Shocking Wealth:
The Long-Term Impact of Housing Wealth Taxation[1]

Matthijs Korevaar Peter Koudijs

--- Preliminary draft ---

November 11, 2022

Abstract

Housing is the main asset through which households accumulate wealth and taxation of it is highly debated. We provide the first empirical estimates of the effects of shocks to property taxation on lifetime wealth accumulation and investment. To do so, we examine a unique 18th-century tax reform in Holland which resulted in large and unanticipated changes in the effective tax rates on real estate wealth, plausibly exogenous to the owners and different for each property. We collect archival data on the wealth and home-ownership of all 18th-century Amsterdam inhabitants and determine their individual exposure to the shock. We find the reform capitalized into house values in the short-run and had large impacts on household wealth that grew substantially over time. On average, a one percent shock increased wealth at death by four percent. We show these effects are consistent with the notion that households do not update housing consumption in response to large tax changes: large positive or negative shocks had little-to-no impact on the likelihood of selling voluntarily, even in a liquid market with low transaction taxes. Instead of altering housing investment, we show changes in taxes primarily affected annual saving. The shock had a very large impact on foreclosure rates and still affected property-level vacancy and owner-occupancy rates 70 years after the reform. Our findings suggest that shocks to property taxation have large and persistent effects on household wealth and the housing stock, which extend far beyond their direct impact on housing valuations.

Keywords: housing wealth, wealth accumulation, housing policy, taxation, foreclosures

[1] Korevaar: Erasmus School of Economics, Erasmus University Rotterdam and Columbia Business School, Columbia University (email: korevaar@ese.eur.nl). Koudijs: Erasmus School of Economics, Erasmus University Rotterdam (email: koudijs@ese.eur.nl). This paper has benefited from funding of the European Research Council (project HISHOUSHOCK). We thank Bram van Besouw for helpful comments and Martijn Sailllard, Nina Wildemast, Laura van der Ziel and Lizette van der Ziel for excellent research assistance.

1. Introduction

Housing is the main asset through which households accumulate wealth. Taxation of it is highly debated. On the one hand, a wide set of tax policies affects the desirability and safety of homeownership and its returns, although the effectiveness of such policies is debated (e.g. Sommer & Sullivan 2018, Goodman & Mayer, 2018). Prominent examples include the exemption of (imputed) rental income from wealth taxation and mortgage-interest deductions. On the other hand, property tax revenue is a major source of revenue for local governments. Effective rates do not only vary across jurisdictions but also across properties within the same jurisdiction. In the United States, property taxes are typically regressive with expensive properties having lower effective tax rates (Amornsiripantich 2021, Berry 2021). In practice, there thus exist sizable differences in the effective tax rates property owners face on otherwise equivalent housing assets. Such differences might have substantial wealth effects as they directly affect the present value of the net-of-cost flows received by property owners.

In this paper, we investigate how tax-driven housing wealth shocks affect the long-term wealth accumulation of households. To identify this effect, we study a unique 18th-century reform of the property wealth tax in Holland that equalized the effective tax rates on properties, resulting in a large shock to tax rates and real estate wealth that was exogenous to the owners and different for each property. Using archival data from Amsterdam, we show the 1732 reform capitalized into house prices and that its large impact on the wealth of households grew over time. We show this is driven by the fact that households do not adjust their housing consumption and investment decisions in response to a large shock, instead adjusting annual saving and non-housing consumption. In line with this, we find the shock had a sizeable impact on foreclosure rates and a persistent impact on the quality of properties and their occupancy.

There are four major reasons why this reform is an ideal experiment to measure the impact of tax-driven shocks on housing wealth. First, the tax shock was exogenous to the current owners and varied at property-level, implying that every property-owning household was affected differently. This allows for much more precise identification compared to estimates based on geographic or time-variation in taxes, which are common in the literature. Before the 1732 reform, property wealth taxes were levied based on the rental value in 1632. Taxation was uniform, so residential (+- 85%) and non-residential (+- 15%) real estate was treated equally and so were owner-occupied properties (+- 20%).[2] Due to the enormous growth of Amsterdam in the mid-17th century, the 1632 values were soon outdated. Later in the 17th century, laws were passed to update tax values when properties were changed or improved but these updates were not consistently applied to divergences in rental values realized earlier. As a result, effective tax rates differed significantly and persistently across properties. Like today, these tax differences were widely considered unfair but also difficult to change because households priced them when purchasing a property. It took a century until a new tax register was designed based on current rental values, which equalized and updated tax rates across properties, resulting in a major but heterogeneous shock to housing wealth.

Second, the tax reform primarily redistributed the burden of property taxation rather than increasing total revenues, which increased only by a small amount at the provincial level. Relative to other cities, Amsterdam its tax burden increased because the city had become larger and relatively more expensive, but it did not benefit from the additional tax payments. Contrary to many modern systems, property tax revenues were not used to fund local expenditures but instead paid for provincial expenditures that mostly consisted of debt service and defense. Absent a social welfare system, there were no compensatory mechanisms in place for households that lost substantial amounts of wealth due to the shock. This implies that any wealth effects we measure were purely the result of the shock.

Third, Amsterdam had extremely well-developed institutions for registering and taxing personal property, implying it had well-functioning housing markets and that there exists plenty of

[2] For the remainder of the paper, we will refer to 'housing' given that nearly all real estate wealth consisted of residential properties.

administrative archival data to track the long-term impact of the shock. To the best of our knowledge, this paper is the first to empirically identify the long-term effects of housing wealth taxation. We have newly-digitized measures of wealth at both death and marriage for all inhabitants of Amsterdam in the 18th century, which we can link to the property tax shock for individuals with unique names. For some individuals, we can also observe their investment portfolios at death. To measure housing market effects, we link the tax shock data on all housing sales in Amsterdam in this period, including foreclosures. We also make use of data on occupancy from the rental censuses in 1805, allowing us to link the tax shock in 1732 to the long-term development of properties.

Fourth, although our shock happened centuries ago, it shares many characteristics with modern discussions about reforming housing taxation. One striking example is the current situation in New York City. Due to a complicated system of exemptions and valuations, properties can face effective tax rates ranging from less than 0.01% to over 2% of market value per year.[3] Politicians and action groups have been calling for reform for decades, but actual reform is not yet in sight. Our historical experiment informs what happens if such a system of inequitable housing wealth taxation does get replaced by a system that taxes housing wealth more equally across owners and properties.

The main results of the paper are as follows. First, we use the records of the tax reform to determine the magnitudes of the tax shock. Relative to actual rental value, Amsterdam households on average annually paid 5% extra in taxes after the reform, with a standard deviation of 7% at property-level and 6% at household-level. Most properties experienced increasing taxes after the shock but a fifth paid less and thus gained value. Linking the tax shock to housing transactions of the same property, we find that properties with relatively low effective tax rates before the reform sold for higher prices but decreased in value after the 1732 reform equalized tax rates across properties. This effect was stable in the decades before the reform, implying that property-owners exposed to the shock experienced parallel trends in house values until the

[3] See: "How a $2 million condo in Brooklyn ends up with a $157 tax bill", *Bloomberg,* October 14, 2021

reform shocked prices. We find an estimated capitalization factor of 0.8, close to and not significantly different from full capitalization, similar to modern estimates (Koster & Pinchbeck, 2022). The changes in taxes thus translated into shocks to housing wealth of roughly the same order of magnitude. In the remainder of the paper, we will refer to the magnitude of the shock as the change in housing wealth assuming full capitalization of the tax changes.

We then move on to study the effects of the reform on household wealth accumulation. We focus our main analysis on the subset of individuals for whom we can link their exposure to the estate tax record containing their exact wealth at death. We find that a one percent increase in house value due to the tax shock increased wealth at death by approximately four percent, controlling for real estate wealth at the time of the shock or wealth at marriage, indicating that the long-term wealth effects of the shock were much larger than the initial price effects. An individual experiencing the average shock (-5%) would thus lose twenty percent of wealth at death. A one standard deviation increase in the shock (6%) would increase wealth by 24 percent. In line with growing wealth effects over time, we show these effects are much larger for individuals that died long after the shock compared to those that died shortly after.

These estimates are all conditional on an individual leaving an estate and thus having registered real or financial assets. In the second part of the analysis, we investigate whether the shock also affected the probability of dying without any financial assets. We show that a one standard deviation increase in exposure to the shock increased the likelihood of dying with any financial assets by 5% relative to a baseline probability of approximately 60%. All these effects are robust to controlling for pre-reform wealth, confirming the clean nature of the tax shock.

The growing wealth effect of the shock over time suggests that households' savings and returns were still affected by the shock long after it materialized. We argue that a key factor driving this effect is the limited adjustment of households' housing consumption in response to the shock. The wealth shock was driven by a change in property tax payments, effectively changing the costs of their housing consumption and affecting their after-tax disposable income. In a frictionless world, this shock would cause households to reoptimize their consumption and savings. However, adjusting housing consumption is costly due to the fixed costs involved in moving, implying few households might be willing to do so. If such frictions deter households from adjusting consumption, annual savings will fall by larger amounts causing the wealth effect of the shock to grow over time. Given that households typically have limited wealth next to their own house, this might lead to reduced funds for the upkeep of the property or even result in foreclosure.

