

Department of Economics Discussion Paper Series

Estimating the Elasticity of Intertemporal Substitution using Dividend Tax News Shocks

Martin B. Holm, Rustam Jamilov, Marek Jasinski, Plamen Nenov

Number 1038
Originally published: March, 2024
Updated: July, 2024

Department of Economics
Manor Road Building, Manor Road
Oxford, OX1 3UQ

ISSN 1471-0498

Estimating the Elasticity of Intertemporal Substitution using Dividend Tax News Shocks[†]

Martin B. Holm Rustam Jamilov Marek Jasinski Plamen Nenov

July 2024

Abstract

This paper studies the spending response to news about a dividend tax reform to estimate the elasticity of intertemporal substitution (EIS). The Norwegian dividend tax reform was proposed in 2003, announced in 2004, and implemented in 2006, raising the dividend tax rate by 28 percentage points. We compare the spending responses of exposed households to a control group with no dividend income. Exposed households increased spending after the news and reduced spending after implementation. We show that this behavior is only consistent with an EIS above one. Using a capitalist-worker framework, we estimate the EIS to be around 1.6.

[†]Holm: University of Oslo, m.b.holm@econ.uio.no. Jamilov: All Souls College, Oxford, rustam.jamilov@all-souls.ox.ac.uk. Jasinski: Statistics Norway, marek.jasinski@ssb.no. Nenov: Norges Bank, plamen.nenov@norges-bank.no. We thank Annette Alstadsæter, Steve Bond, Philipp Dörrenberg (discussant), Andreas Fagereng, Andrea Ferrero, Niels Johannesen, Joseba Martinez, Michael McMahon, Gisle Natvik, Alan Olivi, Franck Portier, Hélène Rey, Federica Romei, Andreas Schaab, Alp Simsek, Stefanie Stantcheva, Paolo Surico, Ludwig Straub, John Vella and seminar participants at various venues for useful comments and suggestions. We want to thank Sondre Engebretsen for excellent research assistance. We also want to thank Janis Berzins and Øyvind Norli for providing us with the Oslo Børs data. This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 851891) and the Norwegian Research Council (grant agreement No. 326419). First version: March 2024.

1 Introduction

The elasticity of intertemporal substitution (EIS) is one of the key behavioral parameters in dynamic macroeconomics and finance. Since it governs the sensitivity of consumption growth to rate of return changes, it is central to the main workhorse equation in modern macroeconomics – the consumption Euler equation. Moreover, since its relation relative to unity governs the relative strength of substitution and income effects, it is also important in understanding the effect of rate of return changes or capital income taxation on the *level* of consumption and saving and hence on the capital accumulation process (Straub and Werning, 2020). Consequently, many recent theoretical and empirical contributions in macroeconomics and finance build on the assumption that the EIS is greater than one.[1]

Despite its central role, estimating the EIS has proven to be challenging, as we explain further in the related literature discussion below. Combined with the stark differences in theoretical and quantitative implications from having an EIS above or below 1, this has resulted in a lack of broad consensus and, in fact, a large debate in modern macroeconomics and finance as to the value of the EIS.

In this paper, we make progress towards ascertaining the value of the EIS relative to unity by leveraging a unique quasi-natural experiment. Specifically, we examine the 2006 Norwegian dividend tax reform and analyze it using the rich Norwegian registry data, allowing us to reliably impute household spending and saving. Our analysis robustly points to the EIS being greater than unity among the households who were exposed to the effects of the reform.

The Norwegian dividend tax reform is unique for several reasons. First, it was a major reform, with dividend income taxation for personal shareholders going up from 0 to 28%. Second, and most importantly for our purposes, there was a substantial delay from the time it was examined by a commission in 2003 and officially announced in 2004, to the time it was implemented in 2006. This allows us to study the spending and saving effects of *news* about a future capital tax change, and hence, to understand the anticipatory saving (or dis-saving) effect that this news induced. Third, the reform took place in a data-rich environment, particularly because wealth taxation was (and still is) present at the time, which means that we have access to highly reliable third-party reported data on household (and firm) balance sheets.

Why does the anticipatory (dis-)saving response to news about the dividend tax reform allow us to identify the value of the EIS relative to unity? When households receive news about a lower *future* rate of return on their savings, this lower future rate of return on

[1] See, e.g., Bansal and Yaron, 2004, Barro, 2009, Barro and Ursúa, 2012, Gabaix, 2012, Kaplan and Violante, 2014, Iachan et al., 2021, Barro and Liao, 2021, and Achdou et al., 2021.

saving induces offsetting income and substitution effects with the relative strength of each effect determined by the value of the EIS relative to unity. If the EIS is higher than one, then the household engages in anticipatory dis-saving and vice versa for an EIS lower than one. This approach to signing the EIS relative to unity due to news about future rates of return was suggested in recent theoretical work by Flynn et al. (2023). In our paper we build on this insight by leveraging the unique Norwegian dividend tax reform, which provided news about future rates-of-return changes.

Using a dynamic difference-in-differences methodology, we identify the spending effects of dividend tax news and the subsequent dividend tax implementation. Specifically, we compare the spending trajectories of a treatment group, defined as prime-age private business owning households with a relatively large share of dividend income out of gross income before the reform, relative to a control group of private business owners with no dividend income. Our identifying assumption is that the spending of the two groups would have evolved similarly in the absence of the dividend tax reform. Our baseline specification addresses a number of potential threats to identification stemming from systematic differences in age, exposure to sectors of the economy, and municipality-level economic shocks, as well as to other aspects of the tax reform.

We find that exposed households responded to the reform by increasing spending after the news and reducing spending after implementation. In terms of magnitudes, relative to the control group, households in the treated group increased their spending by around 6% more in 2004, followed by a persistent decline of around 5%. Therefore, empirically we find evidence for anticipatory dis-saving by the exposed households in response to news about future dividend tax reforms.

To map our empirical findings of a strong anticipatory dis-saving response to tax news to the value of the EIS, we construct a structural model and calibrate it to match the estimated dynamic spending response. Our model builds on a standard two-agent capitalist-worker framework with the addition of tax news shocks. Capitalists own firms and are affected by the dividend tax reform. Workers supply labor and are not affected by the tax reform because they do not receive any dividend income. Firms pay after-tax dividends. Importantly, we allow for potentially limited pass-through of the dividend tax on after-tax dividends via a parameter that governs (in equilibrium) the elasticity of the rate of return on saving for capitalists to dividend tax rate changes. Furthermore, we also allow for short-run heterogeneity in tax incidence through a parameter that, in reduced form, governs the duration of tax avoidance. Our calibration procedure therefore allows us to identify the EIS *conditional* on the degree of tax incidence in the short run and the long run.

We calibrate the model to the details of the dividend tax reform, leaving the EIS, the pass-through, and the short-run tax incidence parameters free. We then jointly calibrate these parameters by targeting the estimated average relative spending response over the 2003-2012 period. Our impulse-response matching exercise returns a value for the EIS of 1.59. In order to account for parameter uncertainty, we compute confidence bands via bootstrapping and find that the baseline model-implied EIS is statistically different from unity at the 95% confidence level. The result of the EIS being greater than unity is confirmed in a number of robustness tests and sensitivity checks, such as changing the structure of the labor market, the share of workers in the economy, or patience of the capitalists.

We argue that our deterministic model represents a simple yet realistic framework for interpreting our empirical findings. Specifically, the private business-owning households in the data, which we equate with the capitalists in the model, face large *returns* risk (see, e.g., Fagereng et al., 2020) but limited *labor income* risk. The presence of idiosyncratic business risk is irrelevant for the identification of the value of the EIS relative to unity via inspection of the anticipatory (dis-)saving response to news (Flynn et al., 2023), and we therefore abstract from it in our model. Labor income risk, on the other hand, could impact the identification of the sign of the EIS relative to unity because it affects the relative strength of income/wealth and substitution effects (Farhi et al., 2022; Holm, 2023). However, we argue that labor income risk does not substantively affect our results because the households in our sample are relatively wealthy, and we make the standard and realistic assumption that capitalists' behavior is consistent with decreasing absolute risk aversion.

An important conceptual issue to clarify is the following: Whose EIS does our approach uncover? Generally, the answer depends on both the *agents* whose intertemporal behavior is impacted and on the specific *policy change*. A significant body of work documents heterogeneity in the EIS across the population (Guvenen, 2006). The consensus seems to be that non-firm-owners generally have a very low EIS, close to zero. On the other hand, firm owners have a high EIS.[2] Since our empirical methodology recovers the average treatment effect on the treated (ATT) from the reform, the EIS we back out is thus representative of the average effect of the reform among private business owners. Jakobsen et al. (2020) recover a large EIS for the wealthiest segments of the population impacted by the 1989 Danish wealth tax reform. In both our and their studies, the treated agents are wealthy, and policies are related to distortions in capital and wealth accumulation. On the other

[2]One theoretical mechanism – that leverages non-homothetic preferences – for this measured heterogeneity was first provided in Browning and Crossley (2000) with a more recent treatment in Andreolli and Surico (2021). Consider a model with consumer goods that differ in income elasticities. The consumption share of luxury goods increases with agent wealth, and since luxuries are easier to delay into the future than necessary goods and services, the EIS can increase with wealth.

hand, Best et al. (2020) estimate a small EIS conditional on the sample of mortgagors and notched interest schedules in the U.K. mortgage market. In contrast to our paper and to Jakobsen et al. (2020), Best et al. (2020)'s treated agents (re-financing mortgagors) are most likely not in the top quintile of the wealth distribution, while the "policy" in question does not immediately impact capital accumulation decisions. Thus, one may conclude that the EIS is heterogeneous across the population and increasing in wealth, and the relevant EIS is policy-dependent because the treated group varies across different policies. Indeed, the average EIS we uncover is the EIS of private business owners and will be relevant for understanding the implications of, e.g., capital tax policies.

Related literature. Our paper contributes to three strands of the literature. First, we contribute to a large literature in macroeconomics that estimates the EIS.[3] Despite being a key behavioral parameter in modern macroeconomics and finance, there is no broad consensus in the literature as to its value, with estimates ranging from 0 to greater than 2. Moreover, estimates vary based on the data and empirical methodology used (i.e., use of aggregate time-series, disaggregated cross-sectional, or panel data, and reduced-form or structural approaches, the underlying structural assumptions on preferences, the estimation method, etc.). In addition to population heterogeneity in the EIS, at least part of this variation in estimates is due to endogeneity issues and biases arising from using aggregate time-series variation in rates of return or from misspecification of the structural equations used.[4]

Our quasi-experimental approach, therefore, brings our paper closer to the more recent advances in the literature. For example, similar to Gruber (2013) we also use cross-sectional variation for identification. However, rather than variation in tax rates, we use exposure to dividend income combined with an arguably unanticipated news shock about dividend income taxation. More recently, Jakobsen et al. (2020) combine administrative wealth data from Denmark together with a sizable tax reform – the 1989 Danish wealth tax reform – and a difference-in-differences methodology to show that wealth taxes have a large effect

[3]See, e.g., Hall (1988), Hansen and Singleton (1983), Campbell and Mankiw (1989), Mankiw and Zeldes (1991), Attanasio and Weber (1993), Blundell et al. (1994), Attanasio and Browning (1995), Beaudry and Van Wincoop (1996), Vissing-Jørgensen (2002), Vissing-Jørgensen and Attanasio (2003), and more recently Gruber (2013) Cashin and Unayama (2016), Best et al. (2020), Ring (2020), Calvet et al. (2021), Crump et al. (2022) among others. Campbell (2003), Hansen et al. (2007), Attanasio and Weber (2010), Havránek et al. (2015), and Havránek (2015) are examples of excellent survey and meta-studies, as well as a detailed treatment of how the estimates depend on the chosen method and the impact of publication bias.

[4]See the discussions in Bansal et al. (2010), Bansal et al. (2011), Gruber (2013), and Schmidt and Toda (2015) on the issues of downward biases arising in estimating EIS using aggregate time-series variation in rates of return. Also see Yogo (2004) for a discussion of the use of weak instruments in the estimation of EIS using time-series variation.

on wealth accumulation, suggesting an important behavioral response to the wealth tax change. Similar to us, they use a calibrated consumption-saving model to back out the implied EIS that is consistent with the reduced-form estimates and obtain relatively large values of around 2 to 4. Our contribution is to analyze a different type of reform, a dividend tax rather than a wealth tax reform. Moreover, we examine both the effects of the announcement and implementation, specifically emphasizing the importance of *news* of future tax changes for identification.

The large implied values for the EIS we find are in stark contrast with other recent empirical studies that use rich administrative data. For example, Best et al. (2020) use interest rate variation due to bunching around loan-to-value mortgage thresholds in the UK combined with a dynamic model of mortgage refinancing choice to back out an EIS of around 0.1. Similarly, Ring (2020) uses Norwegian administrative data but instead focuses on border discontinuities in the pricing of housing used for assessing the wealth tax in Norway and also finds a low value of the EIS. As mentioned before, one reason for having such large differences in the implied values is that much like in Jakobsen et al. (2020), the EIS we back out is for relatively rich households who might have a higher EIS, but also for whom the reform is more relevant because dividends make up a substantial part of their income. Second, the dividend tax reform we use is very salient. Therefore, any optimization frictions that may lead to a relatively low estimate of the EIS are absent in our empirical setting.[5]

We also contribute to a large and growing literature on the effects of capital income taxation, particularly dividend income taxation.[6] This literature, however, is primarily focused on the effects of capital income taxation on investments (e.g., Hall and Jorgenson, 1967). Relative to this literature, we study a complementary effect of dividend income taxation: the spending response of capital income recipients. Additionally, our theoretical model and the importance of the value of the EIS for understanding our empirical results paints a more nuanced picture regarding the investment response to dividend tax reforms. This is particularly relevant for settings where it is hard to decouple consumption from investment decisions, as is the case for closely-held private firms, where the consumption smoothing decisions of firm owners may impact the investment response to the dividend tax. In that case, depending on the value of the EIS, investments may either fall (as would be the case if EIS is relatively large) or *increase* (as would be the case if EIS is relatively small)

[5]See Iachan et al. (2021) for a further discussion of this point.

[6]See, e.g., Harberger (1962), Hall and Jorgenson (1967), Feldstein (1970), Auerbach (1979), Bradford (1981), Poterba and Summers (1983), Cummins et al. (1994), Chetty and Saez (2005), Auerbach and Hassett (2007), House and Shapiro (2008), Chetty and Saez (2010), Yagan (2015), Alstadsæter et al. (2017), Zwick and Mahon (2017), Barro and Furman (2018), Straub and Werning (2020), Boissel and Matray (2022), Furno (2022), Chodorow-Reich et al. (2024).

in response to the dividend tax (Straub and Werning, 2020). Additionally, and related to the insights of Korinek and Stiglitz (2009), we emphasize the importance of the anticipatory aspect of the reform for post-reform spending behavior due to inter-temporal tax arbitrage. Finally, our calibration procedure allows us to also back out the implied pass-through of the dividend tax reform, which is usually assumed to be 1 but turns out to be low in our setting.[7]

Finally, our structural model builds on two distinct literatures. First, our framework features limited asset market participation (Mankiw and Zeldes, 1991). Similarly to the canonical two-agent New Keynesian (TANK) literature, we assume that only an exogenous fraction of households receives dividend income and is thus exposed to the dividend tax reform in the model (Campbell and Mankiw, 1989; Galí et al., 2007; Bilbiie, 2008). Second, corporate dividend payments to capitalists are taxed; in particular, we allow for *news shocks* about future dividend tax rates. In this regard, we are leaning on the vast literature on news-driven business and financial fluctuations (Beaudry and Portier, 2004, 2006).[8]

2 Illustrative model

To illustrate how the spending response to news about future capital taxes depends on the magnitude of the elasticity of intertemporal substitution (EIS), we present a simple framework motivated by Flynn et al. (2023). Consider an agent living for three periods, solving the following problem:

$$\max_{C_0,C_1,C_2} u(C_0) + u(C_1) + u(C_2)$$

$$s.t.\ \ C_0 + \frac{C_1}{R_1} + \frac{C_2}{R_1 R_2} = R_0 A_0.$$

where A_0 is initial financial wealth and R_t is the period t (after-tax) gross return on savings. For simplicity, we assume the agent has no labor income and there is no discounting between periods. Because the portfolio allocation decision of the agent is separable from the consumption-saving decision in the absence of labor income, we also abstract from the portfolio choice (see, e.g., Iachan et al., 2021). Assume further that period utility is $u(C) = \frac{C^{1-1/\psi}-1}{1-1/\psi}$ ($\ln(C)$ if $\psi = 1$) where $\psi > 0$ denotes the EIS.

[7] In addition, the stock price response that we document in Section 3 is consistent with the findings of Poterba and Summers (1984) for capital tax reforms in the UK and with Auerbach and Hassett (2007) and Chetty et al. (2007) for the case of the 2003 U.S. dividend tax reform.

[8] See also, among many others, Lorenzoni (2009); Beaudry and Lucke (2010); Barsky and Sims (2011); Blanchard et al. (2013); Beaudry and Portier (2014).

