Online Appendix to "Inflating Away the Public Debt? An Empirical Assessment"

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

November 2020
For Online Publication

A Proof of proposition 1

While the bulk of the proof is already in section 2, here we fill in some missing steps. First, adding and subtracting $Q_t^1 W_{t+1}$ on the right-hand side of equation (2), and using equation (1) to replace out W_{t+1}, a few steps of algebra deliver the law of motion for the market value of government debt:

$$W_t = Q_t^1 W_{t+1} + s_t + x_{t+1} \tag{A1}$$

where the *revaluation* term x_{t+1} is equal to:

$$\sum_{j=0}^{\infty} \left(Q_t^{j+1} - Q_t^1 Q_{t+1}^j\right) K_{t+1}^j + \sum_{j=0}^{\infty} \left(H_t^{j+1} - H_t^1 H_{t+1}^j\right) \frac{B_{t+1}^j}{P_t} + \sum_{j=0}^{\infty} \frac{H_{t+1}^j B_{t+1}^j}{P_{t+1}} \left(\frac{P_{t+1} H_t^1}{P_t} - Q_t^1\right). \tag{A2}$$

Iterating this equation forward, from date 0 to date $t+1$, delivers equation (3) in the text.

Dividing both sides of the law of motion for W_t by Q_t^1, multiplying by $m_{t,t+1}$ and taking expectations gives:

$$W_t = \mathbb{E}_t(m_{t,t+1} W_{t+1}) + s_t + \mathbb{E}_t\left(\frac{m_{t,t+1} x_{t+1}}{Q_t^1}\right). \tag{A3}$$

For now, assume that the last term on the right-hand side is zero. We will show it shortly. Multiply both sides of (A3) by $m_{0,t}$ and take expectations as of date 0, so that using the law of iterated expectations you get the recursion:

$$\mathbb{E}_0(m_{0,t}W_t) = \mathbb{E}_0(m_{0,t+1}W_{t+1}) + \mathbb{E}_0(m_{0,t}s_t). \tag{A4}$$

Iterate this forward from date 0 to date T, and take the limit as T goes to infinity. With the no-Ponzi scheme condition $\lim_{T\to\infty} \mathbb{E}_0\left(m_{0,T}W_T\right) = 0$, you get the result in expression (5):

$$W_0 = \mathbb{E}_0\left[\sum_{t=0}^{\infty} m_{0,t}s_t\right]. \tag{A5}$$

Finally, for the first equality in expression (5), replace out the bond prices from equation (1) using equation (4).

The missing step was to show that $\mathbb{E}_t(m_{t,t+1}x_{t+1}/Q_t^1) = 0$ for all t. Consider the first element of x_{t+1} and take the following steps:

$$\begin{aligned}
&\mathbb{E}_t\left[m_{t,t+1}\sum_{j=0}^{\infty}\left(\frac{Q_t^{j+1}}{Q_t^1} - Q_{t+1}^j\right)K_{t+1}^j\right] = \\
&\sum_{j=0}^{\infty} K_{t+1}^j\left[Q_t^{j+1}\,\mathbb{E}_t\left(\frac{m_{t,t+1}}{Q_t^1}\right) - \mathbb{E}_t\left(m_{t,t+1}Q_{t+1}^j\right)\right] = \\
&\sum_{j=0}^{\infty} K_{t+1}^j\left[Q_t^{j+1} - \mathbb{E}_t\left(m_{t,t+1}\,\mathbb{E}_{t+1}\,m_{t+1,t+1+j}\right)\right] = \\
&\sum_{j=0}^{\infty} K_{t+1}^j\left[Q_t^{j+1} - \mathbb{E}_t\left(m_{t,t+1+j}\right)\right] = 0.
\end{aligned} \tag{A6}$$

The first equality comes from replacing the expectation of the sum by the sum of the expectations and from taking the prices known at date t out of the expectation; the second from using the result in equation (4) twice; the third from using the law of iterated expectations; and the fourth from using equation (4). Identical steps show that the other terms in $\mathbb{E}_t(m_{t,t+1}x_{t+1}/Q_t^1)$ are also equal to zero.

The final part of the proof to clarify are expressions (6) and (7). First note that in principle, $m_{0,t}$ can depend on many random variables. However, all we need to evaluate is $\mathbb{E}\left(m_{0,t}/\pi_{0,t}\right)$. Therefore, only the dependence of $m_{0,t}$ on $\pi_{0,t}$ will lead to a non-zero term once the expectation of the product of the discount factor and inverse inflation is evaluated. Therefore, we can write $\mathbb{E}\left(m(\pi_{0,t})/\pi_{0,t}\right)$.

Second, let $\hat{f}(.)$ be the probability density function for inflation. Then, we can re-write the expression as: $\int \left(\hat{f}(\pi_{0,t})m(\pi_{0,t})/\pi_{0,t}\right) d\pi_{0,t}$. Define the inverse of the risk-free rate, from the perspective of date 0, as $R_t^{-1} = \int \hat{f}(\pi_{0,t})m(\pi_{0,t})d\pi_{0,t}$. From the definition of the stochastic discount factor, this is the price of a bond that would pay one dollar for sure in t periods, regardless of the realization of inflation. Then, define the risk-adjusted density for inflation as: $f(\pi_{0,t}) = \hat{f}(\pi_{0,t})m(\pi_{0,t})R_t$. Using these results it follows that:

$$\mathbb{E}\left(m(\pi_{0,t})/\pi_{0,t}\right) = \int \left(\frac{\hat{f}(\pi_{0,t})m(\pi_{0,t})}{\pi_{0,t}}\right) d\pi_{0,t} = R_t^{-1} \int \left(\frac{f(\pi_{0,t})}{\pi_{0,t}}\right) d\pi_{0,t}. \quad \text{(A7)}$$

This confirms expression (7).

