

The Zero-Beta Interest Rate*

Sebastian Di Tella

Benjamin Hébert

Stanford and NBER

UC Berkeley, Stanford, and NBER

Pablo Kurlat

Qitong Wang

USC and NBER

USC

January 30, 2024

Abstract

We use equity returns to construct a time-varying 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. The zero-beta rate is high and volatile. In contrast to safe rates, the zero-beta rate fits the aggregate consumption Euler equation remarkably well, both unconditionally and conditional on monetary shocks, and can explain the level and volatility of asset prices. We claim that the zero-beta rate is the correct intertemporal price.

*We would like to thank Marianne Andries, Adrien Auclert, Jonathan Berk, John Cochrane, Darrell Duffie, Lars Hansen, Benjamin Holcblat, Arvind Krishnamurthy, Eben Lazarus, Mortiz Lenel, Martin Lettau, Hanno Lustig, Thomas Mertens, Monika Piazzesi, Mikkel Plagborg-Møller, Shri Santosh, Martin Schneider, Ken Singleton, David Sraer, Chris Tonetti, and Motohiro Yogo for helpful comments, as well as seminar and conference participants. We would particularly like to thank Rohan Kekre and Tyler Muir for excellent discussions of our paper. We would also like to thank Paul Delatte for his research assistance. An earlier version of this paper circulated as “The Zero-Beta Rate.”

Replication code (requires a WRDS account): <https://github.com/bhebert/TheZeroBetaRate>

The interest rate is one of the most important prices in a market economy. It captures the intertemporal price of goods and plays a central role in business cycles and monetary policy. It is typically identified with the return of safe assets such as Treasury Bills. In this paper we explore the hypothesis that this assumption is wrong and responsible for long-standing puzzles in macroeconomics and asset pricing. As an alternative, we use equity returns to construct a measure of the intertemporal price that we call the zero-beta interest rate: the expected return of a stock portfolio orthogonal to the stochastic discount factor. The main contribution of the paper is to show that, in contrast to safe rates, the zero-beta rate fits the aggregate consumption Euler equation remarkably well, both unconditionally and conditional on monetary shocks. In addition, the zero-beta rate is high, volatile, and persistent enough to explain market returns without large risk premia.

The Aggregate Consumption Euler Equation. Our starting point is the aggregate consumption Euler equation:

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

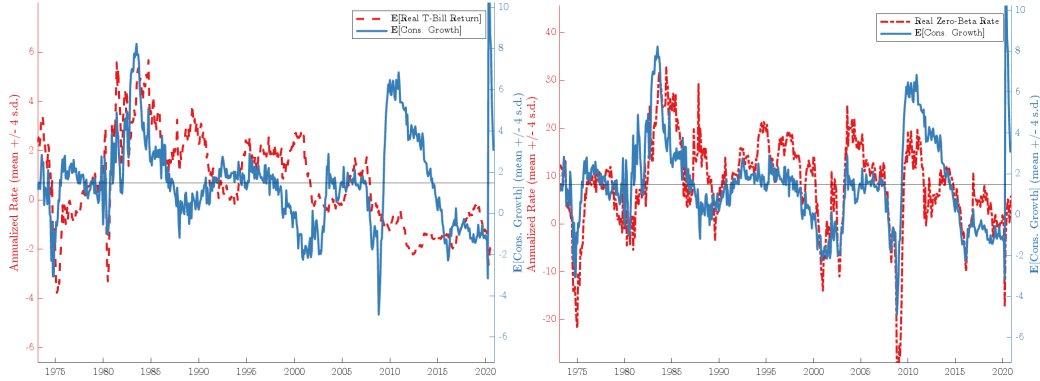
where c_t is aggregate consumption, P_t is the price level, and R_{t+1} is the gross nominal return on any asset. Applying this equation to a risk-free nominal bond ($R_{t+1} = R_{b,t}$) yields the traditional consumption Euler equation, which lies at the heart of macroeconomics.

But, as is well known, the traditional Euler equation does not fit the data.¹ 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 two series are not proportional, indicating that the linearized version of the traditional Euler equation fails to hold.

We conjecture that the failure of the aggregate consumption Euler equation reflects that Treasury Bills and other nominally safe assets have a convenience yield, like cash or deposits. We do not take a stand about the origin of this convenience yield and treat it as a residual. Instead, we assume that equities are not conve-

¹Hansen and Singleton [1983]; Dunn and Singleton [1986]; Hall [1988]; Yogo [2004]

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 3, 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.

nient, and propose to measure the intertemporal price from their returns. Equities, of course, may have risk premia, and empirically these risk premia are not well explained by the traditional consumption Euler equation.² That is, there are factors other than consumption that enter the SDF that prices equities. But if one constructs a portfolio of equities with returns that are orthogonal to the SDF that prices equities, the portfolio will not carry a risk premium, and its expected return (the zero-beta rate) should be the intertemporal price of consumption.

The first step of the paper is to construct the zero-beta rate. We 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. This is our measure of the zero-beta rate. It is not

²See e.g. Mehra and Prescott [1985] on the market risk premium and Fama and French [1993] on the cross-section of equity risk premia

feasible to follow this procedure sequentially because constructing excess returns and estimating betas requires an estimate of the zero-beta rate; for this reason, we estimate all of our parameters simultaneously via GMM. Our procedure allows us to construct a *time-series* for the zero-beta rate, which is essential for our ultimate goal of testing the aggregate consumption Euler equation, and distinguishes our paper from the large literature that estimates the average *level* of the zero-beta rate (following Black et al. [1972]).

The right panel of Figure 1 shows the real zero-beta rate. It is high on average (8.3% annually) and volatile (standard deviation 9.3%). It therefore has a large (around 7.6% per year on average) and volatile spread relative to the expected real return of Treasury bills, which we do not attempt to explain.

The main result of the paper is that the zero-beta rate fits the aggregate consumption Euler equation well. The right panel of Figure 1 shows the result graphically. It shows the real zero-beta rate plotted against expected consumption growth, predicted with the same set of macroeconomic variables. 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 a measure of the intertemporal price of consumption. Our more formal analysis constructs the zero-beta rate using GMM and then tests the non-linear 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 0.5 (if CRRA, risk aversion above 2), but are able to reject higher values of the IES. The aggregate consumption Euler equation also holds conditional on monetary shocks identified using the Romer and Romer [2004] and Nakamura and Steinsson [2018] approaches.

It is important to clarify that our result does not imply that a simple Euler equation describes consumption at the household level. Households may face uninsurable idiosyncratic risk, borrowing constraints, and trading frictions. But a central result in the heterogenous-agent model literature is that this can be consistent with 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.³ 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. In any case, our result is an empirical fact that should be used to discipline and test macroeconomic models. Diamond [2020] is one example of a model in which, consistent with our evidence, the consumption Euler equation holds for equities but not for deposits or bonds.

The zero-beta rate and asset pricing. Taking the view that the zero-beta rate is the correct intertemporal price provides a new perspective on several asset-pricing puzzles.

First, going back to Black et al. [1972], it is well-known that the expected return of an equity portfolio with zero covariance to the market is well in excess of Treasury bill yields.⁴ In other words, the intercept of the security market line (the relationship between market beta and expected returns) is high. One common interpretation of this fact is that the CAPM is wrong—there are priced factors beyond the market—and that once these are incorporated, the zero-beta rate should coincide with safe interest rates. We, however, estimate a zero-beta rate that is high on average despite our use of a rich set of cross-sectional factors. This is consistent with the existing literature, and robust to our specific choices of factors (including factors that are designed to be informative [Ammann et al., 2023]). Lopez-Lira and Roussanov [2020], for example, find that even if one removes almost all of common factors in stock returns, the remainder still has a high average return (see also Kim et al. [2021]). Relative to this literature, our contribution is to construct a time series of the zero-beta rate and show that it fits the aggregate consumption Euler equation.

Second, our time series for the zero-beta rate is high, volatile, and persistent enough to explain the level of market returns and the volatility of price-dividend ratios without requiring an equity risk premium. The average level is roughly equal

³e.g. Krueger and Lustig [2010], Auclert et al. [2020, 2023]

⁴See also Shanken [1986] and more recently Hong and Sraer [2016] and Bali et al. [2017]. The average level of our zero-beta rate is consistent with the findings of Hong and Sraer [2016] and Bali et al. [2017].

to the average return on the market (i.e. the security market line is on average flat), consistent with the existing literature e.g. Hong and Sraer [2016] or Bali et al. [2017].⁵ The time series is volatile and persistent enough to generate sufficient variation in long-term discount rates to match the volatility of price-dividend ratios under the assumption of a constant equity risk premium. That is, contra Campbell [1991], we find that variation in these ratios is not evidence for time-varying risk premia. Taken together, our findings suggest that the equity risk premium (defined relative to zero-beta rates) could in fact be small and constant. However, our estimates are certainly not precise enough to rule out the existence of equity risk premia.

Third, our results suggest that it is instead Treasury bill yields that are puzzling; there is a large and volatile spread between the zero-beta rates and Treasury bill yields. Our results do not depend on any specific theory of this spread. We suspect it reflects convenience of safe bonds, possibly because they can be used by banks to back deposits.⁶ This interpretation is consistent with the aggregate consumption Euler equation working for the zero-beta rate but not for the return of Treasury bills.

Our results are also related to other documented convenience spreads, arbitrages, or quasi-arbitrages between different safe rates. For example, van Binsbergen et al. [2019] documents a ~35 basis point spread between Box rates and overnight index swap (OIS) rates, and Du et al. [2018] find CIP violations of similar magnitudes. The Box rate is effectively safe (i.e. zero variance, not just zero-beta) and therefore potentially “convenient”,⁷ so it does not measure the *inconvenient* zero-beta rate. This point applies to many other safe rate spreads, including the AAA corporate-Treasury spread studied by Krishnamurthy and Vissing-Jorgensen [2012]. If the “less convenient” rate is still somewhat convenient, the spread does not reflect the full difference between inconvenient and convenient rates.⁸ However,

⁵Savor and Wilson [2014] show that the security market line is upward sloping on days with economic announcements and downward sloping on other days, and hence flat on average. However, the intercept (which is what is relevant for the zero-beta rate) is similar across the two sets of days. We elaborate on this point in Section 3.

⁶See Lenel et al. [2019], Diamond [2020], and Piazzesi and Schneider [2021].

⁷As noted by van Binsbergen et al. [2019], “the CBOE recognizes the Box trade is riskless and adjusts its margin requirements appropriately.”

⁸This point has been made by both Diamond [2020] and Angeletos et al. [2023].

these trades do imply a lower bound on the zero-beta rate once we take into account potential leverage, a point we elaborate on in Section 6. For example, Boyarchenko et al. [2018] show that a bank leveraging up the JPY-USD CIP violation can achieve a safe return on equity of over 4% above OIS rates. We indeed find that the zero-beta rate is similar to this levered return in their overlapping sample period.

Our paper is also part of a larger literature on the segmentation of stock and bond markets and the implications thereof. In our view, there are many otherwise-difficult-to-explain phenomena that can be rationalized if one accepts that bonds are not priced by the stochastic discount factor that prices stocks. We discuss these connections in Section 6.

1 Theoretical Framework

The aggregate consumption Euler equation (1) has three distinct implications: (i) a consumption-based SDF prices the cross-section of stocks, (ii) the yield on a safe bond is the inverse of the mean growth of that consumption-based SDF, and (iii) the expected return on a zero-beta portfolio is also the mean growth of that consumption-based SDF. In equations, defining the nominal SDF $\Lambda_t = \delta^t (c_t)^{-\sigma} P_t^{-1}$,

$$\begin{aligned} \text{(i): } 0 &= \mathbb{E}_t \left[\frac{\Lambda_{t+1}}{\Lambda_t} (R_{i,t+1} - R_{j,t+1}) \right], \\ \text{(ii): } R_{b,t}^{-1} &= \mathbb{E}_t \left[\frac{\Lambda_{t+1}}{\Lambda_t} \right], \\ \text{(iii): } R_{0,t}^{-1} &= \mathbb{E}_t \left[\frac{\Lambda_{t+1}}{\Lambda_t} \right]. \end{aligned}$$

Here, $R_{i,t+1}$ and $R_{j,t+1}$ are stock returns, $R_{b,t}$ is the nominal return on a risk-free bond, and $R_{0,t}$ is the expected return on a zero-beta (to the SDF) portfolio.