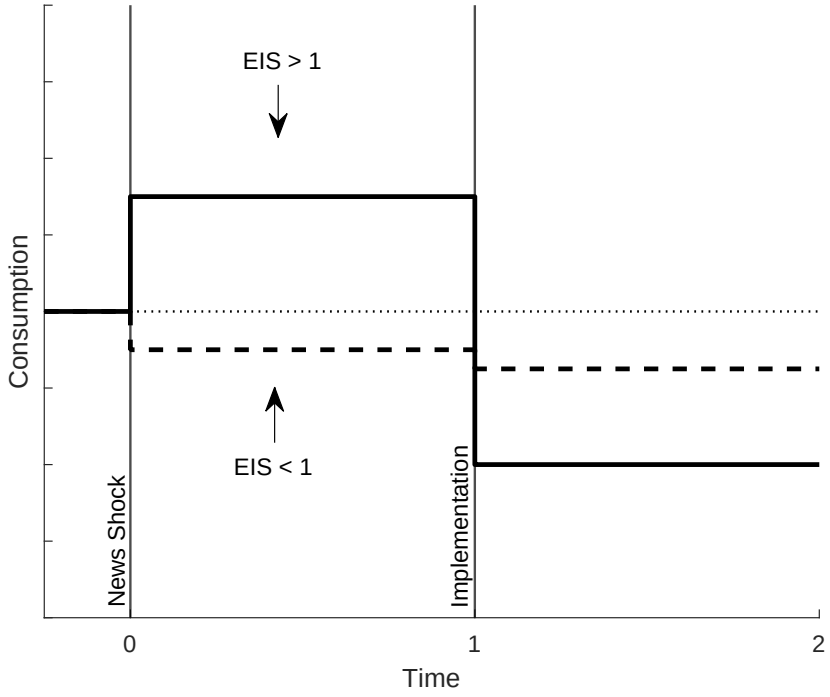


Figure 1: Example consumption paths in response to a lower R_2.

The solution to this problem has to satisfy the Euler equations: $u'(C_0) = R_1 u'(C_1)$ and $u'(C_1) = R_2 u'(C_2)$. Using the definition of the period utility function, the budget constraint, and the two Euler equations, the derivatives of the consumption policy functions with respect to R_2 are

$$\frac{\partial C_0}{\partial R_2} = -(\psi - 1)\kappa_0(R_0 A_0, R_1, R_2), \quad \frac{\partial C_1}{\partial R_2} = -(\psi - 1)\kappa_1(R_0 A_0, R_1, R_2), \quad \frac{\partial C_2}{\partial R_2} = \kappa_2(R_0 A_0, R_1, R_2),$$

where κ_0, κ_1, and κ_2 are functions of $R_0 A_0$, R_1, and R_2 (Appendix A includes the details).

Hence, C_0 and C_1 are strictly decreasing in R_2 iff $\psi > 1$, while C_2 is always strictly increasing in R_2. Intuitively, a higher $t = 2$ return induces $t = 0$ income and substitution effects. As long as the EIS is greater than one, the substitution effect dominates the income effect and the agent decreases her $t = 0$ consumption and increases her $t = 0$ savings. Conversely, a lower R_2, for example, due to an increase in the future capital tax rate, leads to an increase in $t = 0$ consumption if the EIS is greater than one.

Figure 1 illustrates consumption responses to a lower R_2 (the dividend tax reform we study) for different values of the EIS relative to unity. The main idea of this paper is that the *sign* of the consumption response between the news shock and implementation of the reform can allow us to identify whether the EIS is greater or less than one. If spending increases in response to news about a dividend tax increase, the EIS is greater than one.

Conversely, if spending decreases in response to news about a dividend tax increase, the EIS is less than one.

2.1 Robustness

While the model we present here is stylized, the sign comparative statics also hold in richer models with multiple assets, returns risk, and labor income (Flynn et al., 2023). Here we briefly discuss a few notable extensions.

More general preferences. The sign comparative statics hold for more general preferences (e.g., Epstein-Zin) as shown by (Flynn et al., 2023). They also hold for particular non-homothetic preferences used in the macro literature. In Appendix B, we specifically show that the comparative static result holds with non-homothetic preferences as in Straub (2019). In that environment, however, the EIS does not depend on a single parameter but on a combination of parameters and on the level of spending, and may, therefore, change depending on changes in that level.

Idiosyncratic business risk. In the data, private business-owning households face large *returns* risk (see, e.g., Fagereng et al., 2020) which is not accounted for in our simple framework above. However, as argued by Flynn et al. (2023), idiosyncratic business risk does not pose a challenge for the identification of the value of the EIS relative to unity via inspection of the anticipatory (dis-)saving response to tax news (see Section 4.3 and Proposition 5 in Flynn et al., 2023). Intuitively, business risk impacts the risk-adjusted return on wealth but does not change the intertemporal trade-off of the agent in response to the tax news.

Idiosyncratic labor income risk. Unlike business risk, labor income risk could impact the identification of the sign of the EIS relative to unity because it affects the relative strength of income/wealth and substitution effects (Farhi et al., 2022; Holm, 2023). In particular, Holm (2023) shows that the substitution effect to interest rate changes weakens in the presence of income risk if the utility function has the property *temperance*.[9] However, under the standard and realistic assumption that capitalists' behavior is consistent with decreasing absolute risk aversion (DARA), we argue that labor income risk does not substantively affect our results because the households in our sample are relatively wealthy.[10]

[9] Temperance is defined as a negative fourth derivative of the utility function, $u'''' < 0$, see the Theorem in Eeckhoudt and Schlesinger, 2006.

[10] For example, Holm (2023, Corollary 1) shows in a two-period setting that the marginal effect of income risk on the substitution effect is $\frac{1}{8}(1+1/\psi)(2+1/\psi)\frac{1}{C}$ with the power utility function we use here, converging

Revaluation effects on human wealth. With labor income, there may be revaluation effects of human capital and thus countervailing wealth effects. For example, if the lower future portfolio return also induces a positive revaluation of human capital wealth, consumption might increase at $t = 0$ even when the EIS is less than unity. While such wealth effects may be important in some settings, we view them as more limited in the context of a *dividend tax reform* since that reform affected capital and capital income rather than human capital wealth. Nevertheless, our structural model in Section 6 would allow us to discipline any offsetting revaluation effects via labor income.

3 Institutional Setting

This section describes the institutional settings of the dividend tax reform. We first describe the reform in detail before illustrating its impact on the stock market and aggregate savings.

The dividend tax reform. Before 2006, labor and capital income in Norway were taxed at very different rates. Wage earners faced a progressive income tax, with marginal tax rates from around 35% to 64.7%.[11] Capital owners, on the other hand, faced a flat profit tax of 28%.[12] The large difference in marginal tax rates between wage and capital income for high-income individuals was a concern for policymakers because it incentivized inefficient income shifting between wage and capital income. For example, the government introduced a temporary dividend tax of 11% in 2001 (Innstilling til Odelstinget nr. 23, 2000-2001) partly motivated by this marginal tax rate difference, which was subsequently reversed and thus in place only for the fiscal year 2001. On January 11, 2002, the newly-elected government appointed an expert commission to suggest permanent changes in the tax system, specifically motivated by the large difference in marginal tax rates between labor and capital income (NOU 2003:9, 2003, p. 11).

On February 6, 2003, the government-appointed commission published their findings (NOU 2003:9, 2003). Among the key recommendations was the implementation of a 28% dividend tax, raising the effective tax rate on firm owners from 28% to 48.2%. At the same time, the commission recommended reducing the top marginal tax rates on wage income from 64.7% to 54.3%. On March 26, 2004, the government officially announced a tax reform that mostly followed the commission's recommendations, introducing a dividend tax and reducing the marginal tax rate on wage income (Stortingsmelding nr. 29, 2003-2004).

to zero as wealth increases.

[11] This marginal tax rate included the employer's national insurance contributions ("arbeidsgiveravgift").

[12] Additionally, there was a wealth tax with a marginal tax rate of 1.1% of net (taxable) wealth above NOK 580,000 during this period.

However, for administrative reasons, the introduction of the dividend tax was postponed to January 1, 2006.

Henceforth, we will be referring to the interval around the publication of the report – February 2003 – and the official announcement of the tax reform – March 2004 – as the "news shock" period. Additionally, we will be referring to January 2006 as the reform "implementation" date. Finally, the interval between 2004 and 2006 is labeled and referred to as the "transition" period.

The reform – once it was officially implemented in 2006 – introduced a 28% personal tax on dividends and capital gains in excess of a threshold amount based on the riskless returns in any given year.[13] Under the previous tax regime, dividends were tax-exempt for any shareholder, while capital gains were almost always applied to a zero base and hence were tax-exempt as well. Firms paid no taxes on dividends and capital gains both before and after the reform.[14] The reform also decreased the top marginal tax on labor income from 64.7% to 54.3%, while the sum of taxes paid by the firm and the investor on dividends and capital gains increased from 28% to 48.2%, following the recommendations of the commission.[15]

Stock market impact of the reform announcement. The timeline above reveals that the extent to which the tax reform was anticipated is not immediately obvious. In this section, we argue that the 2003 news shock was unanticipated by market participants. To illustrate this point, we examine the behavior of the stock market from 2001 to 2008. Specifically, we compute and track cumulative returns of high and low-dividend-paying stock portfolios using stock-level data from the Oslo Stock Exchange.

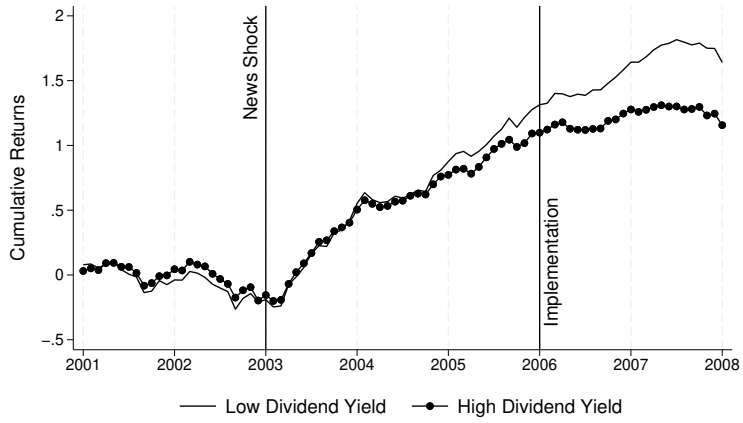
Figure 2 plots cumulative returns of two equal-weighted portfolios, one with above-median (high) and the other with below-median (low) dividend yield stocks, over the period January 2001 - January 2008.[16] Dividend yields are computed prior to the base month. As the figure shows, both high- and low-dividend portfolios behaved similarly prior to the news shock period. It was only in 2004 that the returns of the two portfolios

[13]The annual risk-free rate of return allowance for shareholders/partners (RRA) is computed as the exemption rate multiplied by the sum of the cost price of the share/holding and any unused allowance from previous years. The unused allowance is then carried over to the next year with interest and can be deducted from future dividends and capital gains associated with the same share/holding. The exemption rate for shareholders and partners is the average interest rate on three-month Norwegian Treasury bills in the year for which the allowance is to be calculated. Therefore, the dividend payouts of firms with dividend yields lower than the average yield on 3-month Norwegian Treasury bills fall within the RRA exemption.

[14]During the transition period, stocks could be transferred to a holding company without triggering a capital gains tax.

[15]See Appendix C for further details on the tax reform.

[16]See Appendix D for details on portfolio construction.



Notes: This figure shows the cumulative returns of two portfolios consisting of high- (above median) and low-dividend yielding stocks, all relative to January 2003. Appendix D presents details on how the portfolios are constructed.

Figure 2: Cumulative stock returns for high vs. low dividend stock portfolios.

began to diverge.[17] Divergence accelerated noticeably during the transition period and after the implementation date. This finding suggests that the market did not anticipate the dividend tax reform prior to the news shock.

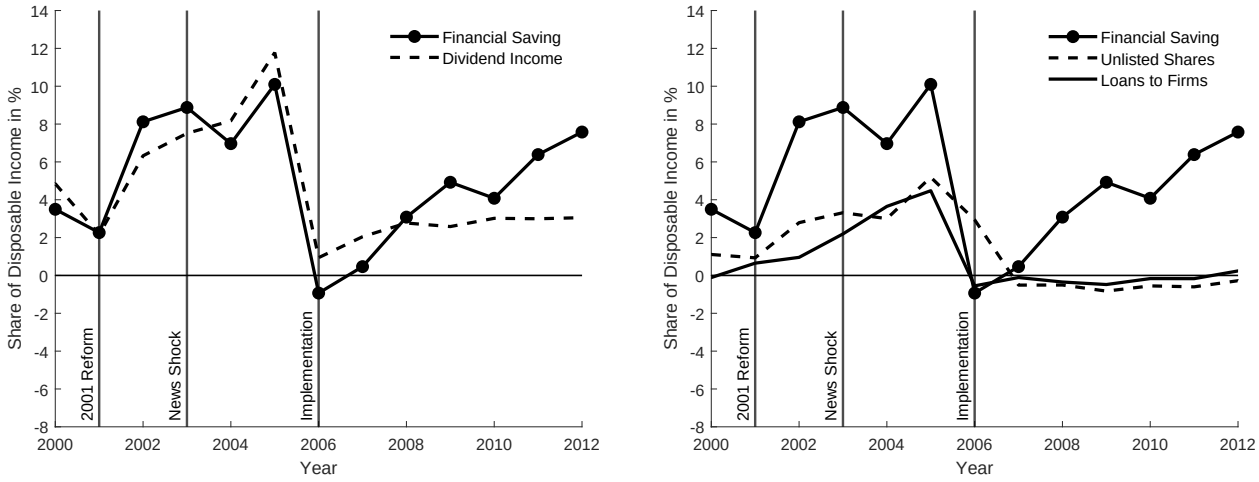
Aggregate impact. To illustrate aggregate implications of the dividend tax reform, Figure 3a displays total dividend income and savings as shares of households' disposable income. There are two notable observations. First, the dividend income of households responded to the tax rate change suggesting that the reform was salient. For example, the temporary dividend tax hike of 2001 reduced dividend income by almost 50% compared to 2000. Moreover, around the implementation of the permanent tax reform, dividend income as a share of disposable income dropped from more than 10% in 2005 to around 1% in 2006. Second, household saving to a large extent follow dividend income prior to 2006. Dividend income was elevated from 2002 to 2005, and most of this income was saved. Hence, while the dividend tax reform affected aggregate dividend income, it is not clear from the aggregate data whether consumption also increased.

A substantial fraction of household saving in the years before the tax reform was due to various methods of intertemporal shifting of dividends to avoid taxes. Figure 3b shows that in the few years prior to 2006, saving in unlisted shares[18] and loans to firms[19]

[17] Appendix D presents several robustness tests for the stock market exercise. Those show that the divergence happened already in 2003, i.e. earlier than what Figure 2 suggests.

[18] We measure saving in unlisted shares as the transactions of unlisted shares in the non-financial sector held by the household sector in the financial accounts.

[19] We measure saving in loans to firms as transactions of loans to the non-financial sector owned by



(a) Saving and dividend income. **(b)** Saving in unlisted shares and loans to firms.

Notes: Figure (a) shows households' financial saving and dividend income as a share of disposable income. Figure (b) shows households' saving in unlisted shares and loans to non-financial firms as a share of disposable income. All numbers are from the national accounts.

Figure 3: Saving and dividend income as a share of disposable income.

increased. Both saving in unlisted shares and corporate loans are examples of ways to shift the dividend tax burden across time.[20] The first method to shift is to pay out dividends before the reform and transfer these back to the firm in the form of paid-in capital. The second method is to pay out dividends before the reform and lend these back to the firm. Because withdrawals of paid-in capital or repayment of debt are exempt from dividend taxation, firm owners paid no taxes when taking out dividends prior to the reform and were also able to extract resources from the firm without paying taxes after the reform. We will return to the issue of intertemporal shifting of tax burdens when we impute the spending of firm owners and when we interpret our main results.

4 Data and Imputed Spending

This section presents the data and describes how we impute spending, specifically focusing on imputed spending among private business owners.

Data sources. We use data from several Norwegian administrative registries from 2000 to 2012, combined using unique personal identification numbers. Because Norway levies both a wealth tax and an income tax, the tax authorities collect information on household

households in the financial accounts.

[20] Alstadsæter and Fjærli (2009) document a similar pattern by looking at firm balance sheet data.

balance sheets and income statements. Most variables in the income and wealth data are third-party reported by employers, financial intermediaries, or the tax authorities (e.g., assessed housing wealth and private business wealth), the main exception being ownership of foreign wealth. In addition, we use the housing transaction data when we impute spending, the stock ownership data to define private business owners, income statements and balance sheets of firms to measure dividends and changes in paid-in capital, information on family status to construct households, birth year, and home addresses, and information on sector of employment from the employer-employee register. We deflate all values to real 2011 U.S. dollars.

Imputed spending. The main variable of interest in our study is *imputed spending*. defined from the budget constraint as income not saved. Because the wealth tax is levied at the household level and imputed spending requires measuring saving, imputed spending is defined at the household level. In the analysis below, however, the unit of observation is the individual, where all variables of an individual in a multi-person household are defined as the average value of that variable for the household. Moreover, we note that imputed spending is all spending on items that do not enter the balance sheet, including durable components such as housing refurbishment and household appliances, but also non-durable spending. Appendix E describes how we impute spending for households that do not own private businesses.

Imputed spending of private business owners. The main focus of this paper is the spending response of private business owners (owners of incorporated firms not listed on the public stock exchange) to the dividend tax reform. We focus on private business owners holding significant positions in the private business (above 25%). In principle, all wealth within the private business should be counted as part of household wealth in accordance with the household's ownership share. Similarly, any saving and income in the private business are part of household saving and income. Importantly, in this way of accounting, resource flows from a firm to its owner, such as dividends, are not part of household income but instead represent a movement of resources between different bank accounts – the private account and the firm account.

