Online Appendix to "How likely is an inflation disaster?"

Jens Hilscher — UC Davis

Alon Raviv — Bar-Ilan University

Ricardo Reis — LSE

April 2022

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

A Constructing the marginal distributions of inflation

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

A.1 Data pre-cleaning

Before starting the construction of the inflation distributions, we pre-process the data. The raw data includes both data errors as well as data points that are based on trades at different times of the day. This lack of simultaneity means that option prices may not pass some basic screens. We only use data if it passes the following requirements: (1) cap and floor premia are monotonic in the strike price, (2) cap and floor premia increase monotonically with maturity, (3) butterfly spreads, which represent one way of constructing nominal Arrow-Debreu security payoffs, have positive prices, and (4) the put-call parity implied real rates are consistent across strike prices.

A.2 Implied volatility smoothing

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

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

A.3 Strike prices

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

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

A.4 Constructing distributions

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

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

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

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

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

B Model of inflation risk

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

B.1 Identifying disasters: baseline

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

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

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

All observations after some preceding trough/peak up to (and including) the next local extremum are classified as one cycle (method C2, see below). These cycles, often spanning several years, are the unit for evaluating whether there is a disaster. The inflation of an entire cycle $C = \{t_C, t_C + 1, \ldots, t_C + T_C\}$ 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.

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

B.2 Alternatives to identifying disasters

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

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

B.3 Results under alternatives, for pooled sample

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

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

Table 1: Unconditional and conditional probabilities, Pareto fits

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

B.4 Results under alternatives, separating high and low

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

C Horizon factor model

The Markov model that we use to model inflation dynamics has six parameters: symmetric movements in the middle of the distribution, p_{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.

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

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

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

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

C.1 Model fit

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

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

Figure 1: 2021 Q1, 5-year cumulative distribution

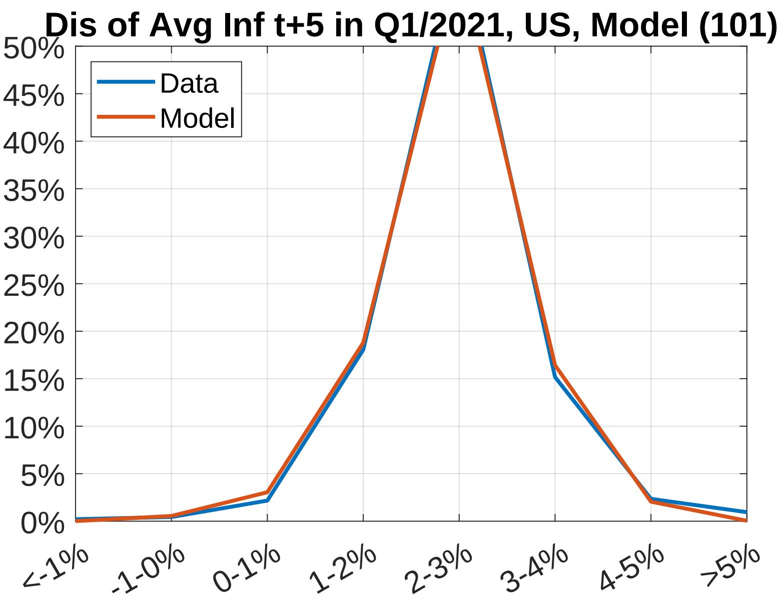

need to choose a value for average inflation in the high and low inflation disaster states (below -1% and above 5%); we set these equal to -2% and 6%.