Our empirical exercise will test (iii) without imposing (i) or (ii). This is important because, as discussed above, neither (i) nor (ii) hold empirically. The first fails because other factors apart from consumption are necessary to price the cross-section of stock returns. The second fails both in the time series (Hansen and Singleton [1982]) and because there are safe bonds with different interest rates.

We first build a “proof-of-concept” monetary model to show that it is possible for (iii) to hold without (i) or (ii). The model features an additional factor in the SDF (to explain why (i) fails) and a non-pecuniary benefit (to explain why (ii) fails). These two features motivate our empirical exercise, in which we require that the zero-beta portfolio have zero beta with respect to many factors (not just consumption) and construct it using stocks and not bonds. Our proof-of-concept model is not intended to be as realistic as possible, but instead to transparently convey the key ideas motivating our empirical exercise. That said, the findings of our empirical exercise are facts about the data that hold irrespective of the validity of our proof-of-concept model.

1.1 Setup

There is a representative household with preferences

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

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. ξ_t is an exogenous stochastic process that will generate fluctuations in the stochastic discount factor that are independent of macro quantities (which is needed to explain why the factors that price of the cross-section of stocks are not macro factors).

This is a standard monetary setting augmented with safe bonds in the utility function and exogenous shocks to the stochastic discount factor. We include both 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. If prices are flexible, money and bonds are neutral and super-neutral. We assume ξ_t is a martingale and independent of M_t , B_t , $\eta_{m,t}$, and $\eta_{b,t}$.

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 (a simplifying assumption that allows

us to equate output and consumption). Denote their nominal return by $R_{i,t+1}$. We assume that at any time there is no portfolio of these risky assets that is risk-free and that there is at least one portfolio of risky assets whose return is uncorrelated with the SDF.⁹

The household's budget constraint is:

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

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 . Note that the representative 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, taking prices P_t , $R_{b,t}$, and $R_{i,t+1}$ as given. Market clearing requires $c_t = y_t$, $B_t = B_t^s$, $X_{i,t} = 0$, and $M_t = M_t^s$. We do not model the supply side of this economy (and in particular do not take a stand on nominal rigidities), so the model does not pin down prices P_t and real output y_t separately. Our results are thus consistent with flexible prices and with different forms of price stickiness.

1.2 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 (note that this expression generalizes the definition used above). 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 $c_t^{-\sigma} / P_t$.

⁹Strictly speaking, the returns $R_{i,t}$ and the SDF are endogenous objects. These assumptions should be understood as assumptions on the underlying payoffs of the risky assets.

The Euler equations for money, safe bonds, and risky assets are:¹⁰

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

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

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

If the household saves in money or safe bonds, it takes into account the convenience those assets provide. That is, the safe bond yield is not the intertemporal price of consumption, but rather the price of a bundle that includes both consumption and some amount of non-pecuniary convenience. As a result, the traditional consumption Euler equation does not work with the safe interest rate $R_{b,t}$ (it would require $\eta_{b,t} = 0$).

Risky assets, on the other hand, do not have convenience, but the household needs to consider the covariance of their return with the SDF, which includes ξ_t . As a result, the traditional consumption Euler equation (with $\xi_t = 1$, i.e. the consumption CAPM) does not hold in general for these assets.

However, the traditional consumption Euler equation holds for a specific type of portfolio of risky assets. Let $R_{p,t+1}$ be the return of a portfolio of risky assets that is conditionally uncorrelated with Λ_{t+1}/Λ_t . We can use (7) to obtain the traditional consumption Euler equation¹¹

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

where the correct interest rate is the $R_{0,t} \equiv \mathbb{E}_t [R_{p,t+1}]$. We call this expected return the “zero-beta rate,” and note that it is the expected return of all non-convenient

¹⁰Here we use the guess 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. Equation (6) then shows that the process for $c_t^{-\sigma}/P_t$ is a function of the process for $M_t/\eta_{m,t}$ and therefore independent of ξ_t , verifying the guess.

¹¹Here we use the fact that the return $R_{p,t+1}$ is uncorrelated with the SDF, and also that ξ_t is a martingale independent of $c_t^{-\sigma}/P_t$.

zero-beta portfolios. It is straightforward to show that the zero-beta rate is the inverse of the mean growth of the SDF, $R_{0,t} = \mathbb{E}_t[\Lambda_{t+1}/\Lambda_t]^{-1}$.

In summary, while the traditional consumption Euler equation does not work with the return of safe assets (because of convenience), nor with risky assets (because of risk premia), it holds with portfolios of risky assets with zero covariance to the SDF, whose expected return is the zero-beta rate. This result is robust to a time-varying and potentially endogenous SDF and convenience on safe assets, and is independent of assumptions on nominal rigidities and the supply environment.

2 Measuring the Zero-Beta Rate

In this section we describe our procedure for measuring the zero-beta rate using stock portfolios. We first impose standard assumptions from the asset pricing literature on the structure of our data, then provide an intuitive description of our estimation procedure, and lastly describe our estimator formally within a GMM framework.

2.1 Asset Pricing Assumptions

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

$$R_{i,t+1} - R_{0,t} = \alpha_i + \sum_{j=1}^K \beta_{ij} F_{j,t+1} + \varepsilon_{i,t+1}, \quad (9)$$

where $\varepsilon_{i,t+1}$ has an unconditional zero mean and α_i and β_{ij} are regression coefficients. We assume the betas are constant over time (specifically, that $\varepsilon_{i,t+1}$ is uncorrelated with the factors conditional on the information at time t).¹² Consistent

¹²We do not assume that the $\varepsilon_{i,t+1}$ are uncorrelated with each other. Equity returns might have co-movement beyond the priced factors.

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 $\varepsilon_{i,t+1}$.

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-traded factors. Our assumption is that

$$\frac{\Lambda_{t+1}}{\Lambda_t} = R_{0,t}^{-1} + \sum_{j=1}^K \omega_{j,t} (F_{j,t+1} - \mathbb{E}_t [F_{j,t+1}]) + \zeta_{t+1}, \quad (10)$$

where ζ_{t+1} is mean zero and uncorrelated with any stock return, conditional on the information at time t , $\mathbb{E}_t [R_{i,t+1} \zeta_{t+1}] = 0$. This assumption implies that our nonlinear SDF is equivalent to a linear SDF for the purpose of pricing equity portfolios,

$$1 = \mathbb{E}_t \left[\frac{\Lambda_{t+1}}{\Lambda_t} R_{i,t+1} \right] = \mathbb{E}_t \left[\left(R_{0,t}^{-1} + \sum_{j=1}^K \omega_{j,t} (F_{j,t+1} - \mathbb{E}_t [F_{j,t+1}]) \right) R_{i,t+1} \right]. \quad (11)$$

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

Up to this point we have a relatively standard asset-pricing setup, except that the mean growth of the SDF and the excess returns are both defined using the zero-beta rate $R_{0,t}$ instead of the Treasury bill yield $R_{b,t}$. This is important: both (9) and (10) embed the assumption that the expected return of any zero-beta portfolio is the zero-beta rate without imposing the traditional assumption that $R_{0,t} = R_{b,t}$.

Instead, we assume that the zero-beta rate is linear in a set of L predictor variables, $Z_{l,t}$ for $l \in \{1, \dots, L\}$, the vector of which we denote Z_t . To simplify our

notation, let $Z_{0,t} = 1$ and center (de-mean) $Z_{l,t}$, for $l \geq 1$. We assume

$$R_{0,t}(\gamma) = \gamma' \cdot Z_t, \quad (12)$$

where $\gamma \in \mathbb{R}^{L+1}$ is a vector of constants. The predictors Z_t include the Treasury bill yield $R_{b,t}$, so we nest as a special case the traditional view that $R_{0,t} = R_{b,t}$.

2.2 An Infeasible Portfolio Procedure

To build intuition, we first describe a portfolio-based procedure to estimate the zero-beta rate $R_{0,t}$. This procedure is infeasible, but it provides the motivation and the moments for the GMM procedure we use.

First note that the zero-beta rate is the expected return of any zero-beta portfolio. Let w be a vector of portfolio weights such that $w' \cdot \beta = 0$, and let $\iota \in \mathbb{R}^N$ be the vector of ones. Equations (9), (11), and (12) imply that

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

Consider then the following procedure to estimate the zero-beta rate:

1. Run time-series regressions on (9) for each test asset i , to obtain α and β . The moment conditions associated with these regressions can be written as $\mathbb{E}[\varepsilon_t \otimes F_t] = 0$, adopting the convention that $F_{0,t} = 1$.
2. Using the estimated betas, choose a zero-beta portfolio, w . Adopt the normalization that the portfolio weights sum to one.
3. Predict the return of the zero-beta portfolio w with an OLS regression on:

$$\underbrace{w' \cdot R_{t+1}}_{R_{p,t+1}} = \underbrace{\gamma' \cdot Z_t}_{R_{0,t}} + u_{t+1}. \quad (13)$$

The moment conditions for this regression are $\mathbb{E}[w' \cdot (R_{t+1} - \iota R_{0,t}) Z_t] = 0$.

This procedure is infeasible because we don't know the zero-beta rate, which is used to construct the excess returns $R_{t+1} - R_{0,t}$ required for the time series regression in

step 1.¹³ That is, step 1 is not feasible until we have done step 3, which requires the weights chosen in step 2, which requires an estimate of the betas from step 1. The solution is to do all steps at once in a GMM procedure that uses the moment conditions of the two regressions. We outline this procedure below.

The choice of the weight vector w is arbitrary, provided that the portfolio has zero-beta. A natural choice, given the goal of accurately estimating γ , is to choose w so as to minimize the variance of u_{t+1} . This is equivalent to choosing the unit-investment, zero-beta portfolio that minimizes the variance of the excess returns $w' \cdot R_{t+1} - R_{0,t}$. We adopt this weighting scheme in our GMM procedure.

2.3 Estimation via GMM

The relevant parameters of our model are $\theta = (\alpha, \beta, \gamma)$, where (α, β) are the regression coefficients associated with (9) and γ is the vector of coefficients in (12). 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. Specifically, we will use the same moments as the portfolio procedure described above.

For given parameter values θ , define the residuals

$$\hat{\varepsilon}_{i,t+1}(\theta) = R_{i,t+1} - \alpha_i(\theta) - (1 - \beta_{i1}(\theta))(\gamma'(\theta) \cdot Z_t) - \beta_{i1}(\theta)R_{m,t+1} - \sum_{j=2}^K \beta_{ij}(\theta)F_{j,t+1} \quad (14)$$

which correspond to regression (9). 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 of $\hat{\varepsilon}_{i,t+1}$.

We use two sets of moments. First, the moments corresponding to the time-series regressions in step 1, $\mathbb{E}[\hat{\varepsilon}_{t+1}(\theta) \otimes F_{t+1}(\theta)] = 0$. There are $N \times (K + 1)$ of these moments, and they can be used to identify the parameters (α, β) .

Second, we impose our assumption that all zero-beta portfolios have an expected return equal to the zero-beta rate. Define the symmetric orthogonal projec-

¹³Constructing the first factor, $F_{1,t+1} = R_{m,t+1} - R_{0,t}$, also requires the zero-beta rate.

tion matrix $H(\theta) = I_N - \beta(\theta)\beta(\theta)'$.¹⁴ Given any weight vector w , $H(\theta) \cdot w$ is a zero-beta weight vector, and all zero-beta weight vectors can be formed this way.¹⁵ We use the moment conditions $\mathbb{E}[H(\theta) \cdot (R_{t+1} - \iota\gamma'(\theta) \cdot Z_t) \otimes Z_t] = 0$, which are the moment conditions from step 3 above, applied to the set of all zero-beta portfolios. Note that this moment condition is not valid without the orthogonal projection matrix $H(\theta)$ —non-zero-beta portfolios can have expected returns that are not equal to the zero-beta rate.

