

How do taxes affect the trading behavior of private investors? Evidence from individual portfolio data *

FLORIAN BUHLMANN
PHILIPP DOERRENBERG
BENJAMIN LOOS
JOHANNES VOGET

This Version: July 29, 2021

Abstract

We exploit a large reform of capital-gains taxation in Germany combined with portfolio-level daily panel data to study the causal effect of taxes on individual stock-trading behavior and the disposition effect. We find substantial spikes in selling probabilities around an intertemporal tax discontinuity, and no such spikes after the abolishment of the discontinuity. Using bunching methods, non-parametric regressions and effective tax rates, we quantify the tax effect and identify interesting patterns of heterogeneity. We further find evidence that the well-established disposition effect is strongly affected by the tax discontinuity through tax motivated selling of both gains and losses.

JEL Classification: G11, D14, H20, C41

Keywords: Private investors, Trading Behavior, Capital-gains, Taxation, Disposition Effect

***Buhlmann:** ZEW Mannheim and University of Mannheim. Email: florian.buhlmann@zew.de. **Doerrenberg:** University of Mannheim, CESifo, IZA and ZEW. Email: doerrenberg@uni-mannheim.de. **Loos:** TU Munich. Email: benjamin.loos@tum.de. **Voget:** University of Mannheim and ZEW. Email: voget@uni-mannheim.de. Buhlmann gratefully acknowledges financial support by MannheimTaxation Science-Campus. We thank Jannis Bischof, John Y. Campbell, Ulrich Glogowsky, Nathan Goldman, Michelle Hanlon, Jeff Hoopes, Eckhard Janeba, Dominika Langenmayr, Rebecca Lester, Christian Leuz, Ernst Maug, Henrik Kleven, Andreas Lichter, Andreas Peichl, Jim Poterba, Stefan Ruenzi, Arthur Seibold, Sebastian Siegloch, Oliver Spalt, Stefanie Stantcheva, Erik Theissen, Christian Traxler, Daniel Waldenstroem as well as various seminar/conference participants for helpful comments and suggestions. The usual disclaimer applies.

1 Introduction

As part of the *American Families Plan*, president Joe Biden recently announced a major reform of capital-gains taxation in the US. The reform plan provides for a sharp increase in the long-term capital-gains tax rates for rich households and eliminates the difference between tax rates that apply to short-term and long-term capital gains of these households (The White House 2021). The primary objective of the reform is to ensure that capital gains are no longer taxed at a lower rate than labor income. However, a favorable tax treatment of assets held for the long run for example gives investors an incentive to stay committed and engage in corporate governance, thereby developing the asset’s full potential (Aghion et al. 2013; Dimmock et al. 2018). Shorter holding periods therefore are a possible unintended consequence of the reform. It is eventually an empirical question how such a reform affects investment behavior. How sensitive are investors to changes in the capital-gains tax schedule? Will they realize gains and losses in a different manner than they do in the current tax regime?

Our understanding of the causal effects of capital-gains taxes and the differential tax treatment of short-term and long-term capital gains on investment behavior is for various reasons limited (see below). In this paper, we aim to contribute to this understanding and shed new light on the effect of capital-gains taxes on trading behavior (including holding durations) of private investors. For this purpose, we leverage a large individual-level panel data set (with daily frequency) combined with the institutional setting of Germany where differences in tax rates between short-term and long-term capital gains were eliminated in 2009.¹

Why is this important? First, investors may respond to the elimination of the tax difference between short-term and long-term capital gains by reducing long-term investments (with the implications described above). In line with this notion, our results imply that investors hold on to their appreciated stocks 0.368% longer for each 1% increase in the (net-of-tax) capital gain which is attained by qualifying for the reduced tax rate on long-term capital gains. Second, a significant causal effect of taxes on investors’ trading would imply that there is scope for tax policy to incentivize investors to cut losses in a timely manner and to facilitate a sustained participation in stock market returns.² More generally, a credible understanding of how taxes (and tax reforms) affect capital-gains realizations of individual investors is important because capital gains are a major component of private savings (Fagereng et al. 2019) and realized capital gains are subject to investor-level capital-gains taxes in most countries around the world.

¹The interesting commonality of the recently proposed US reform and the German 2009 reform is the elimination of the difference in tax rates for short-term and long-term capital gains. However, there are at least two notable differences between the German reform that we study and the current US plans. First, the US reform plan provides for a reduction of the tax differential for households making over 1 million USD (couples 2 million USD). Second, the proposed level of US tax rate for these households would be larger than the German post-reform tax rate. Due to these and other cross-country differences, we do not claim that the empirical results in our setting will have exact explanatory power for the effects that can be expected from the new US reform plans. We use the US reform as a motivating example and think of our paper as being broadly illustrative for the effect that such type of reform can potentially have on private investors.

²With respect to losses, we find that investors shorten the holding period of depreciated stocks by 0.435% for each 1% increase in the (net-of-tax) loss related to realizing losses at the lower long-term rate instead of the higher short-term rate. (The value of the loss tax shield is higher when losses are realized subject to the higher short-term tax rate.)

In theory, realization-based taxes on capital gains induce investors to defer the realization of gains (lock-in effect) and to realize losses as they accrue (because losses can be used to offset taxable gains).³ However, it has been suggested that such effects of taxes on individual trading behavior are often swamped or offset by non-tax considerations (Hanlon and Heitzman 2010). In particular, the well-documented disposition effect, according to which investors are more likely to realize gains than to realize losses (Shefrin and Statman 1985; Odean 1998), runs in opposite direction to the effect of capital-gains taxes on individual investment behavior (Grinblatt and Keloharju 2001).

These considerations constitute the motivation for the research questions in this paper: First, we study the causal effect of capital-gains taxes on individual-level holding periods and selling probabilities of private stock-market investments. Second, we study the causal effect of taxes on the disposition effect. The literature has touched upon these two research questions, but the evidence is relatively scarce and we aim to move beyond existing studies in understanding the role of investor-level taxes in trading markets. In particular, our paper improves upon existing identification approaches and provides more credible empirical evidence.

Our data contain daily information about the entire trading behavior (including purchases and sales of stocks) in a panel of approximately 100,000 individual investors for the period 1999 to 2016. These data are confidential and obtained from a large commercial bank in Germany.⁴ Benchmarking with official statistics and the comparable US data set used in e.g. Odean (1998), we show that our sample of investors is representative for the overall population of German investors and similar to US investors. We focus on the trades of stocks in our analyses and explore the effect of taxes on the holding duration and selling probabilities of stocks.

To identify causal tax effects, we exploit the institutional setting of capital-gains taxation in Germany before and after a large reform in 2009. This reform consisted of the following components: i) Before the reform, short-term gains with a holding period of less than one year were taxed at half of the marginal tax rate of the seller. Long-term gains with a holding period of more than one year were tax exempt. Short-term losses with a holding period of less than a year could be used to offset tax-relevant gains (but, unlike the US, not ordinary income). As a result, the pre-reform tax system created a holding-period based intertemporal tax discontinuity in the taxation of capital gains.⁵ ii) This discontinuity was abolished in the context of the reform, and all capital gains are subject to a flat tax of 25% as of 2009.

We start our empirical analysis with a 'raw-data' investigation of the number of realized sales around the holding-period dependent tax discontinuity. For this purpose, we non-parametrically plot the number of sales by holding duration before and after the reform and

³The theoretical effects of taxes on trading behavior are, for example, discussed in Constantinides (1984) and Ivković et al. (2005). We elaborate on the theoretical predictions in the context of our set-up further below in the Introduction.

⁴The type of data are comparable to the frequently used US data set which is propriety data from a discount brokerage house (e.g. Odean 1998; Barber and Odean 2000; Barber and Odean 2001). Our data have for example been used by Leuz et al. (2017).

⁵The term 'intertemporal tax discontinuity' was coined by Shackelford and Verrecchia (2002) to describe 'a circumstance in which different tax rates are applied to gains realized at one point in time versus some other point in time'. We follow this terminology and use it throughout the paper.

separately for losses and gains. Theoretically (following e.g., Constantinides 1984), we expect that tax-sensitive investors realize losses as long as they can be deducted from the tax base (i.e., before passing the intertemporal tax discontinuity). This implies that we should see an increased number of realized losses *before* the tax discontinuity in pre-reform years. On the other hand, tax-sensitive investors should delay the sale of gains until they qualify for the preferential tax treatment. We thus expect an increased number of realized gains *after* the tax discontinuity in pre-reform years.⁶ As the holding period is not tax relevant in post-reform years, we do not expect to find any irregularities in the number of sold gains or losses around the 365-days holding period.

The empirical findings are consistent with these predictions. We observe in pre-reform years that the number of sold losses spikes sharply just before the 365-days tax discontinuity; the number of sold losses in the seven days before the tax discontinuity is roughly 3.2 times as large as the number of sold losses during the seven-day bin just after the tax discontinuity. We further see that investors defer sales of gains until they have reached the 365-days holding period; there is a sharp spike in the number of sold gains in the weeks just after the 365-days tax discontinuity.⁷ In the post-reform years, when the 365-days tax discontinuity was no longer applicable, we see no spikes or other irregularities around the holding period of 365 days. The absence of any spikes whatsoever in post-reform years clearly indicates that the pre-reform spikes are not driven by any non-tax factors and can indeed be attributed to a causal effect of the tax. We see consistent and substantial effects when we consider the value of trades, rather than the frequencies of trades, around the discontinuity.

We then estimate the elasticity of the length of the holding period with respect to the tax rate using a difference-in-bunching approach that exploits data from a time period without tax discontinuity (the post reform years) to construct the counterfactual distribution (see Kleven 2016). The estimated net-of-tax elasticity (calculated based on a notch formula) ranges in between 0.368 for gains and -0.435 for losses. Another way to frame the behavioral response is the corresponding reduction in effective capital gains tax rates: By adjusting the lengths of their ownership spells, investors manage to reduce their effective capital gains tax by 11.3 percent due to behavioral responses. These results suggest that the length of ownership spells is substantially affected by discontinuities in capital-gains tax schedules. This then implies that capital-gains tax rate schedules with lower rates for longer holding periods indeed have the scope to induce stock investments with a longer horizon. The tax-induced commitment to a stock

⁶The model in Constantinides (1984) further predicts that gains should be realized immediately once they qualify for the lower long-term tax rate (or held until death). This implies that we should see a spike in the number of sales to the right of the tax discontinuity (i.e., during the first week after 365 holding-period days). Losses, on the other hand, should be realized as they accrue, according to the model, and their realizations do not necessarily spike anywhere in the short-term-tax period as long as one abstracts from transactions costs of selling the stock.

⁷It is consistent with the model prediction that we observe a larger number of realized losses to the left and a larger number of realized gains to the right of the tax discontinuity. The spike in realized gains that we see in the first week after the tax discontinuity is also consistent with predictions. For losses, the model does not predict that realizations should spike just before the tax discontinuity (see footnote 6). However, this finding is consistent with the notion that the intertemporal tax discontinuity serves as a commitment device for loss-averse investors – as for example described in Shefrin and Statman (1985). We discuss this notion in more detail in section 6.

can benefit the firm, as suggested by previous evidence on the role of tax-sensitive institutional investors in corporate governance and firms' innovativeness (Aghion et al. 2013; Dimmock et al. 2018).

The subsequent parts of our analysis are based on non-parametric regressions which estimate for each day of the holding period the probability that a given stock is sold on this particular day of the holding period.⁸ The non-parametric regressions confirm our previous results. In pre-reform years, we estimate strongly increased selling probabilities just before holding periods of 365 days for losses and just after 365 days for gains. We see no increased selling probabilities around the 365-days holding period in post-reform years.

While the literature mostly focuses on turn-of-the-year (December) trading of losses (see below for more), our findings provide evidence of tax-induced spikes in selling probabilities which are independent of turn-of-the-year effects: we estimate our regression models separately for stocks sold in December vs. the rest of the year and find that taxes matter in all months and not only December. We also estimate non-parametric regressions separately for each year in our sample period. We see spikes around the 365-days tax discontinuity in all pre-reform years, but we never see any spikes or irregularities around the discontinuity in any of the post-reform years. Our results are thus not driven by a few exceptional years in our sample.

Average effects of taxes potentially mask heterogeneity across different types of investors. Our rich data allow us to study several sources of heterogeneity and to understand which types of investors exhibit the largest tax responses, while controlling for factors such as wealth and income. We focus on three sources of heterogeneity which have received considerable attention in the trading literature (e.g., Barber and Odean 2001; Seru et al. 2009; Korniotis and Kumar 2011): age, experience and gender. We find evidence that tax responsiveness is increasing in trading experience (conditional on age and other covariates) and age (conditional on experience and other covariates). This implies that differences in investment behavior commonly associated with age and experience may weaken considerably in the absence of capital-gains taxes. We further find that women are more likely to sell their losses during the days before the tax discontinuity, which may reflect that they value the certainty of a tax loss shield more as they are less (over)confident that losses will turn into gains any time soon. Following Ivković et al. (2005), we also explore heterogeneity w.r.t the magnitudes of gains and losses and find that the tax responsiveness increases in the size of the gains or losses. Hence, the causal effect of taxes matters particularly when there was a major update about the asset's expected returns and capital gains taxes have the largest effect on ownership spells for the outliers in a given portfolio.

A robust finding in the literature on trading behavior is that investors have a larger propensity to realize gains than to realize losses, the so-called disposition effect. Considering that the disposition effect and tax effects potentially run in opposite directions (see intuition

⁸The non-parametric regressions allow us to include control variables and interactions, which make it possible to study heterogeneous effects. In addition, they allow studying the disposition effect in line with the related literature, which usually uses such a regression framework. We present all regression results in graphs which plot for each day of the holding period the coefficient capturing the probability of selling.

above),⁹ we study how the disposition effect interacts with tax effects. In post-reform years (without intertemporal tax discontinuity), we observe the disposition effect on each single day of the holding period, confirming findings in the large literature that document the disposition effect. In pre-reform years, however, we detect the disposition effect only for holding-period days which are sufficiently distant to the tax relevant 365-days tax discontinuity. In the neighborhood to the left of the tax discontinuity, we observe that gains are sold with a much smaller probability than losses. To the right of the tax discontinuity, gains are sold with a greater probability than losses, but this increased probability is much larger than the 'usual' disposition effect that we see in post-reform years and further away from the tax discontinuity. We also find evidence that the tax discontinuity affects the disposition effect even on holding-period days distant from the tax discontinuity. Compared to the post-reform benchmark (without tax relevant discontinuity), the disposition effect in pre-reform years tends to be lower during the first year of the holding period and higher after 365 days holding period have passed on days distant to the 365-days tax discontinuity.

1.1 Contribution to the Literature

We identify the following main contributions of our paper and its empirical findings. First, we combine high-frequency data on the individual investor level with an institutional set-up that allows for clean identification of tax effects on trading behavior. This combination rarely exists in previous literature. Comparable micro-level data are very scarce and many papers studying the tax effects on investors therefore use data which are not ideally suited for the research questions at hand.

One set of papers uses investor-level data obtained from a brokerage house (e.g., Badrinath and Lewellen 1991, Odean 1998, Grinblatt and Keloharju 2001, Barber and Odean 2004, Ivković et al. 2005). The identification of tax effects in this literature is not based on tax reforms, but on the comparison of trading behavior in taxable accounts and non-taxable accounts or on the comparison between trading patterns in December and the rest of the year (where December differences are attributed to taxes because of end-of-year tax planning). However, the isolation of tax effects is difficult in such settings because there likely exist differences in trading behavior between taxable and non-taxable accounts for non-tax reasons, and because the December effect is for example confounded by the momentum effect, window dressing or an overall tendency of investors to 'clean-up' their portfolios towards the end of the year. We move beyond these papers in that we use similar data, but combine them with powerful, direct causal identification using exogenous variation in tax rates and reduced-forms methods with graphical evidence. For example, we are (to the best of our knowledge) the first paper to use the bunching method with individual-level investor data (including information on unrealized losses and gains) and to exploit a reform eliminating the differential taxation of short-term and long-term capital gains. In addition, we have more recent data and a considerably longer time span than most of

⁹The intuition behind the relationship between taxes and the disposition effect is described by Grinblatt and Keloharju (2001, page 603) as follows: 'The disposition effect can be regarded as the opposite of tax-loss selling in that investors are holding onto losing stocks more than they are holding onto winning stocks'.

the existing work, which allows us to show that the tax effects are robust to different economic conditions.¹⁰

Another set of papers uses tax-return data (e.g., Feldstein et al. 1980, Dowd et al. 2015, Jacob 2018, Dowd and McClelland 2019, Agersnap and Zidar 2021, Lavecchia and Tazhitdinova 2021) or firm- and stock-level data (a review is in Hanlon and Heitzman 2010) to study the link between capital gains and investor-level capital-gains taxes. The data in these studies are not ideally suited for consistent inference of individual-level parameters because they either lack information on unrealized gains and/or they are aggregated data which do not allow studying important individual-level aspects of behavioral responses to capital-gains taxes, such as investor heterogeneity, differential effects on losses and gains, and the disposition effect.¹¹

Second, we investigate the causal effect of taxes on the disposition effect. Papers such as Odean (1998) and Ivković et al. (2005) have considered the relationship between the disposition effect and taxes. However, the findings are mostly based on end-of-the year irregularities in the disposition effect and the literature has therefore not settled whether the existing evidence is exhaustive. For example, Ben-David and Hirshleifer (2012) raise skepticism whether the altered disposition effect toward the end of the year, as for example suggested by Odean (1998), is driven by tax effects. In fact, a causal effect of capital-gains taxes is not at all required to rationalize why the ratio of stocks disposed at a gain versus stocks disposed at a loss (i.e., the disposition effect) drops in December for US data. In contrast to our German setting, capital-gains induced losses can be used to offset ordinary income in the US (see below for more institutional differences). A small share of buy-and-hold investors, who only sell one loss-making position in December to benefit from the deduction against ordinary income, suffices to produce a similar December pattern even if capital-gains taxes had no effect at all.¹² This phenomenon adds to the previously mentioned factors that confound December effects.

