

NBER WORKING PAPER SERIES

TAX NEWS: THE RESPONSE OF HOUSEHOLD SPENDING TO CHANGES IN
EXPECTED TAXES

Lorenz Kueng

Working Paper 20437

<http://www.nber.org/papers/w20437>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

August 2014

I am grateful to Alan Auerbach and Yuriy Gorodnichenko for their encouragement and support throughout this project. David Card, Pierre-Olivier Gourinchas, Dwight Jaffee, Erik Hurst, Patrick Kline, Nick Li, Brian Melzer, Dina Pomeranz, Demian Pouzo, Emmanuel Saez, Adam Szeidl, Mu-Jeung Yang and seminar participants at numerous places all provided valuable feedback. I thank Daniel Feenberg for support with mapping the CEX to TAXSIM, Kevin Moore for sharing his code for the SCF, and Craig Kreisler for patiently answering all my questions about the CEX. All remaining errors are mine. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Lorenz Kueng. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Tax News: The Response of Household Spending to Changes in Expected Taxes
Lorenz Kueng
NBER Working Paper No. 20437
August 2014, Revised August 2016
JEL No. E21,E62,G12,H31,H74

ABSTRACT

Although theoretical models of household behavior often emphasize fiscal foresight, empirical studies of household consumption have yet to document the role of news about tax changes. Using novel high-frequency bond data, I develop a model of the term structure of municipal yield spreads as a function of future top income tax rates and a risk premium. Testing the model using the presidential elections of 1992 and 2000 as two quasi-natural experiments shows that financial markets forecast future tax rates remarkably well in both the short and long run. Combining these market-based tax expectations with data from the Consumer Expenditure Survey, I find that spending of higher-income households increases by close to 1% in response to news of a 1% increase in expected after-tax lifetime (permanent) income. These findings imply that by ignoring anticipation effects, previous estimates of the total effect of a tax change could be substantially biased.

Lorenz Kueng
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
and NBER
l-kueng@kellogg.northwestern.edu

This paper addresses two related questions: How predictable are personal income tax rates in the U.S., and does household spending respond to news about future taxes even before the rates change? These are classic questions with broad implications both for policy makers and for applied economists. For instance, our assessment of the effectiveness of fiscal policy as a tool to stabilize business cycles depends crucially on our ability to capture all behavioral responses to a specific tax reform, including any anticipation effects. Such anticipation effects in turn depend on whether the tax policy is forecastable and whether and to what extent households take such information into account in their decision making. A major source of disagreements among economists concerning these issues has been the difficulty to credibly identify both the timing and the magnitude of changes in expected taxes, and to estimate the household response to such fiscal news.

This study tackles these issues as follows: First, I measure expected future taxes using novel high-frequency data of municipal bond yields with different maturities. Second, I combine these market-based expectations with micro-level data from the Consumer Expenditure Survey (CEX) to quantify the degree of fiscal foresight by estimating the effect of the tax news shocks on household spending.

In order to identify news about future taxes, I exploit the differential tax treatment of two types of bonds. Interest on municipal bonds is tax-exempt (including state bonds), while interest on Treasury bonds is subject to federal income taxes. Thus, relative price changes between municipal and Treasury bonds reflect changes in expected future tax rates, holding fixed other risk factors. I go beyond identifying the timing of news to directly measure the entire path of expected tax rates comparing pairs of bonds of different maturities. The tax news shocks used in this paper therefore measure not only *when* households receive information, but also *how much* information they receive.

Identifying the entire expected path of expected taxes in turn is important when assessing the degree of forward-looking household spending behavior. Combining the market-based tax news with data from the CEX, I calculate changes in expected after-tax lifetime income for each household in the survey. Models with forward-looking households predict that consumption should respond to news about *all* future after-tax income, not just current taxes or income.

Focusing on the sample of higher-income households, for which the identified news shock is most directly related to changes in expected after-tax lifetime income, I find that household spending responds strongly to the bond-based tax news. The hypothesis that higher-income households are not forward-looking is strongly rejected. Interpreting the estimated response through the lens of a standard life-cycle consumption model, spending on nondurables and services increases by about 1% in response to news of a 1% increase in after-tax lifetime income, i.e., with changes in permanent income.

The idea behind the identification approach can be illustrated using a particular short-term and long-term municipal yield spread, say a 2-year and a 15-year spread. All else equal, the 2-year spread reflects the bond market's expectation of marginal tax rates over the next 2 years, while the 15-year spread contains tax expectations over the next 15 years. Suppose that new information

arrives about a potential tax change in 3 years from today. This would affect the 15-year but not the 2-year yield spread. Hence, the two yield spreads are informative about *when* those tax changes are expected to occur. Alternatively, suppose that market participants learn about a tax reform that is expected to be implemented within a year from today. Such information would have a large effect on the 2-year spread. However, the effect on the 15-year spread depends on the *expected persistence* of the reform. If the reform is expected to be only temporary, then it would not affect this long-run yield spread much. Hence, the term structure of municipal yield spreads allows us to infer both the timing and the expected persistence of a tax shock. Using new data on high-quality municipal bond yields going back to the late 1970s, I infer the entire path of expected tax rates over a forecasting horizon of up to 30 years at any given point in time and compare it to the subsequent evolution of the corresponding tax rate.

Of course, there are other factors than expected tax rates that affect municipal yield spreads. In fact, most models of the term structure of municipal yield spreads abstract from changes in expected tax rates and instead focus on other determinants. I therefore begin by providing graphical, non-parametric evidence that changes in expected tax rates are indeed an important determinant of municipal yield spreads, especially at lower frequencies. I then use two quasi-natural experiments—the presidential elections in 1992 and 2000—and daily election probabilities from a political prediction market to formally estimate the degree of fiscal foresight in bond prices during the elections and to identify which marginal tax rate is reflected in the yield spreads. Changes in the election probabilities provide an alternative measure of changes in expected future tax rates because each of the major presidential candidates had a very different tax reform proposal during both elections, and tax reforms were a major part of the candidates’ election platforms. I show that the municipal yield spreads indeed respond strongly to changes in the election probabilities during both elections, both statistically and economically, and with the expected sign, positively with Bill Clinton’s and negatively with George W. Bush’s betting contract, respectively. Moreover, the quantitative response in both cases is consistent with differences in the candidates’ proposals for the top personal income tax rate.

These quasi-natural experiments allow the yield spreads to be decomposed into tax expectations and other factors, but are not available over most of the sample period. As an alternative approach to decomposing the yield spreads, I extend the traditional model of municipal yield spreads to allow for changes in expected future tax rates. The model also takes into account factors other than tax news that can affect the yield spread, such as liquidity and default risk or tax uncertainty. In order to minimize the influence of these other factors in the empirical analysis, I use taxable and tax-exempt bonds that are as similar as possible on all pricing dimensions other than their tax status. For instance, I only use the most liquid municipal bonds with the lowest credit risk to measure the tax-exempt yield curve, specifically AAA-rated general obligation state bonds. For the taxable yield curve, I use the less liquid off-the-run Treasury bonds.

I first use this model to interpret the reduced-form estimates of the two quasi-natural experiments. The structural model shows that the reduced-form coefficients are remarkably consistent

with the evolution of the effective marginal tax rate of the top 1% of the income distribution over the subsequent 30 years or so following the election, both in terms of the timing and persistence of the path of top tax rates. Therefore, bond markets have strong fiscal foresight with respect to both the timing and the magnitude of the shocks. I then use the model to decompose the term structure into tax news and other risk factors for the entire sample. I show that the decomposition performs well for the remaining large tax changes in the sample period for which we do not have quasi-natural experiments, in particular the Reagan tax cuts in the early 1980s.

I directly address concerns that these other risk factors that affect the yield spreads could also affect household spending. First, I show graphically that the low-frequency component of the term structure of yield spreads reflects changes in expected tax rates, while high-frequency movements are driven by other risk factors such as liquidity and credit risk. Consistent with this fact I find that spending of higher-income households only responds to low-frequency changes of the yields, but not to the high-frequency component. Moreover, I show that failing to decompose the yield spreads into low- and high-frequency components attenuates the spending response substantially. This attenuation is consistent with the high-frequency component introducing noise from the econometrician's perspective, i.e., spending-irrelevant innovations in bond yields.

Second, I show that the results are robust to directly controlling for such risk factors with standard measures used in the literature. Both the results from the two quasi-natural experiments and the household spending response are robust to directly controlling for such bond risk factors.

Third, consistent with the hypothesis that those other risk factors introduce noise in the household spending analysis, I find that during periods of bond market stress, when the yield spreads are dominated by other risk factors rather than expected taxes, household spending does not respond to the fiscal news measure based on these bond yields. On the other hand, household spending strongly responds to the yield-based tax news measure when restricting the sample only to periods prior to large tax reforms.

Fourth, I control flexibly for other factors that could simultaneously drive both household spending and expected tax rates or bond yields. In particular, I use time fixed effects (i.e., year-by-month dummies) to flexibly control for macroeconomic shocks, and household income and demographics to control for other factors such as life-cycle effects and intertemporal substitution. Moreover, I focus on post-1980 tax reforms for which the extensive narrative analysis of historical records by [Romer and Romer \(2010\)](#) reveals that they are not driven by countercyclical motives or by contemporaneous changes in government spending.

While [Romer and Romer \(2010\)](#) label these tax changes that are orthogonal to the current state of the economy and to current government spending as exogenous, forward looking households might also care about future changes in government spending, which in the extreme case of an infinitely-lived or dynastic representative agent with rational expectations leads to Ricardian equivalence and no effect on current consumption as shown by [Barro \(1974\)](#). However, note that the focus of this paper is on higher-income households. It seems reasonable that many cuts in future government

spending or transfers, such as the welfare reforms in the 1990s which partially financed the tax cuts of the 1980s, affected higher income households much less than the previous tax cuts. Consistent with this hypothesis I find no difference in behavior between households with and without mortgages, even though the former might fear a repeal of the mortgage interest deduction which mostly benefits higher-income households. Similarly, I do not find a difference in the response of families with and without children. Hence, higher-income parents do not factor in potentially lower government services more than households without children even though their children would most likely be affected by those changes. Finally I show that younger households tend to respond slightly more to the tax changes than older households despite the fact that they are more likely to be affected by any potential off-setting government spending changes in the future.

Although I do not find empirical evidence in favor of Ricardian equivalence, it is possible that the tax changes have effects on future pre-tax income, say through supply-side effects such as higher economic growth, wages and labor supply. The tax reforms in the sample affected marginal rates substantially and are not the lump-sum tax changes required to derive Ricardian equivalence. Hence, higher-income households today might respond to the combined net effects of changes in expected future tax rates on future after-tax income. Consistent with this possibility I find that the tax reforms in the 1980s, which lowered marginal tax rates the most, elicit the largest consumption response with an elasticity of -1.4 , although the precision is too low to statistically reject that the response is the same as over the entire sample period 1980-2001. Nevertheless, we can strongly reject the hypothesis that higher-income households do not take into account the (net) effects of expected future tax rates in their current consumption decisions.

Finally, consistent with the hypothesis that the expected tax changes do affect higher- and lower-income households differently, I show that only higher-income households that are substantially affected by the expected tax changes captured in the municipal yield spreads do indeed respond to the fiscal news. Specifically, households in the top income tax brackets strongly respond to news about future top income tax rates, while households in lower tax brackets whose tax rates exhibit low correlation with the top tax rate do not respond to the same measure of fiscal news.

This paper touches on several strands of literature. To the best of my knowledge, it provides the first direct evidence of the effect of news about future after-tax income on consumption at the household-level. The household consumption response to a well identified news shock estimated in this paper therefore fills an important gap in the large literature that estimates the life-cycle model of consumption; see e.g., [Jappelli and Pistaferri \(2010\)](#) for an excellent survey.

This paper also relates to a more recent literature in macroeconomics, which uses aggregate time series variation to quantify the role of news and uncertainty shocks in explaining business cycles; see [Lorenzoni \(2011\)](#) for a survey of the macroeconomics of news shocks, and [Bloom \(2009\)](#) and [Baker, Bloom and Davis \(2016\)](#) and the literature therein for studies of the effect of uncertainty shocks on the macro economy. This study contributes to the news shock literature by identifying the news component of an important observable shock—changes in top personal income tax rates. The

previous literature instead focuses mostly on news about future Total Factor Productivity (TFP), which is not directly observable. Hence, a particular advantage of this approach is the fact that the performance of the news measure can be visually assessed by comparing it to the evolution of the corresponding realized tax rate.

The paper proceeds as follows. [Section 1](#) first provides non-parametric evidence of fiscal foresight contained in bond prices and then extracts forward tax rates from a model of the term structure of municipal yield spreads. [Section 2](#) uses two quasi-natural experiments to quantify the degree of fiscal foresight in bond prices and to assess the model's ability to relate subsequent legislated tax changes to the observed reduced-form responses of the yield spreads. [Section 3](#) uses the identified tax news shocks to estimate the spending response of higher-income households, and [section 4](#) provides extensive robustness checks. [Section 5](#) concludes.

1 Determinants of Municipal Yield Spreads

This section extracts the path of expected tax rates from the term structure of municipal yield spreads which are used in section 3 to estimate the response of high-income household to fiscal news. I start by providing non-parametric evidence that future tax rates predict changes in relative yield spreads well, especially at lower frequencies. I then analyze three other factors that could affect the yield spread besides expected tax rates, namely default risk, state income taxes, and changes in liquidity risk or tax uncertainty. I limit the effect of liquidity and default risk by using only a subset of high-quality municipal bond yields, and matching them with taxable Treasury yields of off-the-run bonds which are less liquid and hence resemble more closely the municipal bonds than the typical on-the-run Treasury yields.¹ However, these unobserved factors affect the yield spreads at higher frequencies. I therefore extend the traditional model of municipal yield spreads to allow for both changes in expected tax rates and for a latent factor that captures such unobserved risks.²

In order for the implied tax rates—which I call break-even tax rates (BETR) because of the similarity to the concept of the break-even inflation rate—to reflect expected tax rates, it is important to identify who the marginal investor is and hence which tax parameter is identified by the yield spread. Moreover, the marginal investor has to be stable over the sample period in order for changes in break-even tax rates to reflect changes in expected tax rates rather than changes in the marginal investor. Household-level data from the Survey of Consumer Finances and aggregate data from the Flow of Funds suggest that the marginal investor during the sample period from 1977 to 2001 is indeed a high-income household. I formally test this hypothesis in the next section using two presidential elections as quasi-natural experiments, and I estimate the degree of fiscal foresight contained in the municipal yield spreads. While these tests do not rely on the model of the term

¹ Treasury bonds that are issued before the most recently issued bond of a particular maturity are called off-the-run, while the most recently issued bond is called on-the-run.

² Interpreting the model more broadly as a linear approximation of a more general but potentially non-linear model, such as the model proposed by [Green \(1993\)](#) for example, the latent factor also reflects higher-order terms of the linear approximation.

structure of yield spreads developed in this section, I find that the model performs well in relating the reduced-form responses estimated during the two presidential elections to the subsequent legislated tax changes.

This section also relates to a large literature in empirical finance that analyzes the determinants of the municipal yield spreads. Early studies such as Fama (1977), Poterba (1986), Kochin and Parks (1988), Green (1993), or Fortune (1996) use different tax rates to model the municipal yield spread. This paper contributes to this literature by identifying the marginal investor for an important class of assets and shows that the disagreement about the fundamental determinants of the municipal yield spread might be due to changes in the marginal investor over time, in particular before and after the 1970s. While high-income households appear to be the marginal investors since the late 1970s, bank corporations were probably the marginal investors in the 1960s and early 1970s. Moreover, I show that economic fundamentals explain most of the variation in the municipal yield spread over long horizons, i.e., at lower frequencies, while liquidity, credit risk or discount rate shocks are important for high frequency movements. Previous studies that found changes in tax rates to be important determinants of changes in the municipal yield spread include Mankiw and Poterba (1996), Slemrod and Greimel (1999), and Ayers, Cloyd and Robinson (2005).

