The Zero-Beta Rate*

Sebastian Di Tella — Stanford University
Benjamin Hébert — Stanford University

Pablo Kurlat — USC
Qitong Wang — USC

March 17, 2023

Abstract

We use equity returns to construct a time-varying measure of what we call the zero-beta rate: the expected return of a stock portfolio orthogonal to the stochastic discount factor. In contrast to safe rates, the zero-beta rate fits the aggregate Euler equation remarkably well. It has a large and volatile spread with respect to the safe rate. This spread responds to monetary policy shocks, which move zero-beta and safe rates in opposite directions. We claim that the zero-beta rate is the correct intertemporal price and that the safe rate primarily reflects the behavior of a convenience yield on safe assets.

Preliminary and incomplete; please do not distribute.

*We would like to thank Marianne Andries, Adrien Auclert, Jonathan Berk, John Cochrane, Lars Hansen, Benjamin Holcblat, Arvind Krishnamurthy, Eben Lazarus, Martin Lettau, Hanno Lustig, Thomas Mertens, Monika Piazzesi, Shri Santosh, Martin Schneider, Ken Singleton, David Sraer, and Chris Tonetti for helpful comments, as well as seminar and conference participants.

1 Introduction

The interest rate is perhaps the most important price in a market economy. It captures the intertemporal price of goods and plays a central role in business cycles and monetary policy. We use equity returns to construct a measure of the interest rate that we call the zero-beta rate: the expected return of a stock portfolio orthogonal to the stochastic discount factor (SDF). We show that the zero-beta rate fits a stable aggregate consumption Euler equation remarkably well, that it has a large and cyclical spread with respect to the safe rates that central banks control, and that this spread responds to monetary policy shocks, which move safe and zero-beta rates in opposite directions. We claim that the zero-beta rate is the correct intertemporal price and the behavior of safe rates reflects a large and time-varying convenience yield on safe assets.

Our motivation starts with the aggregate consumption Euler equation,

$$1 = \mathbb{E}_t \left[\delta \left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{R_t}{P_{t+1}/P_t} \right], \tag{1}$$

where c_t is aggregate consumption, P_t is the price level, and R_t is an interest rate. This equation lies at the heart macroeconomics, and is the structural relationship that central banks aim to exploit through monetary policy. But, if we interpret the interest rate R_t as the return of safe bonds such as Treasury bills, the Euler equation does not fit the data. The left panel of Figure 1 makes this point graphically, showing the expected growth rate of consumption and the expected real Treasury bill return, both predicted with the same set of macroeconomic variables. The failure of the classic aggregate Euler equation is well-known (Hansen and Singleton 1983; Dunn and Singleton, 1986; Hall, 1988; Yogo, 2004) and there are many possible explanations, the simplest being a time-varying discount factor.

We suggest instead that it is the safe interest rate that is wrong. It is well-understood that the safest and most convenient assets, such as cash and deposits, have a convenience yield, and we do not expect the Euler equation to hold for these assets. To obtain the correct intertemporal price, one would need to account for their convenience. We pursue the hypothesis that this convenience applies not only to cash and deposits, but instead to the category of safe assets more generally. It is not necessary, for our purposes, to precisely define which assets are safe or specify the origin of such convenience.[1] We instead proceed under the assumption that publicly traded equities are not safe and do not provide convenience.

[1]We expect that Treasury bonds, highly rated corporate bonds, highly rated mortgage bonds, and other similar assets are all to some degree convenient, perhaps because they are used to back other convenient assets such as deposits and money market fund shares.

Figure 1: Expected consumption growth vs. Expected Real 1m Treasury Bill Return (left) and Real Zero-Beta Rate (right)

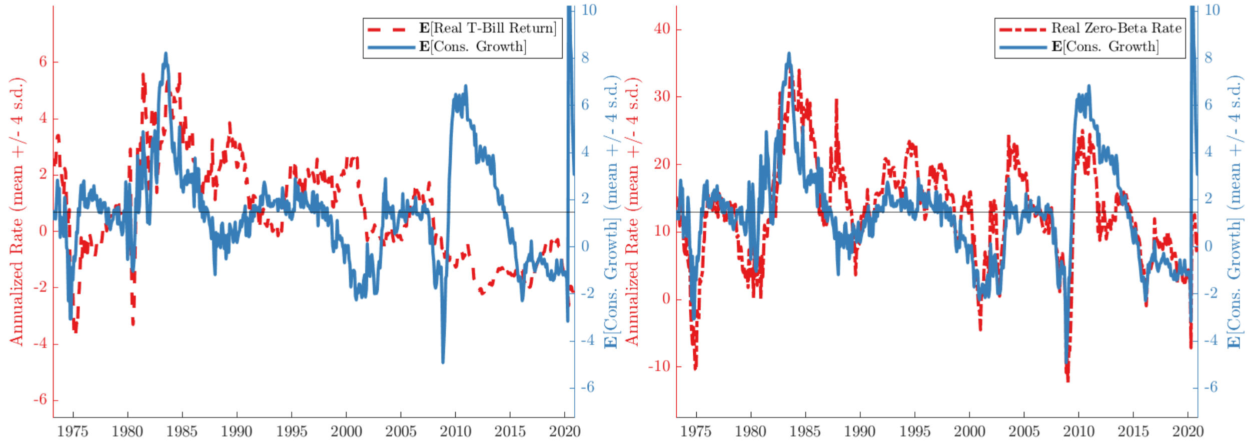

Notes: Both panels of this figure plot expected real returns against expected consumption growth, over time. Expected real returns are constructed using nominal rates (1m bill yields on the left, our zero-beta rate on the right) less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the instruments described in Section 4, which are the same instruments used to construct the zero-beta rate. In both panels, the right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the relevant expected real return. All series are annualized.

Our approach is therefore to measure the interest rate by looking at the expected return on an equity portfolio. But equities, of course, may have a risk premium. Our observation is that if one constructs a portfolio of risky equities whose excess return is orthogonal to the SDF, it should not have a risk premium, and its expected return should be equal to the intertemporal price of goods. We call the expected return of such a portfolio the zero-beta rate.

The first contribution of the paper is to construct a time-series of the zero-beta rate. We first postulate a model of the SDF that is linear in a set of factors that the literature has found to explain the cross-section of equity returns. Next, we estimate the betas of the excess returns of equities with respect to each of the factors, and use these estimated betas to construct a unit-investment, minimum-variance, zero-beta portfolio. Finally, we project the returns of this zero-beta portfolio on a set of macroeconomic predictors to obtain an expected return, which we call the zero-beta rate. This sequential procedure is infeasible because to construct excess returns and estimate betas one needs to know the zero-beta rate, so we instead estimate all of our parameters simultaneously via GMM.

The right panel of Figure 1 shows the real zero-beta rate. It is high on average, and has a large and volatile spread with the expected real return of Treasury bills (around 8% per year on average). The spread is so large that it renders inconsequential the much smaller spread between the returns on different types of safe assets. In what follows, we will use the Treasury bill yield as our measure of *the* safe interest rate.[2] The average level of the zero-beta rate may seem surprising. But it actually reflects a well-known fact, going back to Black et al. [1972], who pointed out, in the context of CAPM, that the expected return of an equity portfolio with zero covariance to the market was well in excess of Treasury bill yields.[3]

One common interpretation of the Black et al. [1972] finding has been that the CAPM is wrong—there are more priced factors beyond the market—and that once we incorporate them, the zero-beta rate should coincide with safe interest rates. But finding factors to explain the spread is very hard. Lopez-Lira and Roussanov [2020], for example, find that even if one removes almost all of common factors in stock returns, the remainder still features a large Sharpe ratio (computed using Treasury bill yields as the safe rate).[4]

We propose instead that the zero-beta rate really is the interest rate, the intertemporal price of goods, and that the time-varying spread with the safe rate reflects convenience on safe assets and not an unexplained risk premium. This claim is testable, via the Euler equation.

The second contribution of the paper is to show that the zero-beta rate fits the Euler equation strikingly well. The right panel of Figure 1 shows the results graphically. It shows the real zero-beta rate plotted against expected consumption growth, predicted with the same set of macroeconomic variables. The results are striking: the two series strongly co-move once they are rescaled. This is essentially a graphical test of the linearized Euler equation, and our results suggest that the zero-beta rate is indeed the intertemporal price of consumption.

Our more formal analysis constructs the zero-beta rate using GMM and then tests the Euler equation moments (which are not used in the construction of the zero-beta rate) using weak-identification-robust methods. We cannot reject an intertemporal elasticity of substitution (IES) below 0.6 (if CRRA, risk aversion above 1.6), but are able to reject higher values of the IES. This contrasts with the results of the same test applied to Treasury bill yields and to the aggregate market return. The Euler equation is rejected for all values of the IES we consider when applied to the Treasury bill yield, and is not rejected for essentially any values when applied to the market return, reflecting weak identification.

The consequences of our results for monetary policy are potentially vast. There is a large spread between the safe rates that central banks control as policy instruments and the zero-

[2]For example, Krishnamurthy and Vissing-Jorgensen [2012] find a spread of around 70 bps between Treasurys and AAA corporate debt.

[3]See also Shanken [1986] and more recently Bali et al. [2017].

[4]See also Kim et al. [2021].

beta rate that goes in the Euler equation, and this spread is time-varying. If this spread were exogenous to monetary policy, the central bank could still exploit the classic Euler equation as a structural relationship. If instead monetary policy affects the spread, then the behavior of the zero-beta rate can differ substantially from the behavior of the safe rate.

Our third contribution is to compute the response of the safe and zero-beta rates to monetary policy shocks (identified using the Romer and Romer [2004] and Nakamura and Steinsson [2018] approaches). Our point estimates suggest that an unexpected monetary tightening that raises the Treasury bill yield will *lower* the zero-beta rate. While, other things being equal, a higher Treasury bill yield is associated with a higher zero-beta rate, an unexpected monetary tightening also flattens the yield curve and increases credit spreads, and these are associated with a lower zero-beta rate. Empirically, these effects dominate.

This result may seem strange at first because it implies that the intertemporal substitution effect of a monetary tightening is the opposite of what conventional wisdom says. It makes current consumption cheaper relative to future consumption, not more expensive. But it is actually in line with empirical facts and theory. Empirically, a contractionary monetary policy shock reduces consumption *growth*, as opposed to generating a lower *level* of consumption but higher growth. The Euler equation implies that the zero-beta rate should fall as households correctly expect lower future income and try to save. In other words, the zero-beta rate fits the Euler equation both unconditionally and conditional on monetary policy shocks. It is rather the short-run rise of safe rates that is puzzling in light of the fall of consumption growth (or the fall of consumption growth that is puzzling in light of the rise of safe rates). The literature has traditionally explained this with adjustment costs such as habits or informational frictions.[5] Instead, we attribute the rise in safe rates to an endogenous fall in convenience yields, and show that this outcome arises naturally in a stylized New Keynesian model augmented with convenience on safe bonds.

1.1 Literature Review

The aggregate consumption Euler equation plays a central role in our analysis. One natural objection is that while this equation is a central feature of representative-agent models, in heterogenous agents models the consumption Euler equation might hold for some individuals but not at the aggregate level. But it's almost exactly the other way around. At the individual level households face uninsurable idiosyncratic risk, borrowing constraints, and trading frictions. Such features are central to most heterogenous-agent models, with and without nominal rigidities. But a central result in this literature is that there is nonetheless

[5]See Christiano et al. [2005] or Smets and Wouters [2007] for the former, and Auclert et al. [2020] for the latter.

an Euler equation at the aggregate level. Werning [2015] explores the issue in detail and provides clean theoretical results, but the result shows up with variations throughout that literature (Krueger and Lustig [2010], Auclert et al. [2018, 2020]). To be clear, one can write models where the aggregate Euler equation fails (see Bilbiie [2021]), but an aggregate Euler equation can be consistent with realistic individual-level consumption behavior. Our empirical evidence supports this view.

Our approach is based on the idea that safe assets may have a convenience yield. It is well-known that some safe assets, such as cash or bank deposits, have a lower yield than safe bonds such as Treasury bills. Krishnamurthy and Vissing-Jorgensen [2012] show that Treasurys have a lower yield than equally safe AAA corporate debt (by around 70 bps), which they interpret as convenience yield. Implicitly they are taking the yield of AAA corporate debt as the interest rate without convenience. Lenel et al. [2019] measure the "shadow spread" between short rates and the return on short-term bonds that can be extrapolated from the yield curve. Their interpretation is that short debt is useful to financial intermediaries for backing inside money. van Binsbergen et al. [2019] identify convenience spread by comparing safe rates with the interest rate implied by options, and find a spread of around 30 bps. Du et al. [2018] document violations of covered interest parity, which is a spread between safe dollar rates and safe rates synthesized from other currencies via foreign exchange transactions, of around 50 bps.[6] Our interpretation is that these papers have only identified the "tip of the iceberg" in terms of convenience yields. We use equity returns as the inconvenient benchmark, and find a convenience spread that is an order of magnitude larger.

We do not take a stance on the origin of such convenience, although we conjecture that it originates with bank deposits and other payment instruments. Safe assets of different sorts are useful for backing these payment instruments. In this we are influenced by Lenel et al. [2019], who propose this explanation for the convenience of safe short-term debt and consider the effects of stickiness in both prices and the supply of reserves. See also Piazzesi and Schneider [2021], who study the effect of convenience yields for monetary policy with flexible prices. In contrast, we do not provide a theory of the convenience yield, and introduce convenience in our modeling framework via safe bonds in the utility function. This is a transparent way of introducing convenience without distorting other aspects of the model, and helps guide the empirical work that is the main contribution of the paper.

Our finding of a large spread between the zero-beta rate and safe rates is consistent with

[6] Note that these relatively small arbitrage spreads can be levered to produce risk-free equity returns that are consistent with our estimates of the zero-beta rate. For example, if banks can lever the CIP arbitrage 8:1, a 50-100bps CIP arbitrage is consistent with a 4.5-9% risk-free equity return over safe rates. See Boyarchenko et al. [2018] for examples of this sort of calculation.

the literature on the existence of the "beta anomaly" and its implications.[7] We innovate, relative to this literature, by studying time-variation in the zero-beta rate and documenting the connection between the zero-beta rate and expected consumption growth. Methodologically, we build on GMM tests of the Euler equation (Hansen and Singleton [1982], Dunn and Singleton [1986]), using weak-instrument-robust methods (Stock and Wright [2000], Yogo [2004]) and a regularized covariance matrix estimator (Ledoit and Wolf [2017]), in a procedure inspired by the maximum likelihood approach of Shanken [1986]. Our choice of assets and factors is informed by the work of Novy-Marx and Velikov [2022] on beta-sorted portfolios, and our choice of instruments is guided by the literature on predicting business cycles (e.g. Kiley [2022]).

2 Model

We build a simple monetary model with: (1) a rich and potentially time-varying stochastic discount factor, (2) time-varying and potentially endogenous convenience on safe bonds, and yet (3) a traditional consumption Euler equation. We work in nominal terms and avoid taking a position on the presence or absence of nominal rigidities.

2.1 Setup

There is a representative household with preferences

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\delta^t\xi_t\left(\frac{c_t^{1-\sigma}}{1-\sigma}+\eta_{m,t}\log\left(M_t/P_t\right)+\eta_{b,t}\log\left(B_t/P_t\right)\right)\right], \tag{2}$$

where c_t is consumption at time t, B_t are holdings of safe, one-period nominal bonds held at time t, M_t are money holdings, and P_t is the price level.We include money and bonds in the utility function as a transparent way of introducing convenience to these assets. We do not provide a deeper theory of the source of this convenience. Bonds and money enter the utility function separably from consumption and from each other, and $\eta_{m,t}$ and $\eta_{b,t}$ are shocks to the demand for money and bonds. This is essentially the model in Golosov and Lucas [2007], augmented with safe bonds in the utility function. In this framework, if prices are flexible, money and bonds are neutral and super-neutral. ξ_t is an exogenous stochastic process that will generate fluctuations in the stochastic discount factor that are independent of macro quantities. We assume ξ_t is a martingale and independent of M_t, B_t, $\eta_{m,t}$, and $\eta_{b,t}$.

[7] E.g. Black [1972], Frazzini and Pedersen [2014], Baker and Wurgler [2015], Hong and Sraer [2016], Bali et al. [2017], Baker et al. [2020].

For simplicity, we treat the supply of safe bonds B_t^s and money M_t^s as exogenous. There are N risky assets in zero net supply, which are meant to capture equities in the empirical work. Denote their nominal return by $R_{i,t}$. We assume that at any time: (1) there is no portfolio of these risky assets that is risk-free, and (2) there is at least one portfolio of risky assets whose return is uncorrelated to the SDF.[8]

The household's budget constraint is:

$$P_t c_t + (B_t - B_t^s) + M_t + \sum_{i=1}^{N} X_{i,t} \leq \left(B_{t-1} - B_{t-1}^s\right) R_{b,t-1} + M_{t-1} + \sum_{i=1}^{N} X_{i,t-1} R_{i,t} + P_t y_t + T_t^m. \quad (3)$$

Here $T_t^m = M_t^s - M_{t-1}^s$ is a government transfer, y_t is real income, and $X_{i,t}$ is the nominal amount the household invests in asset i. $R_{b,t}$ is the nominal return on a risk-free bond. Note that the household is liable for the safe bonds B_t^s and also chooses to hold B_t.

The household's problem is to choose c_t, B_t, M_t and $X_{i,t}$ to maximize (2) subject to (3) and the natural borrowing limit. In equilibrium $B_t = B_t^s$, $X_{i,t} = 0$, and $M_t = M_t^s$. We don't take a stand on nominal rigidities and the production environment, so the model does not pin down prices P_t and real output y_t. Our results are thus consistent with flexible prices and with different forms of price stickiness.

2.2 Equilibrium

The household's first order condition for consumption is

$$\Lambda_t = \delta^t \xi_t c_t^{-\sigma} / P_t, \quad (4)$$

where the Lagrange multiplier Λ_t is the SDF. Here we see the role ξ_t plays in creating a realistic SDF, by allowing for movements in the SDF unrelated to consumption. We will guess and verify that ξ_t is independent of $c_t^{-\sigma}/P_t$.

The Euler equations for money, safe bonds, and risky assets are:

$$c_t^{-\sigma} = \eta_{b,t} \left(B_t/P_t\right)^{-1} + \delta \mathbb{E}_t \left[c_{t+1}^{-\sigma} \frac{R_{b,t}}{P_{t+1}/P_t} \right], \quad (5)$$

$$c_t^{-\sigma} = \eta_{m,t} \left(M_t/P_t\right)^{-1} + \delta \mathbb{E}_t \left[c_{t+1}^{-\sigma} \frac{1}{P_{t+1}/P_t} \right], \quad (6)$$

$$c_t^{-\sigma} = \delta \mathbb{E}_t \left[\frac{\xi_{t+1}}{\xi_t} c_{t+1}^{-\sigma} \frac{R_{i,t+1}}{P_{t+1}/P_t} \right]. \quad (7)$$

[8]Strictly speaking, the returns $R_{i,t}$ and the SDF are endogenous objects. (1) and (2) should be understood as assumptions on the underlying payoffs of the risky assets.

If the household saves in money or safe bonds, it takes into account the convenience those assets provide. Here we have used the fact that ξ_t is independent of $c_t^{-\sigma}/P_t$ and a martingale, so it drops out of the Euler equation for money and bonds. For a risky asset, however, the household needs to consider the covariance of its return with the SDF.

The zero-beta rate is the conditional expected return on a portfolio of risky assets whose returns are conditionally uncorrelated with Λ_{t+1}/Λ_t. We have assumed above that such a portfolio exists, with returns $R_{p,t+1}$. We can use (7) to obtain the Euler equation

$$c_t^{-\sigma} = \delta \mathbb{E}_t \left[c_{t+1}^{-\sigma} \frac{R_{0,t}}{P_{t+1}/P_t} \right], \tag{8}$$

where the correct interest rate is the zero-beta rate, $R_{0,t} = \mathbb{E}_t \left[R_{p,t+1} \right]$. We have used the fact that the return $R_{p,t+1}$ is uncorrelated with the SDF, and also that ξ_t is independent of $c_t^{-\sigma}/P_t$. We can also show that the zero-beta rate is the inverse of the mean of the growth of the SDF, $R_{0,t} = \mathbb{E}_t \left[\Lambda_{t+1}/\Lambda_t \right]^{-1}$. The zero-beta rate $R_{0,t}$ is an *expected* return, so it is known at time t. However, as is standard with nominal returns, there may be risk associated with the one-period-ahead price level. This is a separate issue from the distinction between safe rates and zero-beta rates.

We can rearrange (6) to obtain an expression for money demand,

$$c_t^{-\sigma}/P_t = \eta_{m,t} M_t^{-1} \times \frac{R_{0,t}}{R_{0,t} - 1}. \tag{9}$$

At this point we can verify that ξ_t is indeed independent of $c_t^{-\sigma}/P_t$. Under our guess, (8) holds, which implies that the process for $R_{0,t}$ is pinned down by the process for $c_t^{-\sigma}/P_t$, so it is also independent of the process for ξ_t. Therefore all the terms on the right-hand-side of (9) are independent of ξ_t, which implies that so is $c_t^{-\sigma}/P_t$.

It only remains to solve for the interest rate on safe bonds. Rearranging (5) we obtain an expression for the spread,

$$s_t \equiv 1 - \frac{R_{b,t}}{R_{0,t}} = \frac{\left(B_t/\eta_{b,t} \right)^{-1}}{c_t^{-\sigma}/P_t}. \tag{10}$$

Notice that $s_t \approx R_{0,t} - R_{b,t}$. The model can be solved block-recursively. First, the consumption Euler equation (8) and the equation for money holdings (9) jointly pin down the zero-beta rate $R_{0,t}$ and $c_t^{-\sigma}/P_t$ as functions of the process for $M_t/\eta_{m,t}$. Equation (10) then pins down $R_{b,t}$ as a function of the process for $B_t/\eta_{b,t}$. All of this is independent of nominal rigidities and supply shocks. To pin down c_t and P_t separately we would need to be explicit about nominal rigidities and supply shocks, which we do not take a stand on.

