How likely is an inflation disaster?*

Jens Hilscher — UC Davis
Alon Raviv — Bar-Ilan University
Ricardo Reis — LSE

September 2024

Abstract

Long-dated inflation swap contracts provide widely-used estimates of expected inflation. We develop methods to estimate complementary tail probabilities for persistently very high or low inflation using inflation options prices. We show that three new adjustments to conventional methods are crucial: inflation, horizon, and risk. An application of these methods finds: (i) US deflation risk in 2011-14 has been overstated, (ii) ECB unconventional policies lowered the deflation disaster probability, (iii) inflation expectations deanchored in 2021-22, (iv) and reanchored as policy tightened, (v) but the 2021-24 disaster left scars, (vi) US expectations are less sensitive to inflation realizations than 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, Seyed Hosseini, Anna Pilipentseva, Rui Sousa, and Borui Zhu for excellent research assistance, and to seminar participants at Markus Academy, Johns Hopkins University, 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 expected inflation 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 at the central bank's target (e.g., Gürkaynak, Levin and Swanson, 2010). Policymakers find it useful because it strips out current temporary fluctuations and it averages over a long period of time, thereby providing focus on what monetary policy can achieve. In speeches, they often point to the approximate constancy of the 5y5y over the past 25 years (see figure 1) to claim success at anchoring expectations, and even small changes in this measure can trigger large shifts in policy: a decline in the EZ in 2011-14 justified the start of quantitative easing.[1]

However, the 5y5y rate is a point estimate of an average. The distribution of its values could be extremely tight or dispersed. Making decisions under uncertainty typically requires knowing the whole distribution of future inflation rates, not just their expected value. Especially important for decision making and risk management are probabilities of extreme inflation realizations, which we will refer to as inflation disasters. It is these tail events that are associated with large costs of inflation, both in models of monetary policy and in opinion polls, as happened in 2021-24, or during the German hyperinflation of the 1920s, or the stagflation of the 1970s.

This paper develops the methods to provide counterparts to figure 1 in the form of tail probabilities of inflation disasters using traded option prices and minimal assumptions about preferences for pricing risk or inflation dynamics. To be specific, the objects of

[1] Reporting 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 behavior and policy action."

Figure 1: Expected 5y5y inflation

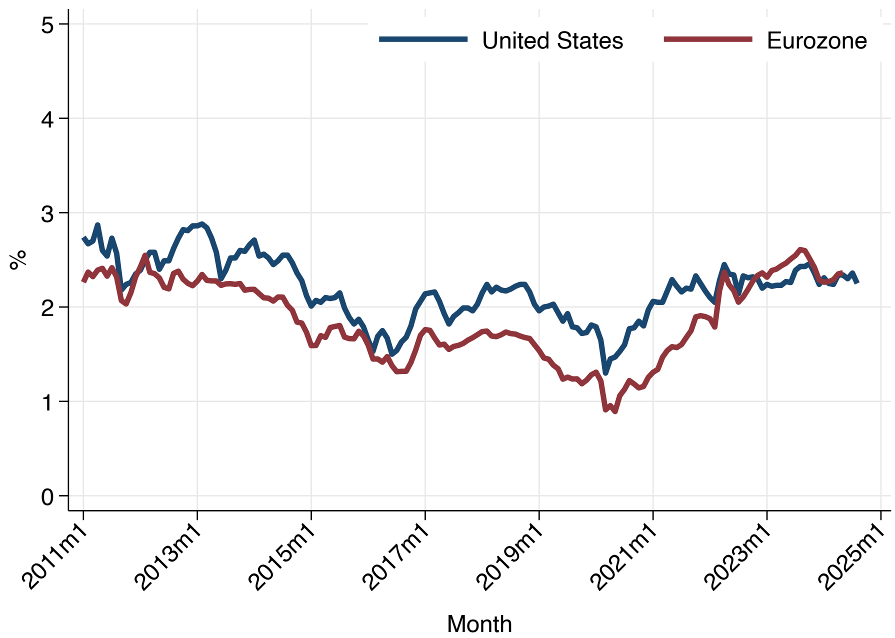

Source: FRED for the US and Bloomberg for the EZ.

interest are the following two probabilities:

$$Prob[\pi_{T,T+H}/H > \bar{\pi} + d] \quad \text{and} \quad Prob[\pi_{T,T+H}/H < \bar{\pi} - d].$$

Starting from the present, a distant future 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 answer the question: What is the current 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% (2pp above the 2% target on average) between 5 and 10 years from now?

In our empirical implementation, we provide 5y5y estimates of these probabilities for the United States (US) and the Eurozone (EZ) starting in October of 2009 and January of 2011, respectively, until April 2024. For disasters, we consider both high inflation and deflation, $d = 0.02$, or severe high inflation and deflation, $d = 0.03$.[2] We use these estimates to measure the success of monetary policy at anchoring market inflation expectations, and to judge the impact of different frameworks and policies on the stability of these anchors.

This paper makes a methodological contribution, which in turn leads to new empirical estimates and a revision of the recent history of inflation expectations.

Methodological contribution: We provide steps to translate the prices of traded inflation derivatives into risk-neutral and physical-measure probabilities of inflation. We show that using standard methods results in inaccurate estimates that can grossly over- or under-state the desired probabilities. These conventional readings must be adjusted in three ways.

First, their units have to be adjusted to match Arrow-Debreu probabilities. When inflation is more likely, this raises the nominal payoff of a call option but it also lowers its real payoff. The conventionally-used nominal state prices are therefore too low for high inflation states since a nominal payoff of $1 in a future high-inflation state is worth less. For low inflation, the opposite is true.[3]

Second, traded options pay out based on realizations of inflation at $\pi_{0,T}$ and $\pi_{0,T+H}$, but not for the desired forward horizon $\pi_{T,T+H}$.[4] If inflation expectations become unanchored gradually, this sluggishness implies that both the 5-year (5y) and the 10-year (10y)

[2]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. Higher choices of d are hard to implement due to little trade of options further in the tails.

[3]This 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 conditional distribution is a trivial point mass, making the adjustment simple.

[4]There 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 focus on and which aligns with policymakers' interest in the 5y5y horizon.

probabilities can understate the probability of a 5y5y inflation disaster. Given our focus on expectations, we innovate by estimating the perceived sluggishness revealed by the pricing of options at different horizons, as opposed to using past sluggishness in inflation realizations as has been done before.[5]

Third, option prices imply probabilities adjusted for risk (or risk-neutral probabilities), but since marginal utility is likely high during disasters, their prices will over-state the actual, or physical-measure, tail probabilities. Building on recent work on rare output disasters, we propose an adjustment that does not require specifying the full dynamics of the stochastic discount factor that prices inflation risk, but instead uses prices of out-of-the-money options.

Empirically, we find that the three adjustments can be large. For instance, during the 2021-23 period of rising US inflation, the median adjustment factors for inflation, horizon and risk were 1.24, 0.38, and 0.66, respectively. As a result, while a simple reading of the 10y option prices would suggest a median 14.0% probability of a 5y5y inflation disaster, the 5y5y actual probability was 4.2%. For some questions, researchers may want to make only one or two of these adjustments: making only the inflation and risk (but not the horizon) adjustments, we provide estimates of inflation disasters between now and five or ten years in the future (so $T = 0$ and $H = 5$ or $H = 10$). Making only the inflation and horizon (but not risk) adjustments, we provide estimates of risk-neutral probabilities.[6]

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. Likewise, we provide tail-focused estimates of the inflation risk premium. We find that periods of high inflation carry a large risk adjustment. In contrast, periods of deflation are associated with a smaller drop

[5]We will shortly discuss the related literature and how our approach differs from previous work.

[6]The website https://r2rsquaredlse.github.io/web-inflationdisasters/ provides time series of the inflation disaster probabilities.

in output, thus resulting in a correspondingly lower risk adjustment.

Contribution to the history of expectations: Our second contribution is an application of the estimates to reassess whether or not inflation expectations have been anchored and how monetary policy has affected anchoring. Our sample includes periods of elevated probabilities for both future deflation and high inflation. It allows us to illustrate our methods while also shining new light on the interpretation of monetary policy and its effects. We reach six conclusions.

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 when using the conventional measures that did not include our three adjustments. In particular, we find that the risk of short-term deflation was at times elevated, but not the risk of a deflation trap at the 5y5y horizon.

Second, we find that the risk of a EZ deflation trap persisted throughout the sample and is significantly higher than in the US. The unconventional monetary policies since 2014 and the ECB's mission review of 2022 succeeded at lowering the probability of deflation in the near future, but not completely at lowering the perceived risk of a deflation trap over the long run.

Third, we find a large increase in the probability of a high-inflation disaster in the US between the third quarter of 2021 and the second quarter of 2022. The same is true in the EZ, but starting later and being more concentrated in the first half of 2022. While the mean of the distribution of expected inflation moved little, leading policymakers at the time to conclude that expectations were anchored and the observed increase in inflation would be temporary (Powell, 2021, Lagarde, 2021), the tails showed a sharp deanchoring. The probability of an inflation disaster peaked in the US at 10% in May 2022 and in the

EZ at 6% in June of 2022.

Fourth, we find that this deanchoring had a U-turn in 2022 that coincided almost to the month with the U-turn in monetary policy and the hiking of policy rates. By the end of 2022, the probability of a disaster had stabilized below 4% in the US, while it has averaged 6.5% in the EZ in the last twelve months of the sample. Since the deanchoring coincided with unusually loose monetary policy, and the reanchoring with tightening policy, this provides support for a tight link between monetary policy and the inflation anchor.

Fifth, the probabilities of an inflation disaster by the end of the sample are two to three times higher than they were between 2011 and 2019. The inflation disaster of 2021-24 has left scars in market perceptions. This supports theories where credibility depends on realized outcomes.

Sixth and finally, we calculate how sensitive are the probabilities of disaster to temporarily high inflation, either in the present or the near future. We find that in the US, expected inflation is well anchored in the sense of being insensitive to inflation realizations, but this was less so until recently in the EZ.

Outline: The paper is organized as follows. Section 1.1 discusses our contributions relative to the existing literature. Section 2 lays out our approach, using both a simple two-period setup and a general model, and defines the three adjustments. Section 3 presents US and EZ data, discusses summary statistics of the disaster probabilities, and quantifies the three adjustments. Section 4 applies our estimates to reassess the extent of anchoring of inflation expectations between 2011 and 2024 in the US and the EZ. Section 5 presents more detailed information about the implementation of our method. Section 6 concludes.

1.1 Connection to the literature

This paper is related to three strands of the literature.

Other uses of inflation options data: Closest to our paper, a small literature has constructed inflation probability measures from market prices of inflation options. Kitsul and Wright (2013) (KW), the first paper in this literature, uses US data over a three-and-a-half year period starting in October of 2009 and estimates probabilities of deflation at 1, 3, 5, and 10 year horizons (starting from today). Methodologically, our paper begins where KW ends: we apply our three adjustments to the state prices that they produce. In the applications, we use a much larger sample that includes the recent inflation disaster and the other largest currency in the world; we revisit the probability of deflation during the common sample (reaching a different conclusion to theirs); and we use the estimates to study the connection between anchoring of expectations and monetary policy.

More concretely, there are five important differences relative to our paper. First, KW uses conventional methods to extract state prices from options prices. We adjust for the effect of inflation on the real payoffs, so that the state prices can match Arrow-Debreu probabilities (our first adjustment).

Second, a main focus of our paper is the construction of forward probabilities (our second adjustment). KW only present probabilities from today onwards. The overwhelming focus of policymakers on the 5y5y measure testifies to the relevance of this forward approach, as do common debates on whether inflation is transitory or persistent.[7] In our applications, we find that this distinction is very important for the key macroeconomic debates: in 2011-20, the probability of a short-lived deflation was very different from the probability of a deflation trap, and when an inflation disaster actually happens, as in 2021-24, assessing whether long-run expectations are anchored must exclude the near future.

[7] In the NBER working paper version, Kitsul and Wright (2012) present distributions of forward inflation for four dates in 2011 and 12, but these are based on the very restrictive assumptions of risk neutrality and inflation innovations being independent over time. These assumptions are not consistent with market participants' expectations nor the data on rare disasters.

Third, KW estimates and uses a statistical model for the stochastic discount factor with respect to inflation risk that depends on realizations of actual US inflation. Instead, we only use historical data from eighteen countries on inflation and GDP growth to measure marginal utilities explicitly during inflation disasters, in order to adjust option-implied probabilities for risk (our third adjustment). Otherwise, we let the options prices at different horizons tell us what the market believes are the dynamics of inflation. KW uses its stochastic discount factor to study its relation with inflation, but do not use it to construct physical probabilities of disasters from the options data, which is our focus.

Fourth, our sample period is approximately four times as long. KW uses data from an options market that had just become active, while our sample includes many more years when this market was mature. The longer sample also allows us to apply the estimates to important debates surrounding inflation and monetary policy in the last ten years, and especially to the recent period of dramatically elevated inflation where the question of inflation expectations becoming unanchored was central. Almost none of our applications could have been studied using the KW sample period.

Fifth, and finally, we construct EZ inflation disaster probabilities and discuss how EZ policy affected them. The challenge the ECB has faced in fighting deflation, as well as the comparison between the US and the EZ are an important part of our paper.

Related, Fleckenstein, Longstaff and Lustig (2017) use US data on inflation swaps and options through October 2015 to estimate a stochastic volatility model of inflation dynamics that allows for time-varying risk premia. It also calculates probabilities of inflation being low or high at various horizons. The differences relative to this paper are the same as for KW: the three adjustments, as well as the longer sample period and the inclusion of the EZ allowing us to discuss a rich set of applications. Moreover, when adjusting for risk, we allow for disaster-specific prices of risk. It instead has one inflation risk premium that is not state-specific.