B Debt holdings

We construct monthly maturity structures, that is B_0^t, for four groups of investors: (1) private total, which is publicly held debt, excluding the Federal Reserve and state and local government holdings, (2) foreign, (3) domestic, which is private minus foreign, and (4) the central bank. The data sources are the CRSP U.S. Treasury database, the Treasury Bulletin, the "Foreign Portfolio Holdings of U.S. Securities" report available from the U.S. Treasury, and the System Open Market Account (SOMA) holdings available from the FRBNY. The numbers for face value of total outstanding debt for different categories in section 3.1 come from the Monthly Statement of Public Debt, the Treasury Bulletin (OFS-2, FD-3), and the SOMA. Data are all for end of December 2009 to 2017 for the United States.

We construct holdings of notes and bonds as follows. For (1), the data is available from CRSP, and we subtract out state and local government holdings assuming they have the same maturity structure. Detailed information for (2) is available for June 2015 and we assume proportionate growth of foreign holdings to construct December 2015 numbers. (3) is the difference of (1) and (2). For (4) we use the security level data available from the FRBNY. We assume that all coupon and principal payments mature in the middle of each month.

CRSP does not have data on Treasury bills. We use the issues of the Treasury bulletin to obtain information on bills and follow the same steps as we did above for notes and bonds. In particular, we construct holding of T-Bills as follows. For (1) we subtract Fed holdings from Treasury Bulletin Table FD-2 T-Bill holdings and then subtract out the proportionate amount of state and local government holdings; for (2) and (3) we do the same as for notes and bonds; for (4) we use the SOMA holdings.

Aside from calculating debt holdings for more categories of investors, our method for constructing the maturity structure of (1) has the following differences relative to Hall and Sargent (2011). First, we construct a monthly term structure and assume that promised payments in a month are paid in the middle of the month (instead of using an annual frequency). Second, we exclude state and local government holdings. Third, we base the T-Bill holdings on Table FD-2 of the Treasury bulletin and Fed holdings (rather than Table FD-5 of the Treasury bulletin).

Real interest rates are constructed according to equation (8) in the paper and the spot curve (nominal zero coupon yield curve) provided by Gürkaynak, Sack and Wright (2007). The harmonic means of inflation are calculated using simulated 30-year distributions based on our fitted restricted model. As a check, we also separately calculate the market value of the debt using the real term structure from Gürkaynak, Sack and Wright (2010) for real interest rates. These only extend up to 20 years. To construct real spot rates for longer maturities we assume that forward rates for years 21 to 30 are equal to the average forward rate for years 18 to 20. Our results are robust to this alternative.

C Estimating the marginal distributions for inflation

We estimate inflation distributions from data on zero coupon and year-on-year inflation caps and floors, which we collect from Bloomberg, as do Fleckenstein, Longstaff and Lustig (2017). Kitsul and Wright (2013) use data provided by an interdealer organization, whereas we use the raw reported numbers.

C.1 Data cleaning

In order to construct distributions, the option data require considerable amount of cleaning. We provide Aside from measurement error, we face the difficulty that the options in general are not traded simultaneously, resulting in option pricing functions that are not always well behaved. To screen out such data, we first drop option prices from the data if they contain simple arbitrage opportunities: (i) if the call (put) premium does not monotonically decrease (increase) in the strike price, (ii) if the call option premium does not increase monotonically with maturity, and (iii) if butterfly spreads that correspond to Arrow-Debreu securities do not have positive prices. Next, before we take differences of the data, we transform prices to implied volatility space: for each maturity, we calculate Black and Scholes (1973) implied volatilities for all strike prices, smooth them, convert back to option prices, and construct distributions using the Breeden-Litzenberg formula.[1]

C.2 Zero-coupon inflation options

The basic methodology for construction of the distributions from zero coupon inflation options is fairly standard and similar to Kitsul and Wright (2013) and Fleckenstein, Longstaff and Lustig (2017). Still, we take several steps in cleaning the data that this section of the appendix clarifies.

The zero coupon inflation call options are traded with strike prices between 1% and 6%, in 0.5% increments, and expiration dates ranging from 1 to 10 years as well as 12 and 15 years, although the data for the 2 and 9 year maturities are of generally lower quality. The zero coupon put options are available for strike prices between -2% and 3% and identical maturities as the caps. For the overlapping range of strike prices we use both option prices to reduce measurement error.

Breeden and Litzenberger (1978) suggest using a butterfly trading strategy to construct Arrow-Debreu securities, claims that pay one unit of currency if at some specific time in the future the underlying asset price is equal to a specific value and zero otherwise. While this method provides a good first approximation to risk neutral probabilities, it does not adjust for irregular options prices, due to, for example, non-synchronous trading (Bahra, 1997). Therefore, while we check that all such prices are positive, we must smooth the data otherwise the results are very inaccurate.

To overcome the drawbacks of the unadjusted butterfly strategy, we do the following. First, we drop data that represent simple arbitrage opportunities (discussed in the text). We next calculate Black-Scholes implied volatilities, and, following Shimko (1993) and Campa, Chang and Reider (1998), for each set of options at any expiry date, we fit a four factor stochastic volatility model, the SABR model, which was developed by Hagan et al. (2002). This calibration minimizes the norm of the difference between the observed data and the candidate SABR function, resulting in a vector of optimal parameters for the SABR model which is used to create a volatility curve. We constrain the estimated implied volatility function to ensure that the smoothing does not re-introduce arbitrage opportunities. This method reduces the weight of irregular data, while preserving its overall form. We convert back to option prices and construct risk neutral distributions.

[1]Note that the implied volatility smoothing using Black and Scholes does *not* mean that we are using this model to price the options. It is simply a non-linear transformation of the pricing function.

To be clear, this method does not assume that we can use the Black-Scholes formula for pricing. Instead, it is simply used as a nonlinear transformation on which smoothing is performed.

C.3 Year-on-year inflation options

The construction of distributions from year-on-year options requires a bootstrapping method where cap and floor contracts, which are portfolios of annual caplets and floorlets, are unbundled to recover prices of the underlying options. First, we use a bootstrapping procedure to extract the caplet and the floorlet prices from the cap and floor prices respectively. Second, when calculating the option's implied volatility we use the Rubinstein (1991) transformation which enables us to price the option as a plain vanilla option with a time to maturity equal to the option tenor between inflation reset times, discounting back using the real interest rate, calculated assuming that put-call parity holds for these options (Birru and Figlewski, 2012).

