// NO.20-047 | 10/2020

DISCUSSION PAPER

// FLORIAN BUHLMANN, PHILIPP DÖRRENBERG, JOHANNES VOGET, AND BENJAMIN LOOS

How do Taxes Affect the Trading Behavior of Private Investors? Evidence from Individual Portfolio Data

Leibniz Association

ZEW

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: May 13, 2020

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 difference-in-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: H20, C41, D14, G11

Keywords: Taxation, Capital-gains, Private investors, Trading Behavior, 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**: University of Technology Sydney. Email: Benjamin.Loos@uts.edu.au **Voget**: University of Mannheim and ZEW. Email: voget@uni-mannheim.de.

Buhlmann gratefully acknowledges financial support by MannheimTaxation ScienceCampus. We thank Ulrich Glogowsky, Nathan Goldman, Jeff Hoopes, Eckhard Janeba, Dominika Langenmayr, Rebecca Lester, Christian Leuz, Ernst Maug, Andreas Peichl, Jim Poterba, Stefan Ruenzi, Arthur Seibold, Sebastian Siegloch, Oliver Spalt, Stefanie Stantcheva, Erik Theissen, Christian Traxler, Daniel Waldenstroem, Michelle Hanlon, John Y. Campbell as well as various seminar/conference participants for helpful comments and suggestions. The usual disclaimer applies.

1 Introduction

Many aspects of the trading behavior of individual investors are well documented in the literature.[1] One aspect of individual trading behavior which is not well understood concerns the causal effect of taxes. Capital gains are an important component of private savings (Fagereng et al., 2019) and realized capital gains are subject to investor-level taxes in most countries around the world. It is therefore important to have a proper understanding of how capital-gains taxes affect capital-gains realizations of private investors.

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).[2] 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 than 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 (for example in Barber and Odean (2004) and Ivković et al. (2005); see below and section 2 for more literature), but the evidence is surprisingly limited and our paper aims to move beyond existing studies in understanding the role of investor-level taxes in trading markets. To address the research questions, appropriate micro-level data need to be combined with an institutional set-up that offers plausible exogenous variation in taxes.

However, micro-level data for individual investors and exogenous variation in capital-gains tax rates are scarce (Poterba (2001), for example, discusses the difficulties of identifying tax effects in investment behavior).[3] As a result, the combination of appropriate micro data with a convincing quasi-experimental institutional set-up does not exist in

[1] See Barber and Odean (2013) for an extensive overview of the behavior of individual investors.

[2] The 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.

[3] For example, in the US setting it is 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, which gives an incentive to sell losses towards the end of the calendar year. In addition, there is a step-up of the basis when stocks are bequested. 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.

the literature.[4]

One strand of literature uses tax-return data to study the link between capital gains and taxes (e.g., Feldstein et al. (1980), Jacob (2018), Dowd and McClelland (2019). An early survey is Poterba (2002)). However, tax-return data usually do not include information that are important for a comprehensive understanding of tax effects on trading behavior; for example, they typically only have aggregated information (and thus lack information about single transactions) and they do not include information about unrealized sales. Studies from a different strand of literature use firm and stock level data to shed light on the effect of investor-level taxes (a review is in Hanlon and Heitzman (2010)). These data are not on the investor level and therefore do not allow studying the individual tax responses of investors.

Another set of papers overcomes these data challenges and uses individual-level investor data obtained from brokerage houses to study the link between taxes and trading behavior (e.g., Odean (1998), Grinblatt and Keloharju (2001), Barber and Odean (2004), Ivković et al. (2005)). The identification of tax effects in this literature is often based on the comparison of trading behavior in taxable accounts and non-taxable accounts. However, there likely exist differences in trading behavior between these accounts for non-tax reasons and it is therefore difficult to isolate the tax effect in such a setting. Another approach is to compare trading patterns in December and the rest of the year, and attribute December differences to taxes because of end-of-year tax planning. However, as we show below, such an approach offers no direct evidence of tax effects and might for example be confounded by the momentum effect, window dressing or an overall tendency of investors to 'clean-up' their portfolios towards the end of the year.

In our paper, we add to the existing literature by combining individual-level investor data with a large tax reform that is exploited for causal identification. This set-up allows us to study the direct causal effect of capital-gains taxes on individual trading behavior. We use confidential portfolio-level data provided by a large commercial bank in Germany.[5] These data contain daily information about the entire trading behavior (including purchases and sales of stocks and other assets) in a panel of approximately 100,000 individual investors for the period 1999 to 2016. 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 U.S. 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.

[4]In the following, we provide a brief overview of different literature strands to illustrate the contribution of our study. An extensive review of the literature is presented in section 2.

[5]The 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, 2001)). Our data have for example been used by Leuz et al. (2017).

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 selling investor. 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.[6] The pre-reform system is thus similar to the tax set-up in the US that also differentiates between long-term and short-term gains, though the German pre-reform system has a larger tax differential (with tax free long-term gains). This intertemporal tax discontinuity tax was abolished in the context of the reform. ii) After the reform, all capital gains are subject to a flat tax of 25%. That is, capital-gains taxes became independent of the individual marginal tax rate and independent of the holding duration of the sold asset.

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 (in bins of seven days) by holding duration before and after the reform and 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 the intertemporal tax discontinuity is crossed). 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.[7] 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 the predictions. We see 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

[6] The 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' (Abstract). We follow this terminology and use it throughout the paper.

[7] The 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.

discontinuity. We further see that investors defer sales of gains until they have reached the 365-days holding period; there is a a sharp spike in the number of sold gains in the weeks just after the 365-days tax discontinuity.[8] Our findings are not driven by a few very tax-sensitive investors; the number of distinct investors who sell share packages also spikes around the tax discontinuity.

In the post-reform years where the 365-days tax discontinuity is not tax relevant, 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 suggests 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 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 (e.g., Brown (2013), Kleven (2016)). In contrast to more conventional bunching methods, this approach has the advantage that we do 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.

The first step in this approach is to explore where the excess mass in the number of sales around the tax discontinuity (in pre reform years) comes 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. We then quantify the elasticities: the tax elasticity of the holding period for gains ranges between 0.185 and 0.56 (depending on the applicable tax rate of investors). This translates to a tax-induced increase in the holding period of 16 days for gains. The results for losses are similar: the elasticity estimates range between 0.195 and 0.59 and the change in the holding period is roughly 17 days.

Our next approach to quantify the behavioral response to the tax discontinuity is

[8] It 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 (also see below where we discuss in the bunching setting where the 'excess mass' comes from). 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 7). 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 7.

based on effective capital-gains tax rates for gains. We estimate effective tax rates (following the procedure in Ivković et al. (2005)) for the actually observed sales distribution in pre-reform years as well as a counterfactual scenario that applies the post-reform sales distribution (which does not include any discontinuity induced behavioral responses) to the tax parameters of the pre-reform period. Using this approach, we find that investors manage to reduce their effective tax rate by 11.3% due to behavioral responses.

The next steps of our analysis are based on non-parametric regressions which estimate for each day of the holding period the probability that a given asset is sold on this day of the holding period.[9] 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 tax discontinuity in the post-reform years, neither for losses nor gains, which is further support of a causal tax effect.

While papers such as Odean (1998), Grinblatt and Keloharju (2001) and Ivković et al. (2005) mostly focus on turn-of-the-year (December) trading of losses, 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. In addition, while evidence of turn-of-the-year trading usually focuses on losses, we show that tax motivated behavior affects both the selling of losses and gains. 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. 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 strong evidence that tax responsiveness is increasing in trading experience (conditional on age and other covariates). This finding is based on the observation that spikes in selling probabilities around the tax discontinuity increase in experience. We further see that the tax response increases in age (conditional on experience and other covariates), in particular in the context of losses. Regarding gender,

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

we find that men are less likely to sell their losses during the days before the tax discontinuity. We further explore heterogeneity w.r.t the magnitudes of gains and losses, which has been shown to be potentially relevant in Ivković et al. (2005). We find that the tax responsiveness increases in the size of the gains or losses and that this effect is about double as large for losses relative to gains.

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 above),[10] 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; that is, gains are always sold with a higher probability than losses. This confirms findings in the large literature that documents 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 intertemporal 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.[11]

Relating to papers such as Odean (1998) and Ivković et al. (2005), our findings provide causal evidence that an intertemporal tax discontinuity affects the disposition effect. This finding speaks to the discussion in e.g. Ben-David and Hirshleifer (2012) who 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. We add to this discussion in that we provide clear causal evidence that the disposition effect is affected by taxes. An additional contribution to the literature on the disposition effect arises because we find that, in a system with an intertemporal tax discontinuity (such as the

[10]The 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'.

[11]Previous literature finds for the U.S. that older and more experienced investors are less prone to the disposition effect. The findings from our heterogeneity analysis (see above) indicate 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 the Conclusion (section 7) for more discussion on this.

U.S.), taxes have an effect on trading behavior and the disposition effect throughout the entire year and not only in December. In addition, our results show that the disposition effect is strongly affected through tax motivated selling of losses *and* gains. Overall, our findings then imply that it is not sufficient to adjust for tax effects by allowing for different December effects – which is a very common approach in the literature. More 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.

The paper proceeds as follows. Section 2 provides an overview of the related literature and discusses how our paper contributes to existing findings. Section 3 describes the institutional background of capital-gains taxation in Germany during our sample period. Section 4 provides information on the data and the calculation of holding periods in this data set. We describe the empirical strategy and causal identification in Section 5. The results are presented in Section 6. Section 7 presents some additional discussions and concludes the paper.

2 Contribution to the Literature

We relate to (empirical) studies in different fields and literature strands. We therefore believe that a systematic and extensive overview of the literature studying the tax effects on trading behavior may be valuable to readers. In the following, we describe the approaches and findings in the related empirical literature and elaborate on our contribution relative to the existing studies. We organize the literature review along the different strands of literature that we identified to be relevant for our paper.

Literature using investor-level trading data from brokerage houses or banks. First, we relate to a stream of papers that study the link between taxes and individual trading behavior using portfolio-level micro data. One of the seminal papers in this literature is Ivković et al. (2005) who use data from a discount brokerage. These data allow the authors to track the single investments of individual investors – their US data are very comparable in spirit to the German data that we use. To shed light on taxation effects, Ivković et al. (2005) compare trading behavior in taxable accounts and trading behavior in tax-deferred accounts (IRAs or Keogh plans). The paper finds a negative relation between accrued gains and the selling probability in taxable accounts for stocks with a holding period of more than one year, while it does not observe such a relation in tax deferred accounts. In light of the presumption that taxes should induce investors to defer the realization of gains (see explanation above) and because this negative relation is only observed in taxable accounts, the authors suggest that their finding is an indication of taxation effects (lock-in effect) on trading behavior. The results of the paper also speak to the interaction between taxes and the disposition effect. The relation between accrued gains and selling probabilities in taxable accounts is (as described above) negative once a stock has been held for more than 12 months, and it is positive in the first few months of the holding period. This suggests that the disposition effect outweighs tax effects only in the first few months after the stock was purchased and that tax effects matter more for longer holding periods.

Another paper in this literature is Barber and Odean (2004) who also use investor-level data from brokerages and compare trading in taxable and tax-deferred accounts. They find that the realization probabilities of gains and losses are very similar across these two types of accounts, except in December when loss realization is more pronounced in taxable accounts than in tax-deferred accounts. The authors attribute this 'December' difference between the two types of accounts to tax-loss selling. This somewhat contrasts the results of Ivković et al. (2005) who find that investors are more likely to realize losses in taxable accounts than in tax-deferred accounts throughout the entire year, not just in December. It is therefore an open question if realization probabilities are always different between these two type of accounts or just in December.

Overall, rather than exploiting exogenous variation in taxes, identification of tax effects in Barber and Odean (2004) and Ivković et al. (2005) thus comes from the comparison of taxable and tax-deferred accounts. However, differences in trading behavior between these accounts are not necessarily fully attributable to taxation effects. Trading behavior might be different between these two types of accounts for non-tax related reasons (even conditional on investor fixed effects and exploiting that many investors have both taxable and tax-deferred accounts). For example, investors usually use tax-deferred accounts to save for retirement, and they might therefore be inclined to invest in different types of assets in these accounts than in taxable accounts. In their tax-deferred retirement accounts, investors might seek to invest in less risky assets or purchase assets for these accounts with a much longer investment horizon and hence with the explicit goal of trading these assets less frequently. This assertion that investments in taxable and tax-deferred accounts are different for non-tax reasons is supported by the literature: e.g., theory contributions on the optimal allocation of assets come to the result that certain assets, such as taxable bonds and actively-managed mutual funds, should be held in tax-deferred accounts, whereas other asset types, such as tax-exempt bonds, passively-managed mutual funds and stocks, should be located in taxable accounts (Huang, 2001; Dammon et al., 2001; Shoven and Sialm, 2004). Consistent with the assertion that non-tax considerations make a difference for investment behavior across these two accounts, Barber and Odean (2004) find that turnover is higher in taxable accounts than in tax deferred accounts. This finding of Barber and Odean (2004) induces Ivković et al. (2005, page 1617) to acknowledge that investors may view taxable and tax-deferred accounts differently and that their estimates for tax-motivated trading might therefore be biased.

We move beyond Barber and Odean (2004) and Ivković et al. (2005) in that we use a similarly rich data set of individual investors, but combine it with quasi-experimental variation in tax rates which comes from the intertemporal tax discontinuity and the abolishment of this tax discontinuity. Another difference to 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 holding-period based intertemporal tax discontinuity.

A further set of papers document in individual-level data sets that trading behavior in December is different than trading behavior in other months of the year. For example, Odean (1998) and Grinblatt and Keloharju (2001) find evidence of the disposition effect in all months of the year, except in December. Ivković et al. (2005) also put a focus on turn-of-the-year (December) trading and find that selling probabilities spike in December. These papers interpret these findings as evidence that i) capital-gains taxes matter for investment behavior and that tax-loss selling is prevalent and ii) that the disposition effect is affected by taxes. However, differences in trading behavior in December vs. other months of the year are only indirect evidence of tax effects. These papers do not rely on exogenous variation in tax rates and it is therefore not clear to which extent the 'December' finding is driven by taxes or other seasonality patterns. For example, it is not clear why tax-loss selling should not occur throughout the year.

In addition, as noted by Grinblatt and Keloharju (2004, pages 52-53), the December effect could also be explained with the momentum effect or window dressing. It has been shown that the momentum effect for losses is much larger in December than in other months of the year and also much larger than the December momentum effect for gains (Grinblatt and Moskowitz, 2004). This then implies that it could be rational for investors to sell losses in December even in the absence of tax considerations. Window dressing may also play a role: December often is the time to recap one's portfolio and investors may be embarrassed to carry on losers to the next year. Considering these concerns, Grinblatt and Keloharju (2004) study if investors sell losers and then immediately repurchase the same stocks (so called wash sales) and indeed find evidence for this behavior. This then is a better indication that tax considerations matter, but is yet no direct evidence of tax effects. Also speaking to the context of tax-induced changes of the disposition effect at the end of the year, Ben-David and Hirshleifer (2012) are skeptical if the altered disposition effect toward the end of the year, as for example found by Odean (1998), is driven by tax effects.

Using exogenous variation in rates and the abolishment of a large intertemporal tax discontinuity, we tie to the discussions in the literature in that we provide clear causal and direct evidence for tax effects on trading behavior and on the effect of taxes on the disposition effect. Our approach does not rely on trading patterns in December and shows that taxes matter throughout the entire year; this implies that December irregularities are not necessarily driven by tax effects. We therefore complement the existing literature in that we study a set-up where the concerns about the roots of differential trading behavior in December do not play a role. Our finding on the interaction between taxes and the disposition effect more generally relates to papers that demonstrate that tax considerations of individual investors are sometimes swamped by non-tax considerations or behavioral aspects (Hanlon and Heitzman, 2010).

Literature on the disposition effect. Second, our paper contributes to the general literature on the disposition effect (in non-tax contexts; see above for the relation between taxes and the disposition effect). This literature usually also uses investor-level data. As summarized in the handbook chapter by Barber and Odean (2013), the evidence is very robust that individual investors sell gains with a higher propensity than losses. An important question in this literature is which factors causally affect the disposition effect. This is potentially relevant because an understanding of the causal determinants of the disposition effect can help to improve investment behavior. However, as stated by Frydman and Wang (2020, page 233), the cause of the disposition effect is still debated. A few recent papers provide causal evidence on the determinants of the disposition effect. Frydman and Rangel (2014), Frydman and Wang (2020) and Loos et al. (2020) show that changes in purchase prices or changes in the salience of purchase prices affect the disposition effect. We relate to these papers on the causal drivers of the disposition effect and provide novel evidence that taxes affect the disposition effect and can even reduce it temporarily.

A further contribution of our paper to the literature on the disposition effect is to show that it is not sufficient to have separate December effects to control for tax effects. Consistent with our finding that taxes affect the disposition effect, the literature has acknowledged that tax effects should be controlled for in disposition-effect settings. We provide evidence that taxes affect trading behavior and the disposition effect throughout the entire year, which then implies that December adjustments will not fully control for tax effects.

Literature using tax-return data. Third, we speak to studies that use data from individual tax returns to study the link between capital gains and taxes (an early survey is Poterba, 2002). This literature usually finds a negative relation between realizing capital gains and taxes (e.g., Feldstein, Slemrod, and Yitzhaki, 1980; Bogart and Gentry, 1995;

Daunfeldt, Praski-Staahlgren, and Rudholm, 2010; Jacob, 2013; Dowd, McClelland, and Muthitacharoen, 2015; Jacob, 2018).[12] Our findings on the behavioral effects of capital gains taxes relate to this literature and we confirm that capital gains taxes induce investors to defer the realization of gains. However, as mentioned above, studies using tax-return data typically only have aggregate annual information on the total amount of capital gains and as such lack information on single realized sales; i.e., whether a single realized sale is a gain or loss or how long the respective asset had been held by the investor. Tax return data also do not include information on unrealized sales because these are not tax relevant, and they do not have information about trading activities in non-taxable accounts. Our paper uses portfolio-level data that allow us to overcome most of the data restrictions in this literature. For example, one main finding in our paper relies on the differentiation between gains and losses of single sales, and our empirical approaches account for unrealized assets and exploit the daily frequency of our data set.