We present various pieces of evidence in line with this mechanism. Linking the wealth shock to the housing sale decisions of exposed households, we show the length of tenure was only weakly positively affected by the shock, implying few households decided to move in response to a negative shock. At the same time, it strongly affected the probability of foreclosing on the property, which was often caused by tax delinquency. Across all properties, a one standard deviation increase in the shock cut the probability of selling in foreclosure by almost 40% relative to the base rate of 2.4 percent.

Next, we investigate the impact on wealth accumulation in real estate and non-real estate assets. We show that the large impact of the shock on wealth accumulation was primarily driven by changes in non-real estate wealth: a one standard deviation increase in exposure to the shock reduced the fraction of wealth in real estate by six percent. For the median individual in the sample, holding 3600 guilders of real estate assets and 1000 guilders in other assets, a one standard deviation shock implied an increase of 300 guilders in non-housing wealth at death relative to a 400 guilder decrease in lifetime taxes. In line with the gradual nature of saving, we again find the effect to be larger for individuals that died long after the shock.

To study whether the shock also affected the quality of properties in the long-term, we link the shock to the occupancy of properties in the rental census of 1805. Our intuition is that if the shock affected investments in the property, properties receiving a negative shock are more likely to be in a bad state in the long-term and thus more likely to be vacant. Because Amsterdam ended up in a major crisis after 1780, when most of the exposed generation had passed away, a renovation was likely not very attractive for later owners. The city only started growing again in the mid-19th century. We find that a six percent increase in value due to the shock (one standard deviation), reduced the likelihood of vacancy 70 years later by about 1 percentage point, relative to an average of 8 percent. Finally, we show the shock also had a small but significant positive effect on home-ownership in 1805, in line with the shock significantly affecting foreclosure sales.

The findings of our paper link and contribute to various literatures. First, this paper contributes to the emerging literature on the impacts of wealth taxation (e.g. Seim 2017, Jakobsen et al. 2020, Ring 2021, and Brulhardt et al. 2022). These papers show the large impacts of wealth taxation on reported wealth. Relative to these papers, we study the impacts of taxation on wealth accumulation over much longer horizons and focus specifically on the taxation of housing wealth. Housing is generally a middle-class asset implying that the effects of taxation might differ from general wealth taxes that typically primarily impact those at the top. More importantly, the lack of significant adjustment in housing consumption and investment in response to changes in housing taxation, a key driver of our long-term effects, is likely specific to housing assets. For other financial assets, the costs of such adjustments are arguably much lower. This suggests housing wealth taxes might have very different impacts relative to general wealth taxes. In line with this, existing literature also points to a distinct role for home-ownership in the process of wealth accumulation (Sodini et al. 2021).

Our focus on the tax treatment of housing closely links to a large literature on the impact of fiscal subsidies and taxation on the housing market. Various theoretical papers argue that the mortgage interest deduction and limited taxation of (imputed) rental income are distortive so that tax reforms are generally welfare-improving for households (e.g. Floetotto et al. 2016, Sommer & Sullivan 2018, Boerma 2019). While our empirical analysis confirms the large distortive effects of differential tax treatments of housing investments, these same effects also imply there are large long-term effects when households face unanticipated increases in taxation, such as an increased likelihood of foreclosing and persistent property depreciation. As a result, the long-term impacts of reform are much larger than would be expected based on the short-term wealth shock alone.

These findings also align with a large and growing literature on the impact of property taxation and tax delinquency. Various papers document that in many localities in the United States, property tax appraisals are biased, implying the tax system is regressive and also puts a higher burden on minorities (Hodge et al. 2017, Avenancio-Leon & Howard 2019, Amornsiripantich 2021, Berry 2021). LaPoint (2022) documents the system of tax lien sales (sales after tax delinquency) and shows how such sales might be a sizeable contributor to neighborhood gentrification. Wong (2020) and Fu (2022) show that property tax hikes increase the likelihood of respectively mortgage foreclosure and tax delinquency in the short-term, in line with our findings. Beyond our identification, the main contribution of our paper is that we can study the long-term impact of tax-driven wealth shocks and link this to household-level outcomes.

2. Data

Nearly all data used in this paper originate from hand-written administrative records kept in the Amsterdam City Archives. Part of this data has been digitized by the Amsterdam City Archives, and other sources we have transcribed ourselves from the original archival records. We use several main sources. The most important source for this paper is the *kohier van redres*, a register made in 1731 and 1732 containing the name(s) of the owner(s), the value of the old tax and the value of the new tax for each of the 25,933 parcels in Amsterdam. We have digitized this register entirely and verified it with similar registers existing in the Dutch National Archives and the Amsterdam Archives. We explain this procedure in more detail in Appendix A. As a result, we can compute the tax shock for each property and owner.

The second main source is a matched database containing the wealth at marriage and death for all individuals in Amsterdam in the 18th century. To construct this, we started from the database of all births, marriages and burials in Amsterdam which has been made available to us by the

Amsterdam City Archives. Although the includes millions of records from the period from 1554 until 1810, we focus on the period from the late 17th century onwards. For individuals that married or died between 1699 and 1805, we then digitized data from a wealth tax at marriage and death that was levied during this period. We transcribed all the names of individuals that had to pay tax and their wealth class, ranging from 1 to 4. The data include 51,403 grooms and brides and 115,413 buried individuals. We linked this to their burial and marriage records based on their names and dates of the event. Individuals that had no wealth did not have to pay taxes. Because the linking between the records can be done accurately, we can identify the wealth status of all people that died and married in this period with high levels of certainty.

To match marriages to burials, we use approximate string matching. The basic intuition is that we can match an individual if we find a unique match between a name in the marriage records and all individuals that died within 50 years of that marriage. To restrict the number of potential matches, we also link marriages to the baptism records of children born out of the marriage, so that we can exclude individuals that died more than a year before a child was baptized. In total, we have repeat-wealth observations for about 133,000 people that died between 1700 and 1805 relative to 755,000 deaths (although many of these are children). The entire procedure of matching burials and marriages to the tax records and to each other is explained in Appendix B.

The third main source is the set of housing transactions and estate tax records in this period, introduced in Korevaar (2022). The housing transaction dataset provides details on the names of buyers and sellers, transaction prices, and approximate locations for properties sold in Amsterdam between the 17th century and 1810. The estate tax records provide detailed information on wealth at death and the composition of it across assets, and we match it to the burial records described above.

To estimate the long-term impacts of the shock on the housing stock, we make use of the rental census of 1805, which we have digitized from the archives. This census indicates for every property the actual rental prices (if leased) and the number of units that are vacant and owner-occupied. They can be easily linked to the 1732 tax register based on their property identifier. Not all census data has survived: data are missing for approximately 25% of neighborhoods.

3. Historical Background: The Tax Reform of 1732

During its existence, the Dutch Republic had an advanced system of wealth taxation in place (Fritschy, 2017). The government kept detailed records of property ownership and other personal wealth in *kohieren* that enabled the taxation of wealth. Wealth taxes were sometimes levied on general wealth (general property tax) but more structurally levied per asset. The property tax should thus be thought of as being part of a general set of wealth taxes. Most wealth was held in government bonds and real estate, which were taxed at similar rates for most of the 18th century. Taxes were generally levied on the cash flows provided by the assets rather than their total value. Relative to total value, taxes averaged around 1.5% of bond or real estate wealth per year.

The main property tax was called the *verponding* and it was levied on the annual rental value of a property. The Province of Holland, which included Amsterdam, developed its previous property tax register in 1632. For each parcel in the province, the government obtained the current rental price or appraised it in case the property was not (entirely) leased. Based on this price, the government levied an annual tax that equaled 12.5% of rental value. In 1667, the tax was raised by a tenth to 13.75% of the rental value per year. If the financing needs of Holland were high, it would levy extra taxes, at half the rate or the regular rate. Double taxation was introduced from the late 17th century onwards, implying that individuals paid 27.5% of the assessed rental value in taxes. Depending on the financing needs of Holland, the tax was levied three times a year occasionally, but the tax had been stable at 27.5% of 1632 rental value for almost two decades before the 1732 reform.

Importantly, the rental values were not updated even if rental prices on a property changed. Over time, this resulted in substantial disparities in actual prices and the rental values in the tax records. Effectively, three key issues were driving these disparities. First and most importantly, newly-built properties were not always accurately taxed. In 1654, the States had agreed that the value of newly built property or substantial property improvements had to be taxed. To stimulate construction, newly-built properties were exempt for the first 14 years after completion, except in case of extraordinary taxes. However, in practice, many properties were not assessed properly even after the law was introduced. Second, and relatedly, urban expansion implied that the relative valuations of properties within cities changed. Properties that were initially located in undesirable neighborhoods might have appreciated significantly in price if urban expansions and new construction implied that the neighborhood gentrified. Finally, some cities had generally experienced increasing rental prices on existing properties, while other cities had experienced decreases. This implied that the burden of taxation was much higher in cities that had declined and many had difficulty raising the required amount of taxes.

The problem of undertaxed property and tax write-offs on existing properties became particularly pressing in the 18th century, when increases in public debt forced Holland to raise more wealth taxes to pay for debt service. On the 24th of May 1721, the States of Holland concluded that it was unlikely that the cities that fell behind on taxes would eventually be able to pay all of it, suggesting that "for the future, the tax registers should be strengthened and the losses supplemented with the value of the newly constructed buildings and the reclaimed lands." In the next years, the States of Holland acted. In June 1723, it decided that all cities should make a list of all newly constructed properties before 1709 and that this should be approved by the States of Holland. Between 1724 and 1726, various mandates extended this issue to also include recent property improvements. While recent changes were incorporated in the tax records, most of the disparities remained: the reforms still did not tackle properties whose rent prices had changed substantially, and many property improvements and changes dated back to early in the 17th century and could not be incorporated.