Mertens and Williams (2021) fits a New Keynesian model to US data on interest rates and option prices. It calculates forward probabilities assuming that inflation follows a Gaussian random walk, an assumption which we show to be inconsistent with the market's expectations. Also, it does not make the inflation adjustment to nominal payoffs, nor adjust for risk. As we show, each of these adjustments matter significantly.[8]

Other measures of inflation risk premia: Our model of inflation risk draws on the literature on equity disasters (Barro, 2006, Gabaix, 2012, Barro and Liao, 2021), while 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.

A large literature extracts information about inflation from market prices by fitting term structure models to real and nominal yield data, sometimes including inflation swap prices, e.g., Christensen, Lopez and Rudebusch (2010, 2015), Haubrich, Pennacchi and Ritchken (2012), Hördahl and Tristani (2012). These papers focus on estimating the inflation risk premium together with the implied expected average inflation rate. Instead, we focus on the tails of the inflation distributions, and as such on estimating a disaster-specific risk adjustment (our third adjustment). Our emphasis on inflation-output disasters sets us apart.

Tail macroeconomic outcomes: A small literature focuses on tail outcomes for inflation disasters, specifically inflation at risk (Kilian and Manganelli, 2007, Banerjee et al., 2020, Andrade, Ghysels and Idier, 2012, Lopez-Salido and Loria, 2020). Instead of outcomes

[8]Gimeno and Ibanez (2018) is closer to us in goal but imposes more restrictive assumptions. Hilscher, Raviv and Reis (2022) makes the inflation adjustment, but does not explain it or quantify its effect, nor does it adjust for horizon and risk. It also discusses a very different set of issues.

of realized inflation, we measure market perceptions of this risk as measured by option prices. 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.

A few papers look at expectations of disasters in surveys (Reis, 2022, Ryngaert, 2022). We instead 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).

2 Constructing probabilities of inflation disasters: theory

We explain the intuition behind our method when there are three periods and three states of the world before deriving its applicability in a general setup.

2.1 The intuition of the method in a simple setup

Consider a simple event-tree world illustrated in the left panel of figure 2. At present, in period 0, inflation is at its normal target level $\bar{\pi}$, at which it stays with a 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 p_m and p_d.[9]

The probability of an inflation disaster that we are interested in is p_d. We have data

[9] We focus on a high-inflation disaster for expositional simplicity, but the arguments would be the same for a low-inflation disaster.

Figure 2: Illustrative inflation paths

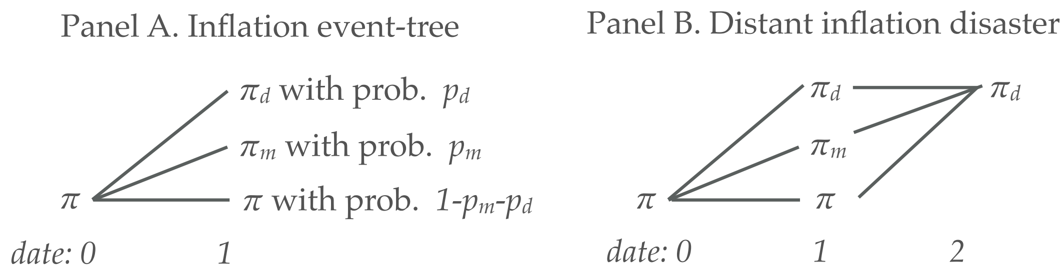

on options that pay one nominal unit if π_d is realized in period 1, and zero otherwise. Assuming no arbitrage, the price of that option is equal to $a_d(1) = p_d \exp(-\pi_d)m_d$: when the event with probability p_d is realized it pays $1, which in real terms requires an inflation adjustment $\exp(-\pi_d)$, that is discounted by the real stochastic discount factor m_d, reflecting the marginal utility of the future payoff.

The conventional approach is to construct state prices by calculating $n_d(1) = a_d(1)\exp(i(1))$, where $i(1)$ is the nominal interest rate, since the nominal rate is the return from buying the three possible options.[10] It then follows that $n_d(1) \geq 0$ and that $\bar{n}(1) + n_m(1) + n_d(1) = 1$. Therefore, the state price $n_d(1)$ can be interpreted as a probability. But what does this measure?

2.1.1 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 a disaster A-D security is: $p_d m_d$. Letting $r(1)$ denote the real interest rate, the associated A-D probability is then $q_d(1) = p_d m_d \exp(r(1))$. This is the real risk-neutral probability.

[10] No arbitrage dictates: $\bar{a}(1) + a_m(1) + a_d(1) = \exp(-i(1))$.

It follows right away that:

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

The approximation comes from the starting assumption that monetary policy is credible, so that break-even inflation—the gap between the nominal and the real interest rates—is equal to the target inflation level. The second equality comes from recalling that the gap between disaster inflation and target inflation is what we earlier called the disaster d. The conventionally-measured probabilities from options $n_d(1)$ must therefore be adjusted by the disaster size d. This is our first adjustment.

Intuitively, when the disaster happens, and the option pays, its $1 is now worth less in real terms. Economic agents therefore pay less for this option than if they were suffering from money illusion. Researchers in turn need to adjust for this effect as well.

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 thus significantly underestimates the risk-neutral probability $q_d(1)$.

2.1.2 Second 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 in that period, would provide an estimate of the risk neutral probability, $q_d(2)$.

However, we do not have the option prices that match this security. Looking at either short-dated or long-dated cumulative inflation options does not capture the desired probability. For instance, a short-dated option that pays if there is a disaster in the first period provides an estimate of p_d, while a long-dated option that pays if there has been a disaster that lasts for two periods gives an estimate of $p_d p_{dd}$. Since inflation moves sluggishly, $p_m p_{md} / p_d$ is large, so that both short and long-dated options understate the desired forward probability.

To calculate the forward probability, we need information on the extent to which inflation is sluggish. In this simple example, given a short-dated and a long-dated option, one more piece of data allows us to calculate $q_d(2)$. Fortunately, there are traded forward-starting options, but for sub-periods of our hypothetical period 2. Namely, in the data, there are forward contracts for annual inflation within our 5-year desired periods. They provide the missing piece of data.

2.1.3 Third adjustment: actual or physical measure (risk-adjusted) probabilities

The next adjustment is the more familiar to financial economists: with an estimate of the stochastic discount factor in the inflation disaster state, m_d, we can go from the risk-neutral to the physical probabilities. Importantly, to answer the question in this paper, one does not need a full model of risk. In our example, 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 the disaster state adjustment will be the largest adjustment of the three states.

Imagine then that the main source of variation in the stochastic discount factor is whether there is a consumption disaster or not. So, $m_d(.)$ is a function of consumption, which can either be normal or in a disaster. Conditional on an inflation disaster, let $\tilde{p}$ be 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 $m_d(.)$ when there is a consumption disaster, to the marginal utility when there is none, call this ratio $\tilde{m}$, is well above 1.

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:

$$q(1) \approx [(\tilde{m} - 1)\, \tilde{p} + 1]\, p_d. \tag{2}$$

Since $\tilde{m} > 1$, the risk-neutral probability will over-state the probability of an inflation disaster.

The rare disasters literature has argued that $\tilde{m}$ can be quite large. However, for inflation, 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, $\tilde{m}$, which measures the marginal utility of both an inflation and a consumption disaster relative to normal times is on average lower than in a 'standard' consumption disaster because, historically, there are several episodes where an inflation disaster came with only a mild (or no) 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 necessary 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 them provides a path forward to identify the two parameters.

2.2 The theoretical result

This section sets out a general framework and derives the key theoretical result on the three adjustment factors to go from option prices to the probability of inflation disasters.

Uncertainty about inflation: Inflation is a random variable and has an associated probability distribution $p(\pi)$.[11]

Inflation securities and inflation risk: The non-negative price in consumption units of an Arrow-Debreu inflation security that pays one unit of the consumption good if inflation is π at the future date is equal to:

$$b(\pi) = p(\pi)m(\pi), \tag{3}$$

where $m(\pi)$ measures the average marginal utility across states of the world where inflation is the same. Since $m(\pi)$ varies *only* with inflation, it has all the information relevant to assess inflation risk.[12]

Risk-neutral Q-probabilities: The real risk-free security pays a constant one unit of consumption. The inverse of its price is e^r, where r is the 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: $e^{-r} = \sum_\pi b(\pi) = \sum_\pi p(\pi)m(\pi)$.[13] Defining

[11] In the appendix we use a setup that allows for different states of the world s that have the same level of inflation but differ in other dimensions, e.g., consumption.

[12] With other sources of risk in the economy, $m(\pi)$ would average across them. The appendix generalizes this. This pattern is present in the data—between 2000 and 2020, the US economy went through booms and busts, but inflation was approximately unchanged.

[13] Also, as is standard, e^{-r} is the expected SDF or marginal utility of consumption growth.

$q(\pi) = b(\pi)e^r$, it is non-negative and adds up to 1 across inflation rates. This is the risk-neutral probability of this inflation rate.

N**-probabilities:** In order to match what is traded in the market, consider a security that pays one nominal unit at the future state-date. Its price is $a(\pi) = b(\pi)e^{-\pi}$. Importantly, 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 Arrow-Debreu inflation security. The nominal interest rate i is likewise defined as the inverse of the price of a security that pays one nominal unit for sure next period $e^{-i} = \sum_\pi b(\pi)e^{-\pi}$. Combining these two, one can define an N-probability (for "nominal risk neutral probability") as $n(\pi) = b(\pi)e^{i-\pi}$, which is itself non-negative and adds up to 1.

Linking $Q-$ **and** $N-$ **probabilities:** Let $\pi^e = i - r$, be expected inflation. It immediately follows that risk-neutral and nominal probabilities are related according to:

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

The Q-probability of average expected inflation coincides with the N-probability. But as we go towards the tails, away from that average they are increasingly apart.

Time and horizons: Starting from the present, the joint risk-neutral probability density of inflation over the following 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, $\pi_{T,T+H}H = \pi_{T,T+1} + \pi_{T+1,T+2} + ... + \pi_{T+H-1,T+H}$, and there is a joint distribution of $q(\pi_{T,T+1}, \pi_{T+1,T+2}, ..., \pi_{T+H-1,T+H})$. Combining all of these probabilities:

$$q(\pi_{T,T+H}) = q(\pi_{0,T+H}) \sum_{\pi_{0,T}} \left[\frac{q(\pi_{0,T})}{q(\pi_{0,T+H})} \times \sum_{\pi_{T,T+1},...,\pi_{T+H-1,T+H}} q\left(\pi_{T,T+1}, ..., \pi_{T+H-1,T+H} | \pi_{0,T}, \sum_{j=1}^{H} \pi_{T+j-1,T+j} = \pi_{T,T+H} \right) \right] \tag{5}$$

The expression in the bottom line takes into account the persistence of inflation across successive periods within the interval of time $(T, T+H)$. On the top line is the adjustment for the sluggishness of inflation over the long horizons.

Final result: Combining all the steps, we get the result as a formula to obtain the desired disaster probabilities:

Proposition 1. *The probabilities of high and low inflation disasters are, respectively:* $\sum_{\pi_{T,T+H} > H(\bar{\pi}+d)} p(\pi_{T,T+H})$ *and* $\sum_{\pi_{T,T+H} < H(\bar{\pi}-d)} p(\pi_{T,T+H})$ *where:*

$$\begin{aligned} p(\pi_{T,T+H}) = & \underbrace{n(\pi_{T,T+H})}_{\textit{Options Data}} \\ & \times \underbrace{\left(e^{(\pi_{T,T+H} - \pi^{e}_{T,T+H})H} \right)}_{\textit{Real Factor}} \\ & \times \underbrace{\left(e^{-r_{T,T+H}H} m(\pi_{T,T+H}) \right)}_{\textit{Risk Factor}} \\ & \times \underbrace{\sum_{\pi_{0,T}} \left[\left(\sum_{...=\pi_{T,T+H}} q(\pi_{T,T+1}, ..., \pi_{T+H-1,H} | \pi_{0,T}) \right) \frac{q(\pi_{0,T})}{q(\pi_{0,T+H})} \right]}_{\textit{Horizon Factor}} \end{aligned} \tag{6}$$

3 Quantifying the adjustment factors in the data

Using data from option prices, we now apply the three adjustments for inflation, horizon and risk to build disaster probabilities.

3.1 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, are often present in these option markets to hedge some of their positions. Therefore, even though trading volumes will differ, these data are as good as those behind figure 1, which are used frequently.[14]

Price data exist for both call and put options for average inflation between the present and up to 15 years in the future 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 the difference is positive, or zero otherwise. The price of that option today is $a(k)$.

We use US options data from October of 2009 to April of 2024 from Bloomberg; for the EZ, the sample starts in January of 2011. While option prices are available daily, sometimes the data quality is low. To be conservative, we construct data at the monthly frequency. We focus on horizons of 5 and 10 years, which are two of the more commonly traded markets for these securities. We also compare put-call-parity real rates with those implied by the inflation swap contracts to confirm that prices are not only consistent within the options market but also across inflation derivative markets. The appendix describes how we construct the data.[15]

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

[15] These 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

3.2 Summary statistics of probabilities and the three adjustments

Section 5 describes the details behind the implementation of the result in proposition 1. Starting from the prices $a(k)$, we first construct $N-$probabilities, like the ones used by Kitsul and Wright (2013), then adjust them for inflation to get the right $Q-$probabilities, then estimate a Markov model to calculate forward-starting (5y5y) probabilities, and finally estimate a state-specific risk adjustment that is different for high inflation and deflation states. Some of these steps are of independent interest, for instance for our understanding of perceived inflation dynamics. Here, we discuss their quantitative relevance.

Table 1 presents summary statistics of the various steps in finding the probabilities. Panels A and B focus on high-inflation disasters, both their final probabilities and the adjustment factors to get to them, respectively. We show the medians for the full sample, and for a period in the US, between September 2021 and August 2023 when realized inflation reached 9.1%, and the high-inflation disaster probabilities were large. Panels C and D show the US deflation probabilities between January 2011 until December 2012, a time when there were heightened deflation fears and the Fed undertook QE rounds two (in November 2010) and three (September 2012) as well as operation twist (September 2011).

