

Inflating Away the Public Debt? An Empirical Assessment

Jens Hilscher

University of California, Davis

Alon Raviv

Graduate School of Business Administration, Bar-Ilan University

Ricardo Reis

London School of Economics

This paper proposes a new method for measuring the impact of inflation on the real value of public debt. The distribution of debt debasement is based on two inputs: the distribution of privately held nominal debt by maturity, for which we provide new estimates, and the distribution of risk-adjusted inflation dynamics, for which we provide a novel copula estimator using options data. We find that inflation by itself is unlikely to lower the U.S. fiscal burden significantly because debt is concentrated at short maturities and perceived inflation shocks have little short-run persistence and are small. (*JEL* E31, E64, G12, G18)

Received August 8, 2018; editorial decision December 1, 2020 by Editor Philip Strahan. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Higher-than-expected inflation can have some benefits, and one of its most celebrated is to erode the real value of outstanding debt. Public debt is at

Daniel Albuquerque, Cynthia Balloch, Fabio Bertolotti, Keshav Dogra, Chao He, Nicola Limodio, and Savitar Sundaresan provided outstanding research assistance, and Eugene Kiselev and Kaiquan Wu helped to assemble the data. We are grateful to Philip Strahan (the editor) and the referees; Patrick Augustin, John Cochrane, John Leahy, Elisa Faraglia, and Carolin Pflueger for conference discussions; Mark Aguiar, Lars Hansen, Anton Korinek, Olivier Jeanne, Alan Taylor, and Pietro Veronesi for useful comments; seminar participants at Boston College, Brandeis, Chicago Booth, Columbia, Drexel, ENSAE, FRB Boston, FRB Minneapolis, FRB San Francisco, Helsinki University, Humboldt, Johns Hopkins, LBS, LSE, McGill, Northwestern, Wharton, Princeton, UC Davis, UCL, and York; and conference participants at the WFA annual meetings, AEA annual meetings, EFA annual meetings, the GEA meetings at DIW, the ENSAI economic day, the Shanghai forum, the CEPR-INET conference on government debt, the BFI conference on government debt, and the NASM of the Econometric Society at UC Davis. This work was supported by grants from the Institute for New Economic Thinking [grant INO13-00025]; the National Science Foundation [grant GG009759]; the US-Israel Binational Science Foundation [grant 2014336]; and the European Research Council in the EU's Horizon 2020 programme [ERC grant INFL, under GA: 682288]. Send correspondence to Ricardo Reis, r.a.reis@lse.ac.uk.

The Review of Financial Studies 00 (2021) 1–43

© The Author(s) 2021. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

doi:10.1093/rfs/hhab018

Advance access publication February 09, 2021

record high levels across many countries, after growing steadily for many years (Reinhart, Reinhart, and Rogoff 2012) and a sharp jump in response to the 2020 pandemic recession (International Monetary Fund 2020). A common way that sovereigns pay for high public debt is by having high, and sometimes even hyper, inflation (Reinhart and Rogoff 2009). Whether this is feasible or likely in the future is an open question. How likely will inflation be sufficiently high to substantially alleviate current heightened debt levels? What are the characteristics of the debt and the stochastic properties of inflation that make substantial debasement more or less likely?

In this paper, we propose a method to quantify the likelihood of future inflation substantially eroding the real value of current public debt. We measure the effect of inflation on the fiscal burden by constructing the distribution of inflation-driven debt debasement. We combine characteristics of the debt (its maturity structure, holders, and nominal or real payouts) and of future inflation (its persistence and the likelihood of large increases) to quantify the probability of inflation-driven debt reduction of various sizes. We then apply our method to U.S. annual data from 2009 to 2017 and thereby inform the debate on how the United States may deal with its outstanding public debt. Finally, we show how to map the distribution of debt debasement into central objects in theories of inflation and its effects.

Both our general goal of quantifying the fiscal impact of inflation and the concrete U.S. application require a debt debasement measure with four properties. First, it must be forward looking, using present data. We ask whether inflation can erode the debt in the future, not if it did so in the past. Second, it must be accurate, using the actual maturity and private holdings of the public debt. We do not provide back-of-the-envelope calculations, but precise estimates that net out holdings of public debt by government agencies and are based on the exact maturity structure. Third, the calculations must rely on plausible scenarios for inflation that are supported by data and reflect uncertainty. We ask whether public nominal debt might be inflated away quantitatively, not if doing so is theoretically possible.¹ Fourth, the calculation must take investors' expectations into account, and how these are reflected in prices, otherwise the measure will dramatically overstate what government can do. We use market prices for inflation in different scenarios extracted from options prices.

Focused on different questions, previous work on this topic has instead looked at the historical experience, it has used approximations for the debt maturity structure, and it has not emphasized using plausible and data-driven distributions of future inflation paths. Considering the historical distribution, Hall and Sargent (2011) characterize the contribution of inflation for the

¹ Of course it is possible; just drive the price level to infinity!

evolution of U.S. public debt in the past.² A different part of the literature studies the many theoretical links between inflation and the public debt, but typically assumes that public debt has a one-period maturity (Cochrane 2011 and Davig, Leeper, and Walker 2011) link high levels of U.S. debt to future higher inflation). Finally, a third set of papers provides back-of-the envelope calculations on how much inflation will lower the debt, but makes ad hoc and implausible assumptions on the dynamics of inflation and expectations (Aizenman and Marion 2011 and Bohn 2011). Our inflation-driven debt debasement distribution is forward looking, accurate, and based on plausible future inflation scenarios.

We start by developing the theory underlying our measure. Based on the law of motion for debt and the intertemporal budget constraint, we derive a simple forward-looking formula showing that the real value of debt is equal to a weighted average of the payments due at different horizons, with the weights given by the expected discounted inverse of compounded risk-adjusted inflation. The formula makes transparent how the distribution of future risk-adjusted inflation paths interacts with the maturity of debt to map onto the distribution of inflation-induced debt debasement. Because the theory underlying our measure requires only the existence of a government budget constraint and a stochastic discount factor to price government bonds, our measure can be easily embedded into equilibrium macroeconomics models, to provide new insights into discussions about inflation targeting, the fiscal theory of the price level, and the maturity of government debt.³

Calculating debasement distributions requires two main inputs. First, detailed data on holders and maturity of nominal public debt. Our new empirical estimates of the evolution of the maturity structure of privately held U.S. debt from 2009 to 2017 are an independent contribution of interest, as they imply a significantly shorter distribution than commonly appreciated. Nominal privately held debt is slightly less than half of total outstanding debt and is disproportionately concentrated in shorter maturities. Second, we need the joint distribution of future risk-adjusted inflation at different horizons. We simulate future inflation paths, which in turn allows us to quantify the risk-adjusted probability of debasement of various sizes. We construct forward-looking density functions for cumulative inflation using data on inflation caps and floors at different maturities, for example, the probability of average inflation over the next two years lying below various cutoffs, like others before us.⁴ We go a step

² Reinhart and Sbrancia (2015) conduct a similar exercise for other developed countries, while Berndt, Lustig, and Yeltekin (2012) and Chung and Leeper (2007) use vector autoregressions to handle expectations of spending shocks, and Giannitsarou and Scott (2008) link spending shocks with governments choosing to inflate.

³ See Cochrane (2017), Del Negro and Sims (2015), Bianchi and Melosi (2017), Missale and Blanchard (1994), Persson et al. (2006), Debortoli et al. (2017).

⁴ Kitsul and Wright (2013) look at the response of the density around monetary policy announcements, while Fleckenstein, Longstaff, and Lustig (2017) assess the risk of deflation in the United States. We instead study the effect of future inflation on the real value of debt.

further by complementing these data with data on forward distributions, that is, probabilities of inflation in a given future year lying below various cutoffs. With both data, we propose a novel econometric method to construct joint distributions of future inflation from option prices that relies on the theory of copulas. These are of independent interest, as they might be useful in other contexts. The data reveal that risk-adjusted inflation is estimated to be generally close to stationary but that the process has long memory, reflected in a slow decay of the autocorrelation function.

Combining our approach and these two inputs, we measure the probability of debasement exceeding cutoffs of various magnitudes for the United States for annual data from 2009 to 2017. We refer to these probabilities, under the forward measure for various maturities, which are appropriate for valuation, as risk-adjusted or risk-neutral probabilities. While objective and subjective probability forecasts for inflation can be extracted in many ways, including financial prices, surveys, and economic and statistical models, these methods tend to focus on forecasting the mean, while placing less emphasis on higher moments. Our goal is to instead provide value-at-risk measures of the likelihood of debt debasement due to inflation. In addition, and crucially, existing measures are not appropriate for pricing. Our goal is to measure the market value of the debt under different scenarios, so we need the risk-adjusted probabilities that are relevant for pricing the government debt. Using these probabilities also allows us to embed reactions of market prices into the debasement distribution. We find that, in 2017, the risk-adjusted probability of debasement lying below a mere 3% of gross domestic product (GDP) is almost 95%. Any loss above 3.7% has less than a 1% probability. Moreover, most of the losses fall on foreign holders of the government debt and the Federal Reserve is also hit. Our results are robust to using real probabilities, instead of risk-adjusted ones, because instances of high inflation are usually those for which marginal utility is high, implying that our risk-adjusted debasement cutoffs represent an upper bound of debasement under real probabilities.

We conduct a series of stress tests, or counterfactuals, disciplined by the data on inflation distributions. We show that our main results are due to two complementary factors. First, the maturity of the relevant U.S government debt, which is sensitive to inflation and held privately, is quite low. Second and related, over short horizons of a few years, market participants place a very low probability on persistently high U.S. inflation. In the near term, there is much debt but little extra inflation, and for longer horizons, there could be significant inflation but little debt. As a result, the total effect is small.

While this interaction roughly explains the broad conclusion on the U.S. application, our goal is to supply methods that provide precise and accurate numerical answers across applications at different times and for different countries. Starting from our accurate calculations, first we identify the main drivers of debasement: debt duration, inflation volatility, and inflation persistence. Then, we propose a new rule of thumb calculation that combines

the three measures. A different concept sometimes used as a rule of thumb is duration (the sensitivity of debasement to a marginal shock to inflation). By itself, any measure of sensitivity would not deliver a distribution. Any such measure would not say anything about the likely magnitude of debasement, because it would ignore the volatility of inflation. Thus, it cannot provide a probability of inflation being high enough to erode the debt, for instance, by more than 1%. More importantly, duration provides an incorrect answer to the question we pose. Duration does not take into account the stochastic properties of inflation, since it assumes a permanent increase in inflation that is parallel across the whole distribution, when, in fact, actual and expected inflation persistence is much lower. Our proposed accurate calculations and associated new rule of thumb take into account both volatility and persistence, which together with the maturity of the debt captures most of the variation in debasement over time in the U.S. data.

Our paper focuses on debt debasement that is driven by *inflation alone*. If high inflation comes associated with low real interest rates and high growth, then these combined effects may indeed lead to a significant decline in debt-to-GDP. We complement our measures with estimates of the joint distribution of output and inflation from an estimated vector autoregression. Realizations of higher inflation are associated with higher-than-expected output growth, because of the Phillips curve, resulting in a higher reduction of the debt. A third driver of debt-to-GDP dynamics is the effect of changes in the real term structure of interest rates. We propose a novel way of decomposing annual changes in the debt-to-GDP ratio into the inflation surprise, the growth surprise, and the real interest rate surprise. All of these components combined, at different times from 2009 to 2017, to affect the debt-to-GDP ratio. In this sense, a joint mix of policy and outcomes that includes high inflation can indeed lead to “inflating away the debt,” although, it would be just as adequate to call this “growing out of the debt.” To move beyond these one-liners, we relate our estimates to several theoretical applications. Many models assume one-period debt for convenience, but then struggle to adequately calibrate the link between the fiscal burden and inflation in the actual data. Our estimates provide the correct targets for these calibrations. In particular, our estimates measure the extent of fiscal backing of an inflation-targeting central bank, they provide the joint distribution of fiscally driven inflation, and they provide the key elasticity of the debt burden to inflation that is needed to calibrate models of tax smoothing and sovereign default.

Inflation also often comes with financial repression, especially in developing countries. Such a policy provides revenue directly, which is easy to estimate and, more interesting for this paper, it substantially increases the scope for inflationary debasement in a way that our new method can assess. We show that extreme financial repression, where bondholders are paid with reserves at the central bank that pay zero interest, yet must be held for a fixed number of periods, is equivalent to ex post extending the maturity of the debt. Under such

circumstances, inflation has a much larger impact, so that if repression lasts for a decade, permanently higher inflation that previously lowered the real value of debt by 2.3% now lowers it by 18.7% of GDP.

In conclusion, it is useful to point out also what we do not do. We study the effect of higher-than-expected inflation on debt debasement. Inflation may affect the state of public finances through other links, like seignorage or tax revenues (Krause and Moyen 2016). This focus allows us to make a minimal set of assumptions of no arbitrage in the government debt market, without having to use a full equilibrium model. Also, inflation leads to redistributions of wealth (Doepke and Schneider 2006, Berriel 2013). We solely focus on the redistribution from private holders of debt, domestic or foreign, to the fiscal authority. Finally, our goal is positive, not normative, so we do not describe whether it is desirable to use inflation given the current maturity of the debt (Faraglia et al. 2013, Leeper and Zhou 2013) or its holders.

1. Theory: The Debt Burden and Risk-Neutral Densities

Our goal is to quantify the effect of future inflation on the real value of debt. This is not straightforward since many future inflation paths and outstanding debt securities exist. This section provides a general but workable definition of the debt burden. We allow for a rich maturity structure of both nominal and real debt, use no arbitrage and a present value approach to aggregate the numerous ways in which future inflation can affect the value of debt, and use market expectations.