Two recent studies use US tax-return data that include information on sales at a less aggregated frequency. Hoopes et al. (2016) have daily data on sales, but their study is not about tax effects. Dowd and McClelland (2019) use American IRS data on capital realizations for directly held assets on the level of the single transaction. Using these data for the tax year 2012, the authors calculate the holding period (in weekly bins) for single assets and study whether the holding period is affected by the intertemporal tax discontinuity in the US. Consistent with our findings, Dowd and McClelland (2019) find that the number of realized gains spikes in the first week in which the lower long-term rate is available. In contrast to our findings, losses spike on both sides of the tax discontinuity. As acknowledged by the authors, this is somewhat surprising since a rational investor should sell losses as short-term in order to offset short-term gains.

Our paper moves beyond Dowd and McClelland (2019) along a number of dimensions. i) In contrast to Dowd and McClelland (2019), we use the institutional setting of the intertemporal tax discontinuity, *and* its abolishment.This setting yields a proper counterfactual and allows us to compare the effects of the tax discontinuity to years where the tax discontinuity did not exist.[13] ii) In addition to plotting the number of sales by holding period around the intertemporal tax discontinuity, we estimate non-parametric regressions with daily frequence which allow to account for sales that are not realized (whereas Dowd and McClelland (2019) use weekly data and do not have information on

[12]Saez (2017) studies the behavioral responses of reported incomes to the 2013 tax reform in the US. Using annual IRS income statistics, the paper finds considerable responses of reported income to the reform, with much of the effect being driven by realized gains.

[13]Dowd and McClelland (2019), who only have one year of data, use the conventional bunching approach to construct a counterfactual. This approach relies on the assumption that the number of realized sales away from the tax discontinuity (which are used to estimate the counterfactual distribution of sales just around the tax discontinuity) are not affected by the tax discontinuity. Theory as well as our results show that this assumption is unlikely to hold.

unrealized sales). iii) We exploit a panel over several years whereas Dowd and McClelland (2019) only have one year of data. iv) We study how capital-gains taxes affect the disposition effect, v) We shed light on tax-motivated trading in December vs. the rest of the year, vi) Looking at the case of Germany, we exploit a set-up where the tax differential between short-term and long-term gains is larger than in the US case.

Literature using firm and stock level data. Fourth, we relate to a strand of literature that uses firm and stock level data to investigate the effects of investor-level capital-gains taxes (see the overview in Hanlon and Heitzman, 2010). This strand of literature for example studies the effects of investor-level taxes on asset prices, end-of-year market irregularities, acquisition premiums, turnover patterns, and the role of taxes in reactions to news disclosures (e.g., Reese Jr, 1998; Lang and Shackelford, 2000; Seida and Wempe, 2000; Poterba and Weisbenner, 2001; Shackelford and Verrecchia, 2002; Blouin, Raedy, and Shackelford, 2003; Ayers, Lefanowicz, and Robinson, 2003; Ayers, Li, and Robinson, 2008; Dai, Maydew, Shackelford, and Zhang, 2008; Blouin, Hail, and Yetman, 2009; Ball, 2013). These papers typically find that capital-gains taxes matter and that they affect asset prices. These findings are indirect evidence that capital-gains taxes affect individual selling behavior, but they do not allow for the identification of tax effects on individual behavior directly. Studying tax effects on individual investors is generally difficult with firm level data.[14]

Literature on the behavioral responses to taxes. Fifth, we also contribute to the large literature on behavioral responses to taxes using individual-level data. This literature studies the causal effects of taxes along many dimensions, often relying on bunching approaches and/or large tax reforms for identification.[15] We contribute to this literature in that we add micro-level evidence on the behavioral effect of taxes along a dimension that has rarely been investigated before, namely individual-level trading behavior.

[14] A related stream of papers studies the tax sensitivity of institutional investors (Blouin et al., 2017; Sikes, 2018).

[15] This large literature for example studies causal effects of taxes on: taxable income (Chetty et al., 2011; Saez et al., 2012; Kleven and Schultz, 2014), investment behavior (Yagan, 2015), dividend payments (Chetty and Saez, 2005), education (Abramitzky and Lavy, 2014), wealth accumulation (Jakobsen et al., 2018), housing prices (Best and Kleven, 2018), wages (Suarez Serrato and Zidar, 2016; Fuest et al., 2018), consumption (Chetty et al., 2009), migration (Kleven et al., 2014; Agrawal and Foremny, 2019) and labor supply (Eissa and Liebman, 1996; Martinez et al., 2018).

3 Institutional Background

Our analysis is based on the system of capital-gains taxation in Germany between 1999 and 2016 (i.e., the time period of our data set). We focus on the trade of shares and describe in this section how capital gains occurring from shares are taxed in Germany. A major reform of capital-gains taxation was implemented in 2009 and therefore falls into our sample period. Both before and after the reform, capital gains from shares are generally only taxed upon realization (i.e., taxes are due when the share is sold).

Taxation of capital gains before 2009. Before the reform, the tax treatment of capital-gains was dependent on the holding period of the underlying asset. 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 identical to the personal income-tax rate (PIT) of the investor. The PIT depends on the sum of all income types (wage income, self-employment, capital-gains, etc). The top income tax rate (PIT rate) was, for example, 42% in 2008 and applied for overall annual taxable income greater than 52,152 EUR. The entry tax rate in 2008 was 15%. Losses from sales with a holding period of $\leq$ 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 $\leq$ 365 days were subject to the tax.

For illustration, consider a fictitious investor who is subject to the top-income tax rate of 42%. She realizes gains worth 2000 EUR from shares that she had held less than 365 days, and she sells other shares within the 365-days holding period which come with losses 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 resulting from assets with a holding period of more than 365 days were not subject to any taxes; the resulting tax liability on gains was zero if the underlying asset was held for more than 365 days. Accordingly, long-term losses resulting from assets with a holding period of more than 365 days could not offset positive capital gains.

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). Overall, the system creates large incentives to realize gains after the relevant holding period of 365 days, while losses should be realized within the 365-days holding period to reduce the tax base.

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%,PIT rate). Losses can be used to offset gains. The half-income method was abolished.

Consider again an 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. Importantly, any tax incentives to hold assets for a certain time period were abolished. The old pre-2009 tax rules applied to all assets bought before January 2009 (resulting in grandfathered assets).

4 Data

4.1 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 out of all of the bank's clients. 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 our paper (in which we focus on the trading of stocks), 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 4.2 below for a definition and more information). Table 1 provides summary statistics for all investors in our analysis 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 turnover[16] 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 (*Bundesbank*) reports to be the average 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: 1.1 (German Federal Bureau of Statistics, 2008b,a).[17] 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.[18] 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

[16]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.

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

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

(1998) and Barber and Odean (2001). Their data are obtained from an U.S. 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).[19]

We further investigate if trading behavior is different in December compared to other months of the year. It is 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 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 expect if sales were equally distributed across all months of the year.

4.2 Unit of Analysis

We are interested in the number of stock realizations around the intertemporal tax discontinuity.[20] To study the number of stock realizations, we use 'share packages' as the unit of analysis throughout most of the paper. 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

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

[20]That 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.

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.

4.3 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 fictive investor buys five shares of some firm on the second of October and sells all of them on the 15^{th} of October, this would result in one observation with a holding period of 13 days. If the first purchase of a share package occured outside our sample period (i.e., prior to January 2009), 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 fictive 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 (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 such that purchase and sale price are comparable.[22] Stocks for which there is a disparity between recorded trades and month-to-month positions are removed from the sample. For each deposit-ISIN combination we keep trades until we detect any difference between stock and accumulated trading volume. We drop all trades for the affected deposit ISIN combination for points in time after the first difference.

[21] We 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.

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

4.4 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 during pre-reform period, and 212,000 gains and 136,000 for the after-reform period.

5 Empirical Strategy

Our empirical strategy aims at identifying the causal effect of capital-gains taxes on trading behavior, in particular on holding periods and the probability to sell an asset. In addition, we shed light on the interaction of taxes and the disposition effect.

5.1 Raw Data: Number of Trades around the Discontinuity

The starting point for our analyses are figures in which we plot the number of sold share packages (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 weekly bins and plot the number of investors in each bin during the one-year window around the tax discontinuity.

5.2 Difference-in-Bunching

We go on and use 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 the recent overview by Kleven (2016)).[23] We use a difference-in-bunching approach where we use 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 Figures). 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 gain or 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 advanategous 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.

The identifying assumption in our set-up then is that the post-reform distribution (without tax discontinuity) is a plausible counterfactual for the pre-reform distribution (which has the tax discontinuity). Looking at our plotted Figures (see below), this assumption seems plausible; the post-reform distribution appears to be very similar to the pre-reform distribution except for the spikes around the tax discontinuity in pre-reform years.

To make the 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. On top of

[23]Bunching applications for example include: Chetty et al. (2011), Chetty et al. (2013), Bastani and Selin (2014), Best et al. (2015), Best and Kleven (2018), and Almunia and Lopez-Rodriguez (2018).

that, 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 size of the causal tax effect is proportional to the excess mass in bunching relative to the counterfactual distribution. To quantify the tax effect, we estimate parameter b which describes the excess mass around the tax discontinuity relative to the counterfactual distribution. This parameter is then used to calculate an implied elasticity which describes the percentage change in holding period in response to a one-percent change in the tax rate. Based on graphical evidence, we define our bunching window for gains to be the first four weeks of the holding period after the 365-days tax discontinuity. The spike for losses is somewhat more concentrated around the last week before the 365-days tax discontinuity. We therefore choose the bunching window for losses to be the three last weeks before the tax discontinuity. Note that an increased bunching window generally simply increases the excess mass and therefore the tax effect.

We calculate a standard error for the excess mass b using a bootstrap procedure. To do so, we randomly draw share packages from our sample with replacement on the investor level to generate a new set of counts and reestimate the excess mass b^j. Repeating this for a thousand times gives us an estimate for the distribution of b^j. We use the standard deviation of the b^j as our estimate for the standard error of the excess mass.

Following for example Chetty et al. (2011) and Glogowsky (2016), we calculate the elasticity parameter based on the excess mass b using the following equation:

$$e = \frac{\frac{\Delta b}{b^*}}{ln(\frac{1-t}{1-t-\Delta t})}. \tag{1}$$

Recall that, in pre reform years with tax discontinuity, the applicable tax rate 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. At the tax discontinuity, the tax rate thus falls from half of the PIT rate of the individual investor to zero. Since we do not have individual tax rates or taxable income in our data, we calculate two sets of elasticities that differ w.r.t. to the PIT rate that we use: using i) half of the top income tax rate in 2008 (42%), ii) and half of the minimum income tax rate (15%).

Note that there are no strictly dominated regions in our set-up. There are at least four reasons why it might be rational to 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.

5.3 Effective Capital-Gains Tax Rates

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):

$$e^{(1-\delta)gT} = e^{gT} - \tau_{cg} * (e^{gT} - 1). \tag{2}$$

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.[24] 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

[24]There 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.

the institutional background in section 3). 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.

5.4 Non-parametric Regressions

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 (similar to e.g., Hartzmark (2014), Chang et al. (2016) and Frydman et al. (2017)).[25] 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.[26] 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.[27] 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:

$$Sell_{ijd} = \beta_0 + \varepsilon_{ijd}, \tag{3}$$

where indices indicate a share package i of individual investor j on calender-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 calender-day

[25] We 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.

[26] To 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.

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

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.

To summarize, our non-parametric regressions estimate separately for each day of the holding period the probability that a share package is sold on this day of the holding period. Our approach implies that we estimate one regression for each day in the one-year window around the 365-days tax discontinuity. We estimate these 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.[28]

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 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 also 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:

$$Sell_{ijd} = \beta_0 + \beta_1 Demogr_{ijd} + X_{ijd}\beta' + \varepsilon_{ijd}, \tag{4}$$

[28] We 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.

where i, j and d again indicate share packages, investors and calender-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 corresponding summary statistics for all variables used here are reported in Table 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. Studying heterogeneity w.r.t. the magnitude of gains and losses also relates to Ivković et al. (2005) who provide a similar analysis. For this purpose, we include an additional variable into the above regressions which 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:

$$Sell_{ijd} = \beta_0 + \beta_1 Change_{ijd} + \varepsilon_{ijd}, \tag{5}$$

where the *Change* variable describes the change between the share price at holding-day t and the purchasing ($\frac{p_{ijtd} - p_{ij0d}}{p_{ij0d}}$). 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 for the one year window around the tax discontinuity.

5.5 Taxes and the Disposition Effect

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 *Sale*-dummy (see above) on a dummy indicating whether a share package is worth more on this day of the holding period than the purchase price. Formally, we estimate the following regression for each day of the holding period and using all sample years and shares with both gains and losses:

$$Sell_{ijd} = \beta_0 + \beta_1 Gain_{ijd} + \varepsilon_{ijd}. \tag{6}$$

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

$$Sell_{ijd} = \beta_0 + \beta_1 Pre + \beta_2 \mathbb{1}(Gain_{ijd}) + \beta_3 Pre \times \mathbb{1}(Gain_{ijd}) + \varepsilon_{ijd}, \tag{7}$$

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.

6 Results

This section presents the empirical results. All of our empirical findings are presented in graphs which aim to visualize the effects and make them approachable. The chapter is organized along the same order as the description of the empirical strategy in section 5.

6.1 Raw Data: Number of Trades around the Discontinuity

Figure 1 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.

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 2 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.

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.

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 3 and 4 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.

6.2 Difference-in-Bunching

The Difference-in-Bunching results are presented in Figures 5 (for gains) and 6 (for losses). As described in section 5, we use the post-reform years (without tax discontinuity) as the counterfactual distribution for the tax-affected pre-reform years. Recall that we divide the number of sales by the respective total number of share packages (including the ones which have not been sold) in order to account for level differences between pre and post reform years. in order to account for differences in levels across pre-reform and post-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, there is no spike in post-reform years.

Where does the excess mass come from? Are the spikes 'fed' by regions to the left or to the right of the tax discontinuity? 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 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.

We estimate an excess mass of 2.32 (standard error: 0.07) for gains (see Figure 5) and an excess mass of 2.43 (standard error: 0.07) for losses (see Figure 6). To derive an elasticity, these excess-mass estimates can be related to the change in tax rates once the holding-period of 365 days is crossed. As we described in Section 3, the applicable tax rate for assets sold within the first 365 days after purchase is the individual personal marginal income-tax rate. Our portfolio data do not include personal marginal tax rates for investors and we therefore calculate the elasticity based on two scenarios: i) the top marginal income-tax rate of 42% applies, ii) the minimum tax rate (lowest bracket) of 15% applies.[29] Recall that, during almost all of the pre-reform years, capital-gains were effectively taxed at half of the applicable marginal tax rate. Using half of the top marginal income tax rate of 42%, we estimate an elasticity of 0.185 for gains and an elasticity 0.195 for losses. Using half of the minimum income-tax rate of 15%, we estimate elasticities of 0.56 for gains and 0.59 for losses. Our estimates translate to a tax-induced change in the holding period of 16 days for gains and 17 days for losses.

[29]The top marginal income tax rate and minimum tax rate were 42% and 15% during most of the years in our data sample. Note that most of the investors in our data are likely to be high earners and closer to the top rate than to the minimum rate.

6.3 Effective Capital-Gains Tax Rates

Using the strategy described above in section 5.3, 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 5 (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 sales and therefore does not include the tax-induced behavioral responses of losses.

6.4 Non-parametric Regressions

We present the results of our main non-parametric regressions in Figures 7 (for gains) and 8 (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 3 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 7 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 8, 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 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 Introdcution and section 2), we study if the effects that we identified before are driven by turn-of-the-year tax planning in December.

We showed in section 4 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 which are to be interpreted just as the regression Figures that we saw before: i) Figure 9: Gains realized in January-November, ii) Figure 10: Gains realized in December, iii) Figure 11: Losses realized in January-November, iv) Figure 12: 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 – see section 2 for a discussion of these confounding factors based on the existing literature).

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 24 to 33. Each of these Figures presents the regression results for three consecutive years.[30] 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 preiod was tax relevant) are presented in Figures 24 to 26. 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.[31] 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 27 and 28. We do not

[30]Note 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).

[31]The 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.

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 29 to 31. 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 32 and 33. As with gains, we do not observe spikes or irregularities around the 365-days holding period in any of the six post-reform years

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.

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 section 5.4).

Figures 13 and 14 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 14). 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 15 (for gains) and 16 (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 17 (for gains) and 18 (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 19 and 20 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 19 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 20 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.[32] 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 20 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.

6.5 Taxes and the Disposition Effect

Figure 21 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 6. 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. This result is consistent with the literature where the disposition effect has been shown to be very robust. 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

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

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 6 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 22 (that is, Figure 22 plots the ratio β_1/β_0). For the purpose of comparing our disposition 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 21, 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 21 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 23 which plots the coefficients of the DiD set-up (β_3 in equation 7). 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.

7 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. The 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 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.

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 heavily 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 (e.g., Feng and Seasholes (2005), Dhar and Zhu (2006) and Seru et al. (2010)). However, as our heterogeneity analyses suggest, 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 U.S. 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.[33]

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 tax discontinuity is crossed. 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.

[33]A complete smoothening of the tax schedule in the U.S. 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.

References

Abramitzky, R. and V. Lavy (2014). How responsive is investment in schooling to changes in redistributive policies and in returns? *Econometrica 82*(4), 1241–1272.