Thus, our estimation method fits precisely into a GMM procedure with moments $\mathbb{E}[g_t(\theta)] = 0$, where

$$g_{t+1}(\theta) = \begin{bmatrix} \hat{\varepsilon}_{t+1}(\theta) \otimes F_{t+1}(\theta) \\ H(\theta) \cdot (R_{t+1} - \iota\gamma'(\theta) \cdot Z_t) \otimes Z_t \end{bmatrix}.$$

As discussed above, we will choose a specific zero-beta portfolio to use in our estimation, which is equivalent to using a reduced-rank weighting matrix. The specific portfolio we choose is the minimum-variance unit-investment zero-beta portfolio. Given an estimate of the covariance matrix of excess returns, $\Sigma_R(\theta)$, we minimize $w' \cdot \Sigma_R(\theta) \cdot w$ subject to the constraints $w' \cdot \beta = \vec{0}$ and $w' \cdot \iota = 1$. 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 minimization problem is given by

$$w(\theta) = \Sigma_R(\theta)^{-1} \cdot \begin{bmatrix} \iota & \beta(\theta) \end{bmatrix} \cdot \left(\begin{bmatrix} \iota' \\ \beta(\theta)' \end{bmatrix} \cdot \Sigma_R(\theta)^{-1} \cdot \begin{bmatrix} \iota & \beta(\theta) \end{bmatrix} \right)^+ \begin{bmatrix} 1 \\ \vec{0} \end{bmatrix}. \quad (15)$$

Our GMM weight matrix is

$$W(\theta) = \begin{bmatrix} I_{N \times (K+1)} & \mathbf{0} \\ \mathbf{0} & w(\theta)w(\theta)' \otimes I_{L+1} \end{bmatrix},$$

¹⁴Here, I_N is the $N \times N$ identity matrix, and $(\cdot)^+$ denotes the Moore-Penrose pseudo-inverse. H is the orthogonal projection matrix with respect to the betas.

¹⁵If w is a zero-beta portfolio, $w = H(\theta) \cdot w$.

which weights our second set of moments by $w(\theta)$ and achieves exact identification. Our weight matrix is not the efficient (in the GMM sense) weight matrix, but is instead akin to a GLS weight matrix, and can also be motivated using a maximum likelihood approach (see Appendix Section I).

Constructing a minimum-variance portfolio requires an estimate of the covariance matrix Σ_R . 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 outline our estimation procedure in Appendix Section B and discuss the details of how we apply the Ledoit and Wolf [2017] estimator to our setting in Appendix Section H.

Casting our estimator in this GMM framework allows us to compute standard errors that take the estimation error of the betas into account. 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 $R_{t+1} - R_{0,t}$ or the first factor $F_{1,t+1} = R_{m,t+1} - R_{0,t}$ without an estimate of the zero beta rate.

The risk prices $\omega_{j,t}$ are not identified by these moments, because we consider only zero-beta portfolios. Standard asset-pricing exercises typically assume one knows the interest rate (the intercept of the SDF) and wants to estimate the risk-prices (the slopes of the SDF). In contrast, we estimate the interest rate and are not interested in risk prices. Our procedure can therefore allow arbitrary time variation in the price of risk. We further discuss the interpretation of our procedure and the potential effects of misspecification of the factors and instruments in Section 6.

3 Data and Results

Our procedure requires a set of equity portfolio returns (the $R_{i,t}$), a set of factors (the $F_{k,t}$), and a set of instruments (the $Z_{l,t}$). When we test the Euler equation we will also need 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 G.

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. For each of these stocks, we compute a five-year trailing beta to the CRSP market return (using monthly data).¹⁶ We then construct 27 (3x3x3) portfolios based on (i) market beta, (ii) market capitalization (i.e. size), and (iii) book-to-market ratios (i.e. value). We construct 27 more portfolios (another 3x3x3 sort) using market beta, size, and investment rates, and another 27 more based on market beta, size, and profitability. These 81 portfolios in total have substantial variation in terms of their exposure to the five factors of Fama and French [2015]. We augment these portfolios with the 49 industry portfolios (based on four-digit SIC codes) from Ken French’s website,¹⁷ and thus consider 130 stock portfolios in total.

There are three 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 (14): 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.

¹⁶Novy-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.

¹⁷Note that the industry 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.

Lastly, and most importantly, we do not include bonds in our set of test assets. This is motivated, as discussed above, by our view that bonds have convenience and hence cannot be priced by the same SDF that prices stocks. If we departed from this choice and included a risk-free bond in the set of tests assets, our procedure would recover the yield on this bond (a convenient rate) instead of constructing a zero-beta rate without convenience.

Factors. Our main specification uses seven 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.¹⁸ 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).

In extensions we also include aggregate consumption growth as a factor. The traditional Euler equation implies that the marginal utility of consumption should be priced, justifying the use of a consumption factor. However, our goal is to first construct the zero-beta rate and then show that it is consistent with an Euler equation; this result is sharper if we do not use consumption data when constructing the zero-beta rate. For this reason, we have opted not to include a consumption factor in our main specification. That said, especially after controlling for the other factors, our equity portfolios have essentially no covariance with consumption growth, and as a result the inclusion or exclusion of the consumption factor has an imperceptible effect on our results. See Appendix Section G for details.

¹⁸Additionally, if the test assets are highly correlated conditional on the factor realizations, the Ledoit and Wolf [2017] shrinkage procedure will produce poor results. For this reason, our main specification uses enough factors to explain a significant portion of the co-movement between our test assets.

Instruments. We have chosen our instruments with the goal of predicting either consumption growth or the intertemporal price of consumption. Our main specification includes five instruments, all of which are available at the monthly frequency starting in 1973.

We include the Treasury bill yield to nest the traditional view that the yield on safe bonds is the intertemporal price. We also 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 variables¹⁹ 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]),²⁰ lagged consumption growth, and using the BAA corporate bond spread (vs Treasuries) 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.

¹⁹Kiley [2022] uses the BAA-Treasury spread instead of the EBP, and the other variables are defined slightly differently.

²⁰We 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.²¹ In our baseline setting, consumption data is not used to construct the zero-beta rate.

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. Appendix Table 2 presents summary statistics for our instruments and consumption data.

3.1 Results

The first column of Table 1 presents the γ coefficients estimated with our GMM procedure, with their associated standard errors. 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 very close to zero. Moreover, all of our predictor variables except the unemployment rate are statistically significant at the 5% level on their own. Collectively, our instruments have a non-trivial ability to predict the zero-beta portfolio return.

Second, the zero-beta rate is high and volatile. The constant in Table 1 is the average monthly nominal return of the zero-beta portfolio (because the instruments have been centered), so our estimates imply an average real zero-beta rate of around

²¹Vissing-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: Constructing the Zero-Beta Rate

	(1)	(2)
	GMM	OLS (inf.)
RF	2.747 (0.706)	2.747 (0.686)
UMP	0.0629 (0.0820)	0.0629 (0.0818)
EBP	-0.993 (0.256)	-0.993 (0.238)
TSP	0.297 (0.0966)	0.297 (0.0958)
CPI_Rolling	-2.196 (0.909)	-2.196 (0.875)
Constant	1.002 (0.112)	1.002 (0.111)
Wald/F	35.74	8.125
p-value	1.1e-06	2.0e-07
RMSE		2.671
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 nominal return of the zero-beta portfolio. 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 on our instruments.

8.3% per year (12.0% annualized nominal, with 3.7% inflation), with a standard deviation of 9.3%.²² In contrast, the Treasury bill yield is low and stable, which implies that there is a large and volatile spread over the Treasury bill yield, roughly 7.6% per year on average. By way of comparison, the mean monthly nominal return of the market in our sample is 0.98% (11.8% annualized), which is to say that the average returns of the zero-beta portfolio are similar to the average return on the market. This finding is consistent with other estimates of the average zero-beta return (Hong and Sraer [2016], Bali et al. [2017]).

Some authors, such as Ammann et al. [2023], find that average zero-beta returns are sensitive to the choice of factors. However, as emphasized above, the level of the zero-beta rate is not the main point of our analysis. Moreover, in our robustness exercises (Appendix Section G), we find that using the Ammann et al. [2023] factors as opposed to the standard Fama and French [2015] factors has little effect on our estimated zero-beta rate. More generally, our results hold across a variety of factor specifications.²³

The high average zero-beta rate is related to, but distinct from, the observation that the security market line (the relationship between market beta and expected returns) is flat. This distinction is best illustrated by the results of Savor and Wilson [2014], who show that the security market line is upward sloping on days with Federal reserve and other economic announcements, and downward sloping on other days. However, a zero-market-beta portfolio earns returns that are similar across announcement and non-announcement days, and are roughly equal to the average market return.²⁴

²²Our econometric analysis occurs at monthly frequencies. For the sake of simplicity and transparency, when discussing annualized returns, we always report monthly returns multiplied by twelve.

²³Because e.g. size is a factor, in the sense that large stocks co-move with other large stocks, a minimum-variance zero-market-beta portfolio will not load substantially on the size factor even if it is not explicitly forced to have zero beta with respect to size. For this reason, explicitly requiring zero beta with respect to size and other additional factors has only a minimal effect on our results.

²⁴See, for example, Panel C of Table 1 of Savor and Wilson [2014]. The intercept (average zero-market-beta excess return) is estimated to be 3.3 and 2.7 bps on announcement and non-announcement days (8.6% and 7% annualized), and the difference is statistically insignificant. It is the market return in excess of the zero-market-beta return that is high on announcement as opposed to non-announcement days.

The standard deviation of the return of the underlying zero-beta portfolio less the zero-beta rate is about 2.7% per month, or 9.4% annualized (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 spread between the zero-beta rate and the Treasury bill yield reflects the convenience of Treasury bills and other safe assets. We treat the spread as a residual and do not propose an explanation for the convenience of safe assets.

Third, the zero-beta rate increases more than one-for-one with Treasury bill yields (i.e. the spread is increasing in bill yields) but is decreasing in the lagged inflation rate. The former is consistent with the observation that Treasury bills have money-like qualities (Nagel [2016]).

In addition, the unemployment rate and term spread positively predict the zero-beta portfolio returns, whereas the EBP negatively predicts those returns.²⁵ 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]. In other words, the zero-beta rate falls when a recession is likely.

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 (i.e. step 3 in the infeasible procedure). 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

²⁵Specifically, a one percentage point increase in the level of the unemployment rate, term spread (annual yield difference), or excess bond premium (annual yield difference), all else equal, leads to an increase in the expected monthly return of our zero-beta portfolio of about .06, .3, and negative one percentage points, respectively.

OLS regression are also moment conditions used in our exactly-identified GMM procedure; as a result, the point estimates are identical. Second, the robust standard errors used in the OLS regression do not account for the fact that betas used to construct the zero-beta portfolio return are themselves estimated. Nevertheless, they are only slightly smaller than 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.

4 The Euler Equation

In this section we show that the zero-beta rate fits the aggregate consumption Euler equation remarkably well, in line with our interpretation of the zero-beta rate as the intertemporal price. The aggregate consumption Euler equation plays a central role in many macroeconomic models. This is true for both representative agent models and in many models with heterogeneous agents (in which the consumption Euler equation may not hold at the individual level but does at the aggregate level; see Werning [2015]).

We start with a linearized version of the Euler equation,


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

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

Let us now revisit Figure 1. This figure compares expected real consumption growth with the expected real return of Treasury bills on the left panel, and with the real 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 panels 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—the zero-beta rate satisfies an aggregate consumption Euler equation.

Figure 2: Results without COVID



Notes: Both panels present results for the “NoCOVID” robustness exercise (Appendix Section G), in which the sample ends in December 2019. They 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 on the right) less expected inflation. Expected inflation is generated from predictive regressions using the instruments described in Section 3, which are the same instruments used to construct the zero-beta rate. Expected consumption growth is generated in the same fashion. In both panels, the right vertical axis is consumption growth, centered at its mean, with limits equal to ± 4 standard deviations. The left vertical axis is the same for the relevant expected real return. All series are annualized.

We find an even more striking congruence between the two series when we end the sample in 2019, excluding the COVID episode, shown in Figure 2. This effect is driven by a change in the coefficients of our regression predicting consumption growth (call them γ_c), which is a result of the unusually large changes in consumption during the COVID period.²⁶ The coefficients that determine the zero-beta rate are largely unchanged (see Appendix Section G).

There is nothing mechanical about these results, but it is important to interpret them correctly. Let $\gamma'_0 \cdot Z_t$ and $\gamma'_c \cdot Z_t$ denote, respectively, the expected real zero-beta rate and consumption growth conditional on Z_t (after demeaning).²⁷ The raw fact is that if one uses the set of macro variables Z_t to separately predict (1) the real

²⁶There are four months in 2020 with more than six standard deviation movements in consumption growth, one of which is a seventeen-standard-deviation event.