For each horizon, for the smallest (largest) bin we report the risk-adjusted probability of inflation lying in or below (above) that bin. If there is considerable mass in the last bin for year-on-year inflation, we use our smoothed estimates to project additional bins.

C.4 Comparing option prices at different dates

For each year, we inspect option prices at the end of December and start of January (our data of interest) and select out of the last five trading days of the year and the first five trading days of the new year, the day that yields the maximum number of option prices that do not violate the threshold criteria, described above and in the text. For end of 2017 data, this was January 9, 2018. Each year, we also look at distributions for several days before and after. All of them look almost identical to the ones we then end up using.

D Estimating the joint distributions for inflation

The marginal distributions for inflation are enough to evaluate the formula in proposition 1. Yet, to calculate inflation paths or counterfactuals, we need joint distributions.

D.1 Proof of proposition 2

Using $F(.)$ to denote the cumulative density function, we have data for one-year inflation $F(\ln \pi_{t+j-1,t+j})$ for $j = 1...J$ where $J = 10$ years, and for cumulative inflation $F(\ln \pi_{t,t+j})$. The data comes in N bins expressed as ranges for inflation.

Sklar's theorem states that there exists a function $c : [0, 1]^J \to [0, 1]$ such that:

$$F(\ln \pi_{t,t+1}, ..., \ln \pi_{t+J-1,t+J}) = c\left(F(\ln \pi_{t,t+1}), F(\ln \pi_{t+1,t+2}), ..., F(\ln \pi_{t+J-1,t+J})\right). \tag{A8}$$

In turn, it follows from the link between marginal and joint distributions and the definition of cumulative inflation that:

$$F(\ln \pi_{t,t+j}) = \int_{\Pi} F(\ln \pi_{t,t+1}, \ln \pi_{t+1,t+2}, ..., \ln \pi_{t+J-1,t+J}) d \ln \pi_{t,t+1} ... d \ln \pi_{t+J-1,t+J} \tag{A9}$$

where the set Π is defined as: $\left\{ \ln \pi_{t,t+j} : \ln \pi_{t,t+n} \sum_{i=1}^{j} \ln \pi_{t+i-1,t+i} \right\}$.

Combining these two results delivers the proposition in terms of cumulative distributions. We maintain the assumption throughout that all distributions are continuous. Therefore, Sklar's theorem also applies to the marginal density functions, with $C(.)$ replacing $c(.)$. For numerical purposes, it was better to work with densities rather than with cumulative distribution functions.

A final point to note is that there are N bins and so N equalities in the distribution, of which one is redundant since probabilities must add up to one. There are J maturities, but for maturity one the equality is trivial. Therefore, in total there are $(N-1)(J-1)$ conditions.

D.2 Parametric copula and method of moments

We use the parametric normal copula, whose formula is:

$$\hat{C}\left(f(\ln \pi_{t,t+1}), ... f(\ln \pi_{t+J-1,t+J}\right) =$$
$$\left(\frac{1}{\det R}\right) \exp\left(-\frac{1}{2} \begin{pmatrix} \Phi^{-1}(f(\ln \pi_{t,t+1})) \\ \vdots \\ \Phi^{-1}(f(\ln \pi_{t+J-1,t+J})) \end{pmatrix} (\rho^{-1} - \mathbf{I}_J) \begin{pmatrix} \Phi^{-1}(f(\ln \pi_{t,t+1})) \\ \vdots \\ \Phi^{-1}(f(\ln \pi_{t+J-1,t+J})) \end{pmatrix} \right) \tag{A10}$$

where $\Phi(.)^{-1}$ is the inverse of the standard normal cdf, and ρ is a correlation matrix of dimension J.

The matrix ρ would only exactly equal the correlation matrix of the variables in the joint distribution if the marginal distributions happened to be normal. Yet, by drawing from the joint distribution using the formula above, and calculating correlation coefficients across many draws, we found that the difference between the actual correlation matrix and ρ was almost always less than 0.01.

To find the unrestricted estimates, we minimize over the $J(J-1)/2 = 45$ independent components of the correlation matrix ρ that lie between -1 and 1. The objective function is the equally weighted average of the $(N-1)(J-1) = 20 \times 9 = 180$ squared deviations from the moments in proposition 2. This is a difficult global minimization over a large parameter space, which we handle through a combination of global and local minimization algorithms and many repeated searches from Sobol-sampled starting points.

D.3 Inflation after 10 years

In drawing paths for inflation, for the first 10 years we use the joint distribution given by the multivariate copula. After that, we assume that inflation is a 9th order Markov process. Therefore, the distribution for inflation in year 11, conditional on inflation in years 2 to 10 is the same as the distribution for inflation in year 10, conditional on years 1 to 9. Since we have the joint distribution for inflation from year 1 to 10, it is easy to derive the conditional distribution for inflation in year 10, conditional on years 1 to 9. Thus, we have the conditional distribution for year 11, conditional on the draws so far. The same applies to year 12, and so on, all the way to 30. Note that, since we assume that the joint distribution of maturities one to ten follows a Gaussian copula, then this procedure implicitly assumes that the joint distribution in maturities one to thirty is likewise a Gaussian copula. The key restriction is that the correlation matrix of $30 \times 29/2 = 435$ elements for the 1-30 copula only have 45 independent separate elements that we estimated for the 1-10 distribution.

D.4 Restricted distribution

We can represent inflation between two successive dates (or maturities) as:

$$\ln \pi_{t,t+1} = \mathbb{E}_t(\ln \pi_{t,t+1}) + p_{t+1} + s_{t+1}, \tag{A11}$$

where p_{t+1} is a non-stationary part, and s_{t+1} a stationary one, with the two being independent and zero mean. From Wold's theorem, the stationary process is fully characterized by its covariance function $v_j = \mathbb{E}_t(s_{t+n}s_{t+n+j})$ for any arbitrary j. The definition of stationary is that this is true for any positive n.