3 Measuring the zero-beta rate

3.1 Empirical Implementation

To construct the zero-beta rate and associated portfolio, we impose several additional assumptions, in effect imposing the structure of a linear factor model. These assumptions allow us to construct an estimate of the zero-beta rate using relatively standard techniques.

We begin with a balanced panel of N assets (equity portfolios, which we refer to as test assets) and T periods. Let R_{t+1} denote the vector of returns across assets $i \in \{1, \dots, N\}$. We assume there is a set of K priced factors, whose values at time t are $F_{j,t}$ for $j \in \{1, \dots, K\}$. The excess returns of each test asset can be projected onto the space of factor returns as

$$R_{i,t+1} - R_{0,t} = \alpha_i + \sum_{j=1}^{K} \beta_{ij} F_{j,t+1} + \epsilon_{i,t+1}, \tag{11}$$

where $\epsilon_{i,t+1}$ has an unconditional zero mean and α_i and β_{ij} are regression coefficients. We assume the betas are constant over time (specifically, that $\epsilon_{i,t+1}$ is uncorrelated with the factors conditional on the information at time t).[9] Consistent with this assumption, we will use portfolios of stocks (beta-sorted portfolios and industry portfolios) that might be expected to have stable betas over time, as opposed to considering individual companies. In our robustness exercises, we allow for time-varying betas, and our results are largely unchanged. Let α be the vector of the α_i coefficients, and let ϵ_{t+1} be the vector of regression residuals $\epsilon_{i,t+1}$.

Our first key economic assumption is that the nonlinear SDF Λ_{t+1}/Λ_t can be well approximated by a linear factor structure for the purpose of pricing equities. The first of these factors is the excess return of the market with respect to the zero-beta rate, $F_{1,t+1} = R_{m,t+1} - R_{0,t}$. The remainder of the factors are assumed to be either zero-investment portfolios that do not explicitly involve the zero-beta rate (such as the SMB and HML portfolios of Fama and French [1993]) or non-tradable factors (such as consumption growth). Our assumption is that

$$\frac{\Lambda_{t+1}}{\Lambda_t} = R_{0,t}^{-1} + \sum_{j=1}^{K} \omega_{j,t} \left(F_{j,t+1} - \mathbb{E}_t \left[F_{j,t+1} \right] \right) + \zeta_{t+1}, \tag{12}$$

where ζ_{t+1} is mean zero and uncorrelated with any stock return, conditional on the information at time t, $\mathbb{E}_t \left[R_{i,t+1} \zeta_{t+1} \right] = 0$. This assumption implies that our non-linear SDF is

[9] We do not assume that the $\epsilon_{i,t+1}$ are uncorrelated with each other. Equity returns might have comovement beyond the priced factors.

equivalent to a linear SDF for the purpose of pricing equity portfolios,

$$1 = \mathbb{E}_t\left[\frac{\Lambda_{t+1}}{\Lambda_t} R_{i,t+1}\right] = \mathbb{E}_t\left[\left(R_{0,t}^{-1} + \sum_{j=1}^{K} \omega_{j,t}\left(F_{j,t+1} - \mathbb{E}_t\left[F_{j,t+1}\right]\right)\right) R_{i,t+1}\right]. \tag{13}$$

Our framework allows for time variation in the zero-beta rate $R_{0,t}$ and the prices of risk for each of the factors, $\omega_{j,t}$. We elaborate on this point below.

Our second key economic assumption is that the zero-beta rate is linear in a set of L predictor variables, $Z_{l,t}$ for $l \in \{1, \ldots, L\}$, the vector of which we denote Z_t. To simplify our notation, let $Z_{0,t} = 1$ and standardize $Z_{l,t}$, for $l \geq 1$, to have mean zero and unit variance. We assume that

$$R_{0,t}(\gamma) = R_{b,t} + \gamma' \cdot Z_t, \tag{14}$$

where $\gamma \in \mathbb{R}^{L+1}$ is a vector of constants and $R_{b,t}$ is the nominal one-period safe bond yield (which we identify with the Treasury bill yield). The conventional view (i.e. that the mean of the SDF is the reciprocal of the return on a safe bond) is $\gamma = \vec{0}$ and therefore $R_{0,t} = R_{b,t}$.

3.2 Estimation via GMM

Our estimation procedure can be described as a GMM version of the MLE procedure of Shanken [1986], modified to produce a time-varying zero-beta rate. We discuss the connection between our approach and the Shanken [1986] approach in Appendix Section G.

The relevant parameters of our model are $\theta = (\alpha, \beta, \gamma)$, where (α, β) are the regression coefficients associated with (11) and γ is the vector of coefficients in (14). Define the residual:

$$\hat{\epsilon}_{i,t+1}(\theta) = R_{i,t+1} - \alpha_i - (1 - \beta_{i1})(R_{b,t} + \gamma' \cdot Z_t) - \beta_{i1} R_{m,t+1} - \sum_{j=2}^{K} \beta_{ij} F_{j,t+1}. \tag{15}$$

Note that the first factor (the market excess return) explicitly involves the zero-beta rate, which is why it is treated differently from other factors in this definition. Let $\hat{\epsilon}_t(\theta)$ denote the vector of residuals.

Following common practice (see Cochrane [2009]), we will use a reduced-rank weighting matrix that selects moments and achieves exact identification, as opposed to using a full-rank weighting matrix with over-identifying restrictions. In our context, this means using a specific zero-beta portfolio (whose weights will end up being parameter-dependent) to estimate our model parameters. Let $w \in \mathbb{R}^N$ denote some cross-sectional weights satisfying the restrictions $w' \cdot \beta(\theta) = \vec{0}$. Any such a vector can be constructed from an arbitrary vector $\tilde{w} \in \mathbb{R}^N$ by $w = H(\theta)\tilde{w}$, where $H(\theta) = I - \beta(\theta)\beta(\theta)^{+}$ is a symmetric orthogonal projection matrix.[10]

We use the projection (OLS) moments $\mathbb{E}\left[\hat{\epsilon}_{it}\left(\theta\right)F_{j,t}\right]=0$ (treating $F_{0,t}=1$ as a constant time series). There are $N\times(K+1)$ of these moments, and they are sufficient to identify the parameters (α,β). We also use instrumented versions of the asset pricing moments, which are conditional,

$$\mathbb{E}_t\left[\left(R_{0,t}\left(\gamma\right)\right)^{-1}+\sum_{j=1}^{K}\omega_{j,t}\left(F_{j,t+1}-\mathbb{E}_t[F_{j,t+1}]\right)\left(R_{i,t+1}-R_{0,t}\left(\gamma\right)\right)Z_{l,t}\right]=0.$$

There are $L+1$ of these asset pricing moments for each asset (recall that $Z_{0,t}=1$ is a constant). When these moments are weighted by the weight vector $w=H\left(\theta\right)\tilde{w}$, they simplify:

$$\begin{aligned}\mathbb{E}_t\left[\left(R_{0,t}\left(\gamma\right)\right)^{-1}+\sum_{j=1}^{K}\omega_{j,t}\left(F_{j,t+1}-\mathbb{E}_t[F_{j,t+1}]\right)\left(\tilde{w}'\cdot H\left(\theta\right)\cdot\left(R_{i,t+1}-R_{0,t}\left(\gamma\right)\right)\right)Z_{l,t}\right]&=\\ \mathbb{E}_t\left[\left(R_{0,t}\left(\gamma\right)\right)^{-1}+\sum_{j=1}^{K}\omega_{j,t}\left(F_{j,t+1}-\mathbb{E}_t[F_{j,t+1}]\right)\left(\tilde{w}'\cdot H\left(\theta\right)\cdot\left(\alpha_i+\epsilon_{i,t+1}\right)\right)Z_{l,t}\right]&=\\ \left(R_{0,t}\left(\gamma\right)\right)^{-1}\mathbb{E}_t\left[\left(\tilde{w}'\cdot H\left(\theta\right)\cdot\left(\alpha_i+\epsilon_{i,t+1}\right)\right)Z_{l,t}\right]&=\\ \left(R_{0,t}\left(\gamma\right)\right)^{-1}\mathbb{E}_t\left[\left(\tilde{w}'\cdot H\left(\theta\right)\cdot\left(R_{i,t+1}-R_{0,t}\left(\gamma\right)\right)\right)Z_{l,t}\right]&=0.\end{aligned}$$

The second equation follows from the first given (11) and the fact that $H\left(\theta\right)\cdot\tilde{w}$ is a zero-beta portfolio. The third equation follows from the second under the assumption of constant betas, and the fourth from the third by the zero-beta property of $H\left(\theta\right)\cdot\tilde{w}$. We use the unconditional version of the fourth equation as moments. In vector form,

$$\mathbb{E}\left[H\left(\theta\right)\cdot\left(R_{t+1}-R_{0,t}\left(\gamma\right)\right)\otimes Z_t\right]=0.$$

There are $N\times(L+1)$ of these moments, but (under the appropriate rotation) $K\times(L+1)$ are trivial as a consequence of the reduced rank of $H\left(\theta\right)$. Intuitively, any zero-beta portfolio must have an expected return equal to the zero-beta rate. However, not all portfolios have an expected return equal to the zero-beta rate—some have a risk premium. The role of the orthogonal projection matrix $H\left(\theta\right)$ is to remove portfolios with a risk premium from the set of moments being considered. This has the effect of eliminating the need to estimate the prices of risk $\omega_{j,t}$.

We weight them by a vector $\tilde{w}\left(\theta\right)$ to produce $L+1$ weighted asset pricing moments,

[10] Here, I is the identity matrix, and $(\cdot)^+$ denotes the Moore-Penrose pseudo-inverse. H is the orthogonal projection matrix with respect to the betas.

which are sufficient to exactly identify the γ parameters. For convenience in what follows, we will define our weights in terms of the vector $w(\theta) = H(\theta) \cdot \tilde{w}(\theta)$ as opposed to $\tilde{w}(\theta)$ directly. This is without loss provided that $w(\theta)' \cdot \beta(\theta) = \vec{0}$.

Summarizing, our GMM analysis uses the moments $\mathbb{E}[g_t(\theta)] = 0$, where

$$g_t(\theta) = \begin{bmatrix} \hat{\epsilon}_{t+1}(\theta) \otimes F_{t+1}(\gamma) \\ H(\theta) \cdot (R_{t+1} - R_{0,t}(\gamma)) \otimes Z_t \end{bmatrix},$$

and the weight matrix

$$W(\theta) = \begin{bmatrix} I & \mathbf{0} \\ \mathbf{0} & w(\theta)\, w(\theta)' \end{bmatrix}.$$

We estimate our model using a continuously updating approach. Our GMM analysis thus solves

$$\hat{\theta} \in \arg\min_{\theta} \underbrace{\sum_{l=0}^{L} \left(T^{-1} \sum_{t=1}^{T} w(\theta)' \cdot (\alpha(\theta) + \hat{\epsilon}_t(\theta))\, Z_{l,t-1} \right)^2}_{\text{Instrumented Asset Pricing Moment Squared}} + \sum_{i=1}^{N} \sum_{j=0}^{K} \underbrace{\left(T^{-1} \sum_{t=1}^{T} \hat{\epsilon}_{i,t}(\theta)\, F_{j,t} \right)^2}_{\text{Projection Errors}}.$$

The risk prices $\{\omega_{j,t}\}$ and preference parameters (σ, δ) are not identified by these moments—the former because we consider only zero-beta portfolios, the latter because we do not include the consumption Euler equation as one of our moments (we consider these moments below). Because our problem is exactly identified, conditional on γ, the (α_i, β_{ij}) point estimates will be the usual OLS estimates. The caveat "conditional on γ" applies because one cannot construct the excess returns $R_{t+1} - R_{0,t}$ without an estimate of the zero beta rate.

To conclude our description of the estimation procedure, we describe how we determine the weight vector $w(\theta)$. Let $\Sigma_R(\theta)$ be an estimate of the covariance matrix of the excess returns $R_{t+1} - R_{0,t}$ given θ. We choose w to minimize the variance of the portfolio, $w' \cdot \Sigma_R(\theta) \cdot w$, subject to the constraints $w' \cdot \beta = \vec{0}$ and $w' \cdot \iota = 1$, where ι is a vector of ones.

We assume that Σ_R is of full rank and that ι does not lie in the span of β, so that the problem is feasible. Under these assumptions, the explicit solution to this problem is given by

$$w(\theta) = \Sigma_R(\theta)^{-1} \cdot \begin{bmatrix} \iota & \beta(\theta) \end{bmatrix} \cdot \left(\begin{bmatrix} \iota' \\ \beta(\theta)' \end{bmatrix} \cdot \Sigma_R(\theta)^{-1} \cdot \begin{bmatrix} \iota & \beta(\theta) \end{bmatrix} \right)^{+} \begin{bmatrix} 1 \\ \vec{0} \end{bmatrix}.$$

We discuss our reasons for using the minimum-variance zero-beta portfolio below.

We use the Ledoit and Wolf [2017] estimator for Σ_R. This covariance-matrix estimator is designed for minimum-variance portfolio problems, and has been shown by those authors to outperform other covariance matrix estimators with respect to the out-of-sample portfolio variance in such problems. Loosely, this is because it avoids over-fitting. We discuss the details of how we apply their estimator to our setting in the Appendix Section F.

3.3 Portfolio Interpretation

We next interpret our procedure in terms of the construction of a portfolio and the prediction of its returns. The portfolio weights $w(\theta)$ are the weights associated with the minimum-variance zero-beta unit-investment portfolio. When we weight the moments

$$\mathbb{E}\left[H(\theta) \cdot (R_{t+1} - R_{0,t}(\gamma)) \otimes Z_t\right] = 0$$

by $w(\theta)$, because $w(\theta)$ is a unit-investment zero-beta portfolio, we are solving for:

$$\mathbb{E}\left[(w(\theta)' \cdot R_{t+1} - R_{0,t}(\gamma)) \otimes Z_t\right] = 0.$$

Using the definition of $R_{0,t}(\gamma)$, (14), these are exactly the moments of the predictive regression

$$w(\theta)' \cdot R_{t+1} - R_{b,t} = \gamma' \cdot Z_t + \kappa_{t+1},$$

in which the excess return of the minimum-variance zero-beta unit-investment portfolio relative to the safe rate is predicted by the variables Z_t.

But for one (rather essential) caveat, our procedure is equivalent to the following steps:

1. For each asset i, run a time series regression to estimate α_i and β_i.

2. Use these betas, along with an estimate of the covariance matrix of excess returns Σ_R, to compute the portfolio weights $w(\cdot)$.

3. Predict the return of this portfolio using the Z_t variables, and call the predicted return the zero-beta rate.

The one caveat is that steps one and two are not feasible without an estimate of the zero-beta rate, because they both involve the excess returns $R_{t+1} - R_{0,t}$. The GMM procedure overcomes this by estimating $\theta = (\alpha, \beta, \gamma)$ simultaneously, in effect performing all three of these steps at once.

This equivalence motivates our use of a minimum-variance portfolio. In theory, all unit-investment, zero-beta portfolios should have the same expected return. Among these, the minimum-variance one is one whose return is the easiest to predict (in the sense that $\mathbb{E}[\kappa_{t+1}^2]$ is minimized), and hence will be the one for which the parameter γ is best estimated. This intuition explains why the minimum-variance portfolio plays a special role in the Shanken [1986] MLE procedure (described in Appendix Section G). That said, our weight matrix is not the efficient GMM weight matrix (which would allow for deviations from the OLS estimates of β to better fit the asset pricing moments). It instead resembles the GLS weight matrix used in standard cross-sectional asset pricing tests (chapter 12.2 of Cochrane [2009]).

It is also worth noting the difference between our procedure and standard cross-sectional asset pricing tests. Standard cross-sectional asset pricing tests take as given a series of excess returns, choose a weight vector that is highly exposed to β, and attempt to identify the prices of risk (the $\omega_{j,t}$ in our notation). In contrast, our procedure treats excess returns as unknown, chooses a weight vector that is orthogonal to β, and makes no attempt to identify the price of risk.

3.4 The Convenient SDF

Our model ascribes all of the spread $s_t \approx R_{0,t} - R_{b,t}$ to the convenience on safe bonds. Our estimation procedure embodies this assumption, via our assumption that the vector of ones, ι, does not lie in the span of the betas. This assumption, which is familiar from arbitrage pricing theory (Chamberlain and Rothschild [1983]) and is often viewed as a technical assumption (see, eg., assumption 3.iii in Kim et al. [2021]), has an important economic meaning in our context.

Our procedure estimates the zero-beta rate, $R_{0,t} = \mathbb{E}_t \left[\Lambda_{t+1}/\Lambda_t\right]^{-1}$, under the assumption that the relevant innovations of the SDF Λ_t are spanned by the factors F_t. However, it is not possible to build a risk-free portfolio using stocks. As a result, by the absence of arbitrage, given any stochastic process $\tilde{R}_{0,t}$, there is an SDF $\tilde{\Lambda}_t$ that prices stocks ($\mathbb{E}_t \left[\tilde{\Lambda}_{t+1}/\tilde{\Lambda}_t \times R_{i,t+1}\right] = 1$) and satisfies $\tilde{R}_{0,t} = \mathbb{E}_t \left[\tilde{\Lambda}_{t+1}/\tilde{\Lambda}_t\right]^{-1}$. In particular, one could choose $\tilde{R}_{0,t} = R_{b,t}$. These other SDFs attribute the spread between the zero-beta expected return our procedure recovers, $R_{0,t}$, and the process $\tilde{R}_{0,t}$ to an omitted factor.[11] All unit-investment, zero-beta portfolios must load equally on the omitted factor to explain why their expected return is $\tilde{R}_{0,t}$ and not $R_{0,t}$. Put another way, the vector of ones must lie in the span of the betas of the assets to the augmented set of factors (including both the omitted factor and F_t). The assumption that the vector of ones does not lie in the span of the betas thus selects, from the set of all SDFs that price equities, the "Convenient SDF" that attributes all of the spread s_t to convenience as opposed to risk premium. This, of course, is the perspective we adopt in this paper, and our econometric method embodies this view.

[11] It is without loss of generality to assume this factor is orthogonal to the included factors F_t, implying that the conditional variance of these other SDFs is higher than the conditional variance of Λ_{t+1}/Λ_t.

This point can be further illustrated using the example of the "betting against beta" model of Frazzini and Pedersen [2014] (but note that consumption is not modeled in their framework, and hence their model cannot speak to the main part of our analysis). In that model, leverage-constrained and unconstrained agents interact in financial markets; the leverage-constrained agents are analogous to the agents who value convenience in our model.[12] The leverage-constrained agents have a single-factor SDF (with the market as the single factor) that prices stocks, whose mean is the inverse of the zero-beta rate. The unconstrained agents have a two-factor SDF, whose factors are the market and a zero-beta portfolio, with a mean consistent with safe bond yields. The betas of each asset to these two factors sum to one (ι lies in the span of the betas). The two SDFs agree on equity prices, but differ in their conditional means. Our procedure, applied to this hypothetical economy, would recover the leverage-constrained investors' SDF given the market as the single factor. It cannot be applied to the two-factor SDF of the unconstrained agents, as our assumption that ι does not lie in the span of the betas is violated.

4 The Zero-Beta rate

4.1 Data

Our procedure requires a set of equity portfolio returns (the $R_{i,t}$), a set of factors (the $F_{k,t}$), a set of instruments (the $Z_{l,t}$), and consumption data (C_t). We will briefly describe the portfolios, factors, instruments, and consumption data we use in our main specification in this section. Additional details can be found in Appendix Section A, and results with alternative portfolios, factors, instruments, and consumption data can be found in Appendix Section E.

Equity Portfolios. Our main equity returns data consists of the equity returns in CRSP which can be matched to a firm in COMPUSTAT, excluding the bottom 20% of stocks by market value in each month. For each of these stocks, we compute a five-year trailing beta to the CRSP market return (using monthly data).[13] We then construct 27 (3x3x3) portfolios

[12] The key difference between these two types of agents is that leverage-constrained agents do not own any safe bonds. In the Frazzini and Pedersen [2014] model, any safe bonds in positive net supply must be owned by an un-modeled third class of agents.

[13] Novy-Marx and Velikov [2022] point out that the smallest deciles of stocks are likely to be less liquid than other stocks, and as result have betas that are attenuated relative to other stocks; this notably affects the conclusions of Frazzini and Pedersen [2014]. For this reason, in our main specification we exclude these stocks, and in our robustness exercises verify that our results are not meaningfully altered by their inclusion. The robustness of our results to the inclusion of these stocks likely stems from our use of betas based on monthly as opposed to daily data, which reduces impact of liquidity on betas.

based on (i) market beta, (ii) market capitalization (i.e. size), and (iii) book-to-market ratios (i.e. value). We augment these portfolios with the 49 industry portfolios (based on four-digit SIC codes) from Ken French's website,[14] and thus consider 76 stock portfolios in total. For additional details on the construction of these portfolios, see Appendix Section A.

There are two considerations that have guided our choices. First, we use beta-sorted portfolios to ensure that there is a wide variation across our portfolios in terms of their beta to the market. Our motive is evident from (15): an equity portfolio with a beta of one to the market is in fact insensitive to the level of the zero-beta rate. Second, we have included a variety of portfolios to ensure that it is possible to form a well-diversified zero-beta portfolio. Commonly used portfolios such as the Fama-French 25 size by value portfolios exhibit a strong factor structure; a portfolio constructed from only the FF25 and forced to have zero beta to the market, size, and value factors would load heavily on poorly estimated residuals.