3.2.1 Inflation adjustment

The first four (two) columns in panel A (B) show the inflation adjustment. During the recent period, the median 5-year N-probability was 20.7%, while the risk-neutral Q-probability was higher at 22.8%. The median adjustment factor to move from N to Q is 1.09.[16] As expected, the effect is larger for the 10y horizon, with a median adjustment

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 indication that it is significant for most of the period that we cover.

[16]The adjustment factor depends on the probability density, which is not constant, so neither is the adjustment factor.

Table 1: Three inflation disaster probability adjustments

This table reports medians of various inflation disaster probabilities for different samples. We focus attention on six measures. In columns 1-4 of Panel A we report probabilities of average inflation lying above 4% over the next five (5y) or 10 (10y) years. N denotes nominal risk-neutral probabilities, Q denotes real risk neutral probabilities, i.e. probabilities after adjusting for the effect of inflation (inflation adjustment). Column five reports foward real risk-neutral probabilites of five-year forward probabilites, that is the probability of inflation lying above 4% in five years for five years (horizon adjustment). The final column adjusts that probability for risk (risk adjustment) and reports the physical probability; it is denoted by P. Panel B reports adjusmtent factors: Adjusting for inflation and therefore moving from N to Q probabilities (5y and 10y); adjusting for horizon, i.e. moving from Q_10y to Q_5y5y probbilities; and adjust for risk, moving from Q_5y5y to P_5y5y probabilities. Panels C and D report deflation probabilities and adjustment factors. Data are monthly and there are 160 observations.

Panel A: High inflation disaster (>4%) probabilities, 9/21 - 8/23

	N_5y	Q_5y	N_10y	Q_10y	Q_5y5y	P_5y5y
US, 9/21-8/23	20.7%	22.8%	14.0%	17.2%	6.3%	4.2%
US	6.0%	6.7%	8.9%	11.0%	5.2%	3.5%
EZ	1.4%	1.7%	2.8%	3.6%	4.9%	3.2%

Panel B: High inflation disaster probability adjustment factors

	N to Q, 5y	N to Q, 10y	Q, 10y to 5y5y	Q to P, 5y5y
US, 9/21-8/23	1.09	1.24	0.38	0.66
US	1.12	1.23	0.41	0.66
EZ	1.17	1.33	0.93	0.66

Panel C: Deflation (<0%) probabilities, 1/11 - 12/12

	N_5y	Q_5y	N_10y	Q_10y	Q_5y5y	P_5y5y
US, 1/11-12/12	6.7%	5.6%	6.9%	4.8%	6.4%	6.2%
US	2.7%	2.3%	2.1%	1.5%	2.5%	2.4%
EZ	5.0%	4.6%	5.0%	4.2%	6.6%	6.3%

Panel D: Deflation probability adjustment factors

	N_5y to Q_5y	N_10y to Q_10y	Q_10y to Q_5y5y	Q_5y5y to P_5y5y
US, 1/11-12/12	0.84	0.69	1.41	0.96
US	0.85	0.72	1.31	0.96
EZ	0.90	0.80	2.26	0.96

factor of 1.24, as the median N-probability was 14%, but the Q-probability was 17.2%.

For deflation, in panels C and D, the adjustments work in the opposite direction. The median 5-year N deflation probability over the 24-month period when it was heightened was 6.7%, compared to 5.6% for the Q-probability, while for the 10y horizon, the difference was larger; median adjustment factors were 0.84 and 0.69 for the two horizons.

In the full sample, the adjustment sizes are a little smaller because extreme inflation and therefore adjustments are less likely. Still, the adjustments are higher the longer is the horizon, and they also turn out to be larger for the US than for the EZ. Overall, N-

probabilities overstate the risks of deflation and understate the risks of high inflation.

3.2.2 Horizon adjustment

The next adjustment is for the horizon; we use it to get the 5y5y forward probability. Panels A and C report the median 10y (column 4) and 5y5y (column 5) disaster probabilities, while Panels B and D report median adjustment factors (column 3).

This adjustment is the largest of the three. In the 2021 to 2023 US high-inflation subsample, the median forward $Q-$probability of an inflation disaster is 6.3% compared to the 10y probability of 17.2%. Similarly, in the full sample, the median forward $Q-$probability (5.2%) is smaller than the 10y probability, which is 11%. Reflecting this, the median adjustment factor is 0.41. The reason is that when US inflation is high, it is not expected to persist.

In contrast, in the EZ, the adjustment factors are much higher, reflecting a higher market-perceived persistence of inflation. As a result, if one were to look at the 10y probabilities, one would think the US is much more likely to have a disaster than the EZ. In fact, at the forward horizon, the two are quite close to each other.[17]

In contrast, for deflation, forward probabilities are higher than 10y probabilities, both when deflation risk is elevated and for the full sample. When inflation is low, there is a worry that it may continue to be low in the future. The median EZ forward deflation $Q-$probability is 6.6% compared to a 10y deflation $Q-$probability of 4.2%, and the median adjustment factor for the EZ is 2.26. These high numbers reflect persistent long run deflation fears by market participants.

[17]Note that 10y and 5y5y probabilities as well as adjustment factors are all time-varying and that we report medians for each measure separately; there is thus not a one-to-one mapping in adjustment factor and e.g. the relative size of the 10y and 5y5y probability.

3.2.3 Risk adjustment

We estimate separate risk adjustment factors (P/Q ratios) for high-inflation and deflation states using data for 18 advanced economies between 1875 and 2015. In the data, the output disasters associated with high and low inflation are of different sizes. Namely, episodes of deflation, like in the late 19th century, have not come with particularly severe depressions.[18] The adjustment factor is a mere 0.96. Instead, high inflation more often came with deep recessions, as in most countries during the 1970s, so the adjustment factor is substantial at 0.66.

If, following the literature, we assumed a common adjustment factor, then we estimate it to be 0.82. This is still large, so that not taking risk into account leads to an overstatement of the physical probability of an inflation disaster.

Of course, like any other empirical estimate, these factors depend on the sample. When we estimate our model with data post 1910, so that we do not include the frequent deflations of the late 19th century, then the adjustment factor for deflation falls to 0.91. Deflations are now associated with more serious recessions. At the same time, in this sample, the adjustment factor for high-inflation disasters is also smaller (0.62). Therefore, the difference between the two tails is almost the same, and high inflation still comes with significantly larger output disasters than deflations.

To compare these estimates to the literature, we calculate the corresponding risk premia rp, defined as $q(\pi + rp) = p(\pi)$: the increase in inflation to equate risk-adjusted and actual probabilities. Note that they are positive for high inflation and negative for deflation, but we will refer to their absolute value. We find only moderately high inflation risk premia averaging 0.23%.

This is in line with the literature that used very different methods. Fleckenstein,

[18] As Atkeson and Kehoe (2004) write using different historical data for 17 countries: "the only episode in which there is evidence of a link between deflation and depression is the Great Depression (1929–1934)." Bordo and Filardo (2014) reach the same conclusion.

Longstaff and Lustig (2017, 2016) estimate risk premia in the range of 0.2-0.25% by taking the difference between subjective expectations from analyst forecasts and market expectations from inflation swap rates. The FRB Cleveland estimates the affine term structure model of Haubrich, Pennacchi and Ritchken (2011) and during our common sample its average was 0.39%. However, because we separate high inflation and deflation episodes, we find a significant variability within this average. The average risk premium for high-inflation disasters is 0.61%, while the risk premium for deflation is close to zero.

3.2.4 Comparing disaster probabilities

Combining all the adjustments leads to the final inflation disaster physical probability in the last column of panels A and C. In the US 2021-23 sample, the median 5y5y inflation disaster physical probability was 4.2%, elevated relative to its median of 3.5% in the US full sample and 3.2% in the EZ full sample. These are all below 5%, indicating the success of the Fed and the ECB at convincing market participants that inflation will hover around its target. These probabilities are asymmetric, but in different directions for the two regions. In the US, the probability of deflation in the full sample is lower, at 2.4%, but in the EZ it is higher at 6.3%. Outside of the short period in 2011-12, the probability of deflation in the US is always small, but for the EZ it is significantly higher. As a result, adding the two, the probability of a disaster is lower in the US, 5.9%, relative to the EZ, 9.5%, entirely driven by the higher probability of a deflation disaster for the latter. The next section decomposes these medians into the evolution of the probabilities over time.

4 A history of the anchoring of inflation expectations

A priority in the pursuit of an inflation target is to anchor inflation expectations. Estimates of the market-perceived probability of inflation disasters give an objective measure of success. Ideally, the estimates should be always close to zero. This section shows that they are not, relates their variation over time to the major events in monetary policy since 2010, and compares the success of the Fed and the ECB at anchoring expectations.

4.1 The fear of deflation

Between 2010 and 2020, both the Fed and the ECB feared that the binding zero lower bound for policy rates would give rise to inflation expectations being stuck persistently below target (Blanchard, Dell'Ariccia and Mauro, 2010). Expectations might be anchored below the 2% target as had likely happened in Japan (Borağan Aruoba, Cuba-Borda and Schorfheide, 2017) and was predicted by some theory (Benhabib, Schmitt-Grohé and Uribe, 2002). This situation justified the used of quantitative easing (Eggertsson and Woodford, 2003) together with many other unconventional monetary and fiscal policies, all with the aim of moving the anchor back to 2% (Eggertsson, 2020). The mission reviews of the Fed in 2020 and the ECB in 2021 were justified in great part by the fear of deflation (Federal Reserve System, 2020, European Central Bank, 2021, Reichlin et al., 2021).

Figure 3 shows the evolution of the probability of deflation and serious deflation (less than -1%) at the 5y5y horizon over time. For the US, on the left-hand side, we zoom in on the 2011-14 period because the probabilities are very close to zero after that (with the exception of the pandemic, as discussed below).

At the start of our sample, the US probabilities were high and rising. Investors were perhaps doubtful of the Fed's ability to steer inflation back on target after actual deflation in 2009. 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 afterwards.

Christensen, Lopez and Rudebusch (2015), Kitsul and Wright (2013), Fleckenstein, Longstaff and Lustig (2017), writing near this time, reported much higher probabilities of deflation. There are three reasons behind the discrepancy in this particular episode.

Figure 3: Probabilities of a deflation disaster

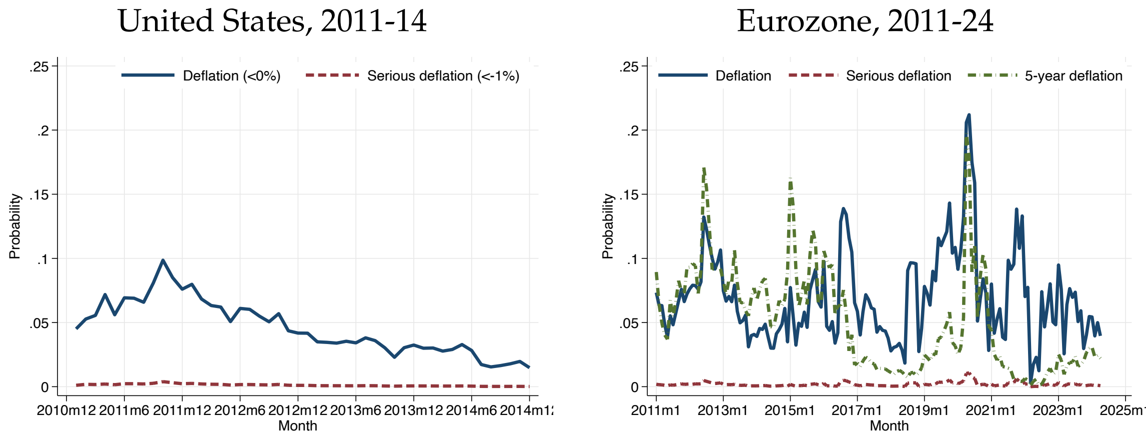

Note: The figure plots 5y5y deflation and extreme deflation probabilities and inflation and risk-adjusted 5y deflation probabilities.

First, those papers mostly focused on deflation in the near horizon, and over a single year, so without the large horizon adjustment factor. Our estimates instead are for the probability of a deflation trap, a persistent period of deflation over 5 years in the long run, the event that policymakers worry most about.[19] Second, and as already indicated in table 1, the inflation adjustment is significant over this long horizon, and without it the probabilities are overstated. Third, as discussed in section 3.2.3, the risk premium for deflation is smaller in our estimates than in the affine models used in previous work that impose a uniform risk-premium. This last adjustment goes in the opposite direction of the other two, since a larger deflation risk premium would make our estimates even smaller, and so even more distant from the previous literature.[20]

[19]During this time, the 10-year actual probability (persistent long-term deflation) was even lower than the 5y5y deflation probability, while the forward risk-neutral probability of deflation in a single year was higher and more volatile. The latter is closer to those in the earlier work. Our estimated Markov model of inflation dynamics shows both strong mean reversion and a high probability of leaving the disaster state as soon as the economy has entered it. Therefore, while deflation was likely, a deflation trap was not.

[20]The appendix shows a version of figure 3 using a pooled risk factor that is constrained to be the same for inflation and deflation. The qualitative conclusions in this section are the same, even if the quantitative estimates are different.

The right panel of figure 3 shows the estimates for the EZ. There are several noticeable differences compared to the US. First, the probability of far-away deflation dropped by less in 2011-14. The ECB justified the use of negative interest rates, quantitative easing, and other unconventional policies starting in 2014 with the aim of fighting deflation.[21] This justification is supported by our estimates, and the policy was initially successful in bringing the estimates persistently below 5% for a little more than one year.