The restriction that we impose is that the non-stationary process is a random walk, so that $\mathbb{E}_t(p_{t+n}p_{t+n+j}) = \sigma n$. This constrains the way in which non-stationarity affects the correlation matrix over time. In particular, the correlation between inflation at date $t+n$ and date $t+j$, or the (n,j) element of the matrix ρ is given by the expression:

$$\rho_{n,n+j} = \frac{v_j + n\sigma}{\sqrt{(v_0 + n\sigma)(v_0 + (n+j)\sigma)}} \tag{A12}$$

While the unrestricted ρ matrix has 45 parameters, with the random-walk constraint, there are only 10 parameters. Nine are the correlations of the stationary part $\{v_1/v_0, ..., v_9/v_0\}$, and one more is the ratio of the relative variances of the permanent and transitory components σ/v_0. We minimize the same objective function but over this smaller parameter space. Note that in this case, the ρ matrix has the easily identified form, $\rho_{n,n+j} = v_j/v_0$, which is the same whatever is n.

Table 1 reports the restricted model parameters. Importantly, note that these are not correlations;[2] instead, the numbers are parameters that enter the restricted copula model. The variance parameter in the bottom row is the ratio of the variance of the shock to the random walk component to the variance of the shock to the stationary component. We can see that, for most years, the process is close to stationary. We can translate these parameters to correlations, which are reported in 2.

D.5 Goodness of fit and restricted versus unrestricted distribution

Table 3 shows the estimated ρ for the unrestricted model in 2017. Because the estimates vary so much, they are a little hard to interpret. Therefore, in Table 4 we calculate the average (off-diagonal) correlation coefficients between short term (1 to 3 years), medium term (4 to 7 years) and long term maturities (8 to 10 years). The two key features that are in common with the restricted estimates are the ones emphasized in the text: the autocorrelation coefficients are not very high, and they do not fall with maturity. But the estimates are quite noisy in that, across contiguous maturities, they jump up and down.

We assess goodness of fit in multiple ways. First, figure 1 shows the data for the risk-

[2] Correlations cannot lie above one, of course. However, within the copula model it is possible for the rho *parameters* to lie above one. Note that when we report corresponding correlations, the numbers are all below one.

Table 1: Estimates of restricted model parameters, 2009 to 2017

Maturity	*2009*	*2010*	*2011*	*2012*	*2013*	*2014*	*2015*	*2016*	*2017*
1	-0.38	-0.04	-0.09	-0.12	-0.12	-0.35	-0.02	-0.20	0.03
2	0.03	-0.62	-0.12	-0.55	-0.14	-0.29	0.42	-0.44	-0.62
3	0.40	-0.57	0.48	-0.04	-0.23	-0.34	0.12	-0.18	-0.55
4	-0.57	-0.07	-0.06	0.11	-0.03	0.02	0.63	0.77	-0.05
5	0.48	0.50	-0.31	0.12	-0.29	-0.17	0.42	0.54	0.30
6	-0.03	-0.02	0.22	-0.27	0.39	-0.16	0.18	-0.42	0.00
7	0.13	-0.08	0.36	-0.29	-0.29	0.47	0.25	0.25	0.13
8	0.46	-0.28	0.05	0.59	0.47	0.07	0.36	0.24	0.50
9	-0.68	0.29	1.23	1.13	0.99	0.65	0.80	0.83	0.03
var	0.00	0.32	0.26	0.16	0.17	0.11	0.00	0.16	0.09

Notes: Estimated parameters from the restricted model, composed of nine autocorrelations parameters of the stationary part and the variance of the random walk process.

Table 2: Restricted model implied correlations, 2009 to 2017

Maturity	*2009*	*2010*	*2011*	*2012*	*2013*	*2014*	*2015*	*2016*	*2017*
1	-0.38	0.19	0.12	0.04	0.04	-0.20	-0.02	-0.03	0.11
2	0.03	-0.19	0.09	-0.30	0.02	-0.15	0.42	-0.21	-0.45
3	0.40	-0.15	0.46	0.09	-0.05	-0.18	0.12	-0.01	-0.38
4	-0.57	0.13	0.12	0.19	0.09	0.10	0.63	0.64	0.04
5	0.48	0.42	-0.03	0.19	-0.08	-0.04	0.42	0.46	0.30
6	-0.03	0.14	0.25	-0.07	0.35	-0.04	0.18	-0.17	0.07
7	0.13	0.11	0.31	-0.08	-0.07	0.40	0.25	0.25	0.16
8	0.46	0.02	0.15	0.45	0.37	0.12	0.35	0.24	0.42
9	-0.68	0.26	0.70	0.74	0.66	0.50	0.79	0.57	0.08

Notes: Estimated correlation coefficients from the restricted models for year-on-year inflation between a given year and that year+j, where j is the maturity in each row.

Table 3: Estimated correlation coefficients of year-on-year inflation in the joint distribution

Maturity	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
1	1.00	0.34	-0.50	0.26	-0.18	0.48	-0.13	0.58	0.76	0.54
2		1.00	0.10	-0.04	0.34	0.42	-0.11	0.38	0.21	-0.17
3			1.00	-0.42	-0.25	-0.44	0.14	0.07	-0.20	-0.31
4				1.00	-0.14	0.29	0.10	0.44	0.52	0.17
5					1.00	0.33	0.18	-0.38	-0.23	-0.28
6						1.00	0.19	0.33	0.38	0.32
7							1.00	0.28	0.35	0.53
8								1.00	0.85	0.35
9									1.00	0.63
10										1.00

Notes: Estimated correlation coefficients for year-on-year inflation between date 2017+j and 2017+l, in column j, row l.

Table 4: Average correlation coefficients of year-on-year inflation in the joint distribution

	Short term (1-3 years)	*Medium term (4-7 years)*	*Long term (8-10 years)*
Short term (1-3 years)	-0.02	0.01	0.21
Medium term (4-7 years)		0.16	0.20
Long term (8-10 years)			0.61

neutral density of cumulative inflation at maturities 2 to 10 against the predicted densities, according to the restricted and unrestricted models. Our method of moments, following proposition 2, consisted of picking parameters ρ to minimize the difference between the data in these 9 plots, and the models. Visually the fit is quite good, and the restriction seems to have almost no effect on the ability of the copula model to fit the data. The patterns the other years, not reporter here, are similar.[3]

Second, we compare the models' prediction for inflation at horizons 12 and 15 with the data for those maturities. Note that this tests not only the normal copula, but also our 9th order Markov assumption to simulate beyond 10 years. Figure 2 shows the model-implied and data distributions. The restricted model fits better.