Agrawal, D. R. and D. Foremny (2019). Relocation of the rich: Migration in response to top tax rate changes from spanish reforms. *The Review of Economics and Statistics*. forthcoming.

Almunia, M. and D. Lopez-Rodriguez (2018). Under the radar: The effects of monitoring firms on tax compliance. *American Economic Journal: Economic Policy 10*(1), 1–38.

Ayers, B. C., C. E. Lefanowicz, and J. R. Robinson (2003). Shareholder taxes in acquisition premiums: The effect of capital gains taxation. *The Journal of Finance 58*(6), 2783–2801.

Ayers, B. C., O. Z. Li, and J. R. Robinson (2008). Tax-induced trading around the taxpayer relief act of 1997. *The Journal of the American Taxation Association 30*(1), 77–100.

Ball, R. (2013). Does anticipated information impose a cost on risk-averse investors? a test of the hirshleifer effect. *Journal of Accounting Research 51*(1), 31–66.

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.

Bastani, S. and H. Selin (2014). Bunching and non-bunching at kink points of the Swedish tax schedule. *Journal of Public Economics*. Forthcoming.

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.

Best, M. C., A. Brockmeyer, H. J. Kleven, J. Spinnewijn, and M. Waseem (2015). Production versus revenue efficiency with limited tax capacity: Theory and evidence from pakistan. *Journal of Political Economy 123*(6), 1311–1355.

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.

Blouin, J., L. Hail, and M. H. Yetman (2009). Capital gains taxes, pricing spreads, and arbitrage: Evidence from cross - listed firms in the u.s. *The Accounting Review 84*(5), 1321–1361.

Blouin, J. L., B. J. Bushee, and S. A. Sikes (2017). Measuring tax-sensitive institutional investor ownership. *The Accounting Review 92*(6), 49–76.

Blouin, J. L., J. S. Raedy, and D. A. Shackelford (2003). Capital gains taxes and equity trading: Empirical evidence. *Journal of Accounting Research 41*(4), 611–651.

Bogart, W. T. and W. M. Gentry (1995). Capital gains taxes and realizations: Evidence from interstate comparisons. *The Review of Economics and Statistics*, 267–282.

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.

Chetty, R., J. N. Friedman, and E. Saez (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the eitc on earnings. *American Economic Review 103*(7), 2683–2721.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review 99*(4), 1145–1177.

Chetty, R. and E. Saez (2005). Dividend taxes and corporate behavior: Evidence from the 2003 dividend tax cut. *The Quarterly Journal of Economics 120*(3), 791–833.

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.

Dai, Z., E. Maydew, D. A. Shackelford, and H. H. Zhang (2008). Capital gains taxes and asset prices: Capitalization or lock-in? *The Journal of Finance 63*(2), 709–742.

Dammon, R., H. Zhang, and C. Spatt (2001). Diversification and capital gains taxes with multiple risky assets.

Daunfeldt, S.-O., U. Praski-Staahlgren, and N. Rudholm (2010). Do high taxes lock-in capital gains? evidence from a dual income tax system. *Public Choice 145*(1), 25–38.

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.

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.

Eissa, N. and J. B. Liebman (1996). Labor supply response to the earned income tax credit. *The Quarterly Journal of Economics 111*(2), 605–637.

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 A. Rangel (2014). Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price. *Journal of Economic Behavior & Organization 107*, 541 – 552.

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.

Fuest, C., A. Peichl, and S. Siegloch (2018). Do higher corporate taxes reduce wages? micro evidence from germany. *American Economic Review 108*(2), 393–418.

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: Vermoegen & schulden [income and consumption sample statistic: Assets & liabillites]. Technical report.

Glogowsky, U. (2016). Behavioral responses to wealth transfer taxation: Bunching evidence from germany. Technical report.

Grinblatt, M. and M. Keloharju (2001). What makes investors trade? *The Journal of Finance 56*(2), 589–616.

Grinblatt, M. and M. Keloharju (2004). Tax-loss trading and wash sales. *Journal of Financial Economics 71*(1), 51 – 76.

Grinblatt, M. and T. J. Moskowitz (2004). Predicting stock price movements from past returns: the role of consistency and tax-loss selling. *Journal of Financial Economics 71*(3), 541 – 579.

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.

Hoopes, J., P. Langetieg, S. Nagel, D. Reck, J. Slemrod, and B. Stuart (2016). Who sold during the crash of 2008-9? evidence from tax-return data on daily sales of stock. NBER Working Paper 22209.

Huang, J. (2001). Taxable or tax-deferred account? portfolio decision with multiple investment goals. In *Mimeo*. MIT Sloan School of Management.

Ivković, Z., J. Poterba, and S. Weisbenner (2005). Tax-motivated trading by individual investors. *American Economic Review 95*(5), 1605–1630.

Jacob, M. (2013). Capital gains taxes and the realization of capital gains and losses: Evidence from german income tax data. *FinanzArchiv: Public Finance Analysis 69*(1), 30–56.

Jacob, M. (2018). Tax regimes and capital gains realizations. *European Accounting Review 27*(1), 1–21.

Jakobsen, K., K. Jakobsen, H. Kleven, and G. Zucman (2018). Wealth taxation and wealth accumulation: Theory and evidence from denmark. Working Paper 24371, National Bureau of Economic Research.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics 8*, 435–464.

Kleven, H. J., C. Landais, E. Saez, and E. Schultz (2014). Migration and wage effects of taxing top earners: Evidence from the foreigners' tax scheme in denmark *. *The Quarterly Journal of Economics 129*(1), 333–378.

Kleven, H. J. and E. A. Schultz (2014, November). Estimating taxable income responses using danish tax reforms. *American Economic Journal: Economic Policy 6*(4), 271–301.

Korniotis, G. M. and A. Kumar (2011). Do older investors make better investment decisions? *The Review of Economics and Statistics 93*(1), 244–265.

Lang, M. H. and D. A. Shackelford (2000). Capitalization of capital gains taxes: evidence from stock price reactions to the 1997 rate reduction. *Journal of Public Economics 76*(1), 69 – 85.

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.

Martinez, I. Z., E. Saez, and M. Siegenthaler (2018, May). Intertemporal labor supply substitution? evidence from the swiss income tax holidays. Working Paper 24634, National Bureau of Economic Research.

Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of Finance 53*(5), 1775–1798.

Poterba, J. M. (2001). Taxation and portfolio structure: Issues and implications. NBER Working Paper 8223.

Poterba, J. M. (2002). Taxation, risk-taking, and household portfolio behavior. In *Handbook of public economics*, Volume 3, pp. 1109–1171. Elsevier.

Poterba, J. M. and S. J. Weisbenner (2001). Capital gains tax rules, tax-loss trading, and turn-of-the-year returns. *The Journal of Finance 56*(1), 353–368.

Protopapadakis, A. (1983). Some indirect evidence on effective capital gains tax rates. *Journal of Business*, 127–138.

Reese Jr, W. A. (1998). Capital gains taxation and stock market activity: Evidence from ipos. *The Journal of Finance 53*(5), 1799–1819.

Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy 2*(3), 180–212.

Saez, E. (2017). Taxing the rich more: Preliminary evidence from the 2013 tax increase. *Tax Policy and the Economy 31*(1), 71–120.

Saez, E., J. Slemrod, and S. H. Giertz (2012). The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of Economic Literature 50*(1), 3–50.

Seida, J. A. and W. F. Wempe (2000). Do capital gain tax rate increases affect individual investors' trading decisions? *Journal of Accounting and Economics 30*(1), 33 – 57.

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.

Shoven, J. B. and C. Sialm (2004). Asset location in tax-deferred and conventional savings accounts. *Journal of Public Economics 88*(1-2), 23–38.

Sikes, S. (2018). Capital gains lock-in and share repurchases. mimeo.

Suarez Serrato, J. C. and O. Zidar (2016). Who benefits from state corporate tax cuts? a local labor markets approach with heterogeneous firms. *American Economic Review 106*(9), 2582–2624.

van Rooij, M., A. Lusardi, and R. Alessie (2011). Financial literacy and stock market participation. *Journal of Financial Economics 101*(2), 449–472.

Yagan, D. (2015, December). Capital tax reform and the real economy: The effects of the 2003 dividend tax cut. *American Economic Review 105*(12), 3531–63.

Tables and Figures

Table 1: Descriptive statistics for all investors in the sample

	N	Mean	Std.Dev.	Min	Max
Unit of observation: Investor					
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
Unit of observation: Share package					
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 4.2 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. 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 monthes 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 1: Raw data: Number of share packages sold around time discontinuity: Gains

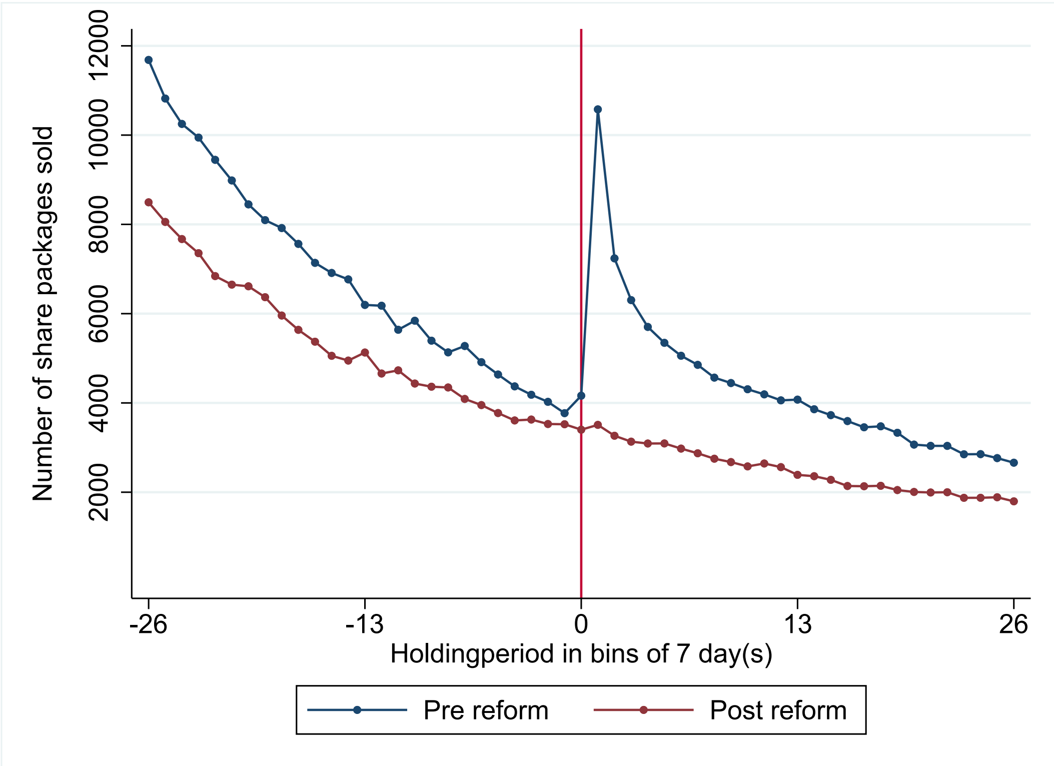

Notes: This figure displays the number of share packages which were sold with a gain in dependency of the holding period. Each dot represents the number of share packages sold in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last week in which gains were taxable. All details are described in section 5.1. 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 x-axis value zero marks the last week in which gains were taxable. 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.

Figure 2: Raw data: Number of share packages sold around time discontinuity: Losses

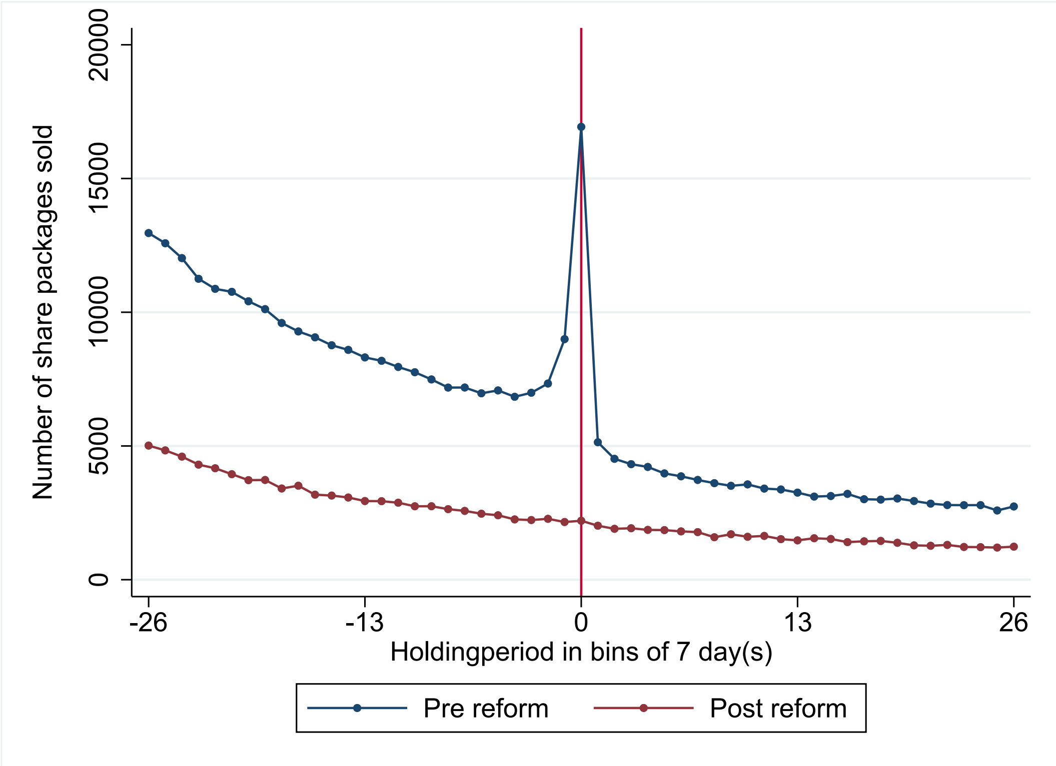

Notes: This figure displays the number of share packages which were sold with a loss in dependency of the holding period. Each dot represents the number of share packages sold in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last week in which losses could be used to offset gains. All details are described in section 5.1. 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 x-axis value zero marks the last week in which losses could be used to offset taxes. 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 3: Number of distinct investors trading around time discontinuity: Gains

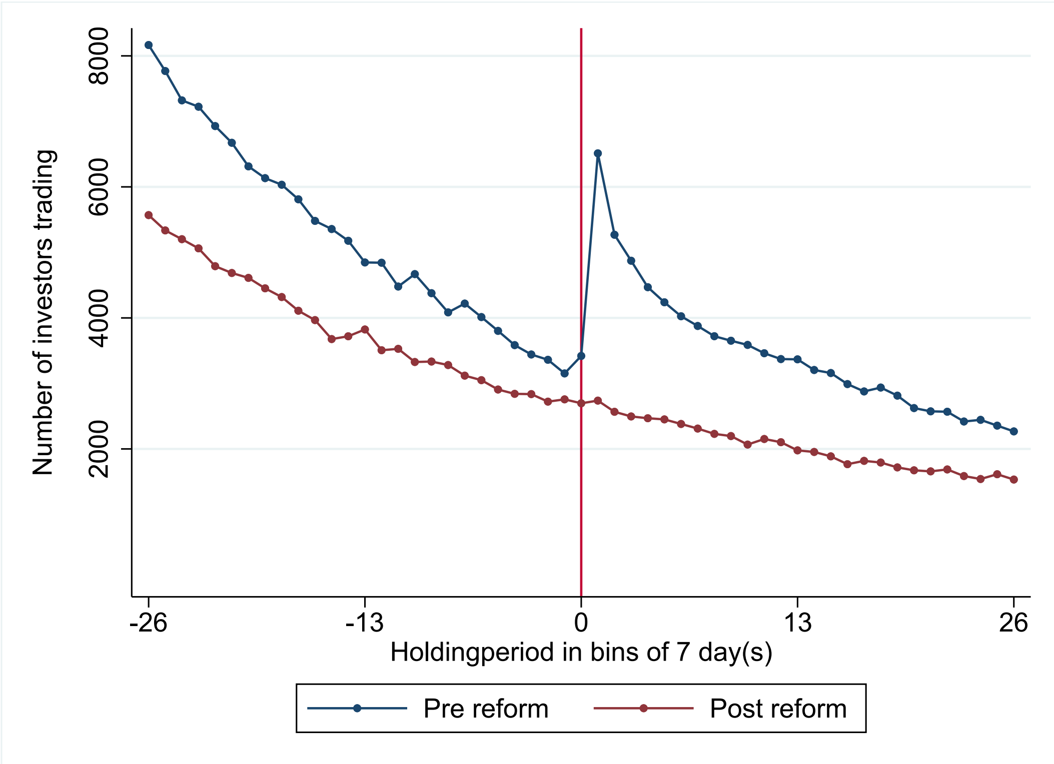

Notes: This figure displays the number of investors who sold an appreciated share package with the respective holding period. Each dot represents the number of investors who sold a share package in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last week in which gains were taxable. All details are described in section 5.1. 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 gains were taxable. Pre reform estimates are based on 44110 investors and 230352 share packages. Post reform estimates are based on 30875 investors and 155603 share packages.