1.1 The public debt

The real market value of government debt at date t is equal to the present discounted value of the full maturity structure of future promised payments:

$$W_t = \sum_{j=0}^{\infty} \frac{H_t^j B_t^j}{P_t} + \sum_{j=0}^{\infty} Q_t^j K_t^j. \quad (1)$$

B_t^j is the face value of zero-coupon nominal debt held at the beginning of period t (carried over from period $t-1$) that has a maturity of j years, so that at date t the government expects to pay B_t^j dollars at date $t+j$. K_t^j is the par value of real debt held at date t that has a maturity of j years, referring to, for example, Treasury inflation protected securities (TIPS). H_t^j is the market price (or inverse yield) of nominal debt with a maturity of j years at date t . Likewise, Q_t^j is the price (or inverse yield) of real liabilities (TIPS) with a maturity of j years at date t . Finally, P_t is the price level, and we will use the notation $\pi_{t,t+j} = P_{t+j}/P_t$ to denote gross cumulative inflation between two dates. The following normalizations apply: $H_t^0 = Q_t^0 = 1$ and $P_0 = 1$.

Modeling the government debt this way involves a few simplifications. First, the government often has a wide variety of outstanding debts. The implicit

assumption above is that their prices are the same as those of all-else-identical traded debt securities. Second, it assumes that coupon-paying bonds can be priced as portfolios of zero-coupon bonds. In this way, we limit the set of debt instruments issued by the government and simply consider promised payments (either principal or coupon payments) at each point in time. Third, unfunded real liabilities of the government like Social Security can be included in K_t^j . Theoretically, they pose no problem. In practice, measuring any of these precisely is a challenge. However, since they are real liabilities, they are unaffected by inflation and therefore do not affect our debt debasement calculation. We therefore focus on nominal public debt instruments issued by the government.

If all debt were due right away, then the expression in Equation (1) would reduce to $B_t^0/P_t + K_t^0$. The simple rule of thumb that an increase in P_t lowers the debt burden proportionately to the privately held nominal debt is accurate. However, with longer maturities, higher future price levels affect yields and so also the value of debt. In addition, a whole sequence of future price levels have the ability to reduce the real value of debt.

1.2 The law of motion for debt

Each period, to pay for the debt, the government must either collect a real fiscal primary surplus of s_t or borrow more from the private sector:

$$W_t = s_t + \sum_{j=0}^{\infty} \frac{H_t^{j+1} B_{t+1}^j}{P_t} + \sum_{j=0}^{\infty} Q_t^{j+1} K_{t+1}^j. \quad (2)$$

Combining the previous two equations provides a law of motion for debt. Looking forward from date 0 for t periods, we can write it as

$$\begin{aligned} W_0 = & W_{t+1} \prod_{\tau=0}^t Q_{\tau}^1 + \sum_{i=0}^t \prod_{\tau=0}^i Q_{\tau}^1 s_i + \sum_{i=0}^t \prod_{\tau=0}^i Q_{\tau}^1 \sum_{j=0}^{\infty} \left(H_i^{j+1} - H_i^1 H_{i+1}^j \right) \frac{B_{i+1}^j}{P_i} \\ & + \sum_{i=0}^t \prod_{\tau=0}^i Q_{\tau}^1 \sum_{j=0}^{\infty} \left(Q_i^{j+1} - Q_i^1 Q_{i+1}^j \right) K_{i+1}^j \\ & + \sum_{i=0}^t \prod_{\tau=0}^i Q_{\tau}^1 \left(\frac{H_i^1 P_{i+1}}{P_i} - Q_i^1 \right) \sum_{j=0}^{\infty} \frac{H_i^{j+1} B_{i+1}^j}{P_{i+1}}. \end{aligned} \quad (3)$$

This equation elucidates that we need some way to aggregate the effect of future inflation realizations. Inflation can affect almost every term on the right-hand side. Worse, to judge how a particular path for inflation $\{\pi_{0,i}\}_{i=0}^t$ affects the fiscal burden, we would need to know how inflation will change the slope of the yield curve at every maturity (the $H_i^{j+1} - H_i^1 H_{i+1}^j$ term), the composition of the debt by maturity, and the future real/nominal split (the B_{i+1}^j and K_{i+1}^j

terms). Likewise, we would need to know the link between inflation and the real yield curve (the $Q_i^{j+1} - Q_i^1 Q_{i+1}^j$ term) as well as the ex post differences between nominal and real returns (the $H_i^1 P_{i+1}/P_i - Q_i^1$ term). Finally, recall that this expression holds for every possible path of inflation as well as for every realization of uncertainty in the economy. Without some aggregation, an unwieldy large number of measures could determine how much the fiscal burden will change in the future.⁵

1.3 Looking forward: The intertemporal budget constraint

We take a present value approach and assume that a stochastic discount factor prices all of these government liabilities. It is well understood that this is equivalent to requiring the absence of arbitrage. Note that all we require is no arbitrage between government bonds of different maturities, not across asset classes. U.S. government bond markets are among the most liquid in the world, have fewer restrictions on short-selling, and serve as the fundamental asset for many traded derivatives. We also allow for varying risk premiums or excess profits across different maturities, as in models of segmented markets or preferred habitats, so we are not ruling out possible liquidity premiums or the effectiveness of quantitative easing policies (Vayanos and Vila 2021).

The stochastic discount factor at date t for a real payoff at date $t+j$ is denoted by $m_{t,t+j}$, and the conventional pricing equations for j -period bonds are

$$1 = \mathbb{E}_t \left(\frac{m_{t,t+j}}{Q_t^j} \right) = \mathbb{E}_t \left(\frac{m_{t,t+j} P_t}{H_t^j P_{t+j}} \right). \quad (4)$$

A nominal discount bond costs H_t^j / P_t in real units at date t , and pays off $1 / P_{t+j}$ real units in j periods; its return times the stochastic discount factor has to have an expectation of 1. The same applies for a real bond. The absence of arbitrage over time implies that the stochastic discount factors across any two maturities, n and j , are linked by: $m_{t,t+j} = m_{t,t+n} m_{t+n,t+j}$ for $1 \leq n \leq j$.

Multiplying by stochastic discount factors at different dates and taking expectations of Equation (3), while taking the limit as time goes to infinity, and imposing that the government cannot run a Ponzi scheme, we obtain the following result:⁶

$$W_0 = \mathbb{E} \left[\sum_{t=0}^{\infty} m_{0,t} \left(\frac{B_0^t}{P_t} \right) \right] + \mathbb{E} \left[\sum_{t=0}^{\infty} m_{0,t} K_0^t \right] = \mathbb{E} \left[\sum_{t=0}^{\infty} m_{0,t} S_t \right]. \quad (5)$$

The first equality provides a workable measure of the time zero debt burden. The government liabilities are fixed at time zero, so this expression depends

⁵ Hall and Sargent (2011) overcome these problems by using a version of this equation that looks backward, instead of forward. Given debt in the present (W_{t+1}) and in the past (W_0), there are historical data on most of the terms above. Our purpose is to look forward to determine how debt depends on future, not past, inflation.

⁶ We could allow for bubbles or Ponzi schemes by the government. As long as their value does not depend on inflation, then Proposition 1 is unchanged.

only on real discount factors looking into the future and on the distribution of all future price levels. Moreover, it shows that because the real payments do not depend on inflation, we can focus on the nominal debt to assess the effect of inflation.

The second equality shows that we can interpret our measure as saying how much fewer taxes the government can collect by lowering the debt burden. Higher inflation may not only lower the real payments on the outstanding nominal debt but also change primary fiscal surpluses. Our purpose in this paper is to measure the effect of inflation on the outstanding public debt, and expression (5) shows that this depends on outstanding nominal debt alone.

1.4 Inflationary debasement: A formula for the debt burden as a weighted average

The price level is the only source of uncertainty in determining how much the outstanding nominal debts that mature in t periods will pay in real terms. Therefore, even though the stochastic discount factor depends, in principle, on all sources of uncertainty in the economy, only its marginal distribution with respect to inflation will lead to nonzero terms once we multiply by inverse inflation and take expectations. Therefore,

$$\mathbb{E} \left[\sum_{t=0}^{\infty} m_{0,t} \left(\frac{B_0^t}{P_t} \right) \right] = \sum_{t=0}^{\infty} B_0^t \mathbb{E} \left(\frac{m(\pi_{0,t})}{\pi_{0,t}} \right), \quad (6)$$

where $\pi_{0,t} = P_t/P_0$ and $P_0 = 1$. Next, recall the standard definition of a risk-neutral density $f(\cdot)$, which is equal to the product of the physical density times the stochastic discount factor divided by the expectation of this product. Recall as well that $R_t = \mathbb{E}(m(\pi_{0,t}))^{-1}$ is the real risk-free return between 0 and t from the perspective of date 0. By definition, it does not depend on future realizations of inflation. Then, it follows that

$$\mathbb{E} \left(\frac{m(\pi_{0,t})}{\pi_{0,t}} \right) = R_t^{-1} \int \left(\frac{f(\pi_{0,t})}{\pi_{0,t}} \right) d\pi_{0,t}. \quad (7)$$

Combining all the results gives our formula for the debt burden as a function of inflation.

Proposition 1. The debt burden is a weighted average of the nominal payments that the government must make at all present and future dates:

$$\sum_{t=0}^{\infty} \omega_t B_0^t \quad (8)$$

with weights given by

$$\omega_t = R_t^{-1} \int \left(\frac{f(\pi_{0,t})}{\pi_{0,t}} \right) d\pi_{0,t}. \quad (9)$$

Our new measure, the distribution of inflationary debt debasement, relies on this formula. The current market value of debt is the expected discounted value of future promised payments. The distribution of debasement results from the dispersion around this mean, specifically the probability of extreme positive shocks to inflation. Recalling the four desired properties of a measure stated in the introduction, this formula can satisfy them all.

First, the forward-looking formula links future inflation to the debt burden today and provides a way to aggregate all the possible different paths for inflation. It forms the basis for computing a value-at-risk measure of debt debasement: the probability that the reduction of debt value via inflation lies above a certain cutoff.

Second, the formula shows what matters, namely, the promised nominal payments at different horizons to the private holders of debt. The maturity of the debt may change in the future in arbitrary ways, but from today's perspective, the ability to inflate current debt depends only on today's distribution, which is included without any approximations.

Third, the formula explicitly takes into account that inflation is stochastic and not perfectly controlled by policy. However, the foreseen paths for inflation are not arbitrary, but rather are disciplined by market expectations as revealed by the joint distribution of inflation which reflects its time-series properties. The probabilities depend on both shocks to fundamentals and potential changes in the policy regime as perceived by the markets.

Fourth, it depends on inflation alone, as all of its relevant effects on prices are captured in the inflation densities. Beyond the direct effect of inflation on the real payments, this also includes the possible effect of inflation on real interest rates through Fisher effects, or changes over time in the compensation for inflation risk, or in liquidity premiums. The relevant densities are for risk-adjusted inflation, so factors that affect bondholders' views and market prices are captured.

2. Data: U.S. Treasuries by Holder and Maturity and Distributions of Future Inflation

The two inputs required by Proposition 1 are the maturity structure of privately held nominal government debt and the joint risk-neutral distribution of future inflation. We now provide data and methods to construct both.

The Treasury reports the total U.S. federal debt at the end of 2017 as \$20.5 trillion, or 103.6% of GDP. Yet, this number does not distinguish between nominal and indexed bonds, nor does it include information about the maturity structure. Moreover, it includes debt held by different branches of the U.S. government. We want to measure government liabilities held by the private sector since any gains of the Treasury at the expense of another government account, will sooner or later have to be covered by the Treasury.

2.1 Holdings of U.S. Treasuries at different maturities

The Internet Appendix describes our multiple data sources. The main source of data is the Center for Research on Security Prices (CRSP) that reports the private holdings of all outstanding marketable government notes and bonds at the end of 2017. We measure monthly total nominal payments, using both face value and coupons at each maturity. We adjust the size of payments to exclude holdings by state and local governments. We then use the nominal yield curve to calculate market values.

At the end of 2017, the total market value of privately held marketable nominal Treasury securities was \$9.6 trillion, or 49.6% of GDP.⁷ This baseline number is much smaller than the face-value \$20.5 trillion headline number reported in the Monthly Statement of the Public Debt because it does not include holdings by other branches of government and real securities. The headline number includes nonmarketable securities (\$6.0 trillion), the vast majority of which are part of intragovernment holdings (\$5.7 trillion). The largest single such holding (\$2.8 trillion) is the Federal Old-Age and Survivors Insurance Trust Fund (Social Security). Excluding nonmarketable debt therefore does not materially affect our results since we would be excluding most of it on account of it being intragovernmental. The headline number also includes the amount of outstanding TIPS (\$1.3 trillion), which are almost exclusively held by the private sector, but which will retain their real value in the face of inflation. We also exclude debt held by state and local governments, especially in state and local pension funds.⁸ Holdings are approximately 7.6% of total privately held debt. Privately held debt also does not include holdings of the Federal Reserve. Any losses on the portfolio of the central bank will map directly into smaller seignorage payments to the Treasury (Hall and Reis 2013), so the same argument that excludes Social Security holdings applies to Fed holdings (\$2.3 trillion face value). Finally, we decompose total private holdings into domestic and foreign by maturity, as debasement may affect the two groups differentially. The market value of foreign-held debt is \$6.3 trillion, or 32.4% of GDP.⁹

Figure 1 shows the maturity structure of privately held nominal debt (B_0^t) using monthly data. A noticeable feature of the distribution is how concentrated it is at the short end. The maturity distribution decays very quickly in the first few years, then flattens between 5 and 10 years, becomes close to zero between 10 to 20 years partly because of the lack of issuance of 30-year bonds between 2001

⁷ All market values are calculated using the zero coupon yield curve provided by [Gürkaynak, Sack, and Wright \(2007\)](#). We ensure that the link between the maturity structure of nominal yields, inflation distributions, and the real yield curve holds by constructing a real yield curve according to Equation (9). Using instead real yields provided by [Gürkaynak, Sack, and Wright \(2010\)](#) does not affect our results.

⁸ We do not include debt issued by state and local governments, estimated by the Census to be \$3.0 trillion in 2016. A large variety of these debt instruments exists and no good source that reports their private holdings. We therefore restrict ourselves to federal debt.