²⁷The γ_0 coefficients are the γ coefficients of our estimation less the coefficients of a regression predicting inflation with the Z_t instruments.

return of the zero-beta portfolio, and (2) aggregate real consumption growth, both after demeaning, the two vectors of prediction coefficients are nearly proportional to each other. Importantly, we are not using consumption data when constructing the zero-beta rate. That is, we are choosing the γ_0 coefficients to best predict the real return of a portfolio of stocks, and separately constructing the γ_c to predict real consumption growth. There is no reason to expect those vectors to be proportional absent economic theory.

To understand this in more detail, let us define

$$\Delta = \gamma_0 - \sigma \times \gamma_c,$$

If the linearized Euler equation held exactly for EIS σ and for all structural shocks, Δ should be equal to a vector of zeroes. What the graphical test in Figure 1 shows is that Δ is very close to zero in the Z_t -covariance sense, that is, $Var(\Delta' \cdot Z_t) \approx 0$. If we had only one predictor Z_t ($L = 1$), the result would be mechanical, because we can always define $\sigma = |\gamma_0|/|\gamma_c|$. As long as these two coefficients have the same sign, we will obtain a positive EIS. But once we have more than one predictor, $L \geq 2$, the result is not mechanical. 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 one illustration of this possibility—it is, in effect, 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 D we also conduct a similar placebo test with long-term Treasury bonds, and find the Euler equation also fails for these bond returns.²⁸

Still, our point estimate for Δ is not exactly a vector of zeros. Figure 1 leaves

²⁸The failure of the Euler equation with respect to long term bonds is consistent with our view of bonds as having some degree of convenience and with the presence of time-variation in bond risk premia.

open the possibility that our point estimates for Δ differ substantially from zero for shocks that are rare and small but of economic interest, so that $Var(\Delta' \cdot Z_t) \approx 0$ even though Δ itself is far from zero. We examine this issue in the next two subsections in two complementary ways, one statistical and the other economic. The statistical approach tests the Euler equation, loosely testing that $\Delta = 0$ (we in fact test the fully non-linear Euler equation as opposed to its linearized version). The null hypothesis is that any deviation of Δ from zero reflects estimation error. The advantage of this approach is that we test the Euler equation for all perturbations of Z_t (i.e. all potential structural shocks). The disadvantage of this test is that it might fail to reject the Euler equation purely due to noise. The alternative economic approach, in contrast, independently identifies economically interesting structural shocks (we will use monetary shocks), and studies if the Euler equation holds, at our point estimates, in response to these specific shocks.

Before moving on, let us point out that 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.²⁹ 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 such similar (up to a scaling factor) predictions.


To further address the possibility of overfitting, in Figure 3 we compare a ridge-penalized zero-beta rate (see Appendix Section C for details) with a similarly constructed forecast of consumption growth.³⁰ The ridge penalty is designed to reduce over-fitting and improve the out-of-sample reliability of our estimate of the zero-beta rate. 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 is not penalized, and the coefficient on the Treasury bill yield, which is attenuated towards 1 (i.e. the traditional view that the zero-beta rate is the Treasury bill yield).

²⁹We 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.

³⁰That is, consumption growth is estimated using a ridge regression, whose penalty parameter is selected via ten-fold cross-validation.

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 3 shows that ridge penalization reduces the scale of both expected consumption growth and the zero-beta rate, and that they remain roughly proportional.

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 C) less expected inflation. Expected inflation is generated from predictive regressions using the instruments described in Section 3, 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 \pm four standard deviations. The left vertical axis is the same for the relevant expected real return. All series are annualized.

4.1 Statistical Tests of the Euler Equation

We will test the instrumented version of the non-linear Euler equation,

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

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 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,³¹ so we will test the Euler equation using weak-instruments-robust methods.

The method we adopt follows Stock and Wright [2000] (see Appendix Section J for the details of how their results can be applied to our problem). 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 the procedure is as follows:

1. Conjecture value of σ (the null hypothesis),
2. Estimate $\hat{\delta}(\sigma)$ using the unconditional Euler moments (16) with $l = 0$ ($Z_{0,t} = 1$),
3. 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 .

We can understand our procedure in the context of the GMM estimator expanded to incorporate the Euler equation moments. Importantly, we do not use the Euler equation moments to construct the zero-beta rate. We construct the zero-beta rate exactly as in the previous section, and use the Euler equation moments only to test. Expressing this procedure as one big GMM estimation allows us to properly account for the estimation error in the zero-beta rate (see Appendix Section J for more details). Let $g_{Test,t}(\theta, \delta, \sigma)$ be the vector of the Euler equation moments

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

for $L = 1 \dots L$, which are not targeted in the estimation of θ and δ . Let $\hat{V}_{Test}(\sigma)$ be the variance-covariance matrix of $\frac{1}{T} \sum_{t=1}^T g_{Test,t}(\hat{\theta}, \hat{\delta}(\sigma), \sigma)$. Under the null hypothesis

³¹It is close to the critical values suggested by, e.g., Olea and Pflueger [2013].

of an IES of σ^{-1} , the test statistic

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

is chi-square distributed with L degrees of freedom. 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.

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 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).

To understand what the test is doing, notice that once we set $\hat{\delta}$ to satisfy the Euler equation on average, what the other Euler moments say is that


$$\frac{c_{t+1}^{-\sigma} R_{t+1}(\theta)}{c_t^{-\sigma} P_{t+1}/P_t} \tag{17}$$

is not predictable by Z_t , where R_{t+1} is one of $R_{0,t}$, $R_{b,t}$, or $R_{m,t+1}$. The test rejects a given σ if the instruments Z_t collectively predict (17).

Our results are shown 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.³² The market

³²This 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.

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 not all values of σ . In our baseline analysis, values of σ below 2 (IES above 0.5) are rejected, and values of σ above 2 (IES below 0.5) 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 $\sigma \leq 10$.³³

4.2 Monetary Policy Shocks and the Euler Equation

In this section we study whether the Euler equation holds conditional on a monetary shock. Specifically, we consider whether the linearized Euler holds at the point estimates of the response of consumption growth and the real zero beta rate to an identified monetary shock (i.e. if the effect of the shock on $\Delta' \cdot Z_t$ is small).

Monetary shocks are interesting on their own, but they are especially relevant in our context because monetary policy aims to exploit the Euler equation as a structural relationship. 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

³³Related to this point (the loss of power as σ grows large), we do not place any special emphasis on the value of σ for which $\hat{S}(\sigma)$ is minimal, as the global minimum is always $\sigma \rightarrow \infty$.

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.

We carry out an exercise analogous to the one behind Figure 1, but conditional on a monetary shock. 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:

$$\gamma_j \cdot (Z_{t+h} - Z_{t-1}) = \phi_{0,h} + \phi_{1,h} \cdot mpshock_t + \varepsilon_{t+h}, \quad (18)$$


where $mpshock_t$ is either the Romer and Romer [2004] shock or the Nakamura and Steinsson [2018] shock,³⁴ 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 C), as these are less likely to suffer from over-fitting. Figure 5 shows our results graphically.

Both of the shocks are normalized to raise nominal Treasury bill yields by one percent on impact. They are both estimated to increase real 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 E 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.

These results 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 in line with facts and theory. It is well-known that a monetary contraction lowers expected consumption growth, also shown in Figure 5. Through

³⁴As 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 (18), 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 C), 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.

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.

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 (Singleton [1994], 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, when applied to the correct intertemporal price of consumption (the real zero-beta rate). Impatience shocks and habits or informational frictions on the consumption side are unnecessary, at least at the aggregate level. Safe interest rates, in contrast, are not the correct intertemporal price of consumption, do not enter the consumption Euler equation, and reflect an endogenous spread with the zero-beta rate. Appendix Section F shows an example of a stylized three-period New Keynesian model where a contractionary monetary policy shock can at the same time raise the safe rate and reduce the spread so that the zero-beta rate falls, consistent with our empirical findings.

5 Is there an Equity Risk Premium?

In Table 1, we reported that average return of our zero-beta portfolio was about 1% per month, which is very close to the average return of the CRSP market index over the same period. Thus, taken at face value, our point estimates imply a roughly zero equity risk premium.

This finding is consistent with prior literature. Baker et al. [2011] and Hong and Sraer [2016] both emphasize that low beta stocks have slightly outperformed high beta stocks in post-1968 data; our result shows that this finding is robust to the inclusion of the factors we consider.³⁵ That said, we certainly cannot claim with any statistical power that the equity premium relative to the zero-beta rate is

³⁵A closely related finding was reported in Blitz and Van Vliet [2007].

small or negative. The strongest claim we can make is that it may well be zero after controlling for consumption risk.

This, however, is a radical claim: it implies that, contra Mehra and Prescott [1985], there may be no equity premium puzzle after all. We are not the first to point out that the zero-beta rate is high, or that convenience might help resolve the equity premium and other asset pricing puzzles (see Bansal and Coleman [1996], Lagos [2010], and Herrenbrueck [2019]). Our contribution is to show that the zero-beta rate is both compatible with the consumption Euler equation and high enough to resolve the equity premium puzzle. That said, our view resolves the equity premium puzzle by creating a “convenience puzzle” (why are spreads so large?) about which we have little to say.

However, the high level of equity returns is not the only evidence supporting the existence of equity risk premia. A long line of literature, summarized in Cochrane [2011], shows that time variation in expected equity returns (i.e. discount rates) is required to explain the volatility of valuation ratios (such as the price-dividend or price-earnings ratios). Campbell [1991], in particular, argues that the required variation in expected equity returns must be variation in expected *excess* returns, as there is simply not enough variation in Treasury bill yields.

We revisit this point using our zero-beta rate.³⁶ Our analysis will focus on the price-dividend ratio of a hypothetical consumption claim (a claim whose dividends are proportional to aggregate consumption). We focus on this claim because it enjoys a clear theoretical connection to the consumption Euler equation and because it avoids the need to explicitly model the relationship between dividends and consumption (see Longstaff and Piazzesi [2004] on this issue). For these reasons, this claim has been studied in a variety canonical asset pricing models, such as Campbell and Cochrane [1999], Bansal and Yaron [2004], and Barro [2006]. This clarity comes at a cost, as the price-dividend ratio for a consumption claim is not directly observable. In the analysis that follows, we use the market cyclically-adjusted price-earnings (CAPE) ratio of Campbell and Shiller [1988] as a proxy, recognizing (as pointed out in the papers above) that we expect the market to resemble a levered version of a consumption claim.

³⁶We would like to thank Monika Piazzesi for encouraging us to pursue this direction.

Consider the Campbell-Shiller decomposition as applied to the log price-divided ratio of a consumption claim. Setting aside constants,

$$pd_t = \sum_{j=0}^{\infty} \rho^j (\Delta c_{t+1+j} - r_{t+1+j}^c) + cons.,$$

where r_t^c is the log real return on the claim and pd_t is the log price-dividend ratio of the claim. If the linearized Euler equation holds when applied to the zero-beta rate (as our empirical work suggests), then, taking expectations and continuing to set aside constants,

$$\mathbb{E}_t [pd_t] = (\sigma^{-1} - 1) \mathbb{E}_t \left[\sum_{j=0}^{\infty} \rho^j r_{0,t+j} \right] - \mathbb{E}_t \left[\sum_{j=0}^{\infty} \rho^j (r_{t+j+1}^c - r_{0,t+j}) \right] + cons.,$$

where $r_{0,t}$ is the log real zero-beta rate. That is, the price-dividend ratio of this claim must predict either future real zero-beta rates, future excess returns, or some combination thereof. If expected excess returns are constant, it follows that all time variation in the price-dividend ratio of the consumption claim must be the result of time-variation in future real zero-beta rates.

We explore this hypothesis in a VAR framework. Let us suppose that our predictor variables Z_t follow a VAR(1) process:

$$Z_{t+1} = \mu + \Phi Z_t + \varepsilon_{z,t+1}, \mathbb{E}_t [\varepsilon_{z,t+1}] = 0.$$

This equation, together with (12),³⁷ the Campbell-Shiller approximation, and the hypothesis of constant expected excess returns implies, continuing to set aside constants,

$$\mathbb{E}_t [pd_t] = (\sigma^{-1} - 1) \gamma_0 \cdot (I - \rho\Phi)^{-1} \cdot Z_t + cons. \quad (19)$$

Note that, following Engsted et al. [2012], we write $\mathbb{E}_t [pd_t]$ and not pd_t , avoiding the assumption that this ratio is observable.

Armed with this equation, we will first compute the standard deviation of $\mathbb{E}_t [pd_t]$.