Overall, the causal effect of taxes on the disposition effect has for these reasons not been established in the literature, and we add to the discussion in that we provide credible empirical evidence for such an effect. We further show that taxes have an effect on the disposition effect (and on trading behavior more generally) throughout the entire year and not only in December. In addition, while the related literature focuses on tax-motivated *loss* selling, our results show

¹⁰Another difference to papers such as Barber and Odean (2004) and Ivković et al. (2005) is that we use data with daily frequency, rather than monthly frequency. The daily data allow us to zoom in the trading behavior along each day of the holding period, which is especially useful in analyzing trading behavior around the intertemporal tax discontinuity.

¹¹Without information on unrealized gains, there is an upward bias in the estimate of the probability to sell a stock and a downward bias in the estimate of average holding periods. With aggregate data, strong assumptions about homogeneous behavior would be required for consistent estimates of individual-level parameters (see, e.g., Gupta 1971, Sasaki 1978 and Pesaran et al. 1989). In addition, tax-return data typically only have aggregated yearly information (and thus lack information about single transactions). Dowd and McClelland (2019) use more frequent tax-return data on sales (weekly frequency whereas we have daily frequency) in a bunching approach. However, they do not have information on unrealized sales (as we do). We additionally contribute relative to Dowd and McClelland (2019) because our long sample horizon allows us to exploit many years and a significant tax reform (whereas Dowd and McClelland 2019 rely on one year of data). Furthermore, we analyze the interaction between taxes and the disposition effect and we investigate within-year heterogeneity of trading behavior.

¹²For example, if the trades of these buy-and-hold investors represent 10% of the sample in December, this would tilt the ratio of disposed gains versus losses from a regular ratio of 55/45 to a December ratio of about 49.5/50.5.

that the disposition effect is strongly affected through tax motivated selling of losses *and* gains.

We also relate to the general (non-tax related) literature on the disposition effect (see Barber and Odean 2013 for an overview). From a methodological perspective, our findings on tax-motivated trading behavior across the entire year imply that it is not sufficient to control for tax effects by allowing for different December effects – which is a very common approach in the large literature. A further contribution relates to the findings in the literature that older and more experienced investors are less prone to the disposition effect (e.g., Feng and Seasholes 2005, Dhar and Zhu 2006, Seru et al. 2010). One can conclude from our heterogeneity analysis that age and gender effects on the disposition effect are driven by tax effects and that heterogeneity in the disposition effect along the age and experience dimensions would be mitigated in the absence of intertemporal tax discontinuities – see our Conclusion for more discussion on this.¹³

Third, our institutional set-up offers various advantages over the institutional rules in the US (which are used by the majority of the related papers). In particular, we are able to compare trading behavior during a time period which featured an intertemporal tax discontinuity (as in the US) to trading behavior during a time period in which this intertemporal tax discontinuity is abolished and replaced with a tax system in which the holding period is not tax relevant. In contrast to the US studies in the literature (e.g., Dowd and McClelland 2019 using one year of data), this allows us to construct a plausible counterfactual and to rule out that irregularities in trading behavior around the 365-day holding period are driven by non-tax factors.¹⁴ In addition, the US setting makes it difficult to isolate the effect of capital-gains taxes because other features of the tax system interfere. In particular, losses can be off-set against ordinary income (up to a certain amount), which gives an incentive to sell losses towards the end of the calendar year (see above for the resulting implications in the context of the disposition effect). Furthermore, there is a step-up of the basis when stocks are bequeathed in the current US system. This implies a zero capital-gains tax rate when holding stocks in the long run (i.e., until death). Our German set-up does not have these features. Finally, the jump in tax rates at the intertemporal tax discontinuity is larger in Germany (before the reform) than in the US, which offers the opportunity to study a situation where transaction costs have relatively less importance (Chetty et al. 2011).

2 Institutional Background

Our analysis exploits features of the capital gains tax regime in Germany between 1999 and 2016 (i.e., the time period of our data set). We focus on equity investment and describe in this

¹³More generally, our findings provide novel evidence on the causal determinants of the disposition effect. As recently suggested by Frydman and Wang (2020), the causes of the disposition effect are still subject to debate, and we are able to add to this debate in that we show that capital-gains taxes have an impact on the disposition effect.

¹⁴This institutional set-up is very appealing for the use of bunching methods: in contrast to more conventional bunching methods, our setting has the advantage that we do not need to estimate a counterfactual distribution that is based on an assumption-intensive extrapolation of regions away from the tax discontinuity to the region in the neighborhood of the discontinuity.

section how the respective capital gains are taxed in Germany. During our sample period, a major reform of capital gains taxation took place in 2009. Both before and after the reform, capital gains from stocks are generally only taxed upon realization (i.e., taxes are due when the stock is sold).

Taxation of capital gains before 2009. Before the reform, the tax treatment of capital-gains was conditional on the duration for which an asset was held before being sold. The gains and losses of assets sold within a holding period of 365 days or less were subject to taxation. This tax was commonly referred to as a 'speculation tax'. The tax rate was a function of the personal income-tax rate (PIT) of the investor. The PIT generally depends on the sum of all income types (wage income, self-employment, capital-gains, etc). The top PIT rate is constant at 42% since 2005, and it applies for overall taxable income (i.e, from all sources) greater than approximately 52,000 EUR (for example, the top tax rate kicks in at 52,152 EUR in 2008). The entry tax rate in 2008 was 15%. Losses from sales with a holding period of ≤ 365 days could be used to offset gains from capital gains. Losses from capital gains could not be used to offset other types of income (such as ordinary income). Between 2001 and 2008 the so-called half-income method applied: one half of the gains/losses from capital gains with holding periods ≤ 365 days were subject to the tax.

For illustration, consider a fictitious investor who is subject to the PIT rate of 42%. She realizes gains worth 2000 EUR from shares that she had held for less than 365 days, and she sells other shares within the 365-days holding period at a loss of 200 EUR. The resulting capital-gains tax liability for this investor then was $1/2 \times (2000 - 200) \times 0.42 = 378$ EUR.

Long-term gains realized from selling assets held for more than 365 days were tax exempt. This also implied that long-term losses from selling assets held for more than 365 days could not be used as a tax shield for capital gains from short-term holdings.

This system of capital-gains taxes applied to assets such as stocks (as long as the investor is not a substantial shareholder), funds, certificates (except guarantee certificate) and capital gains from bonds (except zero bonds).

Taxation of capital gains since 2009. The tax treatment of capital gains was substantially reformed as of January 2009. In stark contrast to the old system, the holding period of assets is not tax relevant anymore. That is, the holding-period based 'speculation tax' was abolished in the context of this reform. It was replaced by a system where all capital gains and capital losses (independent of holding duration) are subject to a flat tax of 25% or, if the PIT rate is smaller than 25%, the PIT rate. That is, the tax on capital gains/losses is $\min(25\%, \text{PIT rate})$. Losses can be used to offset gains. The half-income method was abolished.

Consider again a fictitious investor who is subject to the top PIT rate (which is $> 25\%$) and who has capital gains of 2000 EUR and capital losses of 200 EUR. Her tax liability after the reform is independent of the holding periods of the underlying assets and sums up to: $(2000 - 200) \times 0.25 = 450$ EUR. Any conditionality of the tax schedule on the length of the holding period was abolished. The old pre-2009 tax rules applied to all assets bought before

January 2009 (resulting in grandfathered assets).

3 Data

3.1 Data Description and Summary Statistics

We use individual investor and portfolio data from a large German online bank. The full-service bank has more than half a million customers and operates across the entire country. We obtain a sample of about 110,000 investors which is randomly drawn from the bank’s client pool. Variants of this data set were for example used by Leuz et al. (2017) and Loos et al. (2020). For each investor, we have the complete trading history for the period January 1999 to May 2016. These data allow us to construct an individual-level panel of daily trading activities over almost 18 years. Trading information in the data include type of traded asset, transaction volumes, prices, order types (with or without limit) and dates for purchases and sales. We further have investor information on age, gender, zip code of residence, marital status, employment type, and for how long the investor has had the trading account. In addition, the data include self-reported information about education, income (in categorical ranges), total wealth and risk tolerance.

For the purpose of analyzing equity trading, we restrict the sample to all investors who have purchased at least one stock during the sample period. This leaves us with about 93,000 investors. These investors bought about 8.4 million share packages with an overall purchase value of 49 billion EUR during our sample period (the unit of analysis in most of our analyses will be a share package; see section 3.2 below for a definition and more information). Table A.1 provides summary statistics for all investors in our sample. The average portfolio value (incl. all assets in the portfolio) is 51,725 EUR and the investors in our sample make on average roughly 78 trades in total over the observation period. The average monthly portfolio turnover¹⁵ was 10.86 percent, which implies that most investors have quite active accounts. Most investors in our sample are male (83%) and their average age (by the end of 2015) was 52 years. 6% work in the financial sector and 16% of our sample is self-employed. The average investor in our sample has held the account at this bank for more than 13 years (as of the end of 2015). The share of investors in our sample with a PhD-level degree is 6%, whereas the share over the entire German population is only about 1.5%. This is in line with prior evidence showing that individuals with investment portfolios are more educated than the population average (van Rooij et al. 2011; Cole et al. 2014; Leuz et al. 2017).

We investigate if our sample is representative for the population of investors and does not only include special groups of investors or play-money accounts. For this purpose, we provide several comparisons of our data sample with i) the German population of investors and ii) with other comparable data sets that have been used in the literature (these comparisons build on Leuz et al. 2017 who use a very similar data set). The portfolio value in our sample (51,725 EUR) is very comparable to the number that the German central bank reports to be the average

¹⁵Monthly portfolio turnover is calculated as in Barber and Odean (2001) as one-half of the monthly sales turnover plus one-half the monthly purchase turnover. Sales (purchase) turnover is defined as value of shares sold (purchased) divided by the portfolio value in the beginning of the month.

portfolio value of German equity investors: 48,000 EUR in 2013 (Deutsche Bundesbank 2013). We further construct a variable that measures the ratio of portfolio value over annual income for our data and benchmark this ratio with official statistics reported by the German Federal Office. As income in our data set is reported in several categorical ranges, we use either the midpoint or the lower end of each range as a proxy for investor income. Using the midpoint, the mean ratio of the average portfolio value (over the entire sample period) to annual income is 1.3. Using the lower ends of each income range as a proxy for annual income, this ratio is calculated to be 1.2. These numbers are very close to the ratio of total financial assets to gross household income in the German population, which is 1.1 (German Federal Bureau of Statistics 2008a; German Federal Bureau of Statistics 2008b).¹⁶ In addition, the ratio of the median portfolio value to median gross income for the German investors surveyed by Dorn and Huberman (2005) is 0.6 and it turns out to be 0.6 for our sample as well.¹⁷ Overall, these comparisons allow us to conclude that our investor sample is representative of the population of German investors and should not be significantly biased by play money accounts.

Demographic and portfolio characteristics of the investors in our sample are also well comparable to the well-established investor data set used by, for example, Odean (1998) and Barber and Odean (2001). Their data are obtained from an US online brokerage house and are similar in spirit to the data that we use. For example, average age (50.4 vs 52.26) and the share of males (79% vs 86%) is fairly similar across these data sets. Furthermore, the average portfolio value of about 51,000 EUR is in the same order of magnitude (considering the different time periods) as the average portfolio value of 47,000 USD that is reported in Barber and Odean (2001).¹⁸

We further investigate if trading behavior is different in December compared to other months of the year. It is well-documented in the literature that investors in the US tend to increase their loss selling towards the end of the year (e.g., Odean 1998). We do not observe such a pattern in our data for Germany. Table A.1 shows that we do not observe an increased accumulation of sales in December, neither for losses nor for gains. As the table indicates, seven percent of all annually sold gains and eight percent of all annually sold losses occur in December; this is what one expects if sales were equally distributed across all months of the year.

3.2 Unit of Analysis

We are particularly interested in the number of stock realizations around the intertemporal tax discontinuity.¹⁹ To study the number of stock realizations, we use 'share packages' as the unit

¹⁶We manually calculate this value from total financial assets and the monthly gross income reported in the above sources.

¹⁷We manually calculate this from the values given in Tables 1 and 2 of Dorn and Huberman (2005, pages 443 and 447).

¹⁸The EUR-USD exchange rate throughout our sample period was at 1.16 in Jan. 1999 and 1.11 in May 2016.

¹⁹That is, our primary interest is not regarding the number of investors trading around the tax discontinuity (although we analyze this too in one series of analyses) and we thus do not employ the investor as the unit of analysis.

of analysis throughout most of the paper.²⁰ Our concept of a share package is very similar to 'round-trips' that are frequently used in the literature. In contrast to usual round trips, we also consider packages which are not sold within our sample period. We therefore use a different term to describe our unit of analysis. A key feature is that one 'share package' is independent of the number of shares that are included in this share package. For example, if an investor buys 100 shares and sells this 'package' of 100 shares ten days later, we generate one observation with a holding period of 10 days (see below for more on the measurement of holding period). If another investor buys 10,000 shares and sells her 'package' of 10,000 shares 10 days later, we also generate one observation with a holding period of 10 days. We selected this unit of analysis in order to avoid that our results are driven by the behavior of a few large-scale investors or penny stocks. Our approach reflects that we are eventually interested in the individual behavior of investors and we want to avoid that the individual behavior is weighted with the number of shares that an individual investor moves. In the previous example (one investor selling 10 and one selling 10,000 shares), both of these investors are given the same weight in our analysis because we are interested in the tax-induced trading behavior of both these investors. If single shares were the unit of analysis, the behavior of the smaller 10-shares investor would be almost negligible relative to the behavior of the bigger 10,000-shares investor.

3.3 Measuring the Holding Period

We measure the holding period as the difference in days between purchase date and sales date of a share package. For example, if a fictitious investor buys five shares of some firm on the second of October and sells all of them on the 15th of October, this would result in one observation with a holding period of 13 days. If the first purchase of a share package occurred outside our sample period (i.e., prior to January 1999), we cannot calculate the holding period and have to drop the share package from our analysis.

If there are multiple buys before the first sale occurs we apply the first-in-first out principle (which is in line with the German tax law). For example, if an investor buys two shares on Oct 5, ten shares of the same firm on Oct 10, and then sells all 12 shares on Oct 20, we generate two observations with holding periods of 15 days and 10 days, respectively. If the sale of a share package takes place in parts at different dates, we create one observation for each sale. For example, consider an investor who buys five shares on October 2, then sells three of these shares on October 4 and two shares on October 15. We then create two observations: one with a holding period of two days and the other with a holding period of 13 days.

Sometimes shares change their ISIN (international identification number) or shares are splitted or reverse splitted. We account for this by using hand collected data for ISIN changes and data on splits and reverse splits from datastream.²¹ In cases in which shares have been splitted or reverse splitted, we adjust prices such that purchase and sale price are comparable.²²

²⁰We sometimes also present results for purchase-price weighted analyses; see results section 4.1.2.

²¹We use the data to identify (reverse) splits which were not reported in datastream and ISIN changes. For this purpose cases, we manually check whether there was indeed an ISIN change or (reverse) split.

²²Since the total value of a position is unaffected by the split or reverse split, the price basis before and after

Stocks for which there is a disparity between recorded trades and month-to-month positions are removed from the sample.

3.4 Final Sample

Our analysis is based on several million share packages. For the years before the reform, we include 2.74 million observations of appreciated share packages (gains) and 2.47 million depreciated share packages (losses). In the after-reform years, we have 1.34 million appreciated share packages and 0.85 million depreciated share packages. Restricting the sample to half a year before and after the intertemporal tax discontinuity, we rely on 313,000 appreciated share packages and 380,000 depreciated packages in the pre-reform period, and 212,000 gains and 136,000 depreciations in the post-reform period.

4 Taxes and Trading Behavior

This section presents the various research designs to study the causal effect of taxes on trading behavior (in particular on holding periods and the probability to sell an asset) and the corresponding results. All of our empirical findings are presented in graphs which aim to visualize the effects and make them approachable.

4.1 Raw Data: Number of Trades around the Discontinuity

4.1.1 Research Strategy

The starting point for our analyses are figures in which we plot the number of share packages sold (in weekly bins of seven days) in a one-year window around the holding-period tax discontinuity. That is, we plot the number of sales in each week in the year around the tax discontinuity. We do not yet normalize the data or use any parametric methods here; this exercise thus presents a non-parametric look at the 'raw' and unmanipulated data.

We plot the number of trades around the 365-days tax discontinuity separately for years before and after the 2009 reform and separately for gains and losses. Since the holding period became tax irrelevant in the course of the reform, we expect a smooth distribution of trades around the 365-days tax discontinuity for the years after the reform. A causal effect of taxes on trading behavior would imply that we see, in pre reform years, an increased number of trades of appreciated assets (gains) in the weeks after the 365-days tax discontinuity, and an increased number of trades of depreciated assets (losses) in the weeks before the tax discontinuity.

To investigate if potential tax effects are due to a few large tax sensitive investors who sell many share packages around the tax discontinuity, we also plot the number of *distinct* investors who sell share packages in a given week of the holding period. For this purpose, and analogous to the above strategy, we group the number of distinct investors who sell a share package in

the split is not the same anymore. For example, consider 100 shares with a value of 200 Euro that are splitted by 2. Without adjustments, the price before the split is 2 Euro while it is just 1 Euro after.

weekly bins and plot the number of investors in each bin during the one-year window around the tax discontinuity.