However, second, after a short-lived spike at the end of 2016, the estimates started rising in the middle of 2018 and peaked with the pandemic, exceeding 20% in the second quarter of 2020. A very short-lived spike is also present in the US data for 2020. Both central banks drastically increased quantitative easing and forward guidance in 2020, and the estimates suggest that the fear of a deflation trap caused by the pandemic was reflected in market expectations.

Third, the EZ estimates have since fallen, and have again stabilized below 5% by the end of the sample. However, this is persistently higher than for the US, where this probability has been very close to zero since the start of 2021. Arguably, the market continues to perceive a higher chance that the Eurozone will fall into a deflation trap, in spite of the mission review.

Fourth, the figure shows a third series, for a deflation disaster over the next five years to focus on the near term, as in Boninghausen, Kidd and de Vincent-Humphreys (2018). Until 2016, this tracked the 5y5y probability and was slightly above it. That is, the perception in markets of the probability of an inflation disaster was roughly the same over the next five years as over the following five. But, after that (with the exception of the spike in 2020) the probability of near deflation has been significantly lower than of long-run deflation. Therefore, the fear of deflation in the future is arguably driven not by the cur-

[21] In its 2021 mission review the ECB writes: *The deployment of unconventional monetary policy measures, especially since 2014, has made a significant contribution to countering disinflationary pressures, dispelling deflation concerns and averting a more pronounced downward drift in inflation expectations.*

rent levels of deflation, but by a perception by financial markets that there is something structural in the EZ economy, or in the ECB's mandate and actions, that makes it different from the US and prone to fall into a deflation trap.

4.2 The 2021-24 inflation disaster

After three decades of inflation close to 2%, the three-year period between 2021 and 2024 saw an explosion in the price level in the US and the EZ. Inflation was so high that the 2021-26 five-year period will likely classify as an inflation disaster—average inflation above 4%—in both the US and the EZ, unless average inflation drops below 2% in 2025-26. In December of 2016, the markets had put a probability of his happening at only 2.0% and 1.2%, for the US and the EZ respectively.

The reasons behind the disaster were a combination of supply shocks (Blanchard and Bernanke, 2023), drifting inflation expectations from loose monetary policy (Reis, 2023) and fiscal stimulus (Bianchi and Melosi, 2022), with their relative weights still being debated (Dao et al., 2024). Common to all of them, for the inflation spike to be transitory, it was key that long-run inflation expectations would stay anchored. Policymakers referred to the estimates in this paper in real time (first released in February of 2021, and updated regularly on our websites), and variants of them, to assess this risk (for two examples, see Schnabel (2022) and Gopinath (2022)).

We now have the benefit of hindsight to revisit this turbulent period. The top-left plot of figure 4 shows the 5y5y probabilities of a high-inflation disaster, for the US and the EZ. The estimates tell a story in three stages.

First, throughout 2020, the probability of a high-inflation disaster was low and similar in the two regions, hovering between 0% and 3%. By 2021, the probability started rising for the US alone, reaching 5% by the end of the year, while in the EZ, it stayed low and constant. As both central banks kept monetary policy loose, the market-perceived proba-

Figure 4: Perceptions of a future inflation disaster during the 2021-24 inflation disaster

(a) Probability of a high-inflation disaster

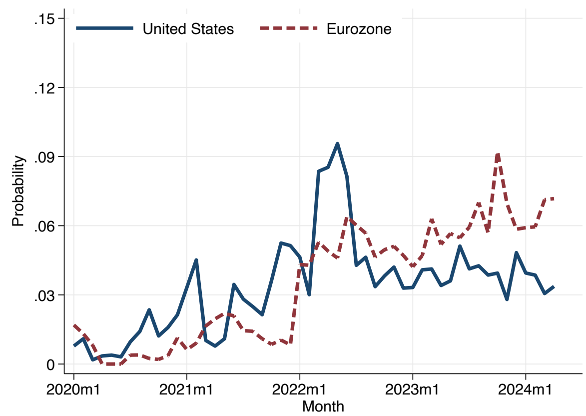

(b) Probability of high inflation in 5y horizon

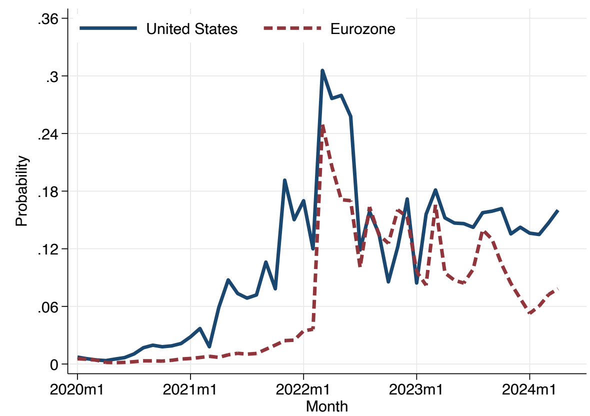

(c) US risk-neutral densities 10-year horizon

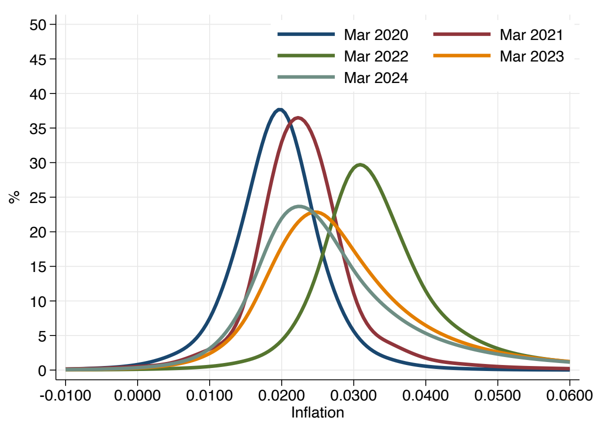

(d) EZ risk-neutral densities

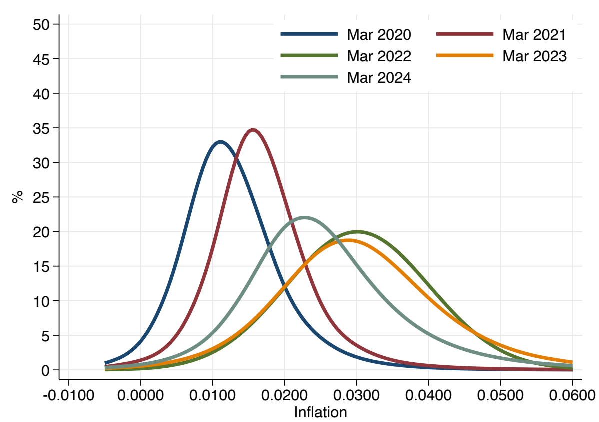

Note: Top panels: 5y5y (forward) and 5y (near term) inflation disaster ($> 4\%$) probability. Bottom panels: 10y horizon risk neutral (Q) densities.

bility of high-inflation in the US rose in tandem with the sharp increase in actual inflation at the time. From the perspective of economic theory, this evolution suggests that even long-horizon expectations are sensitive to extreme current realizations.

Second, in the first half of 2022, the US probability kept rising, peaking at 10% in May. The Fed reacted aggressively to the rise in inflation with a 50 basis point hike on May 5th, that was followed by several more so that rates between the start of May and the end of December that year increased by 400 basis points. At the same time, likely in response, the

US probability of a disaster sharply fell, reaching 3% by the end of 2022. It has remained close to that value since. The sharp policy adjustment, from being highly accommodative to aggressively fighting inflation, came with a corresponding change in inflation expectations. Expectations were unanchored for more than one year, but reanchored once the policy priorities were reestablished.

Turning to the EZ, as inflation took off, so did the EZ disaster probabilities. Since July of 2022, they have been above those of the US, averaging 6%. Similar to the US, this peak coincided with the change in monetary policy, as the ECB increased its policy rates for the first time in 11 years at its 27th of July meeting. Not having peaked as high as the US, the probability of a disaster in the EZ also did not fall much, and increased again in 2023.

Third, at the end of our sample, at the start of 2024, the probability of an inflation disaster is hovering between 3% and 4% in the US, and in the 6-7% range in the EZ. Both are significantly larger than before 2021, when these probabilities had been near 1% in both regions for many years. The inflation disaster has left scars for the future.

To summarize, our estimates show that (i) inflation expectations deanchored in 2021-22, (ii) the tightening of policy had a noticeable effect in stopping or reversing that deanchoring, and (iii) there is a scar from the episode into the future as the probability of an inflation disaster has been permanently higher than before.

The top-right panel of figure 4 digs deeper by looking at the 5y probabilities, to assess the market perceptions of an immediate disaster. Interestingly, they increased in 2021 and 2022 together with actual inflation and with the baseline far-away disaster probabilities, but by much more (the axis scale is doubled). Probabilities peaked at 31% in the US and 25% in the EZ, both in March of 2022. This suggests a tight link between actual inflation, forecasts of inflation, and anchoring of expected inflation. It is consistent with models where the credibility of the central bank depends on its current performance.

The bottom row of figure 4 shows instead the 10-year risk-neutral distributions (so without horizon or risk adjustment) at different points in time, to understand the shifts behind the movements in the tails. In 2021-22, the increase in the tails of the distributions was much more pronounced than the increases in the median. Skewness rose significantly. Then, by 2024, distributions shifted back, but not to where they started in 2020-21. Looking at the prices of the 10-year inflation swaps, which match the mean of these distributions, gave an impression of only a slight unanchoring during this time (recall figure 1). Looking instead at the probabilities of a disaster provided in this paper shows a much more worrying drift, and a clear impact of monetary policy.[22]

Based on the experience of the 5y5y inflation swaps (figure 1) one might assume that inflation expectations stayed anchored throughout. Indeed, Powell (2024) concludes that this contributed to inflation stabilizing quickly after monetary policy tightened, with little impact on unemployment. The estimates in this paper support a different interpretation: expectations deanchored, policy had to and partially succeeded in reanchoring them, but even now they have not gone back to their initial state.

4.3 Measures of anchoring

The simplest definition of anchored expectations is that, unconditionally, they do not move at all. The equivalent for our disaster probabilities is that these are close to zero almost always. A conditional statement of anchoring is instead whether expectations of future inflation are insensitive to realizations of inflation. The discussion of the 2021-24 episode suggests that they are not, and that this depends on the stance of monetary policy. We now explore this further by taking advantage of our methods that give probabilistic measures for the tails.

The baseline estimates of the probability of a high-inflation disaster in a month are conditional on what inflation was in that month (as well as the parameters of the model,

[22] The distribution of survey expectations of inflation in the 1970s shows similar behavior (Reis, 2022).

some of which are time-varying). Given the estimated model of inflation dynamics, we calculated what that probability would have been if this initial inflation was a different value. The top row of figure 5 plots the answers for two hypothetical initial inflation rates: between 2% and 3% and between 3% and 4%. It does so from 2022 onwards, when inflation was well above the two hypothetical ranges: before that, for much of the sample, as inflation was in these ranges, baseline and one of the hypotheticals coincide. The hypotheticals ask to what extent was the heightened probability of an inflation disaster not the result of currently elevated inflation, but of the updated probabilities on future shocks and persistence of inflation embedded in the option prices.

For the US, on the left, the hypotheticals are practically indistinguishable from the baseline. This suggests that initial conditions have little impact on the market-perceived probability of an inflation disaster five years down the line. On the one hand, this implies that the higher perceived probability of a disaster in the US was not just the result of inflation being high, but rather was driven by changes in perceived future shocks and dynamics. This explains why, even as inflation fell back near its target in 2024, the probability of an inflation disaster remained unchanged. On the other hand, this shows that changes in disaster probabilities can be interpreted as reflecting changes in future policy credibility. Inflation expectations in 2021-22 therefore were temporarily unanchored, since the increase in the disaster probability was not driven by high inflation realizations. Overall, one might say that expectations can be quick to unanchor in the US, but also quickly reanchor.

For the EZ, the opposite is true. Had inflation been lower, then the market-perceived probability of an inflation disaster would have been significantly lower. This only stops being the case from the second half of 2023 onwards. Before that, disaster probabilities were high mainly because of high inflation realizations, not because of a change in expected future policy.