Factors. Our main specification uses seven tradable factors: the five equity-related factors of Fama and French [2015], the return of a 6-10y Treasury bond portfolio over a one-month Treasury bill, and return of long-term corporate bonds over long term Treasury bonds (i.e. the Treasury bond and default factors of Fama and French [1993]). We have chosen these factors because they are standard in the literature and because they are thought to explain the cross-section of expected returns in the equity portfolios we study. Our use of the five-factor Fama-French model is motivated in particular by the results of Novy-Marx and Velikov [2022], who find that univariate-beta-sorted portfolios are correlated with the investment and profitability factors of Fama and French [2015]. Our inclusion of the bond return factors is motivated by our use of the term spread and excess bond premium as instruments (discussed below).

We augment these seven tradable factors with a consumption-based nominal SDF,

$$F_{8,t} = \frac{c_{t+1}^{-\sigma}/P_{t+1}}{c_t^{-\sigma}/P_t},$$

which arises from the model. Our main specification uses $\sigma = 5$; in robustness exercises, we present results with $\sigma = 1$, a linear consumption factor, and no consumption factor at all. Note that it is not essential in our framework that the prices of risk $\{\omega_{j,t}\}$ be identified, in contrast to the usual cross-sectional asset pricing exercise. Concretely, even though we include a consumption-based factor, our portfolios have almost no beta to that factor. This lack of variation would be a problem if we wanted to identify the price of consumption risk.

[14]Note that the portfolios include the smallest 20% of stocks; because the portfolios are not beta-sorted and are value-weighted, the inclusion of these stocks in industry portfolios is unlikely to affect our results.

But since our goal is instead to construct a zero-beta portfolio, the presence of irrelevant factors presents no particular difficulties. That is, the typical weak identification problem in cross-sectional asset pricing is not relevant for our procedure. A different weak-identification problem, related to the predictability of returns, is the main challenge. Likewise, our omission of portfolios sorted on investment, profitability, or beta to bond returns does not meaningfully alter our results.

Instruments. We have chosen our instruments with the goal of predicting either consumption growth or the spread between the zero-beta rate and bill yields—an instrument that predicts neither of these is likely to be irrelevant. Our main specification includes five instruments, all of which are available at the monthly frequency starting in 1973.

We include the Treasury bill yield, motivated by the results of Nagel [2016], who argues that spreads between near-money assets should be correlated with the level of nominal interest rates. We include a rolling average of the previous twelve months of inflation (specifically, log-changes in the CPI index), motivated in part by the result of Cohen et al. [2005], who find that the slope of the security market line varies with the level of inflation. Given that finding, it is natural to think that the intercept of the security market line (the zero-beta rate) might also depend on the level of inflation.

We also include three instruments—the term spread (10yr less 3m Treasury yields), the excess bond premium (EBP) of Gilchrist and Zakrajšek [2012], and the unemployment rate (U6)—that have been found to predict recessions. There is an extensive literature on predicting recessions using these and other variables; one recent example is Kiley [2022], who finds that similar variables[15] can be used to predict increases in unemployment rate at the one-year horizon. Insofar as consumption growth is predictable, our prior is that variables that predict recessions are likely to be useful in predicting consumption growth.

We employ these five instruments in our main specification. In robustness exercises, we also consider the cyclically-adjusted price-earnings (CAPE) ratio of Campbell and Shiller [1988], the shadow spread defined by the difference between actual bill yields and those implied by a smoothed term structure (Lenel et al. [2019]),[16] lagged consumption growth, and using the BAA corporate bond spread (vs Treasurys) in the place of the EBP. When using lagged consumption growth and lagged inflation as instruments, we use only $\log(c_{t-1}/c_{t-2})$ and a trailing average of inflation up to P_{t-1} to avoid issues related to measurement error in C_t and P_t, a standard practice in the literature.

[15]Kiley [2022] uses the BAA-Treasury spread instead of the EBP, and the other variables are defined slightly differently.

[16]We would like to thank Monika Piazzesi and Mortiz Lenel for suggesting that we include this spread.

Consumption Data. We use real NIPA non-durable goods and services consumption per capita growth as a our preferred consumption measure. We use this measure both because it is standard in the literature and because our interest lies in studying an aggregate Euler equation. In our robustness exercises, we also generate results with non-durable goods consumption only. Our use of aggregate consumption data should guide the interpretation of our estimates of the IES.[17]

Data Sample. Our data sample begins in January 1973, when all of our instrument variables become available, and ends in December 2020. Because some of our instruments involve lags and changes, our returns series begins in March 1973, and those returns are predicted using data from January and February 1973. Our sample is 574 months long.

4.2 Results

Table 1 presents a rescaled version of the γ coefficients estimated with our GMM procedure, with their associated standard errors. Recall that as part of our procedure, we have standardized our instruments to mean zero and unit variance. The coefficients and standard errors in Table 1 have been rescaled to undo the effects of the standardization. The first column of Table 1 also presents the results of a Wald test of the hypothesis that all of the coefficients (except the constant) are zero.

Several results are immediately apparent. First, our instruments have some ability to predict the return of our zero-beta portfolio—the Wald test p-value is under one percent. Our predictive regression is noisy, and only the EBP, term spread, and inflation coefficients are statistically significant at the 10% level on their own, but collectively our instruments are able to predict the zero-beta portfolio return.

Second, the spread over the Treasury bill yield is significant and economically substantial. The constant in Table 1 is the average monthly return of the zero-beta portfolio (because the instruments have been de-meaned). Our estimate of roughly 0.7% per month corresponds to an annualized excess return of over 8% per year. This estimate is consistent with other estimates of the average zero-beta return (Bali et al. [2017]) and with estimates of the arbitrage-based return on equity available to banks (Boyarchenko et al. [2018]).

[17]Vissing-Jørgensen [2002] and others have shown that the consumption growth of financial market participants is more sensitive to certain shocks than the consumption growth of non-participants. Assessing the extent to which the consumption Euler equation with the zero-beta rate holds for various sub-populations is an interesting direction for future research.

Table 1: Predicting the Zero-Beta Rate

	(1) GMM	(2) OLS (inf.)
Lrf	1.186	1.187
	(0.914)	(0.789)
Lump	0.105	0.105
	(0.0986)	(0.0965)
Lebp	-0.603	-0.603
	(0.342)	(0.309)
Ltsp	0.310	0.310
	(0.118)	(0.119)
L2cpi_rolling	-2.582	-2.586
	(1.175)	(1.048)
Constant	0.718	0.716
	(0.137)	(0.134)
Wald/F	21.46	5.012
p-value	0.000663	0.000167
Observations	574	574

Standard errors in parentheses

Notes: The first column of this table shows our point estimates and standard errors for the γ coefficients from our GMM estimation. The instruments have been centered, so the constant is the average monthly excess return of the zero-beta portfolio over the 1m Treasury bill yield. Standard errors in the first column are robust to heteroskedasticity and account for estimation error in the other parameters. The Wald statistic in the first column is a test of the hypothesis that all coefficients except the constant are zero; the p-value is shown below. The second column shows (for comparison purposes only) the results of a predictive regression of the excess return of the zero-beta portfolio over the 1m Treasury bill yield on our instruments.

The standard deviation of the return of our zero-beta portfolio less the zero-beta rate is about 3.2% per month, or 11% per year (consistent with the variances of the minimum-variance portfolios constructed by Ledoit and Wolf [2017]). This standard deviation is substantially below the standard deviation of the market return, which helps explain why we are able to reject the null of no-predictability for our portfolio despite the notorious difficulty of predicting the market return. The annualized average Sharpe ratio of the portfolio, computed using the excess returns over Treasury bills, is about 0.8, and often (because the return is predictable) above one. Our view is not that this portfolio is a new asset pricing factor uncorrelated with all other factors, but rather that the excess return of our portfolio over Treasury bills reflect the convenience of Treasury bills.

Third, our point estimates suggest that the zero-beta rate increases more than one-for-one with Treasury bill yields (i.e. the spread is increasing in bill yields) and is decreasing in the lagged inflation rate. The former is consistent with the interpretation of our spread as a convenience yield (Nagel [2016]). The latter is interesting in light the finding that inflation also affects the slope of the security market line (Cohen et al. [2005]).

Our point estimates also suggest that the unemployment rate and term spread positively predict the zero-beta portfolio returns, whereas the EBP negatively predicts those returns. That is, times when unemployment is low, the yield curve is inverted, and the excess bond premium is high are times when the return of our portfolio is predicted to be particularly negative. These are exactly the times when an increase in the unemployment rate (and a recession more generally) is particularly likely (Kiley [2022]). Our results thus suggest a connection between expected stock returns and business cycles.

The second column of Table 1 presents the results of an OLS regression in which our instruments are used to predict the return of our zero-beta portfolio. This regression is infeasible on its own: one needs to know the γ coefficients to construct the zero-beta portfolio return. The purpose of this column is to illustrate two points. First, the moment conditions that define this OLS regression are exactly the moment conditions used in our GMM procedure; as a result, up to numerical errors, the point estimates are identical. Second, the robust standard errors used in the OLS regression do not account for the fact that the zero-beta portfolio return is itself estimated. Nevertheless, they are strikingly similar to our GMM standard errors, which do take this into account, suggesting that the main source of uncertainty with respect to the γ parameters is the uncertainty associated with the predictive regression.

Figure 2 presents the nominal zero-beta rates generated by our point estimates for γ, along with Treasury bill yields. It also presents results for the zero-beta rate from a ridge-penalized estimation. In Appendix Section B, we include a ridge penalization for the construction of the zero-beta rate. This ridge penalty is designed to reduce over-fitting and improve the out-of-sample reliability of our estimate of the zero-beta rate.[18] For this reason, the ridge estimate would be our preferred estimate if we wished only to construct the zero-beta rate (as opposed to our exercise in the next section, which tests the consumption Euler equation). The ridge penalty attenuates the γ coefficients towards zero (with the exception of the constant term), which has the effect of moving the zero-beta rate towards the Treasury bill yield. The scale of the penalty is determined via cross-validation, with the goal of minimizing the out-of-sample squared forecast error of the zero-beta portfolio return. Figure 2 shows that this procedure results in a small shrinking of the spread between the zero-beta rate and Treasury bill yields, consistent with a reasonable degree of out-of-sample predictability for

[18] Our use of the ridge penalty explains why we standardize the instruments Z_t.

the zero-beta portfolio return.

Figure 2: Zero-beta rate and the T-Bill Yield

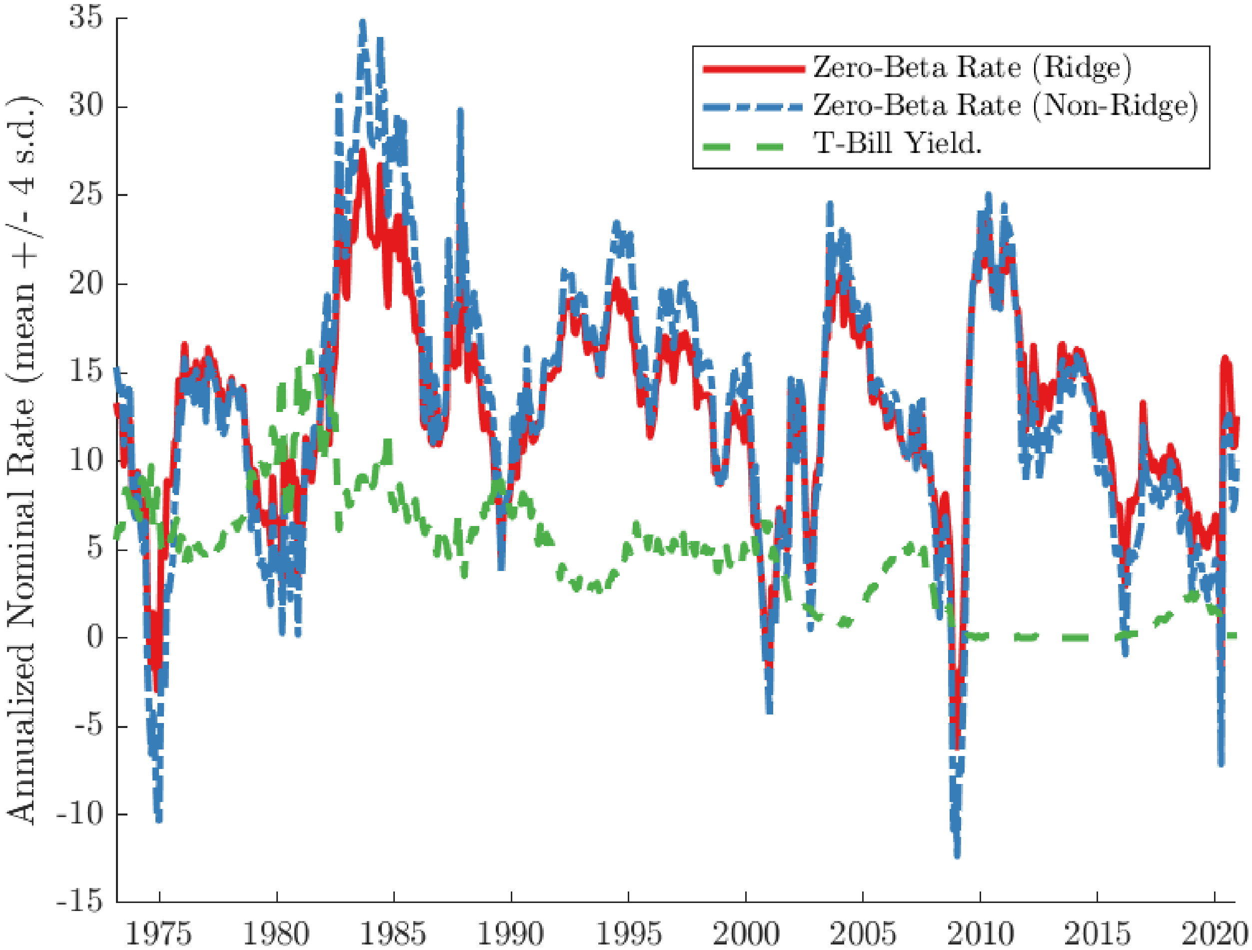

Notes: This figure plots the nominal zero-beta rate constructed from our GMM procedure with cross-validated ridge penalization (see Appendix Section B), the same constructed without ridge penalization, and the nominal Treasury bill yield. The zero-beta rates are constructed using our main specification of factors, instruments, and equity portfolios (as described in Section 4).

4.3 Misspecification of the Factors

Before proceeding, we will briefly comment on how misspecification in the factors would affect our results. If we omit an asset-pricing factor from F_t, we will incorrectly attribute to convenience on safe assets what is really risk premium. If the omitted risk premium is constant, or more generally not predictable by the instruments Z_t, the average level of the zero-beta rate will be wrong but the time-variation will be correct. In other words, our estimate of γ will be unbiased, except for the constant. Our estimates of γ will be biased only if the omitted risk premium is both time-varying and predictable by the Z_t. If we have included unnecessary asset pricing factors, our procedure will generate unbiased estimates of the zero-beta rate, but these estimates will be less precise than if we had excluded the unnecessary factors. We discuss the consequences of misspecification of the instruments below.

5 The Euler Equation

Our model shows that once we take into account the convenience yield of safe assets, the Euler equation should hold with the zero-beta rate, not with the safe rate. In this section we show that indeed the zero-beta rate fits an aggregate consumption Euler equation remarkably well.

We start with a linearized version of the Euler equation,

$$\mathbb{E}_t \left[\log\left(c_{t+1}/c_t\right)\right] = \sigma^{-1}\ln(\delta) + \sigma^{-1}\left(\log\left(R_{0,t}\right) - \mathbb{E}_t\left[\log\left(P_{t+1}/P_t\right)\right]\right),$$

That is, the real zero-beta rate should predict real consumption growth.

Figure 1 compares expected consumption growth with the expected real return of Treasury bills on the left panel, and with the zero-beta rate on the right panel. The two series are plotted on separate axes, which have been aligned in terms of their mean values. The axes have also been scaled to each represent +/- four standard deviations. In effect, these graphs are tests of the linearized Euler equation, with σ defined by the ratio of the standard deviations and δ set to ensure that equation holds at the average values. The contrast between these two figures is striking. The expected real Treasury return bears essentially no resemblance to expected consumption growth. In contrast, the real zero-beta rate co-moves strongly with expected consumption growth.

There is nothing mechanical about this result, but it's important to interpret it correctly. Both time series are projections on the predictors Z_t. If we had only one predictor Z_t, both time series would be perfectly aligned after rescaling, by construction. With only one predictor, the Euler equation has exactly enough degrees of freedom, δ and σ, to ensure it holds. But once we expand the set of predictors Z_t, this stops being a mechanical result. Shocks to different predictors could move the real zero-beta rate and expected consumption growth in different directions. In our main specification, with five predictor variables, there are certainly many other possible time series that could have been generated by our procedure. The failure of the Euler equation with the expected real Treasury bill return provides a placebo test. Our procedure could have concluded that either expected consumption growth or the real zero-beta rate resembled the expected real Treasury bill return. It instead found that those two series were proportional to each other, and neither resembled the expected real Treasury bill return. In Appendix Section C we also conduct a similar placebo test with long-term Treasury bonds, and find the Euler equation also fails for these bond returns.

The fact that the Euler equation holds with the zero-beta rate supports our claim that the zero-beta rate is the correct intertemporal price. Having said this, since the Euler equation does have a free intercept δ, the fact that it holds with the zero-beta rate does not provide much extra support for the average level of the zero-beta rate. That is, if our measurement of the zero-beta rate includes an omitted risk premium, but that omitted risk premium is either constant or not predictable by the Z_t, the zero-beta rate will still fit the aggregate Euler equation, even if its average level is off. The impatience rate δ will absorb the omitted risk premium. But the Euler equation does provide support for the time variation in the zero-beta rate. If the movements in the zero-beta rate mostly reflected predictable movements in the omitted risk premium, we would not expect the linearized consumption Euler equation to hold.

One potential concern with regards to Figure 1 is that the predictive regressions that define the zero-beta rate and consumption growth suffer from over-fitting.[19] In our view, this is almost certainly the case. Nevertheless, it is remarkable that the two predictive regressions, each of which is overfitting a distinct time series, nevertheless generate almost identical (up to a scaling factor) predictions. Figure 3 below compares our ridge-penalized zero-beta rate with a similarly constructed forecast of consumption growth.[20] Ridge penalization reduces the scale of both expected consumption growth and the zero-beta rate, but they remain roughly proportional.

Another potential concern is that the underlying zero-beta portfolio somehow, by coincidence, tracks realized consumption growth, so that whatever predictors Z_t we use, the conditional expectations will mechanically line up. However, since we use consumption as an asset-pricing factor, the innovations in the zero-beta portfolio and consumption growth are orthogonal. That is, consumption growth and the zero-beta portfolio are correlated only through the conditional expectations. We confirm in Appendix Figure 12 that our results are essentially unchanged when we do not include a consumption-related factor when constructing the zero-beta rate.[21] When the consumption-related factor is omitted, the zero-beta rate

[19] We view the potential overfitting of expected inflation as a less serious issue, due to the relative ease of forecasting inflation as opposed to consumption growth or portfolio returns.

[20] That is, consumption growth is estimated using a ridge regression, whose penalty parameter is selected via ten-fold cross-validation.

[21] This is not a surprising result: most stocks have close to zero beta with respect to consumption growth,

Figure 3: Results for Ridge Regressions

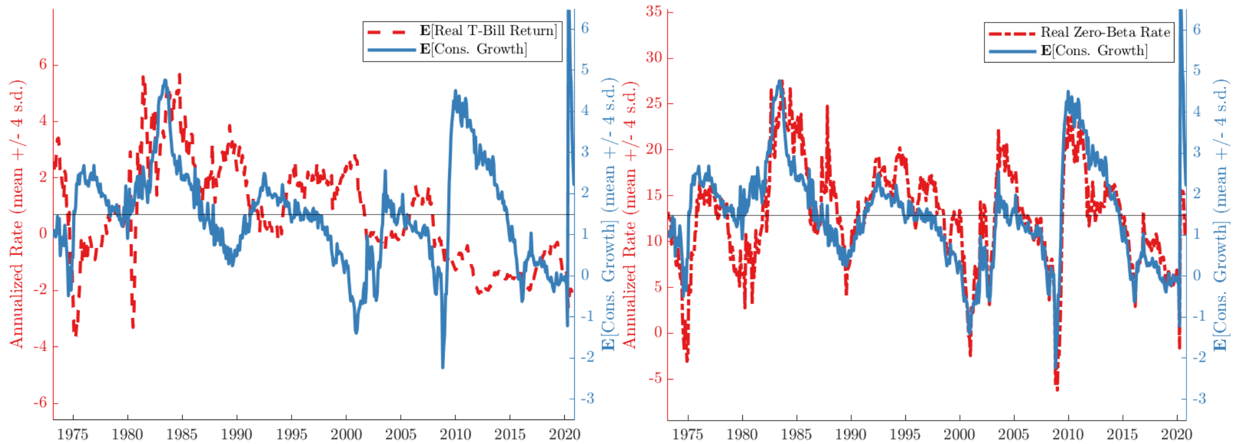

Notes: Both panels of this figure plot expected real returns against expected consumption growth, over time. Expected real returns are constructed using nominal rates (1m bill yields on the left, the zero-beta rate constructed using our main specification and ridge penalization on the right, see Appendix Section B) less expected inflation. Expected inflation is generated from predictive regressions using the instruments described in Section 4, which are the same instruments used to construct the zero-beta rate. Expected consumption growth is generated in the same fashion, but with a ridge regression whose penalty is chosen using ten-fold cross-validation. In both panels, the right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the relevant expected real return. All series are annualized.

is constructed without using consumption data in any way, further emphasizing that the alignment of the two series would not be expected absent theory.