4.1.2 Results

Figure 1a depicts the number of traded *gains* (i.e., appreciated share packages) in weekly bins around the intertemporal (365-days) tax discontinuity separately for pre-reform and post-reform years. The red vertical line at zero marks the 365-days holding period.

In pre-reform years, in which the 365-days tax discontinuity was tax relevant, the number of gains that are sold spikes sharply in the first week after the 365-days tax discontinuity. The number of sold gains in this first week after the discontinuity is more than 2.5 times as high as in the week before the 365-days tax discontinuity. In week 2 after the tax discontinuity the number of sales is roughly 1.8 times as high as in the week before the reform. This trend then continues in subsequent weeks: the number of sold gains remains higher than before the tax discontinuity, but the difference becomes smaller as we move further to the right from the tax discontinuity.²³

Is the spike in the number of realized gains driven by the capital-gains tax discontinuity? In post-reform years, in which the 365-days cut-off is not tax relevant, we see a smooth development of the number of sales around 365 days. Specifically, the number of sold gains does not exhibit a spike just to the right of the tax discontinuity. This is clearly indicative that the large spike in pre-reform years is driven by the capital-gains tax system. We will compare the number of sales in pre and post years further below when quantify the tax effect using the bunching approach.

Figure 1b presents the equivalent plot for the number of sold *losses* (i.e., depreciated share packages). In pre-reform years (with tax-relevant 365-days tax discontinuity), we see a sharp spike in the number of realized losses in the week just before the 365-days tax discontinuity. The number of sold losers is more than 3 times as large in the week before the tax discontinuity than in the week just after the tax discontinuity. In week -2, the spike is still clearly visible but considerably smaller than in week -1; the number of sold losers is about 1.7 times larger in week -2 compared to the week after the tax discontinuity. Importantly, we see a smooth development in the number of realized losses around the 365-days tax discontinuity in post-reform years where crossing the 365-days holding period does not have any tax implications. The spike in pre-reform years, along with the absence of any spike in post-reform years, provides clear evidence that the tax discontinuity affects trading behavior.

²³Why do we see that some gains are sold shortly before the discontinuity? We see the following reasons why investors might sell an appreciated share even on the day before it can be sold tax free. First, loss carryforwards: if the investors has sufficient loss carryforwards, she can sell an appreciated share tax free even if still the long term rate applies. Second, time discounting: for example the investor needs liquidity in the time before the tax discontinuity. Third, expected prices: if the investor assumes that the price will drop strongly on the day after the tax discontinuity, selling on the days before the tax discontinuity might be advantageous for her. Fourth, risk aversion: even if the investor assumed that prices remain constant in expectation, it might be optimal for her to sell before the tax discontinuity in cases for which the expected variance or covariance with the portfolio is sufficiently high. Fifth, inattention or bounded rationality: investors might simply be inattentive to the regulation or do not respond optimally to the regulation (because of limited rationality).

Overall, our findings for both losses and gains are consistent with the notion that investors try to realize losses within the holding period that allows using them as a tax shield, whereas investors defer the realization of gains until they are tax free.

So far, we explore the number of share packages around the tax discontinuity. To shed light on the size of the effects in terms of monetary Euros, we also consider the value of the share packages by looking at purchase-price weighted share packages. Doing so, we find that the tax effects are also substantial in terms of Euro values. For share packages with losses, we find, in pre-reform years, that the value of share packages traded in the week before the discontinuity is 3983 EUR, while this value is substantially lower in the week after the discontinuity (3200 EUR). Considering the value of traded gains, we also find substantial differences around the discontinuity. The value of traded gains in the pre-reform period is 3688 EUR in the week after the discontinuity and substantially lower, 2892 EUR, in the week before the discontinuity.

Number of distinct selling investors in weeks around the tax discontinuity. The previous results showed the number of appreciated and depreciated share packages around the intertemporal tax discontinuity. Are the spikes in the number of sales around the discontinuity driven by a few investors who are tax aware and sell many of their share packages around the tax discontinuity? We shed light on this question by plotting the number of *distinct* investors who trade in a given week. As before, we plot the weekly numbers separately for investors who trade gains and losses, as well as for pre and post reform years.

Figures A.1a and A.1b present the plots for gains and losses, respectively. In pre-reform years, the number of investors selling gains spikes sharply in weeks to the right of the tax discontinuity and the number of distinct investors trading losses spikes sharply in weeks to the left of the tax discontinuity. We see no spikes in post-reform years in which the 365-days tax discontinuity is not tax relevant. The spikes in pre-reform years, along with the absence of spikes in post-reform years, again indicates a causal tax effect.

Overall, this exercise suggests that the sharp spikes in the number of share packages above is not driven by a few tax-sensitive investors selling many share packages around the tax discontinuity. Apparently, many different investors respond to the tax incentives in a way that is consistent with our expectations. We study different sources of potential heterogeneity in tax responses among different investors further below.

4.2 Difference-in-Bunching

4.2.1 Research Strategy

In a next step, we employ bunching methods to quantify the tax effects in a one-year window around the tax discontinuity. Bunching approaches go back to Saez (2010) and are now commonly used (see applications such as Chetty et al. 2011, Kleven and Waseem 2013, Best and Kleven 2018 or Mortenson and Whitten 2020, and the recent overview by Kleven 2016). We apply a difference-in-bunching approach using the sales distribution in the post-reform periods as a counterfactual for the pre-reform distribution (as in e.g., Brown 2013; also see Kleven

2016).

To make the pre and post reform distributions comparable and to obtain a good counterfactual, we account for level differences in the number of sales before and after the reform. We divide all weekly counts by the respective total number of share packages which are held at the start of the one year window (recall that we deliberately did not do this in the previous non-parametric 'raw-data' approach). Unsold share packages are, of course, included in the total count to arrive at unbiased fractions of shares sold. We apply this procedure separately for gains and losses. We therefore need to determine whether an unsold share package is treated as a gain or a loss. Unsold share packages are categorized as a gain or a loss based on the latest price relative to the purchase price.

In many bunching applications, the counterfactual distribution is estimated through predicting the distribution in the region close to the tax discontinuity using the distribution in the region further away from the tax discontinuity. In contrast to this conventional approach, we do not have to estimate a counterfactual and instead rely on actual post-reform data. Our approach is advantageous to the conventional approach because it does not rely on any functional form assumptions and assumptions about 'excluded regions' when calculating the counterfactual. In addition, the conventional way of estimating the counterfactual assumes that the distribution further away from the tax discontinuity is unaffected by the tax discontinuity at 365 days – this assumption then allows to estimate the counterfactual distribution based on points further distant to the discontinuity. Given our data-based counterfactual, we do not need to make this assumption either. This is particularly appealing in our context because it may well be that the salient intertemporal tax discontinuity, for example, induces some investors to realize their gains only after 365-days, although they would have realized the sale considerably earlier in the absence of the tax discontinuity. Our hazard-rate Figures suggest that this is indeed the case; pre-reform hazard rates distant from the tax discontinuity are different from post-reform hazard rates distant from the tax discontinuity.

To make the general point that the spikes that we see in our data are tax effects, it suffices to show that after the reform the distribution becomes smooth around the 365-days threshold – i.e. continuous and without noticeable changes in the derivatives. An appropriate elasticity for the tax effect can be calculated because the counterfactual distribution is available for the complete domain without any further assumptions/ polynomial approximations.

The identifying assumption for calculating an elasticity in our set-up is that the post-reform distribution (without tax discontinuity) is a plausible counterfactual for the pre-reform distribution (to which the tax discontinuity applies). Looking at our plotted Figures (see below), this assumption seems plausible: The post-reform distribution is very similar to the pre-reform distribution except for the spikes around the tax discontinuity in pre-reform years. In addition, we show below in a year-by-year approach that we have a uniform picture across all post-reform years, suggesting that the financial crisis did not affect the sales distribution significantly.

We proceed as follows to calculate elasticities based on a notch formula. For gains, the excess mass of sales after the tax discontinuity is calculated from the difference in the integrals between the pre-reform and the post-reform distribution of sales with a holding period above

52 weeks.²⁴ The counterpart is the missing mass of sales before the tax discontinuity: the integral over the difference between the post-reform and the pre-reform distribution of sales with holding periods of up to 52 weeks. The width of the missing mass integral indicates the extent of holding period adjustment due to the tax discontinuity, which is inferred by equating the missing mass to the observed excess mass. Without further structural assumptions, the net-of-tax elasticity can then be estimated from the reduced form as described by Kleven and Waseem (2013).²⁵ Similarly, one can also arrive at the elasticity of holding periods for stocks with accrued losses - except, of course, that the excess mass is in this case on the left hand side of the tax discontinuity and the missing mass on the right hand side. Standard errors for the bunching analyses are calculated by bootstrapping the estimates clustering stock holdings by their owner.²⁶

Recall from Section 2 that, in pre reform years with tax discontinuity, the applicable tax rate τ for realized stock trades with a holding period of less than one year was one half of the personal income tax (PIT) rate of the investor, and the applicable tax rate was zero for holding periods of more than one year. That is, the tax rate falls from half of the PIT rate to zero at the tax discontinuity. The discontinuity-induced change in the PIT rate that we use to calculate our preferred elasticity is the difference between half of the top PIT rate (i.e., half of 42%) and zero.²⁷

4.2.2 Results

The Difference-in-Bunching results are presented in Figures 2a (for gains) and 2b (for losses). As described above, we use the post-reform years (without tax discontinuity) as the counterfactual distribution for the tax-affected pre-reform years. In the Figures, the vertical red line depicts the 365-days holding period and the blue and red line present the weekly bins for the pre- and post-reform periods, respectively. The patterns in both Figures are (not surprisingly) similar to the patterns that we saw above in the Figures that simply plot the number of sales. In particular, the density of realized gains spikes sharply in the week after the 365-days tax discontinuity in pre-reform years and no such spike is observed in post-reform years. The density of realized losses has a large spike in the week before the tax discontinuity in pre-reform years and, again,

²⁴The upper limit of the integrals is given by the point at which the – intermittently elevated – pre-reform distribution overlaps again with the post-reform distribution.

²⁵In particular, the elasticity is estimated by $\left(\frac{\Delta H^*}{H^*}\right)^2 * \frac{1-\tau^*}{\Delta\tau^*} * \frac{1}{2}$ where τ^* is the tax rate and $\Delta H^* = H^* - H^l$ is the width of the missing mass integral. The integral's upper limit H^* is the 52 week threshold at which the tax discontinuity applies. The integral's lower limit H^l is determined by equating the missing mass with the observed excess mass.

²⁶We draw 1,000 random samples on the investor level from the original data with replacement. We define the standard error of the estimated elasticity e as the standard deviation in the distribution of estimates of e .

²⁷We do not observe the taxable income of investors in our data. We acknowledge that a few investors might not be subject to the top marginal tax rate, and are instead subject to a lower marginal tax rate. However, it is likely that most investors have other sources of income and that the majority of investors are beyond annual taxable income of approximately 50,000 EUR (which is the amount at which the top rate kicks in). In addition, note that the calculated elasticity increases if the elasticity calculations are based on a smaller jump in tax rates at the tax discontinuity. In the results section below, we also report elasticities for the extreme scenario that all investors are subject to half of the marginal entry tax of 15%.

there is no spike in post-reform years.

Where does the excess mass come from? Are the spikes that we see to the right (for gains) and to the left (for losses) of the tax discontinuity 'fed' by sales that investors would have realized before or after the tax discontinuity in the absence of the tax? For gains, we see that the mass mostly comes from the left of the tax discontinuity; this suggests that investors delay the sales of gains until they qualify for tax exemption. For losses, we see that the mass of investors mostly comes from the right side of the tax discontinuity; this suggests that investors move forward the realization of sales in order to count them against their tax-relevant gains.

The main purpose of our Difference-in-Bunching approach is to quantify the magnitude of the tax effect and to estimate an elasticity of the holding duration with respect to the net-of-tax rate. In other words, we aim to calculate the percentage change in holding-period days in response to a one-percent change in the tax rate. Using the notch equation described above and the top PIT rate, we estimate our preferred elasticity of 0.368 (standard error: 0.06) for gains and an elasticity of -0.435 (standard error: 0.03) for losses. Assuming that investors are subject to the entry PIT rate (see footnote 27), we calculate elasticities of 1.206 (standard error: 0.2) for gains and -1.425 (standard error: 0.12) for losses.

4.3 Effective Capital-Gains Tax Rates

4.3.1 Research Strategy

Another approach to quantify the behavioral response to the intertemporal tax discontinuity is based on effective tax rates. This approach uses a similar intuition as the difference-in-bunching approach in that it also uses the post-reform distribution (without tax discontinuity) as a counterfactual. The basic intuition is as follows: i) We calculate effective tax rates for the pre-reform years (with tax discontinuity) based on the actual pre-reform sales distribution for gains. ii) We calculate effective tax rates in pre-reform years absent any behavioral response. To do so, we estimate effective tax rates by applying the tax parameters of the pre-reform years to the post-reform distribution. The difference between these two effective tax rates serves as a quantification of the behavioral response to the tax discontinuity. Note that this approach measures a lower bound for the overall tax-induced behavioral response because it is only based on the realization of gains and does not include the behavioral response in the context of loss selling.

The details of the effective-tax-rate approach are described in the following. We estimate a weighted average marginal effective tax rate (METR) based on the procedure used in Ivković et al. (2005). First, we calculate effective tax rates for each week of the holding period based on the following equation, which goes back to Protopapadakis (1983):

$$e^{(1-\delta)gT} = e^{gT} - \tau_{cg} * (e^{gT} - 1). \quad (1)$$

The effective tax rate δ is implicitly defined by the equation above. It depends on the holding period T , the gain accrual rate g and the statutory tax rate on realized gains τ_{cg} . Second,

we weight the holding period specific effective tax rate with the actual empirical distribution of realized appreciated share packages in pre-reform periods.²⁸ We further assume a gain accrual rate of 0.25% per week, which corresponds to the intermediate rate of 1% per month used in Ivković et al. (2005) (recall that Ivković et al. 2005 use monthly data whereas we use weekly data in this approach). Furthermore, we use a statutory tax rate of 21% for our calculations. This is the statutory rate a top income tax payer in Germany had to pay on capital gains in pre-reform years (see our description of the institutional background in section 2). Consistent with all our analyses, we estimate the METR conditional on holding a share package for at least 26 weeks.

In order to explore by how much the resulting effective tax rate is affected by behavioral responses to the tax discontinuity, we use the same procedure to estimate the METR but basically apply the post-reform distribution to the pre-reform tax parameters. Technically, we weight the pre-reform effective tax rates with the post-reform distribution of sold share packages.

4.3.2 Results

Using the strategy described above, we estimate marginal effective tax rates (METR) for the pre-reform period. To derive a measure for the behavioral response, we do this for the actual sales distribution of gains in pre-reform years – as shown in the discussed Figure 2a (this is the Bunching Figure) – and a counterfactual distribution that abstracts from behavioral responses. We calculate the METR for investors in pre-reform periods to be 6.3%. This is the actual effective tax rate calculated based on the actual pre-reform period that includes behavioral responses. Using the post-reform periods as a counterfactual, our calculations suggest that in absence of a behavioral response the METR in pre-reform periods would be 7.1%. This number is 12.7% ($= (7.1 - 6.3)/6.3$) larger than the METR that includes behavioral responses. In other words, investors would face a METR that is 12.7% larger if they did not exhibit any behavioral responses. Or put differently, the behavioral response reduces the effective tax rate by 11.3% ($= (6.3 - 7.1)/7.1$). Note, again, that this measure of the behavioral response is only based on the sales distribution for gains and therefore does not include the tax-induced behavioral responses of losses.

4.4 Non-parametric Regressions

4.4.1 Research Strategy

We complement our analyses with non-parametric regressions which estimate for each day of the holding period the probability that a given share package is sold on this holding-period day

²⁸There are two typical caveats to this procedure which, as we argue, are not critical in our setting. First, the above formula does not consider the additional tax burden on real capital gains created by inflation. Second, since our observational window is limited for longer holding periods, there is right censoring in the empirical distribution of sold share packages. Both caveats do not constitute big concerns in the pre-reform period (which we focus on here) since the statutory capital gains tax rate, and therefore also the effective tax rate, is zero for holding periods longer than one year.

(similar to e.g., Hartzmark 2014, Chang et al. 2016 and Frydman et al. 2017).²⁹ For this purpose, we set up our data set such that it contains one observation per share package, individual investor and day of the holding period.³⁰ For example, this would give us 11 observations for a share package that an individual investor has held for 10 days (0, 1, 2, ..., 10). We then create a dummy variable – *Sell* – that indicates for each day of the holding period if the asset was sold on this respective day. We merge the resulting dataset with daily price information for all assets, extracted from *Datastream*. For each day of the holding period, we estimate separate regressions in which we regress the *Sell*-dummy on a constant.³¹ The resulting coefficient for the constant then describes the probability that a share package is sold on this particular day of the holding period. We again focus on the year around the tax discontinuity. Formally, we estimate the following regression separately for each day of the holding period:

$$Sell_{ijd} = \beta_0 + \varepsilon_{ijd}, \quad (2)$$

where indices indicate a share package i of individual investor j on calendar-day date d . Note that we would not yet need indices j and d for this regression model here, in which we simply regress the sale dummy for a share package j on a constant. However, further below we will introduce investor-level (j) variables, which partly vary by calendar-day date (d), and we therefore already introduce indices j and d at this point. All standard errors are clustered on the level of the individual investor.