Concretely, we define imputed spending of private business owners as

$$\text{spending}_{i,t} = \text{spending}^{\text{npbo}}_{i,t} + \underbrace{\text{profits}_{i,t}}_{\text{income within the firm}} - \overbrace{(\Delta \text{ book value}_{i,t} - \text{capital gains}_{i,t})}^{\text{saving within the firm}} \tag{1}$$

where spending$^{\text{npbo}}$ is imputed spending for individuals that do not own a private business, defined in Appendix E. In equation (1), profits and the change in book value are observed from the firm's income and balance sheet statements. The remaining task is therefore to impute capital gains within the firm. Imputing capital gains within firms, though, is challenging for two reasons. First, firms' balance sheets include more asset classes than households, for example "plant and machinery" and "ships, rigs, aircraft." Second, the asset categories in the accounting rules for firms are less informative about what the firm owns. For example, "land, buildings and other real property" is a balance sheet category, still, we do not observe whether the wealth in that category consists of land, housing, or commercial real estate. Finding relevant prices for each asset class is challenging and using equation (1) to impute spending is therefore almost infeasible.

Instead, we follow an alternative approach by noting that

$$\underbrace{\text{profits}_{i,t}}_{\text{dividends}_{i,t}+\text{retained earnings}_{i,t}} - \overbrace{(\Delta \text{ book value}_{i,t} - \text{capital gains}_{i,t})}^{\text{retained earnings}_{i,t}+\Delta\text{paid-in capital}_{i,t}} = \text{dividends}_{i,t} - \Delta\text{paid-in capital}_{i,t}. \quad (2)$$

The insight from equation (2) is that we only need to account for the net flows between firm and owner to impute spending of private business owners. Our measure of imputed spending for private business owners is therefore

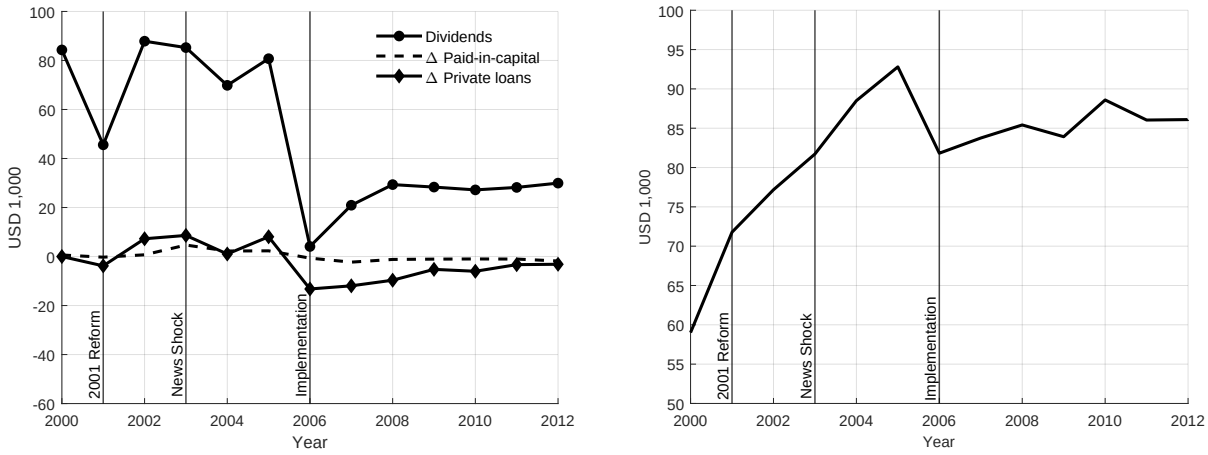
$$\text{spending}_{i,t} = \text{spending}^{\text{npbo}}_{i,t} + \text{dividends from the firm}_{i,t} - \Delta\text{paid-in capital in firm}_{i,t}. \quad (3)$$

where "dividend from the firm" and "Δ paid-in capital in firm" are the dividends flow from the firm to owner i and the change in the owner's paid-in capital in the firm, respectively.

Our imputation procedure makes one additional underlying assumption: we assume that the ownership share is stable from one year to the next. This is necessary for two reasons. First, it allows us to use the definition of spending above without any adjustments for changes in ownership shares. Second, it allows us to compute the owner's share of the change in paid-in capital because we observe the change in paid-in capital from the firm's balance sheet and allocate it to each owner according to their (stable) ownership shares. This assumption is relatively innocuous in our sample, where owners mostly retain a stable ownership share from one year to the next.[21]

To illustrate how our imputed spending approach works, consider the two common ways for owners to react to the announced dividend tax reform: (i) taking out \$100 in tax-

[21] In our sample, 86% of owners have the same ownership share five years later.



(a) Average dividends, paid-in-capital and private loans in the treatment group.

(b) Average spending in the treatment group.

Notes: Figure (a) shows the average dividends from private businesses to the owner, the change in paid-in capital in private businesses due to the owner, and the change in private loans for the owner in our treated sample. Figure (b) shows the average imputed spending in the treatment group.

Figure 4: Imputed spending, dividends, paid-in-capital and private loans

free dividends prior to the reform and transferring $100 back to the firm as paid-in capital, and (ii) taking out $100 in dividends and lending $100 back to the firm. In both cases, household spending should be unaffected. In case (i), we would see $100 in dividends and a $100 change in paid-in capital, which sums to $0. In case (ii), we would see $100 in dividends and a $100 increase in private loans to the firm (part of loans/deposits), also summing to $0. Hence, our approach to impute spending by firm owners accounts for the two common ways of avoiding the future dividend tax, which firm owners used in reaction to the dividend tax reform announcement.

Figure 4 further illustrates our approach by displaying average spending, dividends, change in paid-in capital, and change in private loans among private business owners in our sample (the treatment group, which we define in the subsequent section). First, note that while dividends fluctuate, these fluctuations do not directly translate into spending fluctuations. For example, while dividends dropped by almost $80,000 from 2005 to 2006, spending fell by much less because most of the dividends in 2005 were saved.

Further complications. There are two additional issues when imputing spending for private business owners. First, in many cases, a business owner owns several private businesses. In the ownership registry, we observe all owners of private businesses, both firms and households. The ownership registry thus allows us to compute indirect ownership shares of all firms. We compute indirect ownership shares for households up to layer 10

when we compute ownership shares.

Second, our sample period starts in 2000, but the ownership registry starts in 2004. We therefore impute ownership shares from 2000 to 2003. We restrict attention to significant owners in 2004 (greater than 25% ownership share). We then impute ownership shares in the prior years by assuming that the ownership share of household i in firm j is constant going back in time under the conditions that the firm existed in the firm registry and the owner has non-zero holdings of non-listed stocks in the tax accounts.

Tax evasion. Two additional issues related to tax evasion remain. First, owners of private businesses may use these firms for private consumption. Consuming within the firm is illegal, limiting its scope to how much one can get away with it without raising suspicion among the tax authorities. Nevertheless, it is likely to be prevalent, and the tax data does not allow us to infer the extent of this problem (Leite, 2023). We note that systematic tax evasion leads to level differences in imputed spending (if the bias is constant in logs within groups), which is not a problem for us. However, the tax reform potentially incentivizes owners to evade taxes to a larger extent, as suggested and documented in Alstadsæter et al. (2014). This tax evasion would primarily bias our result after the reform is implemented, not in the period between announcement and implementation, which, as we illustrate in Section 2, is the period when the consumption response to the tax reform varies with the sign of the EIS relative to unity.

Second, households may evade taxes by hiding wealth abroad, which is also illegal. Because our data relies on administrative data collected by the tax authority, hidden wealth is always an issue, especially among wealthy households (Alstadsæter et al., 2019). A concern is that the pattern we observe of heightened spending followed by a permanent decrease may be because owners hide wealth before the reform (resulting in a spending spike) and spend out of this hidden wealth afterward (resulting in a permanent reduction in spending because we do not observe this spending). However, in this case, we would observe a spending spike equal to the value of the hidden wealth and a permanent decline equal to the annuity value of the hidden wealth. Because the estimated spending response before implementation has roughly the same order of magnitude as the permanent decline post-implementation, our observed spending pattern is unlikely to be driven only by hidden wealth responses.

5 Empirical Setup and Results

We now estimate the effects of the dividend tax reform on spending using a difference-in-differences framework. The main goal is to compare the spending of individuals exposed to the reform, the *treated*, with those of an appropriate control group. In this section, we first define the treatment and control groups before exploring the extent to which there are systematic differences between the two groups in pre-determined characteristics. After that, we present the formal estimation framework, discuss threats to identification, and present the main results.

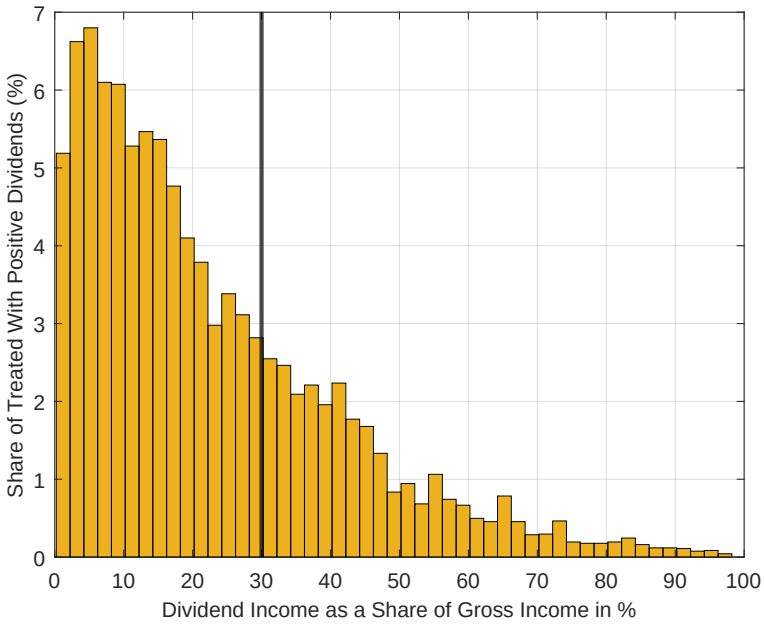
5.1 Treatment and control definitions

We define a household as treated if it satisfies two criteria. First, the mean share of its gross income in 2000 and 2002 - including labor income, transfers, and capital income - attributable to dividends from a private business firm must exceed 30%. We use the average dividend payout from the two years before the news shock when dividends were not subjected to the temporary dividend tax in 2001. We chose this method because dividends from private businesses are "lumpy"; their distribution to individual households can be irregular, with significant variations in both the amount and the timing of payments. Second, as mentioned previously, we restrict attention to households holding significant ownership shares (greater than 25%) in their private business ownership portfolio because these owners tend to retain relatively stable ownership shares. Figure 5 displays dividend income as a share of gross income among private business owners in our sample.[22]

The next step is to define a control group to serve as a counterfactual for the treatment group. We use private business owners whose ownership share exceeds 25% but who did not receive any dividend income from their private business ownership in 2000 and 2002. The identification thus comes from differential exposure to the general reform: Our treated owners received a significant fraction of their income from dividends before the reform, but the control group did not. In Figure 9 and Appendix F, we present an alternative setup in which the control group instead consists of wealthy individuals who do not own private businesses.

Additionally, we impose a few minor restrictions on our sample, which spans annual observations from 2000 to 2012. First, we concentrate on individuals between 25 and 65 years old in 2000. Second, we limit our analysis to households with disposable income

[22]We choose the treatment threshold to ensure that dividends should make up a significant fraction of treated households' income. In Section 6, we show that our results are robust to using different thresholds as well as to using a continuous treatment variable.



Notes: The figure shows the distribution of dividends from private businesses as a share of gross income among private business owners (more than 5% ownership shares) in 2000.

Figure 5: The distribution of dividend income as a share of gross income.

greater than the base amount in the Norwegian social security system (approximately $10,000). Third, we restrict attention to households whose log growth rate of imputed spending did not change by extreme amounts from one year to the next (top/bottom 1%). Our sample consists of 21,758 individuals in the control group and 5,516 in the treatment group in 2000.

Table 1 presents descriptive statistics in 2000 by treatment status in our sample. The descriptive statistics reveal that there are some differences between the two groups. The two groups of households have relatively similar labor income but differ in their disposable income, spending, and wealth. Moreover, treated households are somewhat older. Most of their non-housing wealth is held in private businesses, and dividend income from these businesses is an important part of their income, on average around 47% of gross income in 2000.

5.2 Empirical setup

Our empirical setup estimates the effects of the reform on spending by comparing the relative spending response of treated households, for whom the reform is important, with the control group for whom the reform is less relevant. To capture potential pre-announcement anticipation effects (see the discussion in Section 3), we set 2000 as the base

	Control		Treated	
	Mean	S.D.	Mean	S.D.
Panel A: Household Characteristics				
Age	45.84	9.86	48.34	9.25
Panel B: Spending and Income Statement				
Spending	41.04	85.06	59.01	133.28
Disposable income	45.31	59.77	84.73	96.21
Labor income	56.47	34.57	55.85	27.34
Transfers	5.54	8.65	4.79	8.44
Dividend income from private businesses	0.00	0.00	84.30	148.18
Taxes	19.86	18.99	25.82	27.46
Panel C: Balance Sheet				
Gross wealth	434.78	425.53	638.95	588.69
Housing wealth	377.17	371.88	507.82	492.77
Deposits	23.28	62.46	58.25	129.42
Public Stocks	3.50	29.11	7.15	46.22
Mutual Funds	5.05	22.76	10.91	40.29
Private Business Wealth	37.51	196.81	221.99	490.88
Net Wealth	342.00	400.63	570.79	575.89
Debt	92.78	102.67	68.15	114.03
Panel D: Shares				
Exposure to the reform (dividend share of gross income in %)	0.00	0.00	47.29	20.51
Number of individuals	21,758	.	5,516	.

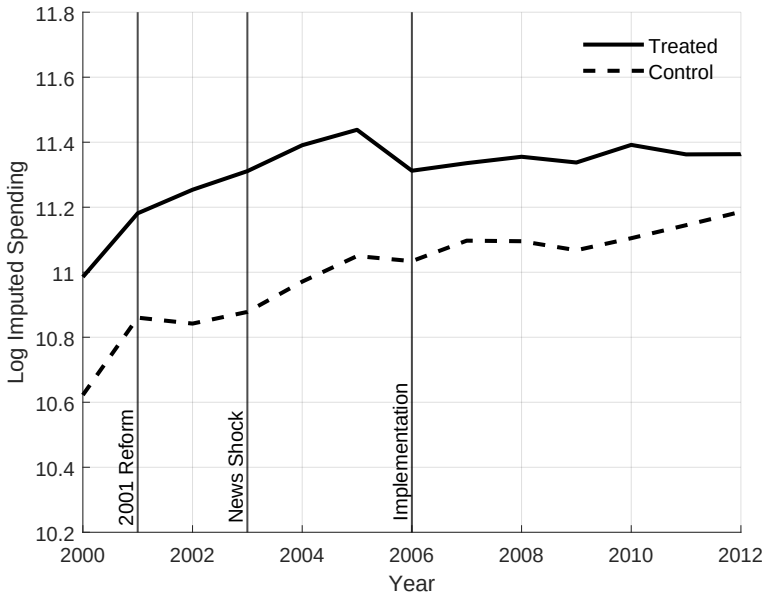
Notes: We define *treated* as having, on average, more than 30% of gross income in the form of dividend income from private businesses in 2000 and 2002. The control group consists of private business owners with no dividend income in 2000 and 2002. Values in Panel B and C are in 1,000 dollars in 2011.

Table 1: Descriptive statistics in 2000.

year and estimate the following dynamic difference-in-differences model

$$c_{i,2000+h} - c_{i,2000} = \alpha + \sum_{h=\underline{h}}^{H} \beta_h \left(D_{i,2000} \times \omega_{2000+h} \right) + \sum_{h=\underline{h}}^{H} \Gamma_h' \left(\mathbf{X}_{i,2000} \times \omega_{2000+h} \right) + \varepsilon_{i,h} \tag{4}$$

for $h = \{0, 1, 2, \ldots, 13\}$, where $c_{i,2000+h}$ denotes log imputed spending for indvidual i in year $2000 + h$, $D_{i,2000} \in \{0, 1\}$ is our treatment variable for household i in year 2000, ω_t denotes a dummy variable for year t, and $\mathbf{X}_{i,2000}$ contains a set of controls for household i in 2000. The controls in the benchmark specification include four groups of pre-reform non-financial income described below, pre-reform 2-digit NACE fixed effects for the firm where the household is employed, and pre-reform age and municipality dummies. Standard errors



Notes: The figure displays average log imputed spending in the treatment and control groups.

Figure 6: Log spending in the treatment and control group.

are clustered at the individual level to address autocorrelation in the error terms $\varepsilon_{i,h}$ within households.

Threats to identification. Our main identifying assumption is that absent the reform, the spending of households in the treated and the control groups would evolve similarly. Figure 6 provides the first look at the raw data used to identify the spending effect of the reform. In years before the reform, average log spending evolved relatively similarly in the treated and the control groups.

Although not necessary for our identifying assumption, one weakness of our setup is the lack of balancing between the two groups. The treated group consists of older individuals with higher disposable income, more spending, and more wealth. We therefore include a set of additional controls for several potential confounders in our baseline empirical specification.

The first potential confounder is the concurrent labor income tax reform. At the same time as the dividend tax reform, the government reduced the top marginal tax rates on wage income. A potential issue is that households in our treatment group are differentially exposed to this income tax reform. To address this concern, we include income controls (four bins), corresponding to four groups of differential changes in marginal income tax rates as in Thoresen et al. (2010, Figure 1).[23]

[23]Specifically, we choose the following three cutoffs in 2000 based on the 2006-NOK marginal tax groups:

Another potential issue in our analysis is that dividend income as a share of total income may vary with age, which in turn can correlate with differences in spending growth across households. It is well established that portfolio shares vary by age (see, e.g., Fagereng et al. (2017) for the case of Norway). Therefore, if more exposed households are older and older households tend to have different spending trajectories than younger households, that could lead to a systematic bias. We therefore include a full set of age controls in our empirical specification.