Figure 4: Number of distinct investors trading around time discontinuity: Losses

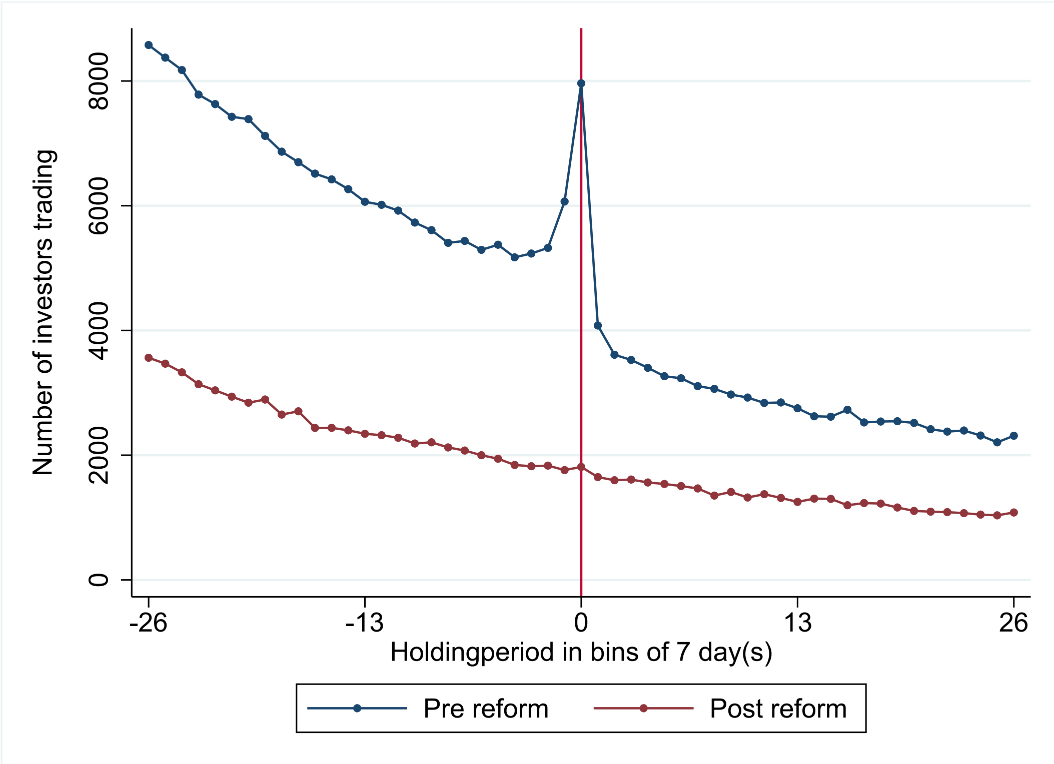

Notes: This figure displays the number of investors who sold a depreciated share package with the respective holding period. Each dot represents the number of investors who sold a share package in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last week in which losses could be used to offset gains. All details are described in section 5.1. 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. 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 5: Difference in bunching: Gains

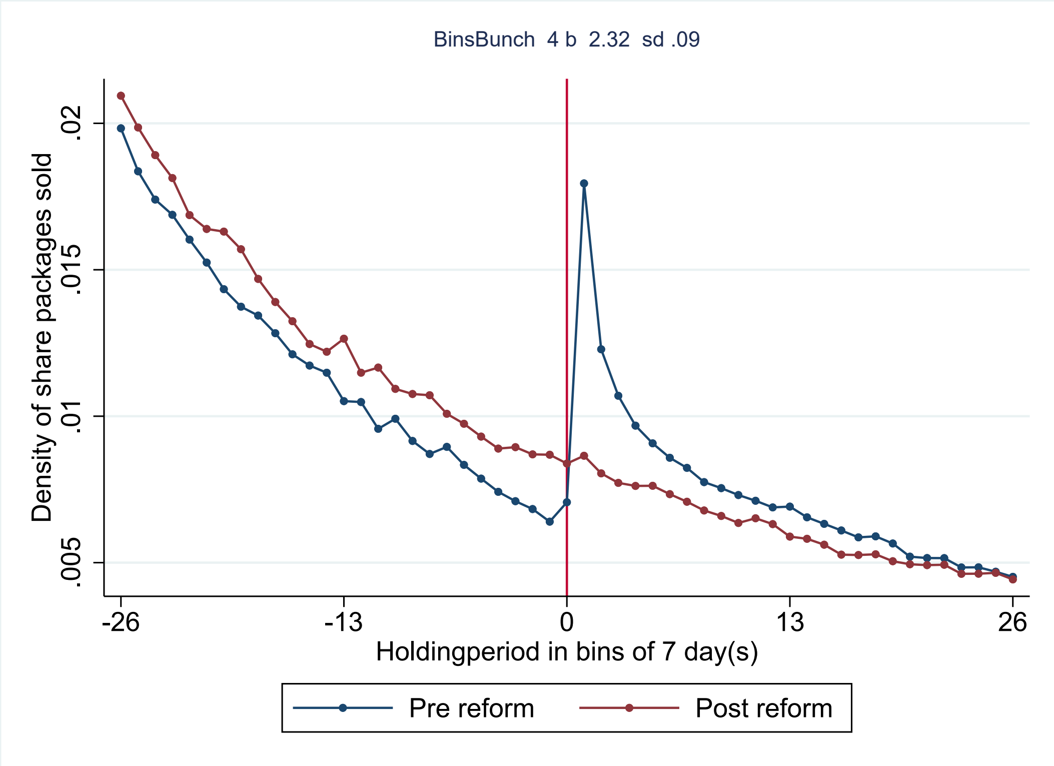

Notes: This figure displays the share of all purchased share packages with a gain in dependency of the holding period. Each dot represents the share of all purchased share packages with a gain which were sold in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last week in which gains were taxable. All details are described in section 5.2. The dotted blue line represents the share of all share packages with a gain purchased prior to 2009 which were sold. The dotted red line represents the share of all share packages with a gain purchased after 2009 which were sold. The vertical red line at x-axis value zero marks the last bin in which gains were taxable. **BinsBunch** denotes the bunching window which in this case includes the 4 bins right after the last week in which losses could be used to offset gains. **b** represents the excess mass and **sd** the standard errors which are bootstrapped on the investor level. Pre reform estimates are based on 57944 investors and 589254 share packages. Post reform estimates are based on 43584 investors and 405628 share packages. These numbers include share packages of shares which have not been sold in the 26 weeks after the tax discontinuity.

Figure 6: Difference in bunching: Losses

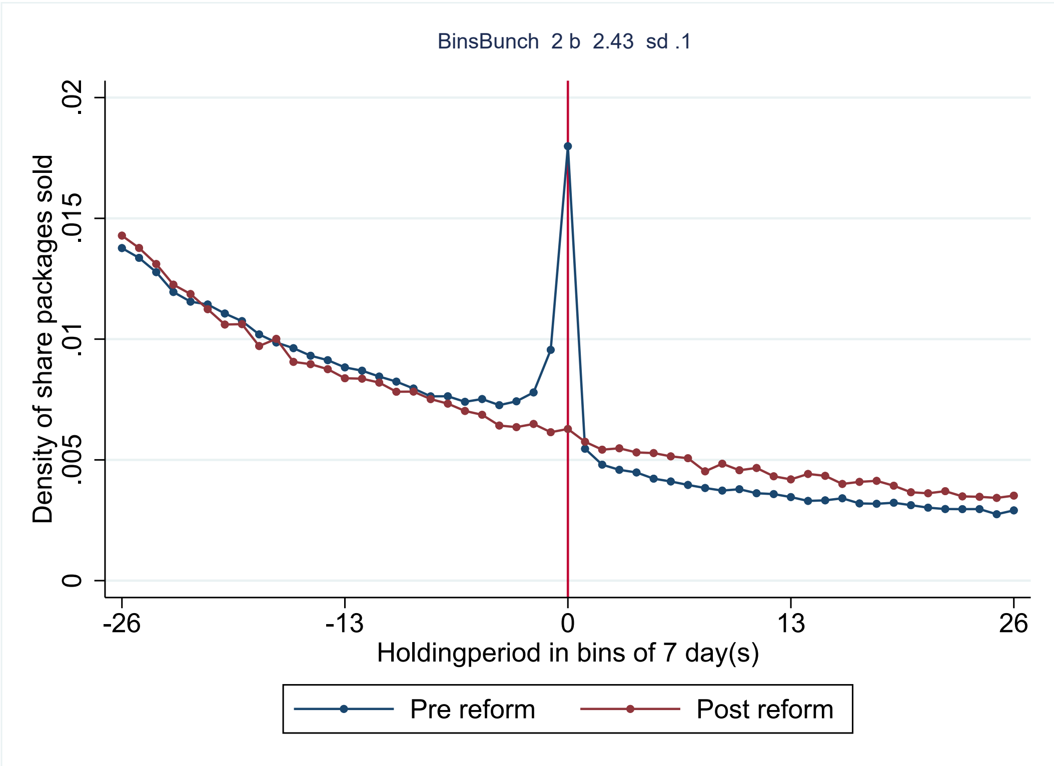

Notes: This figure displays the share of all purchased share packages with a loss in dependency of the holding period. Each dot represents the share of all purchased share packages with a loss which were sold in a 7 days bin of the holding period. Data is shown for 26 weeks before and 26 weeks after the last bin in which losses could be used to offset gains. All details are described in section 5.2. The dotted blue line represents the share of all share packages with a loss purchased prior to 2009 which were sold. The dotted red line represents the share of all share packages with a loss purchased after 2009 which were sold. The vertical red line at x-axis value zero marks the last bin in which losses could be used to offset gains. **BinsBunch** denotes the bunching window which in this case includes the last 2 weeks in which taxes were taxable. **b** represents the excess mass and **sd** the standard errors which are bootstrapped on the investor level. 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.

Figure 7: Non-parametric regressions: Gains

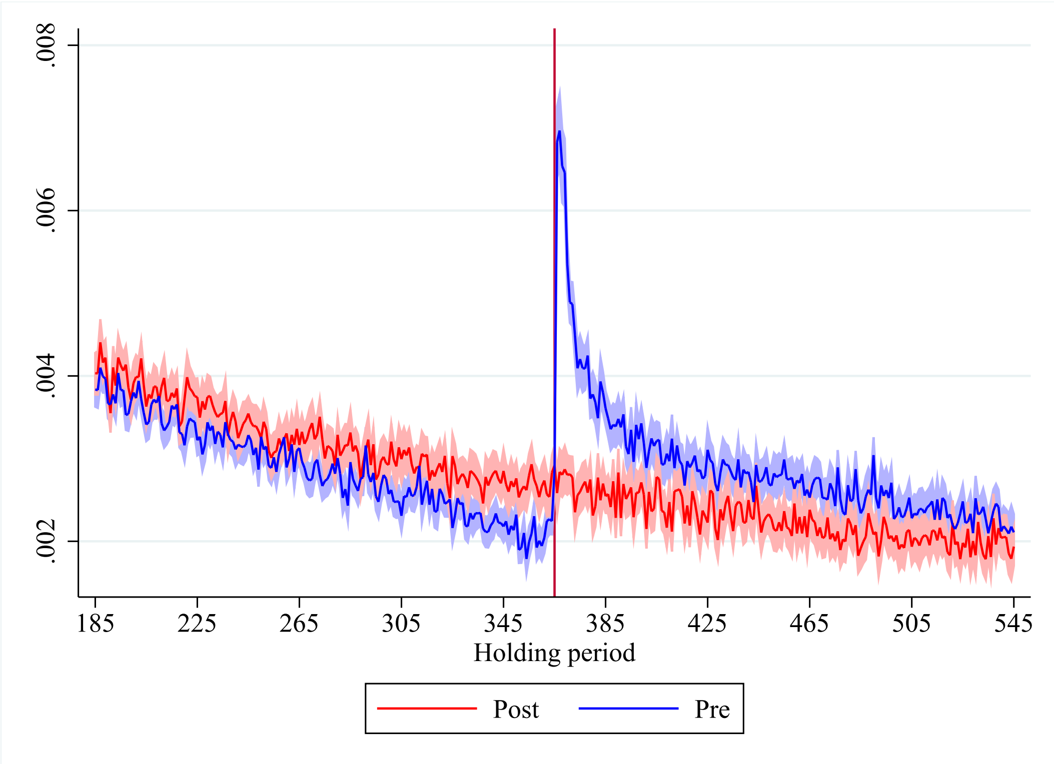

Notes: This figure displays non-parametric regressions estimates for each day of the holding period for share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$. All estimation details are described in section 5.4. The blue line represents estimates for β_0 for share packages which were bought before 2009. The shaded blue area displays 95 percent confidence intervals. The red line represents estimates for β_0 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. 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.

Figure 8: Non-parametric regressions: Losses

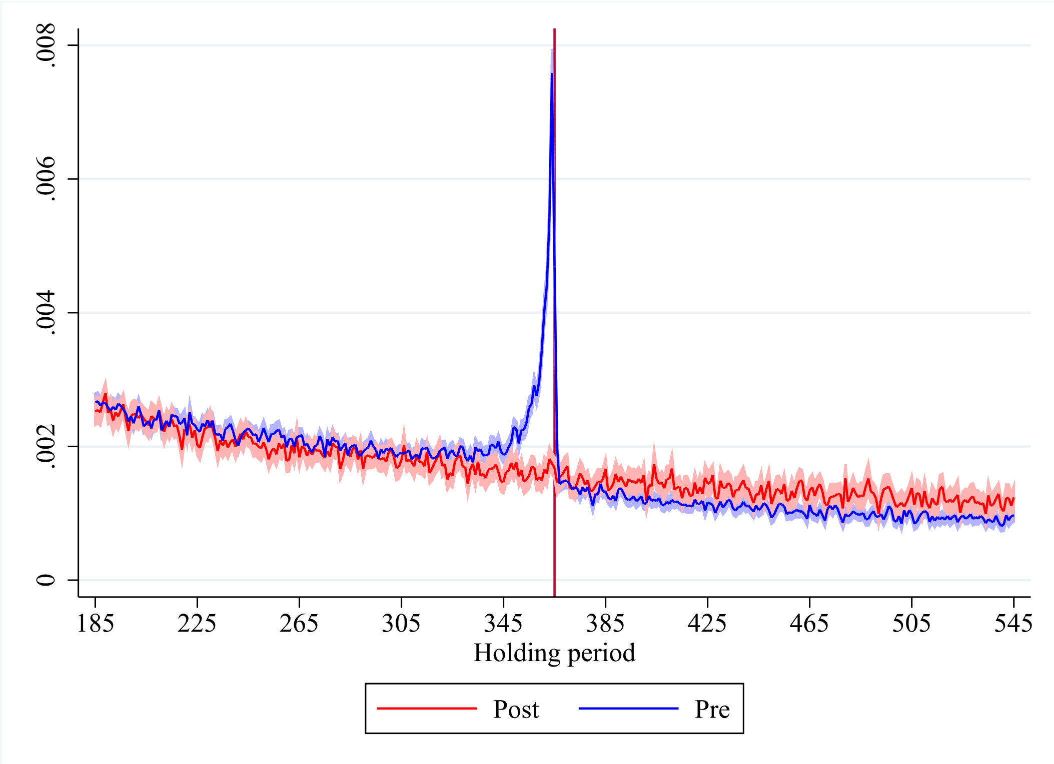

Notes: This figure displays non-parametric regressions estimates for each day of the holding period for share packages with prices below the purchase price at the respective day. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$. All estimation details are described in section 5.4. 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. 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.

Table 2: Non-parametric regressions: Coefficients for 10 day window around the discontinuity.

	361	362	363	364	365	366	367	368	369	370
Gains, Pre-Reform										
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
Gains, Post-Reform										
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
Losses, Pre-Reform										
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
Losses, Post-Reform										
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: $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$. 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 25 for more explanation). All estimation details are described in section 5.4. The table mirrors the results in Figures 7 and 8.

Figure 9: Non-parametric regressions: Gains, January-November

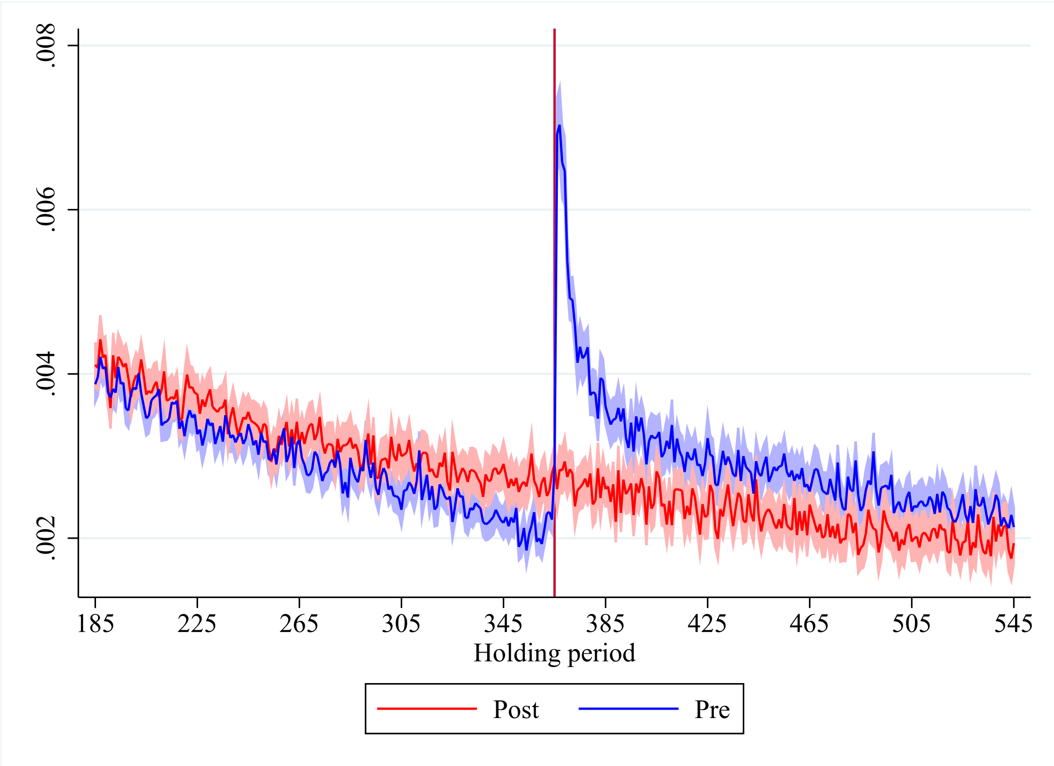

Notes: This figure displays the non-parametric regression estimates for each day of the holding period for share packages with prices above the purchase price. In these regressions we only include share packages on calendar dates not in December. Coefficients indicate the probability that a share-package 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 $\mathbb{1}(Gain_{ijd}) = 1$ & $\mathbb{1}(December_{ijd}) = 0$. All estimation details are described in section 5.4. The blue line represents estimates for β_0 for share packages which were bought before 2009. The shaded blue area displays 95 percent confidence intervals. The red line represents estimates for β_0 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. 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.