Around 1725, the town of Gorinchem started to push for a complete renewal of the register. In 1724, only half of the Gorinchem homes that were in the 1632 register still paid taxes: the other

half had defaulted on taxation and most of these properties were vacated and in ruins. It contrasted its experience to Rotterdam and Amsterdam, whose strong population and rent increases during the 17th century had only resulted in small increases in tax revenues.
Their proposals initially did not get full support. The Gorinchem aldermen wrote in November 1726 that “the proposal to come to a renewal of the property tax register was supported by many members, but some have used the same arguments that were (fortunately unsuccessfully) used to prevent the previous renewal in 1632. That is, due to its effect on financial returns: many old homes are too high in the *verponding,* and many of these properties have been sold, and the price of the sale was affected by the amount of the *verponding*, so that any change will mean a profit to some owners and a loss to others.” In short, reforming the tax implied substantial wealth losses or gains for owners, because the current deviating tax rates have been fully priced in. A reform would thus on average hurt the wealth of cities that had seen strong increases in rental value, such as Amsterdam.

While more than half of Holland’s real estate wealth was located in Amsterdam, it could not prevent reform because every city had one vote. After many reports and meetings, the States of Holland decided on the 24th of March 1729 “that a general update of the tax registers shall be designed, without reducing the total tax revenue”. The decision remained debated: The French philosopher Montesquieu discussed the reform in the memoirs of his travels to Holland in 1729, writing “there were other proposals, but they were not listened to”.

It remained unclear what the new tax rate would be: the actual plans for reform were only approved in May 1730. The States of Holland decided that all properties would be assessed again to determine their actual rental price. The assessment would be based on actual lease contracts or appraised rental values in case those were not entirely available, for example, if a property was (partially) vacant or owner-occupied. In 1731 and 1732, all properties in Holland were assessed by appraisers and by September 1733 the new tax register was available for inspection. The tax was reduced from an eighth of rental value to a one-twelfth of rental value. Because properties had appreciated in price over time, overall tax income increased by 10% in Holland.

From 1735 onwards, all cities levied taxes based on the new tax register. After the reform, tax rates only briefly increased around the Austrian Succession War in the 1740s, as part of a general war-time increase in wealth taxes, but stayed at the previous rates until the late 18th century.

The impact in Amsterdam

In this paper, we focus on the impact of the renewal of the tax register in Amsterdam. First, Amsterdam was by far the largest market, containing more than half of Holland's total real estate wealth in 1732. Second, the city had grown substantially in size since 1632, implying that the valuations of neighborhoods and the properties therein had changed substantially. Amsterdam had engaged in an enormous urban expansion of the city to accommodate its increasing population during the Golden Age (Abrahamse, 2010). Most of this construction activity had taken place in the first decades after the property tax update of 1632. Amsterdam had stagnated in the late-17th century, implying further growth was not necessary. This implies that most changes in rental value were driven by changes in the city and housing market that went back to the first decades after the 1632 reform. Finally, disparities between the old register and the new registers were not driven by more recent property-level improvements, because the States of Holland had already pushed Amsterdam to incorporate recent changes in the register. Altogether, this makes the update an ideal experiment to study the impact of property tax shocks on housing wealth.

In Figure 1, we show an extract from the *kohier van redres*. For each property, it lists the old tax, the tax identifier, the name(s) of the owner(s), the current rent or rental value and the new tax. In total, there are 25,933 properties in the registers. In Appendix A, we explain in more detail how we transcribed these registers and how we dealt with missing observations and parcels that changed over time. Because tax divergence was strongly linked to the 17th century construction boom in Amsterdam, there is some spatial correlation in the shock. The correlation between tax changes in neighboring properties is 0.33 and primarily driven by split properties that receive the same shock. It reduces towards zero at more aggregate levels. Most of the tax shock is thus property-specific. In total, Amsterdam tax revenue increased substantially in the new register, with total taxes increasing by 32% relative to the old amounts. While the formal tax rate had been much higher in the previous tax register, it was in practice lower because most properties had a much higher rent relative to the old taxed rental value.

Figure 1: Extract from the renewal of the property tax register

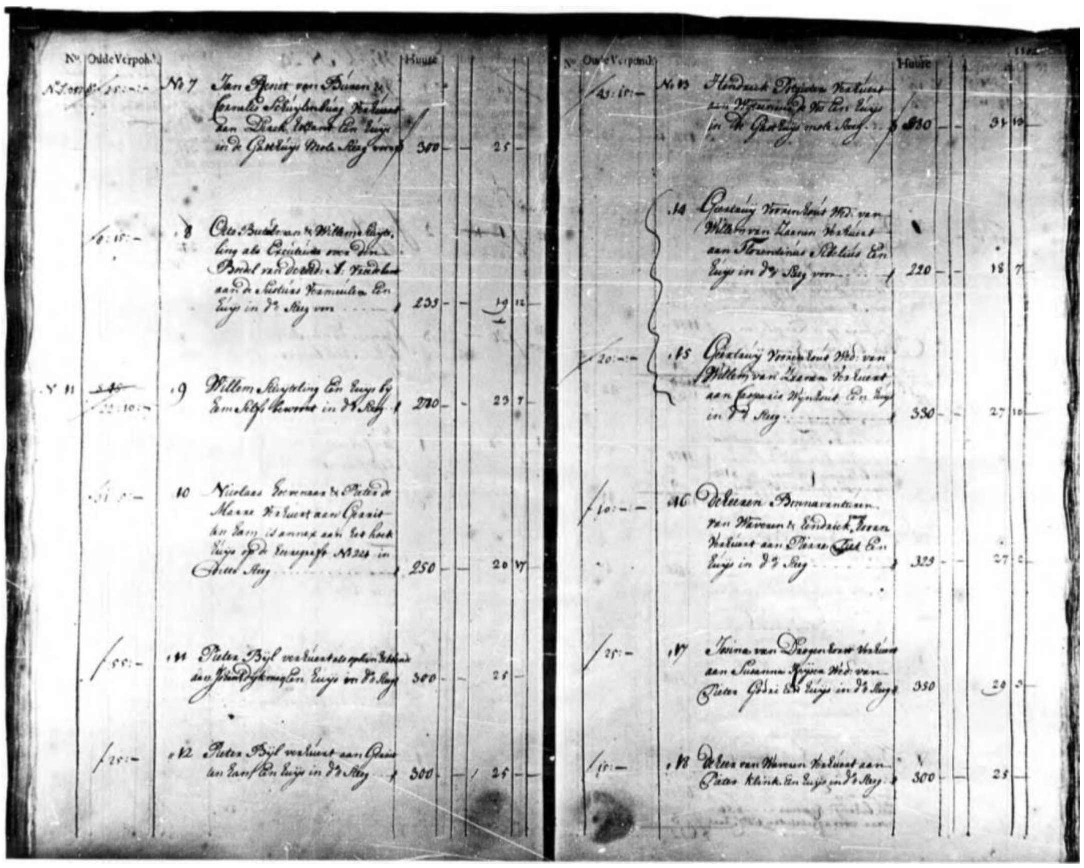

Figure 2 reports the gain or loss due to the change in tax expressed as a percentage of the 1732 rental value of a property. If taxes were fully capitalized, this would equal the loss in home value. Figure 2 excludes properties owned by institutions, leaving us with 24,813 properties. Only 0.6% of properties experienced an implied value gain of more than 20% of rental value. These properties remain in the dataset but are excluded from the plot.

The mean property in the sample lost about 5% in rental value. Excluding extreme outliers, the standard deviation of the shock is 7.9%. The same values are -0.055 and 0.073 when expressing the shock in log terms, which we will use in the empirical part of the paper. In total, 17.3% of properties gained value and about 3.8% of properties lost the maximum value of 16.7% because they were untaxed in the previous register or not correctly identified. Overall, the changes in rental value are sizeable, and if fully capitalized, implied a substantial change in total wealth: for the large majority of households, real estate was the largest asset in their portfolios.