Third, we compare the model's predicted standard deviations of risk-adjusted inflation, with those in the marginal distributions. Figure 3 shows the two models. The restricted model again does a slightly better job when applied to the horizon 12 and 15 distributions, although again it is clear that, in spite of the slightly different estimates of ρ, the two models have similar fits.

E Estimating debt debasement

We draw 500,000 samples for 40 years of inflation using our joint distribution. We convert these into 480 month histories by assuming continuous compounding and a constant inflation rate within each year. We calculate the real value of the nominal payments for each of these draws, and order them, calculating their percentiles in table 2. We repeat the calculations using the B_0^t for each group of investors instead.

For the counterfactuals, we use alternative distributions for inflation from which to take draws, recalculate the value of the debt using the formula in proposition 1, and subtract it from the original number. The alternative distributions are:

1. Permanently higher: Shift all the year-on-year distributions by the difference between the 90th and the 50th percentile at each maturity.

2. Right tail only: We draw from the baseline distribution but discard all histories in which average inflation is below its 90th percentile in any one of the years in the simulation.

[3]Graphs are available upon request.

Figure 1: Marginal distribution for risk-adjusted cumulative inflation: data and models

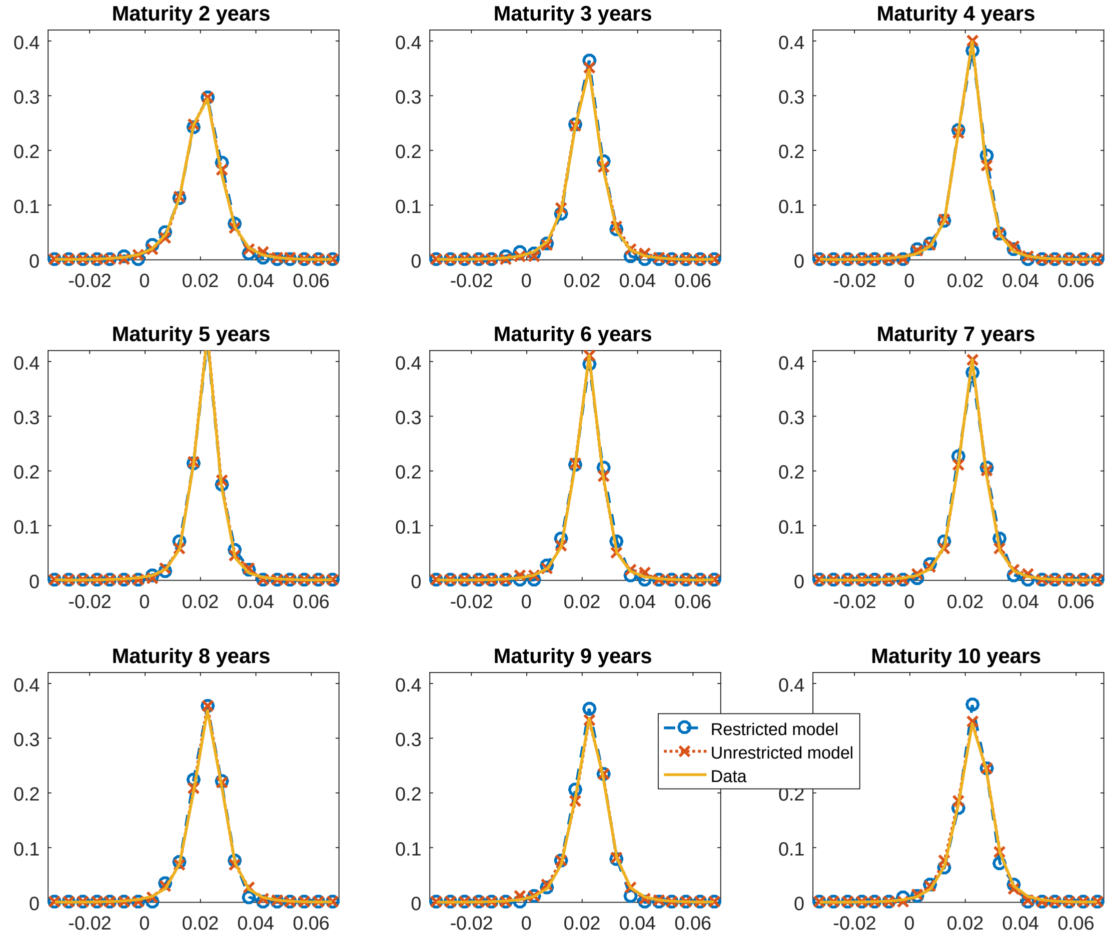

Figure 2: Distribution for $\ln \pi_{t,t+11}$ and $\ln \pi_{t,t+15}$, model and data

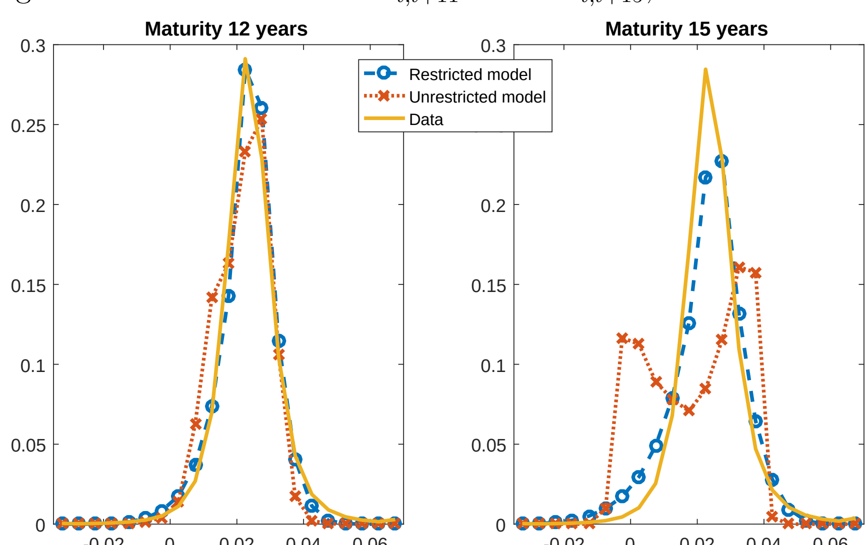

Figure 3: Standard deviation of risk-adjusted inflation at different horizons: data and models