Another potential threat to identification is that households in the treatment and control groups may differ in their exposure to sectors of the economy and to different labor market shocks. To address this concern we control for the pre-reform 2-digit NACE code of the primary employment of the household and municipality-fixed effects.

In several robustness exercises, we address other potential issues. Specifically, since higher dividend income correlates with stock wealth and stock prices were increasing during the 2004-2006 period, more exposed households may be stock-rich households who experience a positive stock wealth shock. To address this concern, we additionally control for the stock share of financial and gross wealth in a robustness exercise.

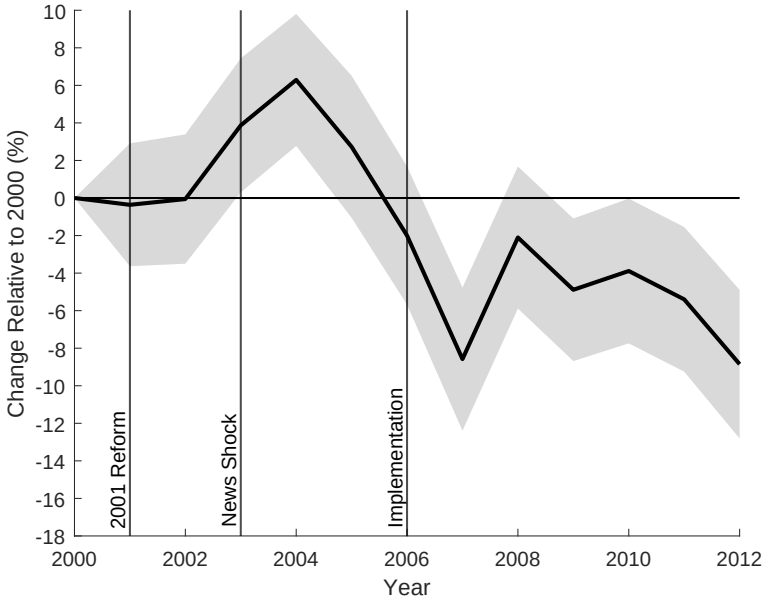
5.3 Results

This section presents our main empirical results. We first present the baseline result before discussing the robustness of our findings.

Main results. Figure 7 displays the spending response of the treated households relative to the control group. Between the announcement and implementation of the reform, spending by treated households increased by around 6% relative to the control group. After the implementation of the reform in 2006, relative spending decreased by about 5% on average. The positive spending response to the reform during the transition phase is consistent with an elasticity of intertemporal substitution bounded below by one. Moreover, the eventual negative spending response is consistent with the tax reform negatively affecting the treated households' permanent income.

Robustness. Figure 8 presents a number of robustness exercises. First, one concern with our empirical setup is that the treatment and control groups differ by pre-determined characteristics. To address this concern, we control for a number of plausible confounders in our baseline specification, as explained above. Figure 8a displays the results for three

394,000, 750,000, 936,800.



Notes: The figure displays the estimated coefficients of equation (4) with 95% confidence bands computed using standard errors clustered at the individual level.

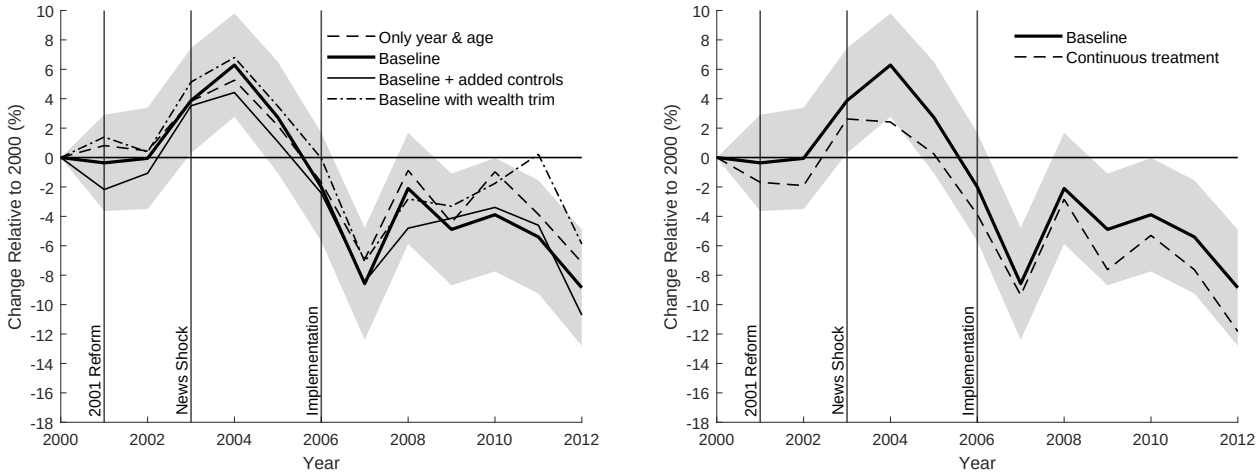
Figure 7: The spending response to the dividend tax reform.

alternative specifications, one with fewer controls (only year and age fixed effects), one where we include additional controls for the stock share of financial wealth and gross wealth, and a third where we additionally restrict the sample restriction to only include individuals with wealth above median with year and cohort to better balance the two groups' levels of wealth. These alternative setups are meant to gauge to what extent different wealth levels or different exposure to financial markets may matter for our results. The results change only marginally.

Second, in our baseline specification, we define households as treated if their dividend income as a share of gross income is sufficiently high. An alternative setup is to use the dividend share of gross income as a continuous measure of treatment exposure. The specification we use is

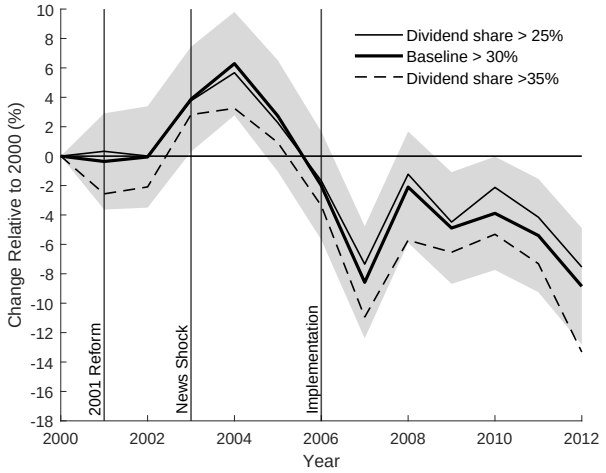
$$c_{i,2000+h} - c_{i,2000} = \alpha + \sum_{h=\underline{h}}^{H} \beta_h \left(\text{exposure}_{i,2000} \times \omega_{2000+h} \right) + \sum_{h=\underline{h}}^{H} \Gamma_h' \left(\mathbf{X}_{i,2000} \times \omega_{2000+h} \right) + \varepsilon_{i,h} \qquad (5)$$

where $\text{exposure}_{i,2000}$ is the average dividend income as a share of gross income for household i in 2000 and 2002. The estimated coefficients are not directly comparable across the two specifications. To make the two lines comparable (same y-axis), we adjust the estimated coefficients such that they denote the effect of an increase in treatment exposure

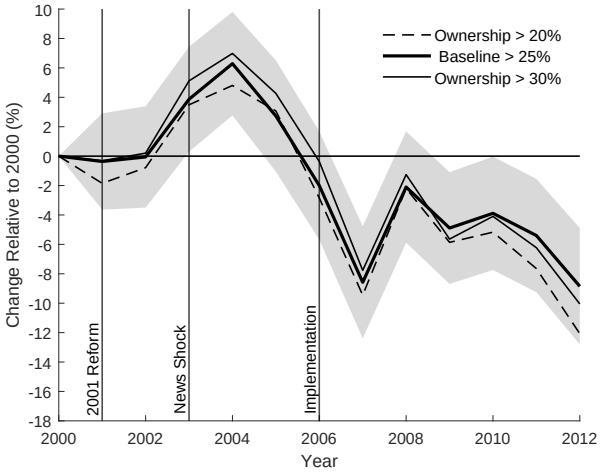


(a) Alternative sets of controls.

(b) Continuous treatment.



(c) Alternative dividend thresholds.



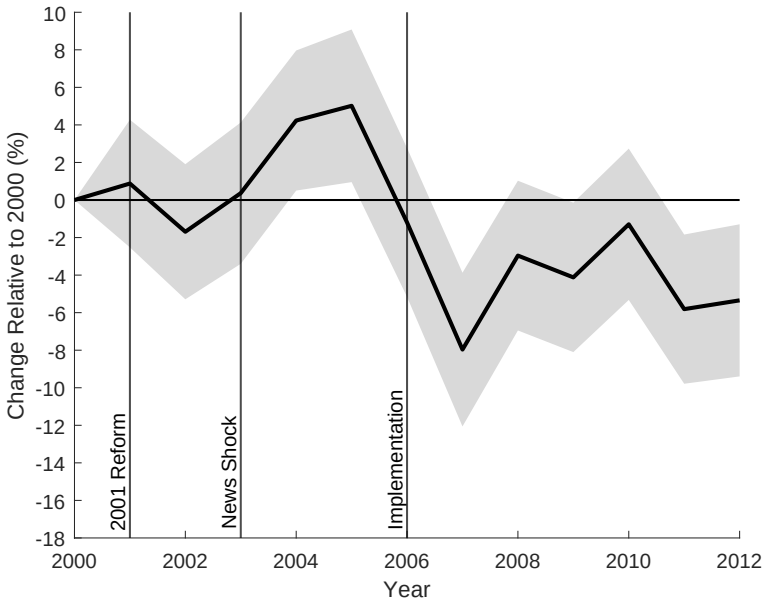
(d) Alternative ownership thresholds.

Notes: The figures display the estimated coefficients of equation (4) with 95% confidence bands computed using standard errors clustered at the individual level.

Figure 8: Robustness of the empirical results.

by 50 percentage points, corresponding to the average exposure (47.29) of the treatment group in the baseline setup. Figure 8b presents the estimated (adjusted) coefficients from equation (5) together with the baseline results. The results using this alternative specification are very similar. There is a positive response after the announcement and a negative response after the implementation of the dividend tax reform. Moreover, this specification also implies that the treatment effect we estimate is approximately the same as an increase in dividend share of gross income by 50 percentage points.

Third, in our baseline specification, we define households as treated if, on average, their dividend income as a share of gross income is more than 30% in 2000 and 2002. The choice of this exact threshold is somewhat arbitrary and represents a trade-off between having a



Notes: The figure displays the estimated coefficients of equation (4) using the alternative sample in Appendix F with 95% confidence bands computed using standard errors clustered at the individual level.

Figure 9: The spending response to the dividend tax reform. Alternative control group.

treatment group that is exposed and a treatment group that is relatively large. Figure 8c displays results when varying this threshold (25% and 35%), illustrating that the results do not change materially.

Fourth, in our sample restrictions, we require owners to hold ownership shares above 25% in all private businesses they own. In Figure 8d, we display robustness to this restriction by either relaxing it (20%) or tightening it (30%). The results do not change materially.

Finally, in the baseline setup, the control group consists of owners of private businesses whose business did not pay dividends in 2000 and 2002. In Appendix F, we present results when using an alternative control group: Wealthy households who do not own a private business. Table A.1 illustrates that this alternative control group also consists of households that are relatively younger and have lower spending, income, and wealth. Nevertheless, we show in Figure 9 that the results are similar when estimating equation (4) using this alternative control group. Indeed, spending among treated households increased by about 5% in 2005 and declined by about 5% on average relative to the control group, similar to the results with our main sample specification.

Taking stock. We find that news about a future permanent dividend tax reform caused a positive response in the spending of households with high dividend income intensity.

Implementation of the reform, in turn, caused a persistent decline in the spending of the same group. The observed front-loading of spending in reaction to a future decline in post-tax income is generally consistent with an elasticity of intertemporal substitution (EIS) greater than unity, as illustrated in Section 2 by our simple model. To pinpoint the exact value of the implied EIS, we now turn to our structural general equilibrium framework that we eventually bring to the data. This framework will also help us account for additional confounding channels not included in the simple model in Section 2.

6 A Model of Capitalists and Workers

In this section we combine our empirical estimates with a structural model in order to find a value for the elasticity of intertemporal substitution implied by the estimated consumption response in the previous section. Our approach introduces dividend tax news shocks into an otherwise standard two-agent, capitalist-worker general equilibrium framework.

6.1 Preferences

There is a continuum of households of measure one. A fraction λ of households are *workers* indexed by w. The remaining fraction $1-\lambda$ are *capitalists* indexed by k. Both household types have the same preferences and supply an exogenous amount of labor with the wage rate W_t normalized to unity for now. Capitalists can save by purchasing claims on firms, which in turn produce capital. Workers, on the other hand, can only save in a risk-free bond and receive no other capital income. This modeling approach is similar to the "stockholder and non-stockholder" setup of Mankiw and Zeldes (1991) and Guvenen (2009).[24]

Capitalists. Capitalists solve the following constrained optimization problem by maximizing utility subject to a sequence of constraints:

$$\max \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j U\left(C_{k,t+j}\right) \quad \text{s.t.} \quad S_{k,t+1}P_t + C_{k,t} \leq (1-\tau_{k,t})N_k + S_{k,t}(D_t + P_t) + T_{k,t}. \tag{6}$$

[24] Alternatively, we can set up a situation where both types of agents own claims on firms, but only one type receives dividend income. In this case, we can label such an environment as a model of dividend-earning and non-dividend-earning capitalists. A dividend distribution function exogenously transfers the total amount of dividend earnings to a fraction of business owners, who then distribute that amount equally among themselves. The fraction of owners receiving dividends would be $1-\lambda$, with the remaining fraction λ representing business owners who do not earn dividend income. As a result, this approach is equivalent to our baseline. The key assumption is that only a fraction of agents receive dividend income, i.e., are exposed to the reform.

The period utility function is of the CRRA form: $U(C_{k,t}) = \frac{C_{k,t}^{1-1/\psi}-1}{1-1/\psi}$, if $\psi \neq 1$, and $U(C_{k,t}) = \ln(C_{k,t})$, if $\psi = 1$. $\beta \in (0,1)$ is the subjective discount factor, $C_{k,t}$ denotes the capitalists' consumption in period t, N_k is labor supply, $S_{k,t}$ are claims on firms, P_t is the market value of those claims, D_t are dividends paid by firms to capitalists, $\tau_{k,t}$ is the proportional labor income tax, and $T_{k,t}$ is a lump-sum tax or subsidy that is paid out to the capitalist. We require $\tau_{k,t}$ due to institutional features of the Norwegian tax reform.

Utility maximization subject to the sequence of period budget constraints implies the standard Euler equation for firm shares:

$$1 = \mathbb{E}_t \left[\beta \left(\frac{C_{k,t+1}}{C_{k,t}} \right)^{-1/\psi} \frac{D_{t+1} + P_{t+1}}{P_t} \right]. \tag{7}$$

Of crucial interest to us is the parameter ψ, the EIS. Absence of arbitrage implies existence of a unique stochastic discount factor $\Lambda_{t,t+1}$ that prices all assets in the economy. $\Lambda_{t,t+1}$ is defined in terms of consumption of the capitalists, because they own all firms and are the marginal investor: $\Lambda_{t,t+1} \equiv \beta \left(\frac{C_{k,t+1}}{C_{k,t}} \right)^{-1/\psi}$.[25] Together with the usual transversality condition, by forward substitution the pricing equation for shares can be obtained as: $P_t = \mathbb{E}_t \sum_{j=t+1}^{\infty} \Lambda_{t,j} D_j$. One then obtains the intertemporal budget constraint: $\mathbb{E}_t \sum_{j=0}^{\infty} \Lambda_{t,t+j} C_{k,t+j} \leq P_t + \mathbb{E}_t \sum_{j=0}^{\infty} \Lambda_{t,t+j} (1 - \tau_{k,t+j}) N_k$.

Workers. Workers save in zero net-supply, risk-free, one-period bonds $B_{w,t}$ that pay an exogenous and state non-contingent gross return R_t^B each period. The period utility is given by:

$$\max \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j U\left(C_{w,t+j}\right) \quad \text{s.t.} \quad C_{w,t} + \frac{B_{w,t+1}}{R_t^B} = N_w + B_{w,t}, \tag{8}$$

where $U(C_{w,t})$ is the same CRRA period utility function as that of capitalists. The Euler equation for the worker is:

$$1 = R_t^B \mathbb{E}_t \left[\beta \left(\frac{C_{w,t+1}}{C_{w,t}} \right)^{-1/\psi} \right]. \tag{9}$$

[25] We can also assume recursive utility (Kreps and Porteus, 1978; Epstein and Zin, 1989, 1991; Weil, 1990). In the absence of aggregate or idiosyncratic uncertainty in the model, our results on the implied EIS in Section 7 are identical to what we find with CRRA preferences. This equivalence can also be shown analytically. The stochastic discount factor (SDF) implied by CRRA preferences is $\beta \left(\frac{C_{t+1}}{C_t} \right)^{-1/\psi}$. With recursive preferences the SDF becomes $\beta \left(\frac{C_{t+1}}{C_t} \right)^{-1/\psi} \left(\frac{V_{t+1}^{1-\gamma}}{\mathbb{E}_t V_{t+1}^{1-\gamma}} \right)^{1-\frac{1}{\xi}}$ with V_{t+1} the continuation value of the household's problem, γ the parameter that controls risk aversion, and $\xi = \frac{1-\gamma}{1-\frac{1}{\psi}}$. In the absence of uncertainty, the last term of the pricing kernel reduces to unity and we recover the CRRA case.