⁹ This value probably overstates how much foreigners ultimately own, since the value includes holdings in offshore financial centers, some of which may be held by American citizens.

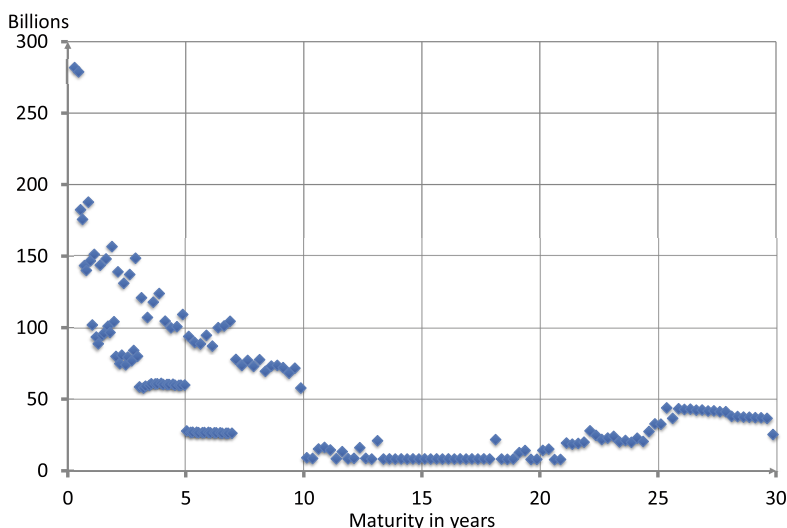


Figure 1
Government payments due to the private sector by maturity (B_0^t)

and 2006, and then picks up starting at 20 and especially from 25 to 30 years. The figure shows that simple approximations, like assuming a single bond with a maturity equal to the average, or assuming an exponential distribution, are not appropriate. The average maturity of the U.S. government debt, weighted by present value of private holdings, is 5.1 years according to our calculations. The maturity structure and average maturity imply that there is a scope for debasement, but the concentration at short maturities means that inflation in the near term will have the highest impact.

Although it is not included in the privately held debt number, we also assemble data on Fed holdings from the System Open Market Account (SOMA). At 12.6% of GDP, the value of the Fed portfolio is a little smaller than domestically held debt in private hands.¹⁰ Domestic and foreign holdings as well as Fed holdings have different maturity structures and will, therefore, have different exposures to inflation. Figure 2 plots foreign holdings (so domestic is the difference between private and foreign), together with the Fed's holdings, at an annual frequency. Private domestic holdings of U.S. Treasuries are concentrated in low maturities, with little holdings above 7 years. Foreigners account for most of the holdings above 5 years, and the Federal Reserve holds a disproportionate share of maturities above 5 years partly due to quantitative easing policies.

¹⁰ We do not include reserves issued by the central bank, not because they are not government debt, which they are, but because their maturity is overnight.

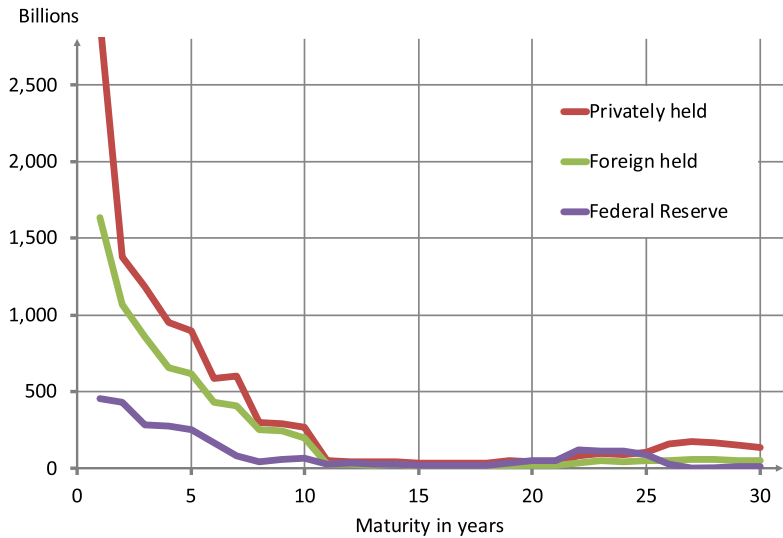


Figure 2
Breakdown of annual payments by holder of debt, by maturity

Finally, for comparison, we also assemble data on those real liabilities that are possible to measure with at least some accuracy. Figure 3 plots annual outstanding TIPS and Social Security liabilities up to 2030. TIPS liabilities follow a similar pattern to nominal liabilities and are much smaller. Social Security liabilities combine Old-Age and Survivors Insurance (OASI) and Disability Insurance (DI) projected shortfalls, the difference between direct program revenues, thus excluding trust fund interest payments, and costs. These numbers are put together by the Social Security Administration (SSA) in order to project forward what will happen to the trust fund. If Social Security benefits are not reduced, these real liabilities may become quite large, especially at longer maturities. However, the liabilities are real since benefits are indexed to inflation and therefore cannot be inflated away.

2.2 The marginal densities of expected inflation

We now turn to constructing the maturity structures of marginal risk-adjusted distributions of cumulative and 1-year forward inflation.

2.2.1 Data on zero-coupon and year-on-year inflation options. The market for over-the-counter U.S. dollar (USD) inflation options emerged in 2002. Since 2009, the market has been liquid enough to reliably reflect market expectations

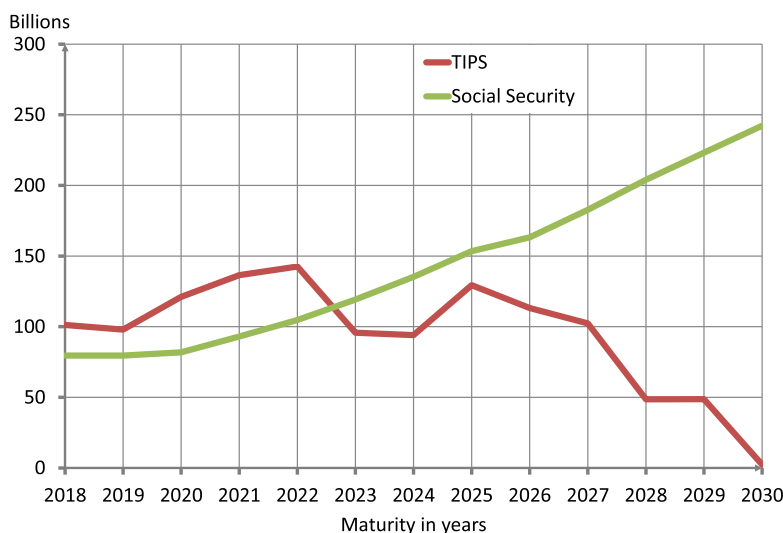


Figure 3
TIPS and projected Social Security liabilities by maturity (K_0^L)

of inflation.¹¹ We use daily data of caps and floors on the CPI for January 9, 2018, from Bloomberg.¹²

Like Kitsul and Wright (2013) and Fleckenstein, Longstaff, and Lustig (2017), we use data on zero-coupon caps and floors that pay off if cumulative inflation between the start of the contract and its maturity lies above or below the annually compounded strike price. The strike price ranges from -2% to 3% (floors) and 1% to 6% (caps), in 0.5% increments. We have data for all maturities between 1 and 10 years, together with data for 12 and 15 years, which we use to check our estimates.

Unlike Kitsul and Wright (2013) and Fleckenstein, Longstaff, and Lustig (2017), we also use data on year-on-year inflation caps and floors. These contracts are portfolios of caplets and floorlets that pay off at the end of each year if inflation during that year is above or below the strike price. Strike prices range from -3% to 6% . These contracts thus depend on inflation realizations during 1-year periods in the future. There seems to be no market demand for instruments with strike prices outside of this range, and contract prices imply that risk-adjusted probabilities of lying outside of these limits are close to zero. We use year-on-year options with maturities that range from 1 to 10 years.

¹¹ Kitsul and Wright (2013) write that by 2011 trading in the interdealer market was \$22 billion; J.P. Morgan (2013) estimates the annual trading volume in inflation derivatives in 2013 was \$50 billion.

¹² The Internet Appendix looks at dates nearby to ensure there was nothing special about this date, and contains more details on the data, the method to estimate the distributions, and robustness checks.

2.2.2 Recovering the risk-adjusted probability density of inflation. As the classic work of [Breedon and Litzenberger \(1978\)](#) noted, given a rich enough set of option contracts with observable prices, it is possible to recover nonparametrically the risk-neutral distribution of inflation without making any specific distributional assumptions about inflation or its link to other asset prices. The formula linking the risk-adjusted distribution of inflation and the price (a_0) and strike ($\pi_{0,t}$) of an option with maturity of t years is

$$f(\pi_{0,t}) = \pi_{0,t} R_t \left(\frac{\partial^2 a_0}{\partial \pi_{0,t}^2} \right). \quad (10)$$

Intuitively, if the price of a call option declines quickly with the strike price, then the outcomes around this strike price are more likely. Because our data give us many option prices for different strike prices and at different intervals, we can estimate this partial derivative by using the differences in these prices.

Instead of point expectations of future inflation, we recover risk-adjusted distributions for future inflation. That is, the distributions reflect the likelihood of different values of inflation, the risk associated with them, and the market price of this risk. This distinguishes risk-neutral distributions from many of the common measures of inflation expectations. Unlike opinion surveys, we are extracting risk-neutral rather than subjective expectations, and we do so from observing profit-making behavior. Unlike the break-even rate of inflation from comparing real and nominal yields, we have a whole distribution for inflation instead of a single number. Moreover, we do not need to worry about the price of the embedded floor which ensures that TIPS always pay back at least par value. Finally, unlike models of the term structure that use the yield curve to extract market-based inflation expectations, these densities do not rely on the associated (often strong) identifying assumptions.

The distributions for expected inflation recovered from options reflect the expectations of participants in this market: pension funds, life insurance companies, banks, and speculators placing bets on inflation. These entities are some of the major holders of the public debt, whose positions would be inflated away. Therefore, theirs are the relevant expectations for our valuation question. As with any measure of financial prices, one may worry about trading and liquidity frictions driving a wedge between prices and expectations. However, for low-frequency questions, such as the one we pose here, any sluggishness of adjustment of prices should not be too important, and we have checked that our results are robust to using options prices from one month before and after. We have also checked that distributions react to news at daily frequencies at least in the right directions. We provide additional details on data cleaning and construction in the Internet Appendix.

2.2.3 Zero-coupon and year-on-year distributions. Using our data on zero-coupon floors and caps, we extract the density $f(\pi_{t,t+j})$, with $j = 1, \dots, 10$. This

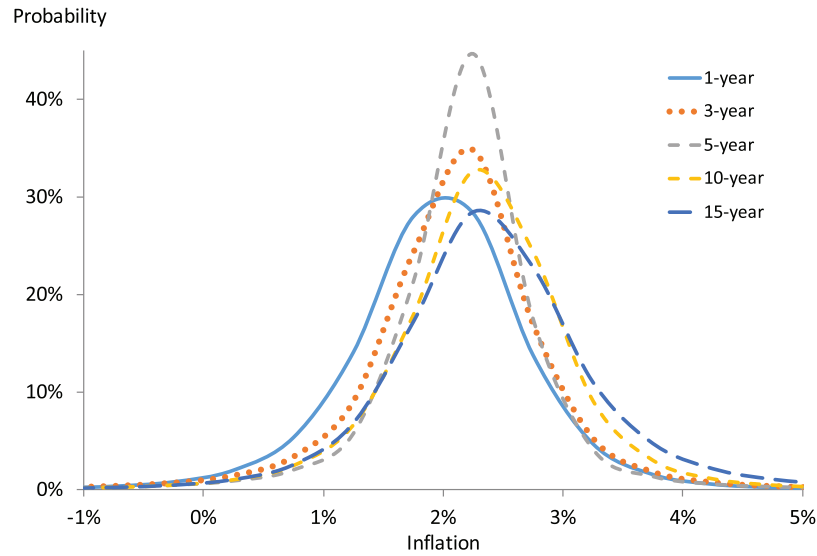


Figure 4
Marginal risk-adjusted distributions for cumulative annualized inflation

gives a term structure of the cumulative risk-adjusted inflation distributions. In turn, using data on year-on-year contracts, we construct 1-year forward risk-adjusted distributions $f(\pi_{t+j-1,t+j})$. Figures 4 and 5 plot the distributions at the end of 2017.

The 1-year mode of inflation is 2%.¹³ All of the distributions in Figure 4 are bell shaped, and their medians and modes are all between 1.75% and 2.5%; the means all lie between 2.0% and 2.4%. All of the distributions have fat tails and are significantly nonnormal. Kitsul and Wright (2013) interpret the tails as reflecting investors' perception that both very high and very low inflation are the costly states of the world. The patterns for the 1-year forward distributions are similar, except that the means at longer maturities are a little higher, rising from 2.0% (1 year) to 2.6% (10 year).

An interesting feature of the cumulative distributions is that, as the horizon increases, the variance is close to constant, and does not fall, something we would expect if the variance in 1-year forward inflation does not rise significantly with the horizon and shocks across years are imperfectly correlated. The standard deviation for cumulative inflation rises slightly from 0.77% at maturity 5 to 0.96% for maturity 15. This could be either because extreme events far in the future are perceived as more costly or because of more uncertainty about inflation. For year-on-year inflation, the standard deviations

¹³ Because of the 0.5% granularity of our data, we report the probability of lying between 1% and 1.5% as a probability mass at 1.25%. The probabilities plotted at 1.75% and 2.25% are almost exactly the same.