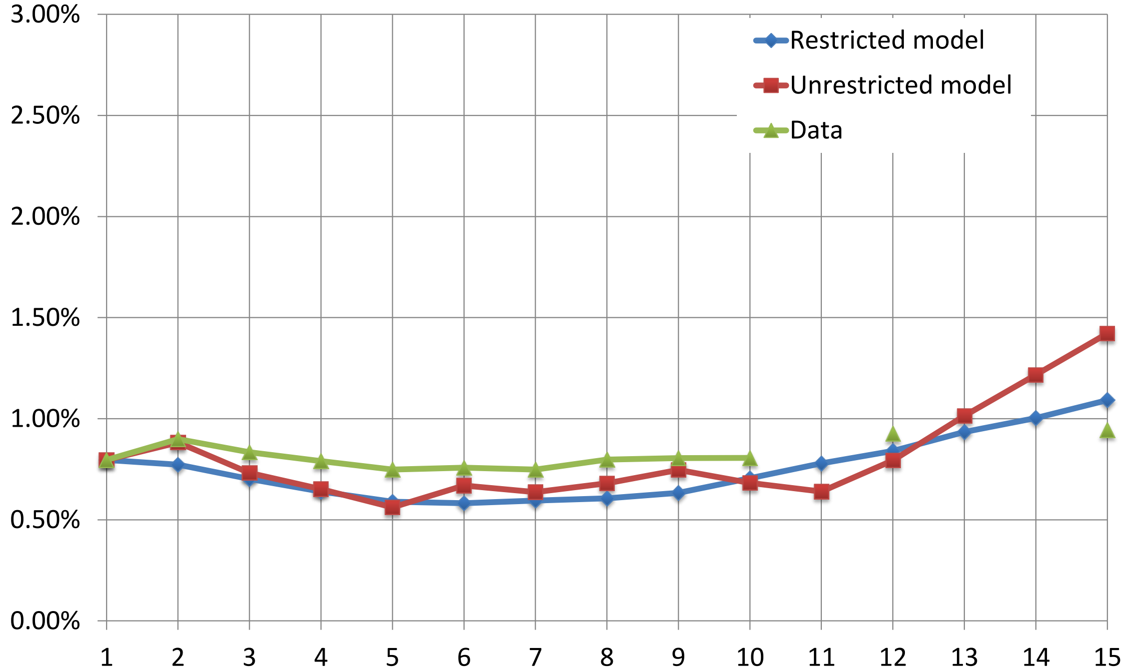

3. Higher and more variable: For each maturity we multiply baseline inflation levels by a scaling factor so that the new mean is equal to the mean in case 1. This results in more variable inflation.

4. Higher for sure: This is the same as case 1 but we now assign all the weight to the mean at each maturity.

5. Partially expected: We set inflation equal to 3% in year 1. For the following histories we use marginal distributions conditional on the year 1 realization. After 10 years we assume that year-on-year inflation is 9th order Markov.

6. Temporary increase: We shift the year-on-year distributions in the same way as case 1 but we now shift them to the 90th percentile in year 1, 80th in year 2, 70th in year 3, and 60th in year 4. There is no change in the year-on-year distributions for maturities equal to and above five years. Note that this is different from the unexpected shock in the previous case since distributions in years 2 to 4 are shifted directly instead of changing only due to the new distribution in year 1.

7. Gradual increase: The one year distribution is unchanged, the 2 year median shifts to the previous 60th percentile, 70th for year 3, 80th in year 4, and 90th for 5 years and

above.

We note that unlike the Value at Risk results in the previous section, but like any counterfactual experiment, these stress tests are potentially exposed to the Lucas critique. However, because these are shifts to risk-adjusted distributions they do incorporate agents' shifting expectations about inflation and risk. It is not the shifts per se that are subject to criticism, but only their interpretation as policy changes.

We also note that we are shifting the risk-neutral distribution of inflation, not the actual inflation target of the central bank. The link between the two depends both on the effectiveness of central bank policy as well as on changes in private assessments of risk. In the extreme case where the central bank controls the distribution of inflation, and where inflation is "pure" in the sense of Reis and Watson (2010), so that changes in inflation are independent of changes in relative prices, the two are the same.

F Risk neutral and physical probabilities

To simplify the discussion, assume that (i) there is a single bond outsanding that pays no coupon but just a principal in T periods ($B_{0,t} = 0, t \neq T$), (ii) the principal is 1 ($B_{0,T} = 1$), (iii) there is a countable set of states of the world, s drawn from a set S with physical probability distribution $p(s)$ so that $p(s) \geq 0$ for all s and $\sum_{s \in S} p(s) = 1$, (iv) the stochastic discount factor in T periods is stochastic, $m_{0,T} = m(s)$, (v) inflation is likewise random $\pi_{0,T} = \pi(s)$. The sole goal of these assumptions is to focus on scalar realizations of inflation at a fixed horizon, instead of sequences of inflation over time. They imply that the burden of the debt is:

$$W_0 = \sum_{s \in S} \frac{p(s)m(s)}{\pi(s)}.$$

With this setup, building the distribution of debasement due to inflation works as follows. First, because the goal is to capture *inflation variation*, not all variation that affects debt, one needs a change of measure. The (marginal) distribution of inflation is given by the standard formula:

$$p(\pi) = \sum_{s:\pi(s)=\pi} p(s).$$

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 $\pi(s') = \pi(s'')$. To obtain the distribution of the SDF as a function of inflation $m(\pi)$ is not as straightforward.