Figure 10: Non-parametric regressions: Gains, December

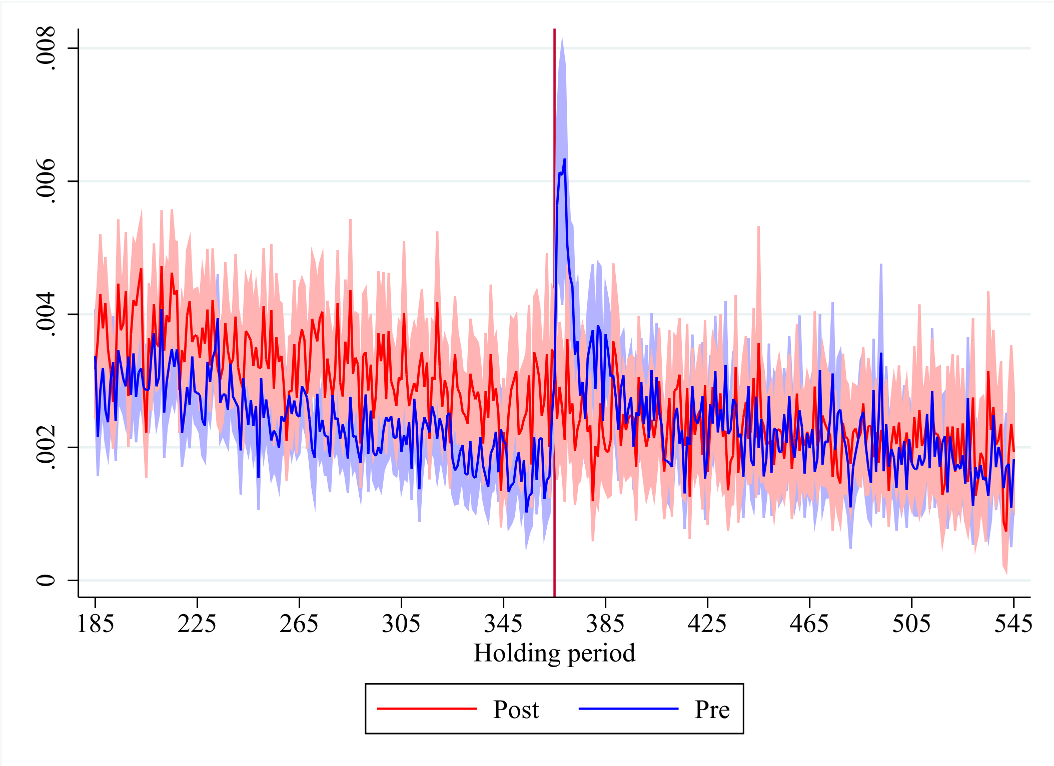

Notes: This figure displays the non-parametric regression estimates for each day of the holding period for share packages with prices above the purchase price. In these regressions we only include share packages on calendar dates in december. Coefficients indicate the probability that a share-package 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 $\mathbb{1}(Gain_{ijd}) = 1$ & $\mathbb{1}(December_{ijd}) = 1$. All estimation details are described in section 5.4. The blue line represents estimates for β_0 for share packages which were bought before 2009. The shaded blue area displays 95 percent confidence intervals. The red line represents estimates for β_0 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. 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 11: Non-parametric regressions: : Losses, January-November

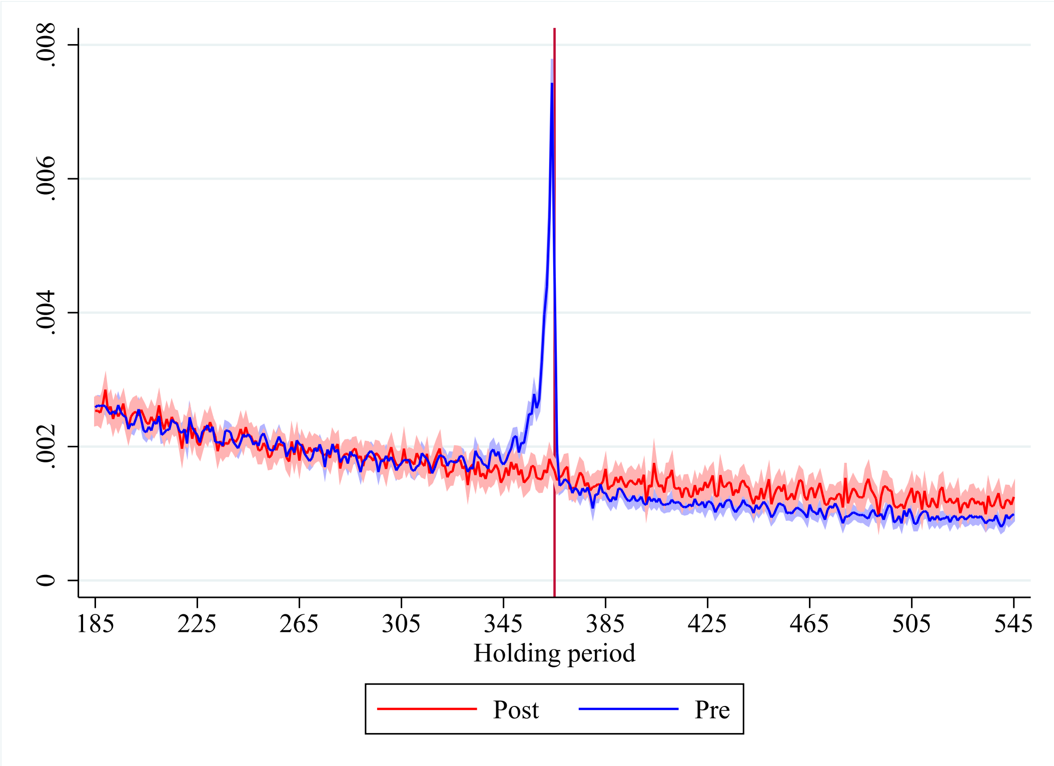

Notes: This figure displays the non-parametric regression estimates for each day of the holding period for share packages with prices below the purchase price at the respective day. In these regressions we only include share packages on calendar dates in december. Coefficients indicate the probability 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 $\mathbb{1}(Loss_{ijd}) = 1$ & $\mathbb{1}(December_{ijd}) = 0$. All estimation details are described in section 5.4. 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. 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.

Figure 12: Non-parametric regressions: Losses, December

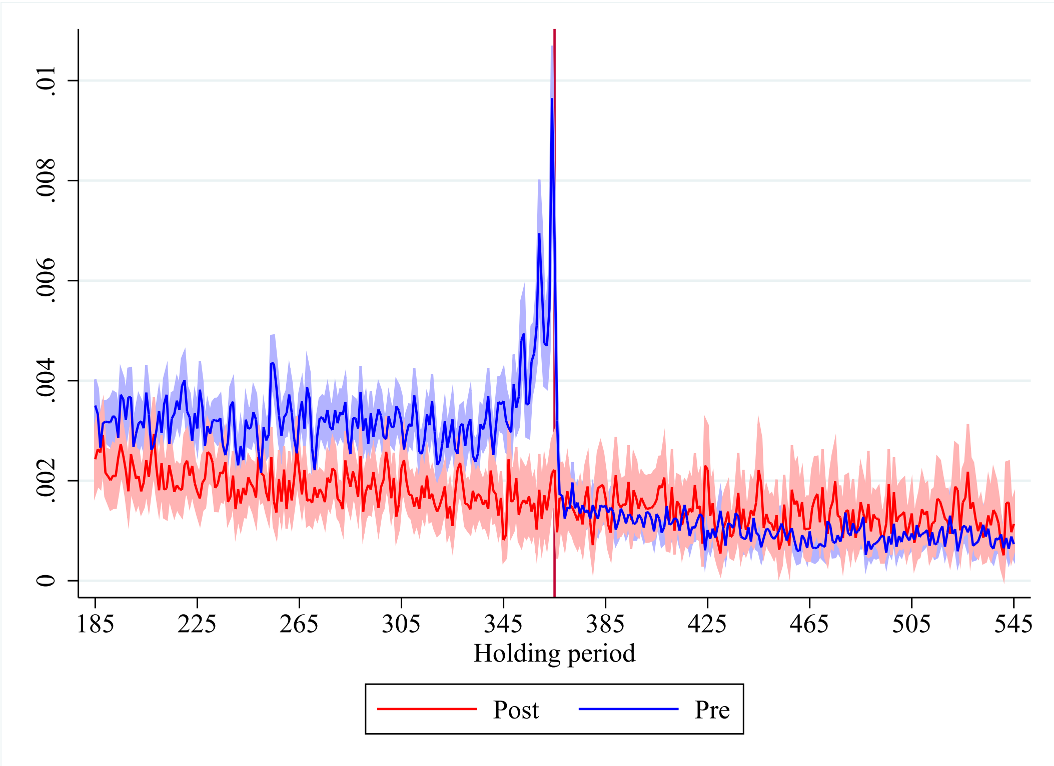

Notes: This figure displays the non-parametric regression estimates for each day of the holding period for share packages with prices below the purchase price at the respective day. In these regressions we only include share packages on calendar dates in december. Coefficients indicate the probability 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 $\mathbb{1}(Loss_{ijd}) = 1$ & $\mathbb{1}(December_{ijd}) = 1$. All estimation details are described in section 5.4. 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. 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 13: Heterogeneity w.r.t. age: Gains

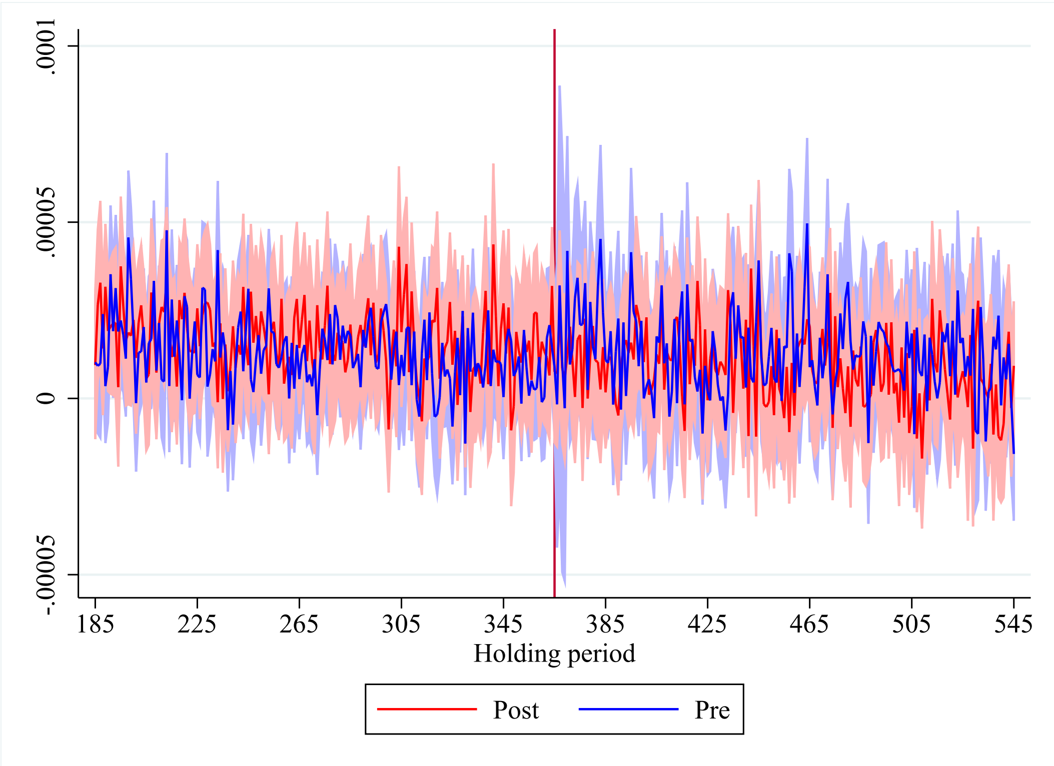

Notes: This figure displays coefficient estimates for investor experience stemming from non-parametric regressions for each day of the holding period. Included are share packages with prices above the purchase price. Coefficients indicate by how much an additional year in age shifts the probability that 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 $\mathbb{1}(Gain_{ijd}) = 1$. Where Age is the age of the investor on a given calendar 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 5.4. The blue line represents estimates for β_1 for share packages which were bought before 2009. The shaded area displays 95 percent confidence intervals. 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. 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.

Figure 14: Heterogeneity w.r.t. age: Losses

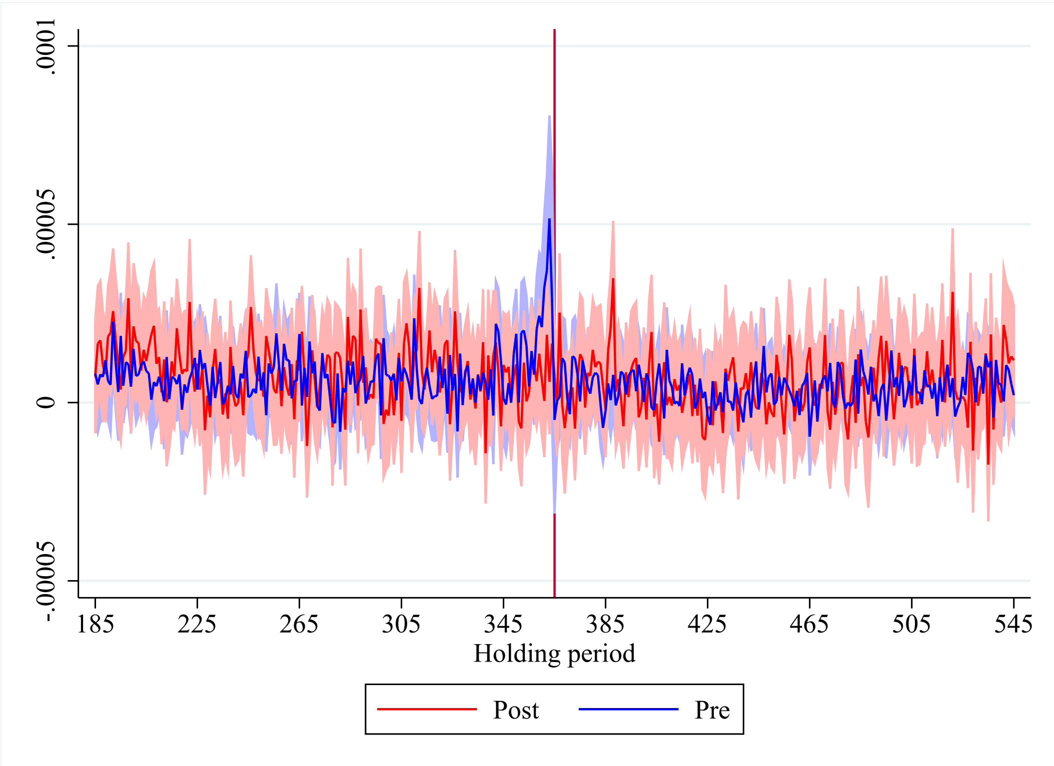

Notes: This figure displays coefficient estimates for investor experience stemming from non-parametric regressions for each day of the holding period. Included are share packages with prices below the purchase price. Coefficients indicate by how much an additional year in age shifts the probability that 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 $\mathbb{1}(Loss_{ijd}) = 1$. 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 5.4. 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. 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 15: Heterogeneity w.r.t. experience: Gains

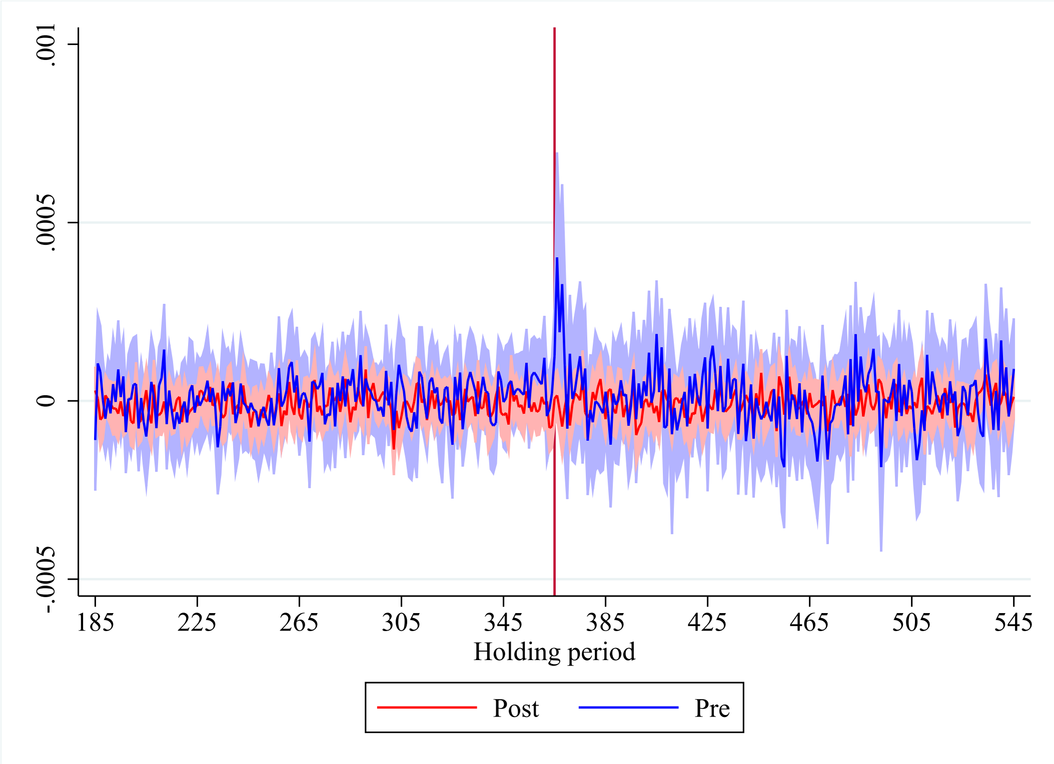

Notes: This figure displays coefficient estimates for investor experience stemming from non-parametric regressions for each day of the holding period. Included are share packages with prices above the purchase price. Coefficients indicate by how much an additional year in experience shifts the probability 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 $\mathbb{1}(Gain_{ijd}) = 1$. 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 5.4. The blue line represents estimates for β_1 for share packages which were bought before 2009. The shaded area displays 95 percent confidence intervals. 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. 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.