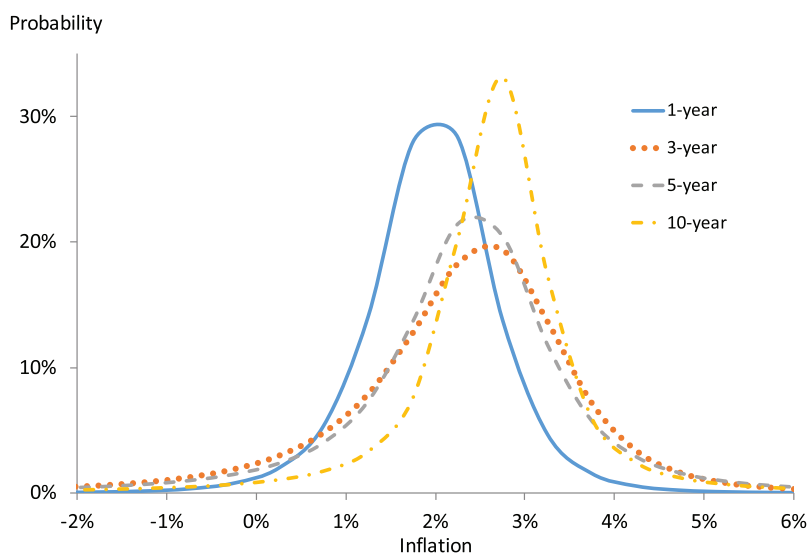


Figure 5
Marginal risk-adjusted distributions for year-on-year inflation

range between 0.8% and 1.9%, but standard deviations from 4 to 10 years all lie between 1.7% and 1.9%.

Overall, the distributions show that the probability of average annual inflation exceeding 4% is below 1.7% across all horizons up to 10 years, and never exceeds 1.7%; the probability of 1-year forward inflation above 5% never exceeds 3.2% and lies below 3% for half of the 10 year maturities. Sustained high risk-adjusted inflation is perceived as being an extremely remote possibility.

2.3 The joint distribution of inflation

From these marginal distributions we next construct joint risk-adjusted distributions of inflation across multiple years. Why do we need the joint distribution and how does the data that we have allow us to construct it? To construct a distribution of future inflation paths, one needs to know the joint distribution of inflation in year one (over the next year) and year two (the following year). Only then can one figure out the conditional distribution in year two based on a given realization in year one. Data on cumulative and 1-year forward distributions provide information on comovement across years. For example, if the correlation across years is high, then the 2-year cumulative distribution will have a relatively higher dispersion compared to the case in which the comovement across years is low. This section turns this simple observation into a method to estimate the joint distribution of inflation over time.

2.3.1 A method of moments copula-based estimator. Understanding how the realizations of a random variable are related over time is, of course, the classic problem in time-series modeling. Our particular data on inflation contracts provides a novel way to approach this problem that has the intriguing feature of only using data at one date in time. We do so by exploiting agents' expectations in financial markets to recover their beliefs about the time series of risk-adjusted inflation.

Consider the problem of obtaining the risk-adjusted joint density between annual inflation over the next 2 years: $f(\ln\pi_{t,t+1}, \ln\pi_{t+1,t+2})$. Sklar (1959) shows that there exists a copula function $C(\cdot): [0, 1]^2 \rightarrow [0, 1]$ such that

$$f(\ln\pi_{t,t+1}, \ln\pi_{t+1,t+2}) = C(f(\ln\pi_{t,t+1}), f(\ln\pi_{t+1,t+2})). \quad (11)$$

This function captures the codependence between the two random variables, so that we can obtain the joint density given information on the marginals. We use a parametric version of the copula function, $\hat{C}(f(\ln\pi_{t,t+1}), f(\ln\pi_{t+1,t+2}), \rho)$, so the unknown copula function is fully characterized by a vector of parameters ρ in the known function $\hat{C}(\cdot)$, where ρ is of dimension M . The typical approach in the literature that estimates copulas would be to use the time series for past inflation to estimate both the marginal densities and the parameters in ρ . Our unusual data allow us to approach the problem differently.

To start, we already have estimates of the marginal densities of forward inflation implied by year-on-year inflation contracts. Moreover, from the zero-coupon options, we also have another marginal distribution: $f(\ln\pi_{t,t+2}) = f(\ln\pi_{t,t+1} + \ln\pi_{t+1,t+2})$. From the definition of the distribution, we write

$$\begin{aligned} & f(\ln\pi_{t,t+2}) \\ &= \int_{\ln\pi_{t,t+1} + \ln\pi_{t+1,t+2} = \ln\pi_{t,t+2}} \hat{C}(f(\ln\pi_{t,t+1}), f(\ln\pi_{t+1,t+2}), \rho) d\pi_{t,t+1} d\pi_{t+1,t+2}. \end{aligned} \quad (12)$$

Since we have N bins for the marginal distributions, this expression gives N moment conditions with which to estimate the M unknown parameters in ρ .

The Internet Appendix extends this logic to show that:

Proposition 2. Given data for the marginal distributions of spot (cumulative) inflation $f(\ln\pi_{t,t+j})$ and forward (year-on-year) inflation $f(\ln\pi_{t+j-1,t+j})$ for $j=1, \dots, J$, one can estimate the joint distribution $f(\ln\pi_{t,t+1}, \ln\pi_{t+1,t+2}, \dots, \ln\pi_{t+J-1,t+J})$ by estimating the M parameters in the ρ vector that satisfy the $(N-1)(J-1) \geq M$ conditions:

$$f(\ln\pi_{t,t+j}) = \int_{\Pi} \hat{C}(f(\ln\pi_{t,t+1}), \dots, f(\ln\pi_{t+J-1,t+J}), \rho) d\ln\pi_{t,t+1} \dots d\ln\pi_{t+J-1,t+J}. \quad (13)$$

The integration set Π is such that $\ln\pi_{t,t+1} + \dots + \ln\pi_{t+j-1,t+j} = \ln\pi_{t,t+j}$, for $j = 1, \dots, J$.

We can use these moments to estimate ρ , akin to GMM, although these are not moments of the distribution of the random variable, as is usual, but rather the distributions themselves.¹⁴

2.3.2 The estimates. Proposition 2 allows one to estimate a very general process for inflation, since there are 180 moments to use with our inflation distributions covering 10 years. We find that a simple process with a few parameters fits the data quite well. First, we use a multivariate Gaussian copula, which has a single parameter to model the codependence of inflation over two horizons. This does not impose normality for inflation, it simply assumes that the joint dependence of inflation over time resembles a normal distribution in the sense that if the marginals were normal, then the multivariate would be normal too. Second, we assume that inflation dynamics have a stationary and a nonstationary component, and that the nonstationary component is a random walk. This is a generalization of the common [Beveridge and Nelson \(1981\)](#) model for inflation. This parsimonious model has only 10 parameters, all of which are easy to interpret: the relative variance of the nonstationary and stationary components of inflation, and the 9 autocorrelations of the stationary part. We search for the parameters that minimize the equally weighted squared deviations of the moments in Proposition 2. The model fits the data quite well, in the sense that the discrepancies between the two sides of the moment conditions in Proposition 2 are small.

Our first finding is that, according to the options data, risk-adjusted inflation is generally close to being a stationary process. We estimate the variance of the innovations to the random walk component of risk-adjusted inflation to be small. Its ratio to the variance of the innovations of the transitory shocks is 0.09 in 2017 and lies above 0.25 only in 2010 and 2011. Thus, in risk-adjusted terms, long-run inflation expectations are well anchored. To understand what drives this result in the options data, one should note that if risk-adjusted inflation were nonstationary, then the variance of year-on-year inflation should tend to increase with the horizon. Yet, Figure 5 shows that, after 5 years, the distributions do not spread out in a significant way: the standard deviation lies between 1.7% and 1.9% for years 4–10.

The second interesting result is that the correlation coefficients do not noticeably fall with the horizon. Table 1 shows the estimated correlation parameters across maturities. Even as far as 8 years into the future, the

¹⁴ Our data on options contracts only go to 10 years, but the debt maturity goes all the way to 30 years. For inflation beyond 10 years, we extrapolate trends by assuming that the joint distribution is a stationary Markov process of order 9 with parameters given by the distribution from 1 to 10 years. The Internet Appendix discusses the details.

Table 1
Estimated correlation coefficients of year-on-year inflation, joint distribution

Maturity	Restricted model	Unrestricted (average)
1	0.11	0.24
2	-0.45	0.09
3	-0.38	0.11
4	0.04	0.15
5	0.30	0.14
6	0.07	0.05
7	0.16	0.16
8	0.42	0.30
9	0.08	0.54

Estimated correlation coefficients for year-on-year inflation between 2017 and 2017+j , where j is the maturity in each row.

autocorrelation is still 0.42. In spite of the evidence for stationarity, risk-adjusted inflation is a process with a long memory, where shocks persist for many years. Intuitively, Figure 4 shows that the distribution of cumulative average inflation does not shrink with the horizon, so risk-adjusted inflation must be expected to revert slowly to its long-run mean.¹⁵ The figure shows that the distribution of inflation only gradually shifts to the right with the horizon, and the expected reversion to the mean is quite slow.

The last column of Table 1 relaxes the restriction of the Beveridge-Nelson model, estimating instead $M = J(J - 1)/2 = 45$ parameters corresponding to the number of cross-horizon correlations for inflation. While the estimates move around more across horizons, the two general conclusions are similar. Interestingly, we find that the in-sample fit is only slightly better than the restricted model. An out-of-sample test of our estimation procedure uses only the maturities 1 to 10. The test compares the copula model's prediction for cumulative inflation in years 12 and 15, for which we have data. We cannot use these data for estimation purposes since the maturity of the traded year-on-year contracts does not extend that far, but we can use the data for this test. We find that the restricted model performs either noticeably better or similarly across all years.

3. Results: Debt Debasement Distribution and Stress Tests

We now combine the joint distribution for risk-adjusted inflation with the data on holders of nominal government debt. Different scenarios have probability weights disciplined by the data giving rise to a distribution of debt debasement that we now characterize.

¹⁵ Curiously, this is consistent with the findings of [Gürkaynak, Sack, and Swanson \(2005\)](#), who also use market expectations of the future in forward contracts to find that 10-year forward interest rates are more correlated with 1-year interest rates than what usual mean-reverting models of short-term rates would suggest.

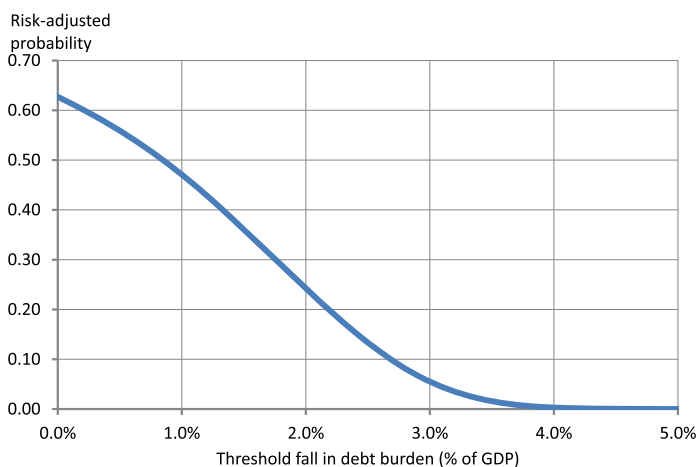


Figure 6
Probability that the fall in the value of debt exceeds a threshold

3.1 Value-at-risk measures of debasement

Proposition 1 provides a way to measure the probabilities of debt debasement of various sizes. Using our estimated risk-adjusted joint distribution for inflation, we draw a large number of paths and, ordering them by their impact on the real value of debt, we estimate probabilities that debt will fall by more than some threshold. This is akin to measures of risk-adjusted value-at-risk (VaR) for each time horizon, calculating potential losses over the entire life of the outstanding payments and then discounting back to today. These measures take into account all possible future scenarios for inflation as perceived today by market participants, including changes in the policy of the central bank, shocks to inflation beyond the control of monetary policy, and current as well as future changes to risk attitudes toward inflation.¹⁶

Figure 6 shows the probability that the fiscal burden will fall by more than a few percentage points of GDP, according to the risk-adjusted distribution for inflation. Strikingly, the numbers are quite small. The probability that debt falls by more than 4% of GDP is 0.3%. The probability of the real value of debt falling by at least 1% of GDP due to inflation is 47%, but anything more than a mere 3% of GDP has the small probability of 5.5%.

Table 2 presents percentiles of the distribution of debasement separately for each investor. Most of the gains are at the expense of foreigners. This happens not just because they hold more debt than domestics, but mostly because they

¹⁶ Because they include perceptions of different policy regimes, these measures are robust to the Lucas critique (Sims 1982). Moreover, Proposition 1 calls for risk-adjusted probabilities, which we use here. Veronesi and Zingales (2010) and Kocherlakota (2012) point out that risk-adjusted probabilities are appropriate when evaluating the effects of policies on welfare (see also Ait-Sahalia and Lo 2000; Jackwerth and Rubinstein 1996). We discuss in Section 3.4.1 how to interpret these and how our results extend to using physical probabilities.

Table 2
Percentiles of the distribution of the fall in the value of debt for bondholders

Percentile	Debt			
	Privately held (49.6%)	Domestic (17.2%)	Foreign (32.4%)	Central bank (12.6%)
90th	2.7%	1.4%	1.3%	1.1%
95th	3.0%	1.6%	1.5%	1.2%
99th	3.6%	1.8%	1.8%	1.4%

Each cell shows the cutoff in the reduction of the real present value of debt as a share of GDP, so that the risk-adjusted probability of a larger loss (borne by the agent in the column) is equal to one minus the percentile in the row.

hold longer maturity debt. Therefore, extreme situations where a succession of high realizations of annual inflation lead to large cumulative inflation affect foreigners more than domestics.¹⁷

The last column in the table shows the effect on a nonprivate holder of debt, the central bank. The Federal Reserve would potentially suffer large losses as a share of its portfolio. At the end of 2017, the central bank held mostly long-term bonds, which are more exposed to inflation risk.

3.2 Inflation counterfactuals

The magnitude of debt debasement in Figure 6 and Table 2 is surprisingly small. To understand why this is the case, we consider a series of counterfactuals that shift the inflation distributions.