The reason is that for states s' and s'' such that $\pi(s') = \pi(s'')$, it may be that $m(s') \neq m(s'')$. To be able to continue using it as a valid SDF that is positive and that has an expectation that is equal to the observed inverse of the safe rate, requires building it by averaging over such states:

$$m(\pi) = \frac{\sum_{s:\pi(s)=\pi} p(s)m(s)}{\sum_{s:\pi(s)=\pi} p(s)}$$

This way, $\sum_{\pi\in\Pi} p(\pi)m(\pi) = \sum_{s\in S} p(s)m(s)$. This is the relevant $p(.)$ measure for the paper, not the one over states.

Second, for every possible outcome for inflation, consider a second stochastic variable:

$$\omega(\pi) = W_0 - \frac{\mathbb{E}^p(m)}{\pi},$$

measuring by how much does debt fall for that draw of inflation. To be completely clear, the expectation is with regards to the distribution of inflation: $\mathbb{E}^p(m) = \sum_{\pi\in\Pi} p(\pi)m(\pi)$. This is still stochastic in that it may be different for different realizations of s, as in the example of s' and s'' above. But it is risk free when it comes to variation in inflation.

We define the risk-adjusted probabilities as: $f(\pi) = m(\pi)p(\pi)/\mathbb{E}^p(m)$. Our distribution is then given by

$$\Phi(\bar{w}) = Prob^f(\omega(\pi) > \bar{w}),$$

that is: the risk-adjusted probability that debt debasement due to inflation variance exceeds a certain amount $\bar{w}$ in dollars or as a share of GDP. Define the set of possible realizations of inflation for which debasement exceeds this threshold to be: $\Pi^f = \{\pi : \omega(\pi) > \bar{w}\}$. Then, we calculate our measure using the simple formula:

$$\Phi(\bar{w}) = \sum_{\pi\in\Pi^f} f(\pi)$$

These are our estimates, which rely on the $f(\pi)$ distribution that we extract from the options. With multiple maturities, one needs instead a risk-adjusted distribution over paths of inflation over time, and a distribution of the maturity of debt coming due held in private hands, as we do in the main paper.

The alternative object:

$$Prob^p\left(W_0 - \frac{m(s)}{\pi(s)} > \bar{w}\right) = \sum_{s\in S^p} p(s)$$

where $S^p = \{s : W_0 - \frac{m(s)}{\pi(s)} > \bar{w}\}$ does not answer the question posed in the paper. It would give the probability of debasement resulting from *all variation* in the data, whether that comes from inflation or not. That is the first major difference between our estimates and the physical probability ones: we focus on inflation risk only. The inflation options provide the right distribution for our question, since they are over Π not S.

Second, consider the natural alternative object under the physical probabilities:

$$\phi(\bar{w}) = Prob^p\left(W_0 - \frac{m(\pi)}{\pi} > \bar{w}\right) = \sum_{\pi \in \Pi^p} p(\pi)$$

where the set is defined as $\Pi^f = \left\{\pi : W_0 - \frac{m(\pi)}{\pi} > \bar{w}\right\}$. Taking into account risk adjustment, that is comparing $\Phi(.)$ with $\phi(.)$ then requires not just comparing $f(\pi)$ and $p(\pi)$ but *also* Π^f and Π^p.

To make this comparison, we consider three cases in the text, each corresponding to different assumptions about $m(\pi)$, which we now state formally:

Case 1: classical dichotomy. In this case, there is no correlation between the SDF and inflation, so there is no inflation risk premium. That is, $m(\pi) = m$ is a constant. Given the way we defined it above, this constant is given by the expectation of the SDF over all the possible states:

$$m = \sum_{s \in S} p(s) m(s)$$

There can still be plenty of variability in $m(s)$ and so lots of other risk in the economy, but the relevant m with respect to inflation risk is just this constant. It then follow immediately that:

$$\begin{aligned} f(\pi) &= m(\pi)p(\pi)/\mathbb{E}^p(m) = p(\pi) \\ \Pi^f &= \left\{\pi : W_0 - \frac{\mathbb{E}^p(m)}{\pi} > \bar{w}\right\} = \left\{\pi : W_0 - \frac{m(\pi)}{\pi} > \bar{w}\right\} = \Pi^p \end{aligned}$$

Therefore physical and risk-adjusted debasement probabilities coincide: $\Phi(\bar{w}) = \phi(\bar{w})$.

Case 2: Additive risk premium in Gaussian models. These are models in which $m(\pi) = k(s)$. This $k(s)$ can vary over time, it is stochastic, but it is constant with respect to inflation. But since the definition of the risk-adjusted distribution $f(\pi)$ is only taking into account inflation risk, it follows that $\mathbb{E}^p(m) = k(s)$. In other words, our risk-free rate is only risk-free w.r.t. inflation risk. Then, by the same steps as in the previous case, it follows that physical and risk-adjusted probabilities of debasement coincide.

Case 3: High inflation commands higher risk premium. Our reading of the estimates of Kitsul and Wright is the following: for $\pi > \bar{\pi}$ we have $m(\pi) > \mathbb{E}^p(m)$. That is, in the right tail of the distributions, it is always the case that $f(\pi) > p(\pi)$. Loosely, the risk-adjusted probability distribution had a fatter right tail than the physical distribution.

Now, we are always looking at high percentiles in the debt debasement distribution. That is, we always consider large $\bar{w}$. Therefore, we are always in the tail of the distribution above this $\bar{\pi}$ where the risk-adjustment is a positive factor over the physical distribution. Now, this immediately delivers that, if one were to keep fixed the same set of inflation paths, the risk-adjusted probability is going to be higher. However to show that

$$\Phi(\bar{w}) = \sum_{\pi \in \Pi^f} f(\pi) > \sum_{\pi \in \Pi^p} p(\pi) = \phi(\bar{w})$$

also requires that Π^f is not too "small" in the sense of this expression relative to Π^p. From the definitions above of these two sets, a sufficient (but not necessary) condition for this to be the case is that:

$$\Pi^p \subset \Pi^f \qquad \equiv \qquad \left\{\pi : \pi > \frac{m(\pi)}{\bar{w} - W_0}\right\} \subset \left\{\pi : \pi > \frac{\mathbb{E}^p(m)}{\bar{w} - W_0}\right\}$$

Since we already know that in this range $m(\pi) > \mathbb{E}^p(m)$, only if $\pi/m(\pi)$ is strongly non-monotonic would this not hold. But note that this is not necessary, and it seems plausible that it holds (see Kitsul and Wright (2013)). Therefore, our (very small) percentiles in the VAR are upper bounds on the physical probabilities.