Figure 2: Change in Tax Rate, as % of rental value

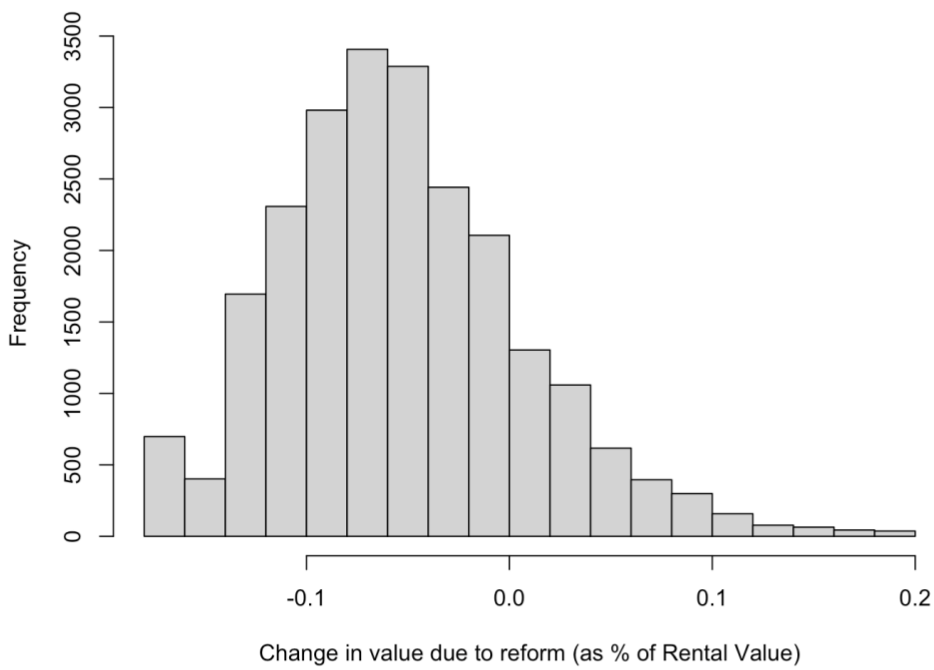

4. The immediate wealth effect: the impact on property prices

The effect in Figure 2 only reflects the distribution of the change in current taxation relative to rental value in 1732. The question is to what extent this affected property prices. As some contemporaries observed, individuals had purchased property before 1732 anticipating that the rental values would not be updated. After 1732, rental values were updated so that effective tax rates did not fluctuate across properties. The question is thus to what extent the prices of properties sold before 1732 were affected by their relative taxation. This depends both on the degree of capitalization of property taxes and the expectations of homeowners about tax rates.

To estimate the price effect, we link the changes in taxes at property-level to repeated transactions of the same properties. We use the same set of repeat-sales identified in Korevaar (2022), focusing on the subset of 18,573 pairs transacted in the five decades before and after the reform, covering the period from 1682 to 1781. Of these 18,573 pairs, we can link 3,299 pairs to the shock in taxes in 1732. About half of these pairs are the purchase and sale transaction executed by the owner that is listed in the 1732 register, the remaining half is based on previous or subsequent sales of the same properties.

The basic intuition of the matching procedure is that we can establish a match if an owner only owns a single property on a given street in the 1732 tax register and the name of that owner appears once as a buyer before 1732 and once as a seller afterwards. Errors in this procedure could be driven by misspellings in names and wrongly identified streets.

To estimate the impact of the shock on property value, we start with some definitions. We define s_i as the log shock in value due to the change in tax, which is equal to the logged value change visible in Figure 2. The value is zero if a property is not matched to a tax change in the 1732, which we indicate here with a matched dummy that takes the value of one if a property is matched:

$$s_{i|matched} = \log\left(\frac{AnnualTax_{i,1632} - AnnualTax_{i,1732}}{TaxValue_{i,1732}} + 1 \right) \times D_i(Matched)$$

To establish the price impact, we estimate the following modified repeat-sales regression:

$$p_{i,t_x} - p_{i,t_y} = \sum_{t=1}^{T} \beta_t D_{t,i} + \sum_{t_{10}=1}^{T_{10}} \gamma_{t_{10}} D_{t_{10},i} \times s_{i|matched} + \alpha F_{i,t_y} + \varepsilon_{i,t_x} - \varepsilon_{i,t_y}$$

p_{i,t_x} is the log price of property *i* for the sale happening at time *x,* with *x* > *y.* $D_{t,i}$ is the matrix of annual repeat-sales dummies, which take the value of 1 if $t = x$ and -1 if $t = y$. β_t is the vector of annual index coefficients. $D_{t_{10},i}$ is again a matrix of repeat sale dummies, but instead uses decadal repeat-sale dummies which we multiply with the estimated shock.

$\gamma_{t_{10}}$ is a set of decadal coefficients that tracks to what extent the evolution of property prices is affected by the property tax shock in 1732. These are the main parameters of interest and we measure them per decade because we do not have enough matched observations to estimate at a higher frequency. If the tax shocks capitalize, properties that paid relatively high taxes before 1732 should have relatively low property prices in decades before 1732 but not afterwards. Additionally, we control for foreclosure discounts using a dummy that takes the value of one if a property is sold in foreclosure (F_{i,t_y}). We take the 1732 period (or the 1732-1741 decade) as baseline.

Figure 3 plots the coefficients of interest of the regression with 95% confidence bands based on heteroskedasticity adjusted standard errors. Importantly, we find that properties that experienced a 1% positive tax shock in 1732 had about 0.8% lower property prices in the period preceding the reform. This point estimate is thus very close to full capitalization of the tax shock: the coefficient is not different from one. Importantly, this effect is also stable in the three decades prior the shock. The insignificant coefficient for the 1692-1701 period likely relates to changes in tax rates. Tax rates were increased in the 1682-1691 decade but were cut again in the 1692-1701 decade, implying that differences in rental values in the register became less important for property value. Tax rates increased again after 1702 but were fixed from 1715 until the reform.

Figure 3: the impact of tax divergence on property prices

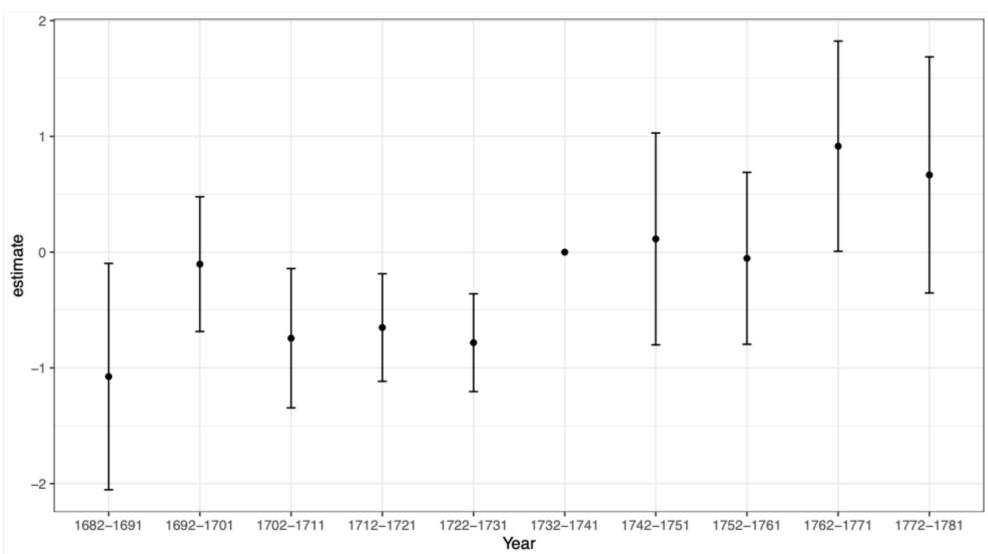

In the three decades after the reform, we do not find that the price impact increased substantially. We find somewhat larger coefficients for the final two decades, with the 1762-1771 also statistically larger than zero. Although the effect is only borderline significant, one factor that could play a role is property improvement and depreciation. Our repeat-sales analysis follows the same properties over time without controlling for quality changes. If households suddenly had to pay more (less) taxes for their properties, this might have affected the savings available for property maintenance. If that would be the case, the effect of the shock on property prices might have strengthened over time. We investigate the evidence for this channel in more detail in the final part of the paper.

5. Wealth effects

For households that were owning a property at the time of the reform the tax shock thus implied a significant shock to their housing wealth. We now turn to analyze how the tax shock affected households' wealth throughout their lifetime. To do so, we link the individuals in the 1732 tax register to their wealth at death and their wealth at marriage. The basic intuition of the linking procedure is that we aim to establish a match between the marriage record and the tax register for individuals with unique names. To do so, we compute Jaro-Winkler distances between each name in the property tax record and all marriages that happened prior to it. If we find a unique match, we include it in the data.[4] To find the wealth at death for these individuals, we use the matches between the burial records, the estate tax records, and the marriage records (see Appendix B.4). In general, we require matches to be very strict to minimize false positives. For matching individuals in the tax register to their marriage and estate tax records, we take less strict values.[5] We can take less strict values here because conditional on owning real estate, the probability of having some wealth at death is much higher relative to that of the population. Thus, in case we can match an individual in the tax register to a couple that later leaves an estate, that match is likely to be correct.

[4] We define a unique match if a match has a summed Jaro-Winkler distance based on the first name plus twice the last name of less than 0.10 and there are no other individuals with scores less than 0.10 away from the minimum score.
[5] The summed Jaro-Winkler distance based on the first name plus twice the last name needs to be less than 0.20.

We only include individuals that died in 1735 or later, the first year in which taxation was based on the new records. In total, we can match 860 individuals based to their marriage and the burials tax they paid and 355 individuals to their marriage and their actual estate. For each individual, we compute its personal tax shock by aggregating the rental values of all properties owned by the individual. Most individuals only own a single property. At person-level *j* the mean log shock s_j is -0.05 with a standard deviation of 0.06. Given that some individuals owned multiple properties, the standard deviation of the person-level shock is smaller than the standard deviation of the property-level shock (0.07).

Both at marriage and death we observe the wealth class of an individual based on the taxation of burials and marriages. There were five classes, with the lowest wealth class containing individuals without wealth, who paid no tax (class 0, *pro deo*), and the top class individual with wealth over 12,000 guilders. Individuals that held a formal office were assessed on income instead of wealth, except if their wealth put them in a higher tax bracket (Hart, 1973). Most people that died with any wealth typically owned a home. 80% of couples and 85% of deceased individuals married or died without wealth. Note that dying without wealth implied dying without any formally registered ownership of any financial assets (such as bonds, equity, real estate or accounts at the Bank of Amsterdam). Some of these individuals likely still had some cash savings. For a subset of individuals where can link to their estate tax records we have an exact measure of wealth as well as information on its composition.