For each counterfactual, we propose a new distribution $\hat{f}(\cdot)$, recalculate the real value of the debt using Equation (8), and subtract it from the market value of debt to obtain our estimate of the fall in real debt. This provides the stochastic equivalent of asking what would happen if inflation was $x\%$ higher but for a set of inflation paths. We discipline the experiments using the data by pinning down the shift x to be consistent with the plausible set of scenarios in our original distribution $f(\cdot)$. Importantly, these are shifts in the risk-adjusted distribution for inflation, not changes in inflation per se. They are useful to understand how different properties of inflation affect the debt debasement effect, but they do not cleanly correspond to actual policy changes.

The first experiment shifts horizontally the marginal distributions for annual inflation at every maturity so that the new median is at the old 90th percentile. One can think of this experiment as capturing an announcement that the inflation target of the Fed is now expected to be higher, as suggested by Blanchard, Dell’Ariccia, and Mauro (2010). The second experiment instead

¹⁷ How much foreigners ultimately lose also depends on how much the exchange rate devalues with the increase in inflation and on whether there are taxes on capital flows. We are not measuring foreigners’ losses. Rather, we are measuring how much the fiscal authority would gain at their expense. Our measure does not depend on whether the holder is domestic or foreign, or what happens to the exchange rate, even if these factors would be important for foreign debtholders.

sets the density below the 90th percentile to zero. Probabilities are then scaled up proportionately to integrate to one. This corresponds to a commitment that inflation will be higher for sure in the future. Only inflation realizations at the right tail of the current distribution become possible. At the same time, because there is no shift to the right as in the first case, very high levels of inflation are also not that likely.

The next three stress tests consider the effect of uncertainty. The third experiment again shifts the marginal density, so the new median is the old 90th percentile, but now this is accomplished by scaling inflation proportionately at every maturity. It is often said that higher average inflation comes with more variable inflation, and this experiment tries to capture this possibility. In the other direction, in the fourth counterfactual, we assume that year-on-year inflation is exactly equal to average inflation in the first stress test. This shows what would happen if inflation became deterministic, so we can understand better the effect of volatility. The fifth counterfactual studies a shift in inflation that is partially expected. We assume that after an initial unexpected jump of inflation upward, the distribution of inflation looking forward is equal to the conditional distribution that we have estimated. Therefore, whereas in the previous experiments all of the changes at all maturities were unexpected, now only the change in the first year catches agents by surprise, but they adjust their expectations right after.

So far, the experiments assumed that the risk-adjusted inflation distribution would change immediately and permanently. The sixth stress test considers instead a temporary increase in inflation, with the distribution for year-on-year inflation shifting rightward so the new median is at the 90th percentile the next year, but only at the 80th percentile the year after, and so on, so that for maturities above 5 years there is no change. The seventh case considers a gradual increase, with the 1-year inflation distribution unchanged, while the 2-year shifts horizontally so the new median is at the old 60th percentile, and so on until the fifth year, after which we have the same permanent shift as in the first case.

Table 3 shows the risk-adjusted harmonic mean of inflation for both the baseline and each of the counterfactuals at different maturities. Table 4 reports the effect of each of the counterfactuals on the real value of the debt. All of the experiments have a relatively small impact on the real value of the debt. Larger inflation, as in experiment 2, lowers the fiscal burden by 3.9%, the highest number in the table. More uncertainty lowers the effectiveness of inflation at debasing debt, because the real value of future nominal payments is convex in inflation, so uncertainty raises the payments' value and lowers the benefits of raising inflation. If investors adjust their expectations after 1 year of surprise inflation, as in case 5, the estimates are significantly smaller. Moreover, the last two rows in Table 4 show that more reasonable time paths for inflation than a sudden and permanent change also significantly cut the effect of inflation on debt. Finally, confirming our previous results, again foreigners absorb a large

Table 3
Expected adjusted average annual inflation for different counterfactuals

Distribution for inflation	Horizon			
	1-year	3-year	5-year	30-year
Baseline	2.0%	2.2%	2.2%	1.8%
Counterfactuals				
1. Permanently higher	2.9%	3.2%	3.4%	2.9%
2. Right tail only	3.6%	3.5%	3.1%	5.9%
3. Higher and more variable	2.9%	3.2%	3.4%	2.5%
4. Higher for sure	2.9%	3.2%	3.4%	2.9%
5. Partially anticipated	3.0%	2.4%	2.2%	2.1%
6. Temporary increase	2.9%	2.8%	2.6%	1.9%
7. Gradual increase	2.0%	2.4%	2.8%	2.8%

Each cell reports $1/E(n/\pi_{0,j})$, that is, the harmonic mean of inflation until horizon j .

Table 4
Counterfactual impact of higher inflation on the real present value of debt

Inflation counterfactual	Debt			
	Privately held (49.6%)	Domestic (17.2%)	Foreign (32.4%)	Central bank (12.6%)
1. Permanently higher	2.3%	0.9%	1.4%	0.8%
2. Right tail only	3.9%	1.9%	2.0%	1.5%
3. Higher and more variable	1.6%	0.5%	1.1%	0.5%
4. Higher for sure	2.8%	1.2%	1.6%	0.9%
5. Partially anticipated	0.4%	0.2%	0.2%	0.1%
6. Temporary increase	0.6%	0.2%	0.4%	0.2%
7. Gradual increase	1.5%	0.7%	0.8%	0.6%

Each cell shows the fall in the real present value of debt as a ratio of GDP.

share of the losses, and the Federal Reserve has likewise large losses relative to the size of its portfolio.

The reason for these small effects is due to the interaction of two features of the inflation data. First, large increases in inflation are unlikely. Even in the most extreme case in Table 3, expected inflation in 1 year never increases by more than 1.6% relative to the baseline. From the perspective of actual market-based distributions, anything larger than this seems unreasonable. But from the perspective of debt valuation, such modest increases in inflation are incapable of generating large debt debasement. Raising the debt debasement effect to 10% would take shifting all year-on-year distributions of inflation in experiment 1 by 6.5% to the right. For the first year, the new median would be 8.5%, and the shift is equal to five and a half standard deviations of the average volatility for the first 3 years. Therefore, in that case, the new and old distributions for inflation would have close to zero overlap, making this scenario literally incredible.

Second, as Table 1 shows, the persistence of inflation is quite low in the first few years, closer to the temporary increase in our sixth counterfactual. While shocks then persist for many years, the increase in the horizon implies that these effects are partially anticipated as in our fifth counterfactual. As the two counterfactuals combined show, this reduces the effect on the real debt. This

Table 5
Counterfactual impact of higher inflation with different maturity distributions

Inflation counterfactual	Including only debt of maturity up to:		
	1 year	5 years	All
1. Permanently higher	0.1%	0.7%	2.3%
2. Right tail only	0.1%	0.8%	3.9%
3. Higher and more variable	0.1%	0.6%	1.6%
4. Higher for sure	0.1%	0.7%	2.8%
5. Partially anticipated	0.1%	0.1%	0.4%
6. Temporary increase	0.1%	0.4%	0.6%
7. Gradual increase	0.0%	0.2%	1.5%

Each cell shows the fall in the real present value of debt as a ratio of GDP.

is due to the interaction of inflation with the debt maturity, which we discuss next.

3.3 Drivers of debasement

In this section, we examine variation in debasement probabilities over time and identify its key determinants: debt maturity, inflation volatility, and inflation persistence.

3.3.1 Debt maturity. Starting with debt maturity, over several years, an additional percentage point of inflation can accumulate to a large effect on debt debasement. If all privately held debt was at long maturities, inflation might significantly reduce its real value. Table 5 investigates the effect of maturity on our estimates by considering the effect of higher inflation only on debt with maturity below 1 year, or only below 5 years. The numbers are significantly lower than when all the debt is included. Moreover, even though in 2017 almost three quarters of the market value of debt has a maturity below 5 years, the debt debasement effect is well below 0.75 of our estimates when all the debt is included. Combining the fact that the maturity of the debt is low, and that the serial correlation of inflation in the first few years is low, gives a low effect on debt debasement. Inflation persists into the more distant future, but by then there is not much debt to debase. In the table, the temporary increase (counterfactual 6) has a larger effect than the gradual (counterfactual 7) on debt with maturity below 5 years, since the former has a greater impact on the short end of the maturity structure. Thus, *ceteris paribus*, higher debt levels and a longer maturity will result in more debasement.

We repeat the work of Section 2 building the distribution of privately held debt across maturities, but now extending it back to 2009. Figure 7 shows the cumulative distributions over time. As is well-known, the distribution has shifted up as the U.S. government issued more debt. Less appreciated, privately held debt has become more concentrated at short maturities over time.

Figure 8 plots debt debasement value-at-risk levels from 2009 to 2017. There is substantial variation. In 2009, debasement of more than 2% was very unlikely; in 2011, only all the way to a debasement of more than 6% was now unlikely.

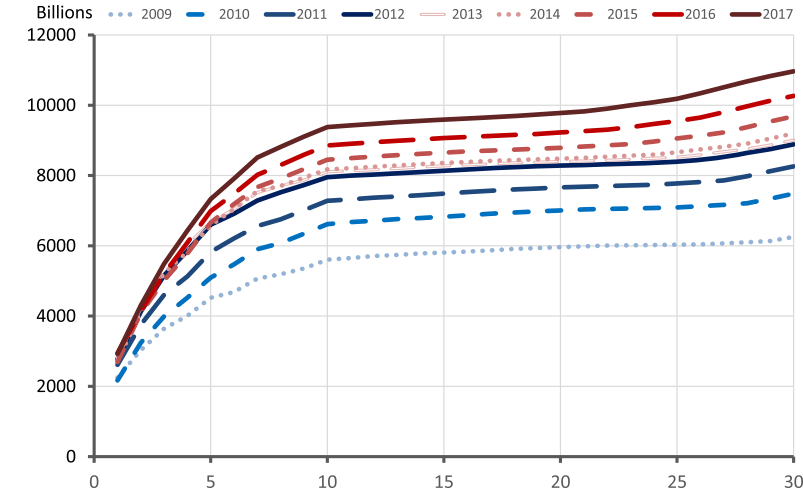


Figure 7
Debt maturity distribution, 2009–2017

Importantly, debasement does not increase primarily with the debt burden. If it did, we would see a monotonic increase in the value-at-risk probability of debasement, reflecting the monotonic increase in the debt burden. Instead, we see that using only debt maturity or variables calculated from debt maturity provides an incomplete picture.

3.3.2 A simple, but misleading, rule of thumb. From Proposition 1, we know that empirically evaluating inflationary debasement requires the simulation of future inflation paths and knowledge of the debt maturity distribution. A related measure is the *debasement sensitivity*, which uses only information on debt maturity and discount rates. This quantifies how much the public debt changes in response to a marginal change in inflation. A common rule of thumb used to evaluate the effect of $\varepsilon\%$ more inflation for sure and forever on the public debt multiplies the debt-to-GDP ratio by the average maturity of the debt.¹⁸ With the headline debt-to-GDP ratio, 103.6% in 2017, and an average maturity of 5.8 years, this common rule of thumb would suggest a sensitivity of 6.0%. Our estimates instead suggest that the probability of such a large sensitivity to an inflation shock today is virtually zero.

¹⁸ This rule of thumb comes from making the very strong assumptions that (a) all debt is nominal, (b) the maturity distribution is exponential with expected maturity m , (c) inflation is a constant π , (d) the nominal interest rate is constant and approximately zero, and (e) the sensitivity is with respect to ε permanently higher inflation. Then the fiscal burden is $\int_0^\infty (B_t/P_t)e^{-rt} dt = \int_0^\infty (B_0(e^{-t/m})) / (P_0 e^{\pi t}) e^{-rt} dt = (B_0/P_0)(1/m - r - \pi)^{-1}$. A first-order Taylor approximation for an ε increase in π at a point at which $r + \pi = 0$ then gives a debasement sensitivity of $(B_0/P_0)m\varepsilon$.

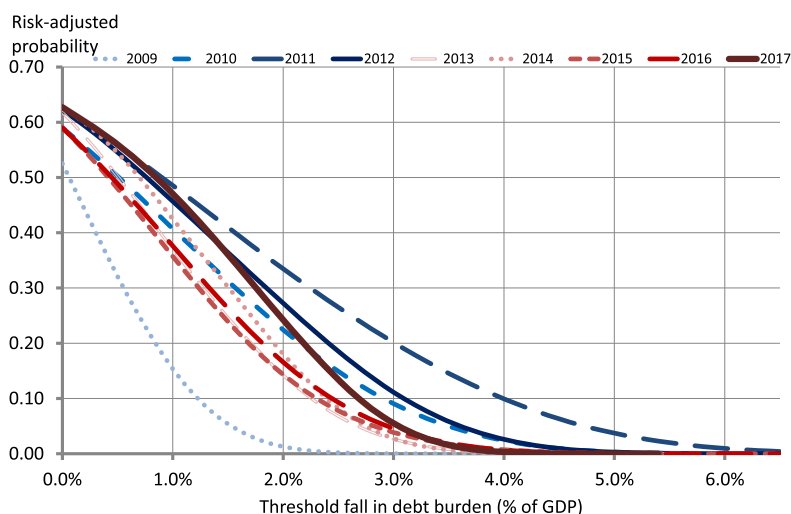


Figure 8
Probability that the fall in the value of debt exceeds a threshold, 2009-2017

The rule of thumb is misleading and incomplete for a few reasons. First, because a sensitivity is not a distribution. At a minimum, the volatility of inflation must also play a role; without such an input it is not possible to move from sensitivity to distribution or value at risk. Simply put, if inflation volatility is zero, debasement sensitivity plays no role. Debasement will always be equal to zero. A back-of-the-envelope calculation might assume that inflation is normally distributed and has a standard deviation of 1% (its historical average in our sample). Then $\varepsilon = 2.33\%$ for the 1% VaR. All combined, we obtain a predicted 1% VaR of 14.1%, much higher than the actual 3.6%.

Second, the estimate above does not use the correct number for debt, which is the privately held one: 49.6% of GDP. By itself, this cuts the estimate in half.