We estimate the sets of (daily) regressions separately for pre-reform and post-reform years to see if selling probabilities around the 365-days holding-period tax discontinuity are different before and after the reform. In light of the differential tax incentives for gains and losses, we further run separate regressions for gains and losses. As a result, we thus have estimates for all four combinations of pre and post years as well as gains and losses.³²

For illustrative purposes, we plot the estimated β_0 coefficients for each day of the holding period (separate plots for gains and losses, and post and pre reform). The β_0 coefficients measure the probability of sale on a given day of the holding period. This procedure provides graphical evidence whether the selling probabilities are affected by taxation. Our main regressions are also displayed in table form with exact coefficients and standard errors. In contrast

²⁹We use data with daily frequency in our regressions. Recall that we used bins of seven days in the non-parametric figures and bunching approaches above. The reason there is a mechanical pattern in the daily data: since it is not possible to trade on weekends, some day-measured holding periods occur more often than others. For example, a seven day holding period is possible for sales made on all five weekdays, whereas a four days holding period is only possible for sales made on Mondays, Thursdays or Fridays. While this is no concern in the regression approach, it is necessary to use weekly data which 'smooth away' this mechanical pattern in the bunching analyses.

³⁰To avoid selection in assets because of missing prices in datastream, we assign the last observed price to shares where the price is missing. This is the case for about 10% of all assets in our sample.

³¹This estimation set-up with a dummy variable being regressed only on a constant motivates the label 'non-parametric regression'.

³²We exclude grandfathered assets from the regressions. Two pieces of evidence suggest that grandfathered assets do not affect the patterns of tax effects that we find. First, the grandfathered assets are included in the preceding analyses ('raw-data' Figures and bunching approach) and the results there are very consistent with our regression results. Second, we estimate our non-parametric regressions separately for each year of our sample period (as we describe below) and find consistent tax effects throughout.

to the bunching approach, the regression approaches allow us to include control variables and estimate heterogeneous treatment effects across investors in an easy fashion with corresponding test statistics. In addition, non-parametric regressions facilitate comparisons with the related literature that typically uses such regression approaches as well.

To complement our main regressions (which bundle all pre-reform or all post-reform years), we also provide non-parametric regressions separately for each year in our data sample. Relating to the large literature focusing on turn-of-the-year trading in December, we further estimate the main non-parametric regressions separately for share packages that are sold in December and share packages that are sold throughout the rest of the year.

Heterogeneity w.r.t. investor characteristics. Our dataset includes several demographic variables which allow us to study heterogeneity across different type of investors. We use the regression setup for this purpose and add investor-level characteristics to the share-package level data. We then run the following type of regression for each day of the holding period to study heterogeneity:

$$Sell_{ijd} = \beta_0 + \beta_1 Demogr_{ijd} + X_{ijd}\beta' + \varepsilon_{ijd}, \quad (3)$$

where i , j and d again indicate share packages, investors and calendar-day dates, respectively. Variable $Demogr$ is the respective variable along which heterogeneous effects may occur. We focus on three different sources of heterogeneity which have received attention in the trading literature (e.g., Barber and Odean 2001; Seru et al. 2009; Korniotis and Kumar 2011): age of the investor, investor experience and gender. To measure experience, we rely on a variable which measures for how many years an investor has held the account at the bank from which we obtain the data (this is comparable to the measure used in e.g. Korniotis and Kumar 2011 or Bhattacharya et al. 2012). In the 'investor experience' regressions, β_1 estimates the increase in selling probability as experience increases by one year. In cases where the focus is on age, β_1 indicates the effect of one additional year of age on the selling probability. Gender is coded such that β_1 in the 'gender regressions' measures the difference in selling probability on a given day for male investors relative to female investors.

In all these regressions, we condition on a set of observable control variables which are all included in vector \mathbf{X} . These control variables include: age, investor experience, gender, birth year, income category, wealth category, dummies indicating employment in the financial sector, having a doctoral degree, and being self-employed. The respective heterogeneity variable of interest, $Demogr$, is of course omitted from vector \mathbf{X} in the respective regression (for example, in cases where we are interested in gender heterogeneity, the gender variable is included in $Demogr$ and not included in \mathbf{X}). In regressions in which we are interested in heterogeneity w.r.t. age (i.e., variable $Demogr$ is age), we exclude birth year from the vector of control variables because age and birth year are strongly correlated and we do not want to 'control away' cohort effects when investigating age heterogeneity. The controls for income and wealth variables also capture potential heterogeneity in tax rates.

The corresponding summary statistics for all variables used here are reported in Table A.1. Including these control variables for example implies that the effect of investor experience is going to be conditional on age. In our results graphs, we plot the β_1 coefficients of this regression. These show if the selling probabilities are different across the groups, and we are of course particularly interested in the differential selling probabilities around the intertemporal tax discontinuity.

Heterogeneity w.r.t. magnitude of gains and losses. The regression set-up also allows us to estimate if potential tax effects depend on the magnitude of the gain or loss of an investor. This is potentially relevant because an investor with a large loss faces larger incentives to sell the share package before the 365-days tax discontinuity because deducting a large loss reduces the tax base by more than a small loss. In addition, if the loss is only small the investor might want to wait and see if share package prices rise. Equivalently, a large gain would trigger a larger tax liability, which increases the incentive to sell a gain after the tax discontinuity. For this purpose, we include an additional variable into the above base regressions. This additional variable measures the percentage change in the value of the share package. In this context, we estimate the following regressions for each day of the holding period t :

$$Sell_{ijd} = \beta_0 + \beta_1 Change_{ijd} + \varepsilon_{ijd}, \quad (4)$$

where the *Change* variable describes the change between the share price at holding-day t and the purchasing ($\frac{p_{ijt} - p_{ij0d}}{p_{ij0d}}$). To avoid that the regression results are driven by extreme outliers which could be caused by mistakes in the price databank, we exclude observations for which the price change is not included within the first and 99th percentile. The constant in these regressions describes the selling probability for share packages without a price change, while $(\beta_0 + \beta_1)$ estimates the selling probability for changes of size 1 (that is, price changes of 100%). β_1 measures the difference between the selling probabilities of share packages without any change and a large change of 100%. We again estimate these models separately for losses and gains and pre and post reform periods. In our graphs, we plot the β_1 coefficients.

4.4.2 Results

We present the results of our main non-parametric regressions in Figures 3a (for gains) and 3b (for losses). The red vertical line again indicates a holding period of 365 days. The blue line plots the daily-estimated coefficients for the selling probability of either gains or losses in pre-reform years. The red line plots the equivalent coefficients for post-reform years. That is, we plot the β_0 coefficients (i.e., the coefficients for the constant) of regression equation 2 in these Figures. The shaded area around the coefficients indicates 95% confidence intervals.

The results are very much consistent with the patterns that we saw in the preceding analyses. In particular, we see in pre-reform years that the probability to sell an appreciated share package spikes sharply during the holding-period days just after the 365-days tax discontinuity, whereas the probability to sell depreciated share packages spikes sharply during the days just

before the 365-days holding period. We do not see any spikes in selling probabilities around the tax discontinuity in post-reform years.

The magnitudes of the spikes are considerably large. As Figure 3a shows, the probability to sell a gain on a given day of the holding period jumps from around 0.002 during the days before a 365-days holding period to approximately 0.007 on the day after the tax discontinuity. No such jump is observed in the post-reform years, again indicating that the tax incentives have a clear effect on trading behavior. Comparing pre- and post-reform sales probabilities further away from the tax discontinuity, the Figure suggests that investors indeed defer the realization of gains until they qualify for preferential tax treatment; the pre-reform selling probabilities tend to be below the post-reform probabilities to the left of the tax discontinuity, and then remain above the post-reform probabilities on days after the 365-days holding period.

For losses, as shown in Figure 3b, the jump is even more considerable than for gains; the selling probability is below 0.002 during the days after the 365-days holding period and stands at 0.008 on day 364. Along with an absent jump in the post-reform period, this is further evidence that the tax discontinuity induces investors to realize their losses as long as they can be used to offset gains. Comparing pre- and post-reform sale probabilities further away from the tax discontinuity, the Figure is suggestive that investors reduce the holding period of losses for tax reasons. The selling probabilities to the right of the tax discontinuity tend to be higher in post-reform years than in pre-reform years. This difference in probabilities could 'feed' the spike to the left of the tax discontinuity. While plausible, these observations (incl. those for gains) rest on a comparison of different time periods and should therefore be viewed with caution.

To complement the graphical evidence, our main regression results are also presented in a table that includes the exact coefficients and standard errors – see Table A.2. In the interest of brevity, the table only shows the coefficients in the 10-day window around the intertemporal tax discontinuity (and not each coefficient in the one-year window that we display in the Figures). Because the coefficients in this table are identical to the coefficients that we discuss above based on the Figures, we do not describe the results displayed in the table in more detail.

Trading in December vs. Rest of the Year. Relating to literature studying differences in trading behavior in December vs. other months of the year (see Introduction), we study if the effects that we identified before are driven by turn-of-the-year tax planning in December.

We showed in section 3 that sales of both gains and losses are evenly distributed around the year. This is a first piece of evidence that trading in December is not fundamentally different than in other months of the year. Relying on the regression approach and our rich data, we now explore selling behavior around the intertemporal tax discontinuity separately for sales realized in December and sales realized throughout the other months of the year (i.e., January-November). This procedure leads to four different Figures (displayed in the Appendix) which are to be interpreted just as the regression Figures that we saw before: i) Figure A.2a: Gains realized in January-November ii) Figure A.2b: Gains realized in December, iii) Figure A.3a: Losses realized in January-November iv) Figure A.3b: Losses realized in December.

The important take-away result of this exercise is that our main effects above are not

driven by turn-of-year trading in December. As the Figures show for both gains and losses, the spikes in selling probabilities around the tax discontinuity are very pronounced all around the year. That is, we can clearly see that trading behavior in the pre-reform years is heavily affected by the tax discontinuity both in the months January-November as well as in December. The Figures also show that there never are any irregularities around the 365-days cutoff in post-reform years (in which the tax discontinuity is abolished), neither in December nor the rest of the year.

Comparing December to the rest of the year, we further observe that December selling probabilities are on a slightly different level than selling probabilities across the rest of the year, especially for losses in pre-reform years. In addition, December trading both in pre-reform and post-reform periods is somewhat noisier across the entire holding-period window than during the rest of the year. This observation could be explained by less number of observations in the December Figures, but it may also point in the direction that non-tax factors affect trading behavior in December.

Overall, these observations imply that an investigation of trading patterns in December cannot separate tax-effects from other non-tax factors (such as window dressing, an overall tendency of investors to 'clean-up' their portfolios towards the end of the year or the momentum effect).

Non-parametric Regressions by Year. To shed light on the yearly dynamics and to examine if a few exceptional years drive our main results above, we estimate the daily selling probabilities separately for each year in our data sample. The resulting Figures, which are to be interpreted as our main regression-based Figures above, are presented in Appendix Figures A.8 to A.17. Each of these Figures presents the regression results for three consecutive years.³³ We again estimate the selling probabilities separately for gains and losses. To make all yearly Figures comparable, the scale of the y-axis is held constant across all Figures.

The results for gains in pre-reform years (i.e., where 365 days holding period was tax relevant) are presented in Figures A.8 to A.10. Overall, we see a spike in selling probabilities to the right of the tax-relevant discontinuity in each pre-reform year of our sample period.³⁴ The results for gains in post-reform years (i.e., where the 365-days tax discontinuity is not tax relevant anymore) are presented in Figures A.11 and A.12. We do not see any spikes or irregularities in selling probabilities around the holding period of 365 days in any of the six post reform years. The results for losses in pre-reform years are shown in Figures A.13 to A.15. We observe clear and substantial spikes in selling probabilities just before the 365-days holding period in each pre-reform year (1999-2007). The results for losses in post-reform years are shown in Figures A.16 and A.17. As with gains, we do not observe spikes or irregularities

³³Note that we do not present results for the year 2008; an analysis of the year 2008 would not be appropriate because the holding period for shares bought in 2008 is still below 365 days when the reform takes place (and the assets attain grandfathered status).

³⁴The spikes are somewhat smaller, though still clearly visible, during the years 2000-2002. The smaller magnitude of the spike during this time period is reasonable given that gains were less prevalent during the burst of the *dot.com* bubble and many investors presumably had losses that they could use to offset gains and which made it less necessary to sell gains in the tax-free domain.

around the 365-days holding period in any of the six post-reform years. Importantly (e.g., for our bunching analyses), the uniform picture across all post-reform years suggests that the financial crisis did not affect the sales distribution significantly; this implies that the pre-reform years are an appropriate counterfactual.

Overall, selling probabilities of both gains and losses spike around the 365-days holding period in all pre-reform years, but we do not see spikes in any of the post-reform years. We interpret this finding as clear evidence that the tax discontinuity affects trading behavior. Remarkably, the tax induced spikes for both gains and losses vanish directly in the year after the intertemporal tax discontinuity has been abolished. This immediate response suggests that the reform was very salient and implies that adjustment frictions and learning (as seen in other settings) play a minor role.

Heterogeneity w.r.t. investor characteristics. We study heterogeneity with respect to three different investor characteristics: age, investor experience (both measured in years) and gender (dummy indicating males). The underlying regression models condition on a set of other investor-level characteristics (see above). In particular, the controls for income and wealth variables serve as effective control variables that capture potential heterogeneity in tax rates.

Figures A.4a and A.4b depict the effect of an additional age year on selling probabilities on each day of the holding period. We particularly see age heterogeneity in the context of loss-selling behavior (see Figure A.4b). The likelihood of selling a loss shortly before the tax discontinuity sharply increases in age in pre-reform years. That is, older workers are more likely to sell gains for tax reasons. We see no such effect in the post-reform years in which the tax discontinuity is abolished. Age heterogeneity is not very pronounced in the context of gains and we cannot conclude from the data that older and younger investors respond differently to the tax discontinuity when it comes to selling gains. In addition, we see no difference in selling probabilities between older and younger investors for holding-period-days further away from the tax discontinuity (this goes for both losses and gains). Importantly, all our age effects are conditional on our measure of experience; that is, they are not confounded by trading experience.

Figures A.5a (for gains) and A.5b (for losses) illustrate the coefficients for investor experience. The result is unambiguous for both losses and gains: experienced investors react stronger to the tax. This is reflected in the finding that selling probabilities around the tax discontinuity sharply increase with each year of experience in pre-reform periods. In other words, the probability to sell a stock for tax purposes around the tax discontinuity increases in trading experience. Further distant to the tax discontinuity, we do not see any significant effects of experience on the probability of selling gains, neither in pre nor in post reform years. This is different for losses: experienced traders are more likely to sell losses throughout the entire set of holding-period days before the tax discontinuity in pre-reform years. This difference disappears for days to the right of the tax discontinuity. Note, again, that these effects of experience are conditional on age of the investor.

Heterogeneity with respect to gender is plotted in Figures A.6a (for gains) and A.6b (for losses). We do not see any conclusive evidence for gender heterogeneity in the context of gains.

For losses, we see a large negative spike just before the tax discontinuity in pre-reform years. This finding indicates that men are less likely to sell their losses on the day before the tax discontinuity, implying that men are less tax responsive in the context of loss realizations.

Heterogeneity w.r.t. magnitude of gains and losses. Figures A.7a and A.7b plot the β_1 coefficients of regression equation for each day of the holding period around the 365-days tax discontinuity. These Figures shed light on the question of whether responses to the tax depend on the magnitude of the loss or gain. The pronounced spike in the blue line in Figure A.7a just after the one year threshold implies that investors become much more likely to dispose those stocks which had the largest gains. This effect then levels off over the subsequent weeks. The pattern disappears completely once the flat tax regime is introduced (red line). The relationship is similar but even stronger for the size of losses: The strong decrease of the blue line in Figure A.7b in the three weeks prior to the one year threshold implies that investors become much more likely to dispose of those stocks which have performed the worst.³⁵ Apparently, the last opportunity to at least preserve some additional value in the form of a tax shield gives an extra impetus to dispose of the more extreme loss makers. This feature may be particularly valuable from an optimal investment perspective because investors are in general more hesitant to dispose of the largest losses as implied by the coefficient plots in the positive range in Figure A.7b after the reform (red line) and before the reform (blue line) – except, as discussed for the blue line, for the last few weeks before the one year threshold.

5 Taxes and the Disposition Effect

This section describes how we study the interaction between taxes and the disposition effect and subsequently presents the corresponding results.

5.1 Research Strategy

We aim to test if taxes affect the disposition effect. The starting point of the analysis is to measure the existence and magnitude of a potential disposition effect in our data. Following the literature (e.g., Chang et al. 2016) and using our previous non-parametric regression framework, we regress a *Sell*-dummy (see above) on a dummy variable – *Gain* – indicating whether a share package is worth more on this day of the holding period than the initial purchase price. Formally, we estimate the following regression for each day of the holding period and using shares with both gains and losses:

$$Sell_{ijd} = \beta_0 + \beta_1 Gain_{ijd} + \varepsilon_{ijd}. \quad (5)$$

If β_1 is greater than zero in this regression, this is evidence of a disposition effect; i.e., gains are sold with larger probability than losses. The coefficient for β_1 measures the magnitude

³⁵Losses are measured as negative values. Hence, a negative coefficient corresponds with an increased likelihood to dispose of larger losses.

of the disposition effect. We plot these β_1 coefficients in our result graphs.

We estimate the above regression separately for pre-reform and post-reform years. Any difference between pre and post reform years, especially around the 365-days tax discontinuity, sheds light on the tax effects of the disposition effect. The difference in the disposition effect between post-reform and pre-reform years can also be estimated in a DiD-type regression of the following form:

$$Sell_{ijd} = \beta_0 + \beta_1 Pre + \beta_2 \mathbf{1}(Gain_{ijd}) + \beta_3 Pre \times \mathbf{1}(Gain_{ijd}) + \varepsilon_{ijd}, \quad (6)$$

where *Pre* indicates years before the reform (when the holding period mattered for the tax liability). The interaction of pre-years and the gain dummy, β_3 , measures the difference in disposition effect before the reform relative to after the reform. We again estimate this regression separately for each day of the holding period, which allows to check if the difference between post and pre years is particularly pronounced around the 365-days tax discontinuity. We plot β_3 for the one year window around the tax discontinuity when we present the graphical results for this approach.