Figure 16: Heterogeneity w.r.t. experience: Losses

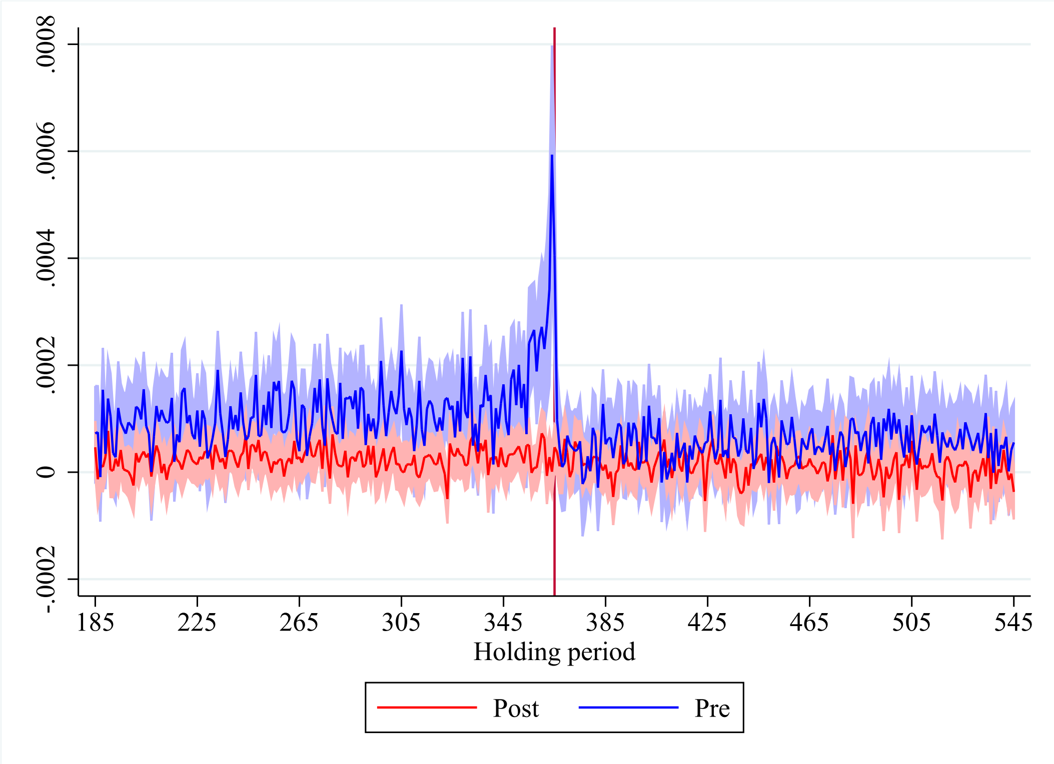

Notes: This figure displays coefficient estimates for investor experience stemming from non-parametric regressions for each day of the holding period. Included are share packages with prices below the purchase price. Coefficients indicate by how much an additional year in experience shifts the probability 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 $\mathbb{1}(Loss_{ijd}) = 1$. 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 5.4. 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. 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 17: Heterogeneity w.r.t. gender: Gains

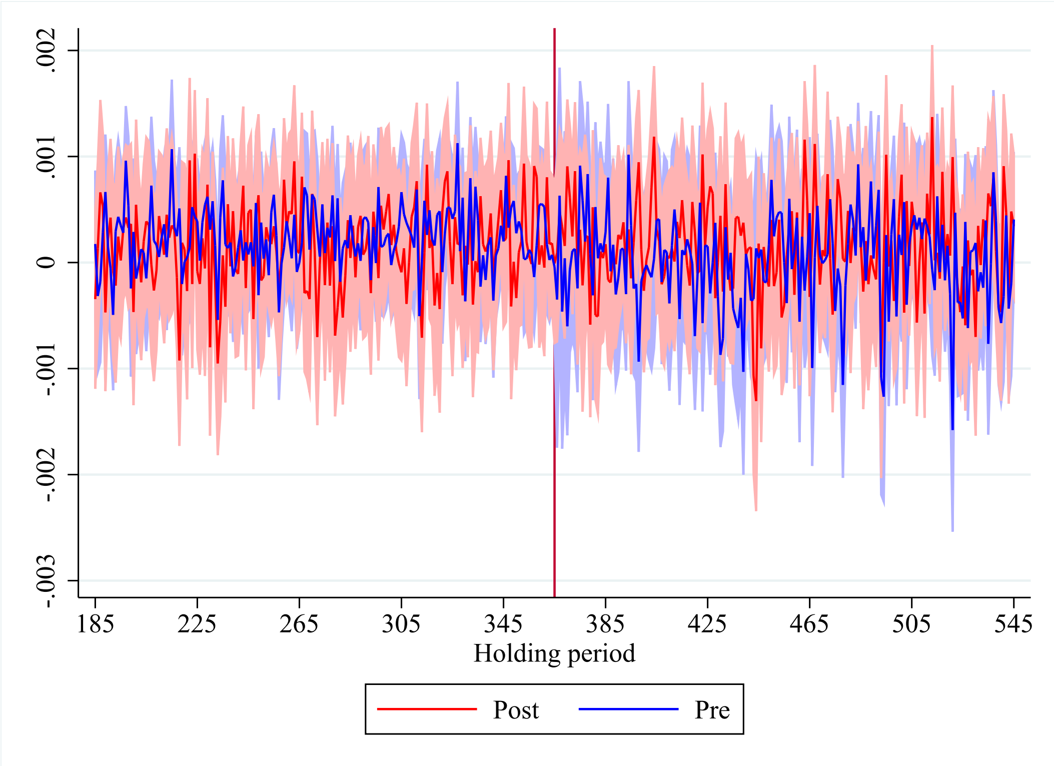

Notes: This figure displays coefficient estimates for a male dummy in the non-parametric regressions for each day of the holding period. Included are share packages with prices above the purchase price. Coefficients indicate the difference in selling probability of a share-package between men and women. 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 $\mathbb{1}(Loss_{ijd}) = 1$ 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 5.4. 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. 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.

Figure 18: Heterogeneity w.r.t. gender: Losses

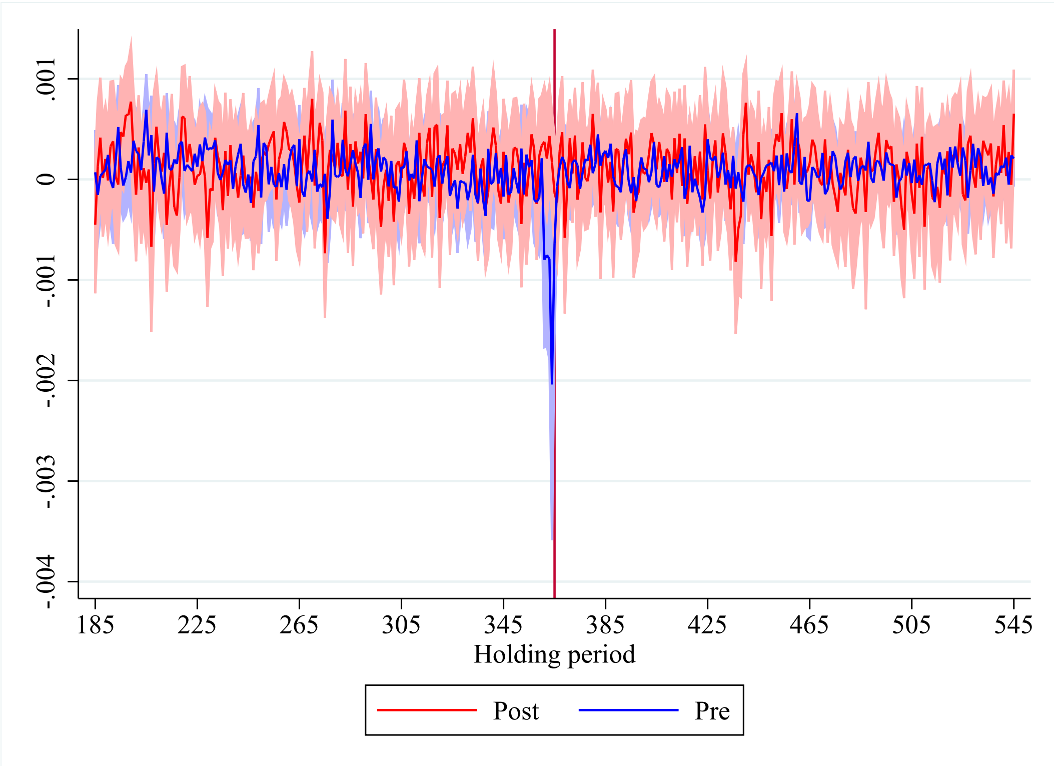

Notes: This figure displays coefficient estimates for a male dummy in the non-parametric regressions for each day of the holding period. Included are share packages with prices below the purchase price. Coefficients indicate the difference in selling probability of a share-package between man and woman. 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 $\mathbb{1}(Loss_{ijd}) = 1$. 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 5.4. 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. 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 19: Heterogeneity w.r.t. price-change magnitude: Gains

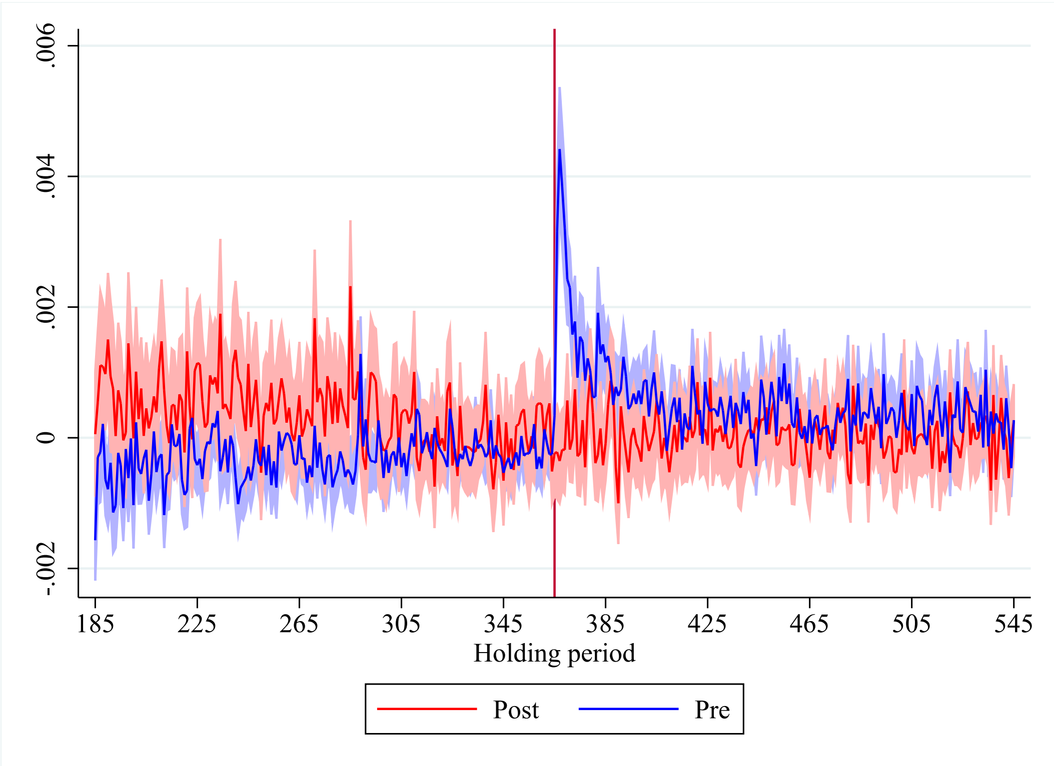

Notes: This figure displays coefficient estimates for the size of a gain from non-parametric regressions for each day of the holding period. Included are share packages with prices above the purchase price. Coefficients indicate by how much an additional percentage point increase in the price increases 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 $Sell_{ijd} = \beta_0 + \beta_1 Change_{ijd} + \varepsilon_{ijd}$ if $\mathbb{1}(Gain_{ijd}) = 1$. Where $Change_{ijd}$ is measured as $\frac{p_{ijtd} - p_{ij0d}}{p_{ij0d}}$. All estimation details are described in section 5.4. 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. 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.

Figure 20: Heterogeneity w.r.t. price-change magnitude: Losses

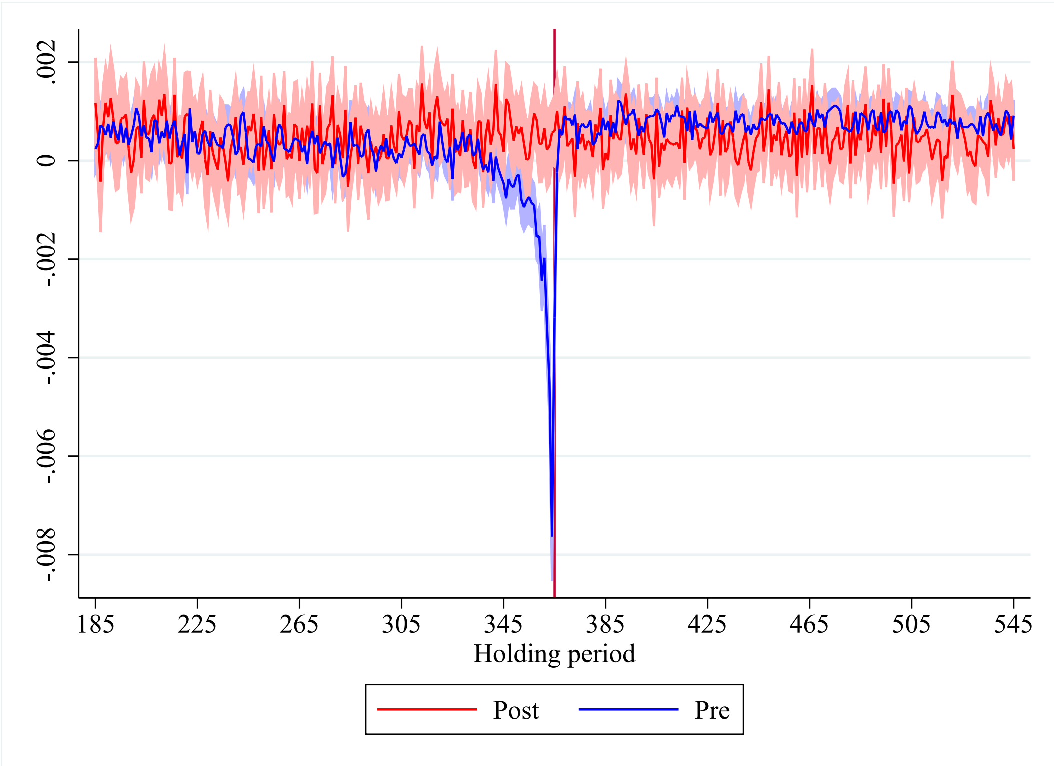

Notes: This figure displays coefficient estimates for the size of a gain from non-parametric regressions for each day of the holding period. Included are share packages with prices below the purchase price. Coefficients indicate by how much an additional percentage point decrease in the price changes 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 $\mathbb{1}(Gain_{ijd}) = 1$. Where $Change_{ijd}$ is measured as $\frac{p_{ijtd} - p_{ij0d}}{p_{ij0d}}$. All estimation details are described in section 5.4. 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. 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 21: Disposition effect around time discontinuity