Third, average maturity is not the correct measure to capture the sensitivity of debt to inflation. A more accurate, but still highly imperfect, rule of thumb measures the response to a ε parallel shift of the inflation distributions at all horizons, using the privately held debt, and using information on future inflation distributions. A first-order Taylor approximation on the size of the shift implies a debasement sensitivity of debt value times ε times the Fisher and Weil (1971) duration of the bond portfolio. The formula is simple but calculating the Fisher-Weil duration requires knowing the entire distribution of privately held bond holdings $\{B_0^i\}_i$ and associated prices. Therefore, calculating it requires the same debt maturity structure data needed for the accurate debasement distributions that we calculate.

Fourth, and even after all the adjustments above, a sensitivity measure is not forward looking. Future inflation, not the historical distribution of

inflation, is relevant for debasement. Using the detailed knowledge of future inflation distributions we have constructed, we can calculate the mean and the 99th percentile of the inflation distribution in each of the next 3 years. This calculation requires almost the same data that we use for our exact calculations. The average difference is a good measure of the parallel shift needed. Using this approach, we get a 1% VaR of 7.0%, based on an inflation shift of 2.7%, a Fisher-Weil duration of 5.1% and the debt-to-GDP ratio of 49.6%. The number is still quite a bit higher than our accurate estimate of 3.6%.

That is because of a fifth and final problem. This measure neglects the joint distribution of inflation over different horizons. It assumes a permanent shift of inflation, that is, a parallel shift of inflation at all maturities. It thus ignores the stochastic properties of inflation, especially its persistence, simply assuming an annual serial correlation of inflation equal to one. In the data, this is 0.11 or 0.24 (restricted or unrestricted model), making parallel shifts extremely unlikely. In our 2009–2017 data on inflation distributions, we never see such parallel shifts. Assuming perfect correlation, as a calculation based on duration does, will result in an overestimate of debasement.

3.3.3 A more accurate rule of thumb. Figure 8 shows value-at-risk distributions. Another way to present the estimates shows instead the variation in value-at-risk levels over time at one level, say 1%. Figure 9 shows these: the debasement cutoffs, over time such that, with a risk-adjusted probability of 1%, debasement will exceed this level. In 2009, it is equal to 2.1%, it then increases to 6% in 2011, and it stabilizes a little below 4% from 2015 to 2017.

The picture also shows a measure based on the Fisher-Weil duration, which assumes a 1% parallel shift of the entire inflation term structure. This measure, by contrast, increases over the sample period, missing the initial strong rise, subsequent fall, and following stabilization of the debasement VaR. The R^2 of a regression of actual VaR on Fisher-Weil predicted VaR is a modest 14%. That rule of thumb is highly inaccurate not only in the level but also in the change of the debasement distribution.

To bring persistence into the picture, we use the average 1-year correlation from our restricted model estimates. There is substantial variation in this measure over the years, which sometimes is close to zero and even negative. In the figure, we add one to the average correlation and divide by 100 to easily plot the numbers in the same graph; for example, in 2017 the average $(t, t+1)$, $(t+1, t+2)$, ..., $(t+9, t+10)$, 1-year correlation is 31%, which in the figure is shown as 1.3%. We then multiply our measure of persistence by the Fisher-Weil duration. This new rule of thumb, that combines information on the debt distribution with information on the perceived persistence of inflation, now explains 46% of the actual measure, according to the R^2 , more than three times what duration itself could explain. Looking at individual years, from 2013 to 2017 duration steadily increases but persistence first declines and then stabilizes; as a result, the debasement VaR is relatively stable.

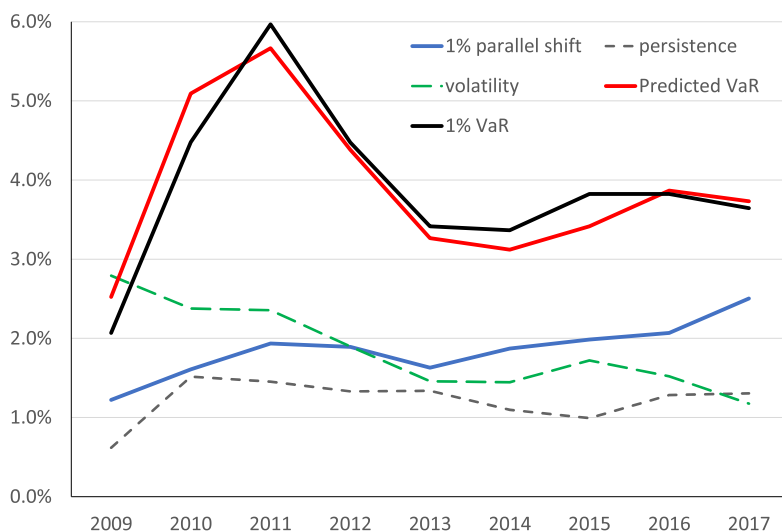


Figure 9
Debasement value at risk and debasement drivers, 2009–2017

Finally, we add the volatility of inflation shocks. We calculate the average annual volatility of year-on-year inflation for years one to three. Multiplying volatility with persistence and duration completes the picture. The predicted VaR, plotted in the figure, now leads to an R^2 of 89%, another substantial increase. In the sample, there has been a large decline in inflation volatility, from 2.8% in 2009 to 1.2% in 2017. This further explains why, in spite of rising debt levels, debt debasement via inflation has remained unlikely through the sample.

In short, our new rule of thumb, the product of a duration-based measure, persistence, and volatility of inflation, captures the variation in debasement probability well. It shows that the three necessary conditions for the presence of significant debasement probabilities are (a) a long maturity structure, (b) a reasonable probability of higher-than-expected inflation, and (c) reasonable levels of inflation persistence.

3.4 Interpreting the probabilities

The probabilities we use are all in risk-adjusted space. This is exactly what Proposition 1 calls for. From a policy perspective, [Veronesi and Zingales \(2010\)](#) and [Kocherlakota \(2012\)](#) make the case that governments should weight different inflation policies according to their effect on welfare. Because they adjust for the growth of marginal utility, risk-adjusted probabilities are appropriate when considering policy. Moreover, the debasement question is a valuation question, so it is appropriate to use valuation weights.

To better interpret our results, one should understand the relationship between risk-adjusted and physical probabilities. First, all probabilities are conditional on information available at the time of measurement, for example, 2017. While we look at paths for inflation into the future, no new information is revealed by the joint distribution of 2018 and 2019 that affects the joint distributions between 2021 and 2022; both are measured given information in 2017.

Second, the relevant multiperiod risk-adjusted expectation of inflation depends on the product of actual inflation and state-dependent risk compensation, which includes anything that changes marginal utility for different realizations of inflation. Our estimates have no way of distinguishing whether the risk-adjusted correlations between inflation at different horizons are due to serial correlation of actual inflation or of the stochastic discount factor. They measure the debt debasement distribution according to the marginal distribution of risk-adjusted inflation, which already includes the correlation of inflation with other factors that affect marginal utility.

A concrete example makes these two points clearer. When we evaluate the expectation of inflation between dates 0 and 2 using the estimated risk-neutral densities, we are computing

$$\begin{aligned}\mathbb{E}\left[\frac{m_{0,2}}{\pi_{0,2}}\right] &= \mathbb{E}\left[\left(\frac{m_{0,1}}{\pi_{0,1}}\right)\left(\frac{m_{1,2}}{\pi_{1,2}}\right)\left(\frac{m_{0,2}}{m_{0,1}m_{1,2}}\right)\right] \\ &= \mathbb{E}\left(\frac{m_{0,1}}{\pi_{0,1}}\right)\mathbb{E}\left(\frac{m_{1,2}}{\pi_{1,2}}\right) + cov\left[\frac{m_{0,1}}{\pi_{0,1}}, \frac{m_{1,2}}{\pi_{1,2}}\right].\end{aligned}\quad (14)$$

Our first point is that all expectations are taken conditional on information at time 0. Therefore, the last term in brackets in the first line is exactly equal to one, no matter what states are realized at dates 1 and 2. Risk-adjustments (distributions of the stochastic discount factor across future states) or real interest rates (conditional means of the stochastic discount factor) may change over time, but that is taken into account by the date 0 expectation. Our second point is that the covariance term can depend on the covariance between $\pi_{0,1}$ and $\pi_{1,2}$, as well as that between $m_{0,1}$ and $m_{1,2}$, and between the two pairs. It takes a model for $m_{0,t}$ and for $\pi_{0,t}$ to separate them.

3.4.1 Physical probabilities. Even though Proposition 1 calls for the use of risk-adjusted probabilities, one might wonder how to translate them to physical probabilities. In that case, to aggregate the effect of different inflation realizations, one would require the joint distribution of future inflation paths and future realizations of the stochastic discount factor. In general, this requires a full general-equilibrium model linking real payoffs and inflation. The Internet Appendix formally shows that our risk-adjusted probabilities likely are either very similar to or provide an upper bound for the physical probabilities, by considering three broad classes of models, which we now verbally describe.

One class assumes that the classical dichotomy holds. This property states that inflation realizations are uncorrelated with risk preferences and discount

rates. In this case, drawing paths only from the marginal distribution of inflation, like we did, delivers the right answer as the physical and risk-neutral distributions of debt debasement coincide.

A second class of models has the property that inflation has a risk premium, and this may vary over time or with shocks, but the risk premium is additive; in this case, to move from one distribution to the other, one needs to adjust the drift of the stochastic process for inflation. Under this approach, the physical probabilities of debasement are still exactly the same as the ones reported in Table 2 and Figure 6. Since the risk-adjustment is additive, the relative ordering of simulated paths for inflation does not change. In other words, the value-at-risk depends on the probability of inflation tail events, and it does not change under this approach.¹⁹

A third class of models has a price of inflation risk, but one that must fit the evidence at the tail of the distribution. Kitsul and Wright (2013) estimate the price of risk in the past using a process for actual inflation fit to the historical data and a flexible parametric pricing kernel, and assume the same model applies to the future. Relevant for our calculations are their estimate at the right tail of the distribution. They find that states of the world with high realizations of inflation have higher risk prices, or risk-adjusted probabilities above the physical ones. In this case, our VaR estimates provide an upper bound for debasement under the physical measure, since they use the higher risk-neutral mass at the tails. Extreme debasement based on physical probabilities is even less likely than what we report.

4. Beyond Inflating the Debt: Growth, Financial Repression, and Links to Theory

Our estimates measure debt debasement *from inflation alone*. However, inflation often comes with changes in real activity or with financial repression. This section shows that adding these into the analysis can lead to a coincidence of high inflation and large reductions in the debt. Whether one should interpret these combined effects as “inflating away the debt” is unclear, so, at the end of the section, we describe how our estimates can be used to inform models of monetary and fiscal policy.

4.1 Variation in GDP

The goal of this paper is to measure inflation-driven debt debasement. Debt-to-GDP may fall for other reasons, notably if GDP rises. Of clear interest is whether states of high inflation are also states of high growth, further lowering the real value of debt compared to GDP. While this does not answer the question

¹⁹ In addition, the inflation risk premium estimated with this approach tends to be small: Fleckenstein, Longstaff, and Lustig (2017) estimate it to be only between 10 and 20 bps. The two distributions are thus close to each other.

of inflation-induced debt debasement directly posed in the introduction, it complements the analysis.

Accurately measuring this perceived risk-adjusted correlation between inflation and GDP would require having traded options on the joint outcomes for inflation and GDP. These do not exist, however. Instead, we proceed by providing estimates based on a sample of historical data and the (strong) assumption that the past empirical correlation of inflation and output roughly estimates the expected risk-adjusted correlation.

To calculate it, first we estimate a structural vector autoregression on inflation, GDP growth, and nominal interests rate using annual data and one lag (results are robust to more lags). At any horizon, this gives the response of inflation as the weighted sum of the three structural shocks. Then, taking a draw from the risk-adjusted inflation distribution, just as we do when constructing our VaR levels, we can back out the underlying vector of shocks that exactly match this draw at three different horizons (we do so at horizons 1, 3, and 5, but the results are robust to other choices). Given this vector of shocks, the vector autoregression then provides the matching impulse response for the corresponding draw of output, that has embedded in it the correlation between inflation and output at all horizons. Combining these output draws with the original inflation draws, we calculate the distribution of debasement of the debt. For GDP growth to materialize, we need to examine the distribution of debt-to-GDP levels in the future, not the present, and so we calculate maturity specific burdens of debt that is currently outstanding relative to future GDP levels.

To take away the effect of the correlation between output and inflation, we next identify one structural shock in the vector autoregression that has the properties that it has no contemporaneous effect on either interest rates or inflation. This is then a pure output shock, accounting for output volatility, but not its correlation with inflation. Again, for each inflation draw, this identification restriction allows us to back out the pure output shock and thus the corresponding output draw.

A third and final exercise considers shutting off both output's correlation and volatility, and simply adds a 2% growth trend to output paths for each inflation draw.

Figure 10 shows the maturity structure of 5% VaR debt-to-GDP reductions for these three cases. Note that adding the whole term structure is similar, but not the same, as calculating the 5% VaR directly.²⁰

When the GDP growth path is fixed at 2%, so that only inflation is stochastic, the maturity structure of inflation-induced debasement reflects the drivers of our main results. At short maturities, because there is a lot of debt, and at long maturities because inflation may be more extreme at longer horizons.

²⁰ Our main calculation compares the discounted debt value relative to GDP today, whereas in this exercise we do not discount and instead compare today's debt to future GDP. Our main calculation also finds the overall 95th percentile, which is not the same as the sum of the period by period 95th percentile levels.