5.1 Testing the Euler equation

While the graphical results in Figures 1 and 3 are suggestive, they are not a proper statistical test. In addition, the linearized Euler equation may be misleading if σ is large (low intertemporal elasticity), so we want to test the instrumented version of the non-linear Euler equation,

$$\mathbb{E}_t\left[\left(\delta\left(\frac{c_{t+1}}{c_t}\right)^{-\sigma}\frac{R_{0,t}}{P_{t+1}/P_t}-1\right)Z_{l,t}\right]=0. \tag{16}$$

In addition to the zero-beta rate, we will test the Euler equation with the Treasury bill yield $R_{b,t}$ and with the market return $R_{m,t+1}$. These are classical tests of the Euler equation and illustrate two polar issues. The Treasury bill return is very easy to predict (the nominal return is known ex-ante), and statistical tests strongly reject the Euler equation. The market return, in contrast, is hard to predict and one cannot reject anything. The latter is a "weak-instruments" problem (Yogo, 2004). The return of the zero-beta portfolio is considerably

especially after controlling for other factors that price the cross-section of stocks.

less volatile than the market and easier to predict, as shown in Table 1. However, the F-test in the (infeasible) OLS regression suggests that we may suffer from a weak instruments problem,[22] so we will test the Euler equation using weak-instruments-robust methods.

The method we adopt follows Stock and Wright [2000], and our main interest lies in testing whether or not, for some value of the IES σ^{-1}, the model cannot be rejected. The essence of our procedure is as follows.

1. Conjecture value of σ (the null hypothesis),
2. Estimate $\hat{\theta}(\sigma)$ (i.e. construct the zero-beta rate) as above,
3. Estimate $\hat{\delta}(\sigma)$ using the unconditional Euler (16) with $l = 0$ ($Z_{0,t} = 1$),
4. Test using the instrumented Euler moments (16) with $l > 0$.

Repeating this procedure for many possible values of σ allows us to construct a confidence set, and to test if the model is rejected for all values of σ. Note that the conjectured value of σ enters our estimation of the zero-beta rate via our use of the consumption risk factor. Because this factor plays a minimal role in the pricing of stocks, the estimated zero-beta rate is largely insensitive to the conjectured value of σ.

Specifically, we continue to construct the zero-beta rate as before, using the asset pricing moments as before as well as the instrumented Euler equation

$$g_t(\theta, \delta, \sigma) = \begin{bmatrix} \hat{\epsilon}_{t+1}(\theta) \otimes F_{t+1}(\gamma, \sigma) \\ H(\theta) \cdot (R_{t+1} - R_{0,t}(\theta)) \otimes Z_t \\ \left(\delta \frac{c_{t+1}^{-\sigma}}{c_t^{-\sigma}} \frac{R_{0,t}(\theta)}{P_{t+1}/P_t} - \right)) \otimes Z_t \end{bmatrix}$$

and the weight matrix

$$W(\theta) = \begin{bmatrix} I & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & w(\theta)w(\theta)' & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & e_0 e_0' \end{bmatrix},$$

where e_0 is a basis vector with one on the element corresponding to $Z_{0,t} = 1$ and zero otherwise. Conditional on σ, the parameters $\theta = (\alpha, \beta, \gamma)$ and δ are exactly identified. In particular, the estimated value of δ, $\hat{\delta}(\sigma)$, will be set so that the consumption Euler equation holds on average. Treating σ as known and minimizing over (θ, δ) results in the estimates $\hat{\theta}(\sigma)$ and $\hat{\delta}(\sigma)$ and the usual GMM standard errors (see Appendix Section H for details).

The moments

$$g_{l,t}(\theta, \delta, \sigma) = \left(\delta \frac{c_{t+1}^{-\sigma}}{c_t^{-\sigma}} \frac{R_{0,t}(\theta)}{P_{t+1}/P_t} - 1\right) Z_{l,t},$$

[22] It is below the critical values suggested by, e.g., Olea and Pflueger [2013].

for $l \in \{1, \ldots, L\}$, are not targeted in the estimation of $\hat{\theta}(\sigma)$ and $\hat{\delta}(\sigma)$. Let $g_{Test,t}(\theta, \delta, \sigma)$ be the vector of these moment conditions, and let $\hat{V}_{Test}(\sigma)$ be the variance-covariance matrix of $\frac{1}{T}\sum_{t=1}^{T} g_{Test,t}(\hat{\theta}(\sigma), \hat{\delta}(\sigma), \sigma)$. Under the null hypothesis of an IES of σ^{-1}, the test statistic

$$\hat{S}(\sigma) = \left(\frac{1}{T}\sum_{t=1}^{T} g_{Test,t}\left(\hat{\theta}(\sigma), \hat{\delta}(\sigma), \sigma\right)\right)' \cdot \hat{V}_{Test}(\sigma)^{-1} \cdot \left(\frac{1}{T}\sum_{t=1}^{T} g_{Test,t}\left(\hat{\theta}(\sigma), \hat{\delta}(\sigma), \sigma\right)\right)$$

is chi-square distributed with L degrees of freedom (Stock and Wright [2000], see Appendix Section H for the details of how their results can be applied to our problem).

Inverting this test statistic allows us to construct confidence sets. Specifically, we construct a 95% confidence set by computing $\hat{S}(\sigma)$ for values of σ between $\frac{1}{4}$ and 10, and comparing the $\hat{S}(\sigma)$ values to the 95th-percentile of a chi-squared distribution with L degrees of freedom.

Importantly, the same procedure can be applied to traditional Euler equations. If we replace $R_{0,t}(\theta)$ in (16) with either the Treasury bill yield $R_{b,t}$ or the CRSP market return $R_{m,t+1}$, we can apply the exact same procedure to estimate $\hat{\delta}(\sigma)$ using the unconditional Euler equation moment and then test on the instrumented Euler equation moments. In these cases θ does not enter the consumption Euler equation, and there is no need to estimate the zero-beta rate while simultaneously testing the consumption Euler equation for a different asset.

The advantage of this approach is that the same set of moments are being used to test the consumption Euler equation as applied to the three different assets (the Treasury bill, market portfolio, and zero-beta rate), consistent with recommendations of Cochrane [2009] (sections 11.5 and 11.6). If the test rejects for a given σ, it is because the instruments Z_t collectively predict

$$\frac{c_{t+1}^{-\sigma}}{c_t^{-\sigma}} \frac{R_{t+1}(\theta)}{P_{t+1}/P_t},$$

where R_{t+1} is one of $R_{0,t}$, $R_{b,t}$, or $R_{m,t+1}$. If the test fails to reject, it is because the null of no predictability cannot be rejected.

Our results are shown below, in Figure 4. Consistent with the findings of Hansen and Singleton [1983], Dunn and Singleton [1986], and Yogo [2004], we are able to reject the hypothesis that the Euler equation holds when applied to the Treasury bill yield. Intuitively, our instruments have some ability to predict consumption growth, and are certainly able to predict the real Treasury bill return. The finding that expected real Treasury bill returns are not proportional to consumption growth (Figure 1) essentially guarantees the rejection of the test.

Figure 4: Tests of Specification by IES

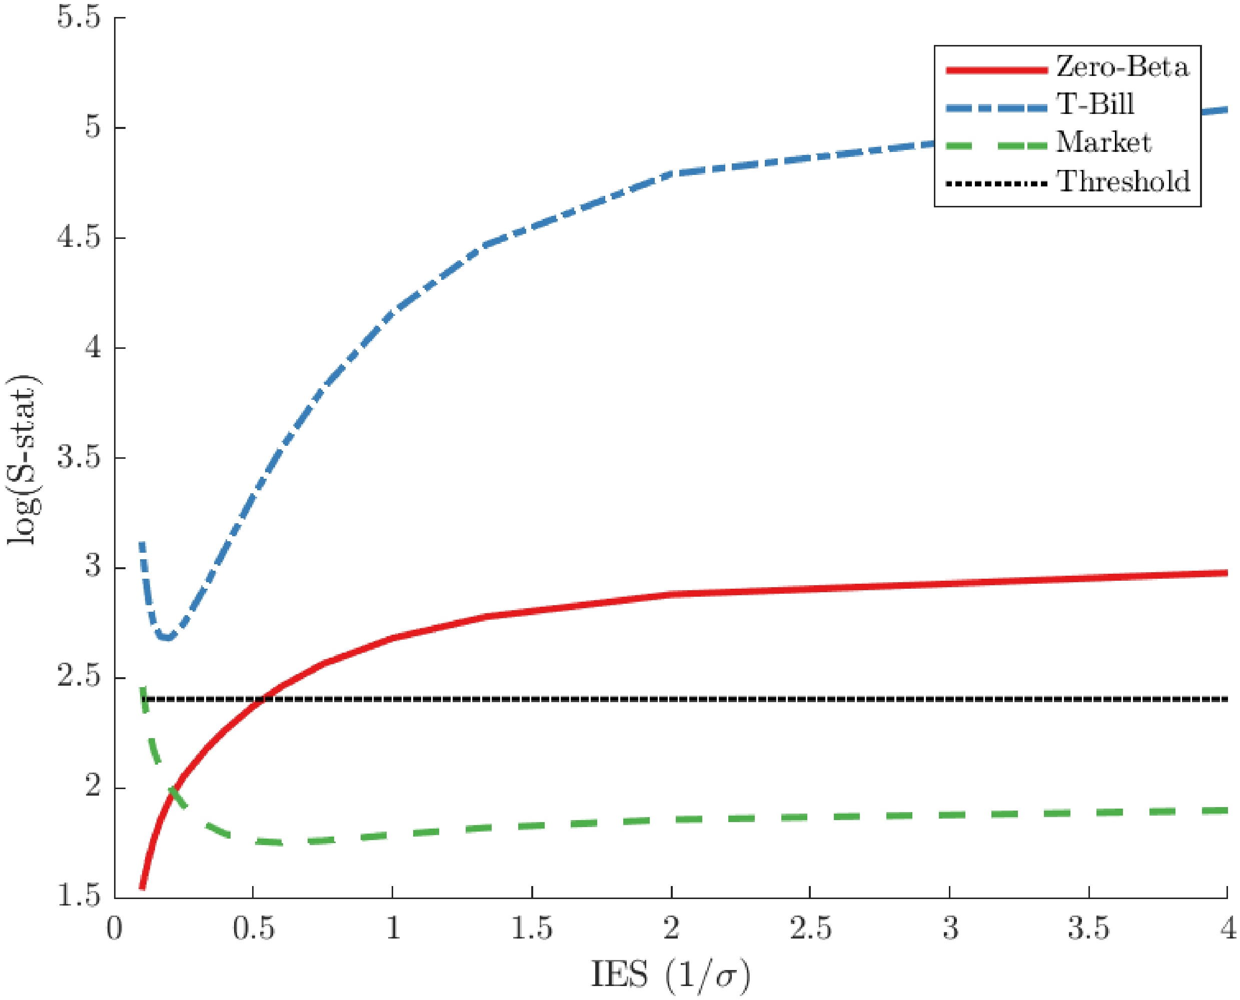

Notes: This figure plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5). The set of σ for which $\hat{S}(\sigma)$ is less than the threshold is the S-set (Stock and Wright [2000]).

Again consistent with the findings of Yogo [2004], our test fails to reject the consumption Euler equation as applied to the market index for almost any values of σ. The market return is volatile relative to consumption growth; for most values of σ, the moment condition is essentially identical to the market return.[23] The market return is difficult to predict, and it is therefore unsurprising the procedure is unable to reject the model. This lack of predictability is the weak-identification problem.

In contrast, the test applied to the zero-beta rate is able to reject the model for some but

[23]This ceases to be true for σ sufficiently large; in this case, the predictability of consumption growth allows the test to reject in some specifications. However, there is a separate issue that arises when σ becomes excessively large, discussed below.

not all values of σ. In our baseline analysis, values of σ below 1.6 (IES above 0.6) are rejected, and values of σ above 1.6 (IES below 0.6) cannot be rejected. This again is unsurprising in light of Figure 1. Our instruments are able to predict both consumption growth and the real zero beta return, and when the latter is scaled down by a factor of five (i.e. $\sigma \approx 5$), the two series are remarkably similar.

A key limitation of our test arises from the non-linear nature of the Euler equation. When σ becomes large (say, $\sigma = 20$), the realized SDF becomes very large on the date with the largest consumption drop in the data sample (in our sample, April 2020). The realization on this date dwarfs all other realizations of the SDF; as a consequence, the variance matrix $\hat{V}_{test}(\sigma)$ becomes almost singular. All of our test statistics (for the Treasury bill, zero-beta rate, and the market) converge towards one in this case. Intuitively, it is as-if there are only two dates in our data set (April 2020 and not-April-2020), and the model is untestable in this case. For this reason, we restrict our analysis to $\sigma \leq 10$.

5.2 Interpretation

The fact that there is a stable aggregate consumption Euler equation with the zero-beta rate is our central result. The zero-beta rate may seem too large and too volatile, but the fact that it fits the Euler equation makes it hard to dismiss. Through the lens of the Euler equation, it seems to be the correct intertemporal price. The fact that the Euler equation holds for any market return at all is surprising in itself. It is well-known that the Euler equation does not hold with the Treasury bill yield, and the conventional view is that there is no stable aggregate consumption Euler equation in the data. The implicit assumption underlying this view is that a safe interest rate is the correct intertemporal price, which implies that convenience yields are small. In contrast, our results suggest that the zero-beta rate is the correct intertemporal price, which implies that safe rates reflect the behavior of large and time-varying convenience yields.

5.3 Misspecification of the Instruments and Alternative Specifications

Our results are robust to the inclusion of extraneous (i.e. weak or irrelevant) instruments, due to our use of weak-instrument-robust inference methods. If we have excluded relevant instruments (ones that predict the zero-beta return, consumption growth, or both), our results will be biased only to the extent that these instruments allow us to reject the consumption Euler equation. Our prior is that it is generally hard to predict either consumption growth or the return on a stock portfolio, but we cannot rule out the possibility of an omitted instrument.

To further allay concerns about misspecification, in Appendix Section E we present the analog of Figures 1 and 4 (our "graphical" and GMM tests) for alternative specifications of our model. Table 4 in that appendix presents our γ coefficient estimates for each of these specifications. We consider the following alternatives:

- With the bottom two deciles of stocks included in the data sample.
- With Fama-French 5-factor sorted + industry portfolios and instead of 3-factor + industry portfolios.
- With a $\sigma = 1$ instead of $\sigma = 5$ used to construct the consumption factor
- With a linear consumption factor instead of a non-linear consumption factor in the SDF.
- With no consumption factor in the SDF.
- With only the market factor, and with only the FF3 factors.
- With our preferred instruments, using the BAA-AAA spread in the place of the excess bond premium.
- With our preferred instruments plus the cyclically adjusted price-earnings (CAPE) ratio.
- With our preferred instruments and a lag of consumption growth.
- With our preferred instruments and the "shadow spread" used by Lenel et al. [2019].
- With instruments-by-factor interactions as factors (allowing for time-varying betas).
- With non-durable goods consumption per capita as opposed to non-durable goods + services per capita.
- With a data sample ending in December 2019 (pre-COVID 19).

Our results are essentially identical across all of these variations, demonstrating that our key conclusion (that the consumption Euler equation cannot be rejected when applied to the zero-beta rate) is robust to various perturbations with regards to our choices of equity portfolios, factors, instruments, and consumption data.

6 Monetary Policy

Monetary policy aims to exploit the Euler equation as a structural relationship. Raising interest rates is meant to lower consumption and output. However, there is a large and time-varying spread between the safe rates that central banks control as policy instruments and the zero-beta rate that enters the Euler equation. If this spread is exogenous to monetary policy, raising the safe rate also raises the zero-beta rate one-to-one. This is what is assumed in applied work such as Smets and Wouters [2003, 2007]. But, as Chari et al. [2009] and Fisher [2015] point out, if the spread is endogenous to monetary policy, movements in the safe policy rates may have surprising effects on the zero-beta rate.

In this section, we explore the effects of a monetary policy shock. We carry out an exercise analogous to the one behind Figures 1 and 3, but conditional on monetary shocks. We regress the predictors Z_t on measures of monetary policy shocks, and then use our predictive regression coefficients to calculate the effect on expected consumption growth, expected real Treasury bill returns, and the real zero-beta rate (call them γ^c, γ^b, and γ^0, respectively). Implicitly, we are assuming that the relationship between Z_t and these variables is structural. We run regressions of the form:

$$\gamma^j \cdot (Z_{t+h} - Z_{t-1}) = \phi_{0,h} + \phi_{1,h} \cdot mpshock_t + \epsilon_{t+h}, \tag{17}$$

where $mpshock_t$ is either the Romer and Romer [2004] shock or the Nakamura and Steinsson [2018] shock,[24] aggregated to the monthly frequency, and $j \in \{c, b, 0\}$. The Romer and Romer [2004] shocks are available from 1973 through 2007; the Nakamura and Steinsson [2018] shocks are available from 2000 through 2019. We use the coefficients from our ridge estimation (see Appendix Section B), as these are less likely to suffer from over-fitting. Figure 5 shows our results graphically.

Both of the shocks are estimated to increase Treasury bill yields on impact by roughly the same amount, but the Romer and Romer [2004] shocks are more transitory. In both cases, the zero-beta rate is estimated to fall in response to the shock, but the effect of the Romer and Romer [2004] shocks is considerably smaller. In Appendix Section D we decompose the impact of each predictor in Z_t. A higher bill yield raises the zero-beta rate, other things equal. But monetary shocks also flatten the term yield curve and increase credit spreads, and these lower the zero-beta rate. These effects dominate the direct effect of the increase in Treasury bill yields.

This result may seem surprising. A monetary contraction is supposed to work by raising the cost of current consumption, but it actually makes it cheaper. However, it is perfectly

[24] As updated by Wieland and Yang [2020] and Acosta [2022], respectively.

Figure 5: Effects of a Monetary Policy Shock on Real Rates and Consumption Growth

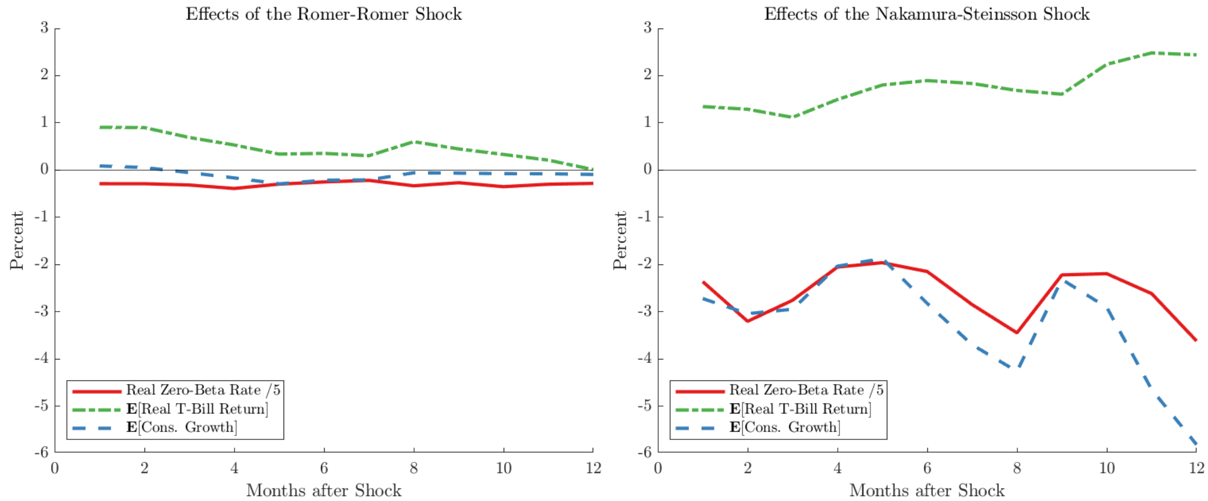

Notes: This figure plots the coefficients $\phi_{1,h}$ from our monetary policy shock regressions (17), for three different γ^j vectors (ones for the real zero-beta rate, the expected real Treasury bill return, and expected consumption growth), with a horizon h from one to twelve months. The γ^j vectors are generated from our ridge specification (as in Figure 3 and Appendix Section B), and the vector for the real zero-beta rate is scaled down by a factor of five, consistent with an IES of 0.2. The left panel plots the results for the Romer and Romer [2004] shocks, the right for the Nakamura and Steinsson [2018] shocks. Both shocks are scaled to represent a one percent increase in the federal funds rate on impact.

in line with facts and theory. It is well-known that a monetary contraction lowers expected consumption growth, also shown in Figure 5. Through the Euler equation, the interest rate should fall as households, correctly expecting lower consumption in the future, try to save. Our point estimates are consistent with the Euler equation holding in response to monetary shocks with an IES of roughly 0.2 ($\sigma = 5$), in line with our previous results.

This is not exactly a restatement of our previous results in Figures 1 and 3—that the Euler equation works with the zero-beta rate—because this is conditional on a monetary shock. That is, while the zero-beta rate and expected consumption growth strongly co-move because they load similarly on Z_t, the relationship is not exact. It could have been the case that consumption growth and zero-beta returns loaded differently on the specific combination of movements in Z_t that results from a monetary policy shock, leading them not to conditionally co-move. To be clear, the GMM analysis in the previous section fails to reject the Euler equation for any source of variation in Z_t, including monetary shocks, once estimation errors are all properly accounted for. But the concern is that because monetary policy shocks are a small part of the overall variation in Z_t, we may fail to reject even if the point estimates of the conditional co-movement are inconsistent with the Euler equation (because those point estimates have large standard errors). Here we show that the point estimates are consistent with the Euler equation conditional on monetary shocks.