1.1 Bond Data

Municipal bond yields are taken from two sources. Weekly data of an index of AAA general-obligation municipal bond yields between 1976 and 1983 come from Delphis Hanover. From 1983 on I use daily data provided by Thomson Municipal Market Monitor. I use the one year of overlap between the two sources to verify that the two series are almost identical. The data reflect generic yield curves of transactions of \$2 million or more that are the general preference of institutional investors and reflect the prices of the highest-grade AAA rated general obligations issued by states. The bonds used to generate the yield curve are not subject to the alternative minimum tax (AMT).

The Treasury term-structure is the off-the-run par yield curve taken from Gürkaynak, Sack and Wright (2007). On-the-run Treasury yields and corporate bond yields are also used as controls for credit and liquidity risk spreads in the quasi-natural experiments and are taken from the Board of Governors of the Federal Reserve System, Historical Data Table H.15. Additional data that reflect bond market conditions include the 30-Day Visible Supply, which is the total dollar volume of new municipal bonds carrying maturities of 13 months or more that are scheduled to reach the market within 30 days and is taken from the Bond Buyer newspaper.

1.2 Non-Parametric Evidence of Fiscal Foresight in Bond Prices

I define the break-even tax rate $\theta_{t,m}$ as the parameter that equalizes the yield of a tax-exempt municipal bond, $y_{t,m}^M$, with the after-tax yield of a taxable Treasury bond of the same maturity m

at date t as

$$(1 - \theta_{t,m}) \cdot y_{t,m}^T = y_{t,m}^M, \quad (1)$$

where $y_{t,m}^T$ denotes the before-tax Treasury yield. If there was no uncertainty and if taxes were constant over the maturity of the two bonds, then the break-even tax rate (BETR) would equal the marginal tax rate of the marginal investor, i.e., $\theta_{t,m} = \tau$. (Below I show how to interpret the BETR when allowing for non-constant expected tax rates and for unobserved risk factors.)

Figure 1 provides non-parametric evidence that municipal yield spreads indeed reflect expectations of future tax rates, especially at lower frequencies, by contrasting the 2- and the 15-year BETR with the marginal tax rate of the top 1% of the income distribution, taken from [Saez \(2004\)](#). The figure shows both the raw data as well as the trend component after applying a low-pass filter explained in more detail in the next section. The 2-year BETR follows the effective marginal tax rate of the top 1% closely, with the exception of the early 1990s, suggesting that the marginal investor is a household in the top of the income distribution. This finding is consistent with the fact that incentives to hold tax-exempt debt increase with the effective marginal tax rate. Importantly, movements in the 2-year BETR anticipate movements in the top tax rate. Moreover, short-term bonds are less affected by other risk factor such as credit risk or liquidity.

The 15-year BETR on the other hand, which averages expected future tax rates over a longer horizon, behaves differently. It sharply decreases during the early 1980s in anticipation of the Reagan tax cuts and stays relatively constant until the late 1990s when it starts to decline again in anticipation of the Bush tax cuts of the early 2000s. The fact that the time series of BETRs with different maturities do not move one-for-one suggests that the bond market forecasts both the timing of future income tax changes and the *expected path* or expected persistence of the tax changes well. The two quasi-natural experiments discussed in the next section confirm this conjecture. In particular, longer-term yield spreads respond strongly to changes in election probabilities and the response is quantitatively consistent with long-term differences between the proposed tax reforms. Hence, the lower level of the 15-year BETR in **Figure 1** does not reflect a disconnect between expected taxes and yields, but instead is consistent with a risk premium that is increasing in the bond maturity.

Figure 1 also shows that the 2-year and the 15-year break-even tax rates are generally below the top marginal tax rate, reflecting a risk premium for other factors than expected tax rates, which is more important for longer-term bonds than for short-term bonds. Hence, longer-term bonds demand a larger compensation for such risks than short-term bonds. The finding that the relative risk premium increases with the maturity of the yield spread is consistent with a large literature on the so-called “muni puzzle,” the observation that the slope of the municipal bond yield curve is almost always steeper than the slope of the Treasury yield curve. There is a large literature in finance that tries to explain this fact, but it seems fair to say that so far no single explanation of this puzzle has emerged.

1.3 Factors other than Expected Tax Rates Affecting Yield Spreads

Of course, there are other factors than expected tax rates that affect municipal yield spreads, especially at higher frequencies. I next discuss the three most important factors.

Credit risk: The first factor that comes to mind is probably default or credit risk. In order to eliminate, or at least minimize, the effect of credit risk on the yield spreads, I use only the highest quality municipal bonds available. In particular, the municipal yield curve used in this paper is based on AAA-rated general-obligation state bonds. State bonds are a subset of all municipal bonds and benefit from the same federal tax exemption. They are more liquid compared to other types of municipal bonds; see e.g., [Harris and Piwowar \(2004\)](#). General-obligation bonds are backed by the full faith and credit of the issuing state, similar to the backing of Treasury bonds. The last default of any state general-obligation bond occurred in 1933 and was only temporary ([Hempel \(1971\)](#)). Moreover, municipal bonds in general and general-obligation bonds in particular have a high recovery rate. For instance, [Fitch Ratings \(2007\)](#) assumes that general-obligation bonds recover 100% of par within one year of default. The prime-grade general-obligation municipal bonds used in this paper are therefore essentially free from default risk.

While municipal bond defaults have been rare historically—especially for state general-obligation bonds—it is nevertheless conceivable that such rare default events affect municipal bond prices. Using pre-refunded municipal bonds, which are escrowed at the Treasury Department and bear the same default risk as Treasury bonds, [Chalmers \(1998\)](#) shows that default risk cannot explain the observed yield spread for such bonds.³

State income taxes: State personal income taxes are another factor that might confound the relationship between the investor’s marginal *federal* tax rate and the municipal yield spreads. While interest on municipal bonds is generally exempt from federal income taxes and interest on Treasury bonds is exempt from state and local income taxes, nothing prevents states from taxing interest on municipal bonds. Many states exempt municipal bond interest from state and local income taxes, either for all or at least for in-state investors, and several states do not collect personal income taxes at all. Moreover, investors have strong incentives to avoid paying state taxes on municipal bonds, for instance by investing in municipal bonds of their state of residence ([Babina, Jotikasthira, Lundblad and Ramadorai \(2015\)](#)).

Using 10-year yields of four states,⁴ which span the spectrum of possible tax treatments of in-state and out-of-state municipal bond investors, I find that the maximum yield spread between

³ Credit default swaps (CDS) data for municipal bonds becomes available only in 2005 and remains illiquid.

⁴ Pennsylvania exempts both in-state and out-of-state municipal bond investors; Massachusetts exempts in-state investors from state taxes but taxes out-of-state investors; Illinois taxes both in-state and out-of-state investors; and Texas has no personal income tax. All state bonds are AAA-rated general obligations, with the exception of Illinois, for which there are no AAA general-obligation bonds available in the sample. Instead, I use AA-rated state bonds that are insured against default risk so that they are comparable to the other AAA-rated general-obligation bonds. The sample starts in 1996, the first year I have data on individual state bonds.

those state bonds is on average 7 basis points and hence an order of magnitude smaller than the spread to similar 10-year Treasury bonds, which is on average 132 basis points over the same period. This strongly suggests that state taxes are not an important determinant of municipal yield spreads relative to federal income tax rates.

Liquidity and other risk factors: The main remaining factors that can affect municipal yield spreads other than expected tax rates are therefore liquidity risk and potentially tax uncertainty. To minimize the effect of liquidity shocks on the yield spread I use off-the-run Treasury bonds, which are less liquid than on-the-run issues and are therefore more similar to municipal bonds, and I use state bonds, which are the most liquid municipal bonds. However, the off-the-run Treasury bond market is still much more liquid than the most liquid municipal bond market. To account for such risk factors, I next extend the traditional model of municipal yield spreads to include a latent risk factor and to allow for changes in the path of expected tax rates.

1.4 A Model of Break-Even Tax Rates (BETR)

Interest income from Treasury bonds is exempt from state and local taxes, but is subject to federal income taxes, while interest on municipal bonds is exempt from federal income taxes and from state income taxes for most investors. In order to interpret the yield data it is important to note that the relative municipal yield spread is different from the expected tax rate. Similarly, the yield spread between nominal and real Treasury bonds—the so-called break-even inflation rate—does not equal the expected rate of inflation. However, in both cases the yield spreads are related to the underlying expectations. To formalize this relationship I start with the definition of the par yield of a Treasury bond. Since Treasury bonds are taxed based on their imputed par yield, the par bond is the natural concept when analyzing the effects of taxes on bond prices, while zero coupon bonds are the starting point of most fixed-income models, which abstract from taxes; see [Kueng \(2012\)](#) for more detail on the tax treatment of bonds since 1970.

The yield $y_{t,m}^T$ on a Treasury bond maturing in m years and selling at par at date t is implicitly defined by⁵

$$1 = \sum_{s=1}^m \mathbb{E}_t[D_s(1 - \tau_s)y_{t,m}^T] + \mathbb{E}_t[D_m]. \quad (2)$$

D_s is the stochastic discount factor of after-tax income s years ahead. In order to satisfy equation (2), the Treasury par yield y^T needs to increase in response to an increase in expected future tax rates $\mathbb{E}_t\tau_s$, holding fixed the discount factor D .

⁵ To simplify notation I abstract here from the fact that coupon payments are semi-annual rather than annual, but I take this into account when analyzing the data. Moreover, whenever possible I use the first subscript—usually t —to denote calendar or “household time” and the second subscript—usually m or s —to denote the forecast horizon in years. For example, $y_{t,m}^T$ is the yield at date t (today) on a Treasury bond (T) that matures in m years. For bond yields, calendar time t is daily or weekly before or after 1983, respectively. “Household time” t in the CEX is quarterly such that $\Delta_t x_t$ is the quarterly change of x_t . However, since the CEX is a monthly rotating panel, the overall sampling frequency of the consumption data is monthly.

Similarly, the par yield $y_{t,m}^{\mathcal{M}}$ of a similar tax-exempt municipal bond is given by

$$1 = \sum_{s=1}^m \mathbb{E}_t[D_s(y_{t,m}^{\mathcal{M}} - \lambda_{s,m})] + \mathbb{E}_t[D_m], \quad (3)$$

where λ is a latent stochastic factor capturing such liquidity shocks and potentially any effects of tax uncertainty. To satisfy equation (3), the municipal par yield $y^{\mathcal{M}}$ has to increase to compensate a positive liquidity shock λ , holding fixed the discount factor D .⁶

The marginal investor is indifferent between investing one more dollar in a Treasury or a municipal bond with the same maturity. Let M be the longest maturity available. Solving (2) and (3) as a function of the relative municipal yield spread $y^{\mathcal{M}}/y^T$ results in⁷

$$\theta_{t,m} \equiv 1 - \frac{y_{t,m}^{\mathcal{M}}}{y_{t,m}^T} = \sum_{s=1}^m \frac{\mathbb{E}_t D_s}{\underbrace{\sum_{i=1}^m \mathbb{E}_t D_i}_{w_{t,s}^{(m)}}} \cdot \mathbb{E}_t \tau_s - \underbrace{\frac{\sum_{s=1}^m \mathbb{E}_t D_s \lambda_{s,m}}{y_{t,m}^T \sum_{i=1}^m \mathbb{E}_t D_i}}_{\Lambda_{t,m}^{\lambda}} + \underbrace{\frac{\sum_{s=1}^m \text{Cov}_t(D_s, \tau_s)}{\sum_{i=1}^m \mathbb{E}_t D_i}}_{\Lambda_{t,m}^{\tau}}, \quad (4)$$

which can be written in vector form as $\theta_{t,m} = w_t^{(m)'} \mathbb{E}_t \tau - \Lambda_t^{(m)}$. The sum of the liquidity premium Λ^{λ} and the tax risk premium Λ^{τ} is $\Lambda_t^{(m)} = \Lambda_{t,m}^{\lambda} - \Lambda_{t,m}^{\tau}$. The expected tax path over the horizon M is given by the vector $\mathbb{E}_t \tau = (\mathbb{E}_t \tau_1 \dots \mathbb{E}_t \tau_M)'$. $w_t^{(m)} = (w_{t,1}^{(m)} \dots w_{t,m}^{(m)} 0 \dots 0)'$ is the vector of annuity weights such that $w_t^{(m)'} \mathbb{E}_t \tau = \sum_{s=1}^m w_{t,s}^{(m)} \mathbb{E}_t \tau_s$ is the annuity value of the path of expected tax rates over the maturity m of the two bonds.⁸ Equation (4) reveals that the BETR is in general a weighted average of expected future tax rates over the maturity of the bonds less a premium Λ .

Since the market for Treasuries is more liquid than the most liquid municipal bond market, and because liquidity demand is high in bad times, the liquidity premium Λ^{λ} is likely non-negative on average. Moreover, marginal income tax rates are low in bad times because of the progressivity of the income tax and the possibility of countercyclical tax policies. After an extensive analysis of the narratives surrounding all major post-war tax changes, [Romer and Romer \(2010\)](#) conclude that all income tax changes from 1980 to 2001—with one minor exception in 2001—are not countercyclical policies or spending related but motivated by concerns about the long-run growth rate or the federal debt. Hence, the tax risk premium Λ^{τ} is likely primarily due to the progressivity of the income

⁶ I add the latent shocks in a linear way to obtain an analytical expression that is linear in both the path of expected tax rates as well as the risk premium; see equation (4) below. Adding the shock in a multiplicative way does not change the conclusions. In an ideal setting we would have two identical bonds with the exception that one is taxable and the other is tax-exempt. $y^{\mathcal{M}} - \lambda$ is the risk-adjusted municipal yield that proxies for such an ideal but unobserved tax-exempt Treasury bond.

⁷ Equating (2) and (3) implicitly assumes that both bonds are held to maturity, thereby abstracting from the timing of capital gains and losses. However, unexpected capital gains on both municipal and Treasury bonds are taxable at the capital gains tax rate. Therefore, the yield spread between the two bonds cancels out any first-order effects of capital gains taxes.

⁸ When I calculate the weights in the empirical section below I take into account that coupon payments are semi-annual and use $\mathbb{E}_t[D_s] = (1 + y_{t,s}^{\mathcal{M}}/2)^{-2s}$. In the absence of discounting, the first m elements of $w_t^{(m)}$ are equal to $1/m$. With discounting, the weights are generally decreasing in m such that $w_{t,m}^{(m)} < 1/m$. If the tax-exempt yield curve steepens, then future income is discounted more heavily, leading the weights on future tax rates to decrease.

tax over the period 1977-2001. The progressivity induces an insurance mechanism by paying larger after-tax interest in bad times and lower after-tax income in good times. The tax premium is therefore likely non-positive. In this paper I do not attempt to separate these two risk factors, and I refer to them jointly as the (liquidity) risk premium.⁹

Stacking equation (4) for the entire term structure of length M I obtain a system of equations that provides a mapping between the M break-even tax rates θ_t and the underlying path of expected forward tax rates $\mathbb{E}_t\tau$ over the forecasting horizon of 1 to M years at any point in time t ,

$$\theta_t = W_t \mathbb{E}_t\tau - \Lambda_t. \quad (5)$$

W_t is the M -by- M lower triangular annuity weighting matrix $[w_t^{(1)} \dots w_t^{(M)}]'$ and the vector of risk premia is given by $\Lambda_t = (\Lambda_t^{(1)} \dots \Lambda_t^{(M)})'$.

Identity and stability of the marginal investor: In order to recover the underlying path of expected tax rates $\mathbb{E}_t\tau$, one needs to know the marginal tax rate of the marginal investor and correct for the risk premium Λ_t . Moreover, as mentioned above, for the analysis of the response of current household consumption to expected future tax rates in the next section it is important to identify the entire path of expected tax rates $\mathbb{E}_t\tau$ from the term structure of break-even tax rates θ_t . Two tax reforms that affect the expected after-tax lifetime income by the same amount should have the same effect on current consumption independent of the timing of the tax changes, abstracting from credit constraints and precautionary saving.