Figure 5: The conditional anchoring of expectations

(a) The influence of initial conditions in the US

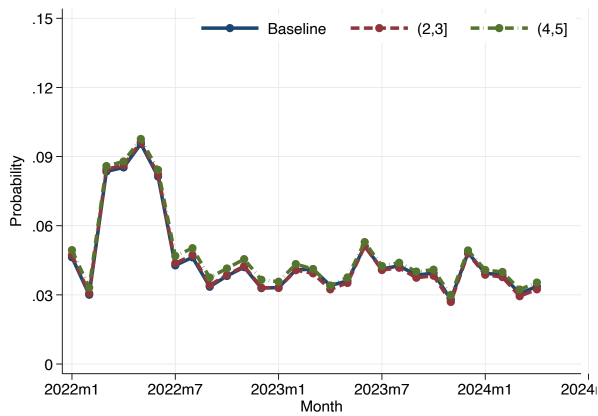

(b) The influence of initial conditions in the EZ

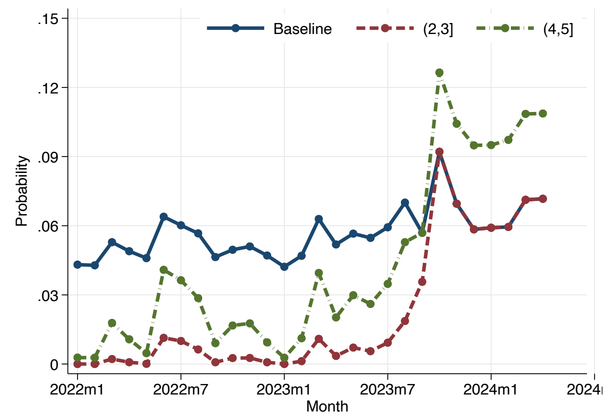

(c) Conditioning on the near future in the US

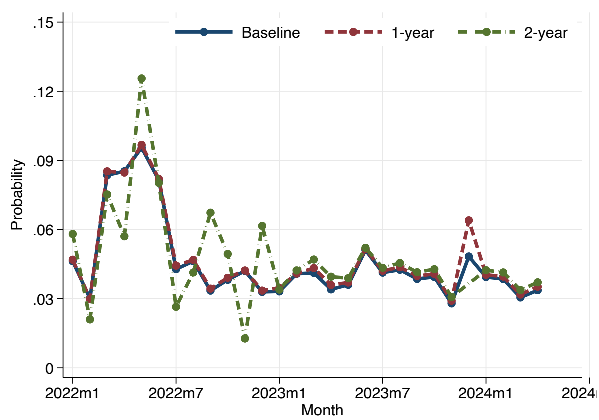

(d) Conditioning on the near future in the EZ

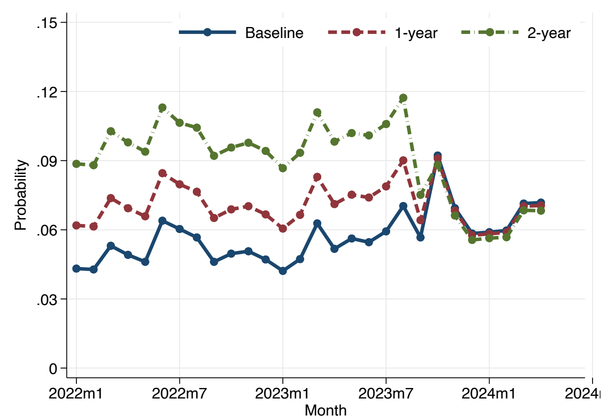

Note: The figure reports various conditional 5y5y inflation disaster ($> 4\%$) probabilities. Top row: Baseline, based on actual current inflation, and varying current inflation (over the previous year) to lying in different ranges, either 2%-3% or 3%-4%; bottom row: Changing, in addition, inflation over the next two years.

The bottom row of figure 5 calculates two separate counterfactuals. Taking as given that inflation was high, it asks what would have been the market-perceived probability of a disaster conditional on knowing that inflation would stay high in the next year, as well as the year after. These conditional probabilities measure anchoring by showing whether the option prices at different horizons expect high inflation to persist.

Again, for the US, the conditional and unconditional probabilities are quite similar.

Even if markets were convinced in 2022 that inflation would stay high for the next two years, that would not change their trust that the Fed would prevent a disaster in 2027-32. But this is not the case for the EZ before the middle of 2023. Markets expect that if inflation persists for one or two years, the willingness or ability of the ECB to prevent a far-ahead inflation disaster is lower. From this perspective as well, inflation expectations were less well-anchored in the EZ than in the US.

5 Empirical implementation of the three adjustments

Finally, we discuss our empirical implementation of the adjustment factors to the N-probabilities in proposition 1.

5.1 Inflation adjustment: risk-neutral probabilities

Recovering Q-probabilities is just as easy as recovering N-probabilities, as it relies on the same methods from options pricing, and it should always be done.[23]

5.1.1 Recovering N-probabilities

The no-arbitrage pricing condition for the traded securities is:[24]

$$a(k) = \sum_{\pi} \left(p(\pi) m(\pi) \max \left\{ \frac{e^{\pi} - k}{e^{\pi}}, 0 \right\} \right) \approx \int_{k}^{\infty} \left(\frac{e^{\pi} - k}{e^{\pi}} \right) b(\pi) d\pi, \tag{7}$$

where the approximation comes from assuming a continuum of inflation states and using the definition of the Arrow-Debreu prices in equation (3).

[23]In the appendix, we compare our daily N-probabilities with those from KW in the overlap sample for the US data (October 2009 to April 2013). They are almost identical.

[24]Note that the payoff of these securities only depends on inflation, not on the entire set of states $\tilde{p}(s)$, defined in the appendix.

Following Breeden and Litzenberger (1978), recalling the definition of $n(\pi) = b(\pi)e^{i-\pi}$, 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:

$$N(k) = 1 + e^{i}a'(k). \tag{8}$$

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 gross inflation k, which is what is usually reported in the financial media.

However, from equation (3) and the definition of $n(.)$, then $n(\pi) = p(\pi)$ 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 which case the exercise of building a distribution of inflation is not interesting. We therefore need to adjust these conventional N-probabilities.

5.1.2 Risk-neutral probabilities (Q)

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.

One could obtain the $q(.)$ from the data on $n(.)$ using equation (4). The adjustment is needed because, even if investors are risk neutral, they still care about receiving a payoff in a high-inflation state that has lower real value.

Alternatively, going back to equation (7), take derivatives with respect to k:

$$e^r a'(k) = -\int_k^\infty e^{-\pi} q(\pi) d\pi. \tag{9}$$

Taking another round of derivatives with respect to k gives:

$$Q(k) = e^r k a''(k). \tag{10}$$

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

5.2 Horizon adjustment: forward probabilities

Obtaining forward expectations of inflation is straightforward (figure 1). Starting with measures of expected inflation between the present and a far-away date, T, and between the present and a farther-away date $T + H$, the joint linearity of the expectations operator and of inflation (as a difference in logs) implies that $\mathbb{E}^q(\pi_{T,T+H}) = \mathbb{E}^q(\pi_{0,T+H}) - \mathbb{E}^q(\pi_{0,T})$. Going from probabilities on cumulative inflation to probabilities over a forward period is harder. As proposition 1 shows, it requires more data beyond the two distributions for cumulative inflation, as well as a model for the time-series properties of inflation as perceived by markets.

5.2.1 Data on forward starting options

There exist markets for forward-dated options at every date that will pay out depending on the realizations of inflation in $\pi_{T,T+1}$. These options are for inflation in one given year, not on the average over a longer period $H > 1$, which is our focus.[25]

[25] 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.

The markets in which these trade are less liquid, so we want to be conservative in using them. In the data, the five distributions covering the one-year ahead inflation starting in 5 to 9 years are quite similar almost always. 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 possible 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 5.1, this provides an estimate of $q(\pi_{5,6}) \approx ... \approx q(\pi_{9,10})$.

5.2.2 A model of inflation persistence

Since the data is for risk-neutral inflation, the model of dynamics is for risk-neutral inflation as well. We assume that inflation is the sum of three parts: a deterministic part, which has been constant at the inflation target during our sample $\bar{\pi}$; a stochastic one capturing the ups and downs during normal times ε; and a stochastic one capturing sharp jumps during disasters, which may be positive d^h or negative d^l. Altogether, letting Δ be a time period:

$$\pi_{t+\Delta} = \bar{\pi} + \varepsilon_{t+\Delta} + d^h_{t+\Delta} - d^l_{t+\Delta}. \tag{11}$$

We assume that d^h_t and d^l_t are two independent common disasters that arrive as Poisson processes. We then make two major assumptions on ε_t. 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 follows a first-order Markov process with a particular set of restrictions on the transition matrix.[26]

Because strike prices for inflation options come in jumps of 0.5%, our data comes in 8 bins: $\pi(i) = \{\leq -1, (-1,0], (0,1], (1,2], (2,3], (3,4], (4,5], > 5\}$. A discrete approximation of this process is a Markov chain over 8 states corresponding to these bins with an 8×8 Markov transition matrix $\mathbf{P}$:

$$\mathbf{P} = \begin{bmatrix} 1-5p_l & p_l & p_l & p_l & p_l & p_l & 0 & 0 \\ p_{dl}+p_{nn} & p_{ml} & p_{mr} & 0 & 0 & 0 & 0 & 0 \\ p_{dl} & p_{nn} & p_m & p_{mr} & 0 & 0 & 0 & p_{dh} \\ p_{dl} & 0 & p_{nn} & p_n & p_{nn} & 0 & 0 & p_{dh} \\ p_{dl} & 0 & 0 & p_{nn} & p_n & p_{nn} & 0 & p_{dh} \\ p_{dl} & 0 & 0 & 0 & p_{mr} & p_m & p_{nn} & p_{dh} \\ 0 & 0 & 0 & 0 & 0 & p_{mr} & p_{mh} & p_{dh}+p_{nn} \\ 0 & 0 & p_h & p_h & p_h & p_h & p_h & 1-5p_h \end{bmatrix}. \tag{12}$$

Starting with the low-inflation disaster state in the first row, the economy exits it with probability $5p_l$, which should be close to 1 to match the Poisson 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

[26] Mertens 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, in 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.[27] 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.[28]

5.2.3 Estimating the model

There are six parameters to estimate: 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 a set of parameters, we simulate many paths to calculate the probability distributions for inflation at the different horizons.

The data consists of 21 numbers per month, 7 for each of the three distributions: the cumulative distributions $q(\pi_{0,5})$ and $q(\pi_{0,10})$, and the average forward distribution $q(\pi_{5,6})$. These are the moments that the model must hit, in a GMM procedure that assigns them equal weight. The overall fit, which we report in the appendix, is quite good.

In principle, we could estimate the model separately at each month, and recover parameters that are specific to each date. For parsimony, instead, we kept three of the parameters fixed over the whole sample, while letting the other three vary across months. We estimated several other candidate models, including models where four of the parameters vary over time, and where all 6 did so, as well as one where some parameters move

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

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

at an annual while others move at a monthly frequency. The appendix discusses these models, and why our results are robust to this choice, and why the parsimonious setup is preferable.

The first constant parameter is p_{mr}, which captures the extent of mean reversion. For the US, the estimate is 0.50, while for the EZ it is 0.47, capturing the strong sluggishness of inflation. The other two are the exit probabilities for disasters, p_l, p_h. The estimates are almost exactly the same for the US (0.1990 and 0.1998), implying that, as soon as the US enters a disaster state, it leaves it with probability $5 \times 0.199 \approx 1$ shortly after. For the Eurozone, instead they are 0.1999 and 0.0617, so that a high-inflation disaster is perceived to persist for more than one year in the EZ with probability of 69%.

The three time-varying parameters are the probabilities of entering a disaster, the main focus of our interest, and p_{nn}, which captures local volatility and captures the changes in inflation volatility over the sample. Figure 6 shows their estimates over time. The decline in p_{nn} for the US since the start of the decade captures a fall in the perceived volatility of inflation, although since the pandemic that trend has reverted. Independently of this, the probability of jumping to a low-inflation disaster was high at the start, but became quite low after mid-2012, although with a significant jump in 2020. More erratic is the pattern in the 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, and a similar uptick since the pandemic. 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 during the pandemic. The probability of a high inflation disaster stays small throughout, but rises at the very end of the sample.

Figure 6: Inflation dynamics: model parameter estimates

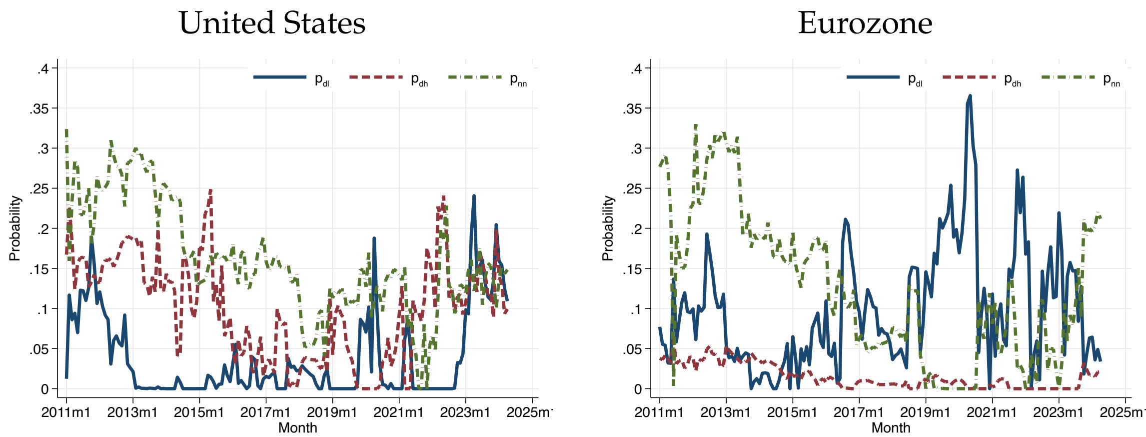

Note: The figure plots time-varying Markov model transition probabilities estimates.

5.3 Risk adjustment: physical probabilities

If the Phillips curve was vertical 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.[29] 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 physical-measure 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.

[29]Note that people may still be arbitrarily risk averse: the stochastic discount factor may still be volatile, so that $\tilde{m}(s)$ over all the states 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.

5.3.1 A model of risk in inflation disasters

We expand on our model of inflation dynamics by considering how it co-moves with consumption c_t. Their joint dynamics are:

$$\pi_{t+\Delta} = \bar{\pi} + \overbrace{u^{\pi}_{t+\Delta} + e_{t+\Delta}}^{\varepsilon_{t+\Delta}} + d^h_{t+\Delta} - d^l_{t+\Delta}, \tag{13}$$

$$\log(c_{t+\Delta}) = \log(c_t) + g + u^c_{t+\Delta} + \beta_0 e_{t+\Delta} - \beta^h d^h_{t+\Delta} - \beta^l d^l_{t+\Delta}. \tag{14}$$

Consumption is expected to grow at rate g subject to some shocks $u^c_{t+\Delta}$ that are independent of inflation shocks $u^{\pi}_{t+\Delta}$, but co-moving in normal times due to the common shock e_t. This may be driven by multiple shocks, and may be correlated over time, but the parameter determining inflation risk premia during normal times is the co-movement scalar β_0.

Our focus is on d^h_t, d^l_t, the disasters that strike inflation, and which are non-zero with probabilities p^h, p^l. (Consumption disasters that do not come with high or low inflation do not trigger the options, so they are included in $u^{\pi}_{t+\Delta}$ and $u^c_{t+\Delta}$, respectively.) The coefficient β^h measures the size of the consumption drop when there is a high inflation disaster, while the coefficient β^l measures the size of the drop following a deflation disaster.