6.2 Technology

There is a continuum of perfectly competitive firms of mass one that produce the final good. The production technology is Cobb-Douglas:

$$F(A, K_t, N_t) = AK_t^{\alpha} N_t^{1-\alpha}, \tag{10}$$

where A is Hicks-neutral aggregate productivity (normalized to unity in the baseline case), K_t is aggregate capital, and N_t is aggregate labor supply in the economy (which is time-invariant in the baseline case). We further assume K_t is pre-determined at time $t-1$. Capital evolves according to $K_{t+1} = I_t + (1-\delta)K_t$, where I_t is firms' investment and δ is the depreciation rate.

Firms take the production function and the law of motion of capital as given and start the period with initial capital K_t. They decide the dividend payout D_t and investment I_t. Dividends are taxed at the rate $\tau_{d,t}$. The actual cost of a dividend payout D_t from the firm's perspective is

$$\varphi(D_t) = D_t(1 + \tau_{d,t})^{\kappa}, \tag{11}$$

where $\kappa \in [0, 1]$ is a parameter that governs (in equilibrium) the elasticity of the rate of return on saving for capitalists to dividend tax rate shocks.[26] The parameter κ captures, in a reduced-form way, various financial frictions, pecuniary costs, portfolio adjustment incentives, preferences for dividend smoothing, the risk-free exemption allowance of the reform, or degree of financial openness – all of which can affect the pass-through from capital income taxation to the rate of return on saving. Furthermore, κ can represent the degree of tax avoidance in the long run (Piketty and Saez, 2013). Many micro-founded environments can be nested in a parsimonious way by this specification. For example, $\kappa = 1$ represents the corner case with full pass-through, such as a closed economy with a single asset and no tax avoidance. The other corner case, $\kappa = 0$, represents an economy with no pass-through from the dividend tax change onto returns, such as an open economy with no portfolio adjustment frictions. Our estimation procedure will allow us to infer κ from the data.

The optimization problem of firms can be written recursively as:

$$V(K) = \max_{\{D,K'\}} [D + \mathbb{E}m'V(K')] \text{ s.t. } \varphi(D) + K' \leq (1-\delta)K + F(A, K, N) - N,$$

where $V(K)$ is the market value of the firm and m' is the stochastic discount factor, equal

[26] A similar reduced-form representation of the dividend payout function is laid out in Jermann and Quadrini (2012) in the context of equity payout adjustment costs and aggregate fluctuations.

to the stochastic discount factor of the capitalists Λ'. The first-order condition with respect to K' is:

$$\mathbb{E}\, m' \underbrace{\left(\frac{\varphi_D(D)}{\varphi_D(D')}\right)[1-\delta+F_k(A,K',N')]}_{\text{Net Return on Investment } R_t} = 1. \tag{12}$$

6.3 General Equilibrium

Fiscal Policy. In the steady state, revenue collection (from dividends or labor endowments) is set to 0. However, our policy experiments that raise dividend income taxes will create surpluses for the government. We therefore assume that the dividend tax and the proportional labor tax get lump-sum rebated back:

$$T_{k,t} = \tau_{d,t} D_t + \tau_{k,t} N_k \tag{13}$$

In the next section, we also allow for alternative ways of closing the fiscal balance, either by allowing government spending to vary or by international lending at a risk-free rate.

Aggregation and market clearing. Since only capitalists own firms, the holdings of each asset holder are pinned down solely by the share of capitalists in the economy: $S_t = S_{t+1} = \frac{1}{1-\lambda}$. Labor market clearing in our environment is trivial and equates the weighted average of endowments to the production function input N: $N = \lambda N_w + (1-\lambda)N_k$. Similarly, aggregate consumption is determined by: $C_t = \lambda C_{w,t} + (1-\lambda)C_{k,t}$. The resource constraint is $Y_t = C_t + I_t$, equal to the goods market clearing condition.

Tax processes. The government has at its disposal two policy instruments: the tax on dividends τ_d and the tax on capitalists' labor endowment τ_k. These instruments are assumed to follow the following exogenous stochastic processes:

$$\begin{aligned} \log \tau_{d,t} &= \log \tau_{d,t-1} + \sigma_d \varepsilon_{d,t-j}, \\ \log \tau_{k,t} &= \log \tau_{k,t-1} + \sigma_k \varepsilon_{k,t-j} \end{aligned} \tag{14}$$

where $\varepsilon_{d,t-j}$ and $\varepsilon_{k,t-j}$ are drawn from $\mathcal{N}(0,1)$.[27] The stochastic processes capture the expectation phenomenon of tax news shocks: an announcement at time *t-j*, captured by an innovation to $\varepsilon_{d,t-j}$, represents a credible signal that the tax rate $\tau_{d,t}$ will change at t. Given

[27] The shock processes are assumed to be unit root. In practice, we compute impulse responses to very persistent shocks, with autocorrelation of all shock processes set to 0.9999. Alternatively, one could solve for transitions from one steady state to another without any material impact on the results.

the context of the Norwegian tax reform, we set the news lag to j=2 because the model is calibrated at an annual frequency.

Definition 1. *A rational expectations general equilibrium, given tax policy innovation shocks* $\{\varepsilon_{d,t}, \varepsilon_{k,t}\}$ *and the tax policy processes, is defined as a set of policies for (i) capitalists:* C_k *and* S_k*; (ii) policies for workers:* C_w *and* B_w*; (iii) policies for firms:* K' *and* D*; (iv) firm market value* $V(K)$*; (v) and aggregate prices* m' *and* R^b*, such that: all policies solve the respective agents' optimization problems,* $m' = \beta \frac{U_c(c_k')}{U_c(c_k)}$*, and all markets clear at any given time* t*.*

6.4 Transitions Dynamics with Tax Avoidance

To allow for slow macroeconomic adjustment to aggregate dividend tax news shocks, we study transition dynamics that can depart from homogeneous and synchronized tax incidence. Households have access to tax avoidance opportunities, which are available upon arrival of the news, but disappear every period with an exogenous and i.i.d. probability $1-\theta$. As a result, following a tax news shock, a share θ of households is able to avoid the full tax burden while the remaining share is not. Over time, the economy converges to the steady state under the new tax regime, which is governed by the long-run pass-through parameter κ.[28] However, in the short run the incidence of dividend taxes is distributed heterogeneously. Intuitively, this approach will make the model-implied spending response to the tax news shock much smoother and thus more in line with the data, e.g. Figure 7.

To model this heterogeneity in a tractable way, we follow Auclert et al. (2021) and define several sequence-space objects of interest. First, define $\{\tau_{k,s}\}_{s=0}^{\infty}$ as the dividend tax *input sequence* that agents take as given. Denote with $C_{i,t} = C_{i,t}\left(\{\tau_{k,s}\}_{s=0}^{\infty}\right)$ the spending *output function* for agent $i \in \{k, w\}$. This function translates the input sequence into the consumption decision at time t. Now, define $\mathcal{J}_i^{HI}$ as the *homogeneous-incidence* Jacobian matrix for a household of type $i \in \{k, w\}$. Each entry of $\mathcal{J}_i^{HI}$ satisfies: $\mathcal{J}_{i,t,s}^{HI} = \frac{\partial C_{i,t}}{\partial \tau_{k,s}}$. That is, it summarizes the optimal consumption response of each agent at time t to exogenous shocks to dividend taxes at horizon s under the special case of $\theta = 0$.

Next, we construct the *heterogeneous-incidence* Jacobian $\mathcal{J}_i$ recursively as a function of $\mathcal{J}_i^{HI}$ for each household type. We compute each entry of $\mathcal{J}_i$ in the following way, focusing on the relevant case of $s > 0$, since the tax reform is announced in advance:

$$\mathcal{J}_{i,t,s} = \begin{cases} (1-\theta)\mathcal{J}_{i,t,s}^{HI} & t = 0, s > 0 \\ \theta \mathcal{J}_{i,t-1,s-1} + (1-\theta)\mathcal{J}_{i,t,s}^{HI} & t > 0, s > 0 \end{cases} \tag{15}$$

[28] Our modelling choice borrows from the literature on deviations from full information and rational expectations (FIRE). Specifically, we build on Mankiw and Reis (2002); Reis (2006a,b); Auclert et al. (2020); Carroll et al. (2020) by assuming an i.i.d. disappearance of tax avoidance opportunities.

Parameter	Value	Description
λ	0.990	Share of workers
β	0.980	Discount factor
δ	0.075	Depreciation rate
α	0.330	Capital share
N	0.300	Labor endowment
A	1.000	Productivity
σ_d	0.280	St. dev., capitalist dividend tax news
σ_l	0.104	St. dev., capitalist labor tax news

Table 2: Externally set model parameters.

Finally, the differential response in spending between the capitalists and the workers can be readily computed as the column-difference between $\mathcal{J}_{k,t,\hat{s}}$ and $\mathcal{J}_{w,t,\hat{s}}$ for all t and a shock at some $\hat{s}$. Parameter θ, as mentioned previously, captures the extent of tax avoidance intensity in the short run. We will infer it from the data during our impulse response matching exercise.

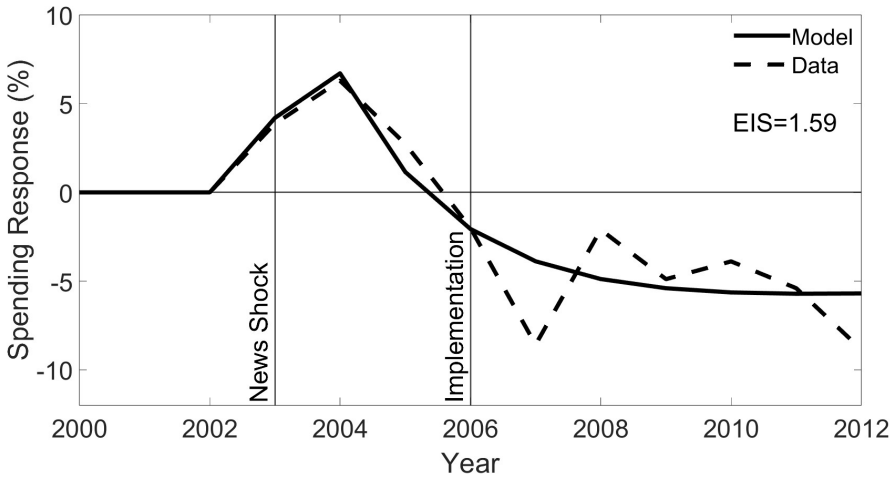
7 Identifying the Elasticity of Intertemporal Substitution

Our modeling approach facilitates a comparison with our empirical strategy. First, capitalists and workers represent the treated and the control groups in our empirical specification, respectively. Second, a dividend tax news shock and implementation in the model correspond to the institutional features of the 2003-2006 Norwegian tax reform. The differential in spending responses of capitalists and workers after the announcement therefore maps directly to our empirical difference-in-differences estimate in the post-announcement/pre-implementation period. We first describe which parameters we calibrate externally. Next, we show how we leverage our structural model to find a point estimate for the elasticity of intertemporal substitution.

7.1 External Parametrization

Table 2 lists all the externally calibrated parameters. The frequency of our calibration is annual. The subjective discount factor and depreciation rates are set to $\beta = 0.98$ and $\delta = 0.075$, respectively. The fraction of capitalists $1 - \lambda$ is set to 0.01, which corresponds to the share of business owners in the data. The capital share is set to $\alpha = 0.33$. Labor endowments $N = N_k = N_w$ are set to 0.3.

Standard deviations of the dividend and labor tax shocks are set in order to represent



Notes: This figure shows the differential response of spending in the model (straight line) and the data (dashed line) in response to the tax reform. Differential spending in the model is defined as consumption by capitalists less consumption by workers. Differential spending in data is defined accordingly in Section 5.

Figure 10: Model and data responses to the dividend tax news shock.

the institutional details of the Norwegian tax reform with $\sigma_d = 0.28$ and $\sigma_l = 0.104$. In our quantitative exercises we will be simulating tax synchronization with a one-standard deviation positive news shock for the dividend tax and a one-standard deviation negative news shock for the labor income tax on capitalists. In combination, these two shocks map exactly to the Norwegian 2003-2006 experience of a simultaneous increase in the dividend tax rate and reduction of the marginal labor income tax rate for the highest income bracket.

7.2 Impulse Response Matching

Our identification procedure is a variant of impulse response matching. First, we take the empirically documented differential response of the treated (capitalists) vs. the control (workers) groups as given in Figure 7. We calibrate a sub-set of model parameters externally, as reported in Table 2. Second, we construct a coarse three-dimensional grid for the pass-through parameter κ, the EIS ψ, and heterogeneous tax incidence θ. Finally, we locate the point on the grid that minimizes the distance between the empirical and the model-implied spending differentials. The final product of this exercise is $\bar{\psi}$: the value of the EIS that corresponds to the global minimal IRF matching error. Further details on the computational procedure are provided in Appendix G.

7.3 Main Result

Figure 10 shows the best-fitting model-generated consumption response of capitalists less the consumption response of workers. The implied $\bar{\psi}$ is 1.59, a value comfortably greater than unity. The model-implied spending response is close to the data and follows a similar shape: the spending differential increases after the announcement and falls gradually after implementation. This pattern is also consistent with our illustrative model laid out in Section 2. An EIS greater than 1 is thus essential for understanding the spending reactions to *news* about future dividend taxes and net portfolio returns, a key result of this paper.

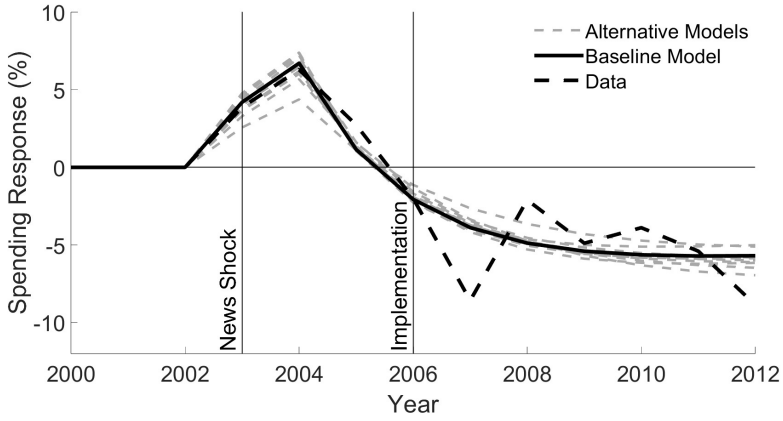
Incidentally, our estimated EIS of 1.59 is very close to what is typically assumed in the literature. Barro (2009) and Gabaix (2012) both set the EIS to 2 in the canonical framework of disaster risk and asset prices. Kaplan and Violante (2014) calibrate the EIS to 1.5 in their model of consumption responses to fiscal stimulus shocks. Bansal and Yaron (2004) also set the EIS to 1.5 in their model of slow-moving long-run risks. Thus, our microeconomic estimates of the EIS provide empirical support for a range of theoretical and quantitative macroeconomic and financial models that build on the assumption that the EIS is greater than unity.

The corresponding values of identified $\bar{\kappa}$ and $\bar{\theta}$ are 0.035 and 0.594, respectively. This suggests that the model is characterized by low long-run reform pass-through and high short-run tax avoidance. The former is indicative of Norway being a fairly open economy with low portfolio adjustment frictions and high substitutability across domestic and foreign assets. The latter points to high levels of private business intertemporal tax shifting, as suggested in the aggregate data in Section 3 and in the administrative data in Section 4, and as argued by Alstadsæter et al. (2014). The sluggish response of stock returns, as shown in Section 3, is also consistent with the above finding.

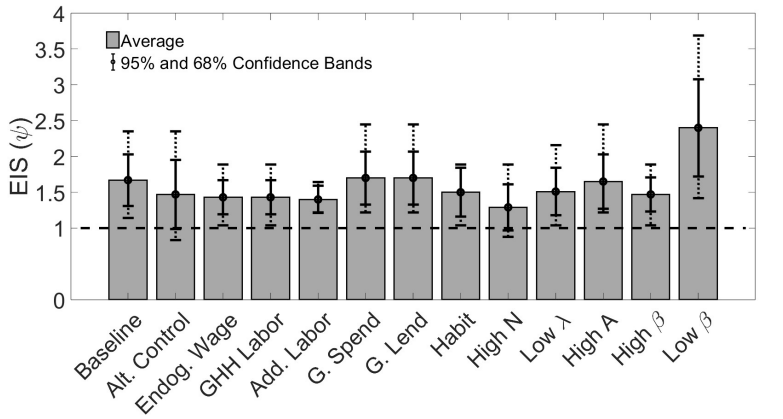
7.4 Model Robustness and Sensitivity Analysis

Our main result that the EIS of capitalists is greater than unity could be sensitive to the uncertainty of estimated responses, a concern that we address with bootstrapping. In principle, it could also suffer from model misspecification or specific parameter choices. We address these concerns in a series of robustness tests. We modify the baseline model in each extension and conduct the impulse response matching exercise as before.

Model bootstrap and confidence bands. Our empirical analysis in Section 5 revealed that the front-loading of relative spending during the 2003-2006 transition period was statistically significant at the 95% level. However, the confidence bands could be perceived



(a) Spending response in alternative models.



(b) Average EIS with 95% and 68% confidence bands in alternative models.

Notes: Panel (a) plots relative spending responses to the Norwegian tax reform in the baseline and 10 alternative models, as described in the text. Panel (b) reports average EIS values as well as 68% and 95% bootstrapped confidence bands implied by each alternative model. Model versions, from left to right, correspond to: the baseline, alternative control group specification, and baseline plus extensions with endogenous wages, endogenous GHH labor supply, endogenous additively separable labor supply, fiscal rule with government spending instead of lump-sum taxes, fiscal rule with government bond lending instead of lump-sum taxes, habit formation instead of heterogeneous tax incidence, high value of the labor endowment, high share of capitalists, high aggregate productivity, high discount factor, and low discount factor.