Through the lens of the Euler equation, it is the rise of the expected real Treasury bill return that is puzzling in light of the fall of consumption growth after a monetary policy shock. Conversely, if one takes the rise in the Treasury yield as the defining feature of a contractionary shock, the fall in consumption growth is puzzling if one has in mind the conventional Euler equation applied to the Treasury return. In other words, if we use the expected real Treasury bill return, the Euler equation fails not only unconditionally, but also conditionally on a monetary shock. A standard way of addressing the conditional failure of the Euler equation is to introduce consumption habit formation into the model (as in Christiano et al. [2005] or Smets and Wouters [2007]), which generates a modified Euler equation and the kind of "hump-shaped" impulse response found in the data. However, this does not resolve the unconditional failure of the Euler equation (Canzoneri et al. [2007]), which is why Smets and Wouters [2007] still need wedges in this equation to match the data. It is also inconsistent with evidence on marginal propensities to consume (Auclert et al. [2020]). Specifically, after receiving a transfer, households tend to spend a substantial fraction of the transfer immediately, with the level of increased spending decaying over time. This pattern, which those authors call "micro jumps," is inconsistent with habit preferences. They propose informational frictions to reconcile the micro evidence with the hump-shaped response of impulse response functions.

Our results suggest instead that the Euler equation holds both unconditionally and conditionally on a monetary shock. Habits or informational frictions on the consumption side seem less important, at least at the aggregate level, although frictions on the supply side or in the transmission of monetary policy might still be important (e.g. it takes time for firms to fire workers and adjust productions plans, and trading frictions may affect money velocity).

6.1 Interpreting the effect of monetary policy shocks

Our results call into question the role of safe rates in the transmission of monetary policy. After a monetary contraction safe rates go up on impact, but this mostly reflects a reduction of convenience spreads. Monetary policy can still be correctly described in terms of policy rules for safe rates, just as it could be described in terms of any other endogenous variable. But the central fact is that the intertemporal price of consumption does not rise on impact, it falls. So how does a monetary contraction cause a reduction in consumption growth in the first place?

We propose an example of how this may come about in the model from Section 2. This is a standard New Keynesian model, augmented with convenience on safe assets. To make the

Figure 6: Effects of a Monetary Policy Shock in the Model

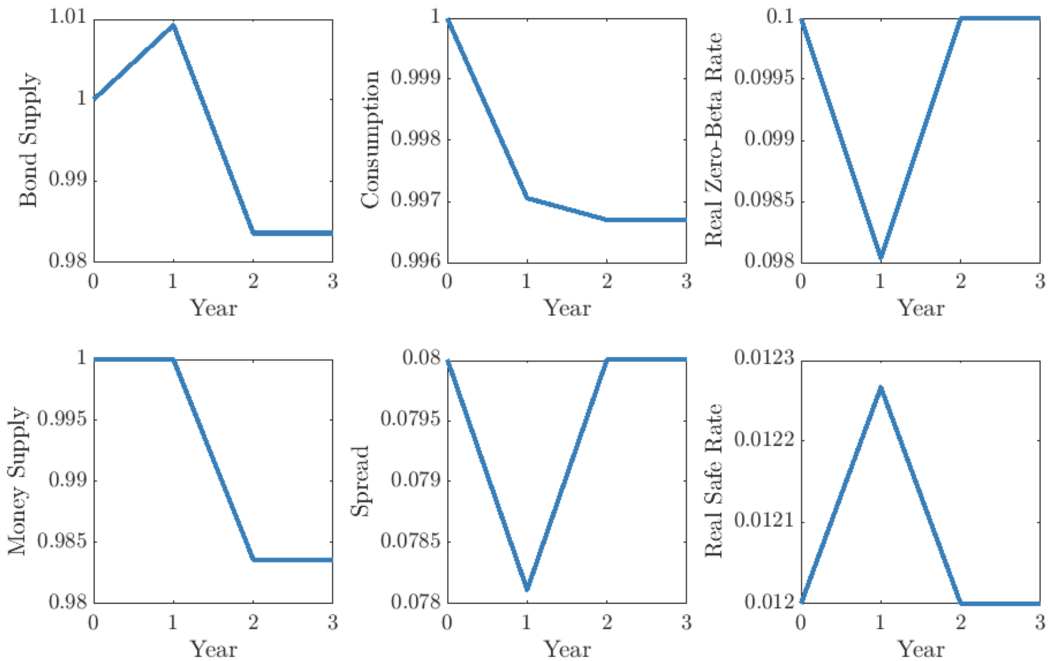

Notes: This figure plots the effects of a monetary contraction. In period one, the central bank sells bonds in exchange for base money. This contracts the supply of broad money (labeled Money Supply) with a delay, in period two, and the supply of safe bonds (Bond Supply) contracts in period two to restore the convenience spread (Spread) to its original value. The impact of this shock on the level of consumption, the zero-beta rate, and the safe bond rate is shown in the remaining three sub-figures.

mechanism transparent, everything happens in three periods, prices are completely sticky and the shock is permanent. The economy starts in period 0 at a steady state. In period 1, the Central Bank conducts an (unexpected) open market operation that increases the supply of safe bonds. The money multiplier takes one period to work its way through the banking system, so M only falls in period 2, and remains permanently low. From period 2 onwards, the supply of safe bonds adjusts to bring the safe rate spread back to steady state, and all quantities and prices are constant. What we have in mind is that the private supply of safe assets is endogenous and adjusts with some delay until the spread returns to its steady state value. We pick parameters to match a steady state spread of 8% and real zero-beta rate 10%, and pick $\sigma = 5$. Figure 6 shows the impulse responses.

The first column shows the shock itself. In period 1 there is an increase in the supply of bonds of just under 1%, and from period 2 onwards there is a permanent fall in the money supply of about 1.5%. The magnitudes are chosen to match the effect of an average-sized Nakamura and Steinsson [2018] shock that raises the safe rate by 2.7 bps.[25] The bond supply

[25]The movements in money and safe bond supply are relatively large compared to the movements in interest rates they generate. This is because the log specification of preferences implies a high interest-elasticity of

falls in period 2 to return spreads back to their steady state.

The second column shows the effect on consumption and spreads. Consumption falls on impact and then falls further when the money supply actually contracts in period 2, so that consumption *growth* also falls by 3.6 bps. The spread falls on impact due to both the fall in consumption (which reduces liquidity demand) and the increase in bond supply.

The third column shows the effect on interest rates. The zero-beta rate falls by 20 bps.With $\sigma = 5$ this is consistent with the magnitude of the fall in consumption growth. The fall in spreads and in zero-beta rates push the safe rate in opposite directions. In this example, the spread falls enough that the safe rate rises by 2.7 bps. All these magnitudes are in line with Figure 5 (scaled by the size of the shock).

Overall, the shock looks like a relatively standard monetary contraction: an open market operation that contracts the money supply with some delay, raises the safe interest rate and lowers consumption and consumption growth. It may seem surprising that the relevant interest rate for intertemporal decisions, the zero-beta rate, falls instead of rising, but this is actually consistent with the intertemporal pattern of money supply and consumption. The assumption that the shock to money supply is permanent and prices completely fixed implies that consumption remains permanently depressed. If we allowed money supply to revert to its original level, or prices to eventually adjust, consumption would return to its original level. In this case, the Euler equation implies that there would be a period of above-steady-state zero-beta rates as consumption recovers, and the contemporaneous impact of the shock on the level of consumption would be different. The purpose of this analysis is to show that a rise in safe rates accompanied by a fall in consumption growth and in the zero-beta rate is consistent with a basic New Keynesian model augmented with convenience on safe assets.

7 Conclusion

The interest rate plays a prominent role in macro and asset-pricing. The conventional view is that safe rates are the correct intertemporal price, which implies that there is no stable aggregate Euler equation and that convenience yields are small. We propose instead that the zero-beta rate is the correct intertemporal price, which implies that there is a stable aggregate Euler equation and that convenience yields are large and volatile. Our perspective is supported by the striking relationship in the time series between zero-beta rates and expected consumption growth.

money and safe bond demand.

References

Miguel Acosta. The perceived causes of monetary policy surprises. *Published Manuscript URL https://www1. columbia. edu/˜ jma2241/papers/acosta_jmp. pdf*, 2022.

Adrien Auclert, Matthew Rognlie, and Ludwig Straub. The intertemporal keynesian cross. 2018.

Adrien Auclert, Matthew Rognlie, and Ludwig Straub. Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Technical report, National Bureau of Economic Research, 2020.

Malcolm Baker and Jeffrey Wurgler. Do strict capital requirements raise the cost of capital? bank regulation, capital structure, and the low-risk anomaly. *American Economic Review*, 105(5):315–20, 2015.

Malcolm Baker, Mathias F Hoeyer, and Jeffrey Wurgler. Leverage and the beta anomaly. *Journal of Financial and Quantitative Analysis*, 55(5):1491–1514, 2020.

Turan G Bali, Stephen J Brown, Scott Murray, and Yi Tang. A lottery-demand-based explanation of the beta anomaly. *Journal of Financial and Quantitative Analysis*, 52(6): 2369–2397, 2017.

Florin Bilbiie. Monetary policy and heterogeneity: An analytical framework. 2021.

Fischer Black. Capital market equilibrium with restricted borrowing. *The Journal of business*, 45(3):444–455, 1972.

Fischer Black, Michael C Jensen, Myron Scholes, et al. The capital asset pricing model: Some empirical tests. 1972.

Nina Boyarchenko, Thomas M Eisenbach, Pooja Gupta, Or Shachar, and Peter Van Tassel. Bank-intermediated arbitrage. *FRB of New York Staff Report*, (858), 2018.

John Y Campbell and Robert J Shiller. Stock prices, earnings, and expected dividends. *the Journal of Finance*, 43(3):661–676, 1988.

John Y Campbell and Robert J Shiller. Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies*, 58(3):495–514, 1991.

John Y Campbell, Andrew W Lo, A Craig MacKinlay, and Robert F Whitelaw. The econometrics of financial markets. *Macroeconomic Dynamics*, 2(4):559–562, 1998.

Matthew B Canzoneri, Robert E Cumby, and Behzad T Diba. Euler equations and money market interest rates: A challenge for monetary policy models. *Journal of Monetary Economics*, 54(7):1863–1881, 2007.

Gary Chamberlain and Michael Rothschild. Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica: Journal of the Econometric Society*, pages 1281–1304, 1983.

Varadarajan V Chari, Patrick J Kehoe, and Ellen R McGrattan. New keynesian models: not yet useful for policy analysis. *American Economic Journal: Macroeconomics*, 1(1): 242–66, 2009.

Lawrence J Christiano, Martin Eichenbaum, and Charles L Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of political Economy*, 113(1): 1–45, 2005.

John H Cochrane. *Asset Pricing:(Revised Edition)*. Princeton university press, 2009.

Randolph B Cohen, Christopher Polk, and Tuomo Vuolteenaho. Money illusion in the stock market: The modigliani-cohn hypothesis. *The Quarterly journal of economics*, 120(2): 639–668, 2005.

Wenxin Du, Alexander Tepper, and Adrien Verdelhan. Deviations from covered interest rate parity. *The Journal of Finance*, 73(3):915–957, 2018.

Kenneth B Dunn and Kenneth J Singleton. Modeling the term structure of interest rates under non-separable utility and durability of goods. *Journal of Financial Economics*, 17 (1):27–55, 1986.

Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1):3–56, 1993.

Eugene F Fama and Kenneth R French. A five-factor asset pricing model. *Journal of financial economics*, 116(1):1–22, 2015.

Jonas DM Fisher. On the structural interpretation of the smets–wouters "risk premium" shock. *Journal of Money, Credit and Banking*, 47(2-3):511–516, 2015.

Andrea Frazzini and Lasse Heje Pedersen. Betting against beta. *Journal of financial economics*, 111(1):1–25, 2014.

Simon Gilchrist and Egon Zakrajšek. Credit spreads and business cycle fluctuations. *American economic review*, 102(4):1692–1720, 2012.

Mikhail Golosov and Robert E Lucas. Menu costs and phillips curves. *Journal of Political Economy*, 115(2):171–199, April 2007.

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

Robert E Hall. Intertemporal substitution in consumption. *Journal of political economy*, 96 (2):339–357, 1988.

Lars Peter Hansen and Kenneth J Singleton. Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica: Journal of the Econometric Society*, pages 1269–1286, 1982.

Lars Peter Hansen and Kenneth J Singleton. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of political economy*, 91(2):249–265, 1983.

Trevor Hastie, Robert Tibshirani, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.

Harrison Hong and David A Sraer. Speculative betas. *The Journal of Finance*, 71(5): 2095–2144, 2016.

Michael T Kiley. Unemployment risk. *Journal of Money, Credit and Banking*, 54(5):1407–1424, 2022.

Soohun Kim, Robert A Korajczyk, and Andreas Neuhierl. Arbitrage portfolios. *The Review of Financial Studies*, 34(6):2813–2856, 2021.

Arvind Krishnamurthy and Annette Vissing-Jorgensen. The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267, 2012.

Dirk Krueger and Hanno Lustig. When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)? *Journal of Economic Theory*, 145(1):1–41, 2010.

Olivier Ledoit and Michael Wolf. Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets goldilocks. *The Review of Financial Studies*, 30(12):4349–4388, 2017.

Olivier Ledoit and Michael Wolf. Analytical nonlinear shrinkage of large-dimensional covariance matrices. *The Annals of Statistics*, 48(5):3043 – 3065, 2020. doi: 10.1214/19-AOS1921. URL `https://doi.org/10.1214/19-AOS1921`.

Moritz Lenel, Monika Piazzesi, and Martin Schneider. The short rate disconnect in a monetary economy. *Journal of Monetary Economics*, 106:59–77, 2019.

Alejandro Lopez-Lira and Nikolai L Roussanov. Do common factors really explain the cross-section of stock returns? *Available at SSRN 3628120*, 2020.

Stefan Nagel. The liquidity premium of near-money assets. *The Quarterly Journal of Economics*, 131(4):1927–1971, 2016.

Emi Nakamura and Jón Steinsson. High-frequency identification of monetary non-neutrality: the information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330, 2018.

Robert Novy-Marx and Mihail Velikov. Betting against betting against beta. *Journal of Financial Economics*, 143(1):80–106, 2022.

José Luis Montiel Olea and Carolin Pflueger. A robust test for weak instruments. *Journal of Business & Economic Statistics*, 31(3):358–369, 2013.

Monika Piazzesi and Martin Schneider. Payments, credit, and asset prices. 2021.

Christina D Romer and David H Romer. A new measure of monetary shocks: Derivation and implications. *American economic review*, 94(4):1055–1084, 2004.

Jay Shanken. Testing portfolio efficiency when the zero-beta rate is unknown: a note. *The Journal of Finance*, 41(1):269–276, 1986.

Frank Smets and Raf Wouters. An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European economic association*, 1(5):1123–1175, 2003.

Frank Smets and Rafael Wouters. Shocks and frictions in us business cycles: A bayesian dsge approach. *American economic review*, 97(3):586–606, 2007.

James H Stock and Jonathan H Wright. Gmm with weak identification. *Econometrica*, 68 (5):1055–1096, 2000.

Jules H van Binsbergen, William F Diamond, and Marco Grotteria. Risk-free interest rates. Working paper, National Bureau of Economic Research, 2019.

Annette Vissing-Jørgensen. Limited asset market participation and the elasticity of intertemporal substitution. *Journal of political Economy*, 110(4):825–853, 2002.

Iván Werning. Incomplete markets and aggregate demand. 2015.

Johannes F Wieland and Mu-Jeung Yang. Financial dampening. *Journal of Money, Credit and Banking*, 52(1):79–113, 2020.

Motohiro Yogo. Estimating the elasticity of intertemporal substitution when instruments are weak. *Review of Economics and Statistics*, 86(3):797–810, 2004.

Online Appendix for "The Zero-Beta Rate," Di Tella, Hébert, Kurlat, Wang

A Data Details

A.1 Equity Portfolios

We use equity returns in CRSP which can be matched to a firm in COMPUSTAT, from 1973 to 2020, excluding the bottom 20% of stocks by market value in each month. The CRSP returns are augmented with the delisted returns also from CRSP. For each firm, we compute the book-to-market ratio, market value, operational profitability and investment according to Fama and French [1993, 2015].

1. Book-to-market ratio

   For portfolio in year t, it is measured with accounting data for the fiscal year ending in year $t-1$ and is the ratio between book equity (BE) and market value (ME). Book equity at t is shareholder equity (SEQ) minus deferred taxes and investment tax credit (TXDITC) minus preferred stock redemption value (PSTKRV).

2. Market value

   For portfolio in year t, it is measured with accounting data for the fiscal year ending in year $t-1$ and is share outstanding (SHROUT) times price (PRC).

3. Operational profitability

   For portfolio in year t, it is measured with accounting data for the fiscal year ending in year $t-1$ and is revenues (REVT) minus cost of goods sold (COGS), minus selling, general, and administrative expenses (XSGA), minus interest expense (XINT) all divided by book equity.

4. Investment

   For portfolio in year t, is the change in total assets (AT) from the fiscal year ending in year $t-2$ to the fiscal year ending in $t-1$, divided by $t-2$ total assets.

We compute the market beta of each stock using rolling 5-year, monthly linear regressions with the market return provided on Ken French's website.[26] We limit the sample to have at least 24 months of data points in the 5-year window.

[26] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/f-f_factors.html

At the end of each June, stocks are allocated to three groups according to NYSE breakpoints, on 30% percentile and 70% percentile, with respect to: Market Beta, Book-to-Market, Market Value, Operational Profitability and Investment. We then take the intersections of these groups to create portfolios. In particular, we construct 27 (3x3x3) portfolios on market beta, size, and book-to-market. We then augmented these portfolios with the 49 industry portfolios provided on Ken French's website.[27] These 76 portfolios are our baseline test assets.

In our robustness exercises, we include two additional sorts on market beta, size, and operational profitability, 27 (3x3x3), and on market beta, size, and investment, 27 (3x3x3).

A.2 Factors

The Fama and French factors are downloaded directly from Ken French's website. The Treasury bond factor is the return of the 6-10y Treasury bonds over the one-month Treasury bill (the latter as defined below).[28] The default factor is the return of long-term corporate bonds less the return of long-term Treasury bonds.[29] The consumption factor is built using the same consumption series used when testing the Euler equation.

A.3 Main Specification Instruments

1. Treasury bill yield

   One-month Treasury bill yield from Fama and French [2015].

2. Rolling average inflation

   Rolling average of the previous twelve months of inflation, which is the log-change in CPI index. CPI index from FRED[30].

3. Term spread

   Difference in the yields of 10-year treasury bond[31] and 1-month treasury bill.

4. Excess bond premium

[27] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_49_ind_port.html

[28] Specifically, we use the Fama maturity portfolio with maturity greater than 60 months and less than 120 months from CRSP Treasury as our long-term bond return measure.

[29] We use the ICE BofA 15+ Year US Corporate Index Total Return Index from FRED (https://fred.stlouisfed.org/series/BAMLCC8A015PYTRIV) and the 10+ Fama bond portfolio from CRSP.

[30] https://fred.stlouisfed.org/series/CPILFESL

[31] https://fred.stlouisfed.org/series/GS10

From Gilchrist and Zakrajšek [2012][32].

5. Unemployment rate

   From FRED[33].

A.3.1 Additional Instruments

1. CAPE

   From Campbell and Shiller [1988][34]

2. Shadow spread

   Following Lenel et al. [2019], we evaluate equation (9) in Gurkaynak et al. [2007] at maturity 1/12 for their estimated parameter values.

   $$f_t\left(1/12,0\right)=\beta_0+\beta_1\exp\left(-1/12/\tau_1\right)+\beta_2\left(1/12/\tau_1\right)\exp\left(-1/12/\tau_1\right)+\beta_3\left(1/12/\tau_2\right)\exp\left(-1/12/\tau_2\right)$$

   We then use the estimated 1-month forward rate to proxy for the one month yield. The shadow spread is then the difference between this estimate and the one-month yield from data.

3. Corporate bond spread

   Difference between the yield of Moody's seasoned BAA corporate bonds [35] and Moody's seasoned AAA corporate bonds[36]

B Ridge Estimation

In this appendix section, we describe the details of our regularized estimation ("ridge") procedure used to construct the zero-beta rate and expected consumption growth presented in Figure 3.

We will treat the projection moments as restrictions in our GMM estimation, which is to say that we require that they hold exactly. Loosely, this can be thought of as putting infinite weight on these moments, relative to other moments. The benefits of this approach are two-fold. First, it allows us to compute the regression coefficients (α, β) analytically,

[32] https://www.federalreserve.gov/econres/notes/feds-notes/updating-the-recession-risk-and-the-excess-bond-premium-20161006.html

[33] https://fred.stlouisfed.org/series/UNRATE

[34] http://www.econ.yale.edu/~shiller/data.htm

[35] https://fred.stlouisfed.org/series/BAA

[36] https://fred.stlouisfed.org/series/AAA

which greatly reduces the time required to compute the estimator. Second, it simplifies the interpretation of our procedure (as described in Section 3.3).