We follow and modify the approaches of Gabaix (2012) and Barro and Liao (2021) to model the size of the disaster. Defining the inverse fall in consumption by $z^h = 1/(1 - \beta^h d)$, we assume that if a disaster strikes, then z^h follows a Pareto distribution:

$$F(z^h) = 1 - \left(\frac{z^h}{z^h_0}\right)^{-\alpha^h} \quad \text{with } \; z^h \geq z^h_0 > 1, \alpha^h > 0. \tag{15}$$

The Pareto distribution has two parameters. The first, z^h_0 is the minimum size of the jumps. The higher it is, the more average consumption falls during inflation disasters.

The second is the exponent α^h capturing how quickly the tail of the distribution thins out. The lower it is, the more likely is a very large consumption disaster. The same applies for deflations, (z^d, z_0^d, α^d).

5.3.2 Estimating the model

We combine data on annual output from Barro (2006) (using real GDP per capita, as it 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, and we use it to identity periods where both inflation and output had disasters and estimate the Pareto distribution as well as the co-movement parameters.

Details of the estimation and alternative estimation approaches are reported in the appendix. In the sample, the unconditional probability of an inflation disaster (10 log points above or below target for 5 years) is 12.9%, and they overlap with output disasters in 20.0% of the cases. Therefore, $\tilde{p} = 0.20$. Separating high and low inflation disasters, then $\tilde{p}^h = 0.356$ and $\tilde{p}^d = 0.085$. For the Pareto distribution, we estimate that $\alpha^h = 5.45$ and $z_0^h = 1.03$ for high-inflation disasters and $\alpha^l = 15.18, z_d^0 = 1.06$ for deflation disasters. That is, deflation disasters more rarely come with output disasters, and when they do, the falls in output are on average higher, but with significantly thinner tails.

Given an estimated model of inflation-consumption disaster co-movement, we follow the standards of the rare-disasters literature (Gabaix, 2012, Barro and Liao, 2021) by using an Epstein-Zin model for marginal utility with a relative risk aversion coefficient of 3.

5.4 Assessing the uncertainty around the adjustments

Two of our adjustments—horizon and risk—required statistical estimates, so they come with estimation uncertainty. For the risk adjustment, we estimated the parameters of the Pareto distribution. A bootstrap over the inflation-output data provides confidence bands. For the horizon adjustment, we estimated a statistical model using GMM. The covariance matrix of the parameter estimates has the standard GMM asymptotic formula. Finally, the delta method translates these to estimates of the uncertainty around the probability of a disaster.

Figure 7 shows the 90% confidence band around the estimates of a high-inflation disaster for the US. On the left panel are confidence bands treating the horizon adjustment as known, so only for the Pareto estimates for the risk adjustment. The right panel zooms in on the more recent period to be clearer, and considers the estimation uncertainty on horizon as well. The bands are relatively tight, between 1.5 and 2.5 percentage points in width in the more recent period. This shows that the estimation uncertainty around our two adjustment factors is relatively minor (there is no estimation needed when applying the inflation adjustment). More relevant is the time-series variation in the estimates, which is driven by changes in the market expectations reflected in the options price data.

Figure 7: Confidence bands from the adjustments

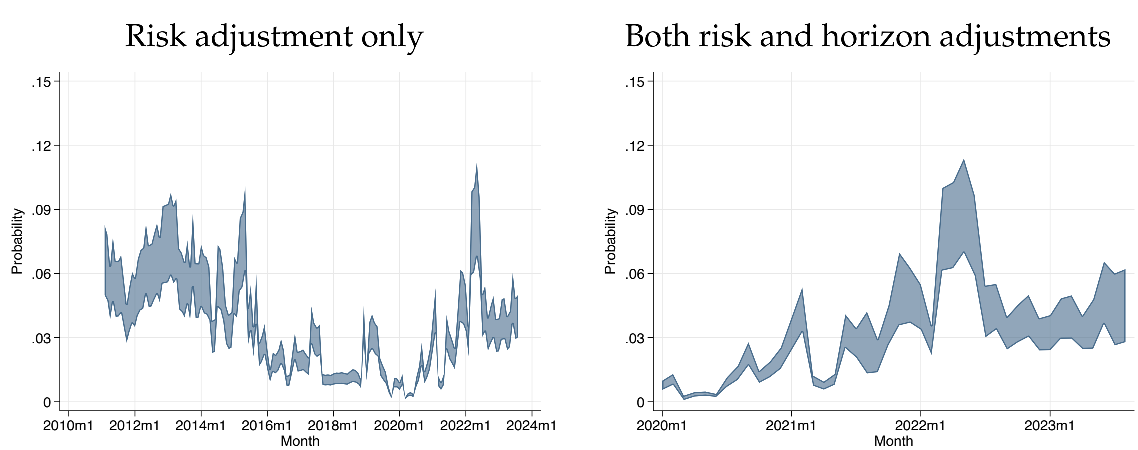

Note: The left panel shows the 90% confidence band for the US 5y5y inflation disaster ($>4\%$) probability when standard errors take into account the uncertainty in the risk adjustment estimate; the right panel adds uncertainty in the horizon adjustment.

5.5 Using market based data and liquidity

Empirical results always depend on data quality. In the case of price data, the common concern is market liquidity. As always, this requires a brief discussion of the source of the data to have the right care in using our estimates and interpreting our results.

The derivatives market for inflation started in 2002, and grew very quickly. Conservatively, we follow several other studies (Kitsul and Wright, 2013, Fleckenstein, Longstaff and Lustig, 2017, Mertens and Williams, 2021, Hilscher, Raviv and Reis, 2022, Nagel, 2016) and only use data from 2009 onwards, when the market was quite liquid. Since then, Chipeniuk and Walker (2021) report that the volume of trading for inflation caps and floors quadrupled between 2009 and 2017. Since 2021, some claim that the US interdealer market has virtually disappeared (Williams, 2023), but others report that since the pandemic, the market has been driven by clients' increased demand for inflation protection, predominantly through inflation caps in the dealer-to-client market.[30] The variation in our estimates during these times is also reasonable and related to policy events. Moreover, since the inflation options are actively used to hedge positions in inflation swaps, their liquidity concerns should be related. Bahaj et al. (2023) estimate liquidity premia in the swaps market, and find that they are moderate at the long horizons that we focus on in this paper, and that fundamentals drive more than 90% of the variation in prices.[31] Nevertheless, to be conservative, in our analysis of the recent history of US and EZ inflation, we focused on trends across many months, rather than month-to-month fluctuations, since these are likely to be more robust to liquidity changes.

Policy institutions use inflation options data regularly, to provide public data (the Fed-

[30] https://www.risk.net/awards/7955889/inflation-derivatives-house-of-the-year-citi report

[31] Ideally, future work would propose a fourth adjustment factor that captures potentially time-varying illiquidity of the option contracts. The literature is still far from delivering this, and the direction of its impact is not obvious: since all option prices move together with the real risk-free rate, only movements in the differential liquidity of options with strike prices that are nearer or more out of the money would affect the constructed probabilities.

eral Reserve Bank of Minneapolis produces a weekly series), research (Kozlowski, 2024), and in speeches justifying policies (e.g., Schnabel, 2022, Lane, 2022). Policy interventions themselves can sustain liquidity in markets (Allen, Carletti and Gale, 2009, Kargar et al., 2021, Falato, Goldstein and Hortaçsu, 2021). An example of such proactive involvement is the creation of the FX options market by the Bank of Israel during the high inflation period of the 1980s (Fischer, 2006).

An alternative to market-based data is to use household and professional surveys (Armantier et al., 2022, Grishchenko and Wilcox, 2024, Fofana, Patzelt and Reis, 2024). However, the tails of surveys reveal the extent of disagreement, while the tails from market options reveal perceptions of rare disasters. The two are conceptually very different. In the same way that market prices may be affected by liquidity, surveys can be affected by biases in answering, so that combining them to extract as much information as possible is likely ideal (Nagel, 2024).

6 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) inflation options' nominal payoffs need to be adjusted to get real Arrow-Debreu probabilities; (ii) disaster probabilities for forward horizons can differ from short or long horizons because of the sluggishness of inflation; and (iii) the risk premium for inflation is not the same at its two tails compared to the center of the distribution. We provide simple, but we hope robust, methods to make all of these adjustments. We show that the adjustments are quantitatively large relative to constructing probabilities using conventional methods.

We apply our methods to data from the US and the Eurozone between 2009 and April of 2024. Starting with the market perceptions of a deflation trap, contrary to previous wisdom, we find that they were low and short-lived in the US 2011-14, but have persisted in the EZ and unconventional monetary policies only provided temporary respite. Turning to high inflation, we find a significant deanchoring of inflation expectations that peaked in mid 2022, and then reanchored as monetary policy tightened. By the end of the sample, we find scars of the high-inflation episode in persistently elevated probabilities of a future inflation disaster. Finally, temporary shocks to inflation, either in the recent past or recent future, have a larger influence on the expectations anchor in the EZ than in the US.

In the future, we hope that our new methods will allow researchers to measure the risk of inflation disasters more accurately. Policymakers will continue to benefit from our measures in assessing the extent to which changes in policy, economic fundamentals, or temporary or permanent shocks affect inflation expectations and the inflation anchor.

References

Allen, Franklin, Elena Carletti, and Douglas Gale. 2009. "Interbank market liquidity and central bank intervention." *Journal of Monetary Economics*, 56(5): 639–652.

Andrade, Philippe, E. Ghysels, and Julien Idier. 2012. "Tails of Inflation Forecasts and Tales of Monetary Policy." Banque de France working paper 407.

Armantier, Olivier, Argia Sbordone, Giorgio Topa, Wilbert Van der Klaauw, and John C Williams. 2022. "A new approach to assess inflation expectations anchoring using strategic surveys." *Journal of Monetary Economics*, 129: S82–S101.

Atkeson, Andrew, and Patrick J. Kehoe. 2004. "Deflation and Depression: Is There an Empirical Link?" *American Economic Review*, 94(2): 99–103.

Bahaj, Saleem, Robert Czech, Sitong Ding, and Ricardo Reis. 2023. "The Market for Inflation Risk." *Bank of England working paper 1028*.

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*.

Benhabib, Jess, Stephanie Schmitt-Grohé, and Martín Uribe. 2002. "Avoiding Liquidity Traps." *Journal of Political Economy*, 110(3): 535–563.

Bianchi, Francesco, and Leonardo Melosi. 2022. "Inflation as a Fiscal Limit." In *Reassessing Constraints on the Economy and Policy*. FRB Kansas City Jackson Hole Symposium.

Blanchard, Olivier, Giovanni Dell'Ariccia, and Paolo Mauro. 2010. "Rethinking Macroeconomic Policy." *IMF Staff Position Note 10/03*.

Blanchard, Olivier J., and Ben S. Bernanke. 2023. "What Caused the US Pandemic-Era Inflation?" *NBER working paper 32497*.

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, Iussue 6/2018*.

Borağan Aruoba, S, Pablo Cuba-Borda, and Frank Schorfheide. 2017. "Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries." *The Review of Economic Studies*, 85(1): 87–118.

Bordo, Michael, and Andrew Filardo. 2014. "Deflation and monetary policy in a historical perspective: remembering the past or being condemned to repeat it?" *Economic Policy*, 20(44): 800–844.

Breeden, Douglas T, and Robert H Litzenberger. 1978. "Prices of state-contingent claims implicit in option prices." *Journal of Business*, 51(4): 621–651.

Chipeniuk, Karsten O, and Todd B Walker. 2021. "Forward inflation expectations: evidence from inflation caps and floors." *Journal of Macroeconomics*, 70: 103348.

Christensen, Jens HE, Jose A Lopez, and Glenn D Rudebusch. 2010. "Inflation expectations and risk premiums in an arbitrage-free model of nominal and real bond yields." *Journal of Money, Credit and Banking*, 42: 143–178.

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.

Dao, Mai Chi, Pierre-Olivier Gourinchas, Daniel Leigh, and Prachi Mishra. 2024. "Understanding the International Rise and Fall of Inflation since 2020." *Journal of Monetary Economics*, forthcoming.

Eggertsson, Gauti. 2020. "A Unified Theory of the Great Depression and the GReat Recession." Brown University manuscript.

Eggertsson, Gauti B., and Michael Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, 34(1): 139–235.

European Central Bank. 2021. "Strategy Review." *https://www.ecb.europa.eu/mopo/strategy/strategy-review/html/index.en.html*.

Falato, Antonio, Itay Goldstein, and Ali Hortaçsu. 2021. "Financial fragility in the COVID-19 crisis: The case of investment funds in corporate bond markets." *Journal of Monetary Economics*, 123: 35–52.

Federal Reserve System. 2020. "Review of Monetary policy Strategy, Tools, and Communications." *https://www.federalreserve.gov/monetarypolicy/review-of-monetary-policy-strategy-tools-and-communications.htm*.

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*.

Fischer, Stanley. 2006. "Reflections on one year at the Bank of Israel."

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.

Fofana, Salomé, Paula Patzelt, and Ricardo Reis. 2024. *Household disagreement about expected inflation.* London School of Economics and Political Science.

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.

Gopinath, Gita. 2022. "How will the pandemic and war shape future monetary policy?" Speech at economic policy symposium of FRB Kansas City, Jackson Hole, Wyoming.

Grishchenko, Olesya V, and Laura Wilcox. 2024. "Tale About Inflation Tails." *FEDS Working Paper 2024-28*.

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, George Pennacchi, and Peter Ritchken. 2012. "Inflation expectations, real rates, and risk premia: Evidence from inflation swaps." *Review of Financial Studies*, 25(5): 1588–1629.

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.

Hördahl, Peter, and Oreste Tristani. 2012. "Inflation risk premia in the term structure of interest rates." *Journal of the European Economic Association*, 10(3): 634–657.

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.