Finally, to conclude, the assumptions we introduced in the first paragraph were just to make the maths reduce to simple operations. Considering a rich maturity of the debt would be the same but now the sets would be defined as sets of sequences of draws of inflation over time, and the assumption in case 3 would be with respect to cumulative inflation at several horizons (which still matches what Kitsul and Wright (2013) report). The assumption of countable sets makes it easier to explain what we do draw by draw instead of using measures and integrating, which requires introducing more elements of real analysis with no change in the underlying logic.

G Proof of proposition 3

Financial repression consists of paying nominal bonds that are due at date t with new N-period special debt that sells for the price $\tilde{H}_t^N$:

$$B_t^0 = \tilde{H}_t^N \tilde{B}_{t+1}^{N-1}. \tag{A13}$$

The value of outstanding debt at date t now is:

$$W_t = \sum_{j=0}^{\infty} \frac{H_t^j B_t^j}{P_t} + \sum_{j=0}^{\infty} Q_t^j K_t^j + \frac{\tilde{B}_{t-N+1}^{N-1}}{P_t}, \tag{A14}$$

while the government budget constraint now is:

$$W_t = s_t + \sum_{j=0}^{\infty} \frac{H_t^{j+1} B_{t+1}^j}{P_t} + \sum_{j=0}^{\infty} Q_t^{j+1} K_{t+1}^j + \frac{\tilde{H}_t^N \tilde{B}_{t+1}^{N-1}}{P_t}. \tag{A15}$$

The text covered the special case where $N = 1$, there is no real debt, and all nominal debt had one period maturity. This appendix proves the general case.

Combining these two equations just as we did in the proof of proposition 1, we end up with a law of motion for debt:

$$W_t = Q_t^1 W_{t+1} + s_t + x_{t+1} + \frac{\tilde{H}_t^N \tilde{B}_{t+1}^{N-1}}{P_t} - Q_t^1 \frac{\tilde{B}_{t-N+2}^{N-1}}{P_{t+1}}. \tag{A16}$$

By precisely the same steps as in the proof of proposition 1, it then follows that:

$$W_0 = \mathbb{E}\left[\sum_{t=0}^{\infty} m_{0,t}\left(\frac{B_0^t}{P_t}\right)\right] + \mathbb{E}\left[\sum_{t=0}^{\infty} m_{0,t} K_0^t\right] = \mathbb{E}\left[\sum_{t=0}^{\infty} m_{0,t}\left(s_t + \frac{\tilde{H}_t^N \tilde{B}_{t+1}^{N-1} - Q_t^1 \tilde{B}_{t-N+2}^{N-1} P_t / P_{t+1}}{P_t}\right)\right]. \tag{A17}$$

Now, replacing the new debt with old debt using equation (A13), we get that the nominal debt burden then becomes:

$$\mathbb{E}\left[\sum_{t=0}^{\infty} m_{0,t}\left(\frac{B_0^t}{P_t}\right)\right] - \mathbb{E}\left[\sum_{t=0}^{\infty} m_{0,t}\left(\frac{B_0^t}{P_t}\right)\right] + \mathbb{E}\left[\sum_{t=0}^{\infty} m_{0,t} \frac{Q_t^1 B_{t-N+1}^0}{\tilde{H}_t^N P_{t+1}}\right]. \tag{A18}$$

Canceling terms and relabeling the limits of the sums (since financial repression started at date 0) we get the nominal debt burden:

$$\mathbb{E}\left[\sum_{t=0}^{\infty} m_{0,t+N-1} \frac{Q^1_{t+N-1} B^0_t}{\tilde{H}^N_t P_{t+N}}\right]. \tag{A19}$$

Finally, recall that $Q^1_{t+N-1} = \mathbb{E}_{t+N-1}(m_{t+N-1,t+N})$. Using the law of iterated expectations we end up with:

$$\mathbb{E}\left[\sum_{t=0}^{\infty} m_{0,t+N} \left(\frac{1}{\pi_{t,t+N} \tilde{H}^N_t}\right) \left(\frac{B^0_t}{P_t}\right)\right]. \tag{A20}$$

To get the equality in the proposition, simply use the upper bound $H^N_t = 1$. To get the equality written in terms of the price of nominal bonds, simply use the law of iterated expectations and the arbitrage condition: $H^N_t = \mathbb{E}\left(m_{t,t+N}/\pi_{t+N}\right)$.

References

Bahra, Bhupinder. 1997. "Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application." *Bank of England working paper 66.*

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.

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

Campa, Jose M., Kevin Chang, and Robert L. Reider. 1998. "Implied Exchange Rate Distributions: Evidence from OTC Option Markets." *Journal of International Money and Finance*, 17: 117–160.

Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig. 2017. "Deflation Risk." *Review of Financial Studies*, 30(8): 2719–2760.

Gürkaynak, Refet, Brian Sack, and Jonathan Wright. 2007. "The U.S. Treasury yield curve: 1961 to the present." *Journal of Monetary Economics*, 54(8): 2291–2304.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2010. "The TIPS yield curve and inflation compensation." *American Economic Journal: Macroeconomics*, 2(1): 70–92.

Hagan, Patrick S, Deep Kumar, Andrew S Lesniewski, and Diana E Woodward. 2002. "Managing smile risk." *The Best of Wilmott*, 1: 249–296.

Hall, George J., and Thomas J. Sargent. 2011. "Interest rate risk and other determinants of post-WWII US government debt/GDP dynamics." *American Economic Journal: Macroeconomics*, 3(3): 192–214.

Kitsul, Yuriy, and Jonathan H. Wright. 2013. "The economics of options-implied inflation probability density functions." *Journal of Financial Economics*, 110(3): 696–711.

Reis, Ricardo, and Mark W. Watson. 2010. "Relative Goods' Prices, Pure Inflation and the Phillips Correlation." *American Economic Journal: Macroeconomics*, 2(3): 128–157.

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

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