³⁷As in the preceding section, we use γ_0 to denote the coefficients on the real zero-beta rate, which equals γ less the predictive coefficients for inflation.

This is a lower bound on $std(pd_t)$, one that is tight if pd_t is a deterministic function of Z_t . We will then plot $\mathbb{E}_t[pd_t]$ and compare it to the log CAPE ratio. We assume an annual ρ of 0.94 and an IES σ^{-1} of 0.2 (consistent with our estimates).³⁸ We estimate Φ using a standard VAR.

Our preferred specification for this exercise is one that includes the CAPE ratio. Including or omitting this predictor variable had only a minimal effect on our earlier results (see Appendix Section G), as the CAPE ratio's one-period-ahead predictive power is small. It is however, important for our VAR system, because it is more persistent than our other predictor variables. This specification delivers a standard deviation $std(\mathbb{E}_t[pd_t])$ of about 28% over our sample. This number is comparable to the standard deviation of the price-dividend ratio of a consumption claim from the Campbell and Cochrane [1999] model calibration (27%), a model that is designed to match the variability of valuation ratios. The Campbell and Cochrane [1999] model, famously, has a constant risk-free rate and attributes essentially all of that variation to variation in excess returns. Our calculations suggest that it is instead possible to attribute all of the variation to variation in the zero-beta rate. Our result can also be compared to the $std(\mathbb{E}_t[pd_t])$ that would come from the same VAR, if one used the Treasury bill rate instead of the zero-beta rate (9%).³⁹

We find similar results when using the γ from our main specification (which omits CAPE) while still including the CAPE ratio in the VAR system, which is to say that the CAPE ratio predicts future zero-beta rates primarily by predicting the future values of other predictor variables. We estimate a $std(\mathbb{E}_t[pd_t])$ of about 21% when omitting the CAPE variable from the VAR, which is consistent with the intuition that the CAPE ratio should be informative about the price-dividend ratio of a consumption claim.


Figure 6 below plots $\mathbb{E}_t[pd_t]$ (from the specification that includes the CAPE ratio) and the log of the CAPE ratio, centering both series to ignore constants. It

³⁸This value of ρ is consistent with the average value of the CAPE over our sample and with the average real zero-beta rate less average consumption growth over our sample (which is the appropriate comparison under the hypothesis of zero excess returns). Campbell [2017] section 5.3 recommends a slightly higher value (0.95 to 0.96) in the context of the market price-dividend ratio; our results are not very sensitive to this difference.

³⁹That is, our VAR is consistent with the findings of Campbell [1991]; it is our use of the zero-beta rate instead of the bill yield that leads us to a different conclusion.

also plots, for reference, a version of $\mathbb{E}_t[pd_t]$ that uses the Treasury bill yield in place of the zero-beta rate. As noted above, we should expect the CAPE ratio, as the valuation ratio of a claim on corporate cash flows, to be more variable than a valuation ratio for a consumption claim. That said, we are able to generate a striking amount of price-dividend variation in a consumption claim, given that we have imposed the assumption of constant expected excess returns.

Figure 6: Predicted Valuation Ratio of a Consumption Claim



Notes: This figure plots $\mathbb{E}_t[pd_t]$ as computed from (19) and a VAR on the Z_t variables, using a specification (“AltCAPE” from Appendix Section G) that includes the CAPE ratio of Campbell and Shiller [1988], as well as the log of that ratio. It also plots a version of $\mathbb{E}_t[pd_t]$ as computed from (19) using the same VAR, but γ_b (corresponding to the expected real bill return) in the place of γ_0 . All series have been centered.

In summary, our estimate of the zero-beta rate is high enough, volatile enough, and persistent enough to rationalize the high level of equity returns and the volatility of valuation ratios under the assumption of a constant and small equity risk

premium. That said, we certainly cannot prove the non-existence of equity risk premia. We also cannot rule out the possibility that equity risk premia are volatile and tend to rise when the zero-beta rate falls. This possibility allows for time variation in equity premia without generating excess variation in valuation ratios, and could be further explored using VAR methods in the style of Campbell [1991].

6 Interpretation and Misspecification Analysis

Our preferred interpretation of our results is that the zero-beta rate is the correct intertemporal price. While its level and volatility may seem surprisingly large, it simultaneously helps resolve macro and asset-pricing puzzles. It comes, however, at the cost of a large and unexplained spread relative to safe bonds. We believe this nonetheless represents progress, because the type of models that can explain this spread are potentially quite different than those designed to explain large and volatile risk premia or failures of the aggregate consumption Euler equation.

We will stress again that the fact that the zero-beta rate fits an aggregate consumption Euler equation does not imply that there is a representative agent or that individual households' consumption is well described by an Euler equation. As Werning [2015] points out, an aggregate consumption Euler equation is consistent with heterogeneous-agent models where individual households are not on their Euler equation due to frictions such as uninsurable idiosyncratic risk or borrowing constraints. Of course, it is possible to write models where the aggregate consumption Euler equation fails with the zero-beta rate. Our results are evidence against these models. That is, even if our interpretation of the zero-beta rate as the correct intertemporal price is not accepted, the fact that the zero-beta rate fits an aggregate consumption Euler equation is a stylized fact that should be used to discipline and test models.

In the rest of this section we first explore the consequences of potential misspecification in our empirical strategy. We then discuss results from the literature that support our interpretation of the zero-beta rate.

6.1 Misspecification Analysis

In this section we explore what happens if we have the wrong set of asset-pricing factors F_t and/or predictors Z_t . We also clarify the relationship between our procedure and the mathematics of stochastic discount factors. In Appendix Section G we carry out several robustness tests using different set of asset-pricing factors F_t and predictors Z_t .

Asset-Pricing Factors F_t . If we include too many factors F_t , our estimate of the zero-beta rate will be unbiased but noisy. The zero-beta portfolio will still be orthogonal to the true SDF, so its expected return will be the zero-beta rate. But since it is also orthogonal to the extra factors, its variance will not be as small as possible.

The problem is more serious if we omit some asset-pricing factors. In this case the zero-beta portfolio may have some covariance to the omitted factors, and its expected return will contain some omitted risk premium. This is, in fact, the conventional interpretation of the high-intercept of the security market line, but the bias can go in either direction. Our use of the minimum-variance zero-beta portfolio ensures that our portfolio will load only weakly on omitted factors that explain a significant fraction of the cross-sectional variation on stock returns. The main concern is if the omitted factor carries a significant risk price without explaining much of the cross-section of returns.

The success of the consumption Euler equation with the zero-beta rate provides some evidence about the possibility of omitted risk premia. If the omitted risk premium is constant or not predictable by Z_t , only the average level of the zero-beta rate will be affected, so it will still fit the Euler equation but our estimate of δ will be biased. Alternatively, if the omitted risk premium is predictable by Z_t , the zero-beta rate will fit the Euler equation only if the omitted risk premium is highly correlated with the true zero-beta rate, in which case our estimate of σ will be biased. If the omitted risk premium is not highly correlated to the true zero-beta rate, the zero-beta rate should not fit the Euler equation. The fact that it does then provides some support for the time-variation in our measure of the zero-beta rate.

Predictors Z_t . If we include too few predictors Z_t , we will obtain not the true zero-beta rate, but its projection on the included predictors. If the consumption Euler equation holds for true set of predictors it will also hold for those we include. But the consumption Euler equation may hold using the included predictors even when it would not for a larger set of predictors. For example, if the economy was hit by true impatience shocks that move δ_t , the consumption Euler equation should not hold (with a constant δ) once these shocks are included. But if the shocks to δ_t are orthogonal to the included predictors Z_t , the consumption Euler equation will appear to hold.

Including unnecessary predictors causes no particular difficulties, although if the number of predictors is large relative to our sample size we will face an econometric problem (“many weak instruments”). If we both include too many predictors and have omitted some asset-pricing factors F_t , the potential for spurious variation in the measured zero-beta rate arises. As explained above, in this case the measured zero-beta rate will include some omitted risk premium, but if this omitted risk premium is not predictable by Z_t , only the average level of the zero-beta rate will be biased. If we now include some variables in Z_t that do not predict the true zero-beta rate but do predict the omitted risk premium, this will create spurious time variation in the measured zero-beta rate.

Incomplete Markets and Multiple SDFs. The assumption that the vector of ones does not lie in the span of the betas, 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, $R_{0,t} = \mathbb{E}_t [\Lambda_{t+1}/\Lambda_t]^{-1}$, under the assumption that the relevant innovations of the SDF Λ_t are spanned by the factors F_t . However, even if this assumption is correct, since it is not possible to build a zero-variance unit-investment portfolio using stocks, by the absence of arbitrage and given any stochastic process $\tilde{R}_{0,t}$, there is an SDF $\tilde{\Lambda}_t$ that prices stocks ($\mathbb{E}_t [\tilde{\Lambda}_{t+1}/\tilde{\Lambda}_t \times R_{i,t+1}] = 1$) and satisfies $\tilde{R}_{0,t} = \mathbb{E}_t [\tilde{\Lambda}_{t+1}/\tilde{\Lambda}_t]^{-1}$. 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.⁴⁰ 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), which means that it is not possible to build a portfolio with zero covariance to these SDFs $\tilde{\Lambda}_t$. To be concrete, if we took the return of the zero-beta portfolio we build in our procedure and added it as an asset pricing factor, $R_{p,t+1} - R_{b,t}$, we could construct an SDF $\tilde{\Lambda}_t$ with $\mathbb{E}_t [\tilde{\Lambda}_{t+1}/\tilde{\Lambda}_t]^{-1} = R_{b,t}$, but we would not be able to apply our estimation procedure to the augmented factors. Our procedure thus recovers, from the set of all possible rates $\tilde{R}_{0,t}$, the $R_{0,t}$ that does not require an additional factor to which all assets are equally exposed.

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.⁴¹ 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.

⁴⁰It 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 .

⁴¹The 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.

6.2 Related Evidence

There are number of other facts documented in the literature that are consistent with the view that the zero-beta rate and not a safe interest rate is the correct intertemporal price of consumption.

- The factors necessary to price the cross-section of bond returns are not the same as those required to price the cross-section of stock returns, and even when a common factor appears relevant for both stocks and bonds, it usually has a price of risk that differs between the two asset classes (see e.g. Fama and French [1993], Choi and Kim [2018], Collin-Dufresne et al. [2022]). Complementing this literature, our evidence suggests that the mean of the SDF the prices stocks is different from the mean of the SDF that prices safe bonds.
- The reluctance of banks to finance themselves with equity instead of debt can be explained by the observation that zero-beta rates are on average higher than Treasury bill yields, and they represent the correct cost of capital [Baker and Wurgler, 2015]. Baker et al. [2020] explore the implications of this idea in the more general context of corporate capital structure decisions.
- This point is particularly stark when dealing with arbitrage opportunities. After the great financial crisis, persistent apparent arbitrage opportunities, such as covered-interest-parity violations [Du et al., 2018] and the spread between Box rates and overnight index swap (OIS) rates [van Binsbergen et al., 2019], appeared. Each of these arbitrage settings involves *three* safe rates: (i) the return of the lower-rate safe asset, (ii) the return of the higher-rate safe asset, and (iii) the *safe* return on equity of a levered investor who exploits the differential. Each of these returns is a safe return, and the Euler equation can hold for at most one of them. The first two of these returns are often similar. For example, van Binsbergen et al. [2019] documents a ~35 basis point spread between Box rates and overnight index swap (OIS) rates, and CIP violations of similar magnitudes. The third return can be much higher. For example, Boyarchenko et al. [2018] show that, under post-financial-crisis regulations,

a bank leveraging up the JPY-USD CIP violation can achieve a safe return on equity of over 4% above OIS rates. Our view is that the third return (the safe equity return) is the one that is comparable to our zero-beta rate, as both are equity returns, and we indeed find that the two rates are similar in their overlapping sample period. We suggest that the lower two returns are each to some degree convenient. The key difference, in this view, between a zero-beta equity portfolio and e.g. a Box rate trade is that the latter is recognized in law and practice as zero-risk and hence can be easily levered by a bank.⁴² That is, the Box rate can be used to back convenient claims such as deposits, and hence potentially is not a measure of the *inconvenient* zero-beta rate. The same point applies to other arbitrage/convenience spreads that can be levered by a bank. We are not the first to make this point; Diamond [2020] develops this idea in a general equilibrium setting,⁴³ and a closely related argument appears in Angeletos et al. [2023]. Those authors build on this point to explain, respectively, the structure of the financial system and the nature of optimal public debt policy when that debt is convenient.