Kargar, Mahyar, Benjamin Lester, David Lindsay, Shuo Liu, Pierre-Olivier Weill, and Diego Zúñiga. 2021. "Corporate bond liquidity during the COVID-19 crisis." *Review of Financial Studies*, 34(11): 5352–5401.

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. 2012. "The Economics of Options-Implied Inflation Probability Density Functions." National Bureau of Economic Research.

Kitsul, Yuriy, and Jonathan H. Wright. 2013. "The Economics of Options-Implied Inflation Probability Density Functions." *Journal of Financial Economics*, 110(3): 696–711.

Kozlowski, Julian. 2024. "Beyond the Mean: Exploring Tail Risks in Inflation Expectations." Federal Reserve Bank of St. Louis.

Lagarde, Christine. 2021. "Monetary policy during an atypical recovery." *Speech at ECB Forum on Central Banking, Sintra, Portugal*.

Lane, Philip R. 2022. "Monetary policy in the euro area: the next phase." Presentation, Remarks at the Annual 2022 meeting of Central Bank Research Association (CEBRA), August 29, Barcelona.

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.

Nagel, Stefan. 2016. "Long-run inflation uncertainty." *International Journal of Central Banking*, 12: 2007–2017.

Nagel, Stefan. 2024. "Stock Market Valuation: Explanations, Non-Explanations, and Some Open Questions." *Comments at NBER summer institute 2024*.

Powell, Jerome H. 2021. "Monetary Policy in the Time of COVID." *Speech at economic policy symposium of FRB Kansas City, Jackson Hole, Wyoming*.

Powell, Jerome H. 2024. "Review and Outlook." Speech at Jackson Hole Symposium.

Reichlin, Lucrezia, Klaus Adam, Warwick McKibbin, Michael McMahon, Ricardo Reis, Giovanni Ricco, and Beatrice Weder di Mauro. 2021. *The ECB Strategy: The 2021 Review and Its Future.* CEPR Press.

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.

Reis, Ricardo. 2023. "The Burst of High Inflation in 2021-22: How and Why Did We Get Here?" In *How Monetary Policy Got Behind the Curve–And How to Get it Back.* , ed. Michael Bordo, John Cochrane and John Taylor, 203–252. Hoover Institution Press.

Ryngaert, Jane. 2022. "Inflation Disasters and Consumption." *Journal of Monetary Economics*, forthcoming.

Schnabel, Isabel. 2022. "Monetary policy and the great volatility." Vol. Speech at economic policy symposium of FRB Kansas City, Jackson Hole, Wyoming. Speech at economic policy symposium of FRB Kansas City, Jackson Hole, Wyoming.

Williams, John C. 2023. "Elementary, Dear Data." Federal Reserve Bank of New York.

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

Jens Hilscher | Alon Raviv | Ricardo Reis
UC Davis | Bar-Ilan University | LSE

September 2024

This appendix consists of four sections. Section A fills in additional details for the theory that, for the sake of brevity, were not included in the main text. Section B discusses how we obtain the probability distributions for inflation from option prices; this section also includes a detailed comparison with probabilities reported by Kitsul and Wright (2013). Section C provides additional detail and discussion on how we estimate inflation dynamics to adjust for the horizon, while Section D does the same for the risk adjustment factor.

A Theory

This section extends and clarifies the setup leading to proposition 1. Note that the following is designed to be read together with the main paper; the discussion in this appendix is not meant as a stand alone summary.

A.1 Multiple sources of uncertainty

In the main text we assumed that inflation was a random variable and the only source of uncertainty. We now generalize this to many sources of uncertainty.

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 $\hat{p}(s) > 0$ for all s and $\sum_{s \in S} \hat{p}(s) = 1$. Inflation is one of many random variables, 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 $p(.)$, not $\hat{p}(.)$, 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.

A.2 Arrow Debreu and inflation securities

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 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:

$$b(\pi) = p(\pi)m(\pi), \tag{A1}$$

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.

To complete the discussion, note that the risk-neutral probabilities then follow the same steps as in the text, in spite of this broader setup. 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. As we describe in detail in the main text, 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 $a(\pi) = b(\pi)e^{-\pi}$.

B Data: constructing marginal distributions of inflation

The paper uses data on two sets of distributions. First, with zero-coupon inflation caps and floors options, we construct distributions of cumulative inflation 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.

B.1 Data pre-cleaning

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. As an additional check we also compared the put-call-parity implied real rates to the inflation swap real rates. Differences were very small, implying across market consistency of pricing.

B.2 Implied volatility smoothing

Next, we 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.

B.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 in 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%.

B.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.

B.4.1 Periods of sparse data 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 in 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 6 to 10-year horizons, this adjustment has little effect.

C Comparison of $N-$probabilities with those in Kitsul and Wright (2013)

We now compare our N-probabilities of disaster with those reported by Kitsul and Wright (2013).

KW obtained their probabilities from BGC Partners, since at the time, as the market was starting these data were not easily available. Instead, we obtained data from a standard Bloomberg terminal. Also, because data quality is much higher today, our data cleaning procedures reported in the previous section were significantly stricter. Finally, they approximated the $n(.)$ density away form the bins using a locally linear regression, while we avoided any approximation and simply stuck to a Markov chain for the data.

Figure 1 compares their estimates with ours. We cannot reproduce their results (the BCG data is not available anymore), but we copy and paste the figures in their paper, which show the probabilities of average inflation lying above 4% or below zero over the 5y and 10y horizons from October 6, 2009 until April 1, 2013. We then also plot our own measure of the $N-$probabilities. In our analysis, to be conservative, we constructed monthly disaster probabilities. For the purposes of this comparison, we have also constructed a daily series.

It is easy to confirm visually that the data are almost the same. Not only the correlation is high, the average level is very similar. For a representative subset of points, which we quantify visually, we find an average absolute difference only of 0.25pp.

D 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_{nn}, 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. 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.

The main model is estimated at the quarterly frequency. Estimating the model using monthly frequency data proved to be computationally too costly relative to the 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

Figure 1: Comparison with Kitsul and Wright's estimates of N-probabilities at 5y and 10y

(a) Kitsul-Wright deflation

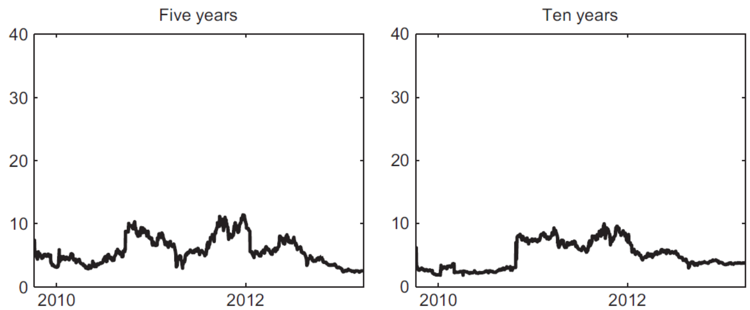

(b) Ours deflation

(c) Kitsul-Wright inflation above 4%

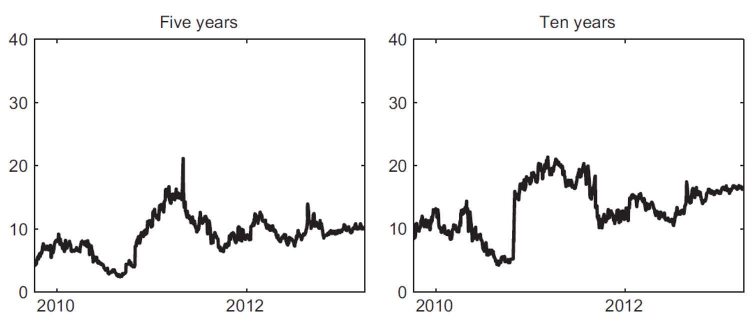

(d) Ours inflation above 4%

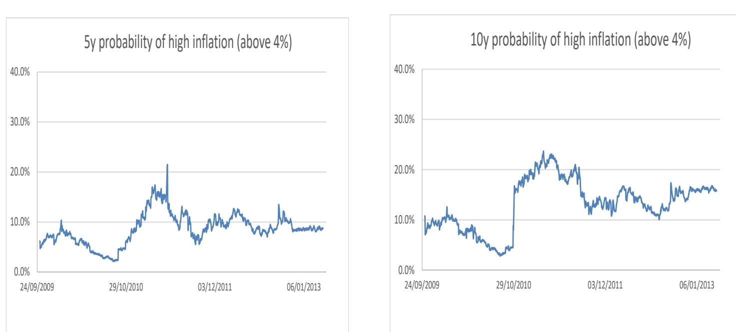

Note: Panels c) and d) are copied and pasted from Kitsul and Wright (2013). Panels (c) and (d) are daily estimates over the same sample using our procedures.

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.

Finally, in order to calculate the model average inflation for 5y and 10y, we 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%.

D.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 illustration, figure 2 present data and model distributions for the first quarter of 2021.

We next compare model fit of our main model (Model 101) to the fit of an alternative model (Model 1) for which all parameters vary freely over time. The top panel of figure 3 plots the root mean squared error of the model and model R^2 over time as well as average model fit. 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.

The bottom panel reports results for the EZ. The pattern is similar to the US, with overall fit being comparable for both models, though again a little better for the flexible

Figure 2: US model-implied densities and actual data in 2021Q1

(a) 5-year cumulative distribution

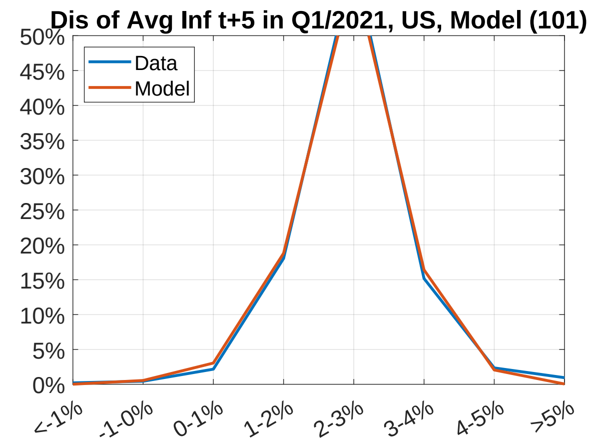

(b) 10-year cumulative distribution

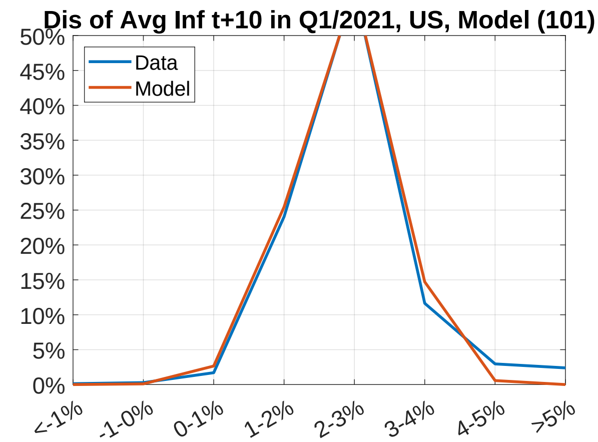

(c) One-year forward distribution

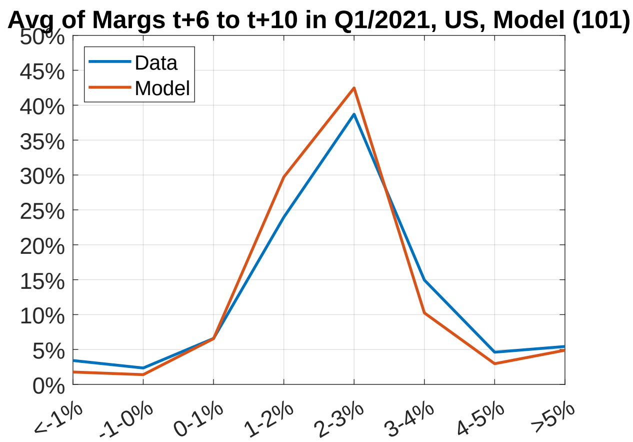

model, as expected. Time variation in model fit is also similar, with the exception of the early days of the pandemic, during which model 101 underperforms by a little more.