Figure 11: Model robustness.

to be wide. To account for coefficient uncertainty, we perform a bootstrapping exercise in which we compute the EIS using 10,000 independent draws from the empirically estimated spending response coefficients. Results are summarized in Panel (b) of Figure 11. Every column and the associated error bars present the mean estimate of the EIS along with 95% and 68% confidence bands. We show that for the baseline model, the average EIS is 1.67 and the 95% confidence band is [1.14, 2.35].[29] We repeat the same bootstrapping procedure

[29]The average estimate turns out to be slightly greater than the baseline of 1.59 due to some (mild) non-linearities.

for every extension and sensitivity test below.

Alternative control group. Our baseline model is calibrated to match the differential spending response as in Figure 7. This setup corresponds to defining the control group as business owners that did not receive dividend income in 2000 and 2002. In Figure 9 we have also presented an alternative empirical setup in which the control group includes individuals who are wealthy but do not own a private business. We now also calibrate our model to match the differential spending response produced by the specification with an alternative control group. The implied EIS for this alternative approach is 1.37, and the values of κ and θ are 0.03 and 0.69, respectively. These values are very close to our baseline results. Moreover, as shown in the second column ("Alt. Control") of Panel (b) of Figure 11, our bootstrap procedure yields an average EIS of 1.47.

Endogenous wages and labor supply. In an attempt to test sensitivity to the specification of the labor market, we perform three robustness exercises. First, we allow for an endogenous unit price of labor, the wage rate W_t, which in the baseline was invariant and set to unity. The competitive wage rate is determined via the marginal product of labor in every period: $W_t = A(1-\alpha)K_t^{\alpha}N_t^{-\alpha}$.

Second, we endogenize labor supply of both capitalists and workers in two separate ways. We first assume non-separability between consumption and leisure in the spirit of Greenwood et al. (1988). The first-order condition with respect to labor supply for agent type x is: $(1-\tau_x)W_t = \phi N_{x,t}^{\chi}$, where τ_x is a proportional labor tax only in the case of capitalists. Both types have the same Frisch elasticity $\frac{1}{\chi}$ that is set to unity. The labor disutility parameter ϕ is set to a value which guarantees that hours equal 0.3 in the steady state for both types, similarly to our baseline case.

Third, we assume that utility is additively separable in consumption and labor. The first-order condition with respect to labor supply for agent type x is now: $(1-\tau_x)W_t C_{x,t}^{-1/\psi_x} = \phi N_{x,t}^{\chi}$. We fix the EIS of workers to unity and allow ψ_k to be determined by the impulse response matching exercise as before.

Alternative fiscal rules. The baseline model assumes that any fiscal surplus is rebated back to the households along the transition path following dividend tax news shocks. We now introduce two alternative fiscal rules. First, we alternatively assume the government uses the surplus to finance productive government spending. The government budget constraint is now $G_t = \tau_{d,t}D_t + \tau_{k,t}N_{k,t}$. And the resource constraint becomes $Y_t = C_t + I_t + G_t$. Second, we allow the government to lend abroad via one-period bonds at the risk-free rate

R_t^F. The government budget constraint becomes: $\frac{1}{R_t^F}B_{t+1}+\tau_{d,t}D_t+\tau_{k,t}N_{k,t} = B_t$. In equilibrium, the risk-free rate is pinned down by the stochastic discount factor: $R_t^F = \frac{1}{\mathbf{E}_t m_{t,t+1}}$.

Consumption habit instead of heterogeneous tax incidence. An alternative way of generating smooth transition dynamics following tax news shocks is through intertemporal non-separability of spending, also known as habits. We therefore introduce internal habit formation in the spirit of Christiano et al. (2005) into the consumption problem of the capitalists and revert the tax incidence parameter θ to 0. The stochastic discount factor becomes: $\Lambda_{t,t+1} = \beta\left(\frac{C_{k,t+1}-\zeta C_{k,t}}{C_{k,t}-\zeta C_{k,t-1}}\right)^{-1/\psi_k}$. The habit parameter ζ replaces θ as the determinant of how slow the post-announcement transition is. Now, we do not parameterize ζ but instead construct a non-linearly spaced grid over the interval $[0.01, 0.99]$ and let the data speak to its value in a three-directional impulse response matching exercise over the grids of ψ, κ, and now ζ.

Parameter sensitivity. Our main result may be sensitive to the values that we assign to externally calibrated parameters. First, the *share* of capitalists $1-\lambda$ directly controls the mass of agents that are going to be affected by the reform experiment. Fortunately, we can pin down the value of λ in the data. However, our baseline λ could be argued to be high; although it is Norway-consistent, the external validity of our findings could be questioned if λ is generally lower in other countries. We therefore lower our λ to 0.95 and re-do our baseline analysis.

Second, another parameter that could potentially impact our results is the labor endowment N. As part of a sensitivity test, we raise $N = N_k = N_w$ to 0.5, a value also normally used in the macro literature (Kaplan et al., 2018). Third, we test the aggregate "state-dependency" of our results by setting aggregate productivity A to a higher value of 1.02 and asking whether the implied EIS of capitalists is lower/higher in booms. This is useful to check also considering that Norway went through a booming phase precisely when the reform was introduced.

The final comparative static exercise is with respect to the discount factor β. Capital income-earning business owners and workers could differ in the degree of patience as patience heterogeneity is known to explain an important fraction of observable inequality in wealth (Fernández-Villaverde and Levintal, 2024). To allow for the possibility that capitalists are either much more patient *or* less patient than workers in our model, we now set the β of capitalists first to 1.03 and then to 0.93 while keeping the β of workers unchanged.

Results from robustness tests. Figure 11 reports all the results in two stages. First, in Panel (a) we plot relative spending responses implied by all 12 alternative models and/or parameterizations that we have described above. Every pattern is quantitatively indistinguishable from the baseline case and tracks the data well. Second, in Panel (b) we report the values of ψ that are implied by each model. Specifically, we report averages along with 95% and 68% bootstrapped confidence bands. The EIS is consistently above 1 in all of our robustness exercises. A notable case is the column "low β" (impatient capitalists), which produces irregularly noisy bands in general but the EIS is still significantly larger than unity at the 95% confidence level.

Additional results. We present and discuss four supplementary results in Figure A.3 Appendix H. First, we show how the model-implied spending response would look if we set the EIS (ψ) to a counterfactually low value such as 0.1, which is occasionally mentioned in the literature. Second, for completeness, we also report the spending response to *surprise* shocks to taxes, as opposed to news shocks. Third, we set $\theta = 0$ and test model behavior without short-run tax avoidance. Fourth, we present the full density of EIS estimates that are produced by bootstrapping the baseline model. Across these supplementary exercises, our main results are upheld, the quantitative value of proper calibration is highlighted, and the role of tax *news* shocks is further emphasized.

8 Conclusion

In this paper, we leverage the combination of a unique policy experiment – the large and salient 2006 Norwegian dividend tax reform – and the rich Norwegian registry data to make progress on identifying the elasticity of intertemporal substitution. Specifically, rather than estimating the EIS via the Euler equation as is common, we look directly at the anticipatory spending response to the reform announcement of exposed vs. less exposed households. Our results show a strong anticipatory spending response, which indicates a value of EIS above 1. This is confirmed via calibration of a standard capitalist-worker model: the model is consistent with the observed relative spending response only if the EIS is greater than 1.

In addition to allowing us to back out the EIS for business owners, our empirical methodology and findings could be useful for the large and growing literature on macroeconomics with heterogeneity. An increasingly popular approach in that literature involves using sequence-space Jacobians for solving and simulating quantitative models (Auclert et al., 2021, 2023). One important object in this general class of models is the matrix that

contains the effects of current and *future* interest rate changes on current and future consumption. Our empirical approach and estimated dynamic spending responses to *news* about future interest rate changes can be used to construct parts of this Jacobian.

References

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *Review of Economic Studies*, 2021, *89(1).*

Alstadsæter, Annette and Erik Fjærli, "Neutral taxation of shareholder income? Corporate responses to an announced dividend tax," *International Tax and Public Finance*, 2009, *16* (4), 571–604.

_ **, Martin Jacob, and Roni Michaely**, "Do dividend taxes affect corporate investment?," *Journal of Public Economics*, 2017, *151*, 74–83.

_ **, Niels Johannesen, and Gabriel Zucman**, "Tax evasion and inequality," *American Economic Review*, 2019, *109* (6), 2073–2103.

_ **, Wojciech Kopczuk, and Kjetil Telle**, "Are Closely Held Firms Tax Shelters?," *Tax Policy and the Economy*, 2014, *28* (1), 1–32.

Andreolli, M. and P. Surico, "'Less is More': Consumer Spending and the Size of Economic Stimulus Payments," *CEPR Working Paper*, 2021, *15918.*

Attanasio, Orazio P and Guglielmo Weber, "Consumption growth, the interest rate and aggregation," *Review of Economic Studies*, 1993, *60* (3), 631–649.

_ **and** _ , "Consumption and saving: models of intertemporal allocation and their implications for public policy," *Journal of Economic literature*, 2010, *48* (3), 693–751.

_ **and Martin Browning**, "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*, 1995, pp. 1118–1137.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub, "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 2021, *89* (5), 2375–2408.

_ **, M. Rognlie, and L. Straub**, "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," *NBER Working Paper*, 2020, *26647.*

_ **,** _ **, and** _ , "The intertemporal keynesian cross," *Journal of Political Economy*, 2023.

Auerbach, Alan J, "Wealth maximization and the cost of capital," *Quarterly Journal of Economics*, 1979, *93* (3), 433–446.

Auerbach, Alan J. and Kevin A. Hassett, "The 2003 Dividend Tax Cuts and the Value of the Firm: An Event Study," in Alan J. Auerbach, James R. Hines, Jr., and Joel Slemrod, eds., *Taxing Corporate Income in the 21st Century*, Cambridge University Press, 2007, p. 93–126.

Bansal, Ravi and Amir Yaron, "Risks for the long run: A potential resolution of asset pricing puzzles," *Journal of Finance*, 2004, *59* (4), 1481–1509.

_ **, Dana Kiku, and Amir Yaron**, "Long Run Risks, the Macroeconomy, and Asset Prices," *American Economic Review*, 2010, *100* (2), 542–46.

_ , _ , **and** _ , "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 2011, *1* (1), 183–221.

Barro, Robert J., "Rare disasters, asset prices, and welfare costs," *American Economic Review*, 2009, *99* (1), 243–64.

_ **and Gordon Liao**, "Rare disaster probability and options pricing," *Journal of Financial Economics*, 2021, *139*, 750–769.

_ **and Jason Furman**, "Macroeconomic effects of the 2017 tax reform," *Brookings Papers on Economic Activity*, 2018, *2018* (1), 257–345.

_ **and José F. Ursúa**, "Rare Macroeconomic Disasters," *Annual Review of Economics*, 2012, *4* (1), 83–109.

Barsky, Robert B and Eric R Sims, "News shocks and business cycles," *Journal of Monetary Economics*, 2011, *58* (3), 273–289.

Beaudry, Paul and Bernd Lucke, "Letting Different Views about Business Cycles Compete," *NBER Macroeconomics Annual*, 2010, *24*, 413–456.

_ **and Eric Van Wincoop**, "The intertemporal elasticity of substitution: An exploration using a US panel of state data," *Economica*, 1996, pp. 495–512.

_ **and Franck Portier**, "An exploration into Pigou's theory of cycles," *Journal of Monetary Economics*, 2004, *51* (6), 1183–1216.

_ **and** _ , "Stock prices, news, and economic fluctuations," *American Economic Review*, 2006, *96* (4), 1293–1307.

_ **and** _ , "News-driven business cycles: Insights and challenges," *Journal of Economic Literature*, 2014, *52* (4), 993–1074.

Best, Michael Carlos, James S Cloyne, Ethan Ilzetzki, and Henrik J Kleven, "Estimating the elasticity of intertemporal substitution using mortgage notches," *Review of Economic Studies*, 2020, *87* (2), 656–690.

Bilbiie, Florin, "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic," *Journal of Economic Theory*, 2008, *140* (1), 162–196.

Blanchard, O., J-P. L'Huillier, and G. Lorenzoni, "News, Noise, and Fluctuations: An Empirical Exploration," *American Economic Review*, 2013, *103.*

Blundell, Richard, Martin Browning, and Costas Meghir, "Consumer demand and the life-cycle allocation of household expenditures," *Review of Economic Studies*, 1994, *61* (1), 57–80.

Boissel, Charles and Adrien Matray, "Dividend Taxes and the Allocation of Capital," *Working Paper*, September 2022.

Bradford, David F, "The incidence and allocation effects of a tax on corporate distributions," *Journal of Public Economics*, 1981, *15* (1), 1–22.

Browning, Martin and Thomas F Crossley, "Luxuries are easier to postpone: A proof," *Journal of Political Economy*, 2000, *108* (5), 1022–1026.

Calvet, Laurent E, John Y Campbell, Francisco Gomes, and Paolo Sodini, "The cross-section of household preferences," *Working Paper*, 2021.

Campbell, John Y, "Consumption-based asset pricing," *Handbook of the Economics of Finance*, 2003, *1*, 803–887.

_ **and N Gregory Mankiw**, "Consumption, income, and interest rates: Reinterpreting the time series evidence," *NBER Macroeconomics Annual*, 1989, *4*, 185–216.

Carroll, C., E. Crawley, J. Sracalek, K. Tokuoka, and M. White, "Sticky Expectations and Consumption Dynamics," *Americal Economic Journal: Macroeconomics*, 2020, *12.*

Cashin, David and Takashi Unayama, "Measuring Intertemporal Substitution in Consumption: Evidence from a VAT Increase in Japan," *Review of Economics and Statistics*, 05 2016, *98* (2), 285–297.

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

_ **and** _ , "Dividend and corporate taxation in an agency model of the firm," *American Economic Journal: Economic Policy*, 2010, *2* (3), 1–31.

_ **, Joseph Rosenberg, and Emmanuel Saez**, "The Effects of Taxes on Market Responses to Dividend Announcements and Payments: What Can We Learn from the 2003 Dividend Tax Cut?," in Alan J. Auerbach, James R. Hines, Jr., and JoelEditors Slemrod, eds., *Taxing Corporate Income in the 21st Century*, Cambridge University Press, 2007, p. 1–32.

Chodorow-Reich, Gabriel, Matthew Smith, Owen M Zidar, and Eric Zwick, "Tax policy and investment in a global economy," Technical Report, National Bureau of Economic Research 2024.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Crump, Richard K, Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa, "Subjective intertemporal substitution," *Journal of Monetary Economics*, 2022, *126*, 118–133.

Cummins, Jason G, Kevin A Hassett, R Glenn Hubbard, Robert E Hall, and Ricardo J Caballero, "A reconsideration of investment behavior using tax reforms as natural experiments," *Brookings Papers on Economic Activity*, 1994, *1994* (2), 1–74.

Eeckhoudt, Louis and Harris Schlesinger, "Putting risk in its proper place," *American Economic Review*, 2006, *96* (1), 280–289.

Epstein, L. and S. E. Zin, "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica*, 1989, *57.*

_ **and** _ , "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis," *Journal of Political Economy*, 1991, *99*.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso, "Asset market participation and portfolio choice over the life-cycle," *Journal of Finance*, 2017, *72* (2), 705–750.

_ **, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri**, "Heterogeneity and persistence in returns to wealth," *Econometrica*, 2020, *88* (1), 115–170.

Farhi, Emmanuel, Alan Olivi, and Iván Werning, "Price Theory for Incomplete Markets," *Working Paper*, 2022.

Feldstein, Martin S, "Corporate taxation and dividend behaviour," *Review of Economic Studies*, 1970, *37* (1), 57–72.

Fernández-Villaverde, Jesús and Oren Levintal, "The Distributional Effects of Asset Returns," Working Paper 32182, National Bureau of Economic Research March 2024.

Flynn, Joel P, Lawrence DW Schmidt, and Alexis Akira Toda, "Robust comparative statics for the elasticity of intertemporal substitution," *Theoretical Economics*, 2023, *18* (1), 231–265.

Furno, Francesco, "The Macroeconomic Effects of Corporate Tax Reforms," *Working Paper*, 2022.

Gabaix, X., "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *uarterly Journal of Economics*, 2012, *127*, 645–700.

Galí, Jordi, J David López-Salido, and Javier Vallés, "Understanding the effects of government spending on consumption," *Journal of the European Economic Association*, 2007, *5* (1), 227–270.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman, "Investment, capacity utilization, and the real business cycle," *American Economic Review*, 1988, pp. 402–417.

Gruber, Jonathan, "A Tax-Based Estimate of the Elasticity of Intertemporal Substitution.," *Quarterly Journal of Finance*, 2013, *3* (1), 1 – 20.

Guvenen, Fatih, "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: a Macroeconomic Perspective," *Journal of Monetary Economics*, 2006, *53*.

_ , "A Parsimonious Macroeconomic Model for Asset Pricing," *Econometrica*, 2009, *77* (6).

Hall, Robert E, "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 1988, *96* (2), 339–357.

_ **and Dale W Jorgenson**, "Tax policy and investment behavior," *American Economic Review*, 1967, *57* (3), 391–414.

Hansen, Lars Peter and Kenneth J Singleton, "Stochastic consumption, risk aversion, and the temporal behavior of asset returns," *Journal of political economy*, 1983, *91* (2), 249–265.

_ **, John Heaton, Junghoon Lee, and Nikolai Roussanov**, "Intertemporal substitution and

risk aversion," *Handbook of econometrics*, 2007, *6*, 3967–4056.

Harberger, Arnold C, "The incidence of the corporation income tax," *Journal of Political Economy*, 1962, *70* (3), 215–240.

Havránek, Tomáš, "Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting," *Journal of the European Economic Association*, 12 2015, *13* (6), 1180–1204.