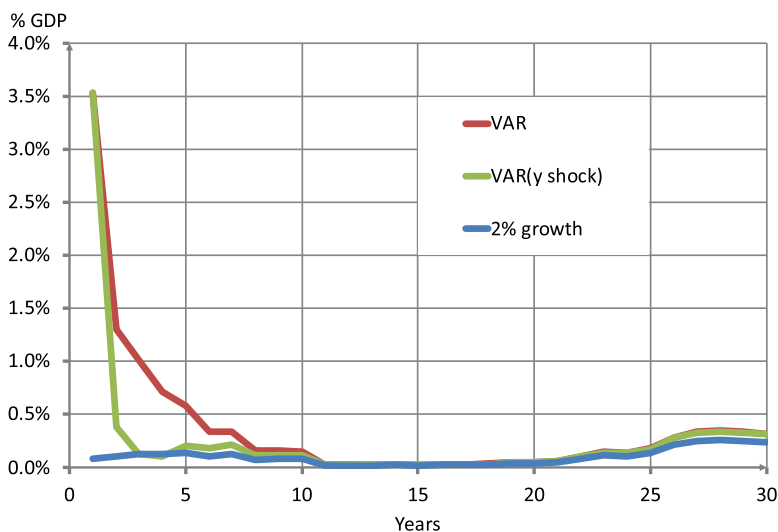


Figure 10
Maturity structure of debt-to-GDP reductions when simulating inflation and GDP

Allowing for GDP growth volatility, but no correlation with inflation, evaluates by how much extreme growth spurts can reduce the debt-to-GDP level for each maturity. Inflation is still allowed to vary, so inflation, growth, or both can spur large debt-to-GDP reductions. The effects are unsurprisingly much larger, since the volatility of GDP through the vector autoregression means that there are extreme cases on the tails, where there are significant growth spurts in the future. Because there is a lot of debt at short maturities, debasement resulting from extreme shocks is large. The third case allows for correlation between shocks to inflation and GDP. Now, the extreme growth spurts can be times of high inflation, given the positive correlation of GDP growth and inflation shocks, especially at short horizons. As a result, the increased extreme debasement relative to assuming no correlation is larger at shorter horizons. Once the horizon rises, the Phillips correlation implicit in the vector autoregression is smaller, and there is less debt, so the difference becomes negligible.

4.2 Isolating the sources of debt-to-GDP changes over 1 year

The debt-to-GDP ratio may decline because of inflation and/or growth, as well as if the real interest rate used to discount future debt payments changes. [Hall and Sargent \(2011\)](#) find that these were all important for the historical evolution of the debt-to-GDP ratio. This section offers a related exercise, but one that takes a forward-looking approach to decompose surprise changes in the debt-to-GDP ratio into surprises on the real return of the debt and the output growth rate every year over our sample period.

Adapting the approach of [Campbell and Shiller \(1988\)](#) and [Campbell \(1991\)](#), start with the unexpected real return of government debt. For expositional purposes, assume that there is a single zero-coupon bond with maturity T . Over time, the cumulative nominal return of the bond is equal to the total discount at which it is currently trading. Breaking the return up into one-period returns and noting that the final nominal payoff of the bond is fixed, the surprise one-period real return is

$$r_{t+1} - \mathbb{E}_t(r_{t+1}) = -(\pi_{t+1} - \mathbb{E}_t \pi_{t+1}) - (\mathbb{E}_{t+1} - \mathbb{E}_t) \sum_{s=t+2}^T (r_s + \pi_s). \quad (15)$$

Lowercase letters denote logs: r_{t+1} is the real return on the debt; π_{t+1} is the realized inflation rate; i_s is the one-period nominal interest rate; and, by assumption, $\sum_{s=t+2}^T (r_s + \pi_s) = \sum_{s=t+2}^T i_s$. The real return on the debt is high if inflation is low or if real rates or expected future inflation have declined. The surprise percent change in the debt-to-GDP ratio is then equal to:

$$(\mathbb{E}_{t+1} - \mathbb{E}_t)(w_{t+1} - y_{t+1}) = (\mathbb{E}_{t+1} - \mathbb{E}_t)(r_{t+1} - g_{t+1}). \quad (16)$$

To obtain the unexpected change of the debt-to-GDP ratio from one period to the next, we scale this expression by the current debt-to-GDP. All combined, note that this is the change in the real value of currently outstanding debt, without including the next period's deficit.²¹

Table 6 reports the 1-year surprise changes for our sample period. We also report one-period surprise inflation and GDP growth, which are the direct inputs into the calculation. Changes in the term structure of future inflation and real rates are the remaining components; for these we report changes at the 10-year point. From the table, we can see that changes in expected inflation and real rates are the main drivers of debt-to-GDP changes. The main reason is that the 1-year horizon, at which inflation and GDP growth are measured, is short relative to the average maturity of the debt. For example, in 2011, despite a large positive inflation surprise, the debt-to-GDP ratio went up, driven by the reduction in expected future inflation and real rates. The correlation of the overall change is equal to 47% with future inflation, and 85% with real interest rates. Other times, 1-year shocks are important. For instance, in 2015, when changes in rates were small, the debt-to-GDP ratio was reduced by positive inflation and growth surprises.

4.3 Financial repression

It may seem surprising that our calculations so far could be silent on the hotly disputed topic of whether inflation has real effects. One might think that if

²¹ Doing so would require a forecast of the next period's deficit since we are interested in the surprise change of the debt-to-GDP ratio. Since the purpose of the current exercise is to understand the effects of inflation, GDP growth, and revaluation, we choose not to model the deficit.

Table 6
Decomposing 1-year surprise changes in the debt-to-GDP ratio

	2010	2011	2012	2013	2014	2015	2016	2017
Surprise debt/GDP change	0.83%	1.89%	0.39%	−0.99%	1.43%	−0.62%	0.05%	0.31%
Inflation surprise	0.30%	1.43%	0.08%	0.13%	−0.91%	0.96%	0.91%	−0.04%
GDP growth surprise	0.66%	−0.32%	0.27%	−0.12%	0.56%	0.95%	−0.32%	0.47%
Change in expected inflation	−0.05%	−0.29%	0.39%	−0.18%	−0.57%	0.10%	0.33%	−0.03%
Change in real rates	−0.60%	−1.27%	−0.55%	1.56%	−0.40%	0.03%	−0.16%	−0.08%

Surprise change in debt/GDP is the one-year surprise real return of current debt adjusted for real GDP growth and scaled by current debt/GDP, inflation surprise is the one-year surprise inflation, and change in expected inflation and change in real rates are the one-year changes in the 10-year expected inflation and real interest rates.

inflation lowers real interest rates, then because the government needs to pay less to roll over its debt, the fiscal burden will be smaller. Yet, if the interest rate is lower, this also means that investors discount the future debt by less. By our assumption of no arbitrage opportunities, the real interest paid on the government bonds and the real interest that private agents use to discount the future are the same. Therefore, in the present value of the fiscal burden, these two effects *exactly* offset, whatever is the impact of inflation on the real interest rate. To be clear, this does not mean that changes in *ex post* realized real interest rates do not affect the value of the debt at intermediate dates prior to maturity (of course they do), but rather that for a forward-looking measure of the market value of the debt, changes in *ex ante* expected real rates are exactly offset by changes in the real discounting factors for these payoffs.

Financial repression is a way to drive a wedge between these two interest rates. This wedge works like a tax on the returns of government debt and as such provides a source of revenue that reduces the fiscal burden. The literature on financial repression, which dates back at least to [McKinnon \(1973\)](#), offers many examples of how this tax is collected and enforced, through channels like caps on interest rates, direct lending to the government by captive domestic savers, or financial regulation, among others. In theory, this would show up as a factor $(1 - \tau_r)$ in each of the terms in our formula in Proposition 1. But, at this general level, we cannot say more empirically about its size or how it varies with inflation. Next, we consider a specific form of financial repression.

4.3.1 Repression as financial regulation. [Reinhart and Sbrancia \(2015\)](#) discuss how many developed countries, including the United States, used a combination of caps on the interest rates on government bonds and inflation between 1945 and 1980 to pay for the World War II debt. One particular way in which this is achieved is by forcing the holders of outstanding debt to roll it over for “special” debt that sells for a higher price (or pays a lower return) than the market price for identical private securities. This is achieved, for instance, by forcing banks to accept this special debt and hold it under the guise of financial regulation and stability. An extreme case of this hidden financial repression is to require banks to hold zero-interest reserves at the central bank. Effectively,

one type of government liability that pays market interest is replaced by another type that pays no interest.

To model this formally, one should assume, for simplicity, that all debt is nominal and has a maturity of one period, and that the holders of maturing bonds are forced to take special bonds as payment that promise to pay \tilde{B}_{t+1} in the next period and sell today for the price of \tilde{H}_t . The price of the bond is higher than the market price for nominal bonds, H_t , capturing repression, and below one, capturing the zero lower bound on interest rates. When it is equal to one, we have extreme financial repression with the special bonds being zero-interest reserves. The law of motion at any date after 0 for debt now becomes

$$W_t = s_t + \left(\frac{\tilde{H}_t - H_t}{P_t} \right) \tilde{B}_{t+1} + W_{t+1} \frac{H_t P_{t+1}}{P_t}. \quad (17)$$

Because $\tilde{H}_t > H_t$, this expression makes clear that financial repression works like a source of tax revenue. Because $B_t^0 = \tilde{H}_t \tilde{B}_{t+1}$, this revenue subtracts from the real value of outstanding debt just like a tax on its holders would. Similar algebra as the one that led to Proposition 1 shows that the debt burden now is

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} m_{0,t} \left(\frac{H_t}{\tilde{H}_t} \right) \left(\frac{B_t^0}{P_t} \right) \right] \leq \sum_{t=0}^{\infty} \frac{B_t^0}{P_0} \mathbb{E}_0 \left[\frac{m_{0,t+1}}{\pi_{0,t+1}} \right]. \quad (18)$$

Since $\tilde{H}_t > H_t$, the debt burden is lower the higher is financial repression. The inequality becomes an equality in the case of extreme financial repression ($\tilde{H}_t = 1$), where the government rolls over its past debt through zero-interest required reserves. This expression shows that the effect of extreme financial repression is essentially equivalent to delaying all payments on the debt for 1 year at a zero interest rate or shifting the maturity structure by 1 year.

Generalizing the previous argument to have financial repression for N periods, we obtain a new version of Proposition 1:

Proposition 3. With financial repression for N periods, the debt burden is still equal to $\sum_{t=0}^{\infty} \omega_t^r B_0^t$ but now the weights are

$$\omega_t^r = R_t^{-1} \int \left(\frac{f(\pi_{0,t})}{\pi_{0,t}} \right) \left(\frac{H_t^N}{\tilde{H}_t^N} \right) d\pi_{0,t} \leq \omega_{t+N}. \quad (19)$$

With extreme financial repression ($\tilde{H}_t^N = 1$), the inequality becomes an equality.

The weights in the formula for the debt burden are lower, as long as nominal interest rates are positive. Higher inflation is now more effective at lowering the real value of the debt. Inflation not only debases the debt but also lowers the real return on the zero-interest reserves.

Table 7
The effect of inflation with financial repression

Duration of repression	Repression	Higher inflation	Total
1 year	1.4%	2.7%	4.1%
5 year	7.5%	4.0%	11.5%
10 year	13.8%	4.9%	18.7%

Each cell shows the fall in the real value of debt as a ratio of GDP as a result first of repression and then of higher inflation under experiment 1. The last column is the sum of the two previous ones.

4.3.2 Estimating the joint effectiveness of inflation and repression.

Table 7 shows the effect of extreme financial repression. To be clear, these are counterfactual measures, not debt debasement distributions. If agents had anticipated repression, any anticipation would have affected the debt outstanding in the first place. We display the counterfactual measures, not because they substitute for or qualify our debt debasement distribution, but because they help to demonstrate how a government may try to inflate the debt.

The second column shows the sole direct effects of financial repression, which can be calculated by using the information in the yield curve. Even with the low nominal interest rates in 2017, repression for 10 years would wipe out more than one quarter of the privately held debt.

Column 4 is our main interest since it shows how repression interacts with inflation. It conducts our experiment 1 within the financially repressed economy, so it measures the sole effect of the inflation distribution shifting to the right. The effect of inflation is much higher than before. This confirms our conclusion of the previous section that a longer maturity of current debt makes inflation effective at lowering the real value of debt. The last column shows the joint effect, at date 0, of both imposing financial repression and shifting the distribution of inflation to the right. The reduction of the real value of debt is substantial.

4.4 Using our estimates in different classes of theoretical models

We have calculated estimates of debt debasement motivated by the high and growing debt levels in many countries and the question of what role inflation can play in reducing debt levels. In addition, our estimates also can be used to measure magnitudes of variables in several classes of theoretical models, including models of monetary policy, fiscal theory of the price level, and fiscal policy. This section expands on these uses.

4.4.1 The fiscal backing of inflation-targeting central banks. The intertemporal budget constraint (omitting indexed debt because it plays no role) is

$$W_0 = \mathbb{E} \left[\sum_{t=0}^{\infty} m_{0,t} s_t \right]. \tag{20}$$

This equation, which we have already seen as Equation (5) only assumes no arbitrage and the absence of Ponzi schemes. It therefore holds in most models of monetary and fiscal policy.

In traditional new Keynesian models, an independent, inflation-targeting, central bank sets a policy tool (like the interest rate) to achieve a desired level of inflation. That desired level is the sum of an inflation target and a term proportional to the output gap, a policy known as flexible inflation targeting (e.g., [Woodford, 2010](#)). Actual inflation is then the inflation target set in the policy regime plus a function of shocks that cause a trade-off between stabilizing inflation and real activity. Completing the equilibrium is the assumption that the fiscal authorities follow a passive, or Ricardian, policy of adjusting fiscal surpluses to satisfy Equation (20).

Our distribution of debt debasement is an estimate of the distribution of how taxes and spending rise and fall to support independent monetary policy. Given shocks to the implementation of monetary policy (like shocks to the Taylor rule), shocks that move the output gap, and shocks to the inflation target in the mandate, there will be a distribution of inflation outcomes. This creates a distribution of W_0 , which, given the fiscal support of independent central banking, implies a distribution for the present value of fiscal surpluses. In short, an independent central bank that targets inflation leaves a fiscal footprint, and our estimated distribution of debasement measures part of this surplus. Our distribution provides a set of estimates under different realizations of shocks for the *fiscal index* proposed by [Cochrane \(2017\)](#) and the extent of *fiscal backing* discussed by [Hall and Reis \(2013\)](#), [Del Negro and Sims \(2015\)](#).

