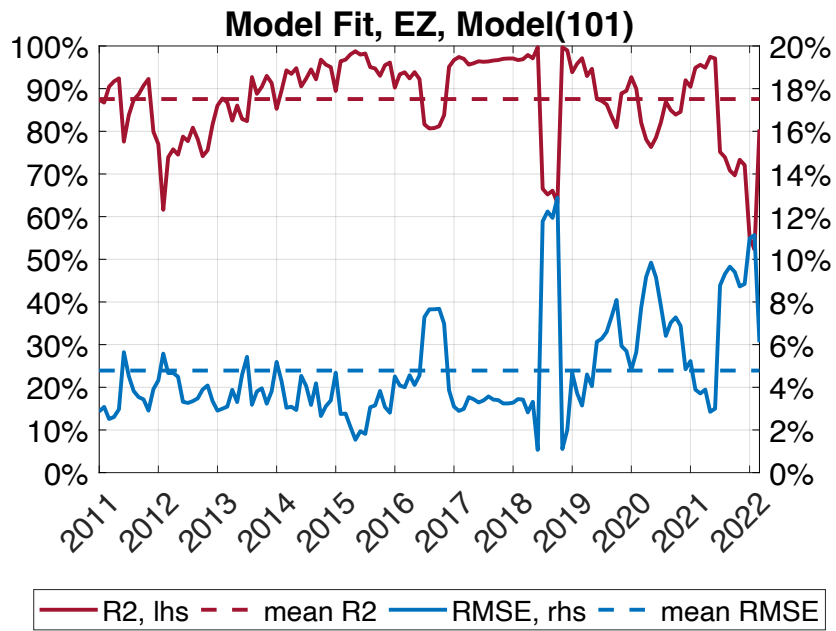


Figure 8: Model fit compared to model with slow-moving mean reversion probability