We use a ridge regression approach to avoid over-fitting with regards to the γ parameters. Specifically, we penalize the square norm of the γ vector, according to a weight $\psi \geq 0$. We choose ψ using cross-validation, in a manner described below. We exclude the γ_0 parameter (which reflects the average difference between the zero beta rate and the safe rate) from the penalty term. This form of penalization has the effect of shrinking our estimate of $R_{0,t}$ towards $R_{b,t}$ plus a constant, biasing us against finding in-sample time variation in the spread between $R_{0,t}$ and $R_{b,t}$. Introducing this kind of bias is useful in that it can reduce the variance of our out-of-sample forecast error, and can be interpreted as imposing a Bayesian prior.[37]

Our GMM analysis thus solves

$$\hat{\theta}_1 \in \arg\min_{\theta} \underbrace{\sum_{l=0}^{L} \left(T^{-1} \sum_{t=1}^{T} w(\theta)' \cdot (\alpha(\theta) + \hat{\epsilon}_t(\theta)) \, Z_{l,t-1} \right)^2}_{\text{Instrumented Asset Pricing Moment Squared}} + \underbrace{\psi \sum_{l=1}^{L} \gamma_l(\theta)^2}_{\text{Ridge Penalty}}$$

$$\text{subject to } \underbrace{\left(T^{-1} \sum_{t=1}^{T} \hat{\epsilon}_{i,t}(\theta) F_{j,t} \right)}_{\text{Projection Errors}} = 0, \forall i, j.$$

The problem is exactly identified (the number of asset pricing moments is equal to the number of predictor variables plus one, and the number of projection moments is equal to the number of (α_i, β_{ij}) parameters). As a consequence of the restrictions, conditional on γ, the (α_i, β_{ij}) point estimates will be the usual OLS estimates, as in the main text.[38]

We select the ridge penalty ψ via cross-validation. Given a candidate ψ, we divide our data sample into ten equal-length, non-overlapping subsets, $\{T_1, \ldots, T_{10}\}$, and estimate our model leaving out one particular T_m, producing a parameter estimate $\hat{\theta}_m(\psi)$. In the left-out subset, we compute the squared moment, $\left(w \left(\hat{\theta}_m(\psi) \right)' \cdot \hat{\epsilon}_{t+1} \left(\hat{\theta}_m(\psi) \right) \right)^2$, which is the out-of-sample variance of the surprise return of zero-beta portfolio. We repeated this process for

[37] See Hastie et al. [2009].

[38] The restriction approach is necessary for this result due to the ridge penalty. Absent the ridge penalty (as in the main text), the asset pricing moments would be zero at the estimated $\hat{\theta}$ and the estimates of $\hat{\beta}$ would coincide with OLS estimates even if these moments received finite weight. With the ridge penalty and finite weight on the projection moments, the asset pricing moments will be non-zero at $\hat{\theta}$ and the estimator will distort $\hat{\beta}$ to reduce asset pricing errors at the expense of larger projection errors. There is nothing incorrect about such an approach, but it complicates the computation and interpretation of the estimator.

each m, computing the sum of squared moments, and choose ψ to minimize this value:

$$\hat{\psi} \in \arg\min_{\psi \geq 0} \sum_{m=1}^{10} \left(w\left(\hat{\theta}_m(\psi)\right)' \cdot \hat{\epsilon}_{t+1}\left(\hat{\theta}_m(\psi)\right)\right)^2 .$$

Once the value of $\hat{\psi}$ is chosen, we compute $\hat{\theta}$ using this value and the full sample.[39]

C Placebo Test

Figure 7 below presents the results of a regression that the predicts the real return of a Treasury bond portfolio (the Fama 6-10y portfolio) using our instruments. The purpose of this placebo test is to demonstrate that there is nothing mechanical about results: our Z_t variables do not perfectly co-move, and there is no guarantee that the same combination of them that predicts bond returns will predict consumption growth. Moreover, our theory predicts that expected real bond returns should (generically) not line up with expected consumption growth, for two reasons. First, longer maturity bonds may inherit some of the convenience of shorter maturity bonds, because they can also be used to back short-dated safe claims (such as deposits or repo). Second, it is well-known (Campbell and Shiller [1991]) that the excess return of bonds over bills is predicted by the term spread (one of our instruments), which is to say that there is a time-varying risk premium for longer-maturity bonds. For both these reasons, we should not expect to two series to be aligned.

[39] This tenfold cross-validation procedure follows the recommendation of Hastie et al. [2009].

Figure 7: Expected Consumption Growth vs. Expected Real Tsy. Bond (6-10y) Returns

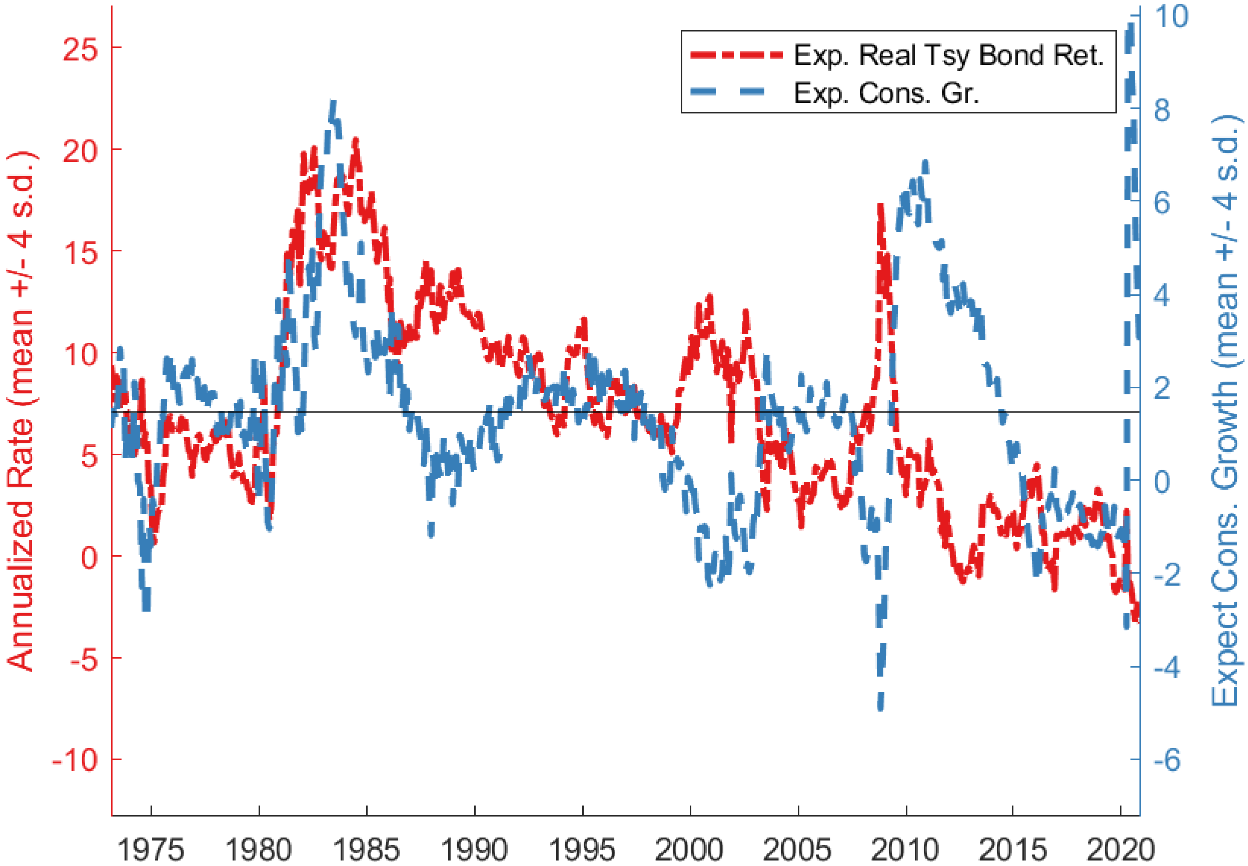

Notes: This figure plots the expected real return of a 6-10y Treasury bond portfolio against expected consumption growth, over time. Expected nominal returns are generated from predictive regressions using the instruments described in Section 4, which are the same instruments used to construct the zero-beta rate, and then converted to real returns by subtracting expected inflation (predicted with those same instruments). The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the expected real bond return. All series are annualized.

D Decomposition of the Effects of a Monetary Shock

In this appendix section, we present in table form some of the point estimates shown in Figure 5. Specifically, in the first four columns of Tables 2 and 3 we present point estimates for six-month changes ($h = 5$), for each of the three variables shown in Figure 5 (the real zero-beta rate, the real expected Treasury bill return, and consumption growth), as well as the spread (Conv) between the zero-beta rate and Treasury bill yield. Columns 5-9 in these tables show the coefficients (ϕ_0^l, ϕ_1^l) of the regression

$$\hat{\gamma}_l \cdot (Z_{l,t+5} - Z_{l,t-1}) = \phi_0^l + \phi_1^l \cdot mpshock_t + \epsilon_{t+5}^l,$$

for each of our L instruments, where $\hat{\gamma}_l$ is the point estimate from our GMM analysis with ridge penalization.

The sum of the coefficients ϕ_1^l, for $l \in \{1, \ldots, L\}$, is the coefficient on the spread regression in column 4, by (14). The tables thus illustrate the key drivers of the result that a monetary shock can simultaneously increase the safe rate while decreasing the zero-beta rate.

Both shocks increase the safe rate (the Treasury bill yield), and for this reason would be expected to increase the zero-beta rate if all else were equal. However, all else is not equal. The Romer and Romer [2004] shock results in a significant flattening of the yield curve, which more than offsets the effect of the increase in short rates (see column 8 of Table 2). The Nakamura and Steinsson [2018] shock has this effect, and also involves a significant increase in the excess bond premium (see columns 7 and 8 of Table 3). The two shocks differ in both their construction and in the periods in which they are available (and the conduct of monetary policy has changed over time), either of which might explain the observed differences between their effects. Note also that neither of these tables includes standard errors.

Table 2: Decomposition of the Effects of a Romer-Romer Shock

	(1) Real Bill	(2) Real Z.B.	(3) Ex. C. Gr.	(4) Conv	(5) g_RF	(6) g_UMP	(7) g_EBP	(8) g_TSP	(9) g_CPI_Rolling
RR_shock	0.338	-1.467	-0.312	-1.805	0.168	-0.227	-0.434	-2.457	0.0574
Constant	-0.0273	-0.0225	-0.0106	0.00474	-0.00283	0.000943	-0.00963	0.0222	0.00229
Observations	413	413	413	413	417	417	417	417	417

Table 3: Decomposition of the Effects of a Nakamura-Steinsson Shock

	(1) Real Bill	(2) Real Z.B.	(3) Ex. C. Gr.	(4) Conv	(5) g_RF	(6) g_UMP	(7) g_EBP	(8) g_TSP	(9) g_CPI_Rolling
NS_shock	1.801	-9.717	-1.945	-11.52	0.178	-1.687	-8.849	-2.746	0.0920
Constant	-0.0856	-0.0550	0.00417	0.0305	-0.00354	-0.00532	0.0202	-0.0235	-0.00171
Observations	230	230	230	230	234	234	234	234	234

E Robustness Exercises

This appendix section contains our robustness exercises. We first provide an index of the various exercises, with a label used to identify the specification.

- NoDrop20: With the bottom two deciles of stocks included in the data sample.
- FF5Industry: With Fama-French 5-factor sorted + industry portfolios and instead of 3-factor + industry portfolios.
- Sigma1: With a $\sigma = 1$ instead of $\sigma = 5$ used to construct the consumption factor
- LinearCons: With a linear consumption factor $F_{8,t} = \Delta c_{t+1}$ instead of a non-linear consumption factor in the SDF.
- NoCons: With no consumption factor in the SDF.
- MktOnly: With only the market and non-linear consumption factors.
- FF3Only: With only the market, size, and value factors of Fama and French [1993], plus the non-linear consumption factor.
- AltBAAS: With our preferred instruments, using the BAA-AAA spread in the place of the excess bond premium.
  - The EBP and BAA-AAA spread are highly correlated conditional on our other instruments, and for this reason we don't include them both.
- AltCAPE: With our preferred instruments plus the cyclically adjusted price-earnings (CAPE) ratio.
- LagCons: With our preferred instruments and a lag of consumption growth.
- Shadow: With our preferred instruments and the "shadow spread" used by Lenel et al. [2019].
- VaryingBetas: With instruments-by-factor interactions as factors (allowing for time-varying betas).
  - i.e. with the seven factors of our main specification (excluding the consumption factor), plus 35 factors $\tilde{F}_{j'',t+1} = F_{j,t+1}Z_{l,t}$ for $j \in \{1, \ldots 7\}$ and $l \in \{1, \ldots 5\}$, plus the consumption factor (43 factors total). This specification is isomorphic to a model in which the betas to the seven main specification factor are linear in the Z_t variables

- NDOnly: With non-durable goods consumption per capita as opposed to non-durable goods + services per capita.
- NoCOVID: With a data sample ending in December 2019 (pre-COVID 19).

Table 4: Predicting the Zero-Beta Rate by Specification

	(1) Main	(2) NoDrop20	(3) FF5Industry	(4) Sigma1	(5) LinearCons	(6) NoCons	(7) MktOnly	(8) FF3Only	(9) AltBAAS	(10) AltCAPE	(11) LagCons	(12) Shadow	(13) VaryingBetas	(14) NDOnly	(15) NoCOVID
RF	1.186	1.543	1.634	1.194	1.215	1.174	1.767	1.138	1.280	1.086	1.104	1.349	3.941	1.248	1.280
	(0.914)	(0.877)	(0.862)	(0.917)	(0.919)	(0.871)	(0.745)	(0.779)	(0.915)	(0.955)	(0.913)	(0.877)	(2.181)	(0.890)	(0.871)
UMP	0.105	0.0758	0.0859	0.108	0.115	0.102	0.0916	0.0940	0.102	0.0838	0.125	0.138	0.230	0.109	0.110
	(0.0986)	(0.0926)	(0.105)	(0.0983)	(0.101)	(0.0979)	(0.0934)	(0.0907)	(0.109)	(0.115)	(0.0977)	(0.104)	(0.324)	(0.0976)	(0.114)
EBP	-0.603	-0.612	-0.885	-0.620	-0.638	-0.593	-0.413	-0.372		-0.605	-0.584	-0.645	-0.740	-0.635	-0.720
	(0.342)	(0.334)	(0.302)	(0.354)	(0.350)	(0.326)	(0.289)	(0.310)		(0.342)	(0.340)	(0.348)	(0.842)	(0.344)	(0.333)
TSP	0.310	0.333	0.277	0.309	0.310	0.310	0.410	0.372	0.298	0.295	0.285	0.251	0.395	0.313	0.331
	(0.118)	(0.114)	(0.121)	(0.118)	(0.117)	(0.118)	(0.110)	(0.110)	(0.120)	(0.131)	(0.118)	(0.127)	(0.302)	(0.117)	(0.128)
CPI_Rolling	-2.582	-2.768	-2.461	-2.565	-2.643	-2.564	-2.644	-2.008	-2.894	-2.785	-2.551	-3.466	-8.045	-2.683	-2.374
	(1.175)	(1.147)	(1.160)	(1.162)	(1.203)	(1.144)	(0.965)	(1.021)	(1.217)	(1.396)	(1.167)	(1.239)	(3.353)	(1.160)	(1.158)
BAAS									0.0546						
									(0.502)						
CAPE										-0.0111					
										(0.0364)					
LagCons											0.252				
											(0.164)				
ShadowSpread												7.713			
												(5.958)			
Constant	0.718	0.720	0.678	0.718	0.717	0.718	0.761	0.748	0.713	0.715	0.717	0.747	0.328	0.717	0.766
	(0.137)	(0.132)	(0.134)	(0.137)	(0.137)	(0.137)	(0.120)	(0.122)	(0.137)	(0.138)	(0.137)	(0.136)	(0.363)	(0.138)	(0.142)

Standard errors in parentheses

Notes: This table presents the γ coefficients from our GMM estimation across the various specifications defined in Appendix Section E. The standard errors are robust to heteroskedasticity and account for estimation error in the other parameters. The instruments have been centered, so the constant coefficient is the average monthly return of the zero-beta portfolio. The "Main" specification of the first column is the one used in the main text (Table 1).

Figure 8: NoDrop20: Results with Bottom Decile Stocks Included

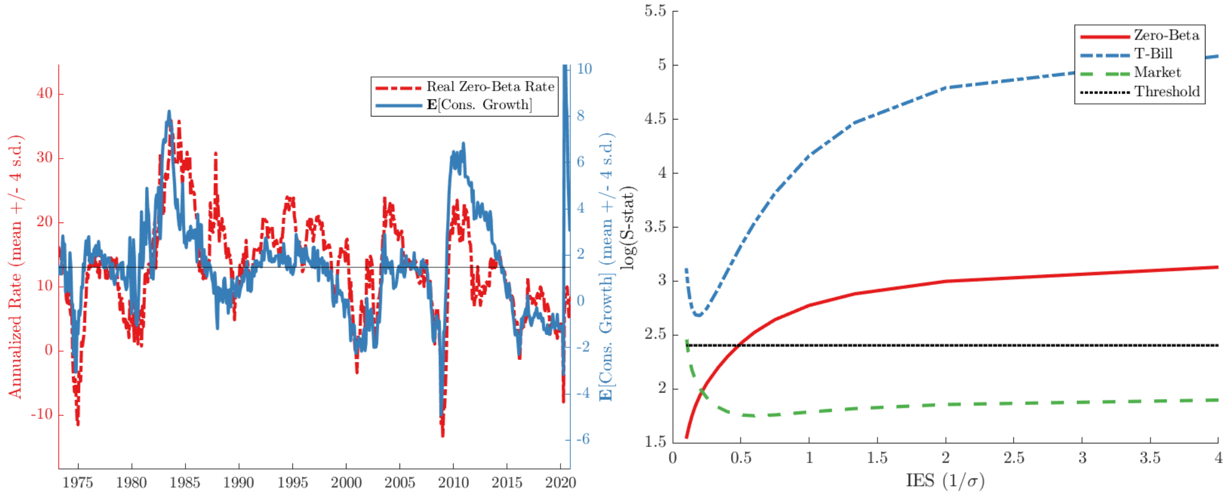

Notes: Both panels present results for the "NoDrop20" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 9: FF5Industry: Results with FF5+Industry Portfolios

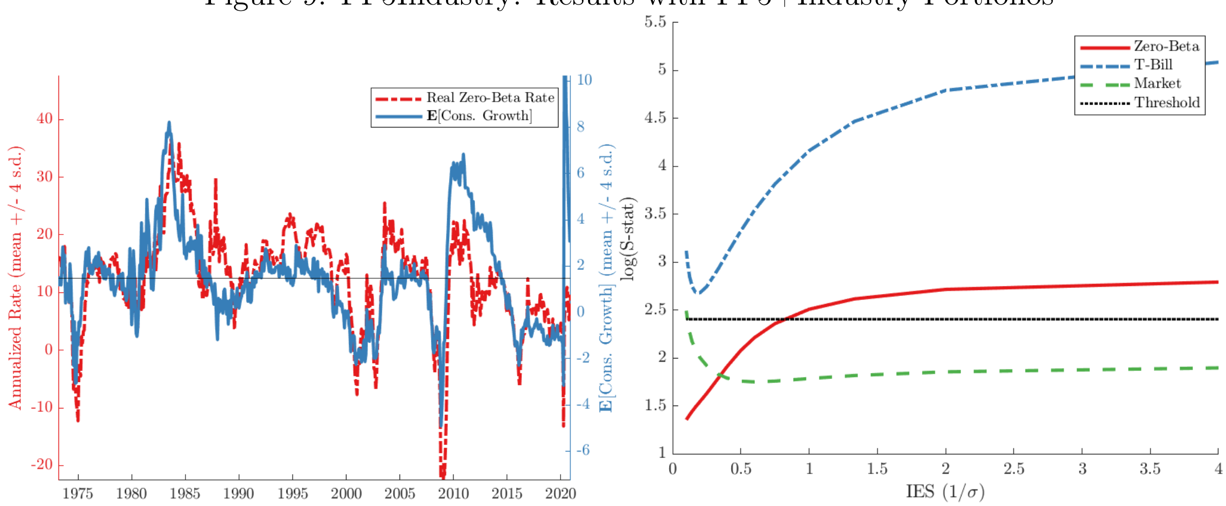

Notes: Both panels present results for the "FF5Industry" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 10: Sigma1: Results with $\sigma = 1$

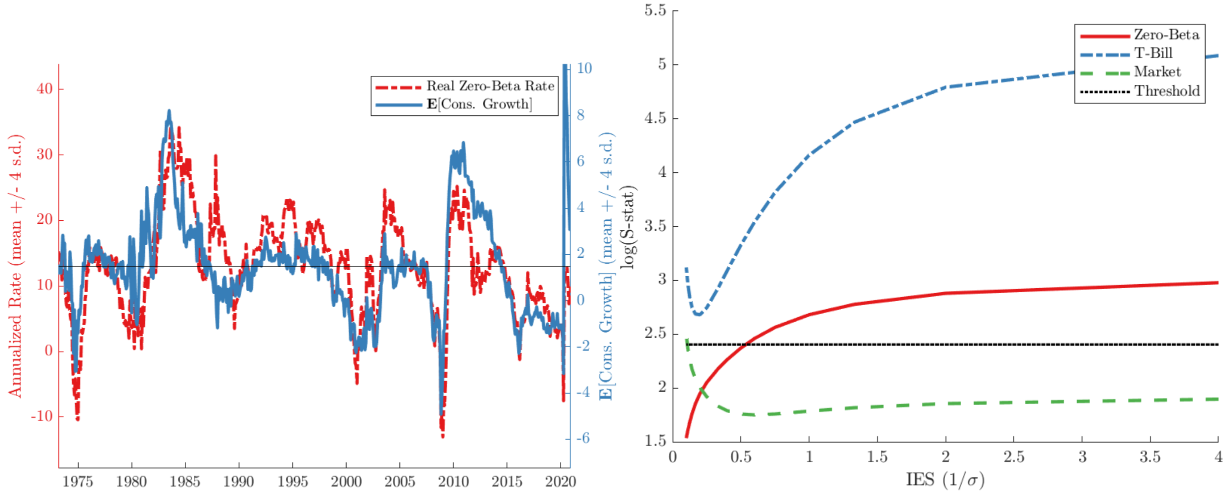

Notes: Both panels present results for the "Sigma1" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 11: LinearCons: Results with a Linear Consumption Factor