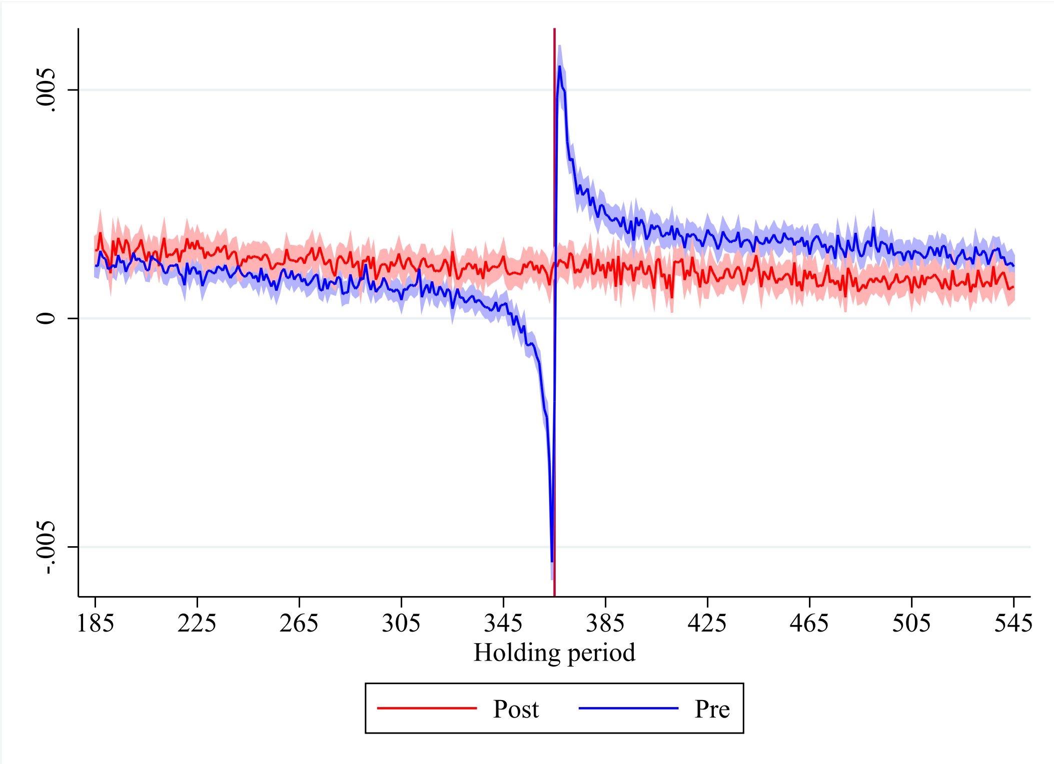

Notes: This figure displays estimates for the average difference in selling probability between gains and losses on each day of the holding period. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \beta_1 Gain_{ijd} + \varepsilon_{ijd}$. All estimation details are described in section 5.5. 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 22: Disposition effect: Gain coefficients relative to loss coefficients

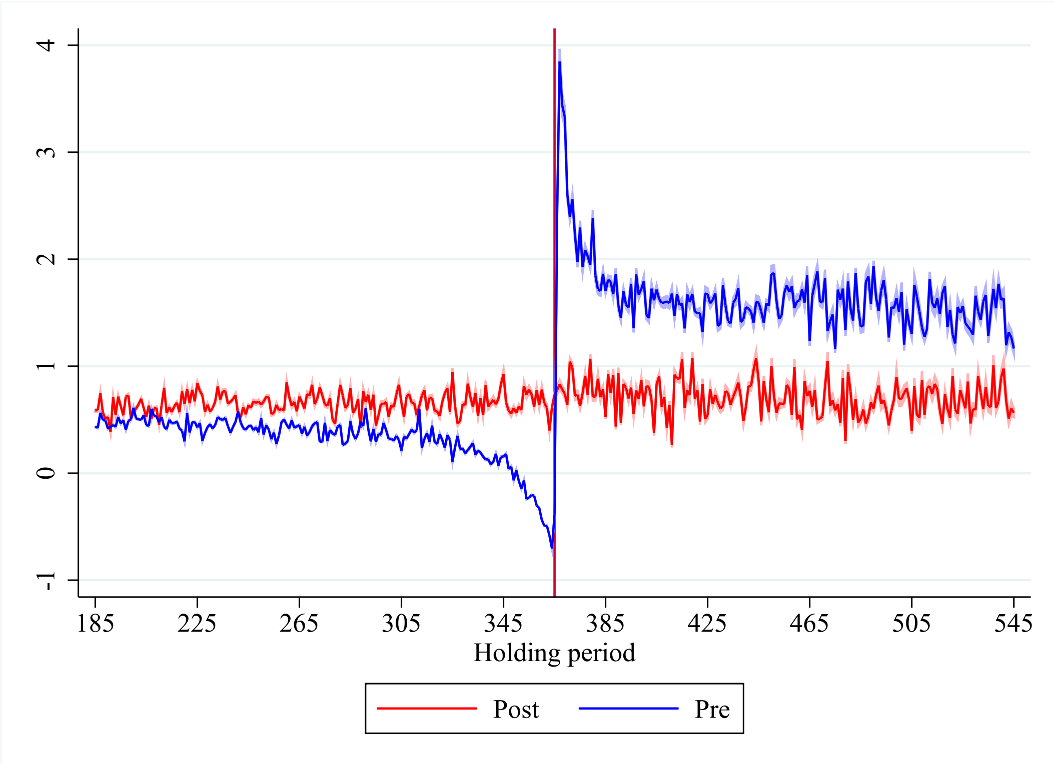

Notes: This figure displays estimates for the relative difference in selling probability between gains and losses on each day of the holding period. That is the the coefficient of the gain dummy is divided by the constant. Standard errors are calculated using the delta method. Coefficients and on the 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 5.5. 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.

Figure 23: Disposition effect around time discontinuity: DiD model

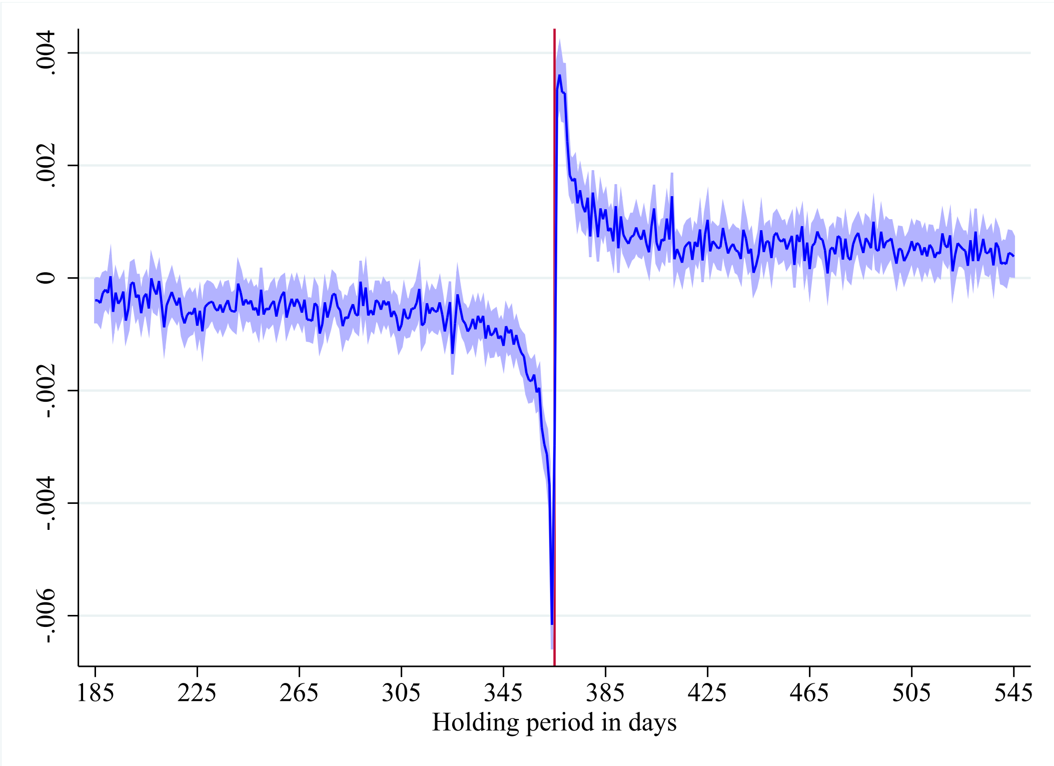

Notes: This figure displays difference in difference estimates for the average difference in selling probability between gains and losses on each day of the holding period. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \beta_1 Pre + \beta_2 \mathbb{1}(Gain_{ijd}) + \beta_3 Pre \times \mathbb{1}(Gain_{ijd}) + \varepsilon_{ijd}$. All estimation details are described in section 5.5. 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.

Appendix

Figure 24: Non-parametric regressions by year: Gains, pre Years 1999-2001

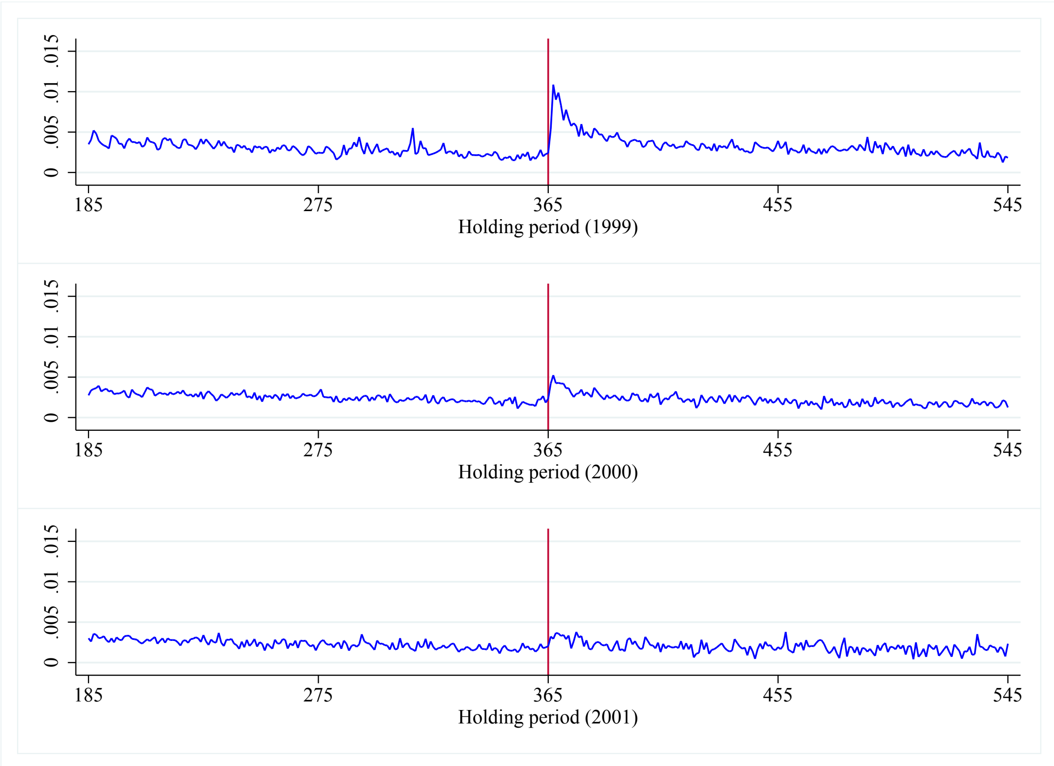

Notes: This figure displays non-parametric regression estimates for each day of the holding period separately for the years 1999-2001. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients 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 5.4. 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 25: Non-parametric regressions by year: Gains, pre years 2002-2004

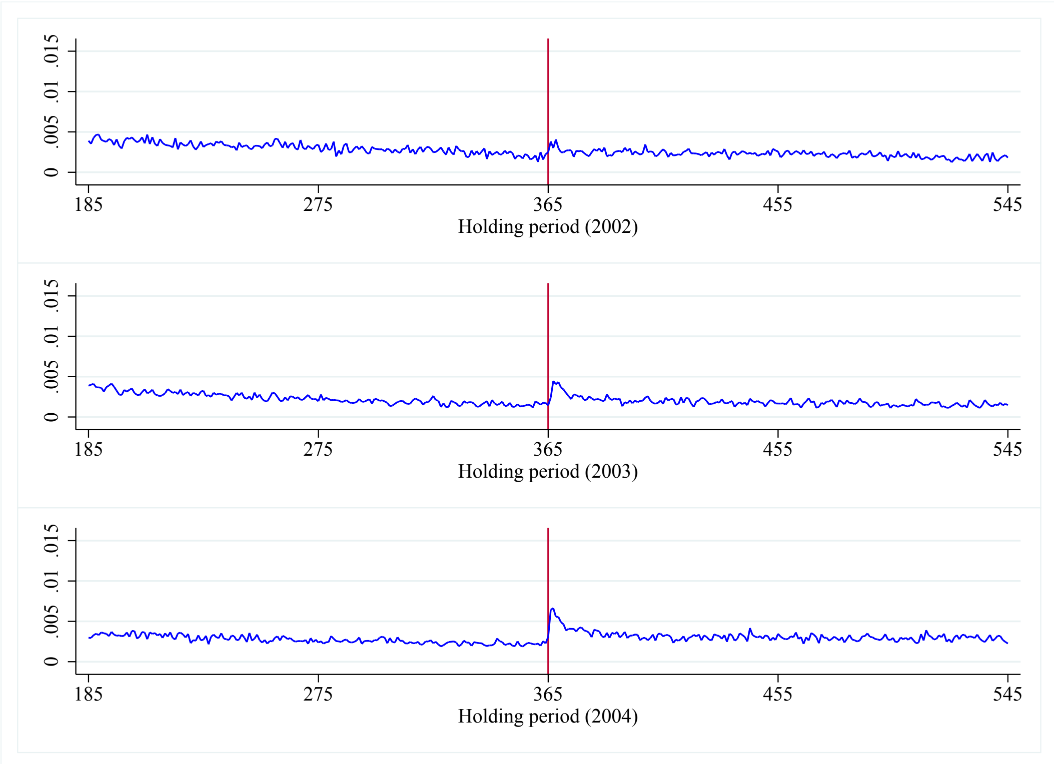

Notes: This figure displays non-parametric regression estimates for each day of the holding period separately for the years 2002-2004. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients 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 5.4. 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 26: Non-parametric regressions by year: Gains, pre Years 2005-2007

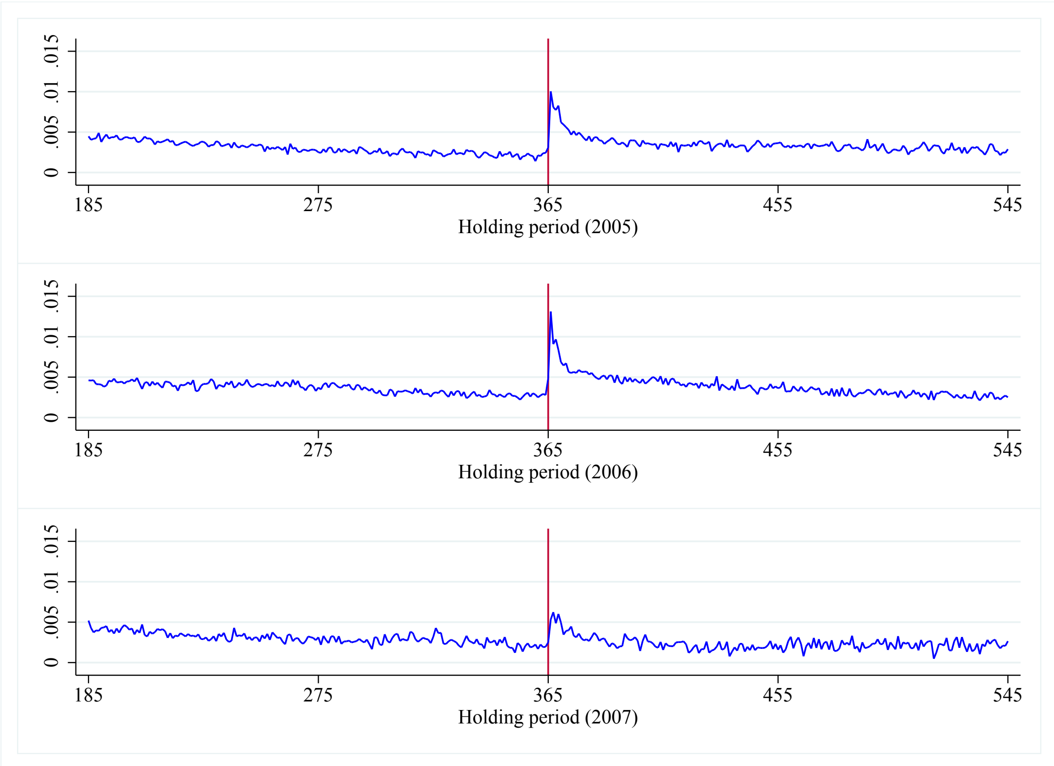

Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2005-2007. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients 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 5.4. 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 27: Non-parametric regressions by year: Gains, post years 2009-2011

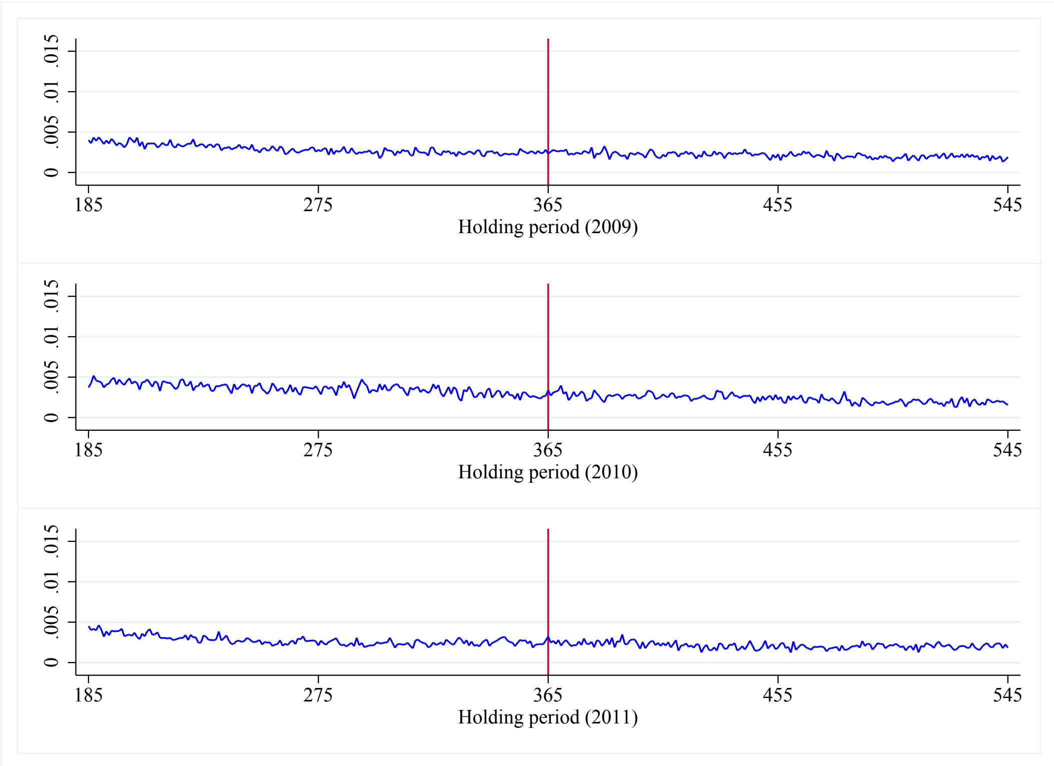

*Notes:*This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2009-2011. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. 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 5.4. 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 28: Non-parametric regressions by year: Gains, post years 2012-2014