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

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

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

Figure 2: 2021 Q1, 10-year cumulative distribution

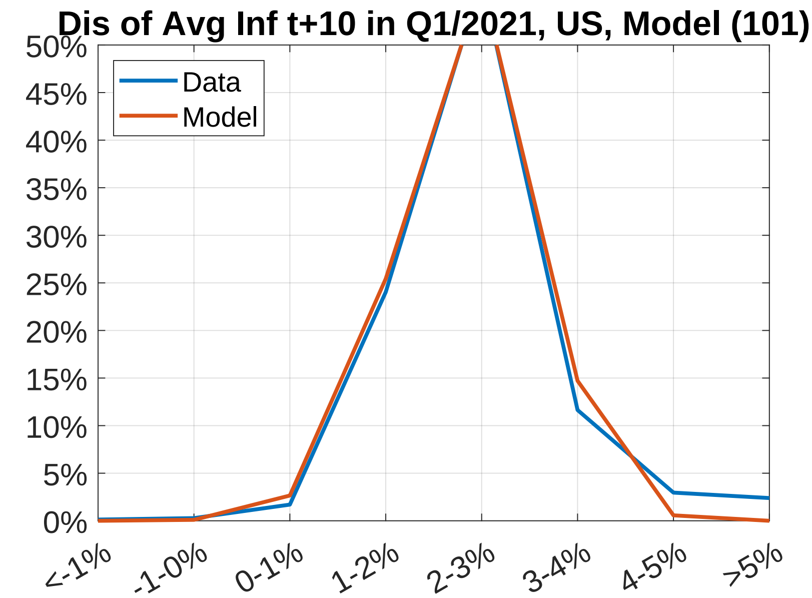

Figure 3: 2021 Q1, one-year forward distribution

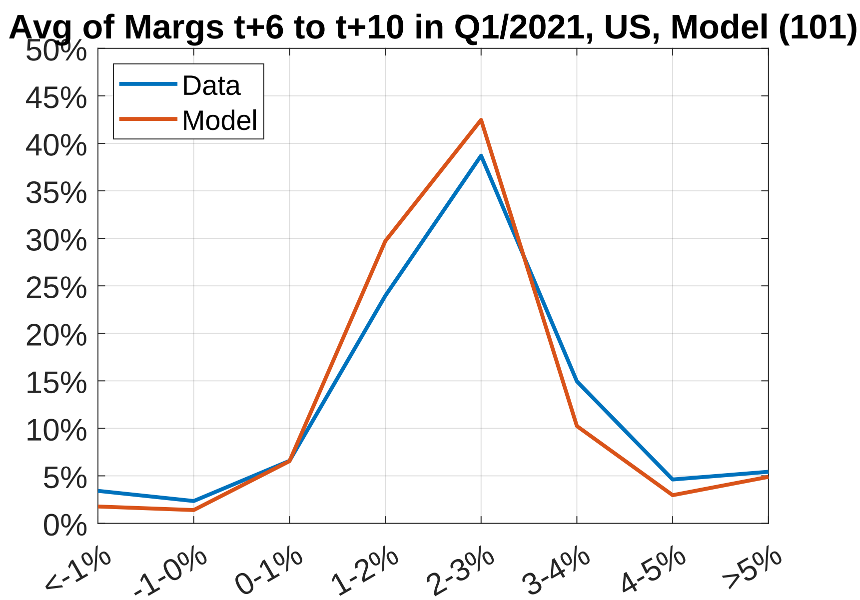

Figure 4: Model fit for US monthly and quarterly models