_ , Roman Horvath, Zuzana Irsova, and Marek Rusnak, "Cross-country heterogeneity in intertemporal substitution," *Journal of International Economics*, 2015, *96* (1), 100–118.

Holm, Martin B, "Monetary transmission with income risk," *The Scandinavian Journal of Economics*, 2023, *125* (2), 441–460.

House, Christopher L and Matthew D Shapiro, "Temporary investment tax incentives: Theory with evidence from bonus depreciation," *American Economic Review*, 2008, *98* (3), 737–68.

Iachan, Felipe S, Plamen T Nenov, and Alp Simsek, "The choice channel of financial innovation," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 333–72.

Innstilling til Odelstinget nr. 23(2000-2001), *Innstilling O. nr. 23 (2000-2001)* 2000.

Jakobsen, Katrine, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman, "Wealth taxation and wealth accumulation: Theory and evidence from Denmark," *Quarterly Journal of Economics*, 2020, *135* (1), 329–388.

Jermann, U. and V. Quadrini, "Macroeconomic Effects of Financial Shocks," *American Economic Review*, 2012, *102(1).*

Kaplan, Greg and Giovanni L Violante, "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ , _ , and Benjamin Moll, "Monetary Policy According to HANK," *American Economic Review*, 2018, *108* (3).

Korinek, Anton and Joseph E Stiglitz, "Dividend taxation and intertemporal tax arbitrage," *Journal of Public Economics*, 2009, *93* (1-2), 142–159.

Kreps, D. M. and E. L. Porteus, "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica*, 1978, *46.*

Leite, David, "The Firm as Tax Shelter: Using Business Resources for Final Consumption," *Working Paper*, 2023.

Lorenzoni, G., "A Theory of Demand Shocks," *American Economic Review*, 2009, *99.*

Mankiw, G. and R. Reis, "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 2002, *117* (4).

Mankiw, N Gregory and Stephen P Zeldes, "The consumption of stockholders and nonstockholders," *Journal of Financial Economics*, 1991, *29* (1), 97–112.

NOU 2003:9, *NOU 2003:9 Skatteutvalget - Forslag til endringer i skattesystemet* 2003.

Piketty, T. and E. Saez, "Optimal Labor Income Taxation," *Handbook of Public Economics*, 2013, *5*.

Poterba, James M and Lawrence H Summers, "Dividend taxes, corporate investment, and 'Q'," *Journal of Public Economics*, 1983, *22* (2), 135–167.

_ **and** _ , "New evidence that taxes affect the valuation of dividends," *Journal of Finance*, 1984, *39* (5), 1397–1415.

Reis, R., "Inattentive Consumers," *Journal of monetary Economics*, 2006, *53* (8).

_ , "Inattentive Producers," *Review of Economic Studies*, 2006, *73* (3).

Ring, Marius A K, "Wealth taxation and household saving: Evidence from assessment discontinuities in Norway," *Working Paper*, 2020.

Schmidt, Lawrence and Alexis Akira Toda, "Do You Save More or Less in Response to Bad News? A New Identification of the Elasticity of Intertemporal Substitution," Technical Report, MPRA Paper No. 78983, 2015.

Stortingsmelding nr. 29(2003-2004), *Stortingsmelding nr. 29 (2003-2004) Om skattereform* 2004.

Straub, Ludwig, "Consumption, savings, and the distribution of permanent income," *Unpublished manuscript, Harvard University*, 2019.

_ **and Iván Werning**, "Positive long-run capital taxation: Chamley-Judd revisited," *American Economic Review*, 2020, *110* (1), 86–119.

Thoresen, Thor Olav, Erlend Eide Bø, Erik Fjærli, and Elin Halvorsen, "Skattereformen 2006–Har Skattesystemet Blitt Mer Omfordelende?," *Statistisk sentralbyrå*, 2010.

Vissing-Jørgensen, Annette, "Limited asset market participation and the elasticity of intertemporal substitution," *Journal of Political Economy*, 2002, *110* (4), 825–853.

_ **and Orazio P Attanasio**, "Stock-market participation, intertemporal substitution, and risk-aversion," *American Economic Review*, 2003, pp. 383–391.

Weil, O., "Nonexpected utility in macroeconomics," *Quarterly Journal of Economics*, 1990, *105*.

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

Yogo, Motohiro, "Estimating the elasticity of intertemporal substitution when instruments are weak," *Review of Economics and Statistics*, 2004, *86* (3), 797–810.

Zwick, Eric and James Mahon, "Tax policy and heterogeneous investment behavior," *American Economic Review*, 2017, *107* (1), 217–48.

Online Appendix for "Estimating the Elasticity of Intertemporal Substitution using Dividend Tax News Shocks"

Martin B. Holm Rustam Jamilov Marek Jasinski Plamen Nenov

July 2024

Contents

A Derivations for the Illustrative Model **1**

B Comparative statics with non-homothetic preferences **2**

C More Details on the Institutional Setting **4**

D Stock Market Analysis Details **6**

E Imputed Spending without Private Businesses **7**

F Details and Results for the Alternative Control Group **9**

G Numerical Details for the Structural Model **11**

H Additional Model Results **13**

A Derivations for the Illustrative Model

The agent's problem is

$$\max_{C_0,C_1,C_2} u(C_1) + u(C_2) + u(C_3)$$

$$s.t.\ \ C_0 + \frac{C_1}{R_1} + \frac{C_2}{R_1 R_2} = R_0 A_0.$$

The corresponding Euler equations are

$$u'(C_0) = R_1 u'(C_1), \quad u'(C_1) = R_2 u'(C_2).$$

Using the definition of the utility function, the budget constraint, and the two Euler equations, the consumption policy functions are

$$C_0 = \frac{R_0 A_0}{1 + R_1^{\psi-1} + R_1^{\psi-1} R_2^{\psi-1}}, \quad C_1 = \frac{R_1^{\psi} R_0 A_0}{1 + R_1^{\psi-1} + R_1^{\psi-1} R_2^{\psi-1}}, \quad C_2 = \frac{R_1^{\psi} R_2^{\psi} R_0 A_0}{1 + R_1^{\psi-1} + R_1^{\psi-1} R_2^{\psi-1}}.$$

The derivative of the consumption policy functions with respect to R_2 are

$$\frac{\partial C_0}{\partial R_2} = -(\psi - 1)\frac{R_1^{\psi-1} R_2^{\psi-2} R_0 A_0}{\left(1 + R_1^{\psi-1} + R_1^{\psi-1} R_2^{\psi-1}\right)^2} \equiv -(\psi - 1)\kappa_0(R_0 A_0, R_1, R_2)$$

$$\frac{\partial C_1}{\partial R_2} = -(\psi - 1)\frac{R_1^{-1} R_2^{\psi-2} R_0 A_0}{\left(1 + R_1^{\psi-1} + R_1^{\psi-1} R_2^{\psi-1}\right)^2} \equiv -(\psi - 1)\kappa_1(R_0 A_0, R_1, R_2)$$

$$\frac{\partial C_2}{\partial R_2} = \frac{\psi R_1^{\psi} R_2^{\psi-1} R_0 A_0 (1 + R_1^{\psi-1}) + R_1^{2\psi-1} R_2^{2\psi-2} R_0 A_0}{\left(1 + R_1^{\psi-1} + R_1^{\psi-1} R_2^{\psi-1}\right)^2} \equiv \kappa_2(R_0 A_0, R_1, R_2)$$

Hence, C_0 and C_1 are strictly decreasing in R_2 if and only if $\psi > 1$, while C_2 is strictly increasing in R_2 always.

B Comparative statics with non-homothetic preferences

We illustrate the robustness of the sign comparative statics for the value of the elasticity of intertemporal substitution relative to unity derived in Flynn et al. (2023) for a simple two-period setting with non-homothetic preferences as in Straub (2019).

Suppose that the agent's period t utility function is given by $u_t(C_t) = \frac{C_t^{1-\sigma_t}}{1-\sigma_t}$ ($\ln(C_t)$ if $\sigma_t = 1$). The agent allocates initial wealth A over period $t = 0$ consumption and period $t = 1$ savings A_1. The period $t = 1$ gross return is R, so that period $t = 1$ consumption is $C_1 = RA_1$.

We follow Flynn et al. (2023) and define the aggregator function between period $t = 0$ consumption c and the period $t = 1$ continuation value v as $f(c, v) = u_0(c) + u_1(v)$. Notice that f and $v = C_1$ define a strongly regular environment according to definition 1 in Flynn et al. (2023), as the aggregator f is strictly increasing and twice continuously differentiable with a non-negative cross-partial derivative, and the continuation value v is also strictly increasing and twice continuously differentiable by the properties of a CRRA utility function. Furthermore, the period $t = 1$ consumption c_0 is interior to the budget set given that the CRRA period utility satisfies Inada conditions. Consequently, Theorem 1 in Flynn et al. (2023) applies to our environment. Specifically, we have the following sign comparative static:

$$\operatorname{sgn}\left(\frac{\partial c}{\partial \log R}\right) = \operatorname{sgn}(1 - \varepsilon\psi),$$

where ψ is the elasticity of intertemporal substitution and ε is the Relative Elasticity of the Marginal Value of Wealth (REMV), an object that measures the impact of wealth effects on the response of consumption.

Following Flynn et al. (2023), we define the elasticity of intertemporal substitution (EIS) as

$$\psi = -\frac{\frac{\partial \log\left(\frac{c}{v}\right)}{\partial \log R}}{\frac{\partial \log\left(\frac{f_c}{f_v}\right)}{\partial \log R}}.$$

Similarly, noting that period $t = 1$ wealth and period $t = 1$ consumption coincide, we define the Relative Elasticity of the Marginal Value of Wealth (REMV) as

$$\varepsilon = \frac{\frac{\partial \log v_A}{\partial \log R}}{\frac{\partial \log v}{\partial \log R}}.$$

For the continuation value defined as $v = C_1$, the value of the REMV is

$$\varepsilon = \frac{\frac{\partial \log v_A}{\partial \log R}}{\frac{\partial \log v}{\partial \log R}} = 1,$$

and we have that

$$\operatorname{sgn}\left(\frac{\partial c}{\partial \log R}\right) = \operatorname{sgn}(1 - \psi).$$

Hence, the sign of the consumption response to the change in log returns $\log R$ is the same as the sign of the EIS relative to unity. Note, however, that in this case, the EIS is not a single parameter as would be the case with time-invariant preferences but instead depends on the values of both σ_0 and σ_1. Specifically, we can re-write the EIS as

$$\psi = -\frac{\partial \log c/\partial \log R - \partial \log v/\partial \log R}{-\sigma_0 \partial \log c/\partial \log R + \sigma_1 \partial \log v/\partial \log R} = \frac{1}{\sigma_0}\frac{\partial \log c/\partial \log R - \partial \log v/\partial \log R}{\partial \log c/\partial \log R - \frac{\sigma_1}{\sigma_0}\partial \log v/\partial \log R}.$$

Hence, with time-varying values of σ, the EIS is a function of the current level of consumption, the level of returns R, and the CRRA parameters for the period utility functions.

C More Details on the Institutional Setting

This appendix presents additional details on the dividend tax reform of 2006.[1]

On February 6, 2003, a government-commissioned committee published an official recommendation for a permanent dividend tax reform. The reform was then announced on March 26, 2004 and implemented on January 1, 2006. The main purpose of the tax reform was to reduce the difference in the marginal tax rates on labor income and capital income. The reform introduced a 28% personal tax on dividends and capital gains in excess of a threshold amount based on riskless returns set by the Ministry of Finance. Under the previous tax regime, dividends were tax-exempt for any shareholder, while capital gains were almost always applied to a zero base and hence were tax-exempt as well. Firms paid no taxes on dividends and capital gains either before or after the reform.[2] The reform also decreased the top marginal tax on labor income from 64.7% to 54.3%, while the sum of taxes paid by the firm and the investor on dividends and capital gains increased from 28% to 48.2%.

To see this, define capital income tax at the corporate level as τ_t^c, which is 28%. The dividend tax above the rate of return allowance (RAA) is 28%. The marginal dividend tax rate τ_t^d is

$$\tau_t^d = \overbrace{0.28}^{\text{Corporate tax}} + \underbrace{(1-0.28)}_{\text{Net of corporate tax}} \cdot \overbrace{0.28}^{\text{Dividend tax rate over RRA}} = 0.482$$

Note that limited companies and partnerships differ in that the profits of limited companies are taxed in companies (at 28%), whereas the profits of partnerships are distributed among and taxed in the hands of the partners (at 28%). However, the withdrawal taxation itself is the same. Personal shareholders and partners pay a 28% tax on the non-exempt portion of dividend/withdrawal from the companies. Added to the general tax of 28% on company earnings, this therefore raises the maximum marginal tax on ownership income from 28% before the reform to 48.2% after the reform.

Tax-free rate of return allowance details. Key to the reform was the exemption from the tax, other than ordinary profit tax of 28%, of a return equivalent to the risk-free interest. This allowance was intended to prevent taxation of dividends from raising the costs of funding for Norwegian equity. The allowance was regarded as particularly important for

[1] This appendix leans heavily on on the descriptions in Sørensen (2005), Alstadsæter and Fjærli (2009), Thoresen et al. (2012), and Alstadsæter et al. (2014).

[2] During the transition in 2005, personally held shares could be transferred to a holding company without triggering a capital gains tax.

start-ups and small companies that cannot fund new investment with retained profits, or which have limited access to credit markets or international capital markets.

The annual risk-free rate of return allowance for shareholders/partners (RRA) is computed as the exemption rate multiplied by the sum of the cost price of the share/holding and any unused allowance from previous years. Unused allowance is then carried over to the next year with interest and can be deducted from future dividends and capital gains associated with the same share/holding. The exemption rate is the average interest rate on three-month Treasury bills in the year for which the allowance is to be calculated. The same RRA forms the basis for calculating the allowance for a sole proprietorship.

Pre-reform split rule. Before 2006, capital gains from the realization of shares were taxable at 28%, though the part of capital gains stemming from withheld profits in the firm was tax free. Dividends were tax exempt before 2006 (Alstadsæter et al., 2014). There was one noteworthy modification to this tax exemption. Under the pre-2006 tax regime, owners who worked in their closely-held firms had tax incentives to withdraw income from their firm in the form of tax free dividends instead of labor income. To avoid such income shifting, a so-called "split model" applied to owners with 2/3 or more of shares in the firm they (or their immediate family) worked for. For these owner-workers, a specific and imputed return to real capital could be distributed as tax-free dividends. Any remaining share of corporate profits was taxed as wage income, independent of how it was distributed to the owner. Due to the imputation rule, owner-managers in firms with low capital and/or few employees had incentives to reduce total ownership in the firm (just) below 2/3, inducing firms to have more dispersed ownership. After the removal of this split model on January 1st 2006, this incentive disappeared.

Tax exemptions. The 2006 tax model applies only to dividends from companies resident in Norway or another EEA country. Dividends from companies resident in non-EEA countries were taxable as before, i.e., fully taxable, but with a deduction in Norwegian tax for taxation at source.

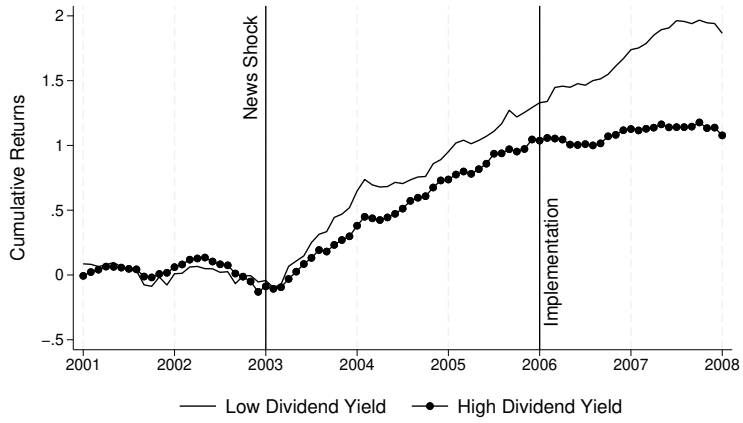
Other details. The system of tax-free inter-corporate dividends and capital gains was maintained to ensure that the tax on capital income would not exceed the tax on labor income. The new 28% tax rate applied to interest, dividends, and realized capital gains, making it more akin to a general capital income tax rather than just a dividend tax. The tax system was neutral regarding dividends and share repurchases both before and after the tax reform. Both payout forms generated the same tax deduction.

D Stock Market Analysis Details

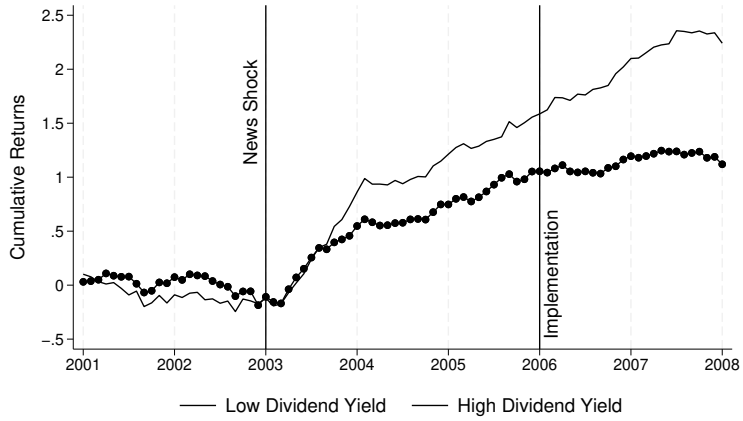
Figure 2 in the main text is constructed in the following way. First, we use monthly data on all publicly traded stocks on the Norwegian stock exchange. The data is comparable to CRSP data for the USA in that it accounts for stock splits and other similar events. We also have data on dividend payouts with the monthly date for the payment, see Ødegaard (2013) for details. Second, following the standard practice in empirical asset pricing, we remove penny stocks and very expensive stocks by dropping stocks with prices less than NOK 1 or greater than NOK 1000. This amounts to roughly the top and bottom 1% of the price distribution. We also drop the top decile of firms by market capitalization in order to focus on a sample that is more comparable to closely-held businesses that we study in the micro analysis. Third, we compute the dividend yield for each publicly traded stock (based on the ISIN number) on the Norwegian stock exchange using annual dividends data up until the reform. The dividend yield is defined as dividends over the price as of December 2002, i.e. prior to the reform news shock, and we partition all stocks in a high and low dividend yield portfolio based on dividend yield being above or below the median. Fourth, we compute portfolio-specific cumulative returns over the period January 2001 - January 2008.