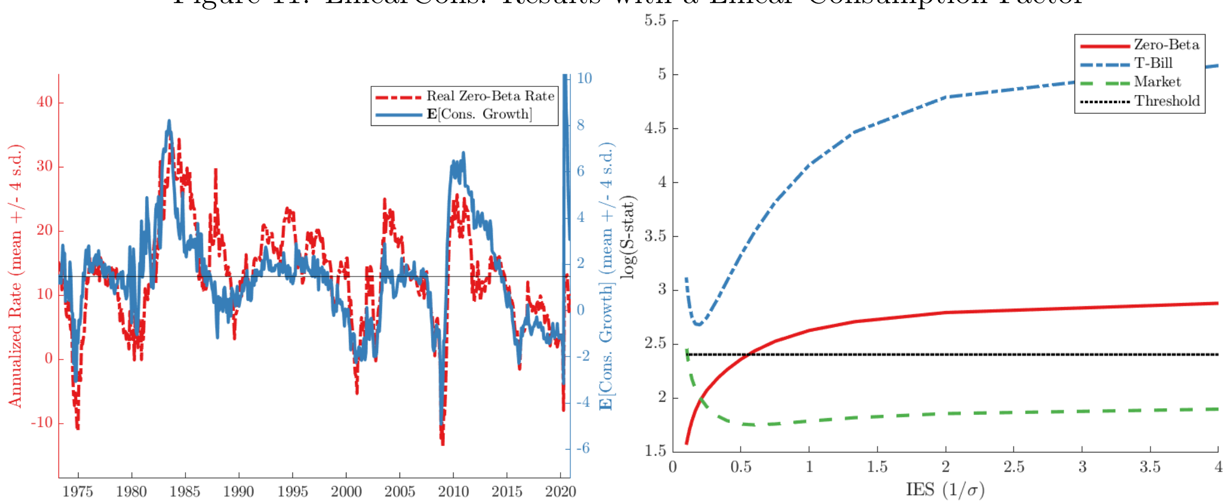

Notes: Both panels present results for the "LinearCons" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 13: MktOnly: Results with only the Market and Consumption Factors

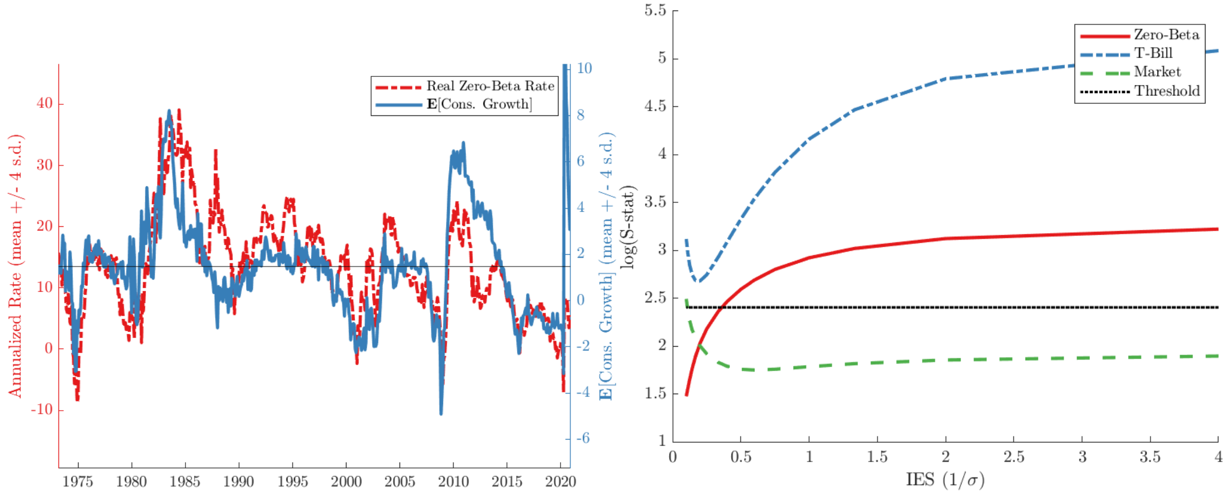

Notes: Both panels present results for the “MktOnly” robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 12: NoCons: Results with No Consumption Factor

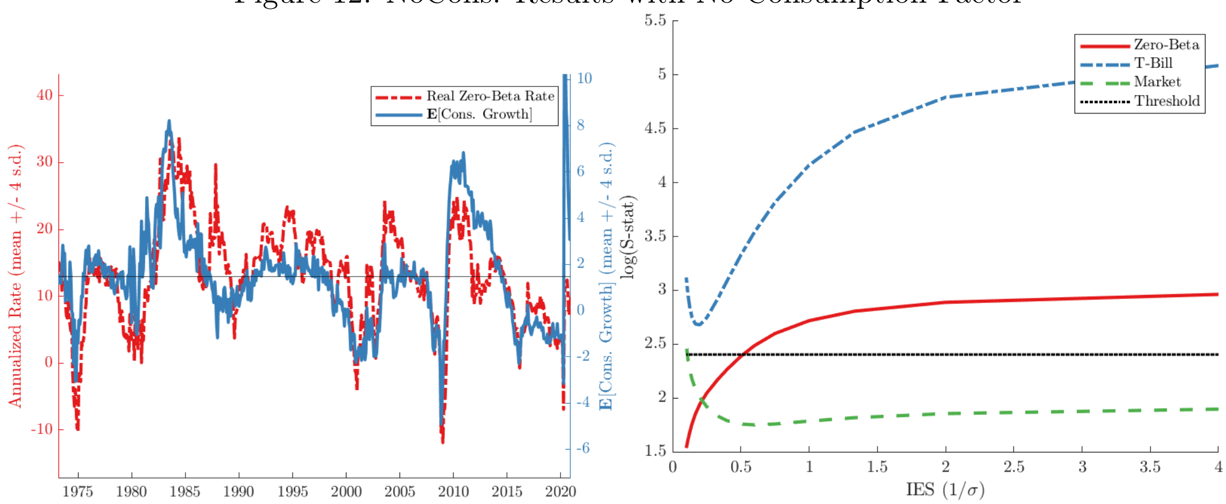

Notes: Both panels present results for the “NoCons” robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 14: FF3Only: Results with only the FF3 and Consumption Factors

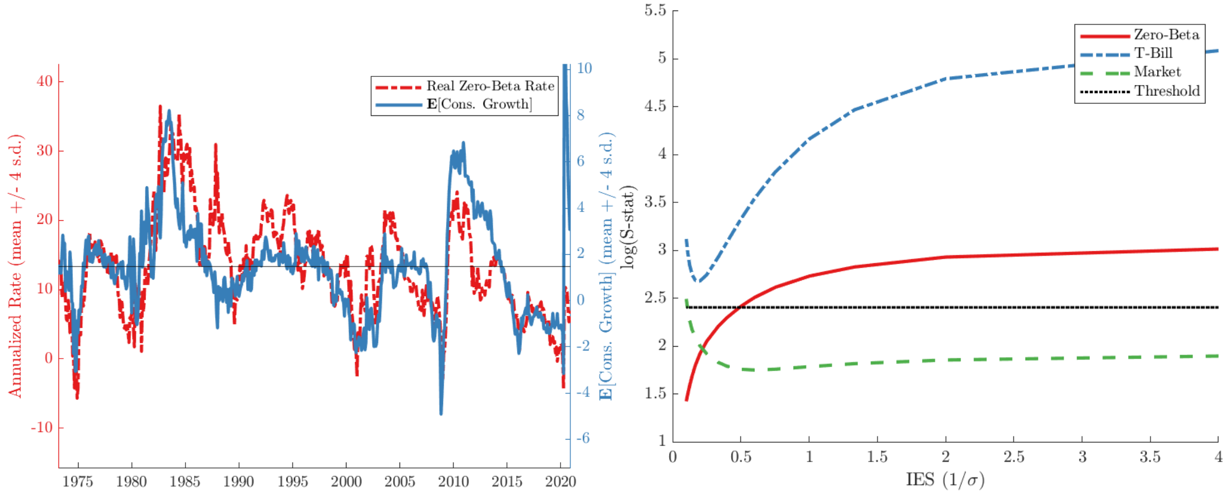

Notes: Both panels present results for the "FF3Only" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 15: AltBAAS: Results with the BAA-AAA Spread instead of the EBP

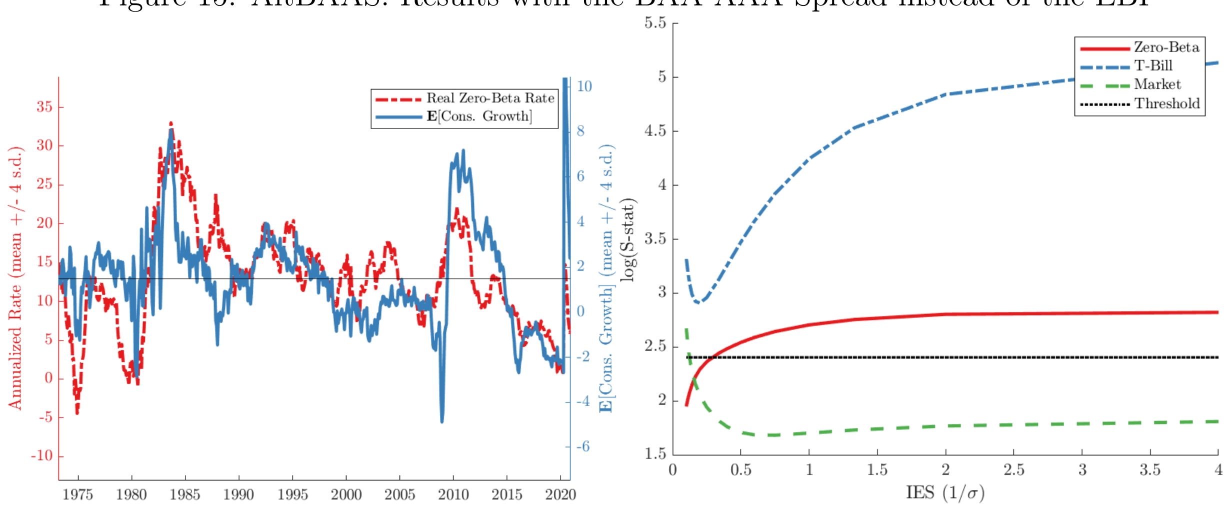

Notes: Both panels present results for the "AltBAAS" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 16: AltCAPE: Results with the CAPE Instrument Included

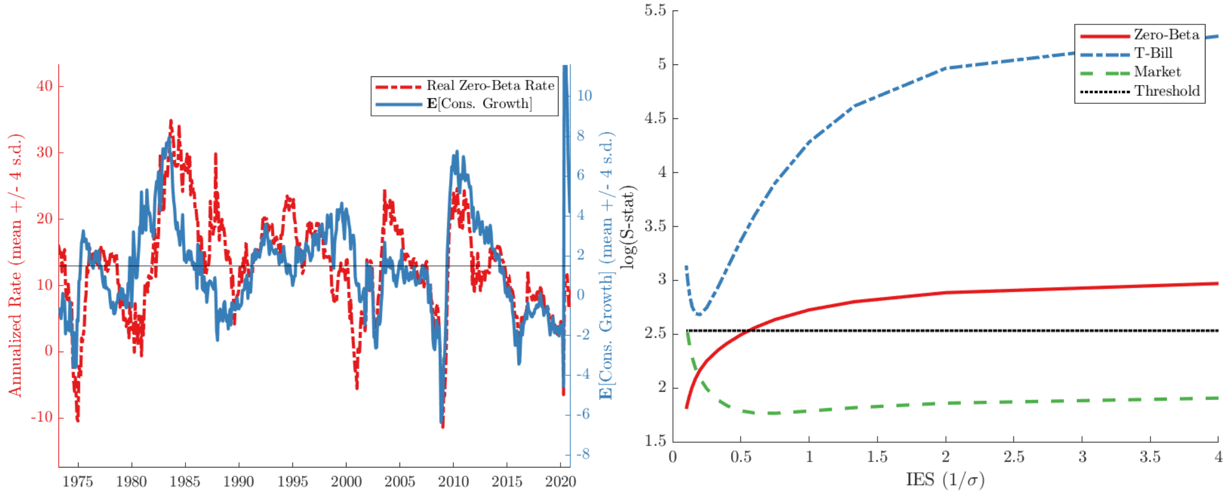

Notes: Both panels present results for the "AltCAPE" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (6).

Figure 17: LagCons: Results with the Lagged Consumption Growth Instrument Included

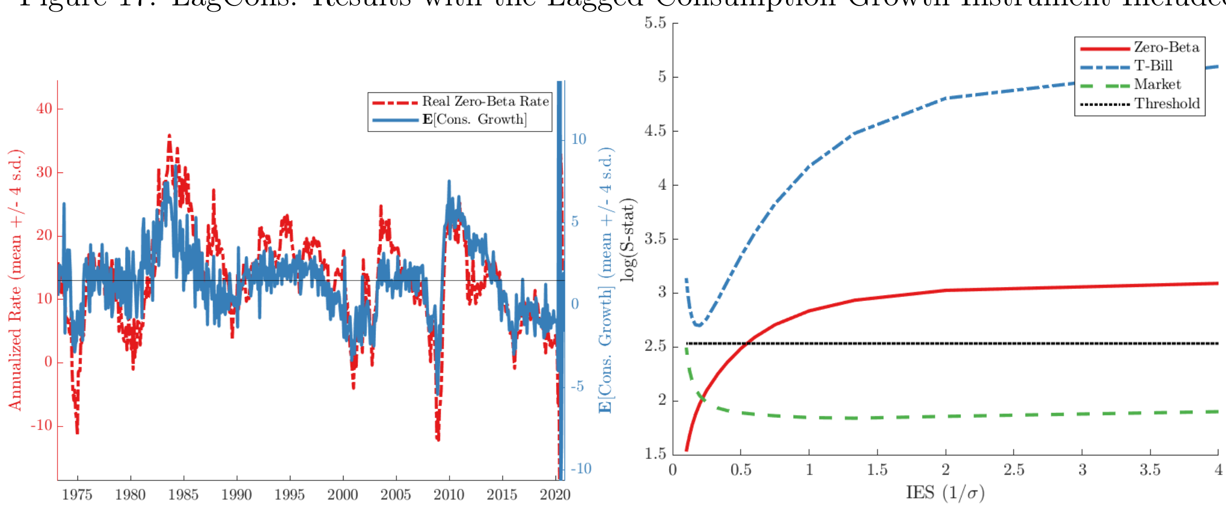

Notes: Both panels present results for the "LagCons" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (6).

Figure 18: Shadow: Results with the Shadow Spread Instrument Included

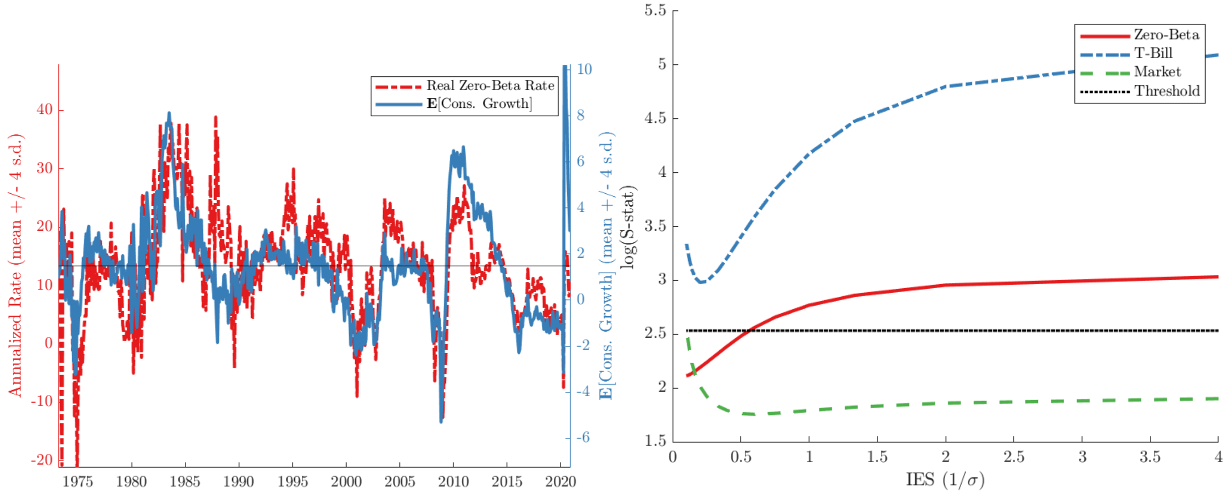

Notes: Both panels present results for the "Shadow" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (6).

Figure 19: VaryingBetas: Results with Linear Betas

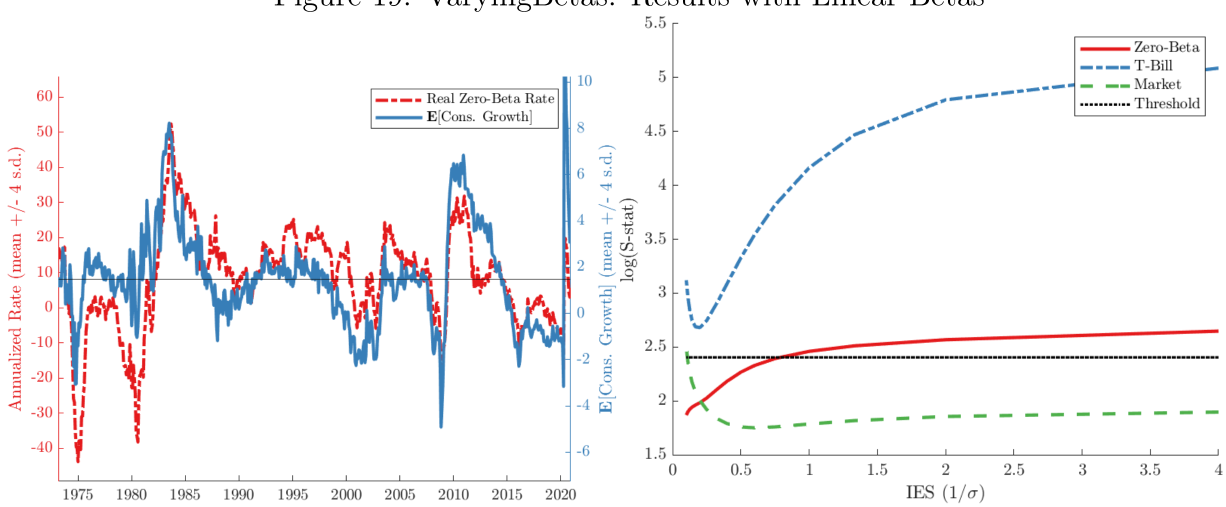

Notes: Both panels present results for the "VaryingBetas" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 20: NDOnly: Results with Non-Durable Goods Consumption

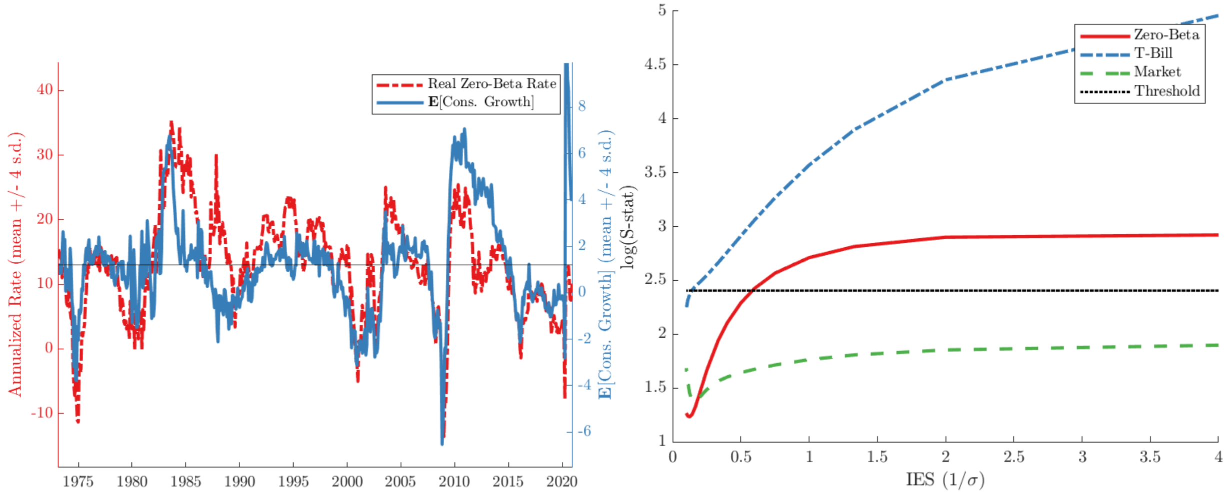

Notes: Both panels present results for the "NDOnly" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

Figure 21: NoCOVID: Results Excluding 2020

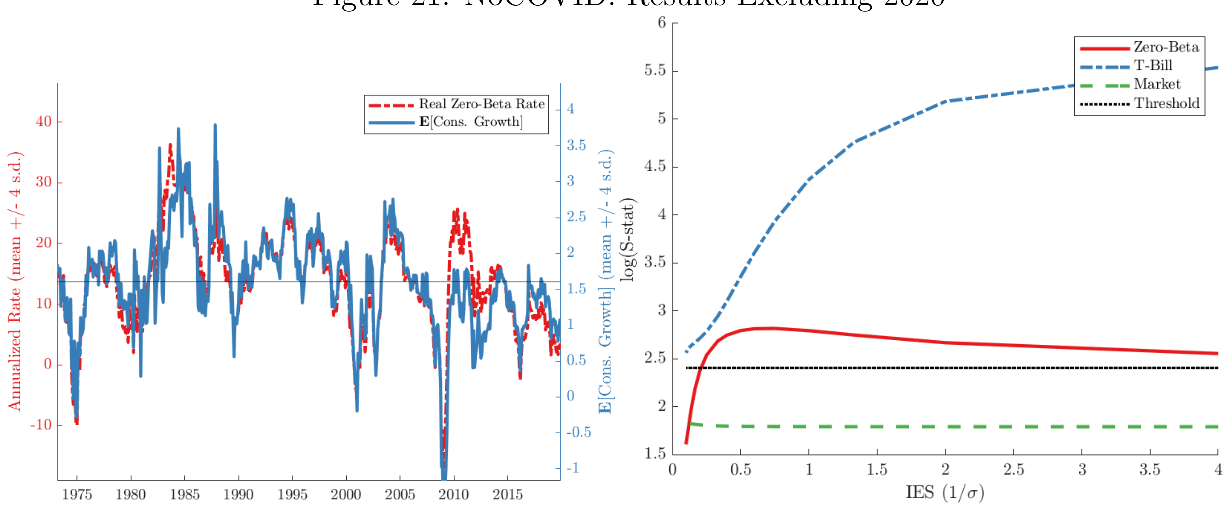

Notes: Both panels present results for the "NoCOVID" robustness exercise (Appendix Section E). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation and expected consumption growth are generated from predictive regressions using the same instruments used to construct the zero-beta rate. The right vertical axis is consumption growth, centered at its mean, with limits equal to +/- four standard deviations. The left vertical axis is the same for the real zero beta rate, and both series are annualized. The right panel plots the log of the $\hat{S}(\sigma)$ statistic for values of σ from 0.25 to 10, which is constructed using the Euler equation moments (16) as applied to the zero-beta rate, Treasury bill yield, and CRSP market return. The dotted threshold line is the log of the 95th percentile critical value of a chi-square distribution with degrees of freedom equal to the number of instruments (5).