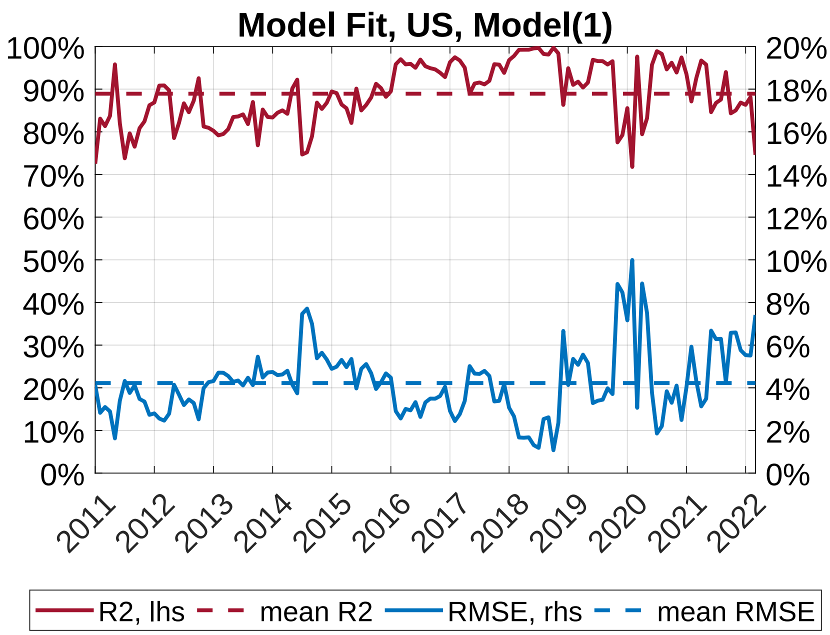

Figure 5: Model fit for US monthly and quarterly models

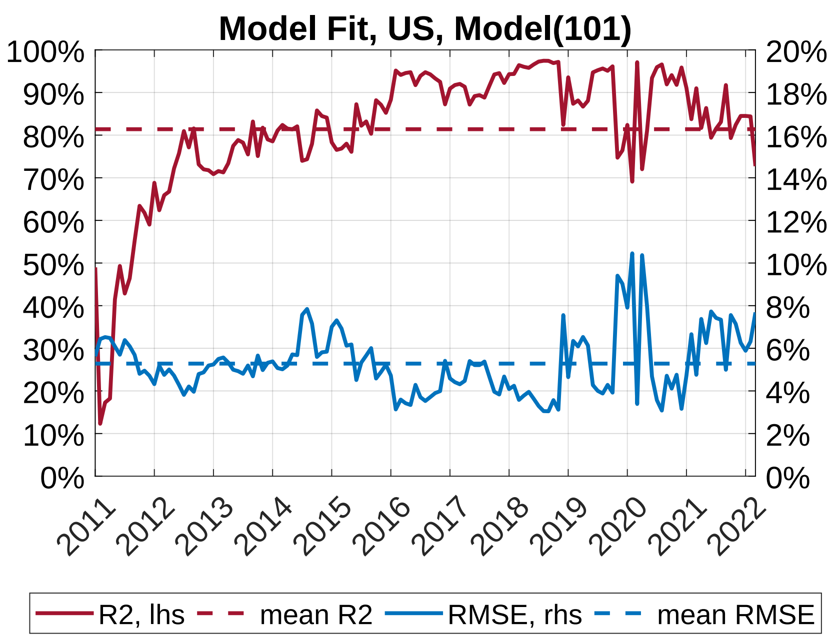

Figure 6: Model fit for EZ monthly and quarterly models

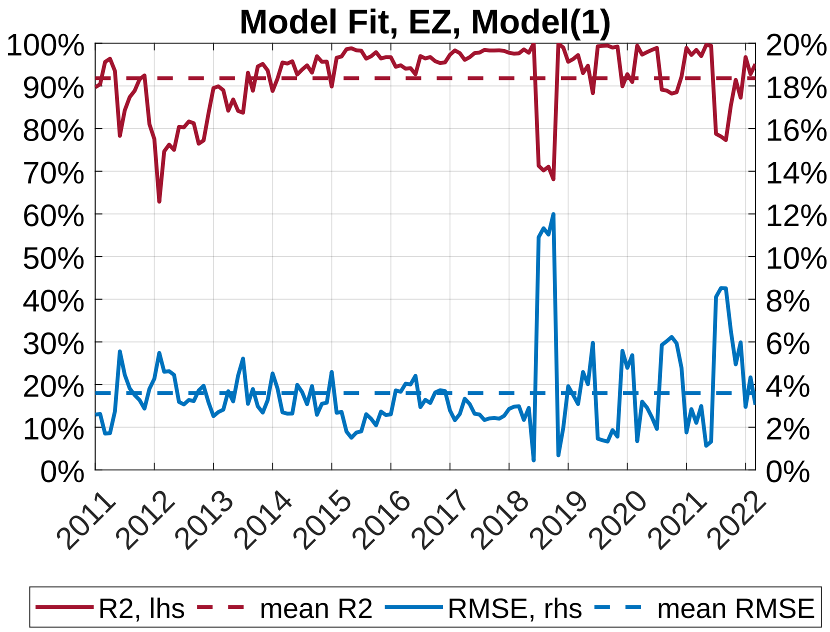

in model fit is also similar, with the exception of the early days of the pandemic, during which model 101 underperforms by a little more.