In our baseline analysis, we regress the log wealth at death of individuals on the shock, controlling for their total real estate wealth in 1732 (measured by rental value). On average, the individuals in the sample own 65% of their wealth at death in real estate, which implies that due to capitalization effects a 1% shock increases wealth by 0.65% upon realization. However, we find that the impact of the shock on wealth at death is much larger. Table 1 reports the main results.

Table 1: Wealth at Death in response to the shock

	Dependent variable:					
	log(Wealth at Death)					
	(1)	(2)	(3)	(4)	(5)	(6)
	Regular	*Strict*	*Regular*	*Regular*	*Pre-1748*	*Post-1748*
Shock	4.239***	3.828***	4.672***	4.622***	1.500	7.657***
	(1.393)	(1.446)	(1.407)	(1.407)	(1.918)	(2.070)
log(Rent Value, 1732)	0.812***	0.758***	0.808***		0.806***	0.800***
	(0.080)	(0.093)	(0.080)		(0.102)	(0.131)
Marriage Class = 1				0.611***		
				(0.195)		
Marriage Class = 2				1.446***		
				(0.263)		
Marriage Class = 3				1.405***		
				(0.319)		
Marriage Class = 4				2.446***		
				(0.218)		
Constant	6.088***	6.551***	6.733***	8.812***	5.926***	6.574***
	(0.325)	(0.365)	(0.487)	(0.369)	(0.414)	(0.624)
Decade-of-Death FE	No	Yes	Yes	Yes	Yes	Yes
Observations	355	252	355	316	196	159
R^2	0.253	0.231	0.270	0.356	0.251	0.293
Note:						*p<0.10**p<0.05***p<0.01

Starting from Column 1, we find that a 1% increase in property value due to the shock increases wealth at death by approximately 4%. These effects are also sizable: a one standard deviation increase in the shock increases lifetime wealth by 24%. Unsurprisingly, we find a close correlation between wealth at death and real estate value in 1732. In Column 2, we report the same results using the strict matching approach that we will use in the analysis in the next table; this reduces the number of observations but does not appear to change our results much. In Column 3, we add decade-of-death fixed effects. We also use estimates of wealth at marriage (for marriage before 1735) to see if our effects strongly depend on how we control for pre-shock wealth. We find similar effects of the shock when using this measure to control for pre-shock wealth but the number of observations reduces slightly as some individuals married before the start of marriage taxation. Note also again that wealth at death correlates strongly with wealth class at marriage. Finally, we split the sample in two parts, covering individuals that died before 1748 and those that died after 1748. The basic intuition is that individuals in the first group died shortly after the shock, on average in 1741, and thus cannot have lost too much wealth on top of the direct capitalization effect. Even if they held onto their properties, they only paid higher-than-expected taxes for several years. Individuals dying after 1748 passed away on average in 1760 and thus paid higher-than-expected taxes for much longer, particularly because additional taxes were levied towards the end of the Austrian Succession War in 1748.

Comparing the results in Columns 5 and 6, we find that the impact of the shock on wealth at death is insignificant for individuals that died before 1748 and economically closer to the direct capitalization effect, although we cannot measure the economic magnitude precisely. More importantly, we find very large and significant effects for individuals that died after 1748. For them, a one standard deviation of the shock increased wealth by over 40%. This suggests the effects were growing over time and largest for individuals that lived long after the shock.

The analysis of Table 1 is conditional on an individual dying with any registered assets and showing up in the estate tax records. It is also possible that the shock affected the probability that an individual died poor, without any assets. We can test this directly by investigating whether the shock affected the probability of an individual paying burials tax, which would imply the individual owned some registered real or financial assets. The results are in Table 2. The outline of Table 2 follows the outline of Table 1 exactly except that all results are based on strict matching given the larger sample and the fact that strict matching is more important given that our main outcome variable is measured with significant noise.

Table 2: Wealthy at Death

	Dependent variable:				
	Wealthy at Death (dummy)				
	(1)	(2)	(3)	(4)	(5)
	Strict	*Strict*	*Strict*	*Pre-1748*	*Post-1748*
Shock	0.889***	0.933***	0.771***	0.765**	1.090***
	(0.259)	(0.261)	(0.275)	(0.373)	(0.367)
log(Rental Value, 1732)	0.100***	0.101***		0.114***	0.084***
	(0.016)	(0.016)		(0.021)	(0.025)
Marriage Class = 1			0.240***		
			(0.040)		
Marriage Class = 2			0.212***		
			(0.052)		
Marriage Class = 3			0.321***		
			(0.071)		
Marriage Class = 4			0.311***		
			(0.049)		
Constant	0.388***	0.481***	0.662***	0.315***	0.556***
	(0.062)	(0.087)	(0.069)	(0.085)	(0.112)
Decade-of-Death FE	No	Yes	Yes	Yes	Yes
Observations	850	850	736	461	389
R^2	0.060	0.063	0.104	0.073	0.055
Note:			*p<0.10**p<0.05***p<0.01		

On average a one percent shock increases the probability of dying with any wealth by 0.9 percentage point. In terms of a standard deviation of the shock, the effect is 5.4 percentage points relative to an average share of the sample dying with any wealth of 64 percent. This effect is somewhat but not significantly smaller when controlling for wealth using wealth class at marriage. Although the economic magnitude is again larger for individuals dying after 1748, we find no statistically significant difference in the likelihood of dying wealthy across the two groups in Columns 4 and 5.

6. Investigating the mechanism

Our results in the previous section show that the tax shock had large effects on wealth at death that exceed the capitalization effect. The wealth effects of the shock appear to be increasing over time, with larger wealth effects on households the longer they live after experiencing a negative shock. We now aim to explain the mechanisms driving this finding.

The starting point of our analysis is the fact that the wealth shock is driven by a change in the future taxes the household pays on its housing consumption and investment. The question is how a household responds to these changes. We start from the assumption that households earn a gross income that is used to fund consumption after paying taxes. Income can be earned from labor or capital and consumption can be in the form of housing- and non-housing consumption. Any money that is not spent on consumption will be saved. Savings are used to earn investment returns and to fund property maintenance. The individual makes decisions on how much to work, save and consume to maximize lifetime utility.

In a frictionless world, we would expect that an increase in effective taxes paid on housing consumption due to the shock would cause the household to reoptimize and cut back on housing consumption. The household will likely also adjust other parameters. However, in the real world, reoptimizing housing consumption is costly, because it involves fixed costs, such as transaction costs and moving costs. Emotional attachment to one's property might also prevent individuals from moving. This might imply that readjustment is much more limited in practice. If that is the case, the burden of increased taxation might have translated into comparatively large shocks in the other parameters. If the shock had large effects on net household savings, it might have substantial and growing effects on long-term wealth accumulation and property investment.

We are going to test each of these different mechanisms directly in this section. First, we look whether the shock affected the likelihood of selling voluntarily and in foreclosure. Second, we look whether the shock affected the fraction of wealth saved in non-real estate assets. Third, we look at long-term effects on property-vacancy and owner-occupancy.

Given that this mechanism highlights the role of frictions in reoptimizing housing consumption, it is important to stress that historical Amsterdam likely serves as a lower bound for such frictions. Transaction taxes amounted to 2.5%, lower than the rate in most modern economies. Brokerage fees were low and homes could be sold at low cost in public auctions organized regularly by the city. The housing market was also fairly liquid. Around 3% of the housing stock traded hand each year, despite a large part of the market being owned by long-term rental investors that traded infrequently. These values are similar to the modern Amsterdam housing market. Finally, few households purchased their property with a mortgage and there were no fiscal frictions affecting the decision to buy or rent. This implies that such frictions did not prevent households from adjusting housing consumption, as is the case in many modern economies.

Effects on sales

The most direct way to test whether households are changing their housing consumption decision in response to the shock is to look at the frequency of housing sales. We use the set of repeat-sales that can be matched to individuals in the 1732 register and test whether the shock affected the selling hazard based on a Cox proportional hazards model. The main dependent variable is the holding period and the main independent variable of interest is the tax shock. We only focus on repeat-sales pairs of owners that were exposed to the shock ($s_{i|exposed}$), which are 1,683 individuals in total. This sample contains both individuals that actively sold their property before they died (599 individuals) and individuals whose heirs sold the property after death (1,084 individuals). Sales typically only list the name of the male owner and only mention the name of his wife if she sells the property after he has been deceased.

The most basic specification is reported in Column 1 of Table 3, which shows that a positive shock weakly reduces the selling hazard of each property *i* with owner *j*, thus increasing the holding period. In terms of the hazard ratio (0.461), a one standard deviation increase in the shock (0.06) reduces the selling hazard by about 3%. This small effect does not change much if we add year of purchase fixed effects and a control for the rental value of the property. In Column 3, we add street fixed effects, which slightly increases the effect but also results in higher standard errors.

In short, there is weakly significant evidence that a positive shock increases the holding period. In Columns 4—6, we repeat the same specifications but add a control and interaction term whether a property is sold by heirs. In that case, the decision to sell the property has not been made by the original buyer, which might have changed the role of the tax shock in the sale decision. For example, for heirs that obtained the property after the tax shock, the shock was likely much less relevant. However, standard errors on the coefficients of interest are very large and the coefficients themselves show no consistent pattern. Unsurprisingly, properties of heirs have lower sale hazards.