5.2 Results

Figure 4 plots the disposition effect on each day of the holding period separately for pre- and post-reform years. That is, we plot the β_1 coefficients of regression equation 5. In the absence of a tax discontinuity in post-reform years, we observe the disposition effect on each day of the holding period. That is, the probability to sell gains is higher on each day of the holding period than the probability to sell losers.

How does the magnitude of our disposition effect compare to estimates in the literature? According to the overview handbook chapter by Barber and Odean (2013), the selling probability of gains is about 20-70% higher than that of losses. To make our estimates comparable with these numbers, we divide the sales probability of gains by the sale probability of losses. Technically, this means we use coefficients from regression equation 5 and divide the coefficient of the gain dummy by the constant (which indicates the probability to sell a loss for each day of the holding period). The results of this exercise are plotted in Figure 5 (that is, Figure 5 plots the ratio β_1/β_0). For the purpose of comparing our disposition tax effect to the estimates in the literature, we mostly consider the post reform period (without tax discontinuity) because, as we see below, the disposition effect in the pre reform period is heavily affected by the intertemporal tax discontinuity. On average over the entire holding period of days 185-545 in the post period, we observe that the probability to sell a gain is 67% higher than the probability to sell a loss. This finding is well in line with the findings in the literature.

Looking at pre-reform years with the tax relevant discontinuity in Figure 4, it is clearly visible that the disposition effect is affected by the capital-gains taxes. To the left of the 365-days tax discontinuity the disposition effect is first reduced and then steadily drops. The disposition effect then turns negative during the days before the tax discontinuity and exhibits a sharp negative spike on the last day before the 365-days holding period is reached. This reveals that

the desire to sell losers before the tax discontinuity for tax reasons dominates the disposition effect. The pattern is reversed for the days just after the 365-days tax discontinuity. The disposition effect is strongly amplified as compared to its usual magnitude; we see a substantial spike in selling probabilities of gains during the days after the tax discontinuity. On subsequent days, the disposition effect remains higher than usually and it takes about 35 holding-period days to go back to the usual level. The findings are consistent with investors selling gains once they are tax free.

Figure 4 provides clear evidence that the disposition effect is affected by the tax around the days of the tax discontinuity. Does the tax discontinuity also impact the magnitude of the disposition effect on holding-period days more distant to the tax discontinuity? To shed light on this question, we require a benchmark against which the disposition effect away from the tax discontinuity can be compared. We use the post-reform periods (without tax discontinuity) as the benchmark. This exercise obviously relies on the assumption that the post-reform disposition effect is a good counterfactual for the pre-reform years. The Figure indicates that, away from the tax discontinuity, the disposition effect tends to be lower during the first year of the holding period and higher after 365 days holding period have passed. This suggests that the tax discontinuity affects the disposition effect even on holding-period days distant to the tax discontinuity.

All above results are also visible in Figure A.18 which plots the coefficients of the DiD set-up (β_3 in equation 6). These coefficients compare the disposition effect between pre-reform and post-reform years. The Figure particularly confirms that the days around the tax discontinuity are substantially different between post and pre years, and additionally adds to the suggestive evidence that the disposition effect is affected by the tax discontinuity even on holding-period days away from the tax discontinuity.

6 Further Discussion and Conclusion

In this paper, we contribute to a better understanding of the role of capital-gains taxes for the stock-market trading behavior of private investors. We provide causal evidence on two interrelated questions: i) How do capital-gains taxes affect the holding period of private stock market investments? ii) How do taxes affect the disposition effect? The existing evidence with regard to these questions is surprisingly limited. This lack of evidence is presumably attributable to the challenge of finding appropriate micro level data on trading behavior in combination with an institutional set up that allows for credible identification of causal tax effects. Our paper overcomes this challenge in that it combines high-frequency portfolio-level data (which we confidentially obtained from a large German bank) with an intertemporal tax discontinuity, and its abolishment, in the German capital-gains tax system. This setting informs policy makers in other countries (such as the US) about the expected effects of similar reforms.

Our findings provide clear and direct evidence that capital-gains taxes affect the trading behavior of individual investors. Selling probabilities, which we estimate on a daily basis, are strongly affected by the tax discontinuity and disappear in years after the abolishment of the

tax discontinuity. Interesting patterns of heterogeneity reveal that more experienced and older investors respond stronger to tax incentives.

We also find that the disposition effect – the tendency to sell gains with a larger propensity than losses – is strongly affected by capital-gains taxes. Depending on the type of sale – gain or loss – the disposition effect is accelerated or mitigated due to the tax. Previous studies have found that more experienced and older investors exhibit smaller disposition effects. However, as can be concluded from our heterogeneity analyses for gains and losses, this is not an intrinsic direct effect of age or experience. We find that it is salient intertemporal tax discontinuities which induce the more experienced investors to dispose of their loss-making positions. When the salient tax discontinuities are removed, there is no difference in the probability to dispose of losses anymore between more or less experienced investors or older and younger investors. This implies that, in the absence of the tax discontinuity, the disposition effect is not different between older and younger and between more and less experienced investors. Hence, if the US were to smoothen the tax schedule for capital gains, the seemingly stronger resistance of more experienced (or older) investors to behavioral biases may disappear as well because it is the tax discontinuity in the tax schedule which helps these types of investors to focus their minds / make up their minds on loss-making positions.³⁶

How do our results relate to the predictions from theoretical models such as Constantinides (1984)? First, our results are consistent with theory in that we see that the tax discontinuity induces investors to delay the sale of gains until they qualify for preferential tax treatment and to realize losses earlier, both relative to a counterfactual without intertemporal tax discontinuity. Second, the sharp spike in selling probabilities of losses shortly before the tax discontinuity is not necessarily consistent with standard theoretical predictions. However, this result is consistent with the notion that the tax discontinuity serves as a self-control device that commits loss averse investors to take care of their losses. The idea of a self-control mechanism to realize losses was first developed by Shefrin and Statman (1985, section I.D.). According to this idea, investors are reluctant to realize losses, and only realize their losses when there is an external self-control mechanism (commitment device) that induces them to sell losses. The tax discontinuity, which is salient and known to investors, potentially serves as such an external self-control mechanism (commitment device) because the accumulated losses lose their valuable tax-shield function once the threshold for the tax discontinuity is passed. As a result, losses are not realized immediately as they accrue (because investors do not like to realize losses) and instead are realized shortly before the tax discontinuity (because of its role as a commitment device). Our results show that investors do not realize losses as they accrue and instead wait until the quickly approaching tax discontinuity nudges them to realize the loss. To this end, our paper provides some indication that taxes can serve as a commitment device for investors with behavioral biases such as loss aversion. This is then an empirical indication for the rational brought forward by Shefrin and Statman (1985), and more generally adds this evidence to the large literature in

³⁶ A complete smoothening of the tax schedule in the US would imply not only the same tax rate on short and long run capital gains but also a loss carry-back option for the deductibility of capital losses against ordinary income or an abolishment of the deductibility against ordinary income.

'behavioral finance', particularly those papers studying loss aversion among individual investors (e.g., Benartzi and Thaler 1995; Berkelaar et al. 2004; Beshears et al. 2018).

References

- Agersnap, O. and O. Zidar (2021). The tax elasticity of capital gains and revenue-maximizing rates. *American Economic Review: Insights*. (forthcoming).
- Aghion, P., J. Van Reenen, and L. Zingales (2013). Innovation and institutional ownership. *American Economic Review* 103(1), 277–304.
- Badrinath, S. G. and W. G. Lewellen (1991). Evidence on tax-motivated securities trading behavior. *The Journal of Finance* 46(1), 369–382.
- Barber, B. M. and T. Odean (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The journal of Finance* 55(2), 773–806.
- Barber, B. M. and T. Odean (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics* 116(1), 261–292.
- Barber, B. M. and T. Odean (2004). Are individual investors tax savvy? evidence from retail and discount brokerage accounts. *Journal of Public Economics* 88(1-2), 419–442.
- Barber, B. M. and T. Odean (2013). The behavior of individual investors. Volume 2 of *Handbook of the Economics of Finance*, pp. 1533 – 1570. Elsevier.
- Ben-David, I. and D. Hirshleifer (2012). Are investors really reluctant to realize their losses? trading responses to past returns and the disposition effect. *The Review of Financial Studies* 25(8), 2485–2532.
- Benartzi, S. and R. Thaler (1995). Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics* 110(1), 73–92.
- Berkelaar, A. B., R. Kouwenberg, and T. Post (2004). Optimal portfolio choice under loss aversion. *The Review of Economics and Statistics* 86(4), 973–987.
- Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2018). Chapter 3 - behavioral household finance. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics - Foundations and Applications 1*, Volume 1 of *Handbook of Behavioral Economics: Applications and Foundations 1*, pp. 177 – 276. North-Holland.
- Best, M. C. and H. J. Kleven (2018). Housing market responses to transaction taxes: Evidence from notches and stimulus in the u.k. *The Review of Economic Studies* 85(1), 157–193.
- Bhattacharya, U., A. Hackethal, S. Kaesler, B. Loos, and S. Meyer (2012). Is Unbiased Financial Advice to Retail Investors Sufficient? Answers from a Large Field Study. *The Review of Financial Studies* 25(4), 975–1032.
- Brown, K. M. (2013). The link between pensions and retirement timing: Lessons from california teachers. *Journal of Public Economics* 98, 1–14.


- Chang, T. Y., D. H. Solomon, and M. M. Westerfield (2016). Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect. *The Journal of Finance* 71(1), 267–302.
- Chetty, R., J. N. Friedman, T. Olsen, and L. 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.
- Cole, S., A. Paulson, and G. Shastry (2014). Smart money? the effect of education on financial outcomes. *Review of Financial Studies* 27(7), 2022–2051.
- Constantinides, G. M. (1984). Optimal stock trading with personal taxes: Implications for prices and the abnormal january returns. *Journal of Financial Economics* 13(1), 65 – 89.
- Deutsche Bundesbank (2013). Statistik ueber wertpapierinvestments [statistic about stock investments]. Technical report.
- Dhar, R. and N. Zhu (2006). Up close and personal: Investor sophistication and the disposition effect. *Management Science* 52(5), 726–740.
- Dimmock, S. G., W. C. Gerken, Z. Ivkovic, and S. J. Weisbenner (2018). Capital gains lock-in and governance choices. *Journal of Financial Economics* 127(1), 113–135.
- Dorn, D. and G. Huberman (2005). Talk and action: What individual investors say and what they do. *Review of Finance* 9(4), 437–481.
- Dowd, T. and R. McClelland (2019). The bunching of capital gains realizations. *National Tax Journal* 72(2), 323–358.
- Dowd, T., R. McClelland, and A. Muthitacharoen (2015). New evidence on the tax elasticity of capital gains. *National Tax Journal* 68(3), 511–544.
- Fagereng, A., M. B. Holm, B. Moll, and G. Natvik (2019). Saving behavior across the wealth distribution: The importance of capital gains. NBER Working Paper No 26588 .
- Feldstein, M., J. Slemrod, and S. Yitzhaki (1980). The Effects of Taxation on the Selling of Corporate Stock and the Realization of Capital Gains*. *The Quarterly Journal of Economics* 94(4), 777–791.
- Feng, L. and M. S. Seasholes (2005). Do investor sophistication and trading experience eliminate behavioral biases in financial markets? *Review of Finance* 9(3), 305–351.
- Frydman, C., S. M. Hartzmark, and D. H. Solomon (2017). Rolling Mental Accounts. *The Review of Financial Studies* 31(1), 362–397.
- Frydman, C. and B. Wang (2020). The impact of salience on investor behavior: Evidence from a natural experiment. *The Journal of Finance* 75(1), 229–276.
- German Federal Bureau of Statistics (2008a). Einkommens- und verbrauchsstichprobe: Einkommen, einnahmen & ausgaben [income and consumption sample statistic: Earnings, income & expenditures]. Technical report.

- German Federal Bureau of Statistics (2008b). Einkommens- und verbrauchsstichprobe: Vermögen & schulden [income and consumption sample statistic: Assets & liabilities]. Technical report.
- Grinblatt, M. and M. Keloharju (2001). What makes investors trade? *The Journal of Finance* 56(2), 589–616.
- Gupta, K. L. (1971). Aggregation bias in linear economic models. *International Economic Review* 12, 293–305.
- Hanlon, M. and S. Heitzman (2010). A review of tax research. *Journal of Accounting and Economics* 50(2), 127 – 178.
- Hartzmark, S. M. (2014). The Worst, the Best, Ignoring All the Rest: The Rank Effect and Trading Behavior. *The Review of Financial Studies* 28(4), 1024–1059.
- Ivković, Z., J. Poterba, and S. Weisbenner (2005). Tax-motivated trading by individual investors. *American Economic Review* 95(5), 1605–1630.
- Jacob, M. (2018). Tax regimes and capital gains realizations. *European Accounting Review* 27(1), 1–21.
- Kleven, H. J. (2016). Bunching. *Annual Review of Economics* 8, 435–464.
- Kleven, H. J. and M. 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.
- Korniotis, G. M. and A. Kumar (2011). Do older investors make better investment decisions? *The Review of Economics and Statistics* 93(1), 244–265.
- Lavecchia, A. M. and A. Tazhitdinova (2021). Permanent and transitory responses to capital gains taxes: Evidence from a lifetime exemption in canada. NBER Working Paper No 28514 .
- Leuz, C., S. Meyer, M. Muhn, E. Soltes, and A. Hackethal (2017). Who falls prey to the wolf of wall street? investor participation in market manipulation. NBER Working Paper 24083.
- Loos, B., S. Meyer, and M. Pagel (2020). The consumption effects of the disposition to sell winners and hold losers. NBER Working Paper 26668.
- Mortenson, J. A. and A. Whitten (2020, August). Bunching to maximize tax credits: Evidence from kinks in the us tax schedule. *American Economic Journal: Economic Policy* 12(3), 402–32.
- Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of Finance* 53(5), 1775–1798.
- Pesaran, H., R. Pierse, and M. S Kumar (1989). Econometric analysis of aggregation in the context of linear prediction models. *Econometrica* 57(4), 861–88.
- Protopapadakis, A. (1983). Some indirect evidence on effective capital gains tax rates. *Journal of Business*, 127–138.

- Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy* 2(3), 180–212.
- Sasaki, K. (1978). An empirical analysis of linear aggregation problems: The case of investment behavior in Japanese firms. *Journal of Econometrics* 7(3), 313–331.
- Seru, A., T. Shumway, and N. Stoffman (2009). Learning by Trading. *The Review of Financial Studies* 23(2), 705–739.
- Seru, A., T. Shumway, and N. Stoffman (2010). Learning by trading. *The Review of Financial Studies* 23(2), 705–739.
- Shackelford, D. A. and R. E. Verrecchia (2002). Intertemporal tax discontinuities. *Journal of Accounting Research* 40(1), 205–222.
- Shefrin, H. and M. Statman (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance* 40(3), 777–790.
- The White House (2021). Fact Sheet: The American Families Plan. White House Statements and Releases, announced April 28, 2021, online at: <https://www.whitehouse.gov/briefing-room/statements-releases/2021/04/28/fact-sheet-the-american-families-plan/>.
- van Rooij, M., A. Lusardi, and R. Alessie (2011). Financial literacy and stock market participation. *Journal of Financial Economics* 101(2), 449–472.