F Details on the Covariance Matrix Estimator

The Ledoit and Wolf [2017] covariance matrix estimator, as we apply it, can be thought of as a function of the estimated β parameters, the sample covariance matrix of factor returns $\hat{\Sigma}_K(\gamma)$, and the sample covariance matrix of excess returns,

$$\hat{\Sigma}(\gamma) = T^{-1} \sum_{t=1}^{T} (R_{t+1} - R_{0,t}(\gamma)) (R_{t+1} - R_{0,t}(\gamma))' .$$

Note that $\hat{\Sigma}_K(\gamma)$ depends on γ via the dependence of first factor (the excess return of the market) on $R_{0,t}$.

We begin by pre-conditioning (as in section 4.2 of Ledoit and Wolf [2017]) using our factor model. Define

$$\hat{\Sigma}_F(\theta) = \beta \hat{\Sigma}_K(\gamma) \beta' + diag\left(\hat{\Sigma}(\gamma) - \beta \hat{\Sigma}_K(\gamma) \beta'\right),$$

where $diag(M)$ is a diagonal matrix whose diagonal is equal to that of M. The covariance matrix $\hat{\Sigma}_F$ can be thought of as the covariance matrix implied by an exact factor model, with our chosen factors.

We next transform the excess return data, to generate

$$Y_t(\theta) = \left(\hat{\Sigma}_F(\theta))^{-\frac{1}{2}} (R_{t+1} - R_{0,t}(\gamma)\right),$$

where $(\cdot)^{-\frac{1}{2}}$ is the symmetric matrix square root. We then apply the Ledoit and Wolf [2017] shrinkage estimator to estimate the covariance matrix of $Y_t(\theta)$ (call this $\hat{\Sigma}_c(\theta)$), and finally generate our estimate of the variance-covariance matrix of returns using

$$\hat{\Sigma} = \left(\hat{\Sigma}_F(\theta)\right)^{\frac{1}{2}} \hat{\Sigma}_c(\theta) \left(\hat{\Sigma}_F(\theta)\right)^{\frac{1}{2}} .$$

This pre-conditioning in effect imposes a uniform prior about the orientation of the eigenvectors of $\hat{\Sigma}_c(\theta)$, as opposed the about the orientation of the same for $\hat{\Sigma}$. The former is more appropriate in light of the co-movement of stocks with, for example, the market factor. Our procedure differs from the empirical exercise in Ledoit and Wolf [2017] in that it uses our K-factor model for pre-conditioning instead of a single factor model, which seems more appropriate for our application.

We also modify their procedure in one additional respect, by using the analytical non-linear shrinkage estimator of Ledoit and Wolf [2020] in the place of the "non-linear" shrinkage estimator. The two methods offer similar out-of-sample performance in the minimum-variance portfolio problem, and the analytical method is substantially faster to compute.

G Relationship to Shanken [1986]

Our starting point when developing our procedure was the MLE approach of Shanken [1986], described in Campbell et al. [1998]. The Shanken [1986] procedure is designed to extract a constant (over the sample period) risk-free rate. The key way in which we have modified the procedure is via our assumption on the structure of the zero-beta rate, equation (14), which replaces the assumption of a constant zero-beta rate. Aside from this difference, our procedure deviates from the MLE estimator by using a regularized covariance matrix estimator.

To begin, let us suppose that the residuals in the projection regressions (15), $\hat{\epsilon}_t$, are Gaussian and i.i.d. with variance–covariance matrix Σ_ϵ, and that all of the factors are tradable. Shanken [1986] derives an MLE estimator for a constant zero-beta rate under these assumptions.[40] Using Σ_ϵ has one particular disadvantage: Σ_ϵ may not be full rank (for example, if the value-weighted sum of the test assets is the market portfolio). In contrast, our procedure handles this case without modification. In the discussion that follows, assume Σ_ϵ is invertible.

Under these assumptions, the log-likelihood function is, ignoring constants,

$$f(R_{t+1}, F_{t+1}, Z_t; \theta, \Sigma_\epsilon) = \frac{1}{2} \ln \left(\det(\Sigma_\epsilon^{-1}) \right) - \frac{1}{2} \hat{\epsilon}_{t+1} (\theta)' \cdot \Sigma_\epsilon^{-1} \cdot \hat{\epsilon}_{t+1} (\theta) \, ,$$

where $\hat{\epsilon}_{t+1}(\theta)$ is defined from (R_{t+1}, F_{t+1}, Z_t) as in (15).

Maximizing the log-likelihood over θ, it follows, given the MLE estimate of γ, that the maximum likelihood (α, β) estimates are exactly the OLS coefficients of the projection regression. Specifically, they solve

$$\mathbb{E} \left[F_{j,t} e_i' \cdot \Sigma_\epsilon^{-1} \cdot \hat{\epsilon}_{t+1}(\theta) \right] = 0$$

for $i \in \{1, \dots, N\}$ and $j \in \{0, \dots, K\}$, where $e_i \in \mathbb{R}^N$ denotes the basis vector that selects the i-th asset. Because this must hold for all i and Σ is full rank, it is equivalent to $E[F_{j,t} \cdot \hat{\epsilon}_{i,t+1}(\theta)] = 0$ for all (i, j), which are the moment conditions associated with the time series regressions (11).

[40] Shanken [1986] in fact assumes a single factor model, but the extension to multi-factor models with tradable factors is straightforward.

The first order condition with respect to γ yields

$$\mathbb{E}\left[\underbrace{(\iota-\beta_{\cdot,1})'\cdot\Sigma_{\epsilon}^{-1}}_{w_{MLE}}\cdot\hat{\epsilon}_{t+1}(\theta)Z_t\right]=0.$$

We will next show that under the stated assumptions, the "portfolio weight" $w_{MLE} = (\iota-\beta_{\cdot,1})'\cdot\Sigma_{\epsilon}^{-1}$ is equal to our $w(\theta)$.

Consider our procedure, applied to an augmented set of test assets that includes the factors themselves (which are now by assumption tradable). Specifically, let $R_{N+j,t+1} = F_{j,t+1}$ for $j = \{1,\ldots,K\}$. Note that, because the non-market factors are assumed to be zero-investment, our procedure would have to be modified by re-defining the $\iota \in \mathbb{R}^{N+K}$ to be equal to one for its first $N+1$ elements and zero otherwise.

The covariance matrix Σ_R for the augmented set of test assets can be written in block form as

$$\Sigma_R = \begin{bmatrix} \beta\Sigma_K\beta' + \Sigma_{\epsilon} & \beta\Sigma_K \\ \Sigma_K\beta' & \Sigma_K \end{bmatrix},$$

where Σ_K is the covariance matrix of the factors.

The minimum-variance zero-beta unit-investment portfolio problem that defines $w(\theta) \in \mathbb{R}^{N+K}$ in this case is equivalent to solving

$$\min_{\tilde{w}\in\mathbb{R}^N} \tilde{w}'\cdot\begin{bmatrix} I & -\beta \end{bmatrix}\cdot\Sigma_R\cdot\begin{bmatrix} I \\ -\beta' \end{bmatrix}\cdot\tilde{w}$$

subject to $\tilde{w}'\cdot\begin{bmatrix} I & -\beta \end{bmatrix}\cdot\iota = 1$. Here, we have defined $w = \begin{bmatrix} I \\ -\beta' \end{bmatrix}\cdot\tilde{w}$ and in effect constructed zero-beta portfolios by hedging out the tradable factors. Straightforward algebra shows that this problem simplifies to a minimum variance portfolio problem, whose solution is $\tilde{w}^* = w_{MLE} = (\iota-\beta_{\cdot,1})'\cdot\Sigma_{\epsilon}^{-1}$. Thus, under the stated assumptions, our portfolio weights are equivalent to the ones implied by the MLE procedure of Shanken [1986] conditional on the estimate of Σ_{ϵ} and Σ_F.

More generally, whenever all factors are tradable and each of those factors lies in the span of the test assets (e.g. using the Fama-French 25 portfolios and three factors), our procedure's $w(\theta)$ and w_{MLE} will coincide (again, conditional on the covariance matrices). Our procedure has the advantages of handling the case of non-tradable factors and of avoiding the assumption that Σ_{ϵ} is of full rank, but is otherwise similar.

The more significant difference between our procedure and the MLE estimator arise from our use of the Ledoit and Wolf [2017] covariance matrix estimator. The MLE estimator for

Σ_ϵ (which can be derived from the first-order conditions) is the sample covariance matrix of the residuals, and use of the Ledoit and Wolf [2017] estimator for Σ_R avoids over-fitting. Our GMM estimator is essentially the MLE estimator, modified to avoid overfitting.

H Details on the Application of Stock and Wright [2000]

In this appendix section, we describe in more detail the procedure we use when constructing the S-sets described in Section 5 of the main text.

At a high level, our follows the approach of Stock and Wright [2000]. The only meaningful modification we make to their approach is to use one set of moments for the purpose of estimating the strongly-identified parameters $(\delta, \theta = (\alpha, \beta, \gamma))$ (in particular, constructing the zero-beta rate) and then using a different set moments (the consumption Euler equation) for the purpose of constructing the test statistic. This approach (which is well-known, see chapter 11.6 of Cochrane [2009]) has the advantage of being easily interpretable. It also has the advantage, in our particular application, of allowing us to analytically compute the (α, β) parameters given any value of γ, which facilitates computation, and it ensures that the zero-beta rate described in Section 4 is the same as the zero-beta rate being tested in Section 5. Stock and Wright [2000] present results under the assumption that the same weight matrix used to estimate the well-identified parameters is also used to construct the test statistic; the purpose of the appendix section is to show that their results can be generalized away from this case. Those authors also assume (for convenience) a positive-definite weighting matrix; our procedure is most naturally cast as involving a positive semi-definite matrix.

Note that we do not prove the standard GMM identification assumptions (global identification, differentiability, etc...) in our setting, and instead assume that they apply. Necessary conditions include that the factors $F_{j,t}$ not be co-linear (as otherwise β cannot be identified) and that the instruments Z_t not be co-linear (as otherwise γ cannot be identified).

Recall the our moment conditions are

$$g_t(\theta, \delta, \sigma) = \begin{bmatrix} \hat{\epsilon}_{t+1}(\theta) \otimes F_{t+1}(\sigma, \gamma) \\ H(\beta) \cdot (R_{t+1} - R_{0,t}(\gamma)) \otimes Z_t \\ (\delta \frac{c_{t+1}^{-\sigma}}{c_t^{-\sigma}} \frac{R_{0,t}(\gamma)}{P_{t+1}/P_t} - 1) \otimes Z_t \end{bmatrix}$$

and that our weight matrix used in estimation is

$$W_T(\theta) = \begin{bmatrix} I & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & w_T(\theta)_T w(\theta)' & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & e_0 e_0' \end{bmatrix}.$$

We have written W_T as a function of the sample size because the zero-beta portfolio weight vector w_T involves an estimate of the variance-covariance matrix. Let Θ be the compact set of possible parameters for (θ, δ, σ), which excludes parameters for which ι lies in the span of β.

Define
$$m_T(\theta, \delta, \sigma) = \mathbb{E}\left[T^{-1} \sum_{t=1}^{T} g_t(\theta, \delta, \sigma)\right]$$
and
$$\Psi_T(\theta, \delta, \sigma) = T^{-\frac{1}{2}} \sum_{t=1}^{T} \left(g_t(\theta, \delta, \sigma) - \mathbb{E}\left[g_t(\theta, \delta, \sigma)\right]\right).$$

We will assume $\Psi_T(\theta, \delta, \sigma)$ converges to a Gaussian process $\Psi(\theta, \delta, \sigma)$ (Assumption B of Stock and Wright [2000]; those authors provide more primitive assumptions in which this holds). Let $\Omega(\theta, \delta, \sigma) = \mathbb{E}\left[\Psi(\theta, \delta, \sigma)\Psi(\theta, \delta, \sigma)'\right]$ be the limiting covariance matrix; we assume it can be consistently estimated using heteroskedasticity-robust methods (Assumption D'' of Stock and Wright [2000]).[41] We will also assume that $W_T(\theta)$ converges uniformly in probability to a symmetric positive semi-definite matrix-valued function $W(\theta)$ that is continuous in θ (Assumption D of Stock and Wright [2000], weakened to required only positive semi-definiteness).

We will treat the parameter σ as weakly identified, and the parameters (θ, δ) as strongly identified. Suppose $(\theta_0, \delta_0, \sigma_0)$ are the true parameters. We decompose

$$m_T(\theta, \delta, \sigma) = m_{1T}(\theta, \delta, \sigma, \sigma_0) + m_2(\theta, \delta, \sigma_0),$$

where in our context,
$$m_2(\theta, \delta; \sigma_0) = \mathbb{E}\left[T^{-1} \sum_{t=1}^{T} g_t(\theta, \delta, \sigma_0)\right]$$
and
$$m_{1T}(\theta, \delta, \sigma; \sigma_0) = \mathbb{E}\left[T^{-1} \sum_{t=1}^{T} \begin{bmatrix} \hat{\epsilon}_{t+1}(\theta) \otimes \left(F_{t+1}(\sigma, \gamma) - F_{t+1}(\sigma_0, \gamma)\right) \\ 0 \\ \left(\delta\left(\frac{c_{t+1}^{-\sigma}}{c_t^{-\sigma}} - \frac{c_{t+1}^{-\sigma_0}}{c_t^{-\sigma_0}}\right)\frac{R_{0,t}(\gamma)}{P_{t+1}/P_t}\right) \otimes Z_t \end{bmatrix}\right].$$

Note that, following Stock and Wright [2000], we have assumed that $m_2(\cdot)$ does not depend

[41] Consistent with the equations of our model, we assume the residuals are serially uncorrelated (following Hansen and Singleton [1982] and chapter 11.7 of Cochrane [2009]).

on T, without imposing this assumption on m_{1T}. In our main specification,

$$F_{t+1}(\sigma,\gamma) - F_{t+1}(\sigma_0,\gamma) = e_8\left(\frac{c_{t+1}^{-\sigma}}{c_t^{-\sigma}} - \frac{c_{t+1}^{-\sigma_0}}{c_t^{-\sigma_0}}\right),$$

which is to say that only the non-linear consumption factor depends on σ. The potential for weak identification is readily apparent: if the difference of the two consumption SDFs (with σ and σ_0) is only weakly related to stock returns and our instruments, the parameter σ will be largely unidentified.

Our key assumption is that the moments $m_2(\cdot)$ satisfy the usual GMM identification conditions. We assume (following Assumption C of Stock and Wright [2000]) that:

1. The function $T^{\frac{1}{2}}(m_T(\theta,\delta,\sigma) - m_2(\theta,\delta;\sigma_0))$ converges uniformly to the function $m_1(\theta,\delta,\sigma)$, which is continuous and bounded on Θ and satisfies $m_1(\theta_0,\sigma_0,\delta_0)=0$.

2. The function $m_2(\theta,\delta;\sigma_0)$ satisfies $m_2(\theta_0,\delta_0;\sigma_0)=0$ and $W(\theta)m_2(\theta,\delta;\sigma_0)\neq 0$ for all $(\theta,\delta)\neq(\theta_0,\delta_0)$. The function $m_2(\theta,\delta;\sigma_0)$ is continuously differentiable with respect to (θ,δ) in the neighborhood of (θ_0,δ_0), with Jacobian $R(\theta,\delta;\sigma_0)$, and $W(\theta)R(\theta,\delta,\sigma_0)$ has full column rank.

The first part of this assumption is exactly part (i) of Assumption C of Stock and Wright [2000]. The second part is a modified version of part (ii) of that assumption: we impose the standard global and local GMM identification conditions on the weighted moments as opposed to the unweighted ones. This modification (which is standard in the GMM literature) allows us to consider positive semi-definite weighting matrices.

First note that, under these assumptions,

$$(\hat{\theta}(\sigma_0),\hat{\delta}(\sigma_0)) = \arg\min_{(\theta,\delta):(\theta,\delta,\sigma_0)\in\Theta}(T^{-1}\sum_{t=1}^{T} g_t(\theta,\delta,\sigma_0))'W_T(\theta)(T^{-1}\sum_{t=1}^{T} g_t(\theta,\delta,\sigma_0))$$

is a $\sqrt{T}$-consistent estimator for (θ_0,δ_0) given σ_0.[42] It follows that the usual GMM formula applies,

$$\sqrt{T}((\hat{\theta}(\sigma_0),\hat{\delta}(\sigma_0))-(\theta_0,\delta_0)) \Rightarrow -[R(\theta_0,\delta_0,\sigma_0)'W(\theta)R(\theta_0,\delta_0,\sigma_0)]^{-1}R(\theta_0,\delta_0,\sigma_0)'W(\theta)\Psi(\theta_0,\delta_0,\sigma_0). \tag{18}$$

Likewise, the usual formula for the moments applies (via the delta method):

$$\sqrt{T}m_2(\hat{\theta}(\sigma_0),\hat{\delta}(\sigma_0),\sigma_0) \Rightarrow \tilde{R}(\theta_0,\delta_0,\sigma_0)\Psi(\theta_0,\delta_0,\sigma_0),$$

[42] This can be proven along the lines of Lemma A1 in Stock and Wright [2000]; the proof must be adapted in relatively straightforward way to the positive semi-definite case.

where

$$\tilde{R}(\theta_0, \delta_0, \sigma_0) = (I - R(\theta_0, \delta_0, \sigma_0)[R(\theta_0, \delta_0, \sigma_0)'W(\theta_0)R(\theta_0, \delta_0, \sigma_0)]^{-1}R(\theta_0, \delta_0, \sigma_0)'W(\theta_0)).$$

Using this formula, we can define the variance-covariance matrix

$$V_{Test}(\sigma_0) = T^{-1}W_{test}\tilde{R}(\theta_0, \delta_0, \sigma_0)\Omega(\theta_0, \delta_0, \sigma_0)\tilde{R}(\theta_0, \delta_0, \sigma_0)'W_{test},$$

where

$$W_{test} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & I - e_0 e_0' \end{bmatrix}.$$

Note that W_{test} selects moments not used in the estimation; as a result, $V_{Test}(\sigma_0)$ will generically have full rank. This matrix can be consistently estimated, conditional on σ_0, as

$$\hat{V}_{Test}(\sigma_0) = W_{test}\tilde{R}(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)\hat{\Omega}(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)\tilde{R}(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)'W_{test}$$

where

$$\tilde{R}_T(\theta_0, \delta_0, \sigma_0) = (I - R(\theta_0, \delta_0, \sigma_0)[R(\theta_0, \delta_0, \sigma_0)'W_T(\theta_0)R(\theta_0, \delta_0, \sigma_0)]^{-1}R(\theta_0, \delta_0, \sigma_0)'W_T(\theta_0)).$$

Using

$$(\frac{1}{T}\sum_{t=1}^{T} g_{Test,t}(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)) = W_{test}(T^{-1}\sum_{t=1}^{T} g_t(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)),$$

it follows via the usual arguments that the statistic $\hat{S}(\sigma_0)$ is chi-squared distributed with L degrees of freedom.

Note also that the standard errors associated with our parameter estimates (as in Table 1) can be computed using the standard GMM parameter covariance matrix,

$$V_\theta((\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma) = T^{-1}G_T(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)\hat{\Omega}(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)G_T(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)'$$

where

$$G_T(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0) = [R(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)'W_T(\hat{\theta}(\sigma_0))R(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)]^{-1}R(\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0), \sigma_0)'W_T(\hat{\theta}(\sigma_0)).$$