Table 3: Hazard rates of selling properties, Cox proportional hazard model

	Dependent variable:					
	Holding Period (hazard model)					
	(1)	(2)	(3)	(4)	(5)	(6)
$s_{i\|exposed}$	-0.775*	-0.768*	-1.048*	0.119	-0.888	-1.659*
	(0.413)	(0.402)	(0.564)	(0.775)	(0.776)	(0.731)
Hazard Ratio:	*0.461*	*0.464*	*0.350*	*1.126*	*0.412*	*0.190*
$log\ (RentalValue_{1732})$		-0.122***	0.044		-0.115***	0.057
		(0.036)	(0.072)		(0.037)	(0.071)
Hazard Ratio:		*0.885*			*0.891*	*1.059*
$Heirs_j$				-0.543***	-0.335***	-0.402***
				(0.066)	(0.067)	(0.078)
Hazard Ratio:				*0.510*	*0.715*	*0.669*
$s_{i\|exposed} \times Heirs_j$				-0.938	0.275	1.076
				(0.922)	(0.911)	(1.041)
Hazard Ratio:				*0.386*	*1.317*	*2.933*
Year of Purchase FE	No	Yes	Yes	No	Yes	Yes
Street FE	No	No	Yes	No	No	Yes
Observations	1,675	1,675	1,675	1,675	1,675	1,675
R^2	0.002	0.298	0.496	0.055	0.315	0.424
Wald Test	3.610	482.58	907.44	98.94	529.29	953.200

In short, while there is some evidence that a negative shock and increase in taxation resulted in a higher likelihood of selling, this effect is economically small and only weakly significant. Individuals that experienced a large negative shock and faced increased taxes generally did not sell their properties. This suggests few households adjusted their housing consumption in wake of the shock.

Effects on foreclosure rates

Households that experienced a negative shock but did not adjust their consumption decision might have been increasingly unable to pay the property tax, particularly if they ran out of savings. If home-owners were delinquent on their taxes, the city would proceed to sell the property in a foreclosure procedure. The proceeds would be used to pay off outstanding tax payments and any other claims creditors had on the property. To test this, we related the property tax shock to the likelihood of a home foreclosing. The results are in Table 4.

Table 4: Foreclosure sales

	Dependent variable:					
	$Foreclosure_{i,j,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$s_{i\|exposed}$	-0.147**	-0.123*	-0.180***	-0.300***	-0.239**	-0.397***
	(0.068)	(0.069)	(0.059)	(0.113)	(0.113)	(0.094)
$log\ (RentalValue_{1732})$		0.005			0.006	
		(0.006)			(0.006)	
$Heirs_j$				-0.057***	-0.071***	-0.043***
				(0.011)	(0.012)	(0.003)
$s_{i\|exposed} \times Heirs_j$				0.265*	0.195	0.346***
				(0.140)	(0.141)	(0.116)
Constant	0.024***			0.061***		
	(0.005)			(0.009)		
Year of Sale FE	No	Yes	Yes	No	Yes	Yes
Year of Purchase FE	No	Yes	Yes	No	Yes	Yes
Street FE	No	No	Yes	No	No	Yes
Observations	1,675	1,675	18,573	1,675	1,675	18,573
R^2	0.003	0.093	0.125	0.042	0.131	0.138

Starting with the first column, properties that experience a one standard deviation of the shock (0.06) reduce their probability of selling in foreclosure by about 0.9 percentage points. We find this effect changes little when adding year of sale and purchase fixed effects (Column 2) as well as street fixed effects (Column 3). To compute the street fixed effects, we now make use of the full sample of transactions in the 1682-1781 period, which allows for increased precision of the estimated effects given that there is little variation in the dependent variable. The effect we find is economically large and amounts to a reduction of around 40% relative to the baseline probability of foreclosing of 2.4 percentage points.

Columns 1—3 combine property transactions executed by individuals during their lifetime and by heirs after death. Inherited properties are much less likely to foreclose because heirs typically either sell the property quickly after death or use it as an investment property. In Columns 4—6 we repeat the same specifications controlling for ownership of heirs. This approximately doubles the economic magnitude of the estimated effects. Unsurprisingly, heirs are extremely unlikely to foreclose, and correspondingly the foreclosure probabilities of sales by heirs are barely affected by the shock.

Effects on savings

The analysis so far shows that individuals experiencing a negative shock were much more likely to foreclose but did not become much more likely to sell their property. This suggests that individuals did not reoptimize their housing consumption in wake of the shock, some even until the point that they were forced to foreclose. This lack of a housing consumption response might imply that households simply save less when forced to pay more taxes and save more when they unexpectedly pay less taxes.

We now study whether the shock affected the share of wealth individuals invested in real estate. If individuals receiving a positive shock did not alter their investment at all, we would expect these households to hold more wealth in real estate assets given the price effect of the tax shock. If individuals instead started saving more in response to a positive shock, we would expect non-real estate investments to grow as well. Note again that for most individuals in the sample, their house was their main investment with an average share of 63% and any other savings would generally be invested in bonds or equities. The records exclude cash. Individuals that lost all their financial assets, potentially due to a negative tax shock, do not show up in the records. Table 5 reports the results of the analysis and follows the outline of Table 1 and Table 2.

Table 5: Impact of the shock on saving

	Dependent variable:				
	Fraction of Wealth in Real Estate				
	(1)	(2)	(3)	(4)	(5)
	Regular	*Regular*	*Regular*	*Pre-1748*	*Post-1748*
Shock	-0.962***	-1.046***	-0.811**	-0.663	-1.432**
	(0.365)	(0.370)	(0.361)	(0.490)	(0.566)
log(Rental Value, 1732)	-0.063***	-0.062***		-0.072***	-0.043
	(0.021)	(0.021)		(0.026)	(0.036)
Marriage Class = 1			-0.134***		
			(0.050)		
Marriage Class = 2			-0.157**		
			(0.067)		
Marriage Class = 3			-0.222***		
			(0.082)		
Marriage Class = 4			-0.410***		
			(0.056)		
Constant	0.388***	0.481***	0.662***	0.315***	0.556***
	(0.062)	(0.087)	(0.069)	(0.085)	(0.112)
Decade-of-Death FE	No	Yes	Yes	Yes	Yes
Observations	355	355	316	196	159
R^2	0.048	0.062	0.188	0.051	0.062
Note:			*p<0.10**p<0.05***p<0.01		

Starting with the first two columns, we find that on average a one percent tax shock led to a reduction in the share of wealth invested in real estate by approximately one percent. This effect is similar when we control for wealth using wealth class at marriage. Given that the tax shock increased real estate wealth by one percent and real estate on average was 65% of a portfolio, non-real estate wealth on average increased by about five percent to accomplish such a reduction in the real estate share.[6] For a one standard deviation shock, this implied an increase in non-real estate wealth by approximately 30%. Note that the median individual in the records owned 3600 guilders in real estate and 1000 guilders in other assets and on average died in 1749, so this would imply 300 guilders in additional savings. For the median individual, a one standard deviation additional tax shock was equivalent to approximately 400 guilders in foregone taxes up to 1748, the median year of death in the sample. This suggest that the tax shock primarily affected the amounts that households could save annually. To further confirm this, we find that the negative effect of a positive shock on the share of wealth invested in real estate grows over time, being larger for individuals that died after 1748. Note that we might still underestimate this effect somewhat if additional savings led household to purchase additional real estate for rental income.

Long-term effects on property occupancy

Beyond the effect of the shock on wealth accumulation, there also might be implicit effects of the shock on properties themselves. If households were facing increased tax burdens they might have reduced the amount of money they invested in the upkeep of the property. In the latter case, housing consumption adjusts a little bit over the long run because the property depreciates more over time, leaving households with less wealth and less housing consumption the longer they undermaintain the property. Relatedly, if the shock led to foreclosure, this distress itself might have accelerated depreciation of the property while also increasing the likelihood of investor ownership, who purchased distressed property in auctions.

To test for any such effects we link the shock to long-term vacancy rates of properties in 1805. In 1805, the city made a register that listed for every housing unit in the city whether it was vacant or not. If the shock correlates positively with any investments made in property improvement or

[6] It is more difficult to accurately estimate this on levels of real estate and non-real estate wealth given that 40% of portfolios exclusively consists of real estate, but doing so results in a similar but less precisely estimated pattern.

maintenance, we would expect that a positive shock results in fewer properties depreciating so much that they become vacant. This channel is in line with some of the anecdotal evidence provided by Gorinchem, which argue that increasing real tax burdens had resulted in increased vacancy and depreciation of the housing stock before the reform. It is important to note that after 1780, when most of the exposed generation had died, Amsterdam house prices and rent prices started declining as the Dutch economy entered a major crisis. This makes it less likely that properties that had significantly depreciated in the decades after the shock were renovated by their subsequent owners.

The results of this analysis are in Table 6, Columns 1—3. We consistently control for the rent value at the time of the shock in 1732, given that it controls for the pre-shock desirability of a property. The vacancy rate of a property is 0 if it is entirely occupied and 1 if it is entirely vacant. If a property contains multiple units, the number of vacancies is scaled by the number of units.