- Recent work by Itskhoki and Mukhin [2021] explores the “disconnect” between exchange rates and the macroeconomy. In their model, exchange rates are disconnected from interest rates as a result of financial frictions (see their equation (16)), and the traditional consumption Euler equation with safe bond yields holds (their equation (3)). Our results suggest instead a disconnect between bond yields and the intertemporal price of consumption. As suggested by Jiang et al. [2021], such a disconnect (wedge) in the traditional consumption Euler equation offer an alternative means of explaining various exchange

⁴²Even though the box rate is constructed from equity options that are themselves risky, from a bank’s perspective lending via the Box trade is equivalent, in terms of risk and leverage, to a margin account deposit at the CBOE (because, as mentioned previously and by van Binsbergen et al. [2019], the CBOE’s margin rules reflect the riskfree nature of the Box trade). Moreover, as documented by Anderson et al. [2021], many global banks can borrow at rates low enough to make the box and CIP arbitrages profitable.

⁴³Quoting Diamond [2020], “the risk-free rate (reflecting the household’s demand for safe assets) lies strictly below the risk-free rates implied by both the household’s and intermediary’s pricing kernels for risky assets. The intermediary’s pricing kernel (which prices debt securities) implies a lower risk-free rate than the household’s pricing kernel (which prices equities).”

rate puzzles.

- Recent “two-account” heterogenous-agent models such as Kaplan et al. [2018] and Auclert et al. [2023] make a distinction between liquid assets/accounts and illiquid ones, which is conceptually related to the distinction between the Treasury bill yield and zero-beta rate emphasized in this paper. Our evidence, which suggest that an aggregate consumption Euler equation holds for the zero-beta rate but not for Treasury bill yields, is potentially informative about the nature of the frictions or preferences that should distinguish the two accounts in this class of models. However, the mapping between the assumed frictions or preferences and the aggregate consumption Euler equation is subtle, in light of the general equilibrium effects emphasized by Werning [2015], and we leave a further exploration of this issue for future work.

7 Conclusion

The aggregate consumption Euler equation holds when applied to the zero-beta rate, both in the time series and conditional on a monetary policy shock. This is consistent with a small and stable equity premium, when the equity premium is defined with respect to the zero-beta rate. That is, the behavior of stock returns and consumption are jointly consistent with the classic approach to macroeconomics and asset pricing. The mystery is instead safe bonds: there is a large and unexplained convenience spread for safe assets, which we make no attempt to explain in this paper.

References

- Miguel Acosta. The perceived causes of monetary policy surprises. *Published Manuscript* URL https://www1.columbia.edu/~jma2241/papers/acosta_jmp.pdf, 2022.
- Manuel Ammann, Tobias Hemauer, and Simon Straumann. Informative value, profitability, and investment factors. In *Informative Value, Profitability, and Invest-*

- ment Factors: Ammann, Manuel | uHemauer, Tobias | uStraumann, Simon.* [SI]: SSRN, 2023.
- Alyssa G Anderson, Wenxin Du, and Bernd Schlusche. Arbitrage capital of global banks. Technical report, National Bureau of Economic Research, 2021.
- George-Marios Angeletos, Fabrice Collard, and Harris Dellas. Public debt as private liquidity: Optimal policy. *Journal of Political Economy*, 131(11):3233–3264, 2023.
- 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.
- Adrien Auclert, Matthew Rognlie, and Ludwig Straub. The intertemporal keynesian cross. 2023.
- 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, Brendan Bradley, and Jeffrey Wurgler. Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. *Financial Analysts Journal*, 67(1):40–54, 2011.
- 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.
- Ravi Bansal and Wilbur John Coleman. A monetary explanation of the equity premium, term premium, and risk-free rate puzzles. *Journal of political Economy*, 104(6):1135–1171, 1996.
- Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The journal of Finance*, 59(4):1481–1509, 2004.
- Robert J Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866, 2006.
- Florin Bilbiie. Monetary policy and heterogeneity: An analytical framework. 2021.

- Fischer Black, Michael C Jensen, Myron Scholes, et al. The capital asset pricing model: Some empirical tests. 1972.
- David C Blitz and Pim Van Vliet. The volatility effect. *The Journal of Portfolio Management*, 34(1):102–113, 2007.
- 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. A variance decomposition for stock returns. *The economic journal*, 101(405):157–179, 1991.
- John Y Campbell. *Financial decisions and markets: a course in asset pricing*. Princeton University Press, 2017.
- John Y Campbell and John H Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of political Economy*, 107(2):205–251, 1999.
- 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.
- Jaewon Choi and Yongjun Kim. Anomalies and market (dis) integration. *Journal of Monetary Economics*, 100:16–34, 2018.

- 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.
- John H Cochrane. Presidential address: Discount rates. *The Journal of finance*, 66(4):1047–1108, 2011.
- 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.
- Pierre Collin-Dufresne, Benjamin Junge, and Anders B Trolle. How integrated are credit and equity markets? evidence from index options. *Evidence From Index Options (August 26, 2022)*. Swiss Finance Institute Research Paper, (20-65), 2022.
- William Diamond. Safety transformation and the structure of the financial system. *The Journal of Finance*, 75(6):2973–3012, 2020.
- 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.
- Tom Engsted, Thomas Q Pedersen, and Carsten Tanggaard. Pitfalls in var based return decompositions: A clarification. *Journal of Banking & Finance*, 36(5):1255–1265, 2012.
- 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.
- Giovanni Favara, Simon Gilchrist, Kurt F Lewis, and Egon Zakrajšek. Updating the recession risk and the excess bond premium. *FEDS Notes*, 2016.
- 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.
- 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.
- Lucas Herrenbrueck. Interest rates, moneyiness and the fisher equation. 2019. Simon Fraser University Working Paper.
- Harrison Hong and David A Sraer. Speculative betas. *The Journal of Finance*, 71 (5):2095–2144, 2016.
- Oleg Itskhoki and Dmitry Mukhin. Exchange rate disconnect in general equilibrium. *Journal of Political Economy*, 129(8):2183–2232, 2021.
- Zhengyang Jiang, Arvind Krishnamurthy, Hanno N Lustig, and Jialu Sun. Beyond incomplete spanning: Convenience yields and exchange rate disconnect. 2021.
- Greg Kaplan, Benjamin Moll, and Giovanni L Violante. Monetary policy according to hank. *American Economic Review*, 108(3):697–743, 2018.
- 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.
- Ricardo Lagos. Asset prices and liquidity in an exchange economy. *Journal of Monetary Economics*, 57(8):913–930, 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.
- Francis A Longstaff and Monika Piazzesi. Corporate earnings and the equity premium. *Journal of financial Economics*, 74(3):401–421, 2004.
- Alejandro Lopez-Lira and Nikolai L Roussanov. Do common factors really explain the cross-section of stock returns? *Available at SSRN 3628120*, 2020.
- Rajnish Mehra and Edward C Prescott. The equity premium: A puzzle. *Journal of monetary Economics*, 15(2):145–161, 1985.
- 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.
- Pavel Savor and Mungo Wilson. Asset pricing: A tale of two days. *Journal of Financial Economics*, 113(2):171–201, 2014.
- Jay Shanken. Testing portfolio efficiency when the zero-beta rate is unknown: a note. *The Journal of Finance*, 41(1):269–276, 1986.
- Kenneth J. Singleton. *Econometric analysis of representative agent intertemporal asset pricing models*, volume 2 of *Econometric Society Monographs*, pages 277–312. Cambridge University Press, 1994. doi: 10.1017/CCOL0521444608.007.
- 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.

For Online Publication: Online Appendix for “The Zero-Beta Interest Rate,” Di Tella, Hébert, Kurlat, Wang

A Data Details

A.1 Equity Portfolios⁴⁴

We use monthly equity returns in CRSP which can be matched to a firm in COMPUSTAT, from 1973 to 2020. The CRSP returns are augmented with the delisted returns also from CRSP. If the return or delisted return reported in CRSP is missing, we replace the missing value with 0. Then we compute the following object:

$$r_t^{\text{adj}} = (1 + r_t) \left(1 + r_t^{\text{delisted}} \right) - 1$$

where r_t is the reported return, and r_t^{delisted} is the delisted return. We use r_t^{adj} as the variable of return used in all further calculations and therefore simply referred it as

⁴⁴This section constructs equity portfolios following the standard procedure of Fama and French [1993]. Qingyi (Freda) Song Drechsler publishes her implementation of the Fama and French [1993] value and size portfolios on her website, as well as on WRDS (<https://wrds-www.wharton.upenn.edu/pages/wrds-research/applications/python-replications/fama-french-factors-python/>). We follow closely her program when computing accounting variables and handling missing values.

return. We replace negative prices (indicating mid as opposed to closing prices) in CRSP with their absolute values. We drop duplicated records (identified by GVKEY and Year) in COMPUSTAT data. For each firm, we compute the book-to-market ratio, market value, operational profitability and investment with annual accounting data according to Fama and French [1993, 2015].

1. Market value

For portfolio in year t , it is measured with market data for the fiscal year ending in year $t - 1$ and is share outstanding (SHROUT) times price (PRC). Since we will later sort stocks into groups based on their firm characteristics, we use the total market value of a firm, by summing across different stock issuances if a firm has multiple stock issuances.

2. 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) plus deferred taxes and investment tax credit (TXDITC) minus preferred stock redemption value (PSTKRV). Depending on availability, we use redemption (item PSTKRV), liquidating (item PSTKL), or par value (item PSTK) for the book value of preferred stock, in this order. If none of these variables are available, we treat the book value of preferred stock as 0. If deferred taxes and investment tax credit (TXDITC) is missing, we use the sum of deferred taxes (TXDB) and investment tax credit (ITCB). If all of these are missing, we treat the value of deferred taxes and investment tax credit as 0.

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.⁴⁵ We limit the sample to stocks that have at least 24 months of data points in the 5-year window.

At the end of each June, the bottom 20% of stocks, in terms of market value, are dropped.⁴⁶ Then stocks are allocated to three groups according to NYSE break-points, common share stock, on 30% percentile and 70% percentile, with respect to: Book-to-Market and Market Value. When matched with COMPUSTAT, we use accounting information at the end of year $t - 1$ to match with CRSP stock records at the end of June of year t .⁴⁷ This way we make sure the accounting information is available to the investors when creating portfolios according to firm characteristics. Then for each June of year t , we compute 30% percentile and 70% percentile break-

⁴⁵https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/f-f_factors.html

⁴⁶Note if a stock drops below the 20% in the following month, it is still included.

⁴⁷For example, if we are observing a portfolio return at the end of March, 1975, this means this portfolio was constructed at the end of June, 1974, using the account information available at that time, which was from the end of Dec, 1973.

points of Market Beta within each Book-to-Market and Market Value group.⁴⁸ 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.⁴⁹ We compute the value-weighted monthly average returns of these portfolios. The weight we use is market value at the end of May, one month before the portfolio is constructed, times the cumulative return without dividend between current month and the end of May.⁵⁰ By doing this, we eliminate the effect of equity issuance on weights. We construct two additional sorts on market beta, size, and operational profitability, 27 (3x3x3), and on market beta, size, and investment, 27 (3x3x3), in similar fashion.

We then augment these portfolios with the 49 industry portfolios provided on Ken French's website.⁵¹ These 130 portfolios are our baseline test assets. In our robustness exercises (and in earlier versions of the paper), we omit the profitability and investment sorts.

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-

⁴⁸For example, the beta breakpoints of the low Book-to-Market, low Market Value group are different from those of the high Book-to-Market, low Market Value group. We do not impose the NYSE restriction and share code restrictions when computing the breakpoints of Beta.

⁴⁹Every stock in the portfolios needs to have non-missing value for Market Beta, Book-to-Market and Market Value. Otherwise the stock is dropped from the portfolio from the end of June, year t to the end of June, year $t + 1$.

⁵⁰For example, to compute the weights used in March, 1973, we multiple the market values of stocks at May, 1972, by the cumulative returns without dividend of these stocks between May, 1972, and March, 1973.

⁵¹https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_49_ind_port.html

month Treasury bill (the latter as defined below).⁵² The default factor is the return of long-term corporate bonds less the return of long-term Treasury bonds.⁵³ The consumption factor (if used) 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.⁵⁴

3. Term spread

Difference in the yields of 10-year treasury bond⁵⁵ and 1-month treasury bill.