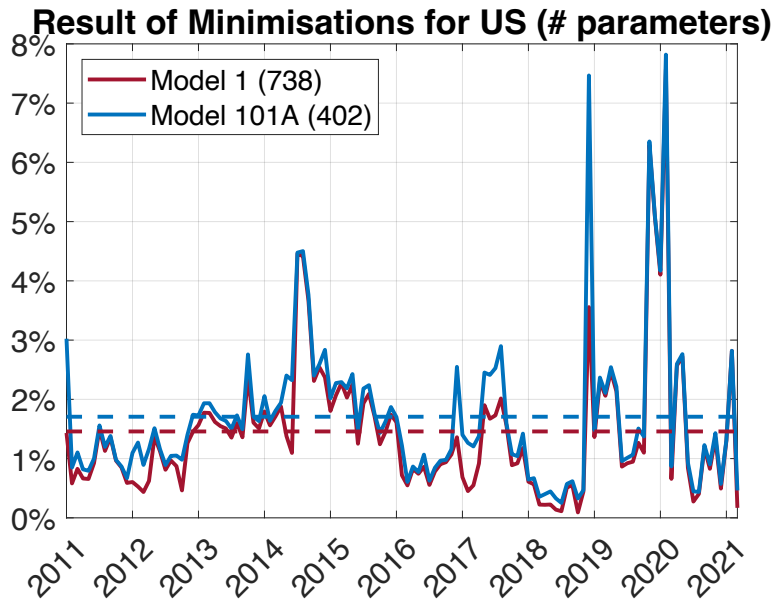


Figure 9: Model parameters: US

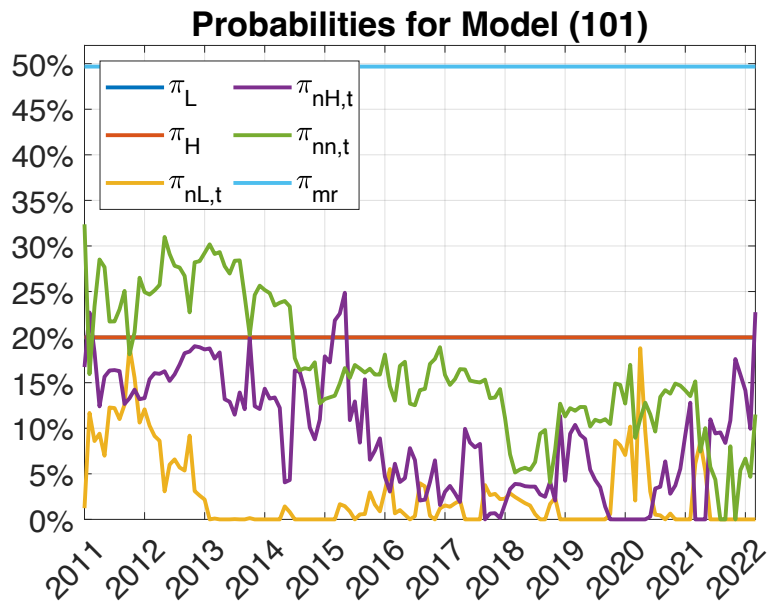
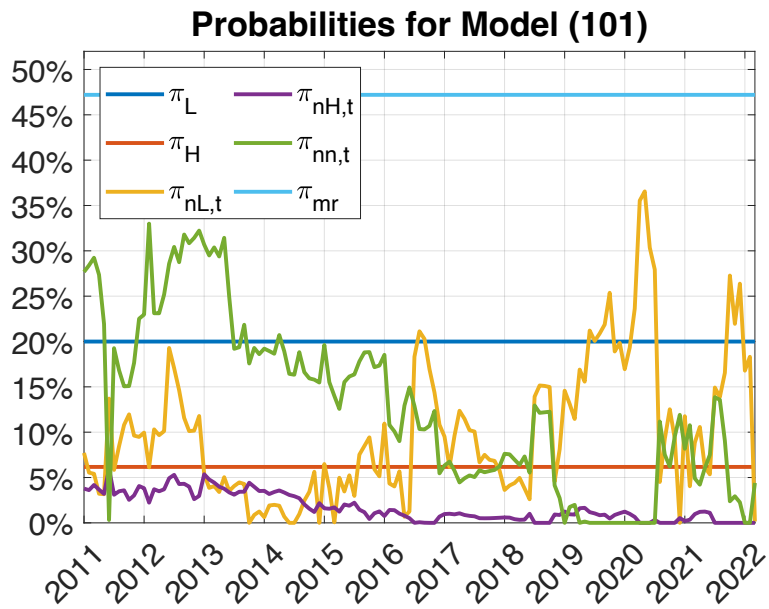


Figure 10: Model parameters: EZ



C.3 Other models

We considered several other candidate models. These models either had too few parameters to have adequate model fit or they had more parameters than were necessary. For completeness, we briefly discuss some of them here.

First, we considered a model with only three parameters – the probability of a local change in inflation, one probability of entering either disaster, and one for leaving it. The model fit was poor. We tried varying the jump size (into and out of disaster) and the number of bins used.

The next model had four more parameters: the probabilities of jumping to either disaster and leaving disaster, one probability of local movements in inflation and a probability capturing mean reversion. As is apparent from the estimated parameters of entering either disaster in our main model or in the flexible six-parameter model, the assumption of the disaster probability being the same for both disasters is too restrictive. It also does not allow us to separately identify disaster probabilities, which is one focus of this work.

In another model, the probability of entering a disaster was allowed to depend on the distance from the disaster state. This added unnecessary flexibility that made little difference in practice.

In another model we allowed the probability of jumping to disaster to depend on the distance from disaster, either by estimating separate probabilities depending on the distance or by assuming that the probability is a function of the number of bins between the current state and disaster. Again this proved to be more complicated than necessary.

Finally, as a separate robustness check we have estimated a model in which parameters vary at different frequencies. The parameters assumed to be constant in our main model vary at annual frequency and the time-varying parameters vary at monthly frequency. The model has the advantage of being able to include monthly data. However, it has the same disadvantage as the fully time-varying model in that low-frequency movements in volatility and disaster probabilities are harder to detect.

To conclude, across models, the different parameter movements were broadly comparable, though, as discussed, the long run decline in volatility cannot be observed as easily since more than one time-varying parameter affects volatility.

References

- Barro, Robert J.** 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–866.
- Birru, Justin, and Stephen Figlewski.** 2012. "Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008." *Journal of Financial Markets*, 15(2): 151–180.
- Black, Fischer, and Myron S Scholes.** 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81(3): 637–54.
- Hagan, Patrick S., Deep Kumar, Andrew S. Lesniewski, and Diana W. Woodward.** 2002. "Managing Smile Risk." *Wilmott*.
- Hilscher, Jens, Alon Raviv, and Ricardo Reis.** 2022. "Inflating Away the Public Debt? An Empirical Assessment." *Review of Financial Studies*, 35(3): 1553–1595.
- Jordà, Òscar, Moritz Schularick, and Alan M. Taylor.** 2016. "Macrofinancial History and the New Business Cycle Facts." In *NBER Macroeconomics Annual 2016*. Vol. 31, , ed. Martin Eichenbaum and Jonathan A. Parker. University of Chicago Press.
- Rubinstein, Mark.** 1991. "Pay Now, Choose Later." *Risk*, 4: 13.
- Shimko, David.** 1993. "Bounds of Probability." *Risk*, 6: 33–37.