Table 6: Long term impacts on occupancy in 1805

	Dependent variable:						
	Vacancy Rate, 1805			Owner-Occupancy Rate, 1805			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
s_i	-0.165***	-0.161***	-0.132***	0.214***	0.119***	0.118***	0.127**
	(0.025)	(0.029)	(0.031)	(0.037)	(0.043)	(0.046)	(0.050)
log(ValNew)	-0.068***	-0.051***	-0.050***	0.097***	0.057***	0.061***	0.063***
	(0.002)	(0.003)	(0.003)	(0.003)	(0.005)	(0.005)	(0.006)
Constant	0.280***	0.078	0.114	-0.088***	0.159	0.105	0.086
	(0.007)	(0.094)	(0.094)	(0.010)	(0.138)	(0.140)	(0.149)
Street & Neighborhood FE	No	Yes	Yes	No	Yes	Yes	Yes
Observations	19,233	19,233	18,727	19,233	19,233	18,727	16,658
R^2	0.057	0.152	0.153	0.055	0.118	0.119	0.117
Adjusted R^2	0.057	0.132	0.132	0.055	0.097	0.098	0.094
Residual Std. Error	0.237	0.228	0.224	0.343	0.335	0.336	0.353
F Statistic	586.091	7.591	7.436	559.625	5.685	5.597	4.991
Note:					$^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$		

The baseline effect in Column 1 suggests that a one standard deviation shock reduced the likelihood of vacancy by approximately one percent. The average vacancy rate in the sample is 8.4 percent, so this is a sizeable economic effect. The effect is similar when we control for location using street and neighborhood fixed effects (Column 2) and only reduces slightly when we remove observations for which no tax was known in the old register, which might have been illegally erected and thus more likely to be vacant later (Column 3).

In Columns 4—7, we extend the argumentation by looking at rates of owner-occupancy of properties. The most direct mechanism for this is that a negative shock increases the probability that a property is sold in foreclosure, implying the property is more likely to end in the hands of a long-term investor, who owned most properties in Amsterdam. Another potential channel is that subsequent owners of a property are less likely to be owner-occupants if it is in a bad condition.

Columns 4—6 follow the same outline as Columns 1—3 except for the change in the dependent variable. We find that the shock positively affects the probability that a home is owner-occupied. This effect somewhat weakens once we control for location. Due to the geographic nature of the shock, homes experiencing a negative shock were more likely to be in newly-developed areas with lower shares of owner-occupied housing. Based on Columns 2—3, a one standard deviation of the shock increased the likelihood of owner-occupancy by 0.7%, relative to an average property-level rate of around 20%. In Column 7, we remove properties that are partially or entirely vacant to show that this effect is not driven by the fact that exposed properties are more likely to be vacant.

7. Conclusion

We studied the long-term consequences of shocks to housing wealth taxation on the wealth accumulation and investment of households. We find large and long-lasting of tax-driven housing wealth shocks on households' wealth accumulation, in particular for households with limited initial wealth. These effects far exceed the short-term impact of the shock on property valuation and taxes. Our findings both indicate that inequitable property taxation has large long-term consequences as well as that reforming property taxation has large negative consequences on households that previously benefited from low rates.

References

Abrahamse, J. E. (2010). De grote uitleg van Amsterdam: stadsontwikkeling in de zeventiende eeuw. *PhD dissertation University of Amsterdam.*

Amornsiripanitch, N. (2020). Why Are Residential Property Tax Rates Regressive?. *Available at SSRN 3729072*.

Avenancio-Leon, C., & Howard, T. (2019). The assessment gap: Racial inequalities in property taxation. *Available at SSRN 3465010*.

Berry, C. R. (2021). Reassessing the property tax. *Available at SSRN 3800536*.

Boerma, J. (2019). *Housing Policy Reform*. Job market paper, University of Minnesota.

Brülhart, M., Gruber, J., Krapf, M., & Schmidheiny, K. (2022). Behavioral Responses to Wealth Taxes: Evidence from Switzerland. *American Economic Journal: Economic Policy*.

Floetotto, M., Kirker, M., & Stroebel, J. (2016). Government intervention in the housing market: Who wins, who loses?. *Journal of Monetary Economics*, *80*, 106-123.

Fritschy, W. (2017). *Public Finance of the Dutch Republic in Comparative Perspective: The Viability of an Early Modern Federal State (1570s-1795)*. Brill.

Fu, E. (2022). The Financial Burdens of Property Taxes: Evidence from Philadelphia. *Available at SSRN 4134172*.

Goodman, L. S., & Mayer, C. (2018). Homeownership and the American dream. *Journal of Economic Perspectives*, *32*(1), 31-58.

Hart, S. (1973). *Een sociale structuur van de Amsterdamse bevolking in de 18e eeuw*. Stadsarchief Amsterdam.

Hodge, T. R., McMillen, D. P., Sands, G., & Skidmore, M. (2017). Assessment inequity in a declining housing market: The case of Detroit. *Real Estate Economics*, *45*(2), 237-258.

Jakobsen, K., Jakobsen, K., Kleven, H., & Zucman, G. (2020). Wealth taxation and wealth accumulation: Theory and evidence from Denmark. *The Quarterly Journal of Economics*, *135*(1), 329-388.

Korevaar, M. (2022). Reaching for Yield and the Housing Market: Evidence from 18th-century Amsterdam. *Available at SSRN 3794782*.

Koster, H. R., & Pinchbeck, E. W. (2022). How do households value the future? Evidence from property taxes. *American Economic Journal: Economic Policy*, *14*(1), 207-39.

LaPoint, C., (2022). Tax Sales, Private Capital, and Gentrification in the U.S. *Yale University Mimeo.*

Ring, M. A. K. (2021). Wealth taxation and household saving: Evidence from assessment discontinuities in Norway. *Available at SSRN 3716257*.

Seim, D. (2017). Behavioral responses to wealth taxes: Evidence from Sweden. *American Economic Journal: Economic Policy*, *9*(4), 395-421.

Sodini, P., Van Nieuwerburgh, S., Vestman, R., & von Lilienfeld-Toal, U. (2021). *Identifying the benefits from home ownership: A Swedish experiment* (No. w22882). National Bureau of Economic Research.

Sommer, K., & Sullivan, P. (2018). Implications of US tax policy for house prices, rents, and homeownership. *American Economic Review*, *108*(2), 241-74.

Wong, F. (2020). Mad as hell: Property taxes and financial distress. *Available at SSRN 3645481*.

Appendix A: Connecting the Old and New Tax

Determining the property-level change in tax required information both on the old tax and the new tax. There are two registers in the archives that detail at property-level both the old tax and the new tax and the names of the owners.[7] There is a third register that contains the final assessment rental values for the new registers, together with the property identifiers (a tax number) and the name of the owner(s).[8] The registers are organized per neighborhood, with Amsterdam divided into 60 neighborhoods within the walls and 5 neighborhoods outside the walls. In total, there are 25,933 properties in the register.

We use the third register to identify the tax values after the tax change, the property identifiers, and the names of the owners. This register is most cleanly written and entirely complete. We use the other two registers to identify the level of the old taxes. There are two caveats. First, in both registers, there are some pages or neighborhoods missing. Second, some parcels that were registered as a single property in the old register would cover multiple parcels in the new register. In most cases this is indicated, but in some cases the old tax value is linked to only one of the parcels in the new register.

Our approach is to use the register that reports combined entries most clearly as a default (no. 33-40) and to resort to the other register (no. 203-268) in case of missing data or unclear entries. Comparing the entirely digitized registers, 10% of tax values differ across the two registers. Most of this is driven by differences in accounting for properties that cover multiple parcels in the new register and a single parcel in the old register. A small fraction of cases might also reflect true errors. Another 20% are missing in one of the two registers, so that we cannot do cross-checks.

[7] Amsterdam City Archives 5045: Archief van de Honderdste en Tweehonderdste Penningkamer of Commissarissen tot de Ontvangst van de Honderdste en Andere Penningen**,** no 33-40 and 203-268

[8] Amsterdam City Archives 5044: Archief van de Thesaurieren Extraordinaris, no 402-405

For properties covering multiple parcels *j = 1,…n* in the new register but a single one *i* in the old register we compute the old tax for property *j* in the new register based on the fraction of rental value that is attributable to the specific property *j:*

$$TaxOld_{i,j} = TaxOld_i \times \frac{RentalValue_j}{\sum_{j=1}^{n} RentalValue_j}$$

In total, we apply this procedure to the 13% of observations in the new register for which we have an old tax and a new tax value and for which the records indicate to which properties the old taxes belong. For some properties, it might be the case that none of the two records correctly link to the old tax: there remain 1176 observations in the data (4% of properties) for which we cannot link to the old tax. This either implied that the property was not taxed before the reform or that the assessors did not write this down. In the latter case, this would give us a biased estimate of the tax, but such cases were likely limited. Taking into account these potential errors and omissions together, our estimate is that for nearly all privately-owned properties in Amsterdam our data correctly identify both the old tax and the new tax payable.