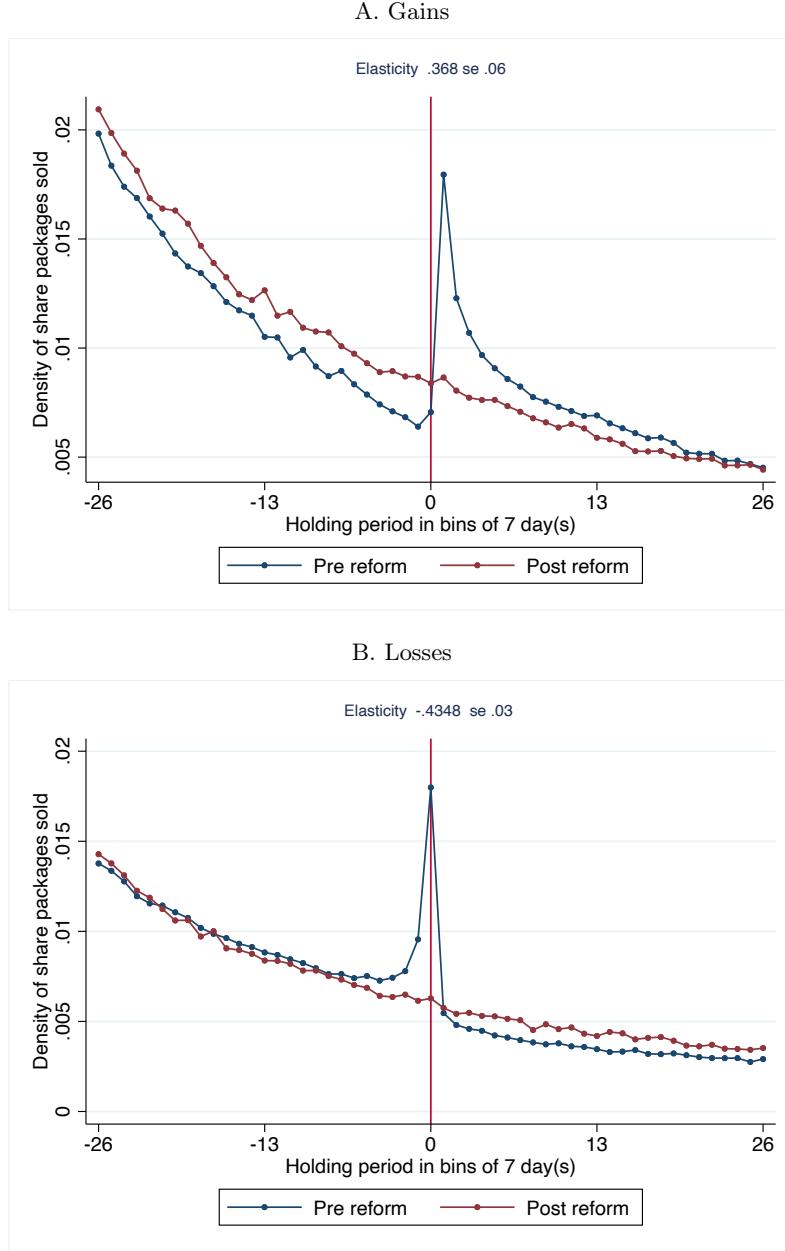
Main Tables and Figures

Figure 1: Raw data: Number of Share Packages sold around Time Discontinuity



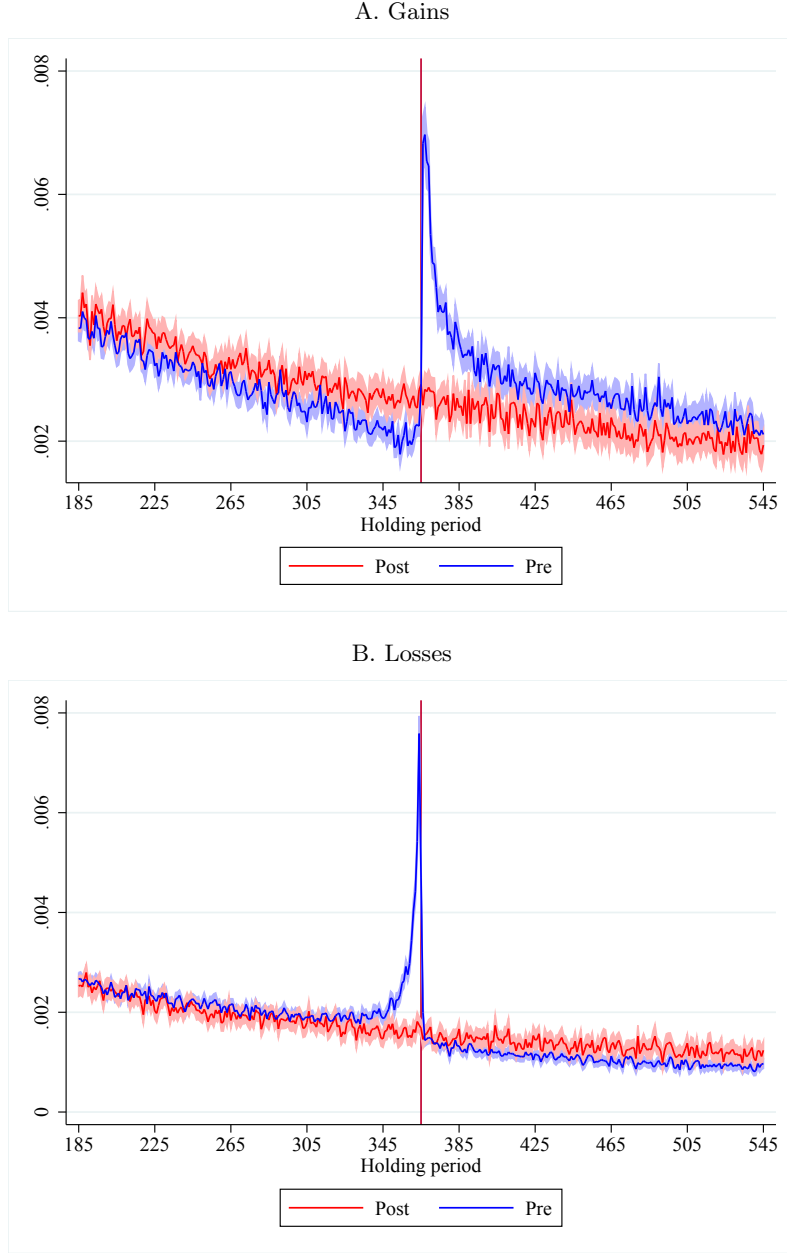
Notes: This figure displays the number of share packages which were sold with a gain (panel A) or with a loss (panel B) in dependency of the holding period. Each dot represents the number of share packages sold in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last week in which gains were taxable. All details are described in section 4.1. The dotted blue line represents sold share packages for which the purchase was made prior to 2009. The dotted red line represents sold share packages for which the purchase was made after 2009. The vertical red line at week 0 marks the last week in which gains were taxable. In panel A, pre reform estimates are based on 44110 investors and 296135 share packages. Post reform estimates are based on 30875 investors and 206263 holding period share packages. In panel B, pre reform estimates are based on 43008 investors and 339970 share packages. Post reform estimates are based on 23757 investors and 126280 share packages.

Figure 2: Difference in Bunching




Notes: This figure displays the share of all purchased share packages with a gain (panel A) or a loss (panel B) in dependency of the holding period. Each dot represents the share of all purchased share packages with a gain (panel A) or a loss (panel B) which were sold in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last bin in which gains were taxable. All details are described in section 4.2. The dotted blue line represents the share of all share packages with a gain (panel A) or with a loss (panel B) purchased prior to 2009 which were sold. The dotted red line represents the share of all share packages with a gain (panel A) or with a loss (panel B) purchased after 2009 which were sold. The vertical red line at week 0 marks the last week in which gains were taxable. The displayed elasticities are calculated based on a notch formula following Kleven and Waseem (2013). In panel A, pre reform estimates are based on 57944 investors and 589254 share packages. Post reform estimates are based on 43584 investors and 405628 share packages. In panel B, pre reform estimates are based on 66396 investors and 941351 share packages. Post reform estimates are based on 43196 investors and 351090 holding period share packages. These numbers include share packages of shares which have not been sold in the 26 weeks after the tax discontinuity. Standard errors are bootstrapped on the investor level.

Figure 3: Non-parametric Regressions




Notes: These figures display non-parametric regressions estimates for each day of the holding period for share packages with gains (panel A) or for share packages with losses (panel B). Coefficients indicate the probability that a share package is sold on this holding-period day. Coefficients and investor-level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$ (panel A) or $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$ (panel B). All estimation details are described in section 4.4. The blue line represents estimates for β_0 for share packages which were bought before 2009. The shaded area displays 95 percent confidence intervals. The red line represents estimates for β_0 for share packages which were bought after 2009. The shaded area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which gains were taxable. In panel A, pre reform estimates are based on 63950 investors and 97 million holding period share package observations. Post reform estimates are based on 51360 investors and 73 million holding period share package observations. In panel B, pre reform estimates are based on 71331 investors and 203 million holding period share package observations. Post reform estimates are based on 52457 investors and 79 million holding period share package observations.

Figure 4: Disposition effect around Time Discontinuity



Notes: This figure displays estimates for the average difference in selling probability between gains and losses on each day of the holding period. Coefficients and investor-level clustered standard errors stem from a series of regressions of the form $Sell_{ija} = \beta_0 + \beta_1 Gain_{ija} + \varepsilon_{ija}$. All estimation details are described in section 5. The blue line represents estimates for β_1 for share packages which were bought before 2009. The shaded area displays 95 percent confidence interval. The red line represents estimates for β_1 for share packages which were bought after 2009. The shaded area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which gains were taxable prior to 2009. Pre reform estimates are based on 72565 investors and 301 million holding period share package observations. Post reform estimates are based on 55847 investors and 152 million holding period share package observations.

Figure 5: Disposition effect: Gain Coefficients relative to Loss Coefficients



Notes: This figure displays estimates for the relative difference in selling probability between gains and losses on each day of the holding period. That is the the coefficient of the gain dummy is divided by the constant. Standard errors are calculated using the delta method. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \beta_1 Gain_{ijd} + \varepsilon_{ijd}$. All estimation details are described in section 5. The blue line represents estimates for β_1/β_0 for share packages which were bought before 2009. The shaded area displays 95 percent confidence interval. The red line represents estimates for β_1/β_0 for share packages which were bought after 2009. The shaded area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which gains were taxable prior to 2009. Pre reform estimates are based on 72565 investors and 301 million holding period share package observations. Post reform estimates are based on 55847 investors and 152 million holding period share package observations.

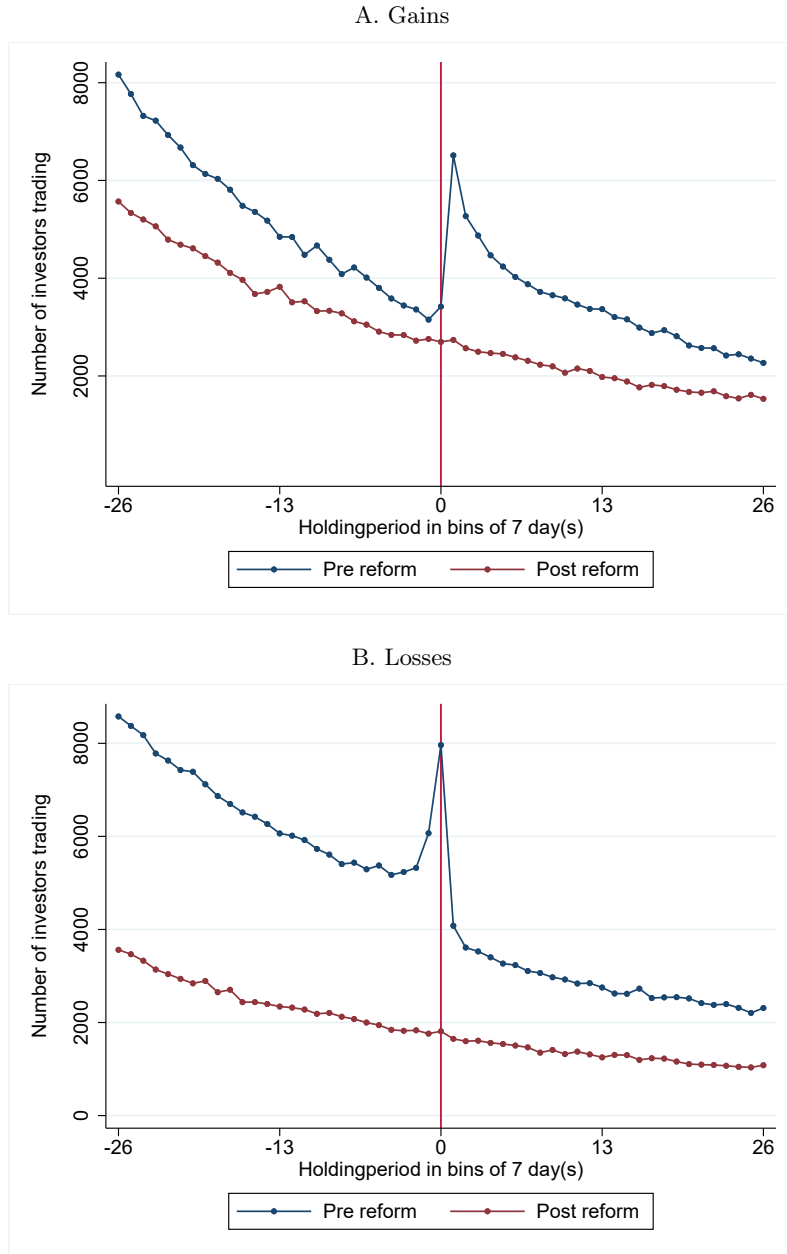
A Additional Figures and Tables

Table A.1: Descriptive Statistics for all Investors in the Sample

	N	Mean	Std.Dev.	Min	Max
<i>Unit of observation: Investor</i>					
Birthyear	93186	1962.74	13.23	1905	2010
Age end of 2015	93186	52.26	13.23	5	110
Trading experience in years end of 2015	93186	13.52	4.28	-0	22
Male	93186	0.86	0.35	0	1
Works in financial sector	93186	0.06	0.24	0	1
Self-employed	93186	0.16	0.36	0	1
Wealth $\leq 30,000$	93186	0.20	0.40	0	1
Wealth $> 30,000 < 100,000$	93186	0.19	0.40	0	1
Wealth $\geq 100,000$	93186	0.07	0.25	0	1
Wealth information missing	93186	0.54	0.50	0	1
Income $\leq 40,000$	93186	0.15	0.36	0	1
Income $> 40,000 < 100,000$	93186	0.30	0.46	0	1
Income $\geq 100,000$	93186	0.04	0.19	0	1
Income information missing	93186	0.51	0.50	0	1
Holding a PhD	93186	0.06	0.24	0	1
Number of trades	93186	77.79	218.29	0	19877
Number of trades 0.5-1.5 years	93186	11.27	24.87	0	876
Average monthly turnover	93109	10.86	15.39	0.00	99.66
Average monthly turnover < 2009	82618	11.80	16.13	0.00	99.41
Average monthly turnover ≥ 2009	87319	9.05	16.12	0.00	100.00
Average portfolio value	93109	51726	239157	0.03	57774533
Average percentage gain per trade	81688	32.63	27.61	0.00	263.64
Average percentage loss per trade	78926	-31.49	18.99	-96.83	-0.01
Average gain (EUR) per trade	86486	9.23	658.07	-5429.97	5345.57
<i>Unit of observation: Share package</i>					
Sale in December	7248978	0.08	0.27	0	1
Sale in December: Gain	3925440	0.07	0.26	0	1
Sale in December: Loss	3323538	0.08	0.27	0	1

Notes: The table depicts the summary statistics for all variables used in our analysis. Variables are defined as follows: Birthyear is the birth year of the investor; Age and trading experience end of 2015 are the age and the trading experience measured by the number of years the investor has a depot at that bank at 12/31/2015; Male, works in the financial sector, holding a PhD and self-employed are dummy variables information comes from the MiFID documentation; Wealth $\leq 30,000$, Wealth $> 30,000 < 100,000$; Wealth $\geq 100,000$ and Wealth missing are 4 mutually exclusive wealth dummies indicating whether the investor belongs to one of the respective wealth groups. Income $\leq 40,000$, Income $> 40,000 < 100,000$, Income $\geq 100,000$ and Income information missing are 4 mutually exclusive income dummies indicating whether the investor belongs to one of the respective income groups. The information for wealth and income stems from the MiFID documentation and is self-reported. Number of trades is the investor average of the total number of share packages (see section 3.2 for a definition) sold; Number of trades 0.5-1.5 years is the investor average of the total number of share packages sold with holding periods in between 185 and 545 days. Average monthly turnover is the investor average of the average monthly portfolio turnover (measured in percent). Monthly portfolio turnover is calculated as in Barber and Odean (2001) as one-half of the monthly sales turnover plus one-half the monthly purchase turnover. Sales (purchase) turnover is defined as value of shares sold (purchased) divided by the portfolio value in the beginning of the month. Average monthly turnover < 2009 and average monthly turnover ≥ 2009 show the average monthly turnover for months prior and after January 2009 respectively. Average portfolio value is the investor average of the average monthly portfolio value as of end of the month. Average percentage gain, average percentage loss and average gain per trade are the investor average of the average gain (loss) of share packages sold by the investor. Sale in December, Sale in December: Gain and Sale in December: Loss, show how many of the sold share packages have been sold in December.

Figure A.1: Number of Distinct Investors Trading around Time Discontinuity



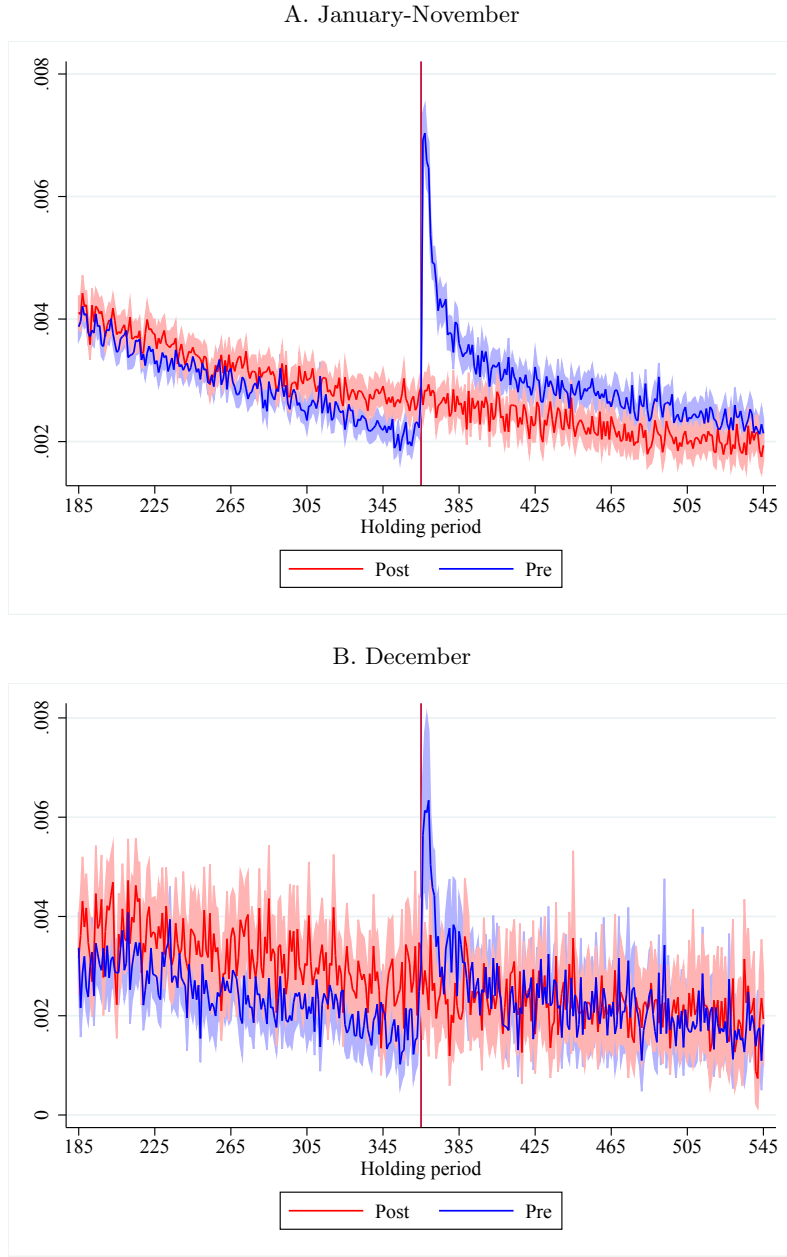
Notes: These figures display the number of investors who sold an appreciated share package (panel A) or a depreciated share package (panel B) with the respective holding period. Each dot represents the number of investors who sold a share package in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last week in which gains were taxable. All details are described in section 4.1. The dotted blue line represents the number of investors who sold share packages for which the purchase was made prior to 2009. The dotted red line represents the number of investors who sold share packages for which the purchase was made after 2009. The vertical red line at x-axis value zero marks the last week in which losses could be used to offset gains. In panel A pre reform estimates are based on 43008 investors and 339970 share packages. Post reform estimates are based on 23757 investors and 126280 share packages. In panel B pre reform estimates are based on 43008 investors and 339970 share packages. Post reform estimates are based on 23757 investors and 126280 share packages.