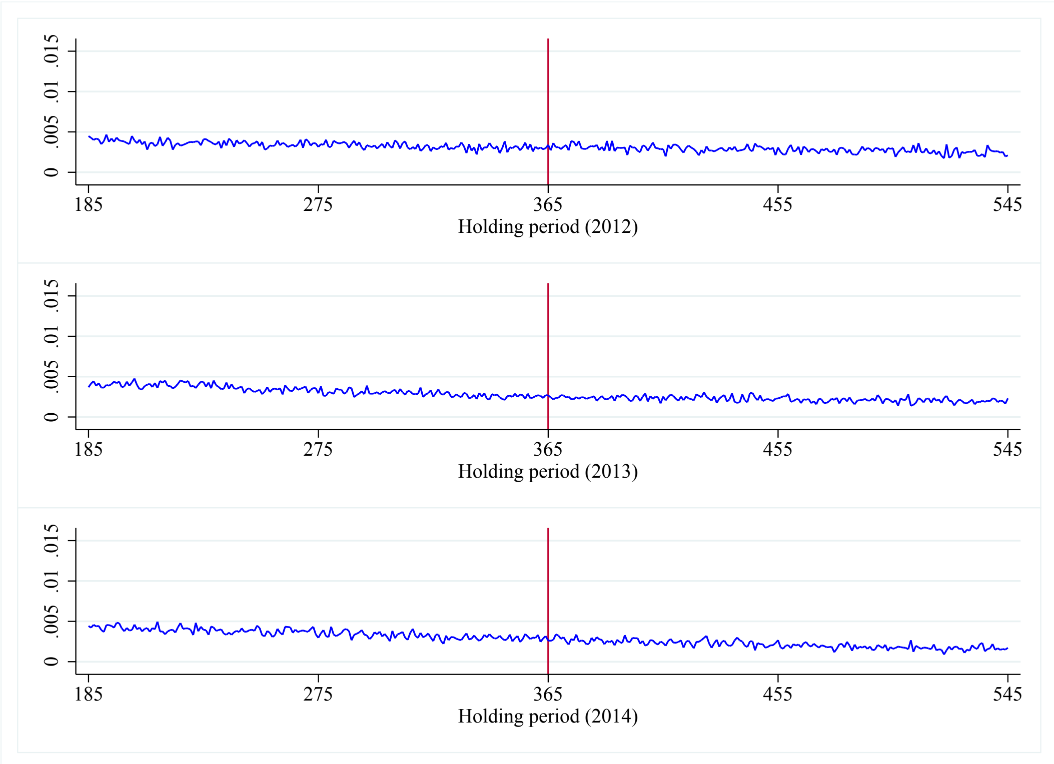

Notes: This figure displays non-parametric regression estimates for each day of the holding period separately for the years 2012-2014. Included are share packages with prices above the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients 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 5.4. 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 29: Non-parametric regressions by year: Losses, pre years 1999-2001

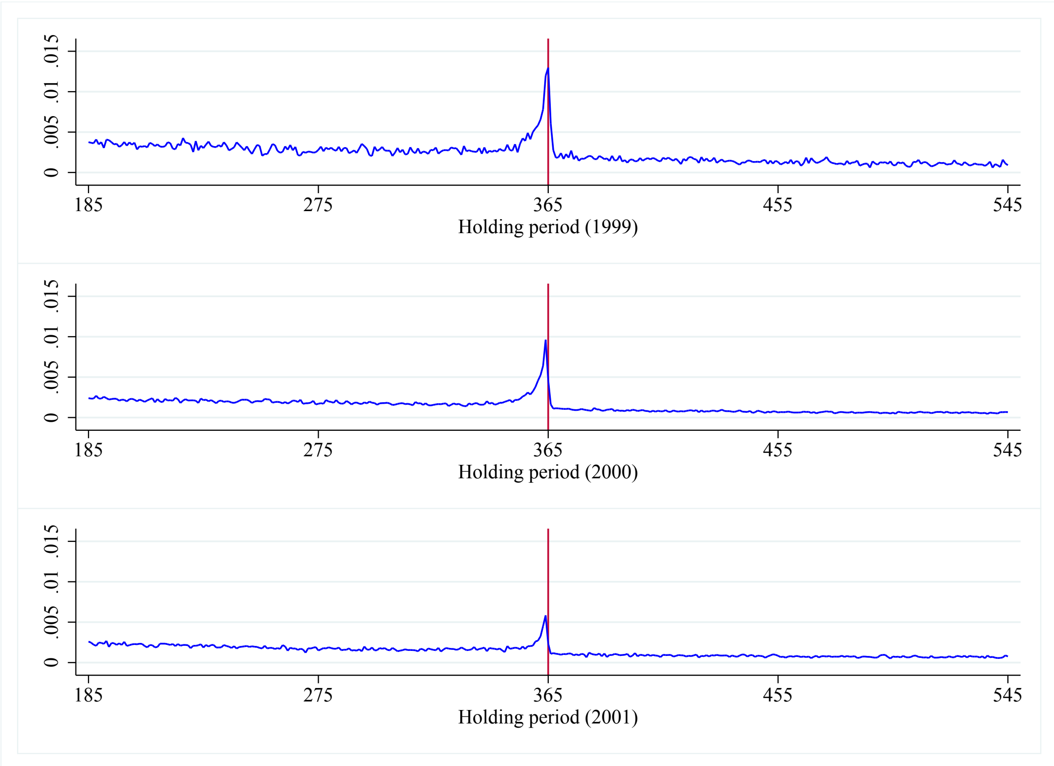

Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 1999-2001. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients 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 5.4. 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 30: Non-parametric regressions by year: Losses, pre years 2002-2004

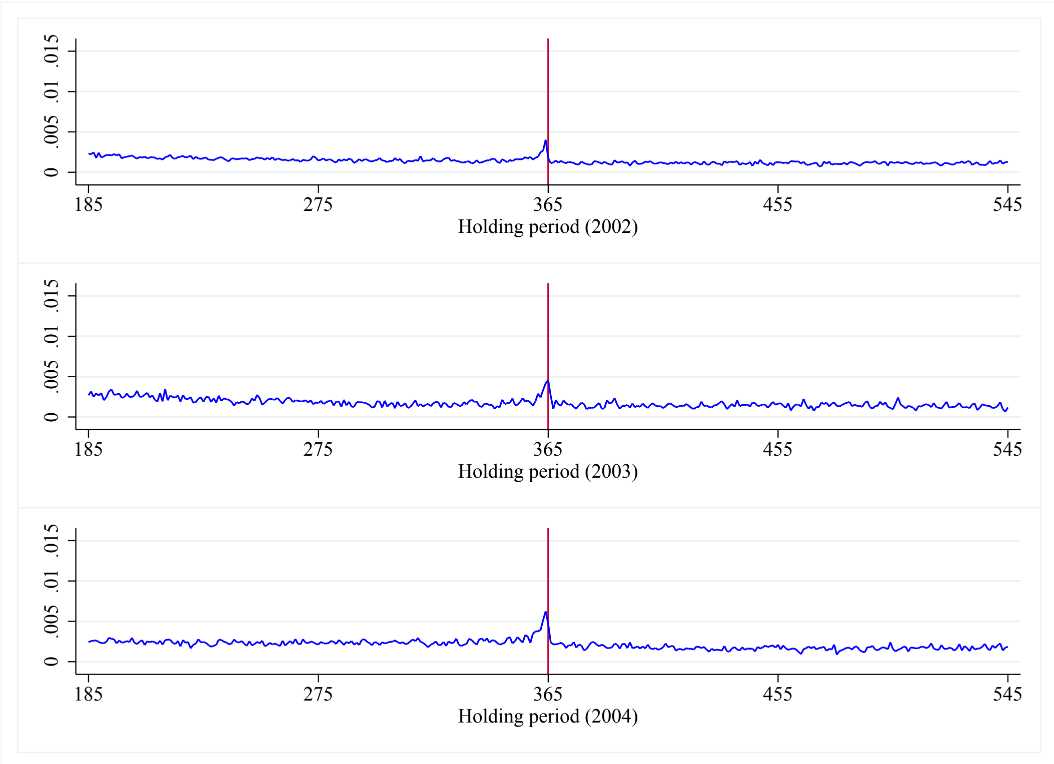

Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2002-2004. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 5.4. 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 31: Non-parametric regressions by year: Losses, pre years 2005-2007

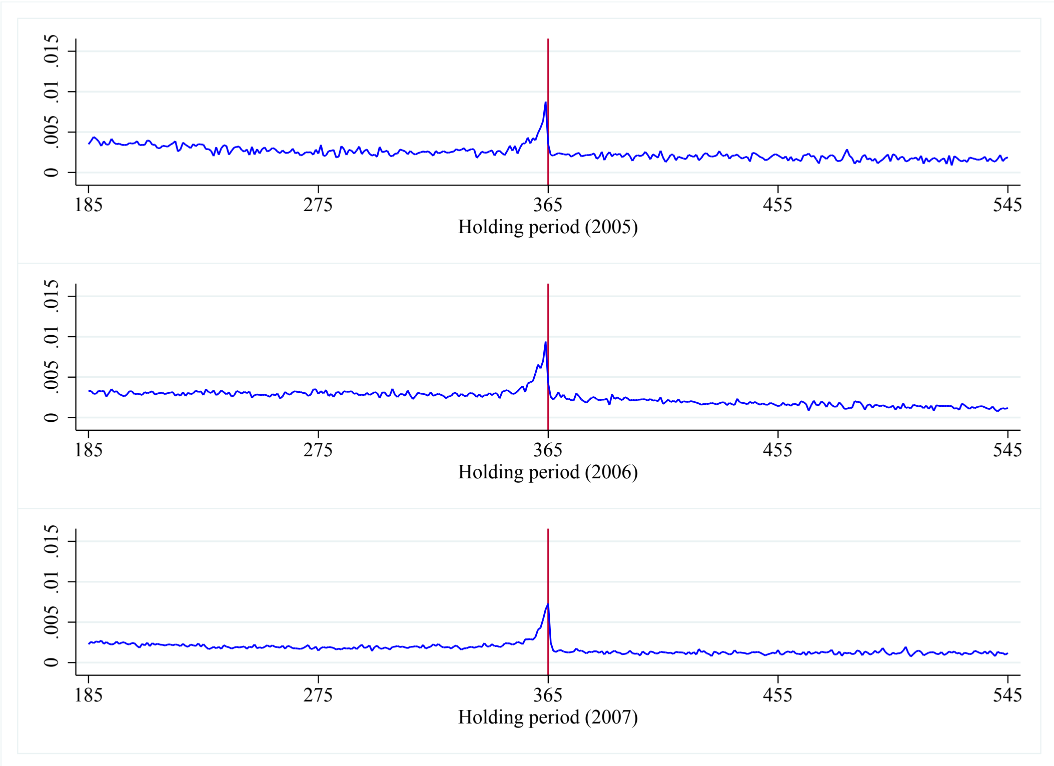

Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2005-2007. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients 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 5.4. 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 32: Non-parametric regressions by year: Losses, post years 2009-2011

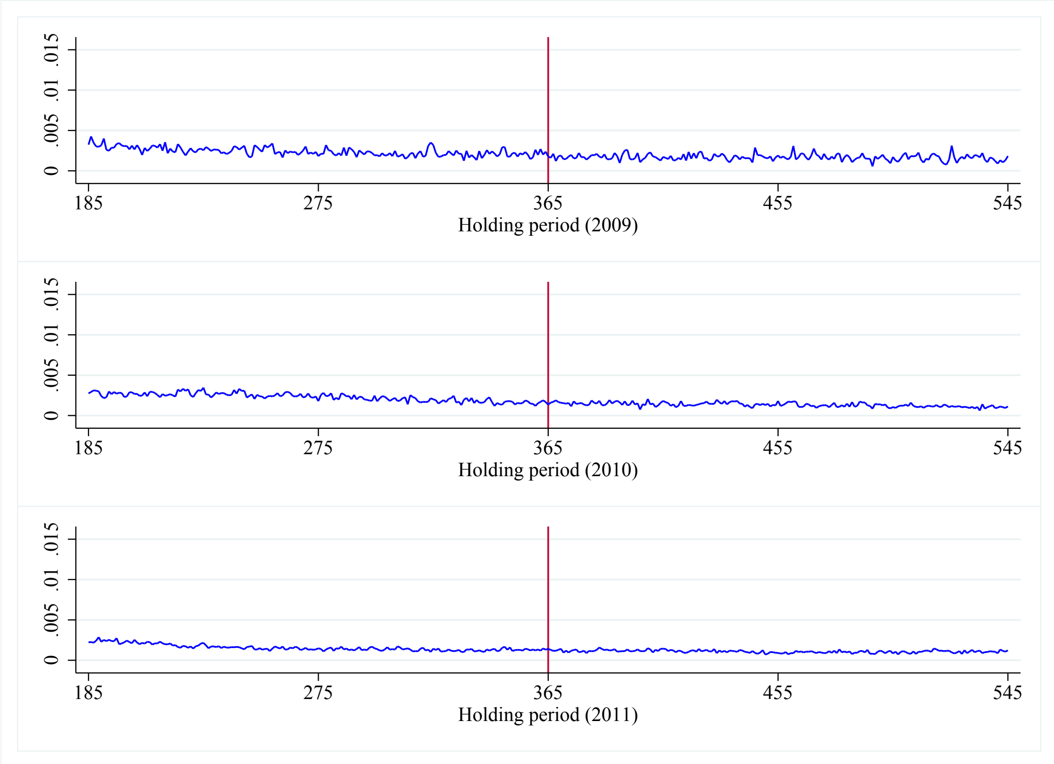

Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2009-2011. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 5.4. 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 33: Non-parametric regressions by year: Losses, post years 2012-2014

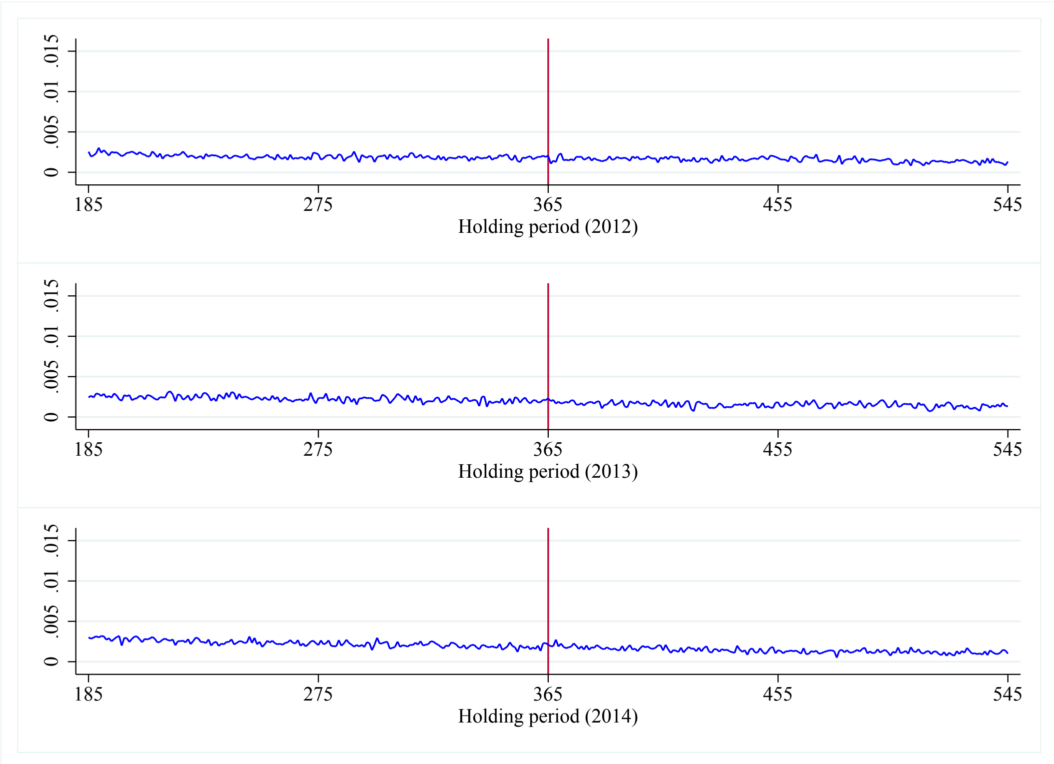

Notes: This figure displays the non-parametric regression estimates for each day of the holding period separately for the years 2012-2014. Included are share packages with prices below the purchase price. Coefficients indicate the probability that a share-package is sold on this holding-period day. Coefficients and on the investor level clustered standard errors stem from a series of regressions of the form $Sell_{ijd} = \beta_0 + \varepsilon_{ijd}$ if $\mathbb{1}(Loss_{ijd}) = 1$. Regressions are estimated for each day of the holding period. All estimation details are described in section 5.4. 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.

Download ZEW Discussion Papers from our ftp server:

http://ftp.zew.de/pub/zew-docs/dp/

or see:

https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html
https://ideas.repec.org/s/zbw/zewdip.html

//
IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de

Discussion Papers are intended to make results of ZEW research promptly available to other economists in order to encourage discussion and suggestions for revisions. The authors are solely responsible for the contents which do not necessarily represent the opinion of the ZEW.