4. Excess bond premium

From Gilchrist and Zakrajšek [2012], as updated by Favara et al. [2016].⁵⁶

5. Unemployment rate

⁵²Specifically, 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.

⁵³We 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.

⁵⁴<https://fred.stlouisfed.org/series/CPILFESL>

⁵⁵<https://fred.stlouisfed.org/series/GS10>

⁵⁶<https://www.federalreserve.gov/econres/notes/feds-notes/updating-the-recession-risk-and-the-excess-bond-premium-20161006.html>

From FRED.⁵⁷

A.3.1 Additional Instruments

1. CAPE

From Campbell and Shiller [1988], as updated on Robert Shiller's website.⁵⁸

2. Shadow spread

Following Lenel et al. [2019], we evaluate a variation of equation (9) in Gurkaynak et al. [2007] at maturity 3/12 for their estimated parameter values.

$$f_t(1/4, 0) = \beta_0 + \beta_1 \left(\frac{1 - \exp\left(-\frac{1/4}{\tau_1}\right)}{\frac{1/4}{\tau_1}} \right) + \beta_2 \left(\frac{1 - \exp\left(-\frac{1/4}{\tau_1}\right)}{\frac{1/4}{\tau_1}} - \exp\left(-\frac{1/4}{\tau_1}\right) \right) + \beta_3 \left(\frac{1 - \exp\left(-\frac{1/4}{\tau_2}\right)}{\frac{1/4}{\tau_2}} - \exp\left(-\frac{1/4}{\tau_2}\right) \right)$$

We then use the estimated 3-month forward rate to proxy for the three month yield. The estimated series is at daily frequency. To aggregate to monthly frequency, we follow the procedure by Lenel et al. [2019]. We first take the average of the yield from last Thursday to this Wednesday. The weekly data is thus measured on each Wednesday. We then take the monthly average of the weekly yields. We download daily 3-month T-bill yield data from the Fred⁵⁹,

⁵⁷<https://fred.stlouisfed.org/series/UNRATE>

⁵⁸<http://www.econ.yale.edu/~shiller/data.htm>

⁵⁹<https://fred.stlouisfed.org/series/DTB3>

then aggregate it to monthly following the same steps as the estimated series. The shadow spread is then the difference between the monthly estimate and the 3-month T-bill yield. In our specification, the shadow spread is divided by 12 to convert into monthly spread.

3. Corporate bond spread

Difference between the yield of Moody's seasoned BAA corporate bonds⁶⁰ and Moody's seasoned AAA corporate bonds.⁶¹

A.4 Consumption

We use the monthly nominal aggregate consumption expenditures on non-durables and services from National Income and Product Accounts (NIPA) Table 2.8.5, the population numbers from NIPA Table 2.6 and the price deflator series from NIPA Table 2.8.4. The consumption growth series used in the estimation is defined as

$$\Delta c_t = 100 \times \log \left(\frac{c_t}{c_{t-1}} \right)$$

where c_t is real consumption expenditures per capita. The sample of consumption growth is from March, 1973 to Dec, 2020. However, when using lagged consumption growth as instrument, the sample is from Jan, 1973, to Oct, 2020 to avoid issues related to measurement error in c_t .

In one robustness exercise, we use non-durable consumption only (excluding services), and otherwise construct the consumption growth variable in an identical

⁶⁰<https://fred.stlouisfed.org/series/BAA>

⁶¹<https://fred.stlouisfed.org/series/AAA>

fashion.

A.5 Summary Statistics

Table 2: Summary Statistics of Instruments and Consumption Growth

	Mean	SD	Min	25%	50%	75%	Max
T-bill Yield (% , Monthly)	0.370	0.288	0.000	0.110	0.390	0.538	1.350
Unemployment (%)	6.287	1.709	3.500	5.000	5.900	7.400	14.700
Excess Bond Premium (% , Annual)	0.075	0.549	-1.085	-0.263	-0.048	0.252	3.472
CAPE	20.761	8.861	6.639	13.015	20.844	26.495	44.198
Term Spread (% , Annual)	1.760	1.366	-3.040	0.750	1.920	2.797	6.940
Shadow Spread (% , Monthly)	0.023	0.021	-0.164	0.012	0.020	0.030	0.137
Corporate Bond Spread (% , Annual)	1.088	0.449	0.550	0.780	0.950	1.270	3.380
Rolling Inflation (% , Monthly)	0.315	0.213	0.050	0.174	0.221	0.380	1.069
Consumption Growth (% , Monthly)	0.124	0.747	-12.909	-0.037	0.139	0.315	6.150

Notes: Summary statistics of the monthly times series included as instruments (for both the main specification and robustness exercises), as well as consumption growth series. The columns 25%, 50%, and 75% are the 25 percentile, median, and 75 percentile of the sample values respectively. Rolling Inflation is the 12-month rolling average of monthly inflation. The rolling inflation sample starts from Jan, 1973, ends at Oct, 2020, to avoid issues related to measurement error, as mentioned in the main text. The sample of consumption growth is from March, 1973 to Dec, 2020. Note, when using lagged consumption growth as instrument, the sample is instead from Jan, 1973, to Oct, 2020. All other instrument variables have sample from Feb, 1972, to Nov, 2020.

B Estimation Algorithm

In this appendix section, we discuss the estimation procedure outlined in the text in more detail. Our procedure takes as input a panel of instruments Z_t , factors F_t , and returns R_t .

The procedure is as follows:

1. Guess a value of γ , and generate $R_{0,t}$ using (12).
 - (a) We use $R_{0,t} = R_{b,t}$ as our initial guess.
2. Run time series regressions for each asset i to estimate α_i and β_i , as in (9).

3. Estimate $\Sigma_R(\theta)$ using our adaptation of the Ledoit and Wolf [2017] procedure; see Appendix Section H for details.
 - (a) Matlab code for the analytical shrinkage estimator we use can be found in Ledoit and Wolf [2020]; their file is called “analytical_shrinkage.m”.
4. Construct portfolio weights for the zero-beta portfolio, using (15).
5. Calculate the GMM objective function

$$\sum_{l=1}^L (T^{-1} \sum_{t=1}^T (w(\theta)' \cdot R_{t+1} - \gamma'(\theta) \cdot Z_t) Z_{l,t})^2.$$

6. Repeat steps 1-5 to minimize the GMM objective.

This procedure takes advantage of the factor that it is possible to analytically compute the (α, β) parameters and weighting matrix $w(\theta)$ given an estimate of γ , which greatly speeds up the computation. However, the procedure is not sufficient to generate standard errors. Standard errors can be calculated using the usual GMM standard error formulas, given our definitions of the moments $g_{t+1}(\theta)$ and weighting matrix $W(\theta)$. This requires calculating the jacobian of the moments with respect to the parameters, which is somewhat involved in our setting but straightforward conceptually. See Appendix Section J for an explicit formula.

C 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, which greatly reduces the time required to compute the estimator. Second, it simplifies the interpretation of our procedure (as described in Section 2.2).

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. We shrink the γ_1 coefficient (which corresponds to $Z_{1,t} = R_{b,t}$) towards one, and the remaining coefficients towards zero. 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.⁶²

⁶²See Hastie et al. [2009].

Our GMM analysis thus solves

$$\hat{\theta}_1 \in \arg \min_{\theta} \underbrace{\sum_{l=0}^L \left(T^{-1} \sum_{t=1}^T w(\theta)' \cdot (\alpha(\theta) + \hat{\varepsilon}_t(\theta)) Z_{l,t-1} \right)^2}_{\text{Instrumented Asset Pricing Moment Squared}} + \underbrace{\psi(\gamma_1(\theta) - 1)^2 + \sum_{l=2}^L \gamma_l(\theta)^2}_{\text{Ridge Penalty}}$$

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

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.⁶³

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, $\{T_1, \dots, T_{10}\}$, 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, $\left(w(\hat{\theta}_m(\psi))' \cdot \hat{\varepsilon}_{t+1}(\hat{\theta}_m(\psi)) \right)^2$, which is the out-of-sample variance of the surprise return of zero-beta portfolio.

We repeated this process for each m , computing the sum of squared moments, and

⁶³The 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.

choose ψ to minimize this value:

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


Once the value of $\hat{\psi}$ is chosen, we compute $\hat{\theta}$ using this value and the full sample.⁶⁴

D Placebo Test

Figure 7 below presents the results of a regression that 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 the two series to be aligned.

⁶⁴This 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 3, 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.

E 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 3 and 4 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). Columns 4-8 in these tables show the coefficients (ϕ_0^l, ϕ_1^l) of the regression

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

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 $l \in \{1, \dots, L\}$, is the effect on the nominal zero-beta rate, by (12). The tables thus illustrate the key drivers of the result that a monetary shock (normalized to increase the nominal safe rate by one percent on impact) 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. Both the Romer and Romer [2004] and Nakamura and Steinsson [2018] shocks result in a significant flattening of the yield curve and an increase in the excess bond premium, which more than offsets the effect of the increase in short rates (see columns 7 and 8 of tables below).. 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 in the scale of their effects. Note also that neither of these tables includes standard errors.

Table 3: Decomposition of the Effects of a Romer-Romer Shock

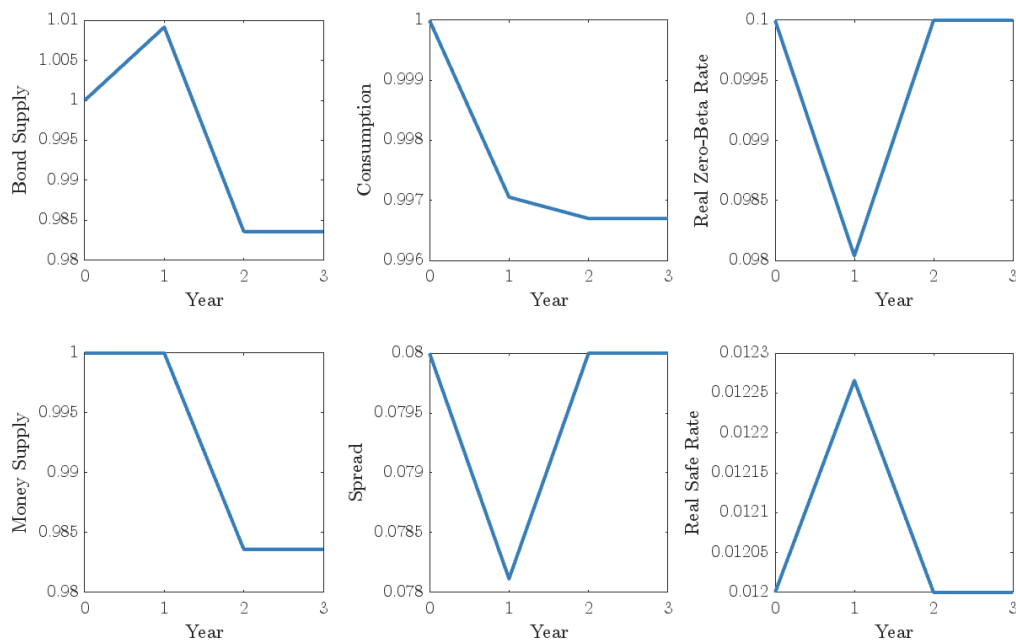
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Real Bill	Real Z.B.	Ex. C. Gr.	g_RF	g_UMP	g_EBP	g_TSP	g_CPI_Rolling
RR_shock	0.338	-1.872	-0.292	1.090	-0.184	-1.450	-0.649	-0.265
Constant	-0.0273	-0.0411	-0.00868	-0.0714	-0.000760	-0.0267	0.0242	0.0116
Observations	413	413	413	413	413	413	413	413

Table 4: Decomposition of the Effects of a Nakamura-Steinsson Shock

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Real Bill	Real Z.B.	Ex. C. Gr.	g_RF	g_UMP	g_EBP	g_TSP	g_CPI_Rolling
NS_shock	1.801	-9.479	-1.873	6.085	-2.264	-3.259	-6.767	-0.877
Constant	-0.0856	-0.00159	0.00492	-0.132	-0.00878	0.167	-0.0375	0.00369
Observations	230	230	230	230	230	230	230	230

F An Example of Monetary Shocks and Endogenous Spreads

Consider a version of the model from Section 1.1 with completely sticky goods prices (leading to standard New Keynesian forces), where the following shock takes place. 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 zero-beta vs. safe rate spread back to steady state, and all quantities and prices are constant thereafter. 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 8 shows the impulse responses.