Using Flow of Funds data, [Ang, Bhansali and Xing \(2007\)](#) show that starting in the 1970s, households hold the vast majority of municipal bonds, both directly and through mutual funds, suggesting that a household is the marginal investor and hence that expected personal income tax rates are reflected in the relative bond prices. However, we need to know which households own municipal bonds in order to determine the marginal tax rate identified by the relative bond spread. Equally important, this marginal investor has to be stable in the income distribution in order for changes in break-even tax rates to reflect changes in expected tax rates rather than changes in the marginal investor holding fixed expected tax rates and any risk premium.

Since equations (2) and (3) are first-order conditions of the marginal investor's portfolio choice problem, they should apply to all households holding both types of bonds. To analyze this claim I use data from the Survey of Consumer Finances (SCF), which is described in more detail in the online appendix. I define the marginal tax rate of the marginal investor as the asset-weighted average of the effective marginal tax rate over all households that own both taxable and tax-exempt bonds.

[Figure 2](#) compares the estimates of the marginal investor's marginal tax rate with the marginal tax

⁹ To quantify Λ^τ I estimate the following population moments: $\min_s \{\text{Cov}(D_s, \tau)\}$, $\max_s \{\text{Cov}(D_s, \tau)\}$, and $\sum_s \mathbb{E}_t D_s$. The estimates are -0.0013 , 0.00128 , and 13.80 , the latter with a standard deviation of 2.02 . Since Λ^τ is only of order $1/1000$, this calculation suggests that the tax risk premium is non-positive and negligible. However, this is only suggestive since I use the current yields to calculate D_s . For a recent study that separately estimates the liquidity and tax uncertainty premium, see [Longstaff \(2011\)](#).

rates of different percentiles of the income distribution taken from [Saez \(2004\)](#). The marginal tax rate of the marginal investor identified in the SCF is close to the tax rate of the top 1%. Since the top two tax brackets move very closely during the sample period it is not important whether the marginal investor is in fact in the top bracket or one bracket below that. The identification of the household spending response to tax news shocks in the next section relies on *changes* in the path of expected tax rates, not the level.

2 Identifying the Path of Expected Tax Rates

In this section I first formally estimate the degree of fiscal foresight contained in municipal yield spreads using two quasi-natural experiments. I then use the model of the term structure of municipal yield spreads derived in the previous section to interpret the reduced-form coefficients obtained from the experiments. Finally, having validated the model using the quasi-natural experiments, I use it to measure expected tax rates over the entire sample period in order to estimate the household spending response to future tax changes in the next section.

2.1 Two Presidential Elections as Natural Experiments

The presidential elections of 1992 and 2000 are close to ideal quasi-natural experiments for testing the model. During both elections the nominees from the Democratic and the Republican Party campaigned on very different proposals concerning the top income tax rates. Furthermore, these tax proposals received extensive coverage by the media and featured prominently in both the primary and presidential debates; see [Seib and Murray \(1992\)](#) and [Calmes \(2000\)](#) for a comparison of the campaign proposals.

In 1992 Bill Clinton proposed to increase the top tax rate by 10% to deal with the high level of government debt. His victory ultimately led to the Omnibus Budget Reconciliation Act (OBRA 1993), which increased the top rate by 8.6% retroactively back to January 1, 1993. Importantly, OBRA 1993 left the dividend and the long-term capital gains tax rates unchanged.¹⁰ President George H.W. Bush, having broken his tax pledge from the 1988 election campaign, promised not to raise any taxes.

Similarly, during the presidential election of 2000 George W. Bush proposed to cut taxes across the board—including the top rate—by about 5%, using the budget surplus that accumulated under President Clinton. Incumbent Vice President Al Gore proposed tax breaks for low and middle income taxpayers while leaving the top rates unchanged. Bush's victory in 2000 ultimately led to the Economic Growth and Tax Relief Reconciliation Act of 2001 (EGTRRA) and the Jobs and Growth Tax Relief Reconciliation Act of 2003 (JGTRRA) which lowered the top income tax rate by 4.6% over three years. Importantly, EGTRRA 2001 leaves the top corporate income, capital

¹⁰ OBRA 1993 also increased the top corporate tax rates, but only by 1% from 34% to 35%. George H.W. Bush also proposed to cut long-term capital gains tax rates from 31% to 15.4%; see [Seib and Murray \(1992\)](#). However, the estimates discussed below are inconsistent with both of those tax rate changes and instead strongly support top personal income tax rates as the main driver of yield spread changes.

gains, and dividend tax rates unchanged.¹¹ Hence, the combination of both elections also tests whether it is indeed changes in the top personal income tax rate that drives changes in municipal yield spreads during periods of fiscal news, rather than other tax rates such as capital gains taxes, dividend income taxes, corporate income taxes, or personal income tax rates in lower tax brackets.

I use data from the Iowa Electronic Markets (IEM), a political prediction market, which provides daily prices of a winner-takes-all contract during the last few months of the presidential races of 1992 and 2000. Those contracts pay \$1 if the specific candidate wins the popular vote and \$0 otherwise. In the following I abstract from the difference between winning the popular vote and winning the election. I checked the robustness of the results for the 2000 election to such differences by dropping the last week of daily election data when such considerations became relevant, and I obtained very similar results. Furthermore, since bets are limited to \$500, market participants cannot use the prediction markets to hedge their income tax risk. Therefore, changes in the prices of such contracts can be interpreted as measuring daily changes in the election probability of the presidential candidates; see [Snowberg, Wolfers and Zitzewitz \(2011\)](#) for a more extensive discussion of the use of prediction markets for economic inference.

Let p_t be the price of a contract that pays \$1 if George W. Bush wins the popular vote in 2000 and 0 otherwise. In the following I use the presidential election of 2000 to explain the analysis, but the same applies for the election of 1992 substituting “Clinton” for “Bush.” Changes in these contracts represent an alternative measure of changes in expected tax rates due to the difference in the candidates’ tax proposals under the assumption that other factors, such as liquidity shocks in the betting and bond markets, are uncorrelated conditional on observables. To address concerns that shocks unrelated to tax expectations are correlated between the betting and bond markets, I include additional variables z_t that control for such factors like default and liquidity risk. I therefore regress the BETR of each maturity $\theta_{m,t}$ on the election probability p_t , while controlling for other factors z_t as well as maturity fixed effects α_m ,¹²

$$\theta_{t,m} = \beta_m \cdot p_t + \alpha_m + \gamma'_m z_t + \varepsilon_{t,m} . \quad (6)$$

[Table 1](#) reports the estimated response $\hat{\beta}_m$ of the BETRs to changes in the election probability in 2000 and 1992 for the eight most commonly traded maturities. Most coefficients are statistically significantly different from zero and have the expected sign, positive for Bill Clinton’s and negative for George W. Bush’s betting contract. [Figure 3\(a\)](#) and [Figure 3\(b\)](#) show the response of the entire term-structure of the 30 break-even tax rates during the two elections. Importantly, we see that the *entire* term-structure of BETRs responds to changes in the election probability. In fact, the

¹¹ Later in his first term, President George W. Bush lowered the dividend tax rate and the long-term capital gains tax rate to 15%, 5%, and 0% (JGTRRA 2003). However, these cuts were not part of his campaign platform; see [Calmes \(2000\)](#).

¹² The controls include measures of (il)liquidity such as the number of repeated prices, muni supply (negotiated and competitive offers), on-the-run/off-the-run Treasury spread, prediction market volume and units traded; measures of credit risk such as the pre-refunded spread, spreads between lower-grade municipal bonds; and the weekly standard deviation of long-term treasuries to control for the effect of call options.

response of long-term rates is typically more precisely estimated than the response at the short end. This shows that the lower *level* of long-term BETRs shown in [Figure 1](#) does not reflect a disconnect between expected taxes and yields, but instead is consistent with the presence of a risk premium that is increasing in the maturity. The results here show that *changes* in BETRs strongly respond to changes in expected tax rates at all maturities as predicted by the theory.

To interpret the magnitudes of the estimated coefficients, note that the contracts pay 100 cents if the candidate wins the popular vote and zero otherwise. Therefore, an increase of the price by 1 cent corresponds to a 1% increase in the perceived probability of the candidate winning. Multiplying the coefficients by 100 yields the predicted change in the BETRs if George W. Bush (Bill Clinton) wins in 2000 (1992) relative to the counter-factual that Al Gore (George H.W. Bush) or the third-party candidate wins.

Interpretation of the regression coefficients: The sign of the coefficients reported in [Table 1](#) and shown [Figure 3\(a\)](#) and [Figure 3\(b\)](#) is consistent with the difference between the tax proposals of the two major presidential candidates in both elections. In order to interpret the magnitudes of the coefficients, it is important to remember that the break-even tax rate responses do not reflect changes in expected tax rates, but instead are weighted averages of differences in the expected path of tax rates under different presidents. However, we can use the model of the term structure of municipal yield spreads derived in the previous section to give the reduced-form coefficients a structural interpretation.

Let $\Pr_t(\text{Bush})$ denote the probability of Bush winning the election conditional on all information available at time t . I assume that the price corresponds to the rational conditional probability measure, i.e., $p_t = \Pr_t(\text{Bush})$. Using the law of iterated expectations, I decompose the conditional expectation of the path of future tax rates $\mathbb{E}_t\tau$ as $\mathbb{E}_t\tau = (\mathbb{E}_t[\tau|\text{Bush}] - \mathbb{E}_t[\tau|\text{Gore}]) \cdot p_t + \mathbb{E}_t[\tau|\text{Gore}]$. Substituting in [\(5\)](#) yields a structural interpretation of the system of the 30 regression equations

$$\theta_t = W_t(\mathbb{E}_t[\tau|\text{Bush}] - \mathbb{E}_t[\tau|\text{Gore}]) \cdot p_t + (W_t\mathbb{E}_t[\tau|\text{Gore}] - \Lambda_t), \quad (7)$$

and in particular of the regression coefficient of interest, $\beta = \mathbb{E}[W_t] (\mathbb{E}_t[\tau|\text{Bush}] - \mathbb{E}_t[\tau|\text{Gore}])$.¹³ This shows that the vector of population regressions β represents the annuity values of the difference in the paths of expected tax rates ($\mathbb{E}_t[\tau|\text{Bush}] - \mathbb{E}_t[\tau|\text{Gore}]$) between a world in which Bush wins the election in 2000 and the counter-factual world in which Gore or the third-party candidate wins.

Letting $\tau_t^{pf} = (\tau_{t,1} \dots \tau_{t,M})'$ denote the perfect-foresight path of realized tax rates at date t over the horizon of 1 to $M = 30$ years, I calculate the hypothetical regression coefficients one would

¹³ $\mathbb{E}[\cdot]$ without a subscript is defined as the average of the conditional expectations over the election sample, i.e. $\mathbb{E}[x] \equiv \frac{1}{T} \sum_{t=1}^T \mathbb{E}_t[x]$. I searched the archives of *The New York Times* and *The Wall Street Journal* for articles that would indicate a change in the tax proposal of the candidates during the sample period but did not find any. Hence I assume that the relative difference between the tax proposals ($\mathbb{E}_t[\tau|\text{Bush}] - \mathbb{E}_t[\tau|\text{Gore}]$) remains constant during the final months of the election, i.e. $\mathbb{E}_t[\tau|\text{Bush}] - \mathbb{E}_t[\tau|\text{Gore}] = \mathbb{E}[\tau|\text{Bush}] - \mathbb{E}[\tau|\text{Gore}] \forall t$. Otherwise, β identifies the average value of the relative difference between the two proposals during the final months of the presidential election, i.e. $\beta = \mathbb{E}[W_t(\mathbb{E}_t[\tau|\text{Bush}] - \mathbb{E}_t[\tau|\text{Gore}])]$.

obtain under perfect foresight, $\beta^{pf} = \mathbb{E}[W_t](\tau_t^{pf} - \tau_t)$. Here I assume that the level of the unobserved counter-factual tax path— $\mathbb{E}_t[\tau|\text{Gore}]$ in 2000 and $\mathbb{E}_t[\tau|\text{H.Bush}]$ in 1992 which is not identified by the regression—is the status quo tax rate during the election year, i.e. $\tau_t = 39.6\%$ in 2000 and $\tau_t = 31\%$ in 1992, respectively. Moreover, I show two scenarios for the tax path of future tax rates beyond 2011, one in which the Bush tax cuts expire in 2011 as scheduled and one in which they become permanent.

Figure 3(a) and **Figure 3(b)** show the term-structure of the perfect-foresight coefficients, β^{pf} . While the regressions do not impose any restrictions on the sign, size, or the shape of the estimated response $\hat{\beta}$, the regression coefficients correspond closely to the perfect-foresight coefficients derived from the model. Hence, the model does a good job at relating current changes in break-even tax rates to the subsequent legislated tax changes.

Inferring the entire path of expected tax rates: Ultimately we are interested in the inverse mapping of equations (5) and (7), i.e., $\mathbb{E}[\tau|\text{Bush}] - \mathbb{E}[\tau|\text{Gore}]$ as a function of β , respectively $\mathbb{E}_t\tau$ as a function of θ_t . These market-based expected tax rates can be interpreted as *forward tax rates*, similar to forward interest rates derived from the term structure of spot rates. Recall that W_t is a lower triangular annuity matrix with its last column vector given by $(0 \dots 0 w_{t,M}^{(M)})'$. In the data, $w_{t,M}^{(M)}$ has a mean of 0.01 with a standard deviation of 0.003 and a minimum of 0.003, so W_t can be close to singular. Inverting this matrix makes the solution sensitive to small perturbations of β or θ_t that are unrelated to tax news. Instead of a direct inverse I use a robust inverse of W_t , known as a first-order ridge regression in the statistics literature.¹⁴ I impose that the expected tax path is a smooth function across maturities $m = 1, \dots, 30$, since it is implausible that the expected tax rate e.g. in 20 years is very different from the expected tax rate in 19 or 21 years. The robust inverse penalizes such non-smooth solutions with a factor μ , called the regularization parameter. The online appendix shows that the parameter μ only significantly affects long-run expectations and discusses how to optimally choose μ .

Figure 3(c) shows the path of expected tax rates during the presidential election of 2000 obtained by inverting the regression coefficients $\hat{\beta}$. The top tax rate is expected to decrease to 35% by the year 2002 and to return quickly back to the initial level of 39.6%. Moreover, the bond markets expect the initial tax cuts to be off-set by later tax increases above the initial level of 39.6%. One interpretation is that the bond markets expect the tax cuts to be unsustainable. Compared to the perfect-foresight tax path, the path of the expected tax rates returns quickly back to rates around 40%. The expected tax rate starts to increase sharply after four years, consistent for example with expectations that President George W. Bush will serve only one term. Turning to the presidential election of 1992, **Figure 3(d)** graphs the path of expected tax rates against the perfect-foresight path. The bond markets correctly anticipate the new level of the top tax rate induced by the Clinton tax

¹⁴ The word *regression* can be misleading in this context since I do not perform statistical inference in the traditional sense of projecting a vector from a larger onto a smaller space. Instead, the first-order ridge “regression” calculates M forward tax rates $\mathbb{E}_t\tau$ from M observed break-even tax rates θ_t . The constraint on the first-derivative of the solution is matched by the additional regularization penalty parameter μ . See the online appendix for more details.

increase in 1993. The path of expected tax rates slightly underestimates the duration of the Clinton tax increase. The path also shows that the bond markets in 1992 expect the tax rates to return back to the initial level in the long run. However, the tax cuts enacted under President W. Bush “only” lowered the top rate to 35% instead of 31%, the level in 1992, or 33%, George W. Bush’s initial campaign proposal.

The two natural experiments show that the model of the BETRs given by equation (5) explains changes in the relative municipal yield spread well in times of large fiscal news. They show that the expected tax rates that underlie the BETRs forecast future top personal income tax rates surprisingly well. The experiments also highlight the necessity of imposing some restrictions on the solution to the inverse problem in order to obtain a smooth and hence reasonable path of expected forward tax rates.

2.2 Deriving Expected Tax Rates from Break-Even Tax Rates

Unfortunately, the additional source of variation provided by the election probabilities is not available for the entire sample period from 1977 to 2001. Instead, I impose two identifying assumptions to recover $\mathbb{E}_t\tau$ from θ_t , and I check the robustness of the household spending response to those assumptions extensively in section 4. The first assumption requires that the bond spreads reflect rational-expectations, i.e., that the BETR forecast error is unbiased,

$$\underbrace{\mathbb{E}[W_t(\mathbb{E}_t\tau - \tau_t^{pf})]}_{\text{BETR forecast error}} = 0. \quad (8)$$

The two natural experiments show that the marginal tax rate of the top 1% of the income distribution is the tax parameter that determines the municipal yield spreads (in the absence of any shocks to the relative risk premium). Rearranging (5) and imposing (8) yields a measure of the average risk premium, $\mathbb{E}[\Lambda_t] = \mathbb{E}[W_t\tau_t^{pf} - \theta_t]$. Figure A1 in the online appendix shows the average risk premium as a function of the maturity m , estimated globally over the entire sample from 1977 to 2001. Consistent with the muni puzzle, the average risk premium is monotonically increasing in the maturity of the BETR. Adjusting for this average risk premium changes the relative *level* of the BETRs, but does not affect changes and hence the estimated response of household spending in the next section.

The second identification assumption deals with temporary shocks to the risk premium. I assume that households and the marginal investors form tax expectations independently of the municipal yield spread. For instance, households read newspapers or follow political campaigns and use all these sources of information to form expectations about future tax rates. The econometrician does not directly observe those news sources, but can infer the aggregate information set by looking at municipal yield spreads and interpret the data through the lens of the BETR model. For instance, suppose the break-even tax rates decrease at date t but immediately rebound the next day, at $t+1$. The econometrician can use this fact to estimate the tax expectations at date t . He will conclude that this change in break-even tax rates was most likely due to say a liquidity shock instead of

tax news. If he uses only past and current prices he will underestimate the rational information set. This way of modeling tax news implies that the econometrician wants to use all prices—past, current, and future—to infer the path of expected tax rates $\mathbb{E}_t\tau$ at any point in time.

Filtering the tax news shocks from the “noise” shocks—i.e., the risk premium shocks—is important since the tax news shocks form the regressor in the consumption analysis in the next section. Liquidity shocks and other shocks that are unrelated to tax news introduce noise and therefore potential attenuation bias of the household spending response coefficient, a fact I demonstrate in the robustness [section 4](#). Not dealing with this attenuation bias would lead to the conclusion that households do not respond to news even if they in fact are forward-looking. To obtain a more precise measure of the expected tax rates I use the two-sided low-pass filter proposed by [Baxter and King \(1999\)](#) that passes all frequencies below two years. While the filter may remove some tax news shocks in addition to risk premium shocks, it reduces potential attenuation bias in the analysis of household consumption. Importantly, when estimating the household spending response to tax news, I always include both the filtered news shock and the residual in the regression. Hence, if the residual did contain information that is relevant to consumers, then we would expect it to have a significant effect on household spending. The two-year low pass filter is motivated both by visual inspection of the time series and by the fact that two years is the shortest period between two income tax reforms in the sample (OBRA 1990 and OBRA 1993). Below I therefore also perform extensive robustness checks of the household spending response to different thresholds as well as different filters (see [Figure 6](#)).

2.3 The Bond Markets Anticipation of the Reagan Tax Cuts

By far the largest change in top marginal rates during the sample period is the first Reagan tax cut, the Economic Recovery Tax Act of 1981, which was part of President Regan’s election platform. Since there does not exist a political prediction market for the election in 1980, I use the methodology derived in the previous section to extract the path of expected tax rates from the municipal yield spreads. Note that while [Figure 1](#) already suggests that the Reagan tax cuts were well anticipated, this is only a conjecture since the time series shown in [Figure 1](#) are break-even tax rates θ_t and not forward tax rates $\mathbb{E}_t\tau$.

[Figure 4](#) shows the path of expected tax rates $\mathbb{E}_t\tau$ at the beginning of each year against the perfect-foresight tax path τ_t^{pf} for each month of January from 1977 to 1982.¹⁵ These paths of expected tax rates are obtained by inverting the break-even tax rates of all available maturities. They confirm the conjecture that the Reagan tax cuts were anticipated well before they were implemented. The sequence of forward tax rates shows that taxes were expected to remain high during Jimmy Carter’s presidency (1977-1980) and to even increase over the foreseeable future.

¹⁵ The web appendix of this paper (www.lorenzkueng.com) contains an animation of the evolution of $\mathbb{E}_t[\tau]$ from January 1977 to August 1982 that shows monthly changes in the path of expected tax rates over a 15-year forecasting horizon. The forecast horizon for this period is 15 years because Treasury yields are not available at longer maturities before 1983.

In fact, the upward sloping path of expected tax rates reflects proposals during the late 1970s by President Carter to raise top income tax rates; see [Poterba \(1986\)](#). The increase is also consistent with expected bracket creep due to rampant inflation during this time when tax brackets were not yet indexed to inflation. Since the marginal tax rate of the marginal investor was well below the top statutory rate (see [Figure 2](#)), expected increases in nominal income would push high-income investors in even higher brackets.

The long-run expectations decreased sharply during the presidential election of 1980 as it became increasingly clear that Ronald Reagan would become the next president and would implement the promised tax cuts. Between 1980 and 1982 as Reagan passed his first tax cut the bond market also started to anticipate the second tax reform, the Tax Reform Act of 1986 (TRA 1986). This figure reveals an astonishing degree of fiscal foresight contained in the municipal yield spreads. In the next section I use these changes in the paths of expected tax rates to estimate the household spending response to tax news as a first case study.

The results in this section show that the time series of market-based tax expectations imply a considerable amount of fiscal foresight. Moreover, the path of expected tax rates $\mathbb{E}_t\tau$ derived from municipal yield spreads does a good job of recovering the underlying rational tax expectations. However, while the wealthy households that invest in municipal bonds have a high degree of fiscal foresight, their expectations may not be representative of (wealthy) consumers as a whole. In the next section I quantify the degree of fiscal foresight of higher-income households by estimating the household consumption response to tax news in order to learn more about consumer behavior and the transmission of tax news shocks in the real economy.

3 Household Spending Response to Tax News

To the best of my knowledge, this study is the first to test the standard life-cycle consumption theory by estimating the spending response to changes in expected life-time resources at the household level before these income changes have been realized. It therefore fills an important gap in the large empirical literature that tests the life-cycle model and its various extensions. Previous studies in this literature have mostly focused on anticipated income shocks, for which the frictionless model predicts no consumption responses (so called excess sensitivity tests), or they estimate the response to unanticipated income shocks for which the model's predicted response depends on the expected persistence of the shock; see [Jappelli and Pistaferri \(2010\)](#). This paper is more closely related to the latter as the tax news shock extracts the persistent component of a tax reform. However, an important distinction of this study is that it estimates the response to a shock that is only expected to be realized in the future and hence represents the present value of future anticipated tax changes.

3.1 Household-Specific Tax News Shocks

The time series of tax expectations described above affects households differently depending on how much of their future taxable income is affected by the expected tax changes. In an ideal

setting I would observe news shocks for each tax rate in each tax bracket as well as news about changes of those tax brackets. Expected tax liabilities for each household could then be computed by integrating over expected brackets and rates. In practice I only observe a time series of paths of expected top tax rates. To take into account the cross-sectional heterogeneity in the exposure of different households to the tax news shock I adopt an approach that combines a research design similar to [Bartik \(1991\)](#) with the traditional Euler equation approach pioneered by [Hall \(1978\)](#). Specifically, to estimate the effect β of changes in expected tax rates, summarized by $\Delta\tau_t^E$, on changes in nondurable and service consumption of household i , Δc_{it} , we would like to estimate a regression equation of the following form,¹⁶

$$\Delta c_{it} = \beta \cdot \omega_i \times \Delta\tau_t^E + \alpha_t + \varepsilon_{it}, \quad (9)$$

where ω_i is a predetermined measure of the household’s exposure to the tax news shock, and α_t is a full set of time fixed effects (i.e., year-by-month dummies).¹⁷

A theory-based measure of household-level tax news: In the case of tax news shocks, we have to deal with an entire vector of time series of expected future tax rates and we have to find household-specific “importance” weights ω_i that map this vector into a scalar household-specific tax news shock. I use the standard life-cycle model to guide the analysis and I then perform extensive robustness checks of the results to the different assumptions that I have to make in order to derive these household-specific tax news shocks.

To be as transparent as possible and in order to obtain an closed-form solution for the weights ω_i I use the certainty-equivalent version of the standard life-cycle model with quadratic preferences, but a similar interpretation can be derived by log-linearizing the inter-temporal budget constraint; see [Campbell and Deaton \(1989\)](#). This simplified model abstracts from three main effects. First, it does not capture the effect of marginal tax rates on consumption via the inter-temporal substitution channel triggered by changes in the effective after-tax interest rate. However, as mentioned above, all regressions include a full set of time fixed effects, which absorb such effects to a first order. Second, certainty equivalence ignores the effect of tax uncertainty on household consumption through changes in precautionary savings. Third, the model ignores potential effects of changes in expected marginal tax rates on current and future labor supply or any other indirect effects of marginal tax rates on pre-tax income such as wages or growth effects in general equilibrium. It also abstracts from non-separability of consumption and leisure. The regression coefficient β can therefore also be interpreted as a reduced form of a more general model that would explicitly model—and more importantly also separately identify—those channels, which is beyond the scope of this paper.

With quadratic preferences, a constant interest rate r equal to the rate of time preference, and household wealth evolving as $A_t = (1+r)(A_{t-1} + Y_{t-1} - C_{t-1})$, consumption changes can be written

¹⁶ Note that the Euler equation approach already differences out household fixed effects.

¹⁷ In the typical application of the Bartik approach, ω_i would be something like the lagged and hence predetermined industry shares of a local economy.

in closed form:

$$C_{it} - C_{i,t-1} = \sum_{s=1}^{H-t} w_s (\mathbb{E}_t - \mathbb{E}_{t-1})(Y_{i,t+s} - T_{i,t+s}).$$

Y is income before taxes T , and the annuity weights are given by $w_s = \frac{r}{1+r} [1 - \frac{1}{(1+r)^{H-t+1}}]^{-1} (1+r)^{-s}$. Following [Campbell and Deaton \(1989\)](#), I normalize consumption changes by (taxable) income, thereby expressing tax liabilities as average tax rates and making the model scale independent, a particularly useful feature when working with micro data. Specifically, I divide by predetermined adjusted gross income Y_i^{lagged} asked in the first interview, thereby avoiding the endogeneity that would arise if we used current or future income, i.e., income asked in the last interview.¹⁸ The robustness section below shows similar results when using lagged consumption (appropriately rescaled) as a normalization of consumption changes.

To bring the model more closely to the data, I replace the constant annuity weights with the weights based on real after-tax discount factors used in the previous section, and I limit the household's planning horizon H to the maximum available maturity M (usually 30 years),

$$\Delta c_{it} \equiv \frac{C_{it} - C_{i,t-1}}{Y_i^{lagged}} \approx - \sum_{s=1}^M w_{t,s}^{(M)} (\mathbb{E}_t - \mathbb{E}_{t-1}) \bar{\tau}_{i,t+s} + \varepsilon_{it}. \quad (10)$$

This equation shows that to a first approximation, changes in expected *average* tax rates determine the growth rate of consumption. Moreover, the tax news shock has a natural interpretation: It is the change in the expected annuity value of average tax liabilities, which is the permanent component of expected future tax changes. The error term ε contains measurement error in consumption growth as well as economic shocks such as news about future income growth and any other shocks that the model ignores and for which I proxy with observable covariates in the empirical analysis.

From expected top marginal rates to expected average rates: To go from the observed expected top marginal tax rates to the unobserved expected average tax rates I have to make two assumptions. First, I assume that changes in the tax base—if they do occur—are perfectly foreseen. With the exception of the Tax Reform Act of 1986 (TRA), which I discuss in more detail in the robustness section below and in the online appendix, this assumption is reasonable for the income tax reforms in my sample because the brackets did not change much. In the robustness section I show that excluding TRA from the sample by focusing on the later period starting in the mid-1990s yields similar results. Second, I scale the perfect-foresight tax rate $\tau_{t+s}(b)$ in each lower income bracket $b < B$ (where B denotes the top tax bracket) by a measure of foresight implied by current expectations of future top tax rates. This measure is the ratio of the market-based expected top tax rate $\mathbb{E}_t \tau_{t+s} \equiv \mathbb{E}_t \tau_{t+s}(B)$ to the perfect-foresight top tax rate $\tau_{t+s}(B) \equiv \tau_{t+s}$, where the latter is

¹⁸ Note that the survey asks households only in the first and last interview about their *annual* income, while expenditures are asked in each quarterly interview for the previous three months. Hence, income asked in the first interview is fully predetermined relative to all quarterly consumption changes a household reports in subsequent interviews.

taken from Saez (2004). We can therefore express the unobserved expected lower-bracket tax rates in terms of observed quantities,¹⁹

$$\mathbb{E}_t \tau_{t+s}(b) = \tau_{t+s}(b) \times \frac{\mathbb{E}_t \tau_{t+s}(B)}{\tau_{t+s}(B)}. \quad (11)$$

This means that if households have perfect foresight about the top marginal tax rate (i.e., if they correctly guess the future tax rate such that the second term is one), they also have perfect foresight about the lower bracket rates. Since the empirical analysis restricts the sample to higher-income households, defined as the top quartile of the income distribution for which policy-induced changes in average tax rates are highly correlated with changes in the top marginal tax rate, this assumption does not affect the results much. In the robustness section I analyze the sensitivity of the results to this assumption. First, I obtain essentially the same results when using only the smaller subset of households in the top 10% percentile who are not much affected by assumptions about expected rates in lower tax brackets, because most of their income falls into the higher brackets. Second, I show that if we instead fix the lower bracket rates at their current predetermined level based on the current tax code (i.e., if we use $\tau_t(b)$ instead of $\tau_{t+s}(b)$) I obtain very similar results. This essentially implements the Bartik approach without using the life-cycle model to derive the correct cross-sectional weights.

With these assumptions, changes in household-specific expected average tax rates can be written as

$$(\mathbb{E}_t - \mathbb{E}_{t-1}) \bar{\tau}_{i,t+s} = \sum_b \frac{B_{i,t+s}(b) \tau_{t+s}(b)}{Y_{i,t+s}} \cdot \frac{(\mathbb{E}_t - \mathbb{E}_{t-1}) \tau_{t+s}}{\tau_{t+s}} = \frac{\bar{\tau}_{i,t+s}}{\tau_{t+s}} (\mathbb{E}_t \tau_{t+s} - \mathbb{E}_{t-1} \tau_{t+s}),$$

where $B_{i,t+s}(b)$ is the income household i receives in tax bracket b . The expected change in the top tax rate $(\mathbb{E}_t - \mathbb{E}_{t-1}) \tau_{t+s}$, which is identified using the municipal yield spreads, can now be interpreted as the signal that the household receives between date $t - 1$ and t . The term $\frac{\bar{\tau}_{i,t+s}}{\tau_{t+s}}$ is a measure of the relevance of the signal for the household's consumption decision. This corresponds closely to the "importance weight" ω_i in the Bartik approach described above: If this ratio is low then the impact of news about the top tax rate in s years has only a small impact on the household's expected after-tax lifetime income, and a rational household should therefore largely ignore the signal. On the other hand, if the ratio is large, then the household should pay close attention to the signal.

It is important to note that this assumption does *not* imply that the expected change in the *average* tax rate is the same for all households. To see this, suppose that the expected future tax schedule in s years from now has only two tax rates, 10% and 50%. Let the first tax bracket range from \$0 to \$10,000 so that all income above \$10,000, which is the second income bracket, is expected to be taxed at the 50% rate. Suppose that the expected top tax rate increases by 10%, i.e. $\frac{\Delta_t \mathbb{E}_{t+1} \tau_{t+s}(B)}{\tau_{t+s}(B)} = 0.1$ such that the lower tax rate increases by 1 percentage point from 10% to 11%

¹⁹ Since the sampling frequency is quarterly at the household-level while expectations are formed over annual intervals, I assume that the perfect-foresight variables do not change between quarters, so that $(\mathbb{E}_{t+1} - \mathbb{E}_t) \tau_{t+s}(b) = \tau_{t+s}(b) \times \frac{(\mathbb{E}_{t+1} - \mathbb{E}_t) \tau_{t+s}}{\tau_{t+s}}$.

and the top tax rate by 5 percentage points from 50% to 55%. The expected average tax rate of a household with an income of \$10,000 increases by 1 percentage point, while the expected average tax rate of a household with an income of \$15,000 increase by $2^{1/3}$ percentage points. Moreover, the expected change of the average tax rate approaches 5 percentage points as income goes to infinity, which equals the expected change of the top tax rate. Figure A2 in the online appendix shows the cross-sectional treatment heterogeneity for each tax reform in the sample. Using the NBER TAXSIM calculator, I compute perfect-foresight average tax rates $\{\bar{\tau}_{i,t+s}\}_{s=0}^M$ for each household i in the CEX that depend on the head of household’s age and the household’s predetermined income percentile. These income profiles allow for predictable changes in average tax rates due to the hump shape of the life-cycle income profile and are described in more detail in the data section and the appendix.

As mentioned above, the assumption in equation (11) is least restrictive for higher-income households for which changes in the top tax rate are closely related to changes in their average tax rate. For this reason I focus on the consumption response of higher-income households to tax news. **Figure 5** shows that changes in average tax rates for households in the top quartile of the adjusted gross income distribution (AGI) are highly correlated over the sample period, although they are not perfectly correlated. I therefore choose the top income quartile of households (by tax filer status and year) in the CEX as my baseline sample, and I then analyze how the results change as I change this threshold. The different income distribution cut-offs trade off measurement error with statistical power: While including more households that are further away from the top tax bracket increases the precision of the estimates, it potentially biases the results if assumption (11) is less appropriate for those additional households, i.e., if tax news shocks are less well measured with municipal yield spreads for those households.

Estimation equation: The basic life-cycle model provides both a prescription for how to construct tax news shocks for each household in the CEX as well as an interpretation of the coefficient β of the following regression of the consumption growth rate on the tax news shock,

$$\Delta c_{it} = \beta \cdot \underbrace{\left(\sum_{s=1}^M w_{t,s}^{(M)} \frac{\bar{\tau}_{i,t+s}}{\tau_s} (\mathbb{E}_t \tau_{t+s} - \mathbb{E}_{t-1} \tau_{t+s}) \right)}_{\text{household-specific tax news shock}} + \alpha_t + \phi' z_{it} + \varepsilon_{it}. \quad (12)$$

α_t is a full set of year-by-month time fixed effects and z_{it} is a vector of additional controls that are standard in the literature. In particular, z_{it} includes a quadratic in the age of the reference person, quarterly changes in the number of adults and children, as well as other household characteristics such as family size, the maximum level of education achieved by either the reference person or the spouse, marital status, and the composition and number of earners. Furthermore, it also contains all the variables (i.e., the “main effects”) used to construct the tax news shock—which is an interaction term—such as federal adjusted gross, the household’s tax filing status, all perfect-foresight average

tax rates ($\bar{\tau}_{i,t+s}$), and the tax news shock formed with the residual of the low-pass filter applied to the break-even tax rates. The robustness section provides an extensive analysis of the effect of different time-series filter choices on the estimated consumption response.

Strictly speaking, the model predicts that $\beta = -1$, i.e., that consumption moves one-for-one with a change in the expected annuity value of life-cycle after-tax income, that is with permanent income. However, one should keep in mind the many simplifying assumptions discussed above that were required to obtain such a sharp prediction. Hence, finding a household response to this measure of tax news that is different from -1 should not be taken as a rejection of the rational-expectation life-cycle model of household consumption. The more relevant null hypothesis is $\beta = 0$, which holds if households are not forward-looking at all. This null hypothesis therefore turns the typical excess sensitivity test of the rational-expectation permanent income hypothesis (PIH) on its head. Instead of testing the sharp prediction of PIH for the consumption response to predictable cash-flow shocks (see e.g., [Kueng \(2016\)](#)), this null hypothesis tests the sharp prediction of models with no forward-looking behavior—such as hand-to-mouth behavior—for the consumption response to news shocks.

3.2 Identification of Forward-Looking Behavior

The tax news response β is identified if news about future tax changes are uncorrelated with other innovations ε that affect the household’s current consumption decision, conditional on observables z_t and time fixed effects α_t , that is if $\text{Cov}(\sum_{s=1}^M w_{t,s}^{(M)} \frac{\bar{\tau}_{i,t+s}}{\tau_{t+s}} (\mathbb{E}_t \tau_{t+s} - \mathbb{E}_{t-1} \tau_{t+s}), \varepsilon_{it} | z_{it}, \alpha_t) = 0$. As discussed above, the year-by-month time fixed effects control for changes in the average interest rate and hence for intertemporal substitution of spending, and they also control for the extent to which fiscal policy is used to counteract aggregate fluctuations in economic activity. While fiscal policy was extensively used prior to the 1980s, it was largely replaced by monetary policy as the main countercyclical policy tool since then, at least until very recently. The extensive narrative analysis of [Romer and Romer \(2010\)](#) therefore concludes that all income tax changes between 1980 and 2003—with one minor exception in 2001—are not countercyclical nor did they coincide with changes in government spending.²⁰ Romer and Romer classify those income tax reforms as exogenous, driven either by attempts to increase long-run economic growth (ERTA 1981, TRA 1986, EGTRRA 2001 and JGTRRA 2003) or by concerns about the federal budget deficit (OBRA 1990 and OBRA 1993).

While the tax reforms in the sample are orthogonal to the current state of the economy—and hence exogenous according to the terminology of Romer and Romer—the corresponding tax news shocks might nevertheless be correlated with news about future pre-tax income. Thus, while it is not possible to fully rule out that the consumption response is at least partially driven by correlated income news shocks, a consumption response to either type of shock still indicates that households

²⁰ The minor exception is the Economic Growth and Tax Relief Reconciliation Tax Act of 2001 (EGTRRA). The countercyclical part of EGTRRA concerns the accelerated implementation of the tax cuts but it does not concern the overall size of the cuts. For instance, [Romer and Romer \(2008\)](#) note that “this [countercyclical] motivation was almost always discussed in the context of making some of the cuts retroactive to January 1, 2001 rather than having them begin on January 1, 2002.” (p.84)

are forward-looking.

One final concern is that other shocks to the bond markets such as liquidity risk are correlated with the business cycle or with bad income news in general. To assess the potential bias introduced, note that when financial markets are under stress—such as during the financial crisis of 2008-2010—the liquidity premium on Treasury bonds tends to increase. Such periods are also associated with lower consumption. An increase in the demand for liquidity provided by Treasury bonds relative to municipal bonds increases the relative price for Treasuries and hence lowers the Treasury yield relative to the yield on municipal bonds. Equation (4) shows that this mechanism causes the BETRs θ_t to decrease and therefore decreases the measured path of expected forward tax rates $\mathbb{E}_t \tau_t$. To the extent that such liquidity shocks are not absorbed by year-by-month time fixed effects, by the filtering of the time series, or by directly controlling for them, they induce a spurious *positive* correlation between changes in the measured path of expected tax rates and consumption changes. This potential correlation leads the response coefficients to be biased upward against the forward-looking consumption model. For these reasons—that income tax reforms in the sample are unrelated to the current state of the economy, and that any remaining liquidity shock biases the results against the basic model—estimating regressions of the form (12) provide a meaningful assessment of the response of higher-income households’ consumption to news about future taxes. Furthermore, the robustness section addresses this concern directly by including the yield residual and by directly controlling for such bond risk factors.

3.3 Household Spending Data

I use data from the Consumer Expenditure Survey (CEX), which is the most comprehensive data source on household consumption in the U.S. A more detailed description of the data and data steps is provided in the appendix. The CEX is a monthly rotating panel and each household is interviewed once per quarter for at most four consecutive quarters. In each interview the reference period for expenditures covers the three months prior to the interview month. I follow the literature and aggregate the expenditures to quarterly spending. Income data is asked only in the first and last interview. The reference period for income covers the twelve months prior to the interview month. The definition of spending on nondurables and services follows the literature and excludes housing services, health care and health insurance, and education services, since these expenditures have characteristics of durable goods.²¹ All nominal variables are deflated using the CPI-U. Due to extreme outliers I winsorize consumption growth at the 10% level; all results are robust to choosing a different threshold. Table 2 provides summary statistics of the main sample of higher-income households used in the analysis. These households have a federal adjusted gross income (AGI) above the 75th percentile by year and tax filer status.

²¹ Following the previous literature, this implicitly assumes linear Engle curves or separability between nondurables and durables.

3.4 The Spending Response to the Reagan Tax Cuts: A Case Study

By far the largest change in marginal and average tax rates occurred in the early 1980s under President Reagan. Soon after Reagan announced his candidacy for President of the United States on November 13, 1979 bond markets started to factor in lower expected tax rates at the longer end of the term structure as seen in [Figure 4](#). Luckily, this event coincides with the start of the CEX sample. I therefore first conduct a case study of these tax changes before extending the sample to include the entire period from 1980 to 2001.

[Table 3](#) estimates the response of household spending on nondurables and services to news about future taxes using regression equation (12) on the sample of households with a federal AGI above the 75th percentile by tax filer status in each year from 1980 to 1982. Column 1 uses a minimalist specification including only year-by-month time fixed effects α_t along the tax news shock. The estimate suggests that household increase spending by about 1.5% in response to news about a decrease in the annuity value of future tax liabilities of 1%. Column 2 adds the residual of the time-series filter discussed in [section 2.2](#). The news shock residual’s effect on spending is economically small and not significantly different from zero consistent with the hypothesis that this variation represents changes in relative bond prices that are not driven by news about future taxes. Column 3 uses variables that are standard in the literature to control for predictable changes in household composition that can break the consumption Euler equation. The increases in spending related to increases in the number of adults or children do not affect the response to the tax news shock. Column 4 controls for the other main effects of the news shocks besides the time fixed effects, the average tax rate and fixed effects for each tax filer status.

The precision of the estimates in columns 1 to 4 is very low given the small sample size. Nevertheless the point estimates suggest a substantial spending response to tax changes that were implemented in later years. By increasing the sample size to include all years under President Reagan we can increase the precision of the estimates and check the robustness of the initial results. Columns 5 and 6 show that the point estimates are very similar when using this longer sample period from 1980 to 1988. Column 6 also controls for other household characteristics including household income (federal AGI), a quadratic of the head’s age, and fixed effects for marital status, family size, number and composition of earners, and the level of education, where I use the maximum number of years of schooling across head and spouse.

3.5 The Spending Response to Tax News in the Full Sample

[Table 4](#) presents the main results of the analysis. Column 1 uses the same minimalist specification as before, but now on the full sample of households with AGI in the top quartile (by tax filer status and year) from 1980 to 2001. The response of household spending to tax news is roughly one-for-one (in absolute value). Adding the residual of the filtered news shock series (column 2) and controlling for predictable changes in family composition (column 3) does not affect the response to the tax news shock. Column 4 adds the main effects of the news shock, which include the 30

perfect-foresight average tax rates $\bar{\tau}_{i,t+s}$ used to construct the tax news shock weights in equation (12) and fixed effects for each tax filer status. The estimated spending response to the tax news shock remains largely unchanged. Column 5 is the preferred specification using the full set of controls, including the other household characteristics described above such as household income and age. The household spending response to the news shocks remains about one-for-one, although still measured with substantial standard errors. Nevertheless, the null hypothesis of no forward-looking behavior is strongly rejected. As mentioned above, this of course does not confirm the basic life-cycle model, and given the relatively large standard errors many other models that also feature forward-looking behavior are consistent with the estimated response.

Finally, column 6 runs the same regression on the smaller sample of households in the top 10% of federal AGI for whom the news shock should be measured with less error since it depends much less on the assumption made in equation (11). While the precision decreases by about 50% as expected based on the substantial reduction in sample size, the point estimate is almost unchanged and remains economically and statistically significant.

3.6 Response Heterogeneity and Ricardian Equivalence

As mentioned above, a model with dynastic households that are altruistically linked or with an infinitely-lived representative agent with rational expectations would (under some assumptions) predict that consumption would not change to changes in expected future taxes if those taxes were lump-sum and if government spending was expected to remain unchanged; see Barro (1974) and Seater (1993). Ideally, we would have a measure of households' expectations of future changes in government spending (and of pre-tax income). Unfortunately, such measures are in general not available; see Ramey (2011) for an exception. In this section I instead analyze three dimensions of cross-sectional heterogeneity that are informative to assess the degree of Ricardian equivalence in the data.

First, I test whether there is a difference in the spending response to tax news shocks between families with and without children. Under Ricardian equivalence we would expect parents to respond less to the tax news since their children will be affected by future changes in government spending or by future changes in taxes not captured by the 30-year forecast horizon used in this paper; see Altonji, Hayashi and Kotlikoff (1992, 1997). Column 1 of Table 5 shows that we cannot reject the hypothesis that households with children respond as strongly as households without children. One might be worried about misclassification if young households plan to have children but currently do not have any yet or if older households had children that are not living in the same household anymore. Column 2 therefore restricts the sample to households between 30 and 50 years old. The point estimates remain largely unchanged suggesting that measurement error is not driving the similarity in the responses. In fact, in both cases the point estimate of households with children is slightly larger than that of childless families.

Second, we would expect younger households to respond less to the tax news since they are

more likely to face off-setting changes in government spending in the future. Column 3 shows that the response is similar for all households below age 55 and that households in the oldest 10-year age bin in fact respond less, not more, rejecting equal responses with a p-value of 0.3%. Again, the households which should respond the least in the presence of Ricardian equivalence effects have in fact the largest point estimates.

Third, I split the sample by whether household have a mortgage. We would expect that households with a mortgage would respond less if they fear that these future tax changes will be financed by reducing or eliminating the mortgage interest tax deduction, which is one instance of government benefits that disproportionately favor richer households.²² Column 4 instead shows that mortgage owners respond similarly to the news as households without a mortgage. Again, the point estimate for mortgage owners is larger, not smaller. Hence, these three tests suggest that Ricardian equivalence is not a first-order concern when interpreting the spending responses to tax news shocks.

4 Robustness of the Spending Response to News Shocks

In this section, I extend the previous analysis by adding two alternative tests of the underlying mechanism and I check the robustness of the finding that changes in expected future taxes have large effects on households' current spending. I investigate the sensitivity of the main result to using different sub-sample periods not previously analyzed and to alternative normalizations of the spending changes. I also test the robustness of the results to different procedures of filtering the tax news shocks and to controlling for bond risk directly. In all specifications I use the baseline sample of households in the top quartile of the AGI distribution by tax filer status in each year and the full set of controls as in column 5 of [Table 4](#).

4.1 Alternative Tests of the Mechanism

[Table 6](#) provides two alternative tests of the underlying mechanism that focus on the absence of a response in cases where we would not expect to find an effect. The first such test focuses on lower-income households. [Figure 5](#) shows that the average tax rate of households with an AGI below the median is only weakly correlated with average tax rates of households in the top of the income distribution. The partial correlation between changes in the average tax rates of households below the median and households in the top 1% is only 3% and is not statistically different from zero. Consistent with news about top tax rates not being informative about future average tax rates for households below the median, column 1 shows no spending response of those households to the bond-based tax news measure.

Second, I analyze the response of household spending during a period in which factors other than tax news dominate the municipal bond market and hence in which the identification approach should break down. As discussed above and seen in [Figure 1](#), in the absence of additional variation

²² Note that any wealthy-hand-to-mouth borrowing constraints would also predict a lower response of households with mortgages and hence bias the test in favor of Ricardian equivalence.

from election probabilities, the bond-based measure of expected tax rates works well in the late 1990s and particularly in the early 1980s when marginal rates change dramatically. However, the break-even tax rates at short maturities behave poorly in the early 1990s. Columns 2 and 3 split the sample accordingly in order to investigate whether the household spending response also reflects this pattern. Column 2 excludes the period from 1989 to 1992 when the relationship between short-term municipal yield spreads and the top tax rate breaks down. The spending response increases slightly but remains similar to the full-sample estimate. However, column 3 shows that the measure of expected tax liabilities does not predict household spending from 1989 to 1992, in line with this disconnect. In fact, the positive but insignificant point estimate of the news residual is consistent with the fact that short-term break-even tax rates move in opposite direction to the top marginal tax rate.

4.2 The Tax Reform Act of 1986 and Changes in the Tax Base

Table 7 provides additional robustness checks of the main results. As discussed above and more extensively in the online appendix, the Tax Reform Act of 1986 (TRA) also changed tax brackets in addition to changing marginal rates. Column 1 analyzes how excluding TRA from the analysis affects the results. Since the bond markets might have predicted part of the rate decreases associated with TRA already several years in advance as suggested by Figure 4, we cannot use data prior to 1986. Furthermore, since the municipal yield spread appears disconnected from tax fundamentals at the short end of the term structure in the early 1990s (Figure 1 and Table 6) I restrict the sample to years 1993-2001. Column 1 shows that despite the substantial reduction in the sample size, the estimated response of household spending to tax news shocks is similar to the response over the full sample period.

Column 2 relates the analysis to the previous literature on fiscal foresight by only using the legislated lags proposed by the narrative analysis of Yang (2008) as a measure of the length of tax policy foresight. These are defined as the periods between the time a tax proposal was first mentioned by a President and when it was enacted, but before rates changed. These periods therefore exclude the times before the formal introduction of a tax bill to Congress, which contains most of the news. In particular, it excludes any information that becomes available during a presidential election, such as the variation used for the two quasi-natural experiments and the variation associated with the Reagan tax cuts shown in the first five panels of Figure 4. Moreover, this restriction also excludes false negatives, i.e., news about possible future tax changes that did not materialize. Despite throwing away a substantial amount of identifying variation household spending strongly responds to the bond-based measure of fiscal news even in this very limited sample.

Finally, column 3 analyzes the robustness of the measure of consumption growth using lagged spending as an alternative normalization of spending changes, which is the specification obtained when log-linearizing the inter-temporal budget constraint. To make the coefficient comparable with

the previous specifications, I scale the series by the sample median of the ratio of consumption to predetermined income. The point estimate is very similar to the main response in column 5 of [Table 4](#), while the standard error is slightly larger.

4.3 Using Bartik-Type Reduced-Form News Shocks

Are the results driven by the structure imposed by the life-cycle model when deriving the household-level tax news shocks in equation (12)? To answer this question I construct an alternative reduced-form tax news shock that follows the Bartik approach outlined in [section 3.1](#) but does not use the life-cycle model to derive the weights. Specifically, I interact the unweighted present value of relative changes in expected top marginal tax rates with the predetermined average tax rate $\bar{\tau}_i^{lagged}$ based on income in the first interview instead of the perfect-foresight average tax rates $\{\bar{\tau}_{i,t+s}\}_{s=1}^M$ implied by the life-cycle model. This leads to the following regression specification,

$$\Delta c_{it} = \beta \cdot \underbrace{\left(\bar{\tau}_i^{lagged} \times \sum_{s=1}^M w_{t,s}^{(M)} \frac{\mathbb{E}_t \tau_{t+s} - \mathbb{E}_{t-1} \tau_{t+s}}{\tau_{t+s}} \right)}_{\text{Bartik-type tax news shock}} + \alpha_t + \phi' z_{it} + \varepsilon_{it}. \quad (13)$$

[Table 8](#) shows very similar result when using this reduced-form tax news shock instead of the more structural news shock in (12). In particular, higher-income households with an AGI in the top quartile (column 1) or in the top 10% (column 2) strongly respond to news about future tax changes, while households with below-median AGI (column 3) are not affected by news about future tax rates in higher income brackets. Hence, the results are not driven by the structural model used to derive the household-level tax news shocks used in the main analysis.

4.4 Effect of Different News-Noise Decompositions

The analysis so far used the two-sided low-pass filter described in the previous section to filter out noise unrelated to tax news from the break-even tax rate series. [Table 9](#) reports results using different filters shown in [Figure 6](#). The two year cut-off is conservative since it probably filters out some tax news shocks. This loss of information lowers the precision of the spending response estimates. On the other hand, this cut-off value lowers the level of noise in the tax news shocks. The reduction of the noise reduces the potential attenuation bias in the spending response coefficients. Therefore, the choice of the frequency cut-off reflects a trade-off between bias and efficiency of the estimates. To see this trade-off, columns 2 and 3 use different cut-offs. Column 2 uses a filter that passes frequencies lower than eight years and hence is even more conservative than the baseline two-year filter. While eight years might appear to be excessive, [Figure 6](#) shows that this filtered series is not very different from the filtered series using the baseline two-year filter. As expected, column 2 reports a less precisely estimated spending response with a higher point estimate, although one cannot reject that the response is the same as in the main analysis reproduced for convenience in column 1.

Column 3 on the other hand passes higher frequencies than in the baseline specification. [Figure 6](#) shows that passing all frequencies lower than half a year makes the filtered series fairly volatile. As expected, the standard error of the estimated spending response falls by more than 60%, but the noisier measure severely attenuates the estimated spending response. To see this point more clearly, column 4 uses the raw series without any filtering. As expected, the substantial noise introduced by the high-frequency shocks biases the spending response to zero. In fact, the attenuation bias in both cases is almost complete, resulting in an estimated response that is economically and statistically insignificant.

Column 5 addresses the concern that a two-sided filter uses not only current and lagged observations, but also future ones. Therefore, the econometrician might overestimate the information set of consumers, thereby underestimating the response to news. This would be the case for example if households tried to infer the path of expected tax rates $\mathbb{E}_t\tau$ from municipal yield spreads θ_t . The households and the econometrician would solve effectively the same signal extraction problem in real time. Therefore, in this scenario the econometrician should only use current and past bond prices to infer the households' information sets. The difference between the two ways of modeling information can be seen from the way the econometrician solves the signal extraction problem of equation (5). The approach taken so far implies that the optimal solution is a two-sided filter, while this alternative view requires the use of a backward-looking one-sided filter. To account for this possibility, I use the optimal one-sided Hodrick-Prescott filter suggested by [Mehra \(2004\)](#), and I follow [Ravn and Uhlig \(2002\)](#) in setting the smoothing parameter. One-sided filters only use current and lagged observations, but induce a phase-shift; that is, the filtered series lags the raw data and hence detects a change in the trend only with some delay. This phase-shift, shown in [Figure 6](#), introduces measurement error in the news shock and might therefore bias the spending response toward zero under the plausible scenario in which households do not use municipal bonds to infer future tax rates. Column 5 shows that the estimated spending response to the tax news shock that uses a one-sided filter is very similar to the main estimate in column 1.

Column 6 instead directly controls for these other bond risk factors by taking residuals from regressing the BETRs on the corresponding Treasury on-the-run/off-the-run yield spreads before applying the filter.²³ The household spending response remains economically and statistically significant and is similar to the baseline result. The smaller point estimate is consistent with the regression residual being more volatile and hence noisier than the original BETR series.

Finally, note that all specifications of [Table 9](#) contain essentially the same information: adding the residual to the filtered news shock reproduces the raw series used in column 4. The residual is never statistically or economically significant. This suggest for example that even if we interpret the model of the term-structure of break-even tax rates as a linear approximation of a more complete but non-linear model such as the one proposed by [Green \(1993\)](#), the higher-order terms of such a model do not significantly affect household spending behavior.

²³ This spread is the only risk measure available going back to the late 1970s.

Overall, the results in this section show that the estimated spending response to tax news is similar under alternative specifications. More importantly, the point estimates change in the direction we would expect if the tax news shocks indeed measure the persistent component of changes in expected future tax liabilities. First, the response is lower in absolute value than predicted by the life-cycle model if we use shocks that are not guided by theory or if we use noisier shocks while it is in line with the theory if we filter out the noise and use theory-based news shocks (Tables 8 and 9). Second, the response is stronger when looking at households which are more affected by the news or when looking at periods during which we would expect the identification approach to work best (Table 6).

5 Conclusion

This paper identifies tax expectations using the yield spread between taxable and tax-exempt bonds with maturities of one to thirty years. Combining these tax expectations with household spending data shows that the basic life-cycle consumption model with forward-looking households describes the behavior of higher-income households remarkably well. This paper is the first to directly measure the response of household spending to news shocks, and thus is the first direct test of the theory's prediction for the response of household spending to new information.

While a full analysis of the macroeconomic implications of these results is beyond the scope of this paper, it is nevertheless useful to discuss two issues related to fiscal policy. First, the response of household consumption to news about future tax policies suggests that studies which ignore anticipation effects do not capture the total effect of a tax change on the economy. Whether the anticipation effect documented in this paper amplifies or dampens the effect that such studies typically find—for example due to households that exhibit excess sensitivity in consumption—depends on how this policy changes households' expectations about other future tax policies. For example, a tax reform that is expected to permanently lower taxes would cause unconstrained forward-looking households to respond to the news, while constrained households (either because of economic or behavioral frictions) would respond to the actual tax change, in which case the total measured effect of a tax reform would be amplified by including the anticipation effects. On the other hand, a tax stimulus that targets constrained households and is expected to be paid for by increasing taxes on unconstrained households in the future would cause the anticipation and excess sensitivity effects to go in opposite directions, thereby dampening the total effect of the stimulus.

Second, while the anticipation effect could partially overcome the long implementation lag of fiscal policies (a major limitation of taxes as short-run policy instruments) the fact that countercyclical policies are usually designed to be budget neutral over the business cycle implies that forward-looking households might realize that a countercyclical tax policy has little effect on their lifetime income and might thus respond neither to the news nor to the actual tax change.

Clearly, more research on the response of households to news shocks needs to be undertaken before such data can offer policy guidance that is empirically well-grounded. Identifying more news

shocks that directly affect household budget sets is clearly desirable in order to verify the results of this study.

References

- Abowd, John M. and David Card, "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 1989, 57 (2), 411–445.
- Altonji, Joseph G., Fumio Hayashi, and Laurence J. Kotlikoff, "Is the Extended Family Altruistically Linked? Direct Tests Using Micro Data," *American Economic Review*, 1992, pp. 1177–1198.
- Altonji, Joseph G., Fumio Hayashi, and Laurence J. Kotlikoff, "Parental Altruism and Inter Vivos Transfers: Theory and Evidence," *Journal of Political Economy*, 1997, 105 (6), 1121–1166.
- Ang, Andrew, Vineer Bhansali, and Yuhang Xing, "Taxes on Tax-Exempt Bonds," *Working Paper*, 2007.
- Ayers, Benjamin C., Bryan C. Cloyd, and John B. Robinson, "'Read My Lips ...': Does the Tax Rhetoric of Presidential Candidates Affect Security Prices?," *Journal of Law and Economics*, 2005, 48 (1), 125–148.
- Babina, Tania, Chotibhak Jotikasthira, Christian Lundblad, and Tarun Ramadorai, "Heterogeneous Taxes and Limited Risk Sharing: Evidence from Municipal Bonds," *Working Paper*, 2015.
- Baker, Scott R., Nicholas Bloom, and Steve J. Davis, "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 2016, 131, x–xx.
- Barro, Robert J., "Are Government Bonds Net Wealth?," *Journal of Political Economy*, 1974, 82 (6), 1095–1117.
- Bartik, Timothy J., "Who Benefits from State and Local Economic Development Policies?," in A. J. Culyer and J. P. Newhouse, eds., *Upjohn Working Papers and Journal Articles*, W.E. Upjohn Institute for Employment Research, 1991.
- Baxter, Marianne and Robert G. King, "Measuring business cycles: approximate band-pass filters for economic time series," *Review of Economics and Statistics*, 1999, 81 (4), 575–593.
- Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 2009.
- Calmes, Jackie, "Bush Promotes Big Tax Cuts in a Move To Attract Middle-Class Swing Voters," *The Wall Street Journal*, 2000, Monday, September 25, A.16.
- Campbell, John Y. and Angus Deaton, "Why is Consumption So Smooth?," *Review of Economic Studies*, 1989, 56 (3), 357.
- Chalmers, John M.R., "Default Risk Cannot Explain the Muni–Puzzle: Evidence from Bonds that Are Secured by U.S. Treasury Obligation," *Review of Financial Studies*, 1998, 11 (2), 281–308.
- Fama, Eugene F., "A Pricing Model of the Municipal Bond Market," *Working Paper*, 1977.
- Fitch Ratings, "Default Risk and Recovery Rates on U.S. Municipal Bonds," *Public Finance Special Report*, 2007.
- Fortune, Peter, "Do Municipal Bond Yields Forecast Tax Policy?," *New England Economic Review*, 1996.
- Green, Richard C., "A Simple Model of the Taxable and Tax-Exempt Yield Curves," *Review of Financial Studies*, 1993, 6 (2), 233–264.
- Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, "The U.S. Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, 2007, 54 (8), 2291–2304.
- Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 1978, 86 (6), 971–987.
- Harris, Lawrence E. and Michael S. Piwowar, "Municipal Bond Liquidity," *Working Paper*, 2004.
- Hempel, George H., *The Postwar Quality of State and Local Debt*, Columbia University Press, 1971.
- Hodrick, Robert J. and Edward C. Prescott, "Post-War US Business Cycles: An Empirical Investigation,"

- Carnegie Mellon University discussion paper no. 451*, 1980.
- Jappelli, Tullio and Luigi Pistaferri, “The Consumption Response to Income Changes,” *Annual Review of Economics*, 2010, 2, 479–506.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 2006, 96 (5), 1589–1610.
- Kochin, Levis A. and Richard W. Parks, “Was the Tax-Exempt Bond Market Inefficient or Were Future Expected Tax Rates Negative?,” *Journal of Finance*, 1988, 43, 913 – 931.
- Kueng, Lorenz, “The Taxation of Bonds: A Short Primer,” *Working Paper*, 2012.
- , “Explaining Consumption Excess Sensitivity with Near-Rationality: Evidence from Large Predetermined Payments,” *Working Paper*, 2016.
- Longstaff, Francis A., “Municipal Debt and Marginal Tax Rates: Is There a Tax Premium in Asset Prices?,” *Journal of Finance*, 2011, 66, 721–751.
- Lorenzoni, Guido, “News and Aggregate Demand Shocks,” *Annual Review of Economics*, 2011, 3, 537–557.
- Mankiw, N. Gregory and James M. Poterba, “Stock Market Yields and the Pricing of Municipal Bonds,” *Working Paper*, 1996.
- Meghir, Costas and Luigi Pistaferri, “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 2004, 72 (1), 1–32.
- Mehra, Yash P., “The Output Gap, Expected Future Inflation and Inflation Dynamics: Another Look,” *The BE Journal of Macroeconomics*, 2004, 4 (1), 1–17.
- Poterba, James M., *Explaining the Yield Spread between Taxable and Tax-exempt Bonds: The Role of Expected Tax Policy*, The University of Chicago Press, 1986.
- Ramey, Valerie A., “Identifying Government Spending Shocks: It’s all in the Timing,” *Quarterly Journal of Economics*, 2011.
- Ravn, Morten O. and Harald Uhlig, “On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations,” *Review of Economics and Statistics*, 2002, 84 (2), 371–376.
- Romer, Christina D. and David H. Romer, “A Narrative Analysis of Postwar Tax Changes,” *online appendix to Romer and Romer (American Economic Review 2010)*, “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks”, 2008.
- and —, “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks,” *American Economic Review*, 2010, 100 (3), 763–801.
- Saez, Emmanuel, “Reported Incomes and Marginal Tax Rates, 1960-2000: Evidence and Policy Implications,” *NBER Working Paper No. 10273*, 2004.
- Seater, John J., “Ricardian Equivalence,” *Journal of Economic Literature*, 1993, 31 (1), 142–190.
- Seib, Gerald F. and Alan Murray, “Where They Stand: A Voters’ Guide to the Candidates and the Issues,” *The Wall Street Journal*, 1992, Monday, September 21, R6.
- Slemrod, Joel and Timothy Greimel, “Did Steve Forbes scare the US municipal bond market?,” *Journal of Public Economics*, 1999, 74 (1), 81–96.
- Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz, “How Prediction Markets Can Save Event Studies,” *NBER Working Paper No. 16949*, 2011.
- Souleles, Nicholas S., “The Response of Household Consumption to Income Tax Refunds,” *American Economic Review*, 1999, 89 (4), 947–958.
- Yang, Shu-Chun Susan, “A Chronology Of Postwar U.S. Federal Income Tax Policy,” *Working Paper*, 2008.

A Appendix

This section provides more details about the household spending and income data used in the analysis. I use sample restrictions that are common in the literature; see for example, Souleles (1999) or Johnson, Parker and Souleles (2006). In particular, I restrict the sample to households with complete income reports and non-zero income where the head’s age is between 25 and 65 years and the head is not a student. In addition, I drop the following cases: interviews with more or less than three monthly observations; households with zero food or total expenditures; non-consecutive interviews; observations with negative expenditures where there should not be any; households with more than one consumer unit; households for which the family size changes by more than three; households for which the age of any member increases by more than one or decreases; and households with negative liquid wealth; and households with positive business or farm income. I correct sample breaks in the following few expenditure categories due to slight changes in the questionnaire: food at home (’82Q1-’88Q1), personal care services (’01Q2), occupation expenditures (’01Q2), and property taxes (’91Q1). As recommended by the BLS, I sum expenditures that occur in the same month but are reported in different interviews.

I impute taxes with the NBER TAXSIM calculator using an iterative procedure to determine the itemization status of each household and to account for deductions that depend on the household’s AGI such as health-care or job expenses. The code is available at www.nber.org/~taxsim/to-taxsim/cex-kueng/cex.do. I compute perfect-foresight average tax rates $\{\bar{\tau}_{i,t+s}\}_{s=0}^M$ for each household i that depend on the head of household’s age and the household’s predetermined income percentile. These profiles allow for predictable changes in average tax rates due to the hump shape of the life-cycle income profile. Other studies also found that household income dynamics are well approximated by a random walk after controlling for the age profile of income; see e.g., Abowd and Card (1989) and Meghir and Pistaferri (2004). Predetermined income summarizes observed household characteristics such as education and experience as well as unobserved heterogeneity such as work effort. Household age has predictive power for future income even conditional on work experience. Having only two dimensions guarantees that there are at least 20 households in each age-income cell in each year.

I estimate future average tax rates non-parametrically; in particular, I discretize the joint distribution of age and income and assume that the household remains in the same age-specific income percentile throughout its life-cycle.²⁴ I use households age 65 to 75 to estimate counter-factual retirement income after age 65. I assume that households expect to receive this level of retirement income for the rest of the planning period. I limit the estimation of the retirement period to households between ages 65 to 75 due to the fact that the quality of the survey answers tends to be poorer for old retirees.

²⁴ More precisely, I use the following income percentile thresholds: 10, 20, . . . , 50, 55, . . . , 95. I use a finer grid for higher incomes to better account for the increasing income inequality during the sample period. I use age bins with a 10-year range to have at least 20 households in each cell. The five age bins 25-34, 35-44, 45-54, 55-64, 65-75 approximate the income life-cycle profiles well.

Table 1: Quasi-natural experiments - Break-even tax rate responses to changes in election probabilities

	maturity of break-even tax rate (BETR):							
	1-year (1)	2-year (2)	3-year (3)	5-year (4)	7-year (5)	10-year (6)	20-year (7)	30-year (8)
A: Election in 2000								
Price of Bush contract [in cents]	0.018 (0.013)	-0.018*** (0.007)	-0.031*** (0.007)	-0.033*** (0.007)	-0.028*** (0.010)	-0.024** (0.011)	-0.006 (0.009)	0.003 (0.013)
Observations	131	131	131	131	131	131	131	131
R-squared	0.892	0.943	0.944	0.938	0.902	0.842	0.833	0.740
B: Election in 1992								
Price of Clinton contract [in cents]	0.140*** (0.048)	0.096** (0.047)	0.140*** (0.040)	0.091*** (0.026)	0.097*** (0.022)	0.103*** (0.021)	0.042*** (0.015)	0.047** (0.018)
Observations	82	82	82	82	82	82	82	82
R-squared	0.958	0.930	0.945	0.969	0.954	0.936	0.974	0.964

Notes: The contract yields 100 cents if the candidate wins the popular vote and zero otherwise. Therefore, an increase of the price by 1 cent corresponds to a 1% increase in the perceived probability of the candidate winning the presidential election. All regressions include controls for (il)liquidity such as the number of repeated prices, muni supply (negotiated and competitive offers), on-the-run/off-the-run Treasury spread, prediction market volume and units traded; for credit risk, such as the pre-refunded spread, spreads between lower-grade municipal bonds; and for call option features, such as the taxable yield curve and the weekly standard deviation of long-term treasuries. Newey-West HAC robust standard errors in parentheses: ***, **, * mark significance at the 1, 5, and 10 percent level, respectively.

Table 2. Summary statistics of CEX sample

Variable	Observations	Mean	Median	Std. Dev.
Nondurable spending growth [in %]	28007	0.08	0.00	6.73
Tax news shock [in %]	28007	-0.03	0.00	0.56
Residual of filtered news shock	28007	-0.01	0.00	2.52
Average tax rate	28007	23.09	22.65	4.32
Federal AGI (in \$1,000 of 1982-84)	28007	55.06	49.67	31.40
Number of earners	28007	1.98	2	0.90
Age of reference person	28007	43.58	43	9.90
Family size	28007	2.88	3	1.41
Change in number of adults	28007	0.01	0	0.28
Change in number of children	28007	0.00	0	0.21

	Fraction (%)		Fraction (%)
<i>Education:</i>		<i>Tax filer status:</i>	
- no school	0.03	- single	25.08
- high school dropout	5.19	- married filing jointly	60.73
- high school	21.08	- head of household	14.19
- college dropout	21.35		
- college	27.73	<i>Composition of earners:</i>	
- graduate school	24.63	- reference person only	27.51
		- reference person and spouse	33.58
<i>Marital status:</i>		- ref. person, spouse and others	15.53
- married	62.15	- reference person and others	19.34
- widowed	3.61	- spouse only	1.48
- divorced	17.64	- spouse and others	0.71
- separated	3.62	- others	0.95
- never married	12.97	- no earners	0.90

Notes: The sample includes households with federal adjusted gross income (AGI) in the top quartile by tax filer status (single, married filing jointly, or head of household) in each year from 1980 to 2001. Education is the maximum level of education achieved by the reference person and the spouse. All summary statistics in this table use sampling weights while all regressions in the following tables are unweighted.

Table 3: A case study of the Reagan tax cuts

Dep. var.: Δc_{it} , "three-monthly" nondurables and services	sample period 1980-82				sample period 1980-88	
	baseline	news residual	household composition	main effects	main effects	other household characteristics
	(1)	(2)	(3)	(4)	(5)	(6)
Tax news shock [in %]	-1.556** (0.785)	-1.555** (0.786)	-1.558** (0.786)	-1.216 (0.960)	-1.423** (0.582)	-1.397** (0.580)
Residual of filtered news shock		0.010 (0.358)	0.036 (0.355)	0.037 (0.354)	-0.054 (0.155)	-0.056 (0.156)
Change in the number of adults			2.340*** (0.593)	2.367*** (0.595)	1.786*** (0.295)	1.640*** (0.305)
Change in the number of children			1.463** (0.686)	1.495** (0.692)	1.162*** (0.372)	0.984*** (0.372)
- Time FE (year-by-month)	YES	YES	YES	YES	YES	YES
- Average tax rate				YES	YES	YES
- Tax filer status FE				YES	YES	YES
- Other household characteristics						YES
Observations	2,886	2,886	2,886	2,886	11,071	11,071
R-squared	0.023	0.023	0.029	0.029	0.031	0.034

Notes: The sample includes households with a federal adjusted gross income (AGI) in the top quartile of the cross-sectional AGI distribution by tax filer status (single, married filing jointly, or head of household) in each year. Columns (1)-(4) use observations for years 1980-82; columns (5) and (6) use observations for years 1980-88. The dependent variable is the growth rate of household spending on nondurable goods and services. The average tax rate in columns (4)-(6) is the household's federal average tax rate based on predetermined income in the first interview. Other household characteristics include household federal AGI, a quadratic in the age of the head of household, and fixed effects for education (maximum of head and spouse), marital status, family size, number of earners, and composition of earners (e.g., single-earner household, two-income family, etc.). Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity. ***, **, * mark significance at the 1, 5, and 10 percent level, respectively.

Table 4: Response of nondurable spending to tax news, 1980-2001

Dep. var.: Δc_{it} , "three-monthly" nondurables and services	AGI > 75 th percentile					AGI > 90 th percentile
	baseline	news residual	household composition	main effects	other household characteristics	
	(1)	(2)	(3)	(4)	(5)	(6)
Tax news shock [in %]	-1.162*** (0.304)	-1.162*** (0.304)	-1.142*** (0.304)	-1.104*** (0.308)	-1.034*** (0.309)	-1.000** (0.488)
Residual of filtered news shock		-0.015 (0.104)	-0.014 (0.104)	-0.017 (0.104)	-0.014 (0.104)	-0.089 (0.190)
- Time FE (year-by-month)	YES	YES	YES	YES	YES	YES
- Changes in family composition			YES	YES	YES	YES
- Perfect-foresight average tax rates				YES	YES	YES
- Tax filer status FE				YES	YES	YES
- Other household characteristics					YES	YES
Observations	28,007	28,007	28,007	28,007	28,007	11,368
R-squared	0.023	0.023	0.027	0.028	0.030	0.040

Notes: The sample includes households with a federal adjusted gross income (AGI) in the top quartile of the cross-sectional AGI distribution by tax filer status (single, married filing jointly, or head of household) in each year from 1980 to 2001 in columns (1)-(5), and with an AGI above the 90th percentile in column (6). The dependent variable is the growth rate of household spending on nondurable goods and services. Changes in family composition control for quarterly changes in the number adults and children separately. The perfect foresight average tax rates are the 30 rates used to construct the tax new shock weights in equation (12). Other household characteristics include household federal AGI, a quadratic in the age of the head of household, and fixed effects for education (maximum of head and spouse), marital status, family size, number of earners, and composition of earners (e.g., single-earner household, two-income family, etc.). Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity. ***, **, * mark significance at the 1, 5, and 10 percent level, respectively.

Table 5: Heterogeneity in the response to to tax news, 1980-2001

Dep. var.: Δc_{it} , "three-monthly" nondurables and services	by having children		by age	by mortgage ownership
	ages 25-65	ages 30-50		
	(1)	(2)		
Tax news shock x Not having children	-0.983*** (0.312)	-0.944** (0.412)		
Tax news shock x Having children	-1.089*** (0.319)	-1.023** (0.422)		
Tax news shock x Age 25-34			-1.062*** (0.314)	
Tax news shock x Age 35-44			-0.934*** (0.329)	
Tax news shock x Age 45-54			-0.955*** (0.358)	
Tax news shock x Age 55-65			-0.458 (0.367)	
Tax news shock x Not having a mortgage				-1.018*** (0.309)
Tax news shock x Having a mortgage				-1.119*** (0.332)
- Time FE (year-by-month)	YES	YES	YES	YES
- Residual of filtered news shock	YES	YES	YES	YES
- Changes in family composition	YES	YES	YES	YES
- Perfect-foresight average tax rates	YES	YES	YES	YES
- Tax filer status FE	YES	YES	YES	YES
- Other household characteristics	YES	YES	YES	YES
- Having children FE	YES	YES		
- 10-year age bins FE			YES	
- Mortgage-ownership FE				YES
Observations	28,007	28,007	28,007	28,007
R-squared	0.030	0.038	0.030	0.030

Notes: The sample includes households with a federal adjusted gross income (AGI) in the top quartile of the cross-sectional AGI distribution by tax filer status (single, married filing jointly, or head of household) in each year from 1980 to 2001. The dependent variable is the growth rate of household spending on nondurable goods and services. Changes in family composition control for quarterly changes in the number adults and children separately. The perfect foresight average tax rates are the 30 rates used to construct the tax new shock weights in equation (12). Other household characteristics include household federal AGI, a quadratic in the age of the head of household except in column (3), and fixed effects for education (maximum of head and spouse), marital status, family size, number of earners, and composition of earners (e.g., single-earner household, two-income family, etc.). Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity. ***, **, * mark significance at the 1, 5, and 10 percent level, respectively.

Table 6: Testing for spending responses where we would not necessarily expect any

Dep. var.: ΔC_{it} , "three-monthly" nondurables and services	AGI < 50 th percentile	municipal market disconnect	
		excluding 1989-92	including only 1989-92
	(1)	(2)	(3)
Tax news shock [in %]	-0.031 (0.230)	-1.132*** (0.335)	-0.093 (1.240)
Residual of filtered news shock	-0.004 (0.072)	-0.098 (0.117)	0.318 (0.217)
- Time FE (year-by-month)	YES	YES	YES
- Tax filer status FE	YES	YES	YES
- Perfect-foresight average tax rates	YES	YES	YES
- Changes in family composition	YES	YES	YES
- Other household characteristics	YES	YES	YES
Observations	51,830	22,430	5,577
R-squared	0.016	0.031	0.035

Notes: The sample in column (1) includes households with a below-median federal adjusted gross income (AGI) by tax filer status (single, married filing jointly, or head of household) in each year from 1980 to 2001. Columns (2) and (3) split the sample of households with AGI in the top quartile by tax filer status in each from 1980 to 2001 into two sub-samples, containing years 1989-92 in column (3) and years 1980-88 and 1993-2001 in column (2). The dependent variable is the growth rate of household spending on nondurable goods and services. The perfect foresight average tax rates are the 30 rates used to construct the tax new shock weights in equation (12). Changes in family composition control for quarterly changes in the number adults and children separately. Other household characteristics include household federal AGI, a quadratic in the age of the head of household, and fixed effects for education (maximum of head and spouse), marital status, family size, number of earners, and composition of earners (e.g., single-earner household, two-income family, etc.). Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity. ***, **, * mark significance at the 1, 5, and 10 percent level, respectively.

Table 7: Additional robustness checks of the main results

Dep. var.: Δc_{it} , "three-monthly" nondurables and services	using post- TRA86 sample	using Yang (2008) dates	using nondurables as basis
	(1)	(2)	(3)
Tax news shock [in %]	-1.048*** (0.372)	-0.737** (0.350)	-0.989*** (0.343)
Residual of filtered news shock	-0.173 (0.166)	-0.089 (0.077)	-0.042 (0.109)
- Time FE (year-by-month)	YES	YES	YES
- Tax filer status FE	YES	YES	YES
- Perfect-foresight average tax rates	YES	YES	YES
- Changes in family composition	YES	YES	YES
- Other household characteristics	YES	YES	YES
Observations	11,359	12,845	28,007
R-squared	0.028	0.021	0.031

Notes: The sample includes households with a federal adjusted gross income (AGI) in the top quartile of the cross-sectional AGI distribution by tax filer status (single, married filing jointly, or head of household) in each year. Column (1) uses the post-TRA 86 sample including years 1993-2001 (i.e., after the municipal market disconnect). Column (2) uses the news dates suggested by the narrative analysis of Yang (2008). Column (3) uses the full sample from 1980 to 2001. The dependent variable is the growth rate of household spending on nondurable goods and services. The perfect foresight average tax rates are the 30 rates used to construct the tax news shock weights in equation (12). Changes in family composition control for quarterly changes in the number adults and children separately. Other household characteristics include household federal AGI, a quadratic in the age of the head of household, and fixed effects for education (maximum of head and spouse), marital status, family size, number of earners, and composition of earners (e.g., single-earner household, two-income family, etc.). Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity. ***, **, * mark significance at the 1, 5, and 10 percent level, respectively.

Table 8: Using reduced-form Bartik-type news shocks, 1980-2001

Dep. var.: ΔC_{it} , "three-monthly" nondurables and services	AGI > 75 th percentile (1)	AGI > 90 th percentile (2)	AGI < 50 th percentile (3)
Average tax rate x PV of expected tax changes	-0.743** (0.335)	-1.217*** (0.464)	0.007 (0.235)
Average tax rate x PV of tax change residuals	-0.046 (0.102)	-0.098 (0.185)	-0.005 (0.070)
- Time FE (year-by-month)	YES	YES	YES
- Average tax rate	YES	YES	YES
- Tax filer status FE	YES	YES	YES
- Changes in family composition	YES	YES	YES
- Other household characteristics	YES	YES	YES
Observations	28,007	11,368	51,830
R-squared	0.029	0.039	0.015


Notes: The sample includes households with a federal adjusted gross income (AGI) in the top quartile of the cross-sectional AGI distribution by tax filer status (single, married filing jointly, or head of household) in each year from 1980 to 2001 in column (1), in the top 10th percentile in column (2), and with a below-median AGI in column (3), respectively. The dependent variable is the growth rate of household spending on nondurable goods and services. The average tax rate is the household's federal average tax rate based on predetermined income in the first interview. Changes in family composition control for quarterly changes in the number adults and children separately. Other household characteristics include household federal AGI, a quadratic in the age of the head of household, and fixed effects for education (maximum of head and spouse), marital status, family size, number of earners, and composition of earners (e.g., single-earner household, two-income family, etc.). Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity. ***, **, * mark significance at the 1, 5, and 10 percent level, respectively.

Table 9: Effect of using different time-series filters for the signal-noise decomposition

Dep. var.: Δc_{it} , "three-monthly" nondurables and services	two-sided BP filter			no filter	one-sided HP filter	control for bond market liquidity
	baseline	pass lower freq.	pass higher freq.			
	(1)	(2)	(3)	(4)	(5)	(6)
Tax news shock [in %] (signal)	-1.034*** (0.309)	-1.236*** (0.366)	-0.082 (0.121)	-0.063 (0.100)	-0.936*** (0.302)	-0.707** (0.300)
Residual of filtered news shock (noise)	-0.014 (0.104)	-0.018 (0.102)	-0.024 (0.175)		0.001 (0.104)	-0.028 (0.104)
- Time FE (year-by-month)	YES	YES	YES	YES	YES	YES
- Changes in family composition	YES	YES	YES	YES	YES	YES
- Other household characteristics	YES	YES	YES	YES	YES	YES
- Perfect-foresight average tax rates	YES	YES	YES	YES	YES	YES
Observations	28,007	28,007	28,007	28,007	28,007	28,007
R-squared	0.030	0.030	0.030	0.030	0.030	0.030


Notes: The sample includes households with a federal adjusted gross income (AGI) in the top quartile of the cross-sectional AGI distribution by tax filer status (single, married filing jointly, or head of household) in each year from 1980 to 2001. The dependent variable is the growth rate of household spending on nondurable goods and services. The perfect foresight average tax rates are the 30 rates used to construct the tax new shock weights in equation (12). Changes in family composition control for quarterly changes in the number adults and children separately. Other household characteristics include household federal AGI, a quadratic in the age of the head of household, and fixed effects for education (maximum of head and spouse), marital status, family size, number of earners, and composition of earners (e.g., single-earner household, two-income family, etc.). Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity. ***, **, * mark significance at the 1, 5, and 10 percent level, respectively.

Figure 1 – Nonparametric evidence of the importance of expected tax rates for the break-even tax rate



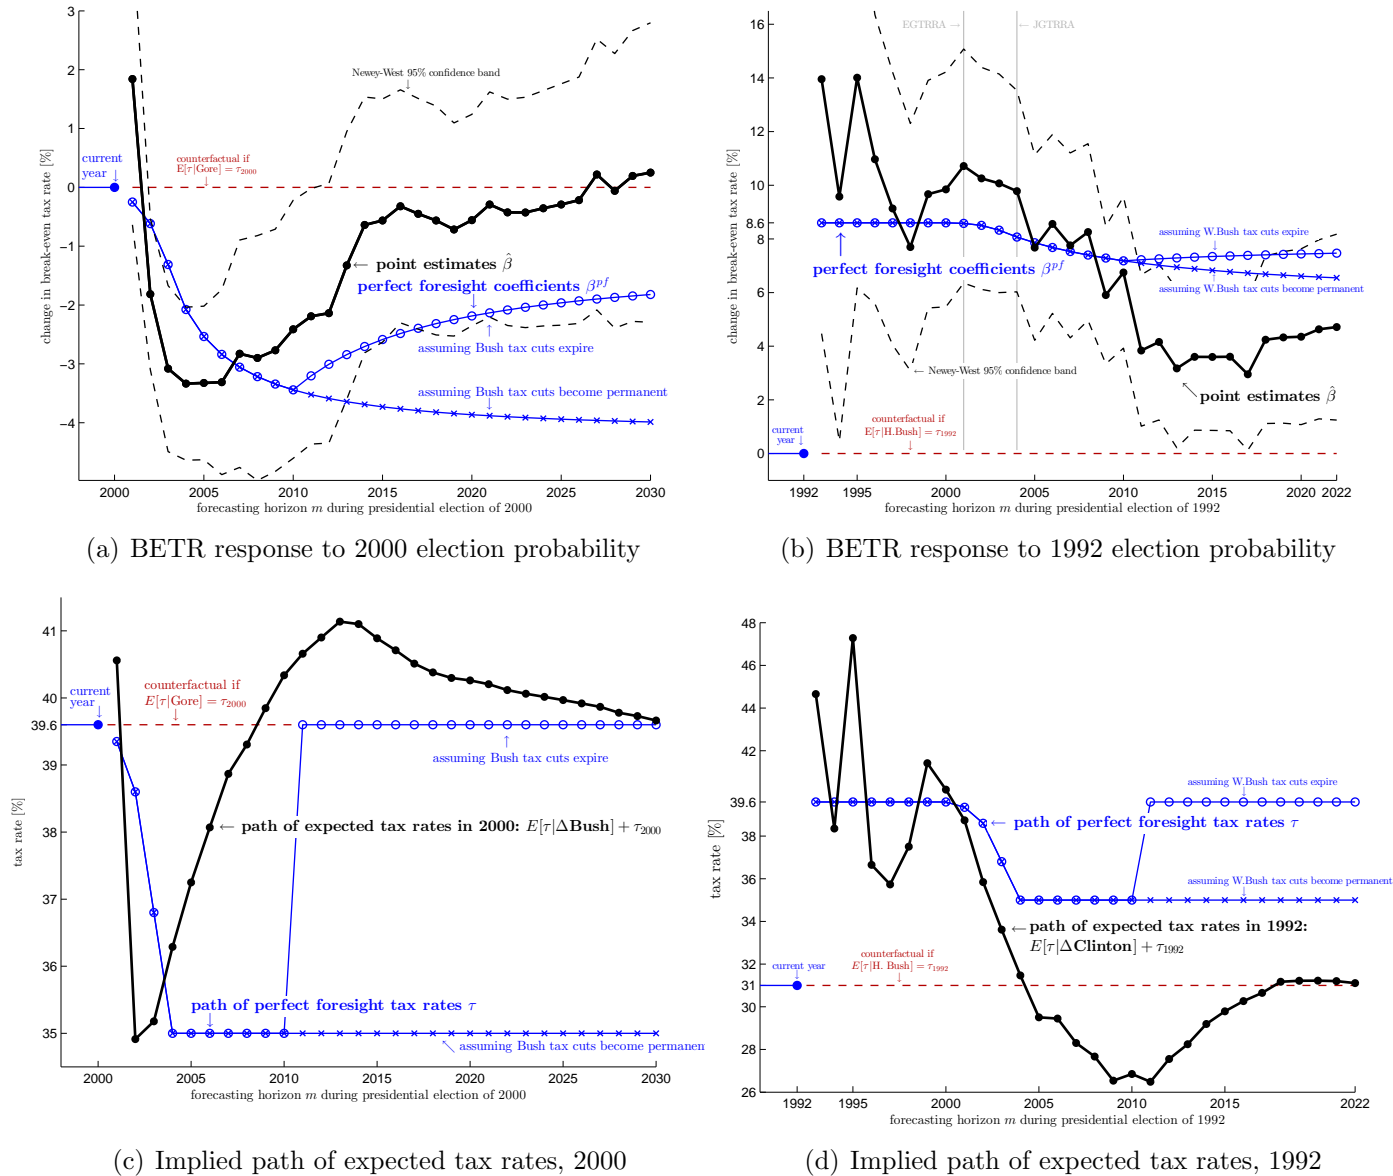
Notes: The thin lines are the raw data and the thick lines are the corresponding low frequency components of the 2-year and the 15-year break-even tax rates, respectively, corresponding to equation (4). The blue thin dashed line is the top 1% tax rate taken from Saez (2004). The solid blue line is the '33% tax bubble' during the years 1988-1990; in this period, the top marginal tax rate is higher than the marginal tax rate of the top 1% of the income distribution.

Figure 2 – Identity and stability of the marginal investor in the taxable income distribution over time



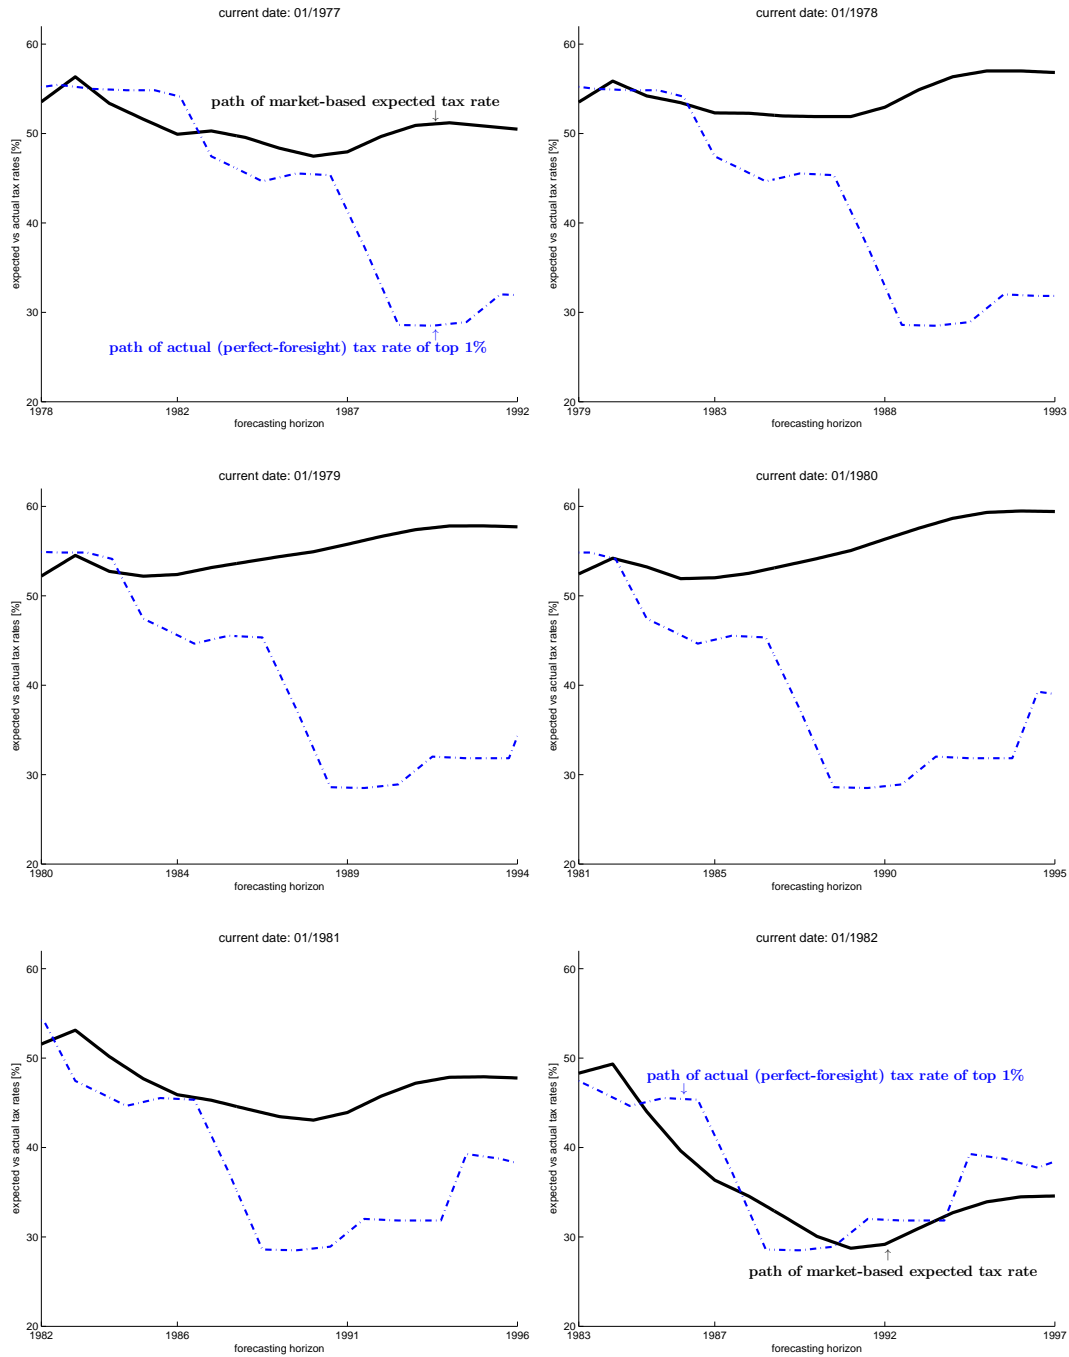
Notes: The blue dots are the estimated marginal tax rate of the marginal investor defined as the asset-weighted average of the effective marginal tax rates over all households that own both taxable and tax-exempt bonds. Two standard error bands are shown around the point estimates of the marginal investor's marginal tax rate. The black lines are the marginal tax rates of different percentiles of the income distribution taken from [Saez \(2004\)](#).

Figure 3 – Response of the BETR and path of expected tax rates to election probabilities in the two quasi-natural experiments




Notes: *Panels (a) and (b)*: The black lines are the estimated response of the BETRs $\hat{\beta}$ in equation (6) to changes in the election probability of George W. Bush and Bill Clinton, respectively, and 95% confidence bands. The blue lines show the population coefficients β^{pf} one would obtain under perfect foresight. *Panels (c) and (d)*: The black line is the expected tax path recovered by inverting equation (7) and adding the top tax rate during the election year.

Figure 4 – Anticipation of the Reagan tax cuts (ERTA 1981 and TRA 1986) between January 1977 and January 1982.




Notes: This figure shows the evolution of the path of expected tax rates between 1977 and 1982 in the run-up to the first Reagan tax cut (ERTA 1981). The dashed line, which represents the perfect-foresight tax path, is the marginal tax rate of the top 1% of the income distribution taken from [Saez \(2004\)](#). The bond market did not anticipate the Reagan tax cuts until the election year of 1980. However, the bond prices already incorporate the second Reagan tax cut (TRA 1986) by the end of 1981. The web appendix of this paper (www.lorenzkueng.com) contains an animation of the evolution of $\mathbb{E}_t[\tau]$ from January 1977 to August 1982 that shows monthly changes in the path of expected tax rates.

Figure 5 – Average tax rates by income percentiles, 1980-2007



Source: Internal Revenue Service's Statistics of Income Division. This figure shows the time series of effective federal average tax rates between 1980 and 2007 for households in different percentiles of adjusted gross income (AGI). It motivates the definition of high-income households used in the paper and the external validity checks using households with incomes in the top 10th percentiles and households with incomes below the median. The average tax rates within the top quartile are highly but not perfectly correlated. The average tax rate of households with incomes below the median is not significantly correlated with the average tax rates of high-income households.

Figure 6 – Robustness: Different decompositions of BETR into news and noise



Notes: The thin black line is the unfiltered 2-year BETR. BP denotes Baxter-King band-pass filters; see [Baxter and King \(1999\)](#). The filter renders the series stationary and passes frequencies lower than 2, 8, and 1/2 years, respectively. HP denotes the optimal one-sided Hodrick-Prescott filter ([Hodrick and Prescott \(1980\)](#)) based on [Ravn and Uhlig \(2002\)](#).