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

C.2 Model parameters

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

Figure 7: Model fit for EZ monthly and quarterly models

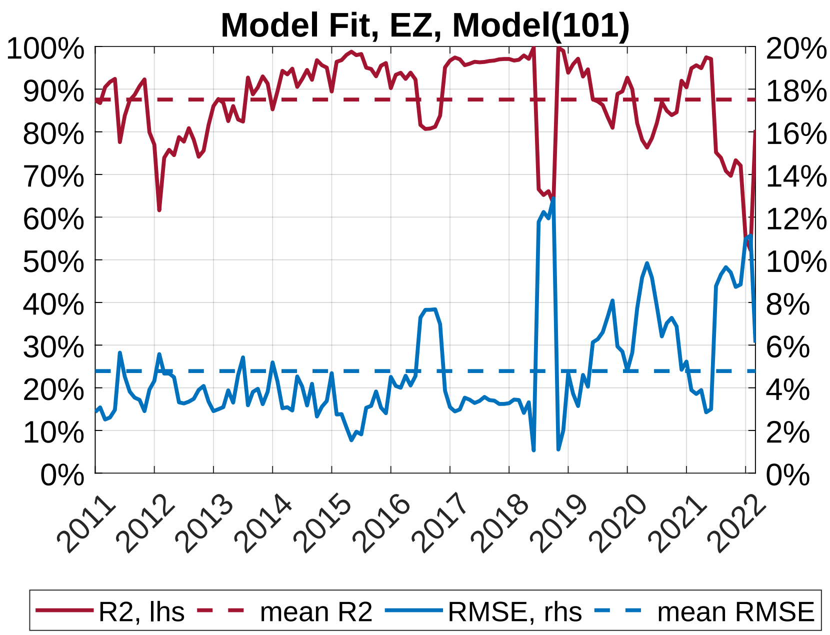

Figure 8: Model fit compared to model with slow-moving mean reversion probability