We also perform two robustness exercises. First, we perform a two-way split based on market beta and the dividend yield. We construct the market beta for each stock using monthly returns data until and including December 2002. For each stock, we run an OLS regression of excess returns on the excess return of the Norwegian stock market index. When computing betas, we remove regressions with less than 24 observations (two years). Figure A.1a shows the results, revealing that the pattern in Figure 2 is driven mostly by firms with low market betas. Second, we perform a two-way split on size (market capitalization) and dividend yield. Figure A.1b reports the results, showing that the pattern in Figure 2 is driven by small market-cap firms.

Importantly for our identification strategy, there are no systematic differences in returns before 2003, suggesting once again that the news shock about the future permanent dividend tax reform was not anticipated.



(a) Low beta.



(b) Low market capitalization.

Notes: This figure shows cumulative returns of two-way sorted portfolios where the second dimension of sorting is the market beta in panel (a) and market capitalization in panel (b), respectively. Appendix D presents further details on how the portfolios are constructed.

Figure A.1: Two-way sorted cumulative portfolio returns.

E Imputed Spending without Private Businesses

This section details how we impute spending for households who do not own private businesses. The first challenge when imputing spending is to define income and saving consistent with the budget constraint. We define income as *disposable income*, the sum of labor income, transfers, business income, capital income, and other income (e.g., inheritances and lottery prizes), net of taxes. We define *saving* as the change in net wealth due to either depositing or withdrawing resources from asset classes. Income, as defined above, is directly observed in the tax accounts. The main challenge in imputing spending is to

compute the relevant measure of saving.[3]

The relevant measure of saving using the budget constraint described above is the sum of active depositing or withdrawing of resources into and from all asset classes. The main challenge is that the tax authorities only report total valuations within broad asset classes at the end of the year, and changes in these values could be due to either saving or capital gains. We compute saving within each asset class differently depending on data availability. For nominal assets, such as debt and deposits, saving during the year is directly observed as the change between end-of-year and beginning-of-year values. For housing, we observe housing transactions in the transaction registry, allowing us to observe the relevant saving measure. For stocks, we compute capital gains on household stocks using the stock ownership register after 2005. This register allows us to observe a household's ownership of specific stocks at the end of each year. We combine this ownership information with price changes in individual stocks to compute capital gains. Before 2005, we only observe total wealth in stocks and impute capital gains for households using capital gains rates from the financial accounts. This approach ensures that capital gains are correct on average but will imply that capital gains for any specific household may be wrong. For stock funds, we use the capital gains rate from the financial accounts to impute capital gains for all years in our sample.

Imputed spending ignoring private businesses, spending^{npbo}, is computed as

$$\text{spending}_{i,t}^{\text{npbo}} = \underbrace{\text{disp. income}_{i,t} - \text{saving}_{i,t}^{\text{nominal assets}} - \text{saving}_{i,t}^{\text{housing}}}_{\text{observed}} \underbrace{-\text{saving}_{i,t}^{\text{stocks/stock funds}}}_{\text{unobserved}},$$

where the main source of measurement errors comes from the unobserved component, saving in stocks and stock funds.

[3] An alternative and consistent way of imputing spending is to include capital gains as part of income and define saving as the change in net wealth. In that case, saving would be directly observed and the challenge would be to compute income. In either case, one must compute a measure of *unrealized* capital gains, which is unobserved in the tax data.

F Details and Results for the Alternative Control Group

In this appendix, we present the treatment and control definition, the sample restriction, the summary statistics, and the main results for the alternative control group.

	Control		Treated	
	Mean	S.D.	Mean	S.D.
Panel A: Household Characteristics				
Age	44.41	11.13	49.04	8.94
Panel B: Spending and Income Statement				
Spending	41.70	158.18	63.03	131.11
Disposable income	45.64	32.43	96.75	106.22
Labor income	55.78	39.16	57.23	28.45
Transfers	7.78	11.75	5.00	9.07
Dividend income from private businesses	.	.	94.75	169.03
Taxes	19.91	22.40	28.53	31.59
Panel C: Balance Sheet				
Gross wealth	522.61	384.65	770.72	612.75
Housing wealth	473.24	333.98	605.73	510.09
Deposits	26.28	63.77	74.17	152.30
Public Stocks	4.50	135.13	8.49	50.51
Mutual Funds	5.24	23.97	13.77	48.37
Private Business Wealth	.	.	256.93	567.19
Net Wealth	453.18	357.40	707.57	595.56
Debt	69.43	90.31	63.16	114.26
Panel D: Shares				
Exposure to the reform (dividend share of gross income in %)	.	.	48.69	20.93
Number of individuals	886,126	.	3,549	.

Notes: We define *treated* as having, on average, more than 30% of gross income in the form of dividend income from private businesses in 2000 and 2002. The control group consists of households with no private business wealth. Values in Panel B and C are in 1,000 dollars in 2011.

Table A.1: Descriptive statistics in 2000 with the alternative control group.

Treatment and control definitions. The treatment definition is the same as in our main sample. A household is treated if it satisfies two criteria. First, the mean share of its gross income in 2000 and 2002 - including labor income, transfers, and capital income - attributable to dividends from a private business firm must exceed 30%. Second, we restrict attention to households holding significant ownership shares (greater than 25%) in

their private business ownership portfolio because these owners tend to retain relatively stable ownership shares.

The next step is to define a control group to serve as a counterfactual for the treatment group. In the main analysis of the paper, we use private business owners whose ownership share exceeds 25% but who did not receive any dividend income from their private business ownership in 2000 and 2002. In this alternative control group, we use households with no private business wealth, but we restrict attention to households with wealth above the median within cohort and year.

Additionally, we impose a few minor restrictions on our sample, similar to the main sample. First, we concentrate on individuals between 25 and 65 years old in 2000. Second, we limit our analysis to households with disposable income greater the base amount in the Norwegian social security system (approximately $10,000). Third, we restrict attention to households whose log growth rate of spending did not change by extreme amounts from one year to the next (top/bottom 1%). Our sample consists of 886,126 individuals in the control group and 3,549 in the treatment group in 2000.

Table A.1 presents descriptive statistics in 2000 by treatment status in the alternative sample. Compared with the main sample, the treatment and control groups in this alternative sample are somewhat more different in average age, disposable income, labor income, spending, and net wealth.

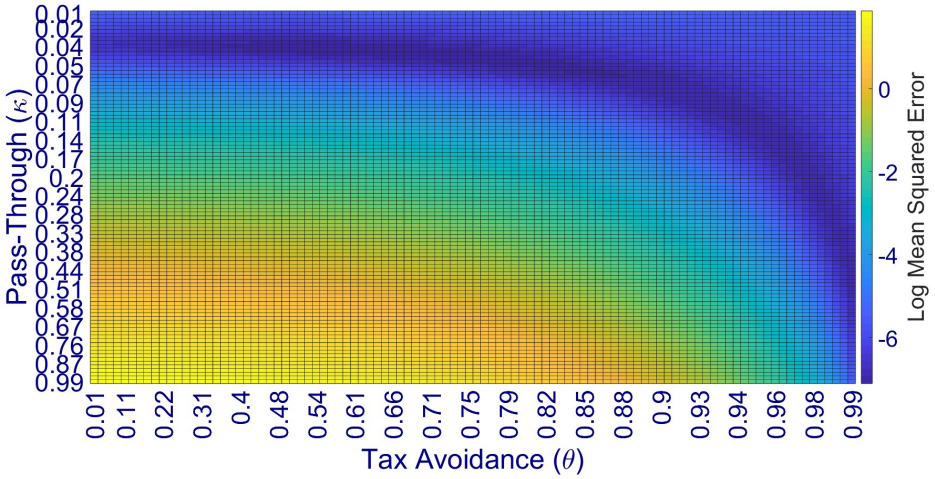
Empirical setup. We estimate the same specification as in the body of the paper, equation (4). The only difference is that the treatment definition $D_{i,2000}$ and the estimation sample differ.

Results. Figure 9 in the paper displays the relative spending response to the dividend tax reform news shock and implementation of the treated households relative to the control group. Similar to the main results in Figure 7, spending increases relative to the control group between announcement and implementation, and declines after implementation. The magnitude of the response is also similar, spending of the treated households increased by about 5% in 2005 and decreased by about 5% on average relative to the control group.

G Numerical Details for the Structural Model

Our impulse response matching approach consists of several steps. First, we take the empirically documented differential response of the treated (capitalists) vs. the control (workers) group. For the baseline case, we use the estimates from Figure 7. For the robustness test with an alternative control group, we use estimates from Figure 9. We calibrate a sub-set of model parameters externally, as reported in Table 2. We then construct a coarse two-dimensional grid for the pass-through parameter κ and the EIS, ψ. The grid for κ is agnostic, ranging from 0.01 to 0.99. The grid for ψ is $[0.1, 4]$, with the lower bound corresponding to the value conjectured by Hall (1988) and the upper bounds being slightly above what is estimated using cross-household differences in after-tax real interest rates (Gruber, 2013). To improve accuracy, all grids are non-linearly spaced, allowing for more points in the region of the parametric space that is most likely to generate low matching errors.

Next, we solve the model for each $\{\kappa, \psi\}$ pair, i.e., for every point on a two-dimensional grid. The grid comprises 100 nodes in each direction. We thus solve the model 10,000 times under different parameter configurations. In every case, we compute and store impulse-response functions to a combination of two news shocks: a positive one-standard deviation shock to ε_d and a one-standard deviation negative shock to ε_l. In particular, we are interested in the model-implied estimates of the consumption response of capitalists less the consumption response of workers following the news shocks. This allows us to construct the homogeneous-incidence Jacobian matrix $\mathcal{J}^{HI}$. Finally, we construct the heterogeneous-incidence Jacobian $\mathcal{J}$. To this end, we first build a grid for the short-run tax avoidance parameter θ, ranging from 0.01 to 0.99 using 100 non-linearly spaced nodes. For each value of θ on this grid, we compute and store a new $\mathcal{J}$. This completes the first step of our approach.



Notes: A heatmap of IRF matching errors produced by the calibration procedure over the three-dimensional grid $\{\psi, \kappa, \theta\}$ with tax avoidance intensity and reform pass-through on the horizontal and vertical axes, respectively. Colder colors correspond to lower mean squared errors.

Figure A.2: Impulse response matching results.

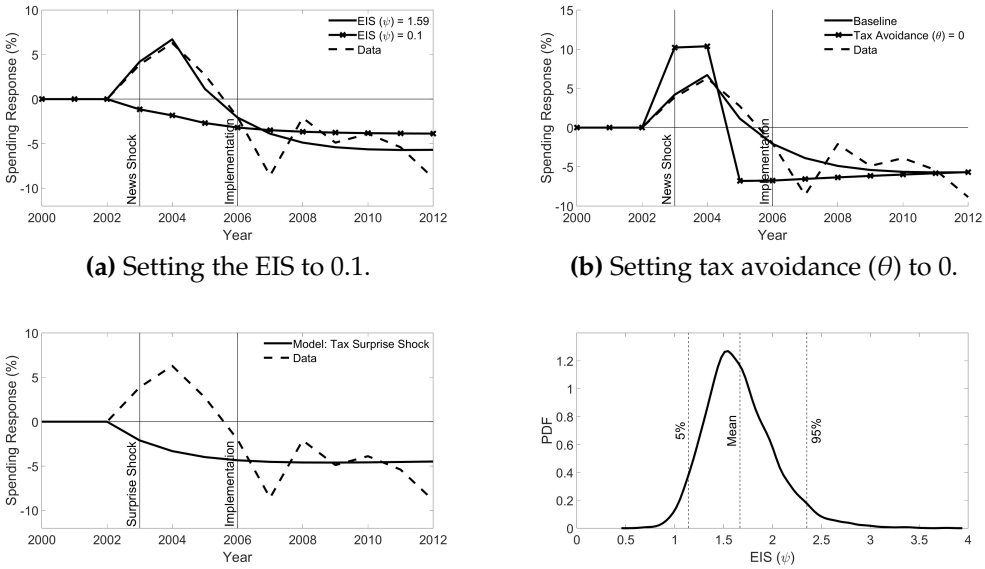
In the second step of our procedure, we locate the point on the grid that minimizes the distance between the empirical and the model-implied spending differentials. In other words, we identify the values of parameters that "match" the empirical impulse responses as closely as feasible. Our candidate model-based sequences are stored in a $10 \times 100 \times 100 \times 100$ array, corresponding to the three-dimensional grid $\{\kappa, \psi, \theta\}$ with ten rows (years). Our target is the full empirical relative spending sequence over 2003-2012, which we denote with $\mathcal{J}^{DATA}$. For each candidate model-based sequence indexed by z, we compute a mean squared error $\mathcal{E}_z$ relative to $\mathcal{J}^{DATA}$, defined as $\mathcal{E}_z = \mathbb{E}\Big[\Big(\mathcal{J}_z - \mathcal{J}^{DATA}\Big)^2\Big]$.

Finally, we identify the index $\bar{z}$ that corresponds to the truple $\{\bar{\kappa}, \bar{\psi}, \bar{\theta}\}$ which produced the model-based sequence $\mathcal{J}_{\bar{z}}$ with the lowest $\mathcal{E}_{\bar{z}}$. The final product of this exercise is thus $\bar{\psi}$, the value of the identified EIS that corresponds to the global minimal IRF matching error.

Figure A.2 plots the outcome of the IRF matching exercise. We present a heatmap with θ and κ on the x-axis and y-axis, respectively. Each colored square on the map represents a (log) mean-square error $\mathcal{E}$ of the corresponding combination of parameters. We see that $\mathcal{E}$ declines as tax avoidance intensity rises and pass-through falls, i.e., the northeastern region is where the best-fitting combinations of parameters are. In fact, there is a clearly visible dark-blue patch that showcases the global minimum area. This area corresponds to the values of the EIS that are generally in the [1.2,1.8] interval with the baseline value of 1.59 being the estimate that produces the globally lowest $\mathcal{E}$.

H Additional Model Results

In this section, we present four additional model results that complement the main text.



(a) Setting the EIS to 0.1.

(b) Setting tax avoidance (θ) to 0.

(c) The impact of surprise tax shocks.

(d) Bootstrap density.

Notes: Panels (a) and (b) plot relative spending responses implied by alternative parametrizations in which we set $\psi = 0.1$ and $\theta = 0$, respectively. Panel (c) reports relative spending responses after a surprise tax shock, rather than a news shock. Panel (d) plots the density of EIS estimates from the model bootstrap procedure as described in the text.

Figure A.3: Additional model results.

First, we test whether picking a wrong model with high matching errors produces empirically inconsistent results. Figure A.3a shows the result when we counterfactually set the EIS (ψ) to 0.1, a very different value from what our calibration suggests. The spending response in this case corresponds to a mean squared error that is at least an order of magnitude above the minimum and such a low EIS fails to match the front-loading of spending before implementation.

Second, Figure A.3b presents relative spending responses from a model with θ equal to 0, corresponding to no tax avoidance in the short run. This shows that the tax avoidance friction does not impact our results qualitatively but is crucial to quantitatively match the relatively smooth response to the tax news shock. Third, in Figure A.3c, we illustrate the importance of the news component by showing the relative spending response to a *surprise* shock to dividend and labor income taxes *without a pre-announcement*. Without the news component, the model is not able to match the front-loading of spending before implementation. Finally, Figure A.3d reports the density of EIS estimates produced by the bootstrapping procedure with 10,000 draws applied to the baseline model. The mean, 5th percentile, and 95th percentile values are 1.67, 1.14, and 2.35, respectively.

References

Alstadsæter, Annette and Erik Fjærli, "Neutral taxation of shareholder income? Corporate responses to an announced dividend tax," *International Tax and Public Finance*, 2009, *16* (4), 571–604.

_ , Wojciech Kopczuk, and Kjetil Telle, "Are Closely Held Firms Tax Shelters?," *Tax Policy and the Economy*, 2014, *28* (1), 1–32.

Flynn, Joel P, Lawrence DW Schmidt, and Alexis Akira Toda, "Robust comparative statics for the elasticity of intertemporal substitution," *Theoretical Economics*, 2023, *18* (1), 231–265.

Gruber, Jonathan, "A Tax-Based Estimate of the Elasticity of Intertemporal Substitution.," *Quarterly Journal of Finance*, 2013, *3* (1), 1 – 20.

Hall, Robert E, "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 1988, *96* (2), 339–357.

Ødegaard, Bernt Arne, "Empirics of the Oslo Stock Exchange. Basic, Descriptive, results 1980-2020," *Working Paper*, 2013.

Sørensen, Peter Birch, "Neutral taxation of shareholder income," *International Tax and Public Finance*, 2005, *12* (6), 777–801.

Straub, Ludwig, "Consumption, savings, and the distribution of permanent income," *Unpublished manuscript, Harvard University*, 2019.

Thoresen, Thor O, Erlend E Bø, Erik Fjærli, and Elin Halvorsen, "A suggestion for evaluating the redistributional effects of tax changes: With an application to the 2006 Norwegian tax reform," *Public Finance Review*, 2012, *40* (3), 303–338.