Figure 3: Model fit

United States

(a) Baseline model

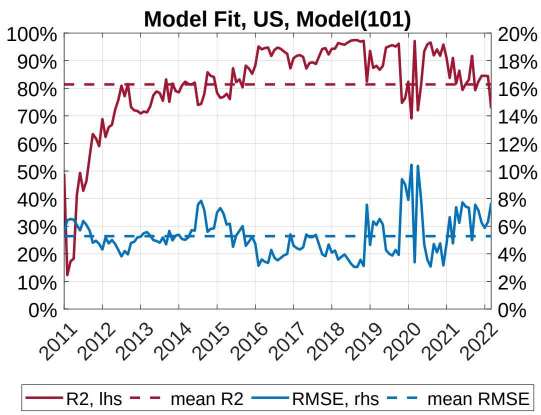

(b) Model with 6 time-varying parameters

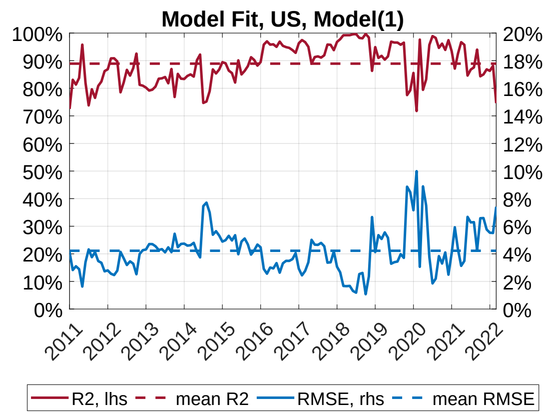

Eurozone

(c) Baseline model

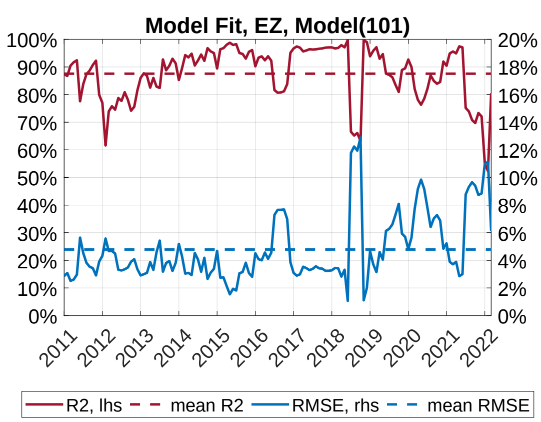

(d) Model with 6 time-varying parameters

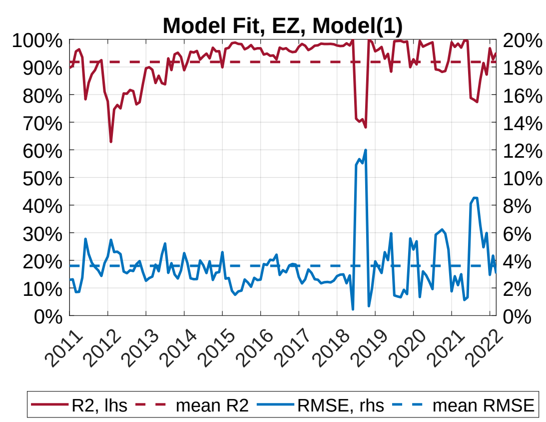

We also explore another model in which the parameters that are held fixed in model 1 are allowed to vary at the annual frequency, rather than at the monthly frequency, which is what we assume in the fully flexible model. This approach, which we refer to as model 101A, results in a substantial increase in parameters relative to model 101,

Figure 4: Model fit compared to model with slow-moving mean reversion probability

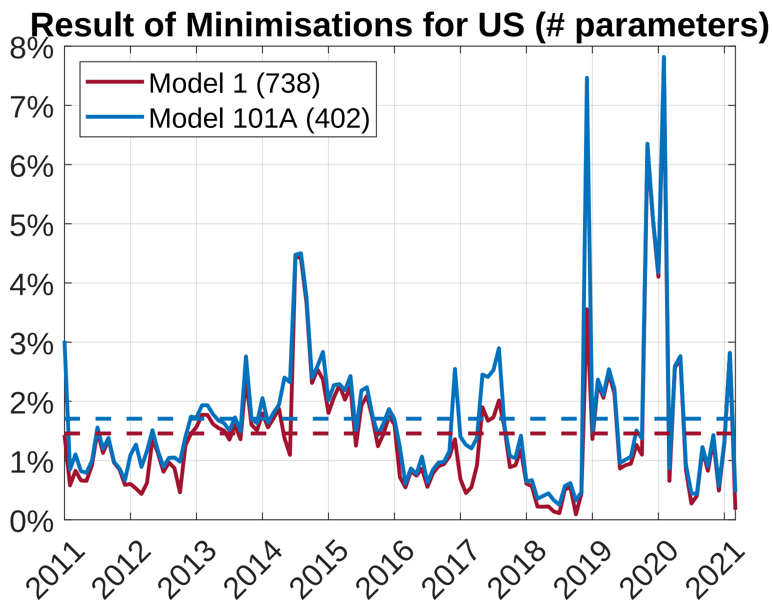

our main model, and it also improves fit a bit, but it has the same feature as the monthly model, which is an inability to clearly identify long-term trends in volatility through the probability of local changes. This is because this model allows for slow-moving changes in mean reversion, which also affects long-run volatility, itself slow-moving. Figure 4 shows model fit compared to the fully flexible monthly model.

D.2 More alternative models

We considered several other candidate models. These models either had two 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, with no clear improvement.

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 a third 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 a fourth 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 estimated a model in which parameters varied at different frequencies. The parameters assumed to be constant in our main model now 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.

E 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.

E.1 Identifying disasters: baseline

Starting with one country's inflation series, we date sequential peaks and troughs by looking for local maxima and minima in 5-year rolling windows: if the midpoint is lower (higher) than all other values in the window, that point is classified as a trough (peak). More formally:

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

All observations after some preceding trough/peak up to (and including) the next local extremum are classified as one cycle. These cycles, often spanning several years, are the unit for evaluating whether there is a disaster. The inflation of an entire cycle $C = \{t_C, t_C + 1, \ldots, t_C + T_C\}$ 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.

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

Figure 5: 150 years of inflation disasters

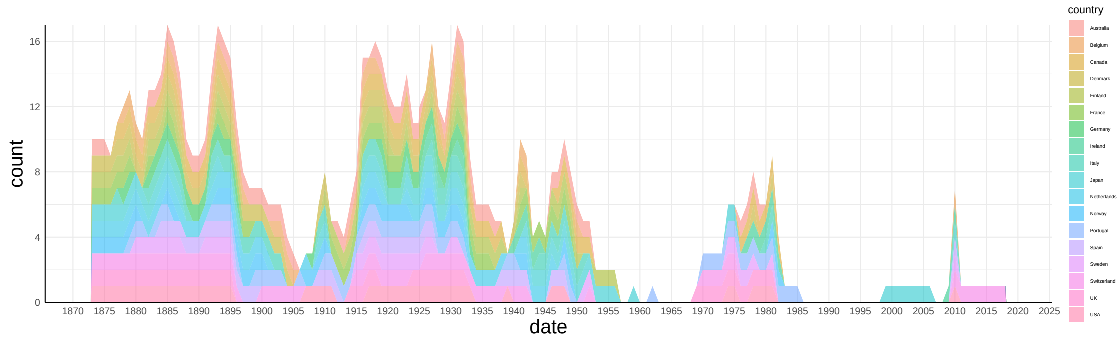

filter that isolates fluctuations of frequency lower than 20 years. If inflation is sufficiently away from the target we call this a disaster. More specifically, a cycle is classified to be in a disaster state π_d if this inflation value deviates from some target by some threshold. For the baseline results, the target is given by applying a 20 year-Butterworth square-wave highpass filter on the inflation series.

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.

Finally, 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).

E.2 Estimating the Pareto distribution

Pooling positive and negative joint inflation-consumption disasters, figure 6 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

Figure 6: Inverse fall in GDP (z) during an inflation and consumption disaster

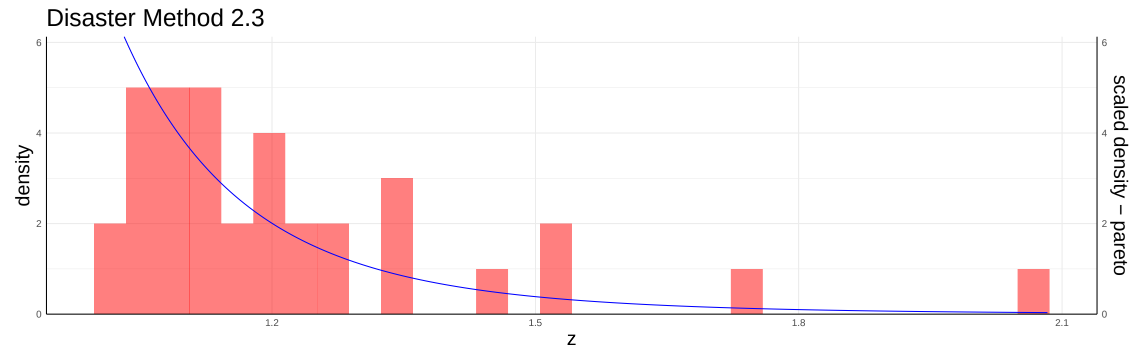

results from fitting the Pareto distribution, but pooled for both high and low inflation disasters (so imposing that z^h and z^l have the same distribution). The resulting estimates are $\alpha = 6.38$ and $z_0 = 1.03$. For comparison, Barro and Liao (2021) report α's in the range of 6 to 8, and set $z_0 = 1.03$.

Separating high and low inflation disasters, the estimates are reported in the main text (and further reported below in table 2).

E.3 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 sub-methods 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 disasters; 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.

E.4 Results under alternatives

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} x {T1, T2, T3}. Then, the table also reports estimated parameters of a Pareto fit on the (transformed) changes in output $z = 1/(1+g)$ during joint disasters. Table 2 presents conditional probabilities where a distinction was made between high and low inflation disasters.

Across the 6 possible methods that result from combining these, the results are quite similar. Recall that the baseline is in the last column, method C2.T3.

E.5 Risk premia

As discussed in the text, following Gabaix (2012), Barro and Liao (2021), we use an Epstein-Zin model for marginal utility, with a relative risk aversion coefficient of 3. We then calculate risk premia defined as $q(\pi + rp) = p(\pi)$: the increase in inflation to equate risk-adjusted and actual probabilities (see main text). Figure 7 plots the resulting estimates of

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

Figure 7: US inflation risk premia

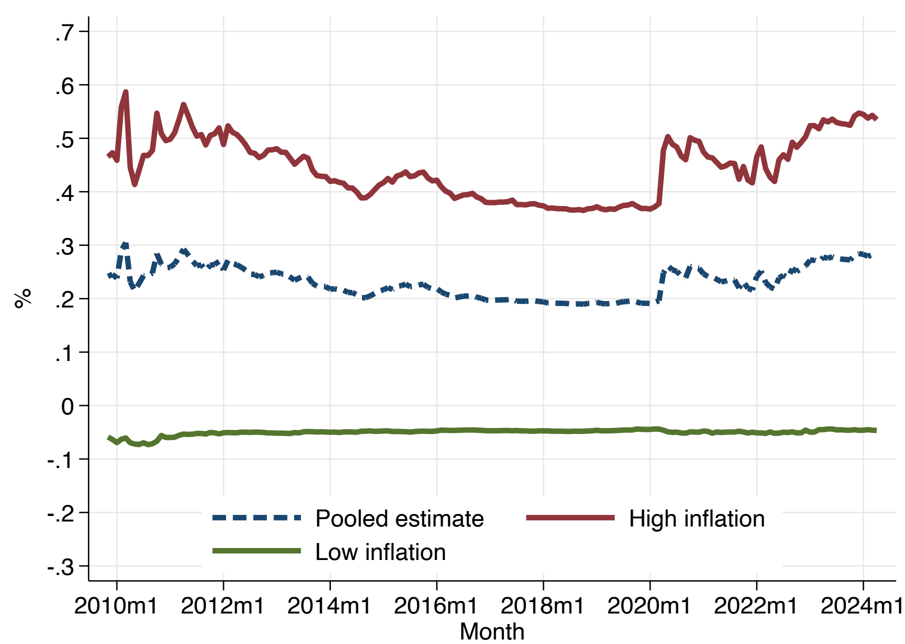

the 10-year US inflation risk premia that come from this procedure.

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
High-inflation disasters only						
unconditional probability of a high-inflation disaster	13.3%	12.9%	6.3%	12.4%	13.1%	5.7%
probability of output disaster conditional on low-inflation disaster	20.3%	23.0%	30.3%	22.7%	25.1%	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
Low-inflation disasters only						
unconditional probability of a low-inflation disaster	10.5%	10.2%	7.8%	10.5%	10.8%	8.6%
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

E.6 Deflation probabilities using a pooled risk factor

Figure 8 shows a version of the deflation probabilities but with a risk factor that uses the pooled estimates from the previous section. Note that the conclusions on the trends in deflation probabilities, their comparison with previous estimates in the literature, and the impact of monetary policy remain the same.

Figure 8: Probabilities of a deflation disaster with a common risk factor

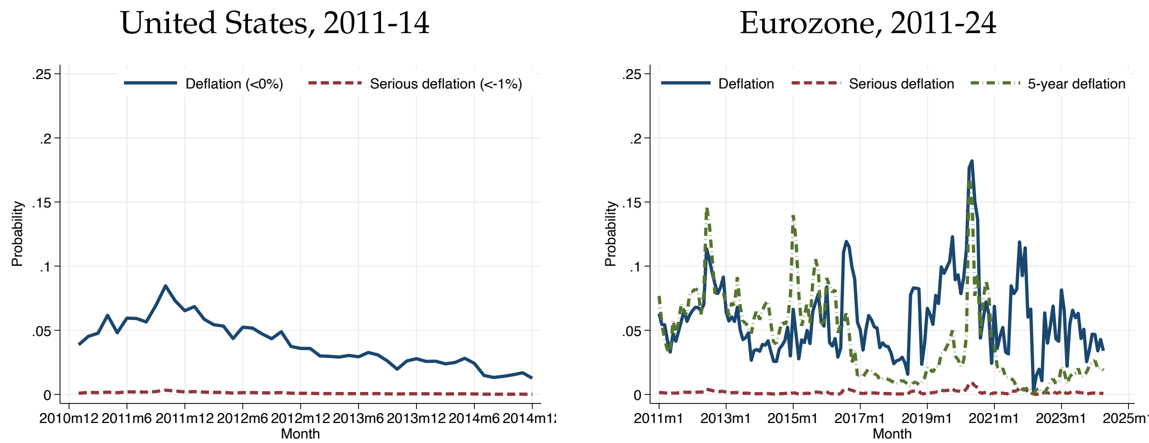

Note: The figure plots 5y5y deflation and extreme deflation probabilities and inflation and risk-adjusted 5y deflation probabilities. However, unlike in the main text, we use a pooled estimate for the inflation risk factor.

References

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.

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.

Gabaix, Xavier. 2012. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." *Quarterly Journal of Economics*, 127(2): 645–700.

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.

Kitsul, Yuriy, and Jonathan H. Wright. 2013. "The Economics of Options-Implied Inflation Probability Density Functions." *Journal of Financial Economics*, 110(3): 696–711.

Rubinstein, Mark. 1991. "Pay Now, Choose Later." *Risk*, 4: 13.

Shimko, David. 1993. "Bounds of Probability." *Risk*, 6: 33–37.