

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

Martin B. Holm Rustam Jamilov Marek Jasinski Plamen Nenov

March 2024

Abstract

This paper studies the spending response to news about a dividend tax reform in order to estimate the elasticity of intertemporal substitution (EIS). The Norwegian dividend tax reform was proposed in 2003, announced in 2004, implemented in 2006, and raised the dividend tax rate by 28 percentage points. We compare the spending responses of exposed households with a high share of dividends to income before the reform to a control group. Exposed households responded to the reform by *increasing* spending after the news and reducing spending after implementation. We interpret our findings using a capitalist-worker framework with dividend tax news shocks. The model can replicate the spending response to the dividend tax news only if the EIS is greater than one, with a baseline estimate of around 2.

[†]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 Andreas Fagereng, Andrea Ferrero, Joseba Martinez, Michael McMahon, Gisle Natvik, Alan Olivi, Franck Portier, H el ene Rey, Federica Romei, Andreas Schaab, Alp Simsek, Stefanie Stantcheva, Paolo Surico, Ludwig Straub 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  yvind Norli for providing us with the Oslo B rs 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).

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 examine 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.

Using a dynamic difference-in-differences methodology, we identify the spending effects of dividend tax news and the subsequent dividend tax implementation.² Specifically,

¹See, 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.

²The idea of signing the value of EIS relative to unity by looking at the relative spending response of more

we compare the spending trajectories of a treatment group, defined as prime-age households with a relatively large share of dividend income out of gross income (labor income plus capital income) before the reform, relative to a control group of wealthy households 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 to 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 5% more in 2004 and 2005, followed by a gradual but persistent decline of around 8%. 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. Firms in the model 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 avoidance 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 between reform news and its implementation (2003-2005), as well as the estimated average relative spending response post-implementation (2006-2013). Our impulse-response matching exercise returns a value for the EIS of 2.06. In order to account for parameter uncertainty, we compute confidence bands via bootstrapping and find that the model-implied EIS is statistically different from unity at the 95% confidence level. The result of the EIS being

vs. less exposed households to news about future rates of return was first suggested in recent theoretical work by Schmidt and Toda (2019) and Flynn et al. (2022).

greater than unity is confirmed in a number of robustness tests and sensitivity checks, for example changing the labor market structure or the share of workers in the economy.

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 returns 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., 2022](#)), 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](#)). 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 (DARA). We therefore abstract away from modeling labor income risk, retaining a tractable yet realistic model of the consumption-saving behavior of capital owners.

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 ([Güvenen, 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.³ 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 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

³One 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.

EIS is heterogeneous across the population and increasing in wealth, and the relevant EIS is policy-dependent, partly because the treatment group varies across different policies. In sum, 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.⁴ 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 vs. disaggregated cross-sectional vs. panel data, reduced-form vs. 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 mis-specification of the structural equations used.⁵

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 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

⁴See, 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.

⁵See the discussions in Bansal et al. (2010), Bansal et al. (2011), Gruber (2013), and Schmidt and Toda (2019) 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.

the announcement and implementation, specifically emphasizing the importance of *news* of future tax changes for identification.

The large implied values for the EIS that we find are in stark contrast with other recent empirical studies that similarly to us use rich administrative level 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 that they use for interpreting their reduced-form estimates 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 to also back out 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, such as firm owners, who rely more on dividend 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.⁶

We also contribute to a large and growing literature on the effects of capital income taxation, particularly dividend income taxