The Supply Elasticity of Municipal Debt: Evidence from Bank-Qualified Bonds

Travis St.Clair*

November 10, 2022

Abstract

This paper provides estimates of the supply elasticity of municipal debt and explores heterogeneity across different types of local governments by exploiting a discrete jump in interest rates created by the Tax Reform Act (TRA) of 1986. In order to qualify for bank financing of tax-exempt debt, governments can issue no more than $10 million of nominal debt per year. Using bunching methods, I quantify the intensive margin responses to the notch for counties, municipalities, townships, special districts, and school districts. The estimates indicate that the average marginal bunching government lowers its borrowing by approximately 5 percent in response to a 8-17 percent increase in interest costs, implying an overall price elasticity of -0.3 to -0.6. The behavioral response of special purpose governments is nearly twice as large as that of general purpose governments. The results have implications for the optimal financing of public infrastructure.

JEL Codes: H74, H76, H21

Keywords: tax-exempt debt, infrastructure, municipal finance

*Robert F. Wagner School of Public Service, New York University, 295 Lafayette Street, 2nd Floor, New York, NY 10012 (email: travis.stclair@nyu.edu). I would like to thank Tatiana Homonoff, Andy Haughwout, Ben Hyman, and Kate Yang for helpful feedback on earlier drafts of this paper. I also thank conference and seminar participants at the Federal Reserve Bank of New York and the National Tax Association. All errors are my own.

1 Introduction

The United States faces an enormous gap between the infrastructure needed to support economic growth and its current rates of spending. According to the American Society of Civil Engineers, the "infrastructure gap" – the difference between total funding on infrastructure and the investment needed to main a state of good repair – stands at approximately $2.6 trillion dollars and continues to grow (American Society of Civil Engineers, 2021). By 2039, the costs resulting from aging utilities and inadequate transportation infrastructure will cost the average American household approximately $3,300 a year. The scope of these challenges call for creative policy solutions from all levels of government.

One of the primary ways that the federal government has attempted to stimulate capital investment is through the federal tax exemption on municipal bonds. In place since the federal income tax was enacted in 1913, the "muni exemption" is typically justified on the grounds that it stimulates infrastructure investment by lowering the cost of borrowing for state and local governments. Due to restrictions on the ability of state and local governments to finance operating expenditures with debt (McNichol and Mazerov, 2020; Gordon and Metcalf, 1991), the overwhelming majority of municipal debt is issued for the acquisition and renovation of public infrastructure (Marlowe, 2015). Consequently, the federal government's ability to stimulate infrastructure investment depends crucially on the supply elasticity of municipal debt, the extent to which state and local governments increase their borrowing – and by extension, their capital spending – in response to borrowing subsidies. Despite the importance of this parameter, there are few existing estimates in the research literature, likely due in large part to data limitations as well as a lack of plausibly exogenous variation in tax-exempt interest rates.[1]

In this paper, I estimate the supply elasticity of municipal debt by exploiting a discrete jump in interest rates created by the Tax Reform Act (TRA) of 1986. Prior to 1986, commercial banks were the largest investor group in the municipal bond market, holding 39 percent of all

[1] An older set of papers investigated the efficiency of the muni exemption, however these papers primarily rely on panel methods rather than quasi-experimental variation (Gordon and Metcalf, 1991; Metcalf, 1993; Holtz-Eakin, 1991; Coronado, 1999; Poterba and Ramirez Verdugo, 2011). Joulfaian and Matheson (2009) is the only paper that I am aware of that directly studies the supply elasticity. They use fixed effects models to conclude that a one percentage point drop in interest rates is associated with an increase in bond issuance of $8.7 billion (2009 dollars).

outstanding issues (Marlin, 1994). The TRA removed the ability of banks to take advantage of tax-exempt bond interest, causing banks' demand for tax-exempt securities to plummet. The TRA did however preserve deductibility for a certain class of securities; banks could still deduct 80% of the carrying cost of securities designated as "bank-qualified." In order to meet the requirements for bank qualification, the security had to be issued by a "qualified small issuer," an issuer that does not reasonably expect to issue more than $10 million in tax-exempt obligations during the year. This in effect created a debt notch; governments issuing less than $10 million per year would be able to reap interest rate savings because of the demand from commercial banks, while those above the notch would not. The notch induces some government borrowers who would otherwise borrow in excess of $10 million to instead bunch at the limit. Figure 1 provides evidence for the behavioral response, showing how the density distribution of tax-exempt borrowing is distorted at the $10 million threshold. Governments bunch to one side of the limit, creating excess mass below the notch and a region of missing mass above it.

To estimate the supply elasticity, I combine estimates of the average behavioral response to the notch, obtained through standard bunching methods, with an estimate of the interest cost differential at the notch. To estimate the average behavioral response, I first quantify the extent of bunching. I then use standard assumptions to translate the intensive margin response into an estimate of the amount of debt foregone by the average buncher (Chetty et al., 2011; Saez, 2010; Kleven and Waseem, 2013; Kleven, 2016). I repeat this procedure for all five types of local governments - counties, municipalities, townships, special districts, and school districts. To estimate the interest cost differential at the notch, I pursue two different approaches. The first uses a difference-in-differences approach to compare governments that were and were not exposed to bank financing by exploiting a temporary increase in the small issuer limit that occurred in 2009-2010. The second approach uses a donut estimator to model the distribution of interest costs at the notch, in the spirit of a regression discontinuity (RD) design, while excluding observations in a narrow band around the threshold to account for the selection bias that would problematize the standard RD approach.

I find that the excess mass represents approximately 0.4 percent of all governments in the sample. This is equivalent to the average government operating along the intensive margin lowering their debt issuance by 4.8 percent, or approximately $500,000, in response to the notch. This response, however, varies across different types of governments, with municipalities lowering their debt issuance by 3.7 percent and special districts lowering by 6.8 percent. I estimate the interest cost differential at the notch to be on the order of 8-17 log points. Combining these two estimates together yields an overall supply elasticity of -0.3 to -0.6, indicating that governments in the vicinity of the bank qualification threshold are fairly inelastic in their response to lower borrowing costs. Special purpose governments are approximately 30 percent more price elastic than general purpose governments.

This paper build on a strand of literature in public economics that exploits bunching at kinks or notches created by the tax code to estimate policy-relevant elasticities. While most of the early literature studied the individual income tax schedule (Chetty et al., 2011; Saez, 2010; Kleven and Waseem, 2013), researchers have since extended the scope of inquiry to examine the behavioral responses of private firms (Liu et al., 2019; Chen et al., 2018) and nonprofits (Marx, 2018; St. Clair, 2016). This paper is among the first to examine bunching among governments.

In concurrent work, Dagostino (2022) also studies the small issuer threshold and demonstrates bunching in response to the notch. Whereas this paper is focused primarily on the supply elasticity and its implications for infrastructure investment, Dagostino (2022) focuses on the macroeconomic multiplier effects of the change in borrowing. In order to understand the extent to which lower financing rates stimulate investment among government borrowers, this paper estimates the interest cost differential at the notch and calculates elasticities across different types of governments. In contrast, Dagostino (2022) uses an instrumental variable strategy to estimate the effect of bank financing on employment growth, finding that every million dollars of extra bank-financed spending generates around 25 jobs per year in the private sector.

This paper proceeds as follows. Section 2 provides background on bank-qualified bonds, while section 3 provides a conceptual framework for understanding the small issuer threshold. Section 4 describes the data. Section 5 discusses the bunching methods and provides estimates of the excess mass and the behavioral response to the notch. Section 6 investigates the interest cost differential at the notch. Section 7 combines these two sets of results to provide estimates of the supply elasticity of borrowing. Section 8 concludes.

2 Background on Bank-Qualified Bonds

Since the federal income tax was enacted in 1913, the interest on state and local bonds has been excluded from taxation. Prior to 1986, commercial banks were among the largest holders of tax-exempt obligations, holding approximately 39 percent of outstanding municipal issues (Marlin, 1994). The Tax Reform Act of 1986 (TRA) significantly scaled back the deductability of the interest expense on a bank's own borrowings in an amount proportional to the interest it receives on tax-exempt bonds, effectively preventing banks from taking advantage of tax-exempt bond interest. However, the Act carved out an exception for securities designated as "bank qualified." The exception allows banks to continue to deduct 80% of the carrying costs of the tax-exempt securities; however, in order for bonds to be bank-qualified, they must 1) not be private activity bonds, 2) be issued by a qualified small issuer, 3) issued for a public purpose, and 4) designated as qualified tax-exempt obligations. Importantly, qualified small issuers were defined as issuers that reasonably expect to issue no more than \$10 million of tax-exempt obligations during the calendar year. As a result of the TRA, the demand by commercial banks for tax-exempt securities declined considerably, with holdings decreasing from approximately \$235 billion in 1985 to \$99 billion by 1992 (National Association of Bond Lawyers, 2017).

Since 1986, the demand by commercial banks for tax-exempt securities has been almost entirely limited to bank-qualified bonds. These provisions remained in place until 2009, when the American Recovery and Reinvestment Act (ARRA) temporarily raised the qualified small issuer limit from \$10 million to \$30 million for obligations issued in 2009 or 2010.

The interest rate savings to issuers depends on the spread between private and tax-exempt bonds, however the Government Finance Officers Association has estimated that the interest rate differential is equivalent to 25-40 basis points (Government Finance Officers Association, 2020). The savings make it clearly beneficial for governments issuing debt around \$10 million dollars, but even for governments planning on larger issues, there are potential advantages to splitting the issue, assuming that the additional costs of issuance or the risk of interest rate swings do not outweigh the savings. In 2019, two members of the House Committee on Ways and Means introduced The Municipal Bond Market Support Act of 2019, which would permantly increase the the annual limit from \$10 million to \$30 million and require it to be adjusted for inflation.

3 Conceptual Framework

In equilibrium, governments will issue debt to finance a preferred level of investment until they are indifferent between financing the remaining costs through borrowing or through taxation (Gordon and Metcalf, 1991). Assuming a balanced budget requirement, the government's budget constraint is $g = t + d - c$, where g represents government expenditure, t is the current level of taxation, and $d - c$ represents new debt issuance net of the end of period cost. Capital investment, k, is financed by debt and a portion of current taxes: $k = d + t_k$. Operating expenditures, o, are financed by the remaining portion of taxes: $o = t - t_k$. The amount of capital investment, k, will be equal to $k = d + o - t$, or in other words, new borrowing plus any residual operating surplus that is allocated to capital investment.[2]

Now consider the effect of introducing a notch in the interest rate schedule at the small issuer threshold. Borrowing above this amount is ineligible for bank financing and thus subject to a higher interest rate. This leads to the new interest rate schedule where interest costs for government g issuing debt d in period t will equal

$$c_g = \begin{cases} r_g * d & \text{if d} \leq \$10 \text{ mil} \\ (r_g + \Delta r_g) * d & \text{if d} > \$10 \text{ mil} \end{cases} \tag{1}$$

where Δr is the average interest rate savings from issuing bank-qualified bonds.

Consider first governments that operate along the intensive margin, i.e. those that would borrow more if all debt were bank-financed, but that adjust the amount of their borrowing in response to the discrete jump in interest rates at the notch. Figure 2a shows how the budget constraint changes at the notch. The marginal bunching government borrows $d + \Delta d$ in the counterfactual in which all debt is bank-qualified. When borrowing above the notch is not eligible for bank-financing, it is indifferent between locating at point d_1 and locating at the notch (\$10 mil). The marginal

[2]Figures A1a and A1b show the change in capital spending and operating expenditures following the issuance of bank qualified debt, confirming that the debt issued around the bank qualification threshold is used for capital rather than operating expenditures.

buncher that moves to the notch issues less debt but also faces lower interest payments.

Thus, the amount of observed/reported debt issued, d, is equal to

$$d = \begin{cases} d^* & \text{if d} \leq \$10 \text{ mil} \\ d^* - \Delta d, \text{where } \$10 \text{ mil} \leq d^* \leq \$10mil + \Delta d & \text{if d} = \$10 \text{ mil} \\ d^* - \gamma, \quad \text{where } d^* > \$10mil + \Delta d & \text{if d} > \$10 \text{ mil} \end{cases} \tag{2}$$

where d^* is the amount of debt the government would issue in the counterfactual in which all municipal borrowing is eligible for bank-financing, Δd is the amount by which governments just above the notch lower their debt issuance in response to the introduction of the notch, and γ is the marginal amount by which governments originally located above $\$10mil + \Delta d$ reduce their borrowing. In the presence of frictions, bunching governments may not locate directly at the threshold but within some interval just below it.

Figure 2b depicts the observed and counterfactual density distribution. In a world of perfect information and homogeneous elasticities, all governments originally locating within the interval (\$10 mil, $d + \Delta d$) bunch at the notch. With heterogeneous elasticities and imperfect information, not all governments adjust their borrowing, and there are some that appear in the manipulation region just above the notch. As a result, the empirical quantity of interest is the average behavioral response, $\bar{\Delta} d$, rather than the location of the marginal buncher. Under the assumption that governments only operate along the intensive margin, then excess mass below the notch will be equal in size to the missing mass above the notch (the so-called "integration constraint") (Chetty et al., 2011).

Now consider the possibility that some governments operate along the extensive margin, i.e. they are willing to borrow with bank financing but drop out of the (tax-exempt) market in the absence of bank financing.[3] In the case that the extensive margin response is non-negligible, then the missing mass above the notch will consist of governments operating along both the intensive and extensive margins. On the other hand, the excess below the notch will continue to represent only the intensive margin response and will thus be strictly smaller in size than the missing mass.

[3]Extensive margin responses might include shifting from tax-exempt to taxable debt or alternatively delaying borrowing until k years in the future such that year t+k is out of sample.

Thus, the size of the excess mass can be used to infer the average behavioral response, $\bar{\Delta d}$, of governments operating along the intensive margin.

4 Data and Summary Statistics

To conduct the empirical analysis, I use government financial data from the Census of Governments and the Annual Survey of State and Local Government Finances. For supplemental analyses, I also use bond data from Refinitiv. The Census has collected data on government revenues and debt issuance since 1967 and is "the only comprehensive source of information on the finances of local governments in the United States" (Pierson, Hand and Thompson, 2015). Every five years the Census collects a full survey of state and local governments, asking questions about the range of government financial activities (revenues, expenditures, debt, and financial assets). Census workers clean the responses and compare them to audited financial statements. In non-census years, the surveys are stratified by government type, with the probability of selection proportional to size.[4] The Census data is especially advantageous in this setting because it contains information on the total amount of debt issued by governments each year as well as the total interest expense. Because the small issuer threshold is based on the total amount of annual borrowing, data on total government borrowing is more informative than data on individual bonds.

I place two restrictions on the sample. First, because the difference-in-differences design I pursue in section 6 requires a true panel, I restrict the sample to governments that have at least seven consecutive years of observations. This removes very small governments that appear only intermittently in the data, many of whom would be unlikely to borrow on the bond market and appear in the vicinity of the notch. Second, I limit the time period to 1998-2015. Not only does this limit the number of governments with missing panel data, but it also excludes the period immediately following the TRA in which fewer governments were limited by the threshold.[5]

Table 1 provides summary statistics. Each record in the dataset represents a government's annual totals. The median government in the sample collects \$6.5 million per year (2015 dollars)

[4]See the Census of Governments for more details: https://www.census.gov/govs/local/

[5]Because the threshold is fixed in nominal terms, its real value has declined over time. Figure A2 show the extent of bunching by census year.

in own-source revenues and has $8.1 million in outstanding debt. As noted above, there are five types of governments in the data: counties, municipalities, townships, special districts, and school districts. Special districts and school districts constitute the majority of government borrowers and issue 46 percent of aggregate debt.

5 Bunching at the Small Issuer Threshold

In this section, I quantify the extent of bunching at the notch and use this to estimate the average behavioral response along the intensive margin. I estimate the average behavioral response overall as well as by type of government. Figure 1 presents the density distribution of governments near the $10 million debt notch for the period 1998-2015. The figure excludes private activity debt as well as debt issued during 2009-2010 when the ARRA temporarily raised the limit. The figure shows a sharp spike in the density distribution at a borrowing level of $10 million, consistent with governments borrowing up to a level that still enables them to qualify as "small issuers." There is no "hole" in the density distribution above the notch as would be expected in a notch analysis with homogenous elasticities (Kleven, 2016). However, this is consistent with a model of heterogeneous elasticities, in which not all municipalities are equally price sensitive in their debt issuance. There is also a small amount of round-number bunching, with smaller peaks visible at 5 and 15 million dollars.

In order to validate that governments have adjusted their borrowing in response to the notch and confirm that the bunching is not simply in response to a reference point, Figure 3 plots the density distribution of debt between 2009 and 2010 when the borrowing limit for small issuers was temporarily raised.[6] The figure shows no evidence of bunching at $10 million, confirming that the bunching observed in Figure 1 occurs in response to the discontinuity in debt costs at the threshold rather than simply in response to a reference point.

[6]The fiscal year for most state and local governments does not correspond to the calendar year, and subsequently there is a lot of partial overlap between government fiscal years and calendar years 2009-2010. This is further complicated by the fact that the "survey years" reported by the census do not always correspond to the fiscal year of the government. When referring to 2009-2010, I include only those fiscal years that fall entirely in the 2009-2010 calendar year window. Specifically, I include survey year 2010 for governments with fiscal years that end prior to July 1 and survey year 2011 for governments with fiscal years that end after June 30.

5.1 Size of excess mass

The graphical results presented in Figures 1 and 3 provide evidence of governments adjusting their debt issuance in response to the notch. In this section, I estimate the size of the excess mass and the size of the missing match at the notch. Further, I use information on the excess mass to infer the size of the behavioral response of governments that operate along the intensive margin.

First, I use a standard bunching design to quantify the extent of bunching. Borrowing the notation of Kleven (2016), I estimate a regression of the following form:

$$n_j = \sum_{i=0}^{p} \beta_i \cdot (d_j)^i + \sum_{i=d_L}^{d_U} \gamma_i \cdot 1[d_j = i] + \nu_j \tag{3}$$

where n_j is the number of governments in bin j, d_j is the level of long-term borrowing (exclusive of private activity debt) in bin j, $[d_L, d_U]$ is the excluded range, and p is the order of the polynomial. I fit a polynomial to the density distribution on both sides of the threshold, but I do not attempt to satisfy the "integration constraint" that is a common feature of bunching analyses due to the possibility of extensive margin responses. I provide estimates for both total manipulation and in-range manipulation. Total manipulation is the excess/missing mass as a percent of the total sample size. In-range manipulation is the excess/missing mass as a percent of the number of governments in the counterfactual range in the region of missing mass (bins \$10 million to d_U) (Dee et al., 2019). I estimate the standard error by block bootstrapping the entire procedure over 500 draws, sampling at the government level. The identifying assumption is that the density distribution would be smooth in the absence of the notch.

Figure 4 shows the empirical and counterfactual density distributions. The baseline specification fits an eighth order polynomial to the distribution and uses a bin size of \$500k and an excluded range of \$9.5-\$14.5 million. I constrain the choice of the excluded range such that it minimizes the difference between the excess mass and the missing mass, as per Kleven and Waseem (2013). For consistency with the model, I impose the further restriction that the excess mass not exceed the missing mass, as this would require that governments crowd in to the region below the notch. Table 2 presents measurements of the extent of bunching using a variety of specifications. In the baseline specification, the excess mass is equal to 0.36% of all governments, or alternatively 15%

of the governments just above the notch under the counterfactual. As would be expected in the presence of extensive margin responses, the missing mass is slightly larger than the excess mass: equal to 0.44% of all governments or 18% of governments in range of the notch. The estimates are fairly robust to the choice of specification, with the estimates for the in-range manipulation ranging from 0.14 to 0.15 for the excess mass and 0.15 to 0.22 for the missing mass.[7]

Based on the size of the missing mass, it is possible to infer $\bar{\Delta d}$, the average behavioral response of governments that lower their debt issuance in response to the notch. While the missing mass reflects both intensive and extensive margin responses, the excess mass is due solely to governments that reduce their debt and thus operate along the intensive margin. I calculate the average amount by which governments operating along the intensive margin lower their borrowing by measuring how far the area represented by the excess mass would extend into the counterfactual distribution above the threshold.

$$\hat{\bar{\Delta d}} = \frac{\left(\sum_{i=d_L}^{i=10Mil} \gamma_i \cdot 1[d_j = i]\right) \cdot \rho}{f^*(10Mil)} \tag{4}$$

Specifically, I multiply the number of excess organizations (represented by the summation) by ρ, the bin width, and divide by $f^*(10Mil)$, the height of the counterfactual density distribution at the notch. This calculation follows the practice in other studies of assuming that the counterfactual density distribution is approximately flat in a narrow range around the notch (Homonoff, Spreen and StClair, 2020; Marx, 2018; Kleven, 2016). Using a bin width of $500k and the estimate of the excess mass from column 1 in Table 2, the average government operating along the intensive margin lowers their debt issuance by $503k in response to the notch, or 4.8 percent.

[7]One concern with these estimates is that they don't account for the possibility that governments are splitting their issue over separate years. This sort of response would be very different than one that leads to higher total borrowing (and higher capital spending). To investigate, I explored the effects of placing an additional restriction on my sample, limiting it to government-years in which a government issued debt in year t but did not issue debt in either year t-1 or t+1. This yielded an in-range estimate of the excess mass of 12%, slightly smaller than the estimates in the main sample (14-15%). Thus, while splitting is likely the cause of some of the bunching, it does not appear to be the main driver. I return to the question of splitting in the conclusion when I discuss the generalizability of the findings.

5.2 Heterogeneity

Figure 5 shows how bunching differs according to the type of government. Table 3 shows estimates of the excess mass by type of government. While all governments show a spike in the density distribution at \$10 million, the extent of bunching is smaller among general purpose governments (counties, municipalities, townships) and larger among special purpose governments (school districts and special districts). The size of the excess mass varies from 12 percent for municipalities and townships to 21 percent for special districts; the corresponding behavioral responses vary from \$385 (3.5 percent) to \$725k (6.8 percent). Pooling general and special purpose governments together reveals that the average behavioral response for general purpose governments is 4.0 percent and 5.2 percent for special purpose governments, a difference that is significant at the 5 percent level.[8]

As Adelino, Cunha and Ferreira (2017) point out, special purpose governments, and special districts in particular, are more financially constrained than general purpose governments insofar as they are unable to increase taxes or raise revenue through alternative channels, leaving them more sensitive to the price of debt. The pattern of results in Table 3 is consistent with this reasoning. Not only are special purpose governments more responsive to the interest rate differential than general purpose governments, but even among special purpose governments, there is more bunching among special districts, who are typically authorized to provide a very limited set of functions, than school districts. In addition, since some special districts are specifically created as vehicles to issue debt, their specialized function may also play a factor in their greater responsiveness.

6 Interest Cost Differential

In order to convert an estimate of the behavioral response to an elasticity, it is also necessary to calculate the average difference in price at the notch for the marginal buncher, i.e. the average difference in cost between issuing debt with and without bank financing. This exercise is complicated by the borrower selection that occurs around the threshold, documented in the previous section. Governments in the manipulation region may have unobserved characteristics that are correlated with interest rates, thereby biasing a comparison of interest rates on either side of the notch that

[8] I calculate the standard error of the difference (0.005) using 500 bootstrapped samples.

conditions only on observables. To address this challenge, I pursue two approaches. First, I use a difference-in-difference (DiD) approach that leverages the temporary increase in the small issuer threshold in 2009-2010. Intuitively, I compare interest costs for governments that issued less than \$30 million (but more than \$10 million), pre- and post-reform, to the interest costs of governments that issued more than \$30 million. The latter group helps to establish a counterfactual of what would have happened to interest costs had the temporary increase in the notch not occurred. The assumption is that any difference in interest costs between the two groups can be attributed to bank-financing. In the second approach, I use a donut estimator that models the distribution of interest costs around the threshold in the same vein as a regression discontinuity design but excludes observations within the manipulated range. This approach uses information about the size of the manipulatation region from the bunching analysis and benefits from the fact that governments bunch in a relatively small band around the notch. In both cases I measure $\hat{\Delta c}$ as a difference in log interest costs, which is approximately equal to the percentage change in interest costs. I pursue these approaches for the sample as a whole, however I also present evidence suggesting there is not a large difference in price across types of governments.

6.1 Difference-in-Differences

The intuition behind the DiD approach is to compare the evolution in interest costs for a treatment group that is eligible to issue bank-financed debt with the evolution in interest costs for a comparison group that is not. Fortunately, the temporary increase in the small issuer threshold in 2009-2010 offers a plausibly exogenous change in the eligibility for bank financing. As part of the American Recovery and Reinvestment Act, Congress raised the cutoff for the bank-qualified designation from \$10 million to \$30 million, allowing a much larger proportion of municipal issuers to capitalize on bank financing. The change went into effect in February 2009 and expired nearly two years later on December 31, 2010. Thus, governments that would not previously have been able to issue bank-qualified debt at their preferred level of borrowing were able to do so for a short window of time.

To exploit the temporary increase, I compare interest costs among governments that borrowed less than \$30 million (but more than \$10 million), both prior to as well as "post" reform (i.e. during

2009-2010), with governments that borrowed more than \$30 million during the same periods.[9] To measure the cost of debt (c), I calculate the difference between log interest paid in the year that the debt is issued and log interest paid in the subsequent year ($c_t = log\_interest_{t+1} - log\_interest_t$). Because in some cases government may pay interest on debt in the same year that it issues the debt, I also include alternative specifications in which I measure the cost instead as ($c_t = log\_interest_{t+1} - log\_interest_{t-1}$).

I then estimate a DiD design of the following form:

$$c_{gt} = \alpha + \beta_1 \cdot Treat_g + \beta \cdot Post_t + \gamma \cdot Treat_g \cdot Post_t + \psi_{gt} + \theta_g + \epsilon_{gt} \tag{5}$$

where c_{dt} represents the log interest cost for government g issuing debt in period t, $Treat_g$ represents governments that issue less than \$30 million of debt pre- and post- reform, and $Post_t$ represents the period in which the threshold was temporary increased (2009-2010).[10] The right hand side also includes vectors of time-varying, ψ_{gy}, and fixed, θ_g, covariates. The covariates include the amount of (log) debt issued, log capital spending, log total debt outstanding, and indicator variables for the type of government (school district, etc.). Standard errors are clustered at the level of the government.

The DiD results (equation 5) are presented in Table 4. The first two columns show results when the dependent variable is measured as the difference in interest costs between year t+1 and year t. The third and fourth columns show results when the dependent variable is measured as the difference in interest costs between year t+1 and year t-1. If the first interest payment of a bond is due in the same fiscal year in which it is issued, then measuring the increase in interest costs from the bond as $log\_interest_{t+1} - log\_interest_t$ may understate the true increase. On the other hand, if a government issues debt every year, then measuring the increase in interest costs as $log\_interest_{t+1} - log\_interest_{t-1}$ may overstate the true costs. However, so long as these differences are fixed across the treatment and comparison groups, this measurement error need not bias the results.

[9] I use survey year 2007 as the pre-reform year because it was a full-census year and thus affords more observations.

[10] As in footnote 6, I specifically use survey year 2010 for governments with fiscal years that end prior to July 1, and survey year 2011 for governments with fiscal years that end after June 30.

The estimates in columns 1-2 imply that bank-financing lowers interest costs by 7-9 log points, which is approximately equal to a decrease of 7-9 percent. The estimates in columns 3-4, which use a slightly different measure of the dependent variable, imply that bank-financing lowers interest costs by 14-16 log points, or approximately 14-16 percent. To put these numbers in perspective, for a tax-exempt bond with a coupon rate of 3%, a 12% decrease in interest costs (the mid-point of the DiD estimates) is equivalent to 36 basis points. This is within the range of 25-40 basis points assumed by the Goverment Finance Officers Association (Government Finance Officers Association, 2020).

6.2 Donut Estimator

As an alternative approach, I also estimate the interest cost advantage of bank-financing by using a donut-RD estimator. The donut approach excludes observations in the manipulated region in order to address the selection bias that would result under the standard regression discontinuity approach (Barreca et al., 2011; Barreca, Lindo and Waddell, 2016). It has the advantage of offering some of the transparency of the standard RD design while also utilizing information on the size of the manipulation region as revealed through the bunching estimation in section 5. Moreover, in this instance, because bunching is confined to a rather narrow region (at least below the threshold), the approach has more credibility than if the excluded region were larger. In addition, the panel nature of the data affords the use of fixed effects, which is useful for addressing unobserable characteristics that are correlated with interest rates and also fixed over time. I measure interest costs as log interest in year t+ 1 (the year following a debt issue) since governments may not pay the full interest expense until the year after it issues debt.

I estimate regressions of the following form:

$$c_{gt} = \alpha + f(b) + Small_{gt} + f(b) \cdot Small_{gt} + \psi_{gt} + \eta_g + \delta_t + \epsilon_{gt} \tag{6}$$

where c_{dt} represents total log interest costs for government g in year t+1, $f(b)$ represents a polynomial function in the amount of borrowing in year t, $Small_{gt}$ represents an indicator variable for a government falling under the small issuer threshold in year t, ψ_{gt} represents a set of time-varying covariates, η_g represents government fixed effects, δ_y represents year fixed effects, and ϵ_{gt} is the error term. In the baseline specification, I use linear polynomials, estimated separately on both sides of the threshold. Importantly, observations within the manipulated region are excluded.

Figure 6 plots residuals from the baseline specification that includes linear polynomials, year and government fixed effects, but no covariates. Each circle represents the average amount of borrowing within bins of $500,000. The figure omits only one bin on either side of the threshold (representing the range $9.5-$10.5 million). The figure indicates a discontinuity in log interest costs of approximately 10 log points. Although not the focus of the analysis, the figure also indicates a change in slope at the notch, suggesting that interest costs rise at a slower rate as the amount of principal increases among bonds that are bank-financed. The results from donut estimation are presented in Table 5. The specifications in the table vary the size of the excluded region as well as the order of the polynomials.

The results using a linear functional form indicate that bank financing confers a cost advantage of 8-17 log points. When the excluded region is limited to $9.5-12 million in debt, the estimates range from 8-11 log points. With a wider excluded region, the estimates increase slightly to 13-17 log points. The specifications using a quadratic polynomial indicate a differential of at least twice this size (20-40 log points), but as these specifications show a poor fit to the data (and do not yield statistically significant coefficients), I do not place much weight on the estimates. Thus, the results from the donut estimator indicate an interest cost differential of 8-17 log points, in line with the DiD estimates above.

For a robustness check, Figure A3 uses bond data from Refinitiv to plot yield as a function of the amount of borrowing. Because small issuers cannot issue more than $10 million of tax-exempt debt per year, the data have been aggregated to the level of the issuer-year, as in Figure 6. The figure shows a discontinuity at the bank qualification threshold of approximately 40 basis points, providing further support for an estimate in the range of 12-13% of borrowing costs.

Neither the DiD nor the donut approach are without flaws. In particular, the difference-in-difference approach assumes parallel trends, which cannot be tested in this context. The extremely small number of governments that issue between $10-$30 million debt every year over the period 2006-2010 precludes a proper panel analysis covering a wider range of years. The donut estimator departs from standard RD assumptions by dropping observations near the cut-off point. Nevertheless, both methods use a variety of different specifications to arrive at similar estimates of the interest cost differential: approximately 8-17 log points.

6.3 Heterogeneity

The relatively small number of observations of certain types of governments make it challenging to estimate a price differential separately for each type of government. In particular, the DiD estimator does not have sufficient power to explore heterogeneity. However, because the donut estimator pools observations from multiple years, there is enough data to obtain a rough estimate of whether bank financing confers a greater interest rate advantage on certain types of governments. Table 6 presents donut RD results for general purpose and special purpose governments separately (using the same specification as in column 1 of table 5). The results show a differential of 10 percent for general purpose governments and 12 percent for special purpose governments. The difference between the estimates is not statistically significant. Although the analysis suffers from a lack of power, due to the similarity of the estimates I assume that the interest rate differential associated with bank financing is equivalent across all types of government in the calculation of elasticities below.

7 Elasticities

7.1 Estimates

In this section, I use the results from the previous two sections to estimate the supply elasticity of municipal debt. From section 5, I use the average debt response of the marginal buncher, which I measure as a percentage change. From section 6, I use the interest cost differential at the notch, which is equivalent to the change in price facing the marginal buncher. Note that the estimates in Section 6 are not estimates of a change in interest *rates*; they are estimates of the percentage change in interest *costs*. This simplifies that analysis because, unlike DeFusco and Paciorek (2017), I do not need to differentiate between an average change in interest rates and the marginal cost facing the marginal bunching borrower. I calculate the elasticity as

$$\epsilon = \frac{\hat{\Delta d}}{\hat{\Delta c}} \tag{7}$$

Table 7 reports the elasticities for a range of estimates of $\hat{\Delta d}$ and $\hat{\Delta c}$. Each elasticity is calculated from the estimate of $\hat{\Delta c}$ at the top of that column and the estimate of $\hat{\Delta d}$ reported at the beginning of that row. Since the estimates of the excess mass in Table 2 are so consistent, I use only one value for $\hat{\Delta d}$, but I vary the estimates of $\hat{\Delta c}$ from a low of 0.081 to a high of 0.171, based on the specificiations that yield statistically significant coefficients. The elasticity estimates range from -0.28 to -0.59, implying that a one percent increase in debt costs results in a reduction of municipal debt supply of 0.3 to 0.6 percent.

7.2 Refunding

The elasticity that is most relevant to policy is the supply elasticity of new debt. This provides information to policymakers about new projects undertaken for capital purposes. If bunching in response to the small issuer threshold primarily reflects debt issued for refunding purposes, then it would suggest that new debt may even be less elastic than suggested by the elasticities presented above. This is a potential concern, as refunding obligations can be designated as bank-qualified.

To investigate the extent to which the borrowing at the notch reflects refunding obligations as opposed to new debt, Figure 7 plots the density distribution of bank qualified debt, separating out refunding debt from new issues. The data come from Refinitiv and include bank qualified debt aggregated to the level of the issuer-year. While the density distribution of refundings appears flatter than the overall distribution due to the smaller number of refundings at lower dollar amounts, in the vicinity of the notch the proportion of total debt that represents refundings remains fairly constant at 40-50%, i.e. the extent of bunching among refundings does not appear noticeably different than in the overall sample. This suggests that refunding and bank qualification are approximately independent of each other. Thus, while refundings represent a nontrivial share of bank qualified debt, especially in the region of the notch, they do not appear to be driving the elasticity results presented in the previous section.

7.3 Heterogeneity

Table 8 reports the elasticities by type of government. Based on the finding in Table 6, I assume that all governments face the same interest rate differential at the notch, and thus the variation in elasticities arises from variation in the behavioral response alone. The elasticities for general purpose governments range from -0.28 for townships and municipalities to -0.33 for counties, for an average of -0.30 across all general purpose governments. The elasticies for special purpose governments are -0.37 for school districts and -0.51 for special districts, for an average of -0.39. As noted above, special purpose governments may be more price sensitive because they are financially constrained in a way that general purpose governments are not; they are limited in their ability to increase taxes or raise revenue through alternative means. In addition, some special districts have functions that are more narrowly focused on debt issuance.

8 Conclusion

This paper estimates the supply elasticity of municipal debt by exploiting a discrete jump in interest rates created by the Tax Reform Act of 1986. First, I document bunching at the small issuer threshold, both in aggregate and across different types of governments. I calculate the size of the excess mass at the notch and use this information to infer the average reduction in borrowing for the marginal borrower, concluding that the margin buncher reduces their borrowing by $500,000, or approximately 5 percent. This behavioral response varies between 3.7 percent for municipalties to 6.8 percent for special districts. Next, I calculate the interest cost differential at the notch using two different approaches, one that exploits the temporary suspension of the notch and another that models the distribution of interest costs around the notch. These approaches yield estimates of 8 to 17 log points for the interest rate differential. Finally, I combine these two estimates to calculate the price elasticity of municipal debt supply. The results indicate that local governments lower their debt supply by 0.3-0.6 percent in response to a 1 percent increase in borrowing costs. General purpose governments in particular are fairly insensitive to price, with an elasticity in the range of -0.3.

Recent work on the muni market has reached mixed conclusions regarding the responsiveness of debt supply. While Adelino, Cunha and Ferreira (2017) find a large supply response to Moody's recalibration of its ratings scale in 2010, Garrett et al. (2017) find no quantity response to large tax exemption induced decreases in borrowing costs. This paper produces findings more in line with those of Garrett et al. (2017), however it departs from these other works by virtue of its identification strategy and its focus on smaller governments near the bank qualification threshold. For governments near the threshold, bank financing does not appear to be an efficient subsidy; governments are not responding strongly to the favorable interest rates offered by bank financing. More generally, the results in this paper cast doubt on the efficiency of federal tax subsidies aimed at stimulating borrowing among smaller governments.

This conclusion is further bolstered when one considers additional factors highlighted above suggesting that the elasticity estimates are, if anything, overstated. To the extent that bunching at the notch results from governments' splitting their debt issues across years rather than decreasing the amount of their borrowing, then the true elasticity of new debt may be even lower. Moreover, issuance costs and other book building expenses are likely lower when placing debt with a bank rather than a non-bank investor, suggesting that the interest rate differential measured above, if anything, understates the total cost differential, and again biases the elasticity upward.

Despite the limited responsiveness overall, the results do point to some important sources of heterogeneity. Special districts in particular are much more responsive to the notch, reducing their debt issuance by 6.8 percent, compared to 3.7 percent for municipalities. Moreover, Figure A2 and Table A1 show that the extent of bunching at the notch has nearly doubled since it was introduced in 1986 as the real value of the notch has fallen. It's unclear whether this change is due to the increase in the number of special districts or other factors, but the results are consistent with a model in which institutions that facilitate bunching, such as special districts, need time to develop.

There are several limitations to this article's approach that are worth noting. First, as noted above, the results described here are local to the notch. Small municipal governments may not react the same way to price changes as larger governments with more sophisticated debt management strategies, who are, after all, responsible for a large portion of capital spending. This may explain why the findings in this paper differ from some of the previous work that has found larger elasticities in the muni market(Joulfaian and Matheson, 2009; Adelino, Cunha and Ferreira, 2017), and it suggests a potentially important source of heterogeneity among government borrowers. While most outstanding muni debt was issued by states and large municipal borrowers, most governments are small; in 2015, 85% of governments collected less than $10 million in tax per year. This underscores the potential benefits of targeting subsidies to larger governments that will be more responsive.

Given the deteriorating condition of infrastructure in the United States, it is more important than ever to understand the policy levers that are available to stimulate infrastructure investment and to do so at minimal cost to the US taxpayer. The municipal debt market remains central to the ability of subnational governments to finance infrastructure investment, but market participants are extremely heterogeneous. Understanding the causes and consequences of this heterogeneity will be central to devising optimal policy responses in the future.

References

Adelino, Manuel, Igor Cunha, and Miguel A Ferreira. 2017. "The economic effects of public financing: Evidence from municipal bond ratings recalibration." *The Review of Financial Studies*, 30(9): 3223–3268.

American Society of Civil Engineers. 2021. "2021 infrastructure report card."

Barreca, Alan I, Jason M Lindo, and Glen R Waddell. 2016. "Heaping-induced bias in regression-discontinuity designs." *Economic inquiry*, 54(1): 268–293.

Barreca, Alan I, Melanie Guldi, Jason M Lindo, and Glen R Waddell. 2011. "Saving babies? Revisiting the effect of very low birth weight classification." *The Quarterly Journal of Economics*, 126(4): 2117–2123.

Chen, Zhao, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Yi Xu. 2018. "Notching R&D investment with corporate income tax cuts in China." National Bureau of Economic Research.

Chetty, Raj, John N Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records." *The quarterly journal of economics*, 126(2): 749–804.

Coronado, Julia Lynn. 1999. "Tax exemption and state capital investment." *National Tax Journal*, 473–481.

Dagostino, Ramona. 2022. "The impact of bank financing on municipalities' bond issuance and the real economy." *Working Paper, University of Rochester.*

Dee, Thomas S, Will Dobbie, Brian A Jacob, and Jonah Rockoff. 2019. "The causes and consequences of test score manipulation: Evidence from the New York regents examinations." *American Economic Journal: Applied Economics*, 11(3): 382–423.

DeFusco, Anthony A, and Andrew Paciorek. 2017. "The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit." *American Economic Journal: Economic Policy*, 9(1): 210–40.

Garrett, Daniel, Andrey Ordin, James W Roberts, and Juan Carlos Suárez Serrato. 2017. "Tax advantages and imperfect competition in auctions for municipal bonds." National Bureau of Economic Research.

Gordon, Roger H, and Gilbert E Metcalf. 1991. "Do tax-exempt bonds really subsidize municipal capital?" *National Tax Journal*, 71–79.

Government Finance Officers Association. 2020. "Bank-qualified debt." *https://www.gfoa.org/bank-qualified-debt*.

Holtz-Eakin, Douglas. 1991. "Bond market conditions and state-local capital spending." *National Tax Journal*, 105–120.

Homonoff, Tatiana, Thomas Luke Spreen, and Travis StClair. 2020. "Balance sheet insolvency and contribution revenue in public charities." *Journal of Public Economics*, 186: 104177.

Joulfaian, David, and Thornton Matheson. 2009. "The supply elasticity of tax-exempt bonds." *Proceedings of the Annual Conference of the National Tax Association*, 102: 136–142.

Kleven, Henrik Jacobsen. 2016. "Bunching." *Annual Review of Economics*, 8: 435–464.

Kleven, Henrik J, and Mazhar Waseem. 2013. "Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan." *The Quarterly Journal of Economics*, 128(2): 669–723.

Liu, Li, Ben Lockwood, Miguel Almunia, and Eddy HF Tam. 2019. "VAT notches, voluntary registration, and bunching: Theory and UK evidence." *Review of Economics and Statistics*, 1–45.

Marlin, Matthew R. 1994. "Did tax reform kill segmentation in the municipal bond market?" *Public Administration Review*, 387–390.

Marlowe, Justin. 2015. "Municipal bonds and infrastructure development-past, present, and future." *A Policy Issue White Paper Prepared on behalf of the ICMA Governmental Affairs and Policy Committee.*

Marx, Benjamin M. 2018. "The Cost of Requiring Charities to Report Financial Information."

McNichol, Elizabeth, and Michael Mazerov. 2020. "State borrowing no substitute for additional direct aid to help states weather covid downturn." *Center on Budget and Policy Priorities.*

Metcalf, Gilbert E. 1993. "Federal taxation and the supply of state debt." *Journal of Public Economics*, 51(3): 269–285.

National Association of Bond Lawyers. 2017. "Direct purchases on state and local obligations by commercial banks and other financial institutions."

Pierson, Kawika, Michael L Hand, and Fred Thompson. 2015. "The government finance database: A common resource for quantitative research in public financial analysis." *PloS one*, 10(6): e0130119.

Poterba, James M, and Arturo Ramirez Verdugo. 2011. "Portfolio substitution and the revenue cost of the federal income tax exemption for state and local government bonds." *National Tax Journal*, 64(2): 591–614.

Saez, Emmanuel. 2010. "Do taxpayers bunch at kink points?" *American economic Journal: economic policy*, 2(3): 180–212.

St. Clair, Travis. 2016. "How do nonprofits respond to regulatory thresholds: Evidence from New York's audit requirements." *Journal of Policy Analysis and Management*, 35(4): 772–790.

Figure 1: Bunching at the Small Issuer Threshold: 1998-2015

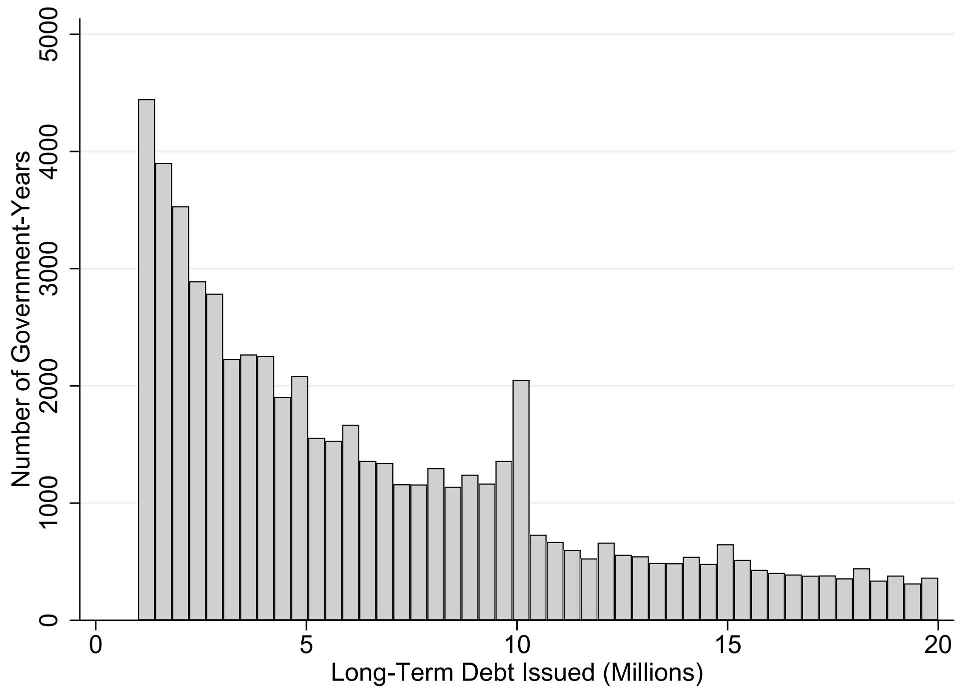

Note: The figure shows the density distribution of long-term debt between 1998-2015. The sample excludes private purpose debt as well as debt issued between 2009 and 2010 when the small issuer threshold was temporarily increased due to the American Recovery and Reinvestment Act.

Figure 2a: Budget set diagram

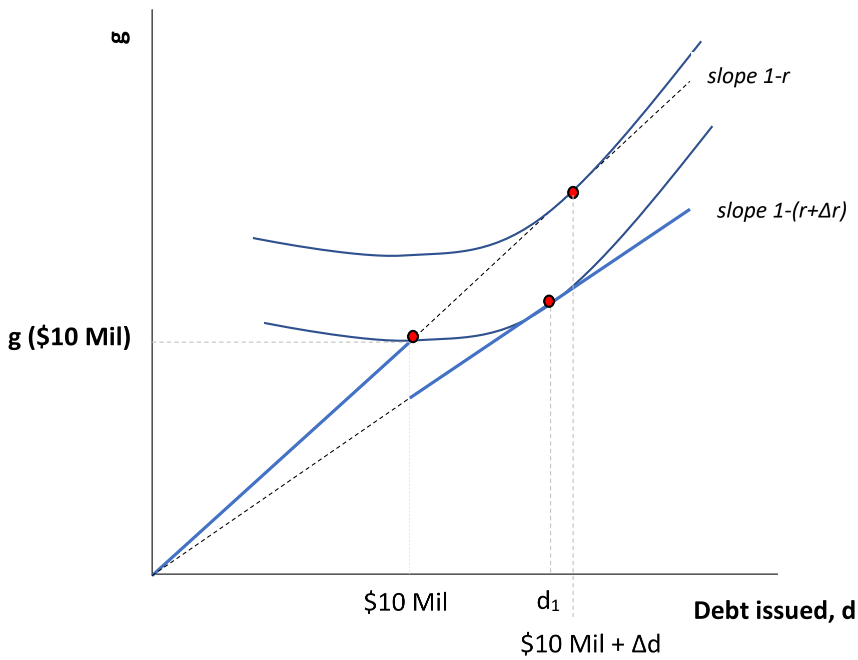

Figure 2b: Density distribution diagram

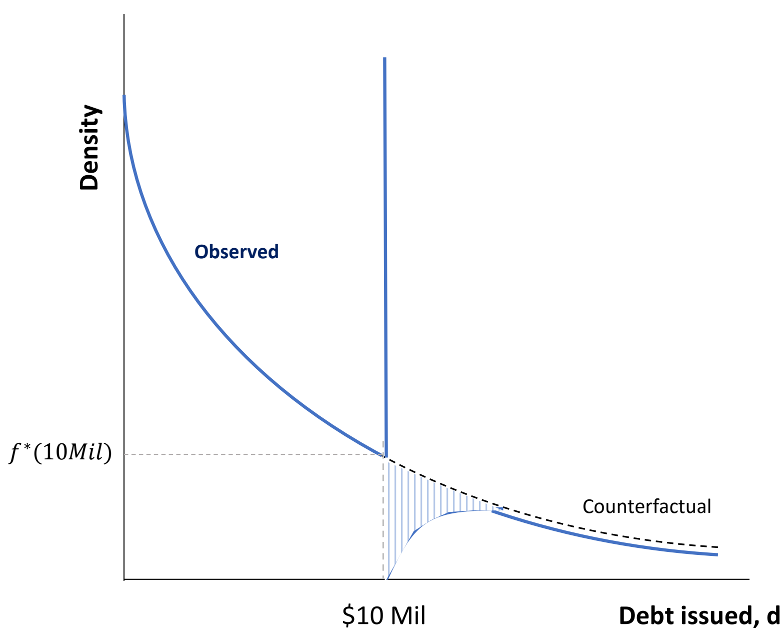

Figure 3: Long-Term Debt Issued in 2009 and 2010

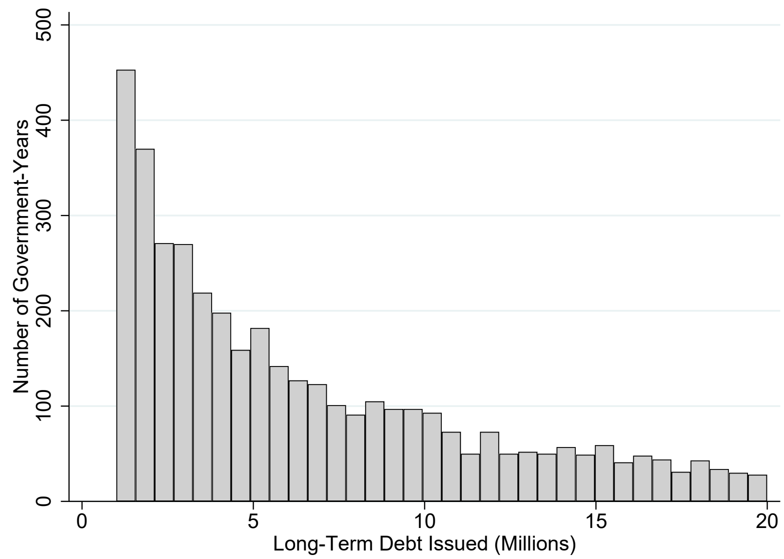

Note: The figure shows the density distribution of long-term debt in 2009 and 2010 when the ARRA temporarily raised the $10 million limit for small issuers.

Figure 4: Bunching Estimation

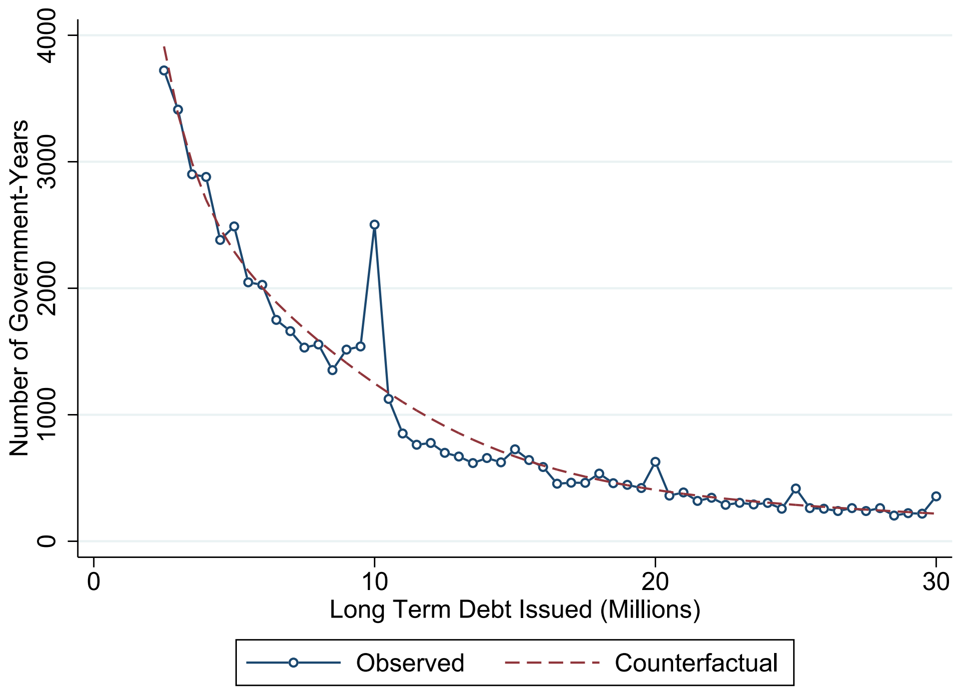

Note: The figure depicts the observed distribution of long-term debt between 1998-2015 (excluding 2009-2010), shown as the mean of bins of size $500,000, and the modeled counterfactual, based on an 8th order polynomial. The excluded range is $9.5 - $14.5 million.

Figure 5: Bunching By Type of Government

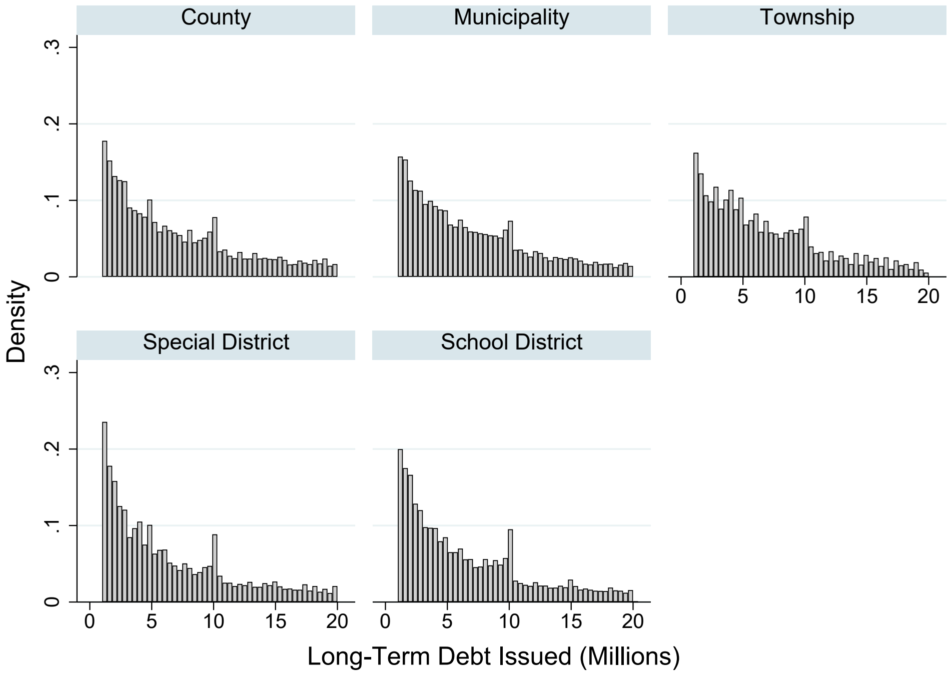

Note: Figure 5 shows the distribution of debt by type of government. The sample excludes private purpose debt as well as debt issued between 2009 and 2010 when the small issuer threshold was temporarily increased due to the American Recovery and Reinvestment Act.

Figure 6: Mean Interest Expense by Amount of Borrowing

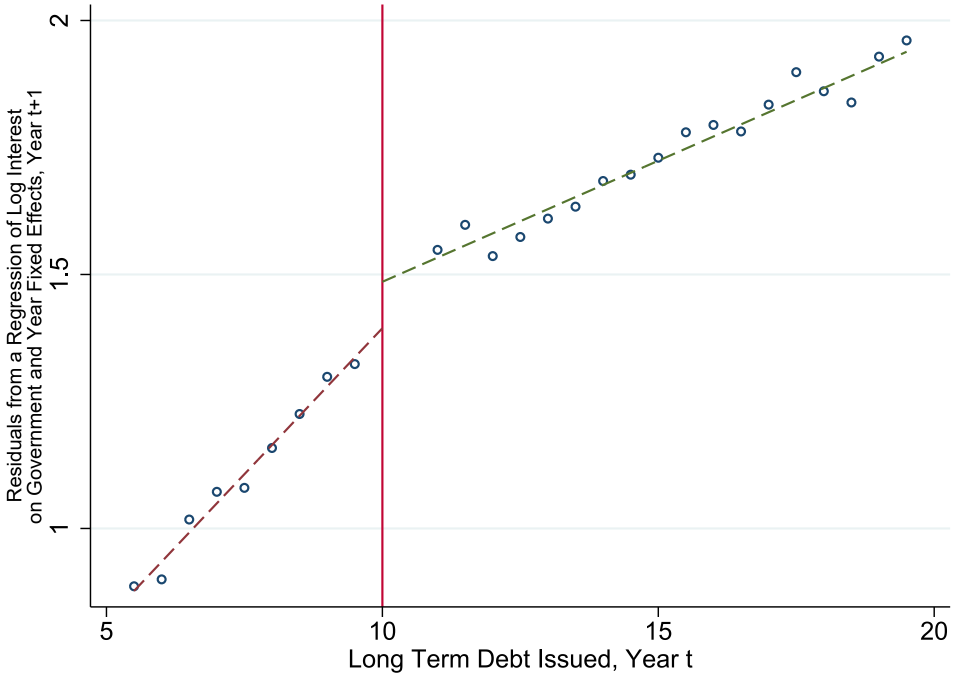

Note: The figure plots residuals from a regression of log interest on government and year fixed effects as a function of the amount of annual long-term debt issued. Each circle represents the mean amount of log interest payments within bins of $500,000. The dashed lines are predicted values from a regression fit to the binned data, allowing for changes in the slope and intercept at the $10 million threshold. One bin on either side of the threshold ($9.5-$10.5 mil) is omitted.

Figure 7: Refunding

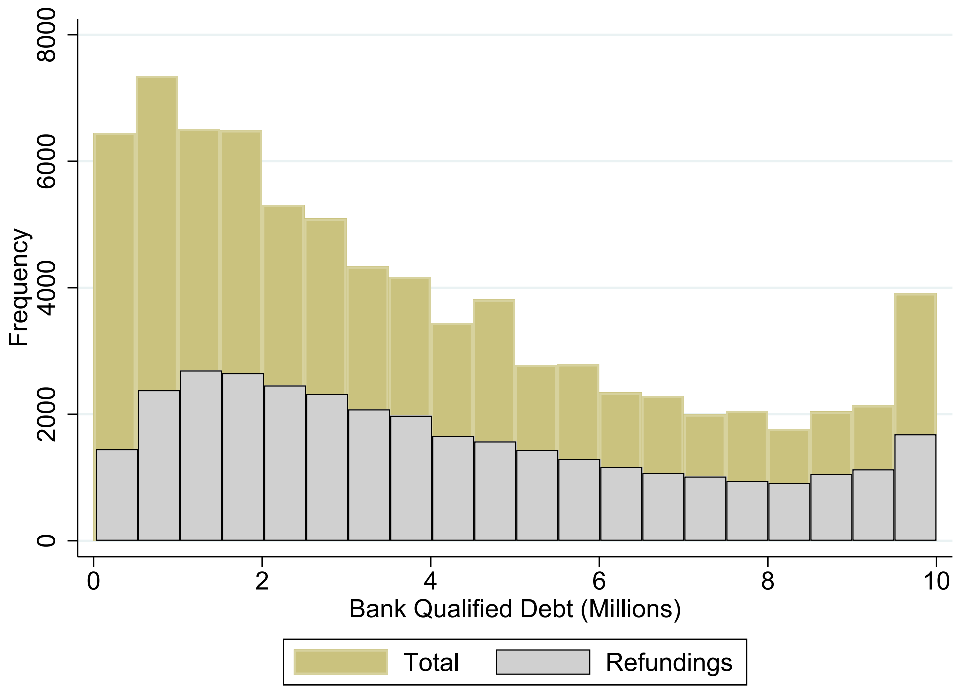

Note: The figure shows the density distribution of bank qualified debt using bond data from 1998-2015 (excluding 2009-2010), separating out refundings from the total. The data come from Refinitiv and include all tax-exempt, bank-qualified debt, aggregated to the level of the issuer-year.

Table 1: Summary Statistics

	(1)	(2)	(3)	(4)
VARIABLES	mean	sd	min	max
Long Term Debt Issued	13	151	0	26,679
Total Debt Outstanding	77	910	0	138,027
Total Interest	3.6	43	0	6,211
Total Taxes	26	307	0	52,398
Own-Source Revenues	45	512	0	96,343
Cash and Securities	70	1,284	0	215,601
Total Expenditures	78	764	0	119,203
County	0.08	0.26	0	1
Municipality	0.11	0.31	0	1
Township	0.04	0.19	0	1
Special District	0.13	0.33	0	1
School District	0.66	0.478	0	1

Note: Financial variables are in millions of 2015 dollars. Data come from the Census of Governments and the Annual Survey of State and Local Government Finances. Summary stats are for 1998-2015, excluding 2009-2010. The variable "Long Term Debt Issued" excludes private purpose debt. The sample is restricted to governments with at least seven consecutive years of observations.

Table 2: Size of Excess Mass

	(1) Total Manip.	(2) In-Range Manip.	(3) Total Manip.	(4) In-Range Manip.	(5) Total Manip.	(6) In-Range Manip.	(7) Total Manip.	(8) In-Range Manip.
Excess Mass	0.0036***	0.15***	0.0035***	0.14***	0.0036***	0.15***	0.0037***	0.14***
	(0.0002)	(0.008)	(0.0002)	(0.009)	(0.0002)	(0.008)	(0.0002)	(0.008)
Missing Mass	0.0044***	0.18***	0.0056***	0.22***	0.0043***	0.18***	0.0037***	0.15***
	(0.0006)	(0.019)	(0.0007)	(0.022)	(0.0006)	(0.019)	(0.0007)	(0.022)
Size of bins	500k		500k		250k		500k	
Polynomial order	8th		9th		8th		8th	
Excluded range	9.5-14.5 Mil		9.5 - 14.5 Mil		9.5 - 14.5 Mil		9.5-15 Mil	

Note: *** $p < 0.01$. The table presents estimates of the size of the excess mass and the missing mass. The specifications vary the size of the bins, the order of the polynomial, and the excluded range. Total manipulation is the excess/missing mass as a percentage of all governments in the sample. In-range manipulation is the excess/missing mass relative to the counterfactual distribution in the range of the missing mass. Block bootstrapped standard errors in parentheses.

Table 3: Size of Excess Mass by Type of Government

	(1) County	(2) Municipality	(3) Township	(6) All General Purpose Governments	(4) Special District	(5) School District	(7) All Special Purpose Governments
Excess Mass	0.14***	0.12***	0.12***	0.13***	0.21***	0.16***	0.16***
	(0.024)	(0.015)	(0.031)	(0.012)	(0.034)	(0.010)	(0.011)
$\Delta\hat{d}$	-0.044***	-0.037***	-0.037***	-0.040***	-0.068***	-0.050***	-0.052***
	(0.007)	(0.005)	(0.009)	(0.003)	(0.010)	(0.003)	(0.003)

Note: * $p < 0.1$, ** $p > 0.05$, *** $p < 0.01$. Table 2 presents estimates of the excess mass by the type of government. The extent of bunching is measured in terms of the in-range manipulation, the excess mass relative to the counterfactual distribution in the range of the missing mass. All estimates use bins of $500k, 8th order polynomials, and an excluded range of $9.5-14.5 million. General purpose governments include counties, municipalities, and townships. Special purpose governments include special districts and school districts. Block bootstrapped standard errors in parentheses.

Table 4: DiD Results - Interest Cost Differential

	(1)	(2)	(3)	(4)
	$\Delta_{t+1,t}$		$\Delta_{t+1,t-1}$	
Treat*Post	-0.073	-0.086*	-0.144*	-0.159**
	(0.049)	(0.048)	(0.079)	(0.080)
Covariates	No	Yes	No	Yes
N	1,042	1,042	1,044	1,044

Note: * p < 0.1, ** p > 0.05, *** p < 0.01. This table shows estimates of γ based on equation 5. The outcome variable is the change in log interent costs. Columns 1-2 measure the outcome variable as the change in interest costs between years $t + 1$and year t. Columns 3-4 measure the outcome variable as the change in interest costs between years $t+1$ and $t-1$. The covariates include the amount of (log) debt issued, log expenditures, log total debt outstanding, log own-source revenues, and indicator variables for the type of government (school district, etc.). Standard errors are clustered at the level of the government.

Table 5: Donut RD Results - Interest Cost Differential

	(1)	(2)	(3)	(4)	(5)	(6)
$Small_{gt}$	-0.113***	-0.081*	-0.209	-0.171***	-0.134**	-0.393
	(0.044)	(0.043)	(0.223)	(0.055)	(0.053)	(0.300)
Covariates	No	Yes	Yes	No	Yes	Yes
Functional Form	Linear	Linear	Quadratic	Linear	Linear	Quadratic
Excluded Region	$9.5-12 Mil	$9.5-12 Mil	$9.5-12 Mil	$9-14.5 Mil	$9-14.5 Mil	$9-14.5 Mil
N	25,090	25,083	25,083	21,035	21,028	21,028

Note: * $p < 0.1$, ** $p > 0.05$, *** $p < 0.01$. This table presents estimates of the interest rate differential at the notch based on OLS regressions on log interest in year t+1 as a function of the amount of long-term borrowing in year t, allowing for changes in slope and intercept at the $10 million borrowing threshold. The sample includes all government years between 1998-2015 (excluding 2009-2010). The regressions are estimated over the range 5-30 million in debt issued, with observations in the excluded region omitted. All specifications include government and year fixed effects. Covariates include log expenditures and log own-source revenues. Standard errors clustered by government in parentheses.

Table 6: Donut RD Results - Interest Cost Differential by Type of Government

	(1) General Purpose Governments	(2) Special Purpose Governments
$Small_{gt}$	-0.095*	-0.117*
	(0.056)	(0.063)
N	9,213	15,877

Note: * $p < 0.1$, ** $p > 0.05$, *** $p < 0.01$. This table presents estimates of the interest rate differential at the notch by type of government, using the Donut RD approach outlined in section 6.2. The specifications mirror column 1 of Table 6 and include linear polynomials, covariates, and an excluded region of $9.5-12 million. The sample includes all government years between 1998-2015 (excluding 2009-2010). The regressions are estimated over the range 5-30 million in debt issued, with observations in the excluded region omitted. All specifications include government and year fixed effects. Covariates include log expenditures and log own-source revenues. Standard errors clustered by government in parentheses.

Table 7: Supply Elasticity of Municipal Debt

	$\Delta\hat{c}$		
	(1)	(2)	(3)
	0.081	0.134	0.171
$\Delta\hat{d}$			
-0.048	-0.59	-0.36	-0.28

Note: This table reports estimates of the interest rate elasticity of municipal debt supply for a range of different estimated parameters. The three columns represent low, mid-range, and high estimates of the interest cost differential ($\Delta\hat{c}$). Each cell reports the elasticity implied by the estimated behavioral response ($\Delta\hat{d}$) and corresponding interest cost differential ($\Delta\hat{c}$).

Table 8: Elasticity by Type of Government

(1) County	(2) Municipality	(3) Township	(6) All General Purpose Governments	(4) Special District	(5) School District	(7) All Special Purpose Governments
-0.33	-0.28	-0.28	-0.30	-0.51	-0.37	-0.39

Note: This table reports estimates of the interest rate elasticity by type of government. Each cell reports the elasticity implied by the estimated behavioral response ($\Delta\hat{d}$) reported in Table 3 and the interest cost differential ($\Delta\hat{c}$),assumed to be uniform across governments and equal to the mid-point estimate from Table 6 (0.134).

Figure A1: Use of Borrowed Funds

Figure A1a: Change in Log Capital Outlays   Figure A1b: Change in Log Operating Expenditures

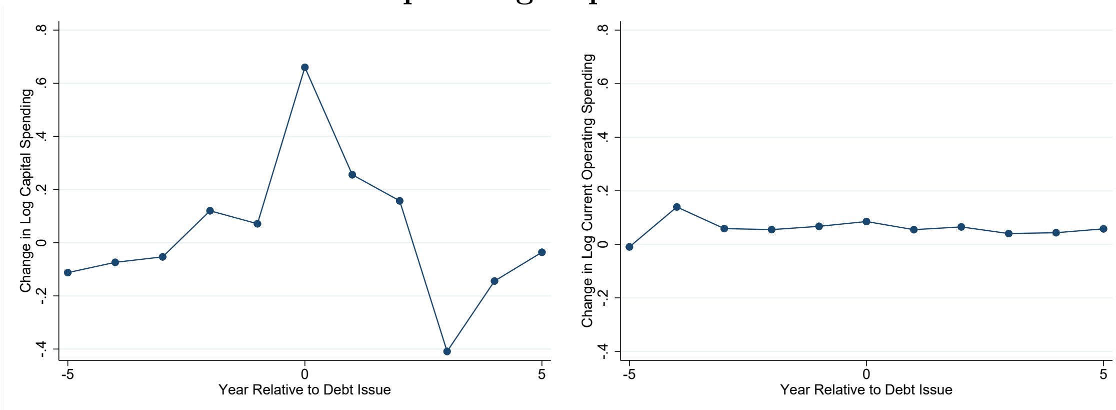

Note: Figures A1a and A1b show the change in log capital spending ($c_t = log\_capital\_spending_t - log\_capital\_spending_{t-1}$) and the change in log current operating expenditures following a debt issue. The sample is limited to government years between 1986-2008 in which a government borrowed between 9-10 million dollars and did not borrow in the previous five years. To ensure a balanced panel, the sample is restricted to governments that report observations for five years before and after the debt issue.

Figure A2: Bunching By Census Years

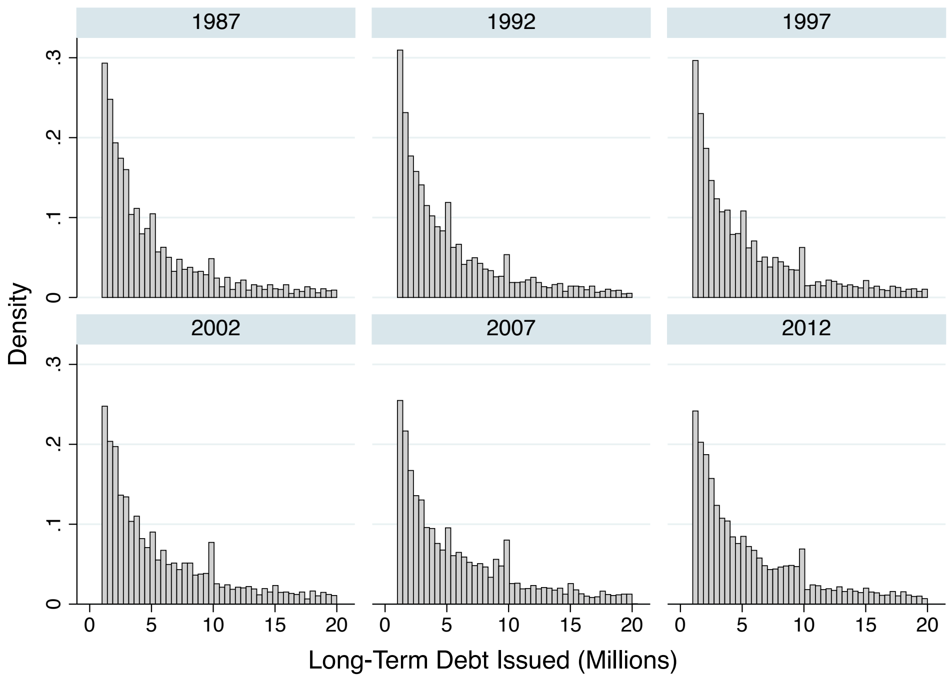

Note: Figure A2 shows the density distribution of long-term debt during the years of a full census. Excludes private activity debt.

Figure A3: Yield by Amount of Borrowing

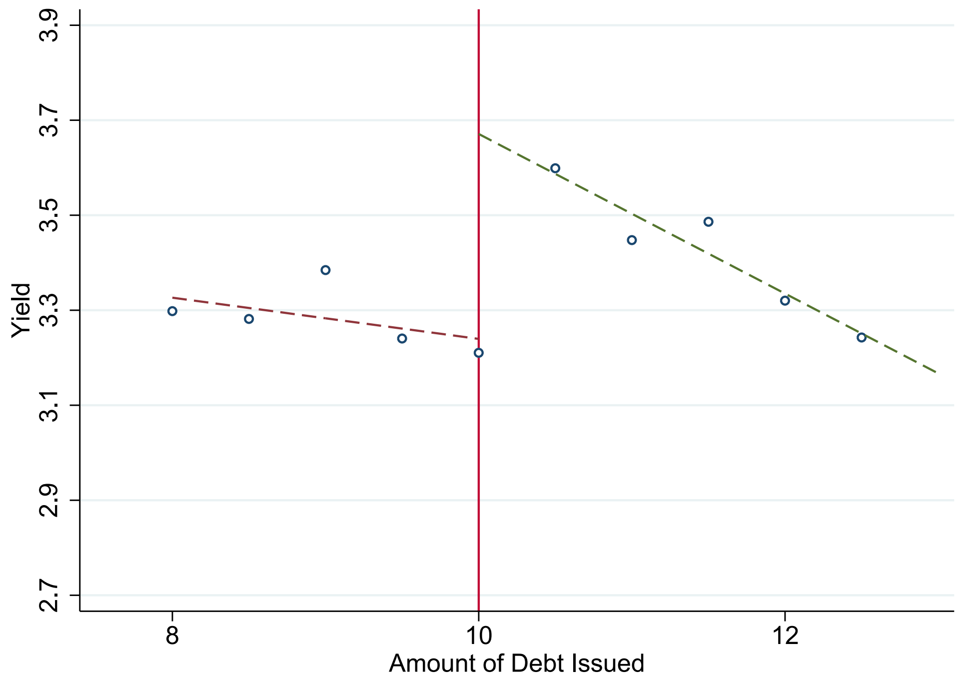

Note: The figure uses bond data to plot yield as a function of the amount of annual debt issued. The data come from Refinitiv and have been aggregated to the level of the issuer-year. The sample covers 1998-2015 (excluding 2009-2010). Each circle represents the mean within bins of \$500k.

Table A1: Size of Excess Mass by Census Year

	(1) 1987	(2) 1992	(3) 1997	(6) 2002	(4) 2007	(5) 2012
Excess Mass	0.11	0.09	0.10*	0.12**	0.24***	0.23***
	(0.085)	(0.076)	(0.058)	(0.049)	(0.059)	(0.055)

Note: * $p < 0.1$, ** $p > 0.05$, *** $p < 0.01$. Table A2 presents estimates of the excess mass by the year of the Census. The extent of bunching is measured in terms of the in-range manipulation, the excess mass relative to the counterfactual distribution in the range of the missing mass. All estimates use bins of \$500k, 8th order polynomials, and an excluded range of \$8.5-14.5 million. Block bootstrapped standard errors in parentheses.