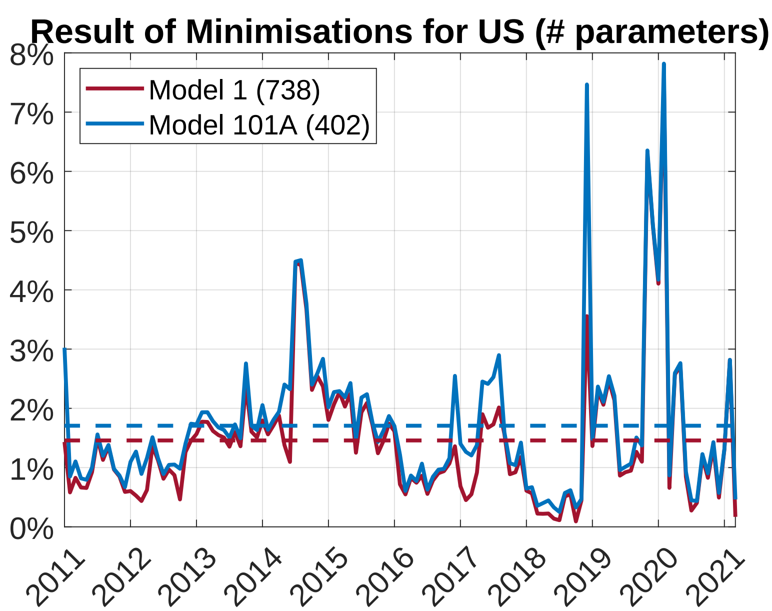

Figure 9: Model parameters: US

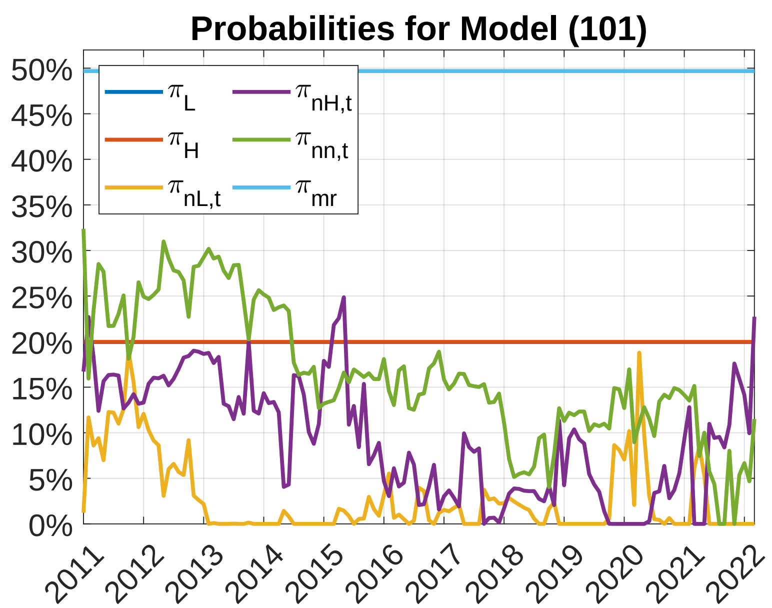

Figure 10: Model parameters: EZ

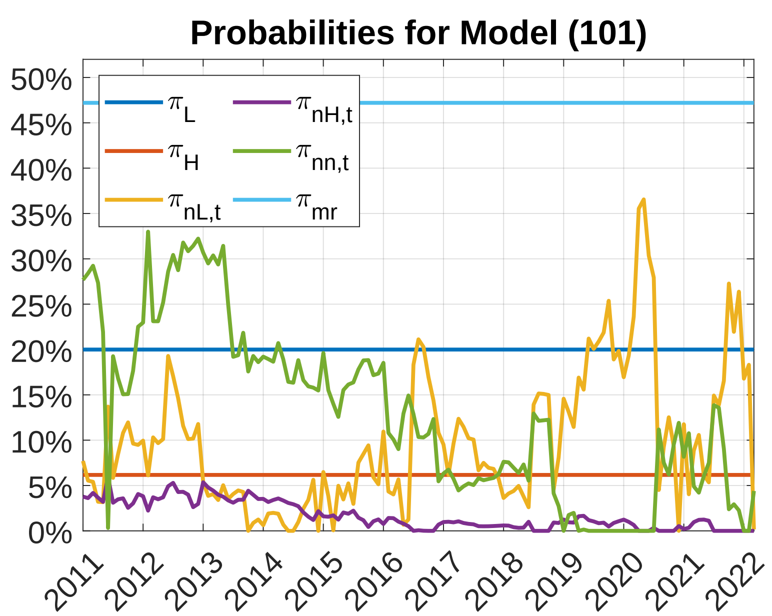

C.3 Other models

We considered several other candidate models. These models either had 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.

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

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

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

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

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

References

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–866.

Birru, Justin, and Stephen Figlewski. 2012. "Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008." *Journal of Financial Markets*, 15(2): 151–180.

Black, Fischer, and Myron S Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81(3): 637–54.

Hagan, Patrick S., Deep Kumar, Andrew S. Lesniewski, and Diana W. Woodward. 2002. "Managing Smile Risk." *Wilmott*.

Hilscher, Jens, Alon Raviv, and Ricardo Reis. 2022. "Inflating Away the Public Debt? An Empirical Assessment." *Review of Financial Studies*, 35(3): 1553–1595.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2016. "Macrofinancial History and the New Business Cycle Facts." In *NBER Macroeconomics Annual 2016*. Vol. 31, , ed. Martin Eichenbaum and Jonathan A. Parker. University of Chicago Press.

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

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