Appendix B: Data Overview and Matching Strategies

General matching approach

A key element of our paper is to match individuals across different datasets. To do so requires identifying unique individuals across databases based on fuzzy string matching. Our general approach follows the approach outlined in Korevaar (2022) for comparable datasets from 17th-18th century Amsterdam, in line with similar procedures used on US census data (Abramitzky et al. 2022). To match individual *i* in dataset *x* to individual *i* in dataset *y,* we compute Jaro-Winkler distances between individual *i* in dataset *x* and a set of 'candidate' matches in dataset *y*. The individual in dataset *y* with the smallest JW-distance to that individual is selected as a match. To assess the uniqueness of a match, we also construct a score that increases in value if there any other *near* matches to that particular name. We vary the tuning parameters for the JW-distances and the computation of the uniqueness scores depending on the characteristics of each dataset. In general, name distances are based on the JW-distance between first name plus twice the JW-distance between the last name. As tuning parameter, we set p=0.10 (see Winkler, 1999). In the remainder of this appendix, we briefly discuss each dataset and the various matching procedures.

B.1 Marriage Data

To reconstruct individual's marriages and their wealth at marriage we make use of three different register. We start by using digitized marriages banns provided by the Amsterdam City Archives, which contain information on 497,569 marriages between the 1565 and 1811. This data provides information on the date of the marriage, the name of the groom and bride and their witnesses, as well as whether it was a protestant or non-protestant wedding. The records also list the name of the previous partner in case the groom or bride was remarrying. Divorce was extremely rare: most remarried individuals were widowed. For a fifth of marriages, we can obtain more detailed data including information about background and the age of the groom and bride. This data comes from De Moor & Van Weeren (2021) and covers 94,303 marriages every five years between 1580 and 1810. Replacing the data from the Amsterdam City Archives with data from De Moor & Van Weeren (2021) provides data on 489,447 marriages between 1565 and 1811.

We clean the names of both grooms and brides to remove any special characters or letter combinations that can be written in multiple ways. We then search for duplicates in the data by computing JW-distances between the bride and groom names and bride and groom names in other records. Some marriages were recorded double, for example if it was registered both in the protestant or non-protestant register (e.g. in case only one of the weds was protestant) or in case a marriage was cancelled to be executed later. After removing duplicates, there remain 461,119 marriages in the dataset.

To obtain information about wealth at marriage, we collect data on mandatory marriage taxes that were wealth-dependent (and introduced in the main text). We only obtain information on 51,403 couples that paid tax between 1699 and 1805. For each couple, we obtain the name of the bride and groom, the year and month of registration and the wealth class of the couple. We only obtained the names of the couples that actually paid tax, which is 20.3% of the total number of newly married individuals in Amsterdam in this period.

We match these individuals to the entire set of marriage records based on their first and last names, adding the JW-scores of the couple (for both groom and bride 1x the score of the first name and 2x of the last name). This matching is relatively straightforward as long as the spelling of the names of the groom and bride is consistent across the records: couples paying marriage tax must appear in the marriage banns in the same period or slightly earlier. We match a couple in case the JW-score is below the value at which the match is more likely to be correct than false.

To determine this value, we compare the distribution of JW-scores including actual matches to a distribution of 'false' matches that compute for each couple in the marriage tax the nearest match to couples that married in a completely different year, and thus do not contain the actual couple.[9] As a cut-off score, we take the value where the expected number of false matches is more than half of the number of expected correct matches. In total, we can match 97% of couples that paid taxes to an entry in the marriage banns in this way. Nearly all of these matches are accurate: only in 0.1% of cases two couples in the marriage tax register are matched to the same individual marriage banns. In that case, we remove the match that has the highest JW-score. Of course, we only observe such double matches in case a couple that paid taxes is falsely matched to another marriage record that also paid marriage taxes. Because only 20% of married couples had to pay marriage tax, this implies about 0.5% of individuals that paid taxes were incorrectly matched. Thus, for 3.5% of married couples with wealth, we either have not identified a match in the records or identified the wrong match. Given that 80% of married couples did not possess any wealth, we in the end correctly observe wealth at marriage for >99% of the couples that married in Amsterdam and appear in the marriage banns.

[9] Plots distributions are available upon request

B.2 Burials Data

For the burials data, we apply a very similar procedure as for the marriage data. To obtain information on the number of burials, we use digitized burial records from the Amsterdam City Archives covering the period from 1554 until 1810, containing in total 1,422,668 persons. For each burial, we know the date, the name of the registered person and the location of the burial site. Not each name in the burial records corresponds to the actual name of the person being buried. For example, when children were buried they were often registered under the name of the father or both parents ("child of …"). In some cases, this also applied to women ("housewife of …" or "…, partner of …"). If such a relationship status was explicitly mentioned, we identified this.

After cleaning all names and removing duplicate observations, we focus on the 755,126 individuals buried between 1701 and 1805, since we also have digitized data on the burials tax in this period. Similar to the marriage records, we only digitized data from the burials tax records for individuals that had to pay tax and thus possessed wealth. For each individual, we obtain the first and last name, the month and year of registration and the wealth class. In total, we digitized information for 115,413 individuals paying burials tax. Only 15% of individuals that died in the 1701-1805 period paid tax. This number is likely lower than the number in the marriage records because a large fraction of buried individuals were children. Some buried children were still taxable because they owned wealth through inheritance or because their parents held a certain office that was taxable based on income (Hart, 1973).

To match individuals in the burial tax register to individuals in the burial registration we use the same procedures as in the case of the marriage records. We compute JW-distance for individuals in the tax records with individuals in the burial records that died in the same month or in an earlier month and find the nearest match. The main difference is that burial records only list a single name whereas marriage records list two names, making it slightly more difficult to obtain a unique match. We thus use lower cut-off values to determine correct matches.

In total, we find a match for 85% of individuals in the burial tax records, lower than the marriage records. We also find much more cases where multiple tax records are matched to the same burial record (0.6% of matches). This implies that an estimated 82% of deceased individuals that paid taxes were matched to the correct burial record. Because 15% of households died with any wealth, this implies we observe the wealth class at death correctly for about 97% of buried individuals.

The lower match rate for the burials dataset relative to the marriage dataset is only for a small extent driven by the fact that is more difficult to find a unique match using a single rather than a double name. Only in about 2% to 3% of cases one or more individuals die in the same month with the same or a highly similar name. A much more important factor is that the name in the burial tax record can differ from the name in the burial record itself when it constitutes a child or partner. For example, a child might be identified by its true name in the burial record but by the name of its parent in the tax records (and vice versa). Duplicate matches are also partially explained by cases where multiple children or a child and wife die in the same month and all are identified as 'child' or 'partner of' the father/husband.

Although the fraction of false matches is small, any that are included in the analysis will bias down our coefficients towards zero. However, the probability of a match being correct is much larger conditional on being included in the sample. Our analysis focuses on individuals that are owning real estate at the time of the reform in 1732, which are mostly men, and who have sufficiently unique names. For this subset of individuals, the linkage between the tax records and the actual marriage and burial records is likely near-perfect.

B.3 Baptism Data

For the period from 1554 to 1811, records digitized by the Amsterdam City Archives provide information on 1,236,573 baptism in the city. For each of these baptisms, we have information on the name of the child, the date of the baptism or the birth date, the church and the name of the parents and witnesses. We focus on observations after 1698 because the marriage tax data no not start prior to 1699. After cleaning and removing duplicates, this leaves us with 707,944 children baptized between 1699 and 1810.

B.4 Matching Marriage Records to Burial Records

To match individuals in their marriage records to their burial records we start by matching the marriage date to children that were born out of the marriage. Because the baptism records report both the name of the mother and father as well as the birth data, we can match the baptism records to the marriage records using similar strategies as presented in section B.2 for matching marriage records to marriage taxes. The main difference is that the pool of potential matches is larger, as we look at all marriages that happened in the preceding 30 years, when nearly every bride should have reached an infertile age.

For every birth, we compute JW-scores of the father and mother with the names of couples in earlier years (for both groom and bride 1x the score of the first name and 2x of the last name). We match individuals when the summed JW scores are below a cut-off and in case there are no other near matches with similar scores. We select the cut-off (0.75) based on the value where a match is more likely to be a placebo match than a correct match. In this way, we can link 462,924 baptisms to the marriage records of their parents. The remaining 38% of newborns either have parents with a non-unique combination of names or were born to parents that did not marry in Amsterdam.

Next, we start matching marriages to earlier marriages. In case an individual remarries after his or her partner passed away, the new marriage record would stipulate the name of the previous partner. We again use approximate string matching to link couples, except that we now require the remarriage to happen within 40 years of the original marriage. 18% of grooms in the data were remarrying of which 57% could be matched to their previous marriage (and vice versa). Similarly, 15% of brides were remarrying of which 72% could be matched to previous marriages.

We then start matching the marriage data to the burials data. For a subset of women, the burial records also report the name of their husband. We can again use the name of the deceased women and her husband to match to the marriage records using the earlier presented strategies. We can match 27% of brides, about 67,000 individuals in total precisely to their wealth at death in this fashion. For grooms, we can only match 0.2% of individuals given that their death records rarely mention the name of their wife.

To match marriages to burials, we look for each bride or groom in the burial records and find the nearest match in the period up to 65 years after marriage, or below 85 years if the age of the groom or bride is known. As starting period for the set of potential matches we use the year of the marriage or, if available, the year in which the final child was born. If a groom or bride remarries, we use the remarriage date of the groom or bride as maximum death year for the earlier partner. For brides, we exclude burials and marriages that are already matched in the previous step. We then compute the best match using JW distances on the first name plus two times the last name (p=0.10) for the set of potential matches. We also compute the second best match to rule out other good matches. We define a match if the summed JW distance of the first best match is below 0.10 and the second best match at least 0.10 above that.