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.

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.⁶⁵ The bond supply 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

⁶⁵The 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 money and safe bond demand.

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 goods 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.

G 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 (by market value) included in the data sample.
- FF3Industry: With Fama-French 3-factor sorted (27 size by book-to-market by beta portfolios) + industry portfolios (76 total) and instead of 5-factor + industry portfolios (130 total).
 - An earlier version of this paper used this as the main specification.
- LinearCons: With a linear consumption factor $F_{8,t} = \Delta c_{t+1}$.
- WithConsSigma1: With a $\sigma = 1$ non-linear consumption factor, $F_{8,t} = \frac{c_{t+1}^{-\sigma}/P_{t+1}}{c_t^{-\sigma}/P_t}$.
 - The $\hat{S}(\sigma)$ statistic re-estimates the zero-beta rate for each σ in this case, as it depends on σ . The assumption of $\sigma = 1$ applies only to the analog of Figure 1.
- WithConsSigma5: With a $\sigma = 5$ non-linear consumption factor, $F_{8,t} = \frac{c_{t+1}^{-\sigma}/P_{t+1}}{c_t^{-\sigma}/P_t}$.
 - The $\hat{S}(\sigma)$ statistic re-estimates the zero-beta rate for each σ in this case, as it depends on σ . The assumption of $\sigma = 5$ applies only to the analog of Figure 1.
- MktOnly: With only the market factor.

- FF3Only: With only the market, size, and value factors of Fama and French [1993].
- 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, $Z_{6,t} = \Delta c_{t-1}$.
 - We use the second lag to reduce the problem of measurement error in c_t , following standard practice in the literature.
- 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, plus 35 factors $\tilde{F}_{j'',t+1} = F_{j,t+1}Z_{l,t}$ for $j \in \{1, \dots, 7\}$ and $l \in \{1, \dots, 5\}$ (42 factors total). This specification is isomorphic to a model in which the betas to the seven main specification factors are linear in the Z_t variables.

- NDOOnly: 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).
- InformativeFactors: using the five “informative” equity factors of Ammann et al. [2023] in the place of the Fama and French [2015] factors (and still including the two bond factors). Note that the Ammann et al. [2023] data we received from those authors ends in December 2019, and hence the results should be compared with the NoCOVID robustness exercise.


Table 5: Predicting the Zero-Beta Rate by Specification

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Main	NoDrop20	FF3Industry	LinearCons	WithConsSigma	WithConsSigma5	MktOnly	FF3Only	AltBAAS	AltCAPE	LagCons	Shadow	VaryingBetas	NDOOnly	NoCOVID	InformativeFactors
RF	2.747 (0.706)	2.844 (0.671)	2.464 (0.857)	2.745 (0.721)	2.743 (0.720)	2.727 (0.722)	2.695 (0.634)	2.700 (0.637)	2.759 (0.719)	2.856 (0.731)	2.720 (0.707)	2.578 (0.719)	2.744 (1.018)	2.747 (0.706)	2.585 (0.695)	2.597 (0.693)
UMP	0.0629 (0.0820)	0.0163 (0.0768)	0.106 (0.0993)	0.0618 (0.0860)	0.0601 (0.0838)	0.0587 (0.0836)	0.0707 (0.0723)	0.0549 (0.0727)	0.131 (0.0887)	0.0853 (0.0892)	0.0710 (0.0853)	0.0813 (0.0823)	0.0386 (0.157)	0.0629 (0.0820)	0.0422 (0.0986)	0.0460 (0.0964)
EBP	-0.993 (0.256)	-0.999 (0.253)	-0.915 (0.329)	-0.986 (0.260)	-0.981 (0.263)	-0.973 (0.258)	-0.657 (0.221)	-0.694 (0.224)	-0.989 (0.256)	-0.989 (0.256)	-0.986 (0.256)	-1.033 (0.257)	-1.021 (0.443)	-0.993 (0.256)	-1.011 (0.261)	-1.042 (0.240)
TSP	0.297 (0.0966)	0.361 (0.0931)	0.232 (0.122)	0.296 (0.0966)	0.295 (0.0969)	0.293 (0.0966)	0.378 (0.0875)	0.389 (0.0887)	0.292 (0.0980)	0.315 (0.105)	0.287 (0.0990)	0.301 (0.0964)	0.390 (0.161)	0.297 (0.0966)	0.310 (0.105)	0.326 (0.106)
CPL_Rolling	-2.196 (0.909)	-2.084 (0.880)	-2.575 (1.143)	-2.191 (0.926)	-2.199 (0.913)	-2.173 (0.919)	-1.815 (0.809)	-1.791 (0.807)	-2.038 (0.965)	-1.981 (1.063)	-2.193 (0.907)	-2.759 (0.966)	-2.063 (1.287)	-2.196 (0.909)	-1.926 (0.915)	-2.011 (0.920)
BAAS									-0.601 (0.411)							
CAPE									0.0121 (0.0259)							
LagCons											0.0938 (0.115)					
ShadowSpread												13.24 (7.769)				
Constant	1.002 (0.112)	1.039 (0.106)	1.031 (0.137)	1.003 (0.112)	1.004 (0.112)	1.006 (0.113)	1.045 (0.101)	1.072 (0.103)	1.023 (0.112)	1.008 (0.112)	1.001 (0.112)	1.016 (0.111)	0.968 (0.189)	1.002 (0.112)	1.029 (0.113)	1.018 (0.112)

Standard errors in parentheses


Notes: This table presents the γ coefficients from our GMM estimation across the various specifications defined in Appendix Section G. 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 9: NoDrop20: Results with Bottom Decile Stocks Included




Notes: Both panels present results for the “NoDrop20” robustness exercise (Appendix Section G). 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 \pm 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: FF3Industry: Results with FF3+Industry Portfolios




Notes: Both panels present results for the “FF3Industry” robustness exercise (Appendix Section G). 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 \pm 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 G). 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 \pm 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: WithConsSigma1: Results with Consumption Factor, $\sigma = 1$




Notes: Both panels present results for the “WithConsSigma1” robustness exercise (Appendix Section G). 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 \pm 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: WithConsSigma5: Results with with Consumption Factor, $\sigma = 5$




Notes: Both panels present results for the “WithConsSigma5” robustness exercise (Appendix Section G). 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: MktOnly: Results with only the Market Factor




Notes: Both panels present results for the “MktOnly” robustness exercise (Appendix Section G). 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: FF3Only: Results with only the FF3 Factors




Notes: Both panels present results for the “FF3Only” robustness exercise (Appendix Section G). 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 \pm 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: AltBAAS: Results with the BAA-AAA Spread instead of the EBP




Notes: Both panels present results for the “AltBAAS” robustness exercise (Appendix Section G). 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 \pm 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 17: AltCAPE: Results with the CAPE Instrument Included



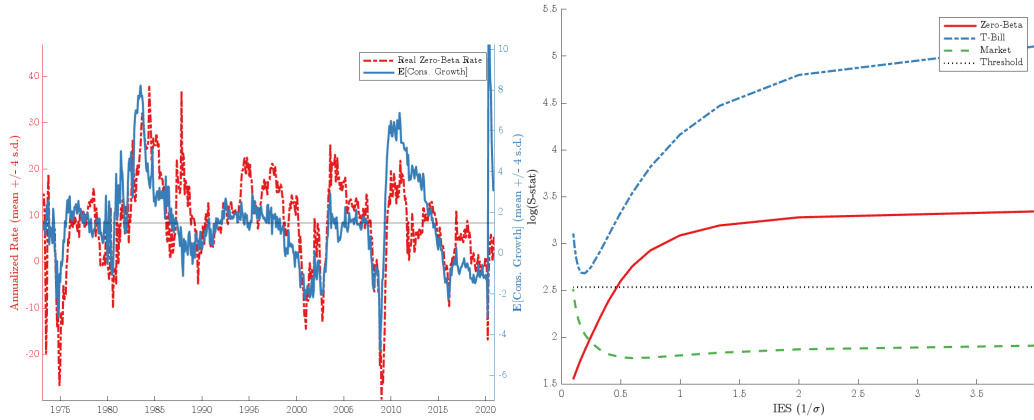
Notes: Both panels present results for the “AltCAPE” robustness exercise (Appendix Section G). The left panel plots the estimated real zero-beta rate against expected consumption growth. Expected real returns are constructed using the nominal zero-beta rate less expected inflation. Expected inflation 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 ± 4 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: LagCons: Results with the Lagged Consumption Growth Instrument Included




Notes: Both panels present results for the “LagCons” robustness exercise (Appendix Section G). 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 ± 4 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: Shadow: Results with the Shadow Spread Instrument Included




Notes: Both panels present results for the “Shadow” robustness exercise (Appendix Section G). 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 \pm 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 20: VaryingBetas: Results with Linear Betas




Notes: Both panels present results for the “VaryingBetas” robustness exercise (Appendix Section G). 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 \pm 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: NDOOnly: Results with Non-Durable Goods Consumption




Notes: Both panels present results for the “NDOOnly” robustness exercise (Appendix Section G). 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 ± 4 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 22: NoCOVID: Results Excluding 2020



Notes: Both panels present results for the “NoCOVID” robustness exercise (Appendix Section G). 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 ± 4 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 23: InformativeFactors: Results with “Informative Factors,” Excluding 2020



Notes: Both panels present results for the “InformativeFactors” robustness exercise (Appendix Section G). 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 ± 4 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).

H 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,

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

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

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

where $\text{diag}(M)$ 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

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

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

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

This pre-conditioning in effect imposes a uniform prior about the orientation of the eigenvectors of $\hat{\Sigma}_c(\theta)$, as opposed to about the orientation of the same for $\hat{\Sigma}_R(\theta)$. 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.

I 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 (12), 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 (14),

$\hat{\varepsilon}_t$, 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.⁶⁶ 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,

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

where $\hat{\varepsilon}_{t+1}(\theta)$ is defined from (R_{t+1}, F_{t+1}, Z_t) as in (14).

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

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

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 $E[F_{j,t} \cdot \hat{\varepsilon}_{i,t+1}(\theta)] = 0$ for all (i, j) , which are the moment conditions associated with the time series regressions (9).

⁶⁶Shanken [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

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

We will next show that under the stated assumptions, the “portfolio weight” $w_{MLE} = (\iota - \beta_{\cdot,1})' \cdot \Sigma_{\varepsilon}^{-1}$ 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 $j = \{1, \dots, K\}$. 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

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

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. In summary, our GMM estimator is essentially the MLE estimator, modified to avoid overfitting.

J 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 4 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 3 is the same as the zero-beta rate being tested in Section 4. 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

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

and that our weight matrix used in estimation is

$$W_T(\theta) = \begin{bmatrix} I_{N \times (K+1)} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & w_T(\theta) w_T(\theta)' \otimes I_{L+1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & e_0 e_0' \end{bmatrix}.$$

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

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

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 $\Omega(\theta, \delta, \sigma) = \mathbb{E}[\Psi(\theta, \delta, \sigma)\Psi(\theta, \delta, \sigma)']$ be the limiting covariance matrix; we assume it can be consistently estimated using

heteroskedasticity-robust methods (Assumption D'' of Stock and Wright [2000]).⁶⁷ 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

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

Note that, following Stock and Wright [2000], we have assumed that $m_2(\cdot)$ does not depend on T , without imposing this assumption on m_{1T} . In our specifications with

⁶⁷Consistent 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]).

a non-linear consumption factor,

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

which is to say that only the non-linear consumption factor depends on σ . In our other specifications, $F_{t+1}(\sigma, \gamma)$ is invariant in σ , but in both cases, the Euler equation moments depend 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 .⁶⁸ It follows that the usual GMM formula applies,

$$\begin{aligned} \sqrt{T}((\hat{\theta}(\sigma_0), \hat{\delta}(\sigma_0)) - (\theta_0, \delta_0)) &\Rightarrow -[R(\theta_0, \delta_0, \sigma_0)' W(\theta) R(\theta_0, \delta_0, \sigma_0)]^{-1} \times \\ &R(\theta_0, \delta_0, \sigma_0)' W(\theta) \Psi(\theta_0, \delta_0, \sigma_0). \end{aligned} \quad (20)$$

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

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

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

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

⁶⁸This 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

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

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

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

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

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