Table A.2: Non-parametric Regressions: Coefficients for 10 day Window around Discontinuity.

	361	362	363	364	365	366	367	368	369	370
<i>Gains, Pre-Reform</i>										
Coefficient	.0021711	.0023936	.0024182	.002389	.0030595	.0072665	.0074098	.0069455	.0068833	.0057054
S.E.	.0001034	.0001135	.0000987	.0000871	.0001176	.0002339	.0002402	.0002403	.0002227	.000168
No share packages	216018	196353	262597	349515	284684	220600	216066	212513	192351	256074
No investors	42943	40470	44416	48705	46263	43071	42802	42648	40065	44055
<i>Gains, Post-Reform</i>										
Coefficient	.0027823	.0026041	.0025716	.0028747	.0029687	.0026218	.0028727	.0028437	.002914	.0028411
S.E.	.0001383	.0001387	.0001174	.0001086	.0001257	.0001352	.0001426	.0001389	.0001454	.0001241
No share packages	164250	155141	206094	265071	211207	160959	163959	162467	152370	202738
No investors	32621	31919	35207	38145	35329	32250	32561	32440	31612	34900
<i>Losses, Pre-Reform</i>										
Coefficient	.0046413	.0051143	.0062907	.0087702	.0054285	.0023066	.0016577	.0016904	.0017309	.0017203
S.E.	.0001334	.0001447	.0001655	.0002088	.0001517	.0000838	.0000671	.0000676	.000073	.0000618
No share packages	413896	380700	500902	656084	523900	400152	402369	402863	370338	489439
No investors	58673	56443	59864	63230	61166	58391	58547	58290	56036	59499
<i>Losses, Post-Reform</i>										
Coefficient	.0016998	.0016377	.0018639	.0017927	.0017302	.0014965	.0015843	.001612	.0016896	.0017526
S.E.	.0001007	.0001027	.0000948	.0000839	.0000881	.0000934	.0000985	.000099	.0001078	.0000925
No share packages	174139	165477	224797	288951	232920	179088	177363	171839	163352	221382
No investors	35679	34629	38689	41880	39491	36490	36333	35483	34468	38477

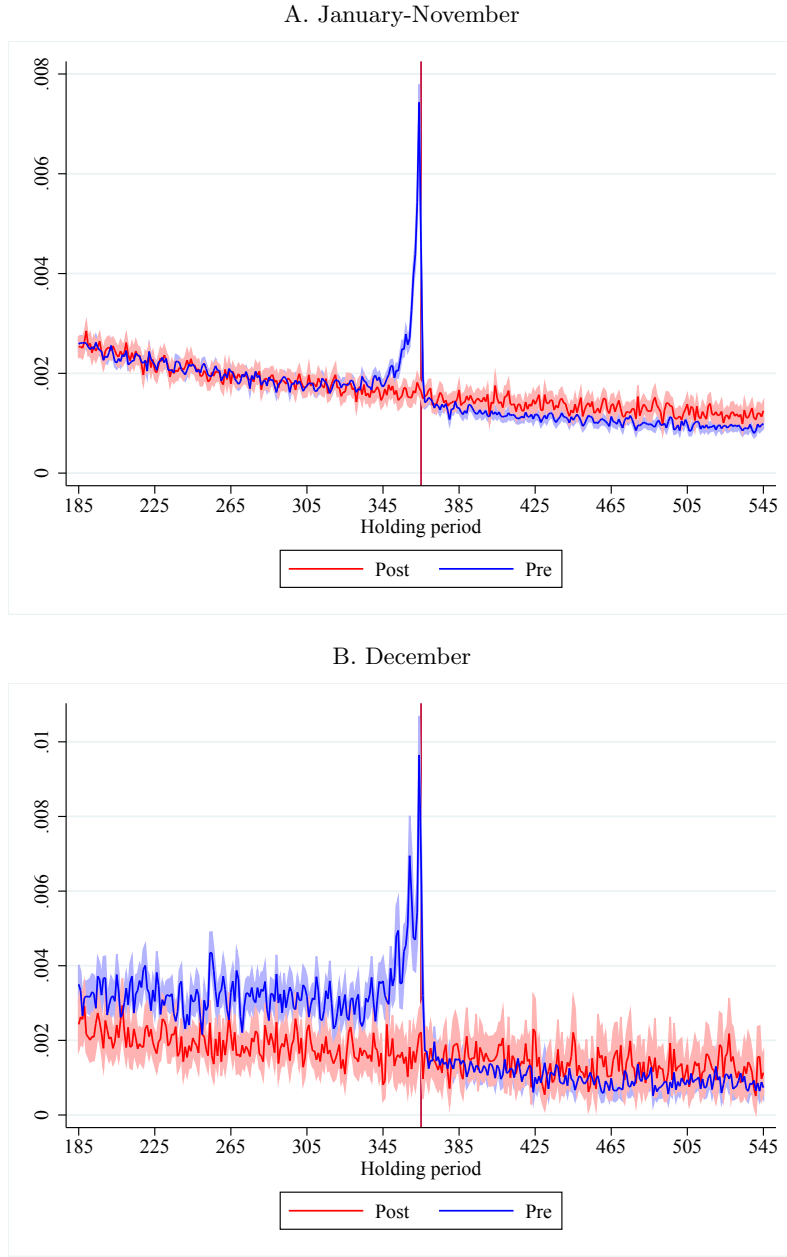
Notes: This table displays the coefficient estimates of the non-parametric regressions. Table shows estimates for the daily selling probabilities in the 10-day window around the intertemporal tax discontinuity. Numbers 361-370 indicate holding-period days. The tax discontinuity is between days 365 and 366. Sample period: 1999-2016. Separate regressions for pre-reform and post-reform periods, as well as gains and losses. The regression equation reads: $Sell_{i,j,d} = \beta_0 + \varepsilon_{i,j,d}$. Standard errors (S.E.) are clustered on the investor level. *No share packages* and *No investors* indicate the number of share packages and number of trading investors that are included in the respective regression, respectively. Note that the number of observations mechanically varies across holding-period days; the number of observations is highest for holding-period days that are multiples of seven, such as 364 (see footnote 29 for more explanation). All estimation details are described in section 4.4. The table mirrors the results in Figures 3a and 3b.

Figure A.2: Non-parametric Regressions: Gains, December vs. January-November



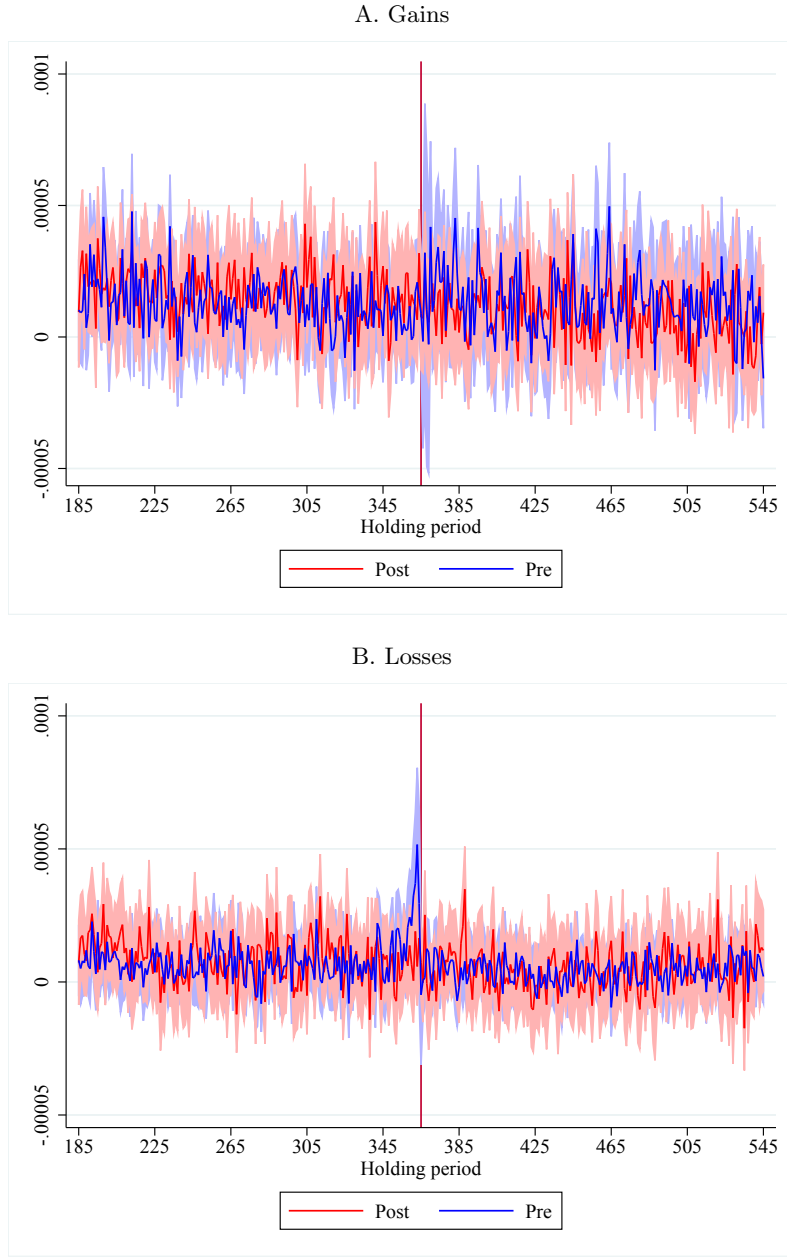
Notes: These figures display the non-parametric regression estimates for each day of the holding period for share packages with prices above the purchase price at the respective day. In panel A we only include share packages on calendar dates not in December. In panel B we only include share packages on calendar dates in december. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$ & $\mathbb{1}(December_{ijd}) = 0$ (panel A) or $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$ & $\mathbb{1}(December_{ijd}) = 1$ (panel B). All estimation details are described in section 4.4. The blue line represents estimates for β_0 for share packages which were bought before 2009. The shaded area displays 95 percent confidence intervals. The red line represents estimates for β_0 for share packages which were bought after 2009. The shaded area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which losses could be used to offset gains. In panel A pre reform estimates are based on 63779 investors and 89 million holding period share package observations. Post reform estimates are based on 51301 investors and 67 million holding period share package observations. In panel B pre reform estimates are based on 54723 investors and 8.2 million holding period share package observations. Post reform estimates are based on 42163 investors and 6.2 million holding period share package observations.

Figure A.3: Non-parametric Regressions: Losses, December vs. January-November



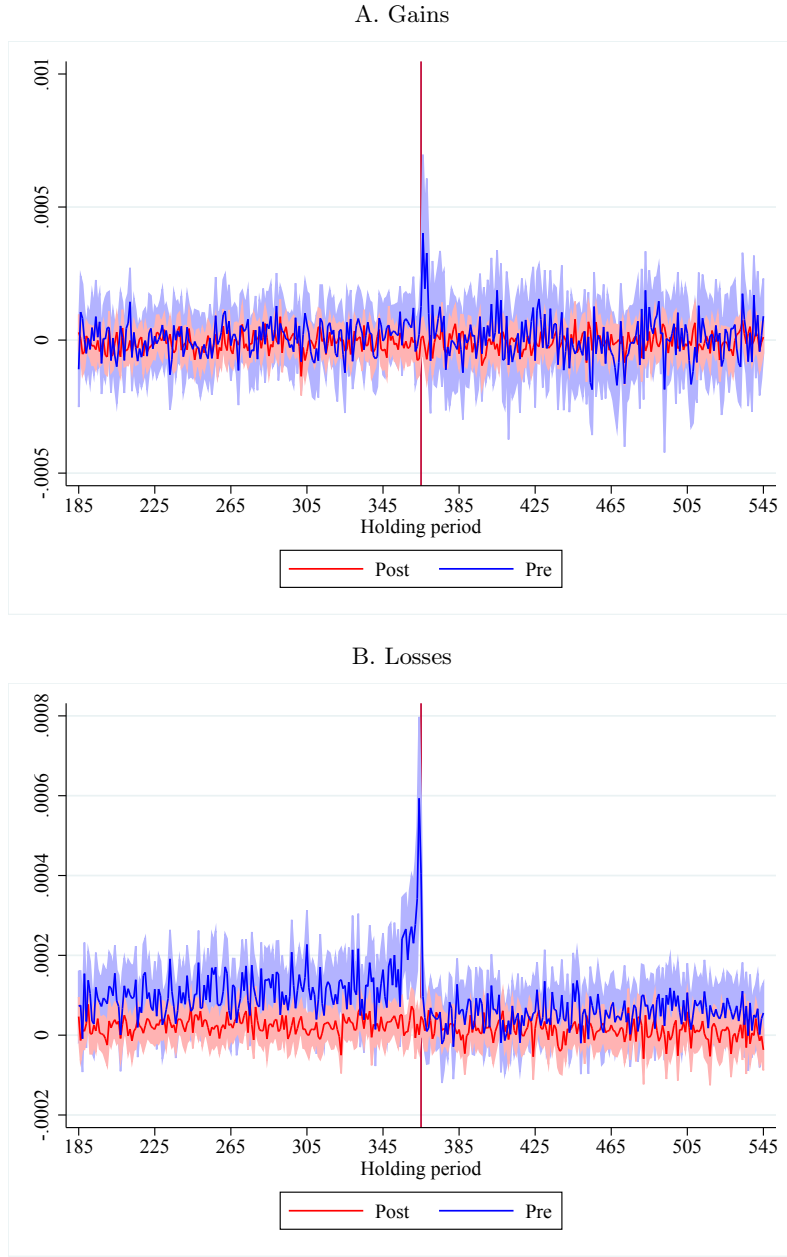
Notes: These figures display the non-parametric regression estimates for each day of the holding period for share packages with prices below the purchase price at the respective day. In panel A we only include share packages on calendar dates not in December. In panel B we only include share packages on calendar dates in december. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$ & $\mathbb{1}(December_{ijd}) = 0$ (panel A) or $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$ & $\mathbb{1}(December_{ijd}) = 1$ (panel B). All estimation details are described in section 4.4. The blue line represents estimates for β_0 for share packages which were bought before 2009. The shaded area displays 95 percent confidence intervals. The red line represents estimates for β_0 for share packages which were bought after 2009. The shaded area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which losses could be used to offset gains. In panel A pre reform estimates are based on 71128 investors and 185 million holding period share package observations. Post reform estimates are based on 52381 investors and 71 million holding period share package observations. In panel B pre reform estimates are based on 68183 investors and 18.5 million holding period share package observations. Post reform estimates are based on 45511 investors and 7.3 million holding period share package observations.