4.4.2 Fiscal-driven inflation. A non-Ricardian, or active, fiscal policy does not provide fiscal backing to the central bank. In simple models, this is modeled as an exogenous present value of fiscal surpluses, the right-hand side of Equation (20). When it falls, then W_0 must fall. Since W_0 depends on B_0^t , which is predetermined, and the price level, then a negative fiscal shock raises inflation. This is the fiscal theory of the price level. The equilibrium of this class of models is a joint distribution of inflation and fiscal surpluses, $(\pi_{0,t}, W_0)$ (see, e.g., [Sims 2013](#)). Our estimates of the distribution of debt debasement map onto the marginal distribution for W_0 in this large class of models.

Recently, [Bianchi and Melosi \(2017\)](#) assume that fiscal policy alternates between active and passive regimes, and estimate these regimes using data on U.S. inflation. They simplify the government budget constraint by assuming that all debt is in nominal terms, one-period, privately held bonds, which requires them to assume an exogenous “term premium” shock that turns out to have a significant posterior variance. Our estimates can map, with some adjustments, to the joint posterior distribution of these shocks and inflation, providing a validation test of their estimates.

Both this example and the previous one make clear that our distribution of debasement includes both shocks to fundamentals *and* shocks to the policy

regime, in terms of the inflation target and the fiscal regime. Rational agents anticipate these policy changes and incorporate them probabilistically into the risk-adjusted distributions of inflation that we measure. To answer “what-if” questions, in the sense of the effect on the fiscal burden of inflation policy changes that surprise agents, one must use our estimates of counterfactuals in Section 3. The distribution of debasement gives an assessment of the consequences of policy changes when markets may get it wrong but are not completely surprised by policy in the sense of having put zero probability on it *ex ante*.²²

4.4.3 The maturity of debt and inflation, tax smoothing, and sovereign default. Missale and Blanchard (1994) write a model where the government can either inflate away the debt or raise taxes to pay for it, and bondholders form expectations of inflation before the policy maker sets actual inflation. The government faces a temptation to generate unexpected inflation to reduce socially costly taxes, and this causes a time inconsistency problem in the tradition of Barro and Gordon (1983). Missale and Blanchard (1994) solve for a reputation equilibrium where the government keeps expected inflation low by choosing the maturity of its debt *ex ante*. A shorter maturity reduces the *ex post* ability to inflate the debt away, and thus controls the temptation to do so. The model has predictions for inflation, debt debasement, and the optimal maturity of debt.

Their model assumes a reduced-form sensitivity of the real value of debt to inflation, which is decreasing in the maturity of privately held debt. Our estimates in the previous section accurately measure this sensitivity. Adding shocks to inflation in their model would produce a distribution of debasement that corresponds to our estimates. Moreover, Section 3 links our estimates of the distribution of debasement to the maturity distribution of privately held debt in a numerically accurate way. Finally, while Missale and Blanchard (1994) use the simple average maturity of public debt to link their theory to the data, we show in Section 3 that this rule of thumb is inaccurate and provide more accurate alternatives.

Persson, Persson, and Svensson (2006) implement the optimal fiscal and monetary policy under commitment by choosing a maturity structure for debt that exactly balances the marginal benefit of a surprise inflation with its marginal cost. They show that a crucial determinant of the marginal benefit is the debt debasement distribution at the optimal Ramsey outcome. This is precisely the object that we measure. Future research can accurately quantify and take these models to the data (e.g., extending Debortoli, Nunes, and Yared, 2017, to consider inflation).

²² Sims (1982) is an early argument for using distributions like we do in our VaR estimates in order to deal with the Lucas critique of counterfactuals.

Finally, [Aguilar et al. \(2013\)](#) study the incentives of governments to default on sovereign debt or to inflate it away, when there may be self-fulfilling crises. They assume a grace period for the repayment of debt in order to proxy for the maturity of the debt, and this parameter plays an important role in their theoretical results. Taking their model to the data requires calibrating this parameter. This can be done by matching it to our estimates of debt debasement.

5. Conclusion

This paper presents and implements a new method to evaluate the effect of higher-than-expected future inflation on the real value of debt. We identify four main drivers of debasement: who holds which debt, its maturity structure, and the stochastic properties of future inflation, including its volatility and persistence. For the United States in 2017, (1) the value of nominal debt held by the private sector is only about half the size of the total outstanding government debt. (2) The maturity structure, especially of debt held domestically, is skewed toward the short end. (3) The likelihood of near-term inflation rising significantly is quite low. And (4) inflation persistence at short horizons is small and so a sustained increase in inflation is unlikely. Taken together, these four features imply that inflationary debt debasement by more than a few percentage points of GDP is unlikely. Rules of thumb, like duration, that ignore the stochastic properties of inflation, provide inaccurate estimates. We propose and validate a new rule of thumb for inflationary debasement that takes into account duration, inflation persistence and volatility. Finally, we point out that our estimates of both the sensitivity of the fiscal burden to inflation and their joint distribution provide numerical targets for theoretical models of fiscal and monetary policy to hit in their calibrations.

In the process, we assemble new data and develop new methods that are of independent interest. We compile data measuring the maturity structure of debt held by different sets of investors for the United States from 2009 to 2017. We show that the debt maturity structure has changed significantly over time, and that its distribution is poorly approximated by a simple exponential formula. We use new data on derivatives contracts for both cumulative and forward inflation. We show that combining both sets of distributions allows us to extract information about the time-series behavior of inflation. Finally, we propose a new estimator for the risk-adjusted joint distribution of inflation using copulas. We find that inflation expectations are well anchored. The market expects risk-adjusted inflation to be stationary but also believes that the process has long memory. We find that inflation volatility has declined significantly from 2009 to 2017 and that there is important variation in inflation persistence.

Can one rescue an important role for inflation in debt dynamics? We show that if GDP or real interest rates are strongly correlated with inflation, then joint movements in all three can result in significant falls in the debt-to-GDP ratio, as was the case in the United States. Financial repression, an active policy

tool that extends debt maturity and is often used in emerging economies, can substantially increase the likelihood of large inflationary debasement. These confirm that it is not inflation, by itself, that can debase the real debt, but other policies that come with it.

Increasing the direct debasement of debt by inflation would require a change in one of the drivers of our low estimates. One way would be to have significantly higher inflation that completely surprises the market even though informed, profit-maximizing agents view this outcome as being virtually impossible. At least in our sample period, this has not happened. Indeed, as noted, expected inflation volatility has declined. Another way would be to have the private sector hold more long-term public debt. Yet, perhaps it is precisely by holding short-term debt that private agents are reducing the incentive of the government to inflate away the debt, so that low maturity holdings and low inflation expectations are mutually consistent. If so, issuing more long-term debt might result in lower debt prices.

Substantial debasement would arise if inflation were combined with financial repression, which could be implemented by forcing an ex post extension of maturities. Of course, such a policy would likely impose significant costs on the economy by impairing financial intermediation. None of these possibilities appears likely, and market prices today put the probability of the United States defaulting at close to zero. Therefore, the markets seem to be expecting economic growth, a favorable evolution of real interest rates, or perhaps primary budget surpluses, so that the debt-to-GDP ratio stabilizes. What is sure and inescapable is that, one way or another, the budget constraint of the government will have to hold.

References

- Aguiar, M., M. Amador, E. Farhi, and G. Gopinath. 2013. Crisis and commitment: Inflation credibility and the vulnerability to sovereign debt crises. Working Paper, Princeton University.
- Ait-Sahalia, Y., and A. W. Lo. 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94:9–51.
- Aizenman, J., and N. Marion. 2011. Using inflation to erode the US public debt. *Journal of Macroeconomics* 33:524–41.
- Barro, R. J., and D. B. Gordon. 1983. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy* 91:589–610.
- Berndt, A., H. Lustig, and S. Yeltekin. 2012. How does the US government finance fiscal shocks? *American Economic Journal: Macroeconomics* 4:69–104.
- Berriel, T. C. 2013. Nominal portfolio choice and the redistribution effects of inflation surprises. Working Paper, Princeton University.
- Beveridge, S., and C. R. Nelson. 1981. A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary Economics* 7:151–74.
- Bianchi, F., and L. Melosi. 2017. Escaping the great recession. *American Economic Review* 107:1030–58.

- Blanchard, O., G. Dell’Ariccia, and P. Mauro. 2010. Rethinking macroeconomic policy. *Journal of Money, Credit and Banking* 42:199–215.
- Bohn, H. 2011. The economic consequences of rising US government debt: privileges at risk. *FinanzArchiv: Public Finance Analysis* 67:282–302.
- Breeden, D. T., and R. H. Litzenger. 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51:621–51.
- Campbell, J. Y. 1991. A variance decomposition for stock returns. *Economic Journal* 101:157–79.
- Campbell, J. Y., and R. J. Shiller. 1988. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1:195–228.
- Chung, H., and E. M. Leeper. 2007. What has financed government debt? Working Paper, Board of Governors of the Federal Reserve.
- Cochrane, J. H. 2011. Understanding policy in the great recession: Some unpleasant fiscal arithmetic. *European Economic Review* 55:2–30.
- . 2017. Michelson-Morley, Fisher and Occam: The radical implications of stable inflation at the zero bound. *NBER Macroeconomics Annual* 32:113–226.
- Davig, T., E. M. Leeper, and T. B. Walker. 2011. Inflation and the fiscal limit. *European Economic Review* 55:31–47.
- Debtoli, D., R. Nunes, and P. Yared. 2017. Optimal time-consistent government debt maturity. *Quarterly Journal of Economics* 132:55–102.
- Del Negro, M., and C. A. Sims. 2015. When does a central bank’s balance sheet require fiscal support? *Journal of Monetary Economics* 73:1–19.
- Doepke, M., and M. Schneider. 2006. Inflation and the redistribution of nominal wealth. *Journal of Political Economy* 114:1069–97.
- Faraglia, E., A. Marcet, R. Oikonomou, and A. Scott. 2013. The impact of debt levels and debt maturity on inflation. *Economic Journal* 123:F164–F192.
- Fisher, L., and R. L. Weil. 1971. Coping with the risk of interest-rate fluctuations: Returns to bondholders from naive and optimal strategies. *Journal of Business* 44:408–31.
- Fleckenstein, M., F. A. Longstaff, and H. Lustig. 2017. Deflation risk. *Review of Financial Studies* 30:2719–60.
- Giannitsarou, C., and A. Scott. 2008. In *NBER international seminar on macroeconomics*. 393–442. Chicago: Chicago University Press.
- Gürkaynak, R. S., B. Sack, and E. Swanson. 2005. The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American Economic Review* 95:425–36.
- Gürkaynak, R., B. Sack, and J. Wright. 2007. The U.S. treasury yield curve: 1961 to the present. *Journal of Monetary Economics* 54:2291–304.
- . 2010. The TIPS yield curve and inflation compensation. *American Economic Journal: Macroeconomics* 2:70–92.
- Hall, G. J., and T. J. Sargent. 2011. Interest rate risk and other determinants of post-WWII U.S. government debt/gdp dynamics. *American Economic Journal: Macroeconomics* 3:192–214.
- Hall, R. E., and R. Reis. 2013. Maintaining central-bank solvency under new-style central banking. Working Paper, Stanford University.
- International Monetary Fund. 2020. World economic outlook update. Report, Washington, DC.
- Jackwerth, J. C., and M. Rubinstein. 1996. Recovering probability distributions from option prices. *Journal of Finance* 51:1611–31.

- J. P. Morgan. 2013. Spotlight on inflation derivatives. Investment insights q1. Report, New York.
- Kitsul, Y., and J. H. Wright. 2013. The economics of options-implied inflation probability density functions. *Journal of Financial Economics* 110:696–711.
- Kocherlakota, N. 2012. Optimal outlooks. Speech Financial Intermediation Research Society Conference. <https://www.minneapolisfed.org/speeches/2013/optimal-outlooks>.
- Krause, M. U., and S. Moyer. 2016. Public debt and changing inflation targets. *American Economic Journal: Macroeconomics* 8:142–76.
- Leeper, E. M., and X. Zhou. 2013. Inflation's role in optimal monetary-fiscal policy. Working Paper, University of Virginia.
- McKinnon, R. I. 1973. *Money and capital in economic development*. Washington, DC: Brookings Institute.
- Missale, A., and O. J. Blanchard. 1994. The debt burden and debt maturity. *American Economic Review* 84:309–19.
- Persson, M., T. Persson, and L. E. O. Svensson. 2006. Time consistency of fiscal and monetary policy: A solution. *Econometrica* 74:193–212.
- Reinhart, C. M., and M. B. Sbrancia. 2015. The liquidation of government debt. *Economic Policy* 30:291–333.
- Reinhart, C. M. and K. S. Rogoff. 2009. *This time is different: Eight centuries of financial folly*. Princeton, NJ: Princeton University Press.
- Reinhart, C. M., V. R. Reinhart, and K. S. Rogoff. 2012. Public debt overhangs: Advanced-economy episodes since 1800. *Journal of Economic Perspectives* 26:69–86.
- Sims, C. A. 1982. Policy analysis with econometric models. *Brookings Papers on Economic Activity* 1:107–152.
- . 2013. Paper money. *American Economic Review Papers and Proceedings* 103:563–84.
- Sklar, M., 1959. Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut Statistique de Université de Paris* 8:229–31.
- Vayanos, D., and J.-L. Vila. 2021. A preferred-habitat model of the term structure of interest rates. *Econometrica* 89:77–112.
- Veronesi, P., and L. Zingales. 2010. Paulson's gift. *Journal of Financial Economics* 97:339–68.
- Woodford, M. 2010. Optimal monetary stabilization policy. In *Handbook of monetary economics*, eds. B. M. Friedman and M. Woodford, vol. 3, chap. 14 723 – 828. North Holland: Elsevier.