Figure A.4: Heterogeneity w.r.t. Age



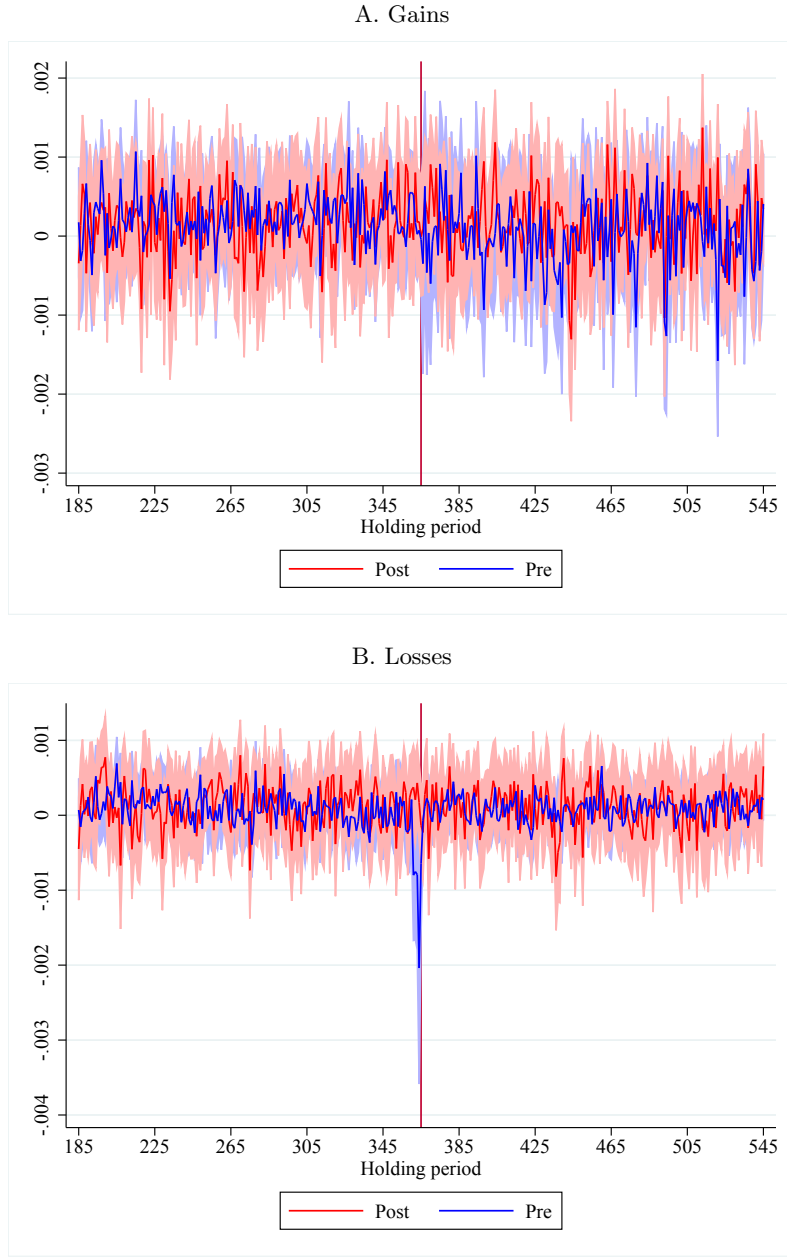
Notes: These figures display coefficient estimates for investor age stemming from non-parametric regressions for each day of the holding period. Panel A includes share packages with prices above the purchase price. Panel B includes share packages with prices below the purchase price. Coefficients indicate by how much an additional year in age shifts the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \beta_1 Age_{id} + Covariates_{ijd}\gamma + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$ (panel A) or $Sell_{ijd} = \beta_0 + \beta_1 Age_{id} + Covariates_{ijd}\gamma + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$ (panel B). Where Age is the age of the investor on a respective calendar-day date. *Covariates* include controls for experience, gender, income category, wealth category, working in the financial sector, having a doctoral degree, and being self-employed. All estimation details are described in section 4.4. The blue line represents estimates for β_1 for share packages which were bought before 2009. The shaded blue area displays 95 percent confidence intervals. The red line represents estimates for β_1 for share packages which were bought after 2009. The shaded red area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which gains were taxable. Panel A Pre reform estimates are based on 63743 investors and 91 million holding period share package observations. Post reform estimates are based on 51244 investors and 72 million holding period share package observations. Panel B pre reform estimates are based on 70783 investors and 176 million holding period share package observations. Post reform estimates are based on 52290 investors and 76 million holding period share package observations.

Figure A.5: Heterogeneity w.r.t. Experience



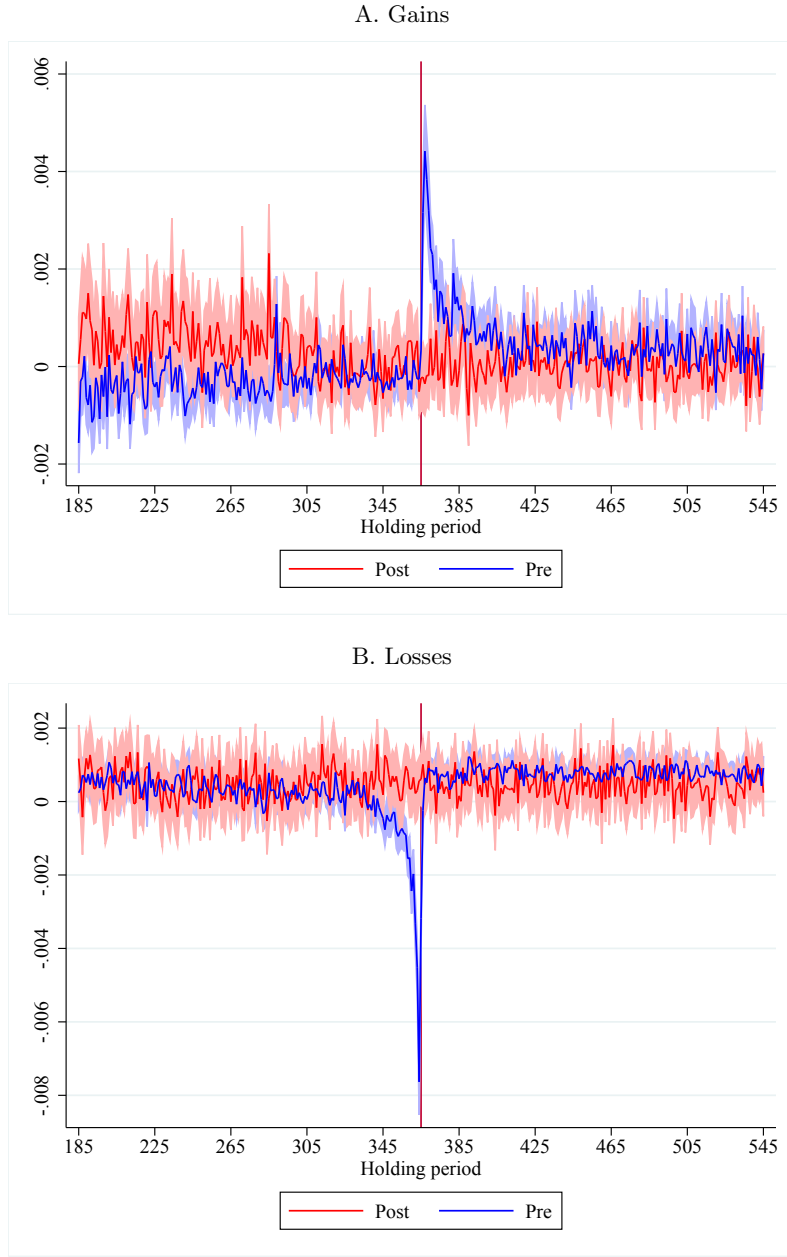
Notes: These figures display coefficient estimates for investor experience stemming from non-parametric regressions for each day of the holding period. Panel A includes share packages with prices above the purchase price. Panel B includes share packages with prices below the purchase price. Coefficients indicate by how much an additional year in experience shifts the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \beta_1 Exp_{id} + Covariates_{ijd}\gamma + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$ (panel A) or $Sell_{ijd} = \beta_0 + \beta_1 Exp_{id} + Covariates_{ijd}\gamma + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$ (panel B). Where Exp is measured by the number of years the investor has a depot at that bank. *Covariates* include controls for age, birthyear (i.e. cohort), gender, income category, wealth category, working in the financial sector, having a doctoral degree, and being self-employed. All estimation details are described in section 4.4. The blue line represents estimates for β_1 for share packages which were bought before 2009. The shaded blue area displays 95 percent confidence intervals. The red line represents estimates for β_1 for share packages which were bought after 2009. The shaded red area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which gains were taxable. Panel A pre reform estimates are based on 63743 investors and 91 million holding period share package observations. Post reform estimates are based on 51244 investors and 72 million holding period share package observations. Panel B pre reform estimates are based on 70783 investors and 176 million holding period share package observations. Post reform estimates are based on 52290 investors and 76 million holding period share package observations.

Figure A.6: Heterogeneity w.r.t. Gender




Notes: These figures display coefficient estimates for a male dummy in the non-parametric regressions for each day of the holding period. Panel A includes share packages with prices above the purchase price. Panel B includes share packages with prices below the purchase price. Coefficients indicate the difference in selling probability of a share-package between man and woman. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \beta_1 Male_i + Covariates_{ijd}\gamma + \varepsilon_{ijd}$ if $1(Gain_{ijd}) = 1$ (Panel A) or $Sell_{ijd} = \beta_0 + \beta_1 Male_i + Covariates_{ijd}\gamma + \varepsilon_{ijd}$ if $1(Loss_{ijd}) = 1$ (Panel B). Where *Male* is a dummy variable indicating whether an investor is male or not. *Covariates* include controls for age, birthyear (i.e. cohort), experience, income category, wealth category, working in the financial sector, having a doctoral degree, and being self-employed. All estimation details are described in section 4.4. The blue line represents estimates for β_1 for share packages which were bought before 2009. The shaded blue area displays 95 percent confidence intervals. The red line represents estimates for β_1 for share packages which were bought after 2009. The shaded red area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which gains were taxable. In panel A pre reform estimates are based on 63743 investors and 91 million holding period share package observations. Post reform estimates are based on 51244 investors and 72 million holding period share package observations. In panel B pre reform estimates are based on 70783 investors and 176 million holding period share package observations. Post reform estimates are based on 52290 investors and 76 million holding period share package observations.

Figure A.7: Heterogeneity w.r.t. Price-change Magnitude




Notes: These figures display coefficient estimates for the magnitude of a price-change from non-parametric regressions for each day of the holding period. Panel A includes share packages with prices above the purchase price. Panel B includes share packages with prices below the purchase price. Coefficients indicate by how much an additional percentage point change in the price changes the probability that a share-package is sold on this holding-period day. Note since the change for losses is negative, negative values mean that share packages with higher losses are sold with a higher probability. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \beta_1 Change_{ijd} + \varepsilon_{ijd}$ if $\mathbf{1}(Gain_{ijd}) = 1$ (Panel A) or $Sell_{ijd} = \beta_0 + \beta_1 Change_{ijd} + \varepsilon_{ijd}$ if $\mathbf{1}(Loss_{ijd}) = 1$ (Panel B). Where $Change_{ijd}$ is measured as $\frac{P_{ijt}d - P_{ij0d}}{P_{ij0d}}$. All estimation details are described in section 4.4. The blue line represents estimates for β_1 for share packages which were bought before 2009. The shaded blue area displays 95 percent confidence intervals. The red line represents estimates for β_1 for share packages which were bought after 2009. The shaded red area displays 95 percent confidence intervals. The vertical red line at day 365 marks the last day in which gains were taxable. In Panel A pre reform estimates are based on 63887 investors and 94 million holding period share package observations. Post reform estimates are based on 51309 investors and 72 million holding period share package observations. In Panel B pre reform estimates are based on 71283 investors and 199 million holding period share package observations. Post reform estimates are based on 52438 investors and 78 million holding period share package observations.

Figure A.8: Non-parametric Regressions by Year: Gains, Pre Years 1999-2001




Notes: This figure displays non-parametric regression estimates for each day of the holding period separately for the years 1999-2001. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which gains were taxable. Estimates for 1999 are based on 20612 investors and 12.2 million holding period share package observations. Estimates for 2000 are based on 28495 investors and 8.3 million holding period share package observations. Estimates for 2001 are based on 20889 investors and 4.9 million holding period share package observations.

Figure A.9: Non-parametric Regressions by Year: Gains, Pre Years 2002-2004




Notes: This figure displays non-parametric regression estimates for each day of the holding period separately for the years 2002-2004. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which gains were taxable. Estimates for 2002 are based on 19197 investors and 7.9 million holding period share package observations. Estimates for 2003 are based on 21058 investors and 13.3 million holding period share package observations. Estimates for 2004 are based on 24235 investors and 13.3 million holding period share package observations.

Figure A.10: Non-parametric Regressions by Year: Gains, Pre Years 2005-2007




Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2005-2007. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which gains were taxable. Estimates for 2005 are based on 25330 investors and 16.2 million holding period share package observations. Estimates for 2006 are based on 28439 investors and 15.0 million holding period share package observations. Estimates for 2007 are based on 21639 investors and 5.6 million holding period share package observations.

Figure A.11: Non-parametric Regressions by Year: Gains, Post Years 2009-2011




Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2009-2011. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which gains were taxable. Estimates for 2009 are based on 19434 investors and 11.8 million holding period share package observations. Estimates for 2010 are based on 22948 investors and 10.2 million holding period share package observations. Estimates for 2011 are based on 22217 investors and 10.4 million holding period share package observations.

Figure A.12: Non-parametric Regressions by Year: Gains, Post Years 2012-2014




Notes: This figure displays non-parametric regression estimates for each day of the holding period separately for the years 2012-2014. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which gains were taxable. Estimates for 2012 are based on 19542 investors and 10.0 million holding period share package observations. Estimates for 2013 are based on 22151 investors and 13.7 million holding period share package observations. Estimates for 2014 are based on 22391 investors and 12.7 million holding period share package observations.

Figure A.13: Non-parametric Regressions by Year: Losses, Pre Years 1999-2001




Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 1999-2001. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which losses could be used to offset gains. Estimates for 1999 are based on 20057 investors and 10.5 million holding period share package observations. Estimates for 2000 are based on 44730 investors and 62.6 million holding period share package observations. Estimates for 2001 are based on 33830 investors and 34.4 million holding period share package observations.

Figure A.14: Non-parametric Regressions by Year: Losses, Pre Years 2002-2004




Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2002-2004. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbf{1}(Loss_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which losses could be used to offset gains. Estimates for 2002 are based on 26874 investors and 20.6 million holding period share package observations. Estimates for 2003 are based on 17801 investors and 7.7 million holding period share package observations. Estimates for 2004 are based on 23738 investors and 12.1 million holding period share package observations.

Figure A.15: Non-parametric Regressions by Year: Losses, Pre Years 2005-2007




Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2005-2007. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which losses could be used to offset gains. Estimates for 2005 are based on 21693 investors and 8.5 million holding period share package observations. Estimates for 2006 are based on 28834 investors and 15.7 million holding period share package observations. Estimates for 2007 are based on 31098 investors and 30.1 million holding period share package observations.

Figure A.16: Non-parametric Regressions by Year: Losses, Post Years 2009-2011



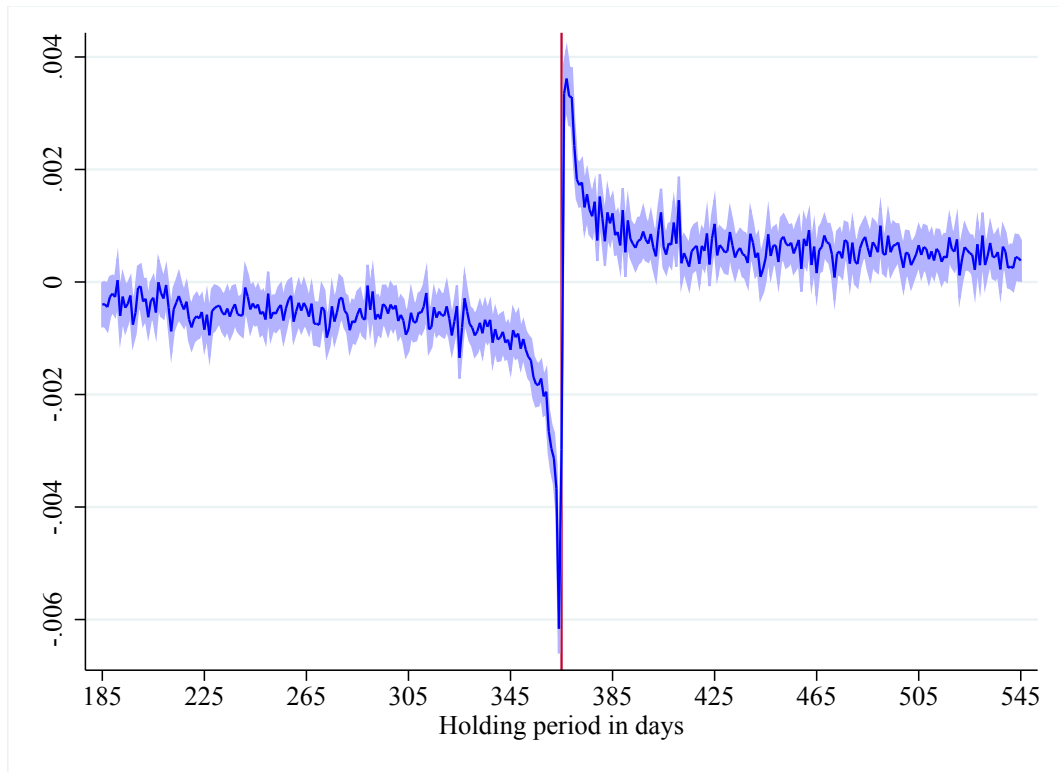
Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2009-2011. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbf{1}(Loss_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which losses could be used to offset gains (prior to 2009). Estimates for 2009 are based on 15224 investors and 5.6 million holding period share package observations. Estimates for 2010 are based on 23815 investors and 13.4 million holding period share package observations. Estimates for 2011 are based on 27218 investors and 21.3 million holding period share package observations.

Figure A.17: Non-parametric Regressions by Year: Losses, Post Years 2012-2014



Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2012-2014. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbf{1}(Loss_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 4.4. The blue line represents estimates for β_0 . The vertical red line at day 365 marks the last day in which losses could be used to offset gains (prior to 2009). Estimates for 2012 are based on 19353 investors and 10.1 million holding period share package observations. Estimates for 2013 are based on 19844 investors and 8.1 million holding period share package observations. Estimates for 2014 are based on 21851 investors and 11.0 million holding period share package observations.

Figure A.18: Disposition Effect around Time Discontinuity: DiD Model



Notes: This figure displays difference in difference estimates for the average difference in selling probability between gains and losses on each day of the holding period. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \beta_1 Pre + \beta_2 \mathbb{1}(Gain_{ijd}) + \beta_3 Pre \times \mathbb{1}(Gain_{ijd}) + \varepsilon_{ijd}$. All estimation details are described in section 5. The blue line represents estimates for β_3 . The shaded area displays 95 percent confidence interval. The vertical red line at day 365 marks the last day in which gains were taxable prior to 2009. Estimates are based on 87948 investors and 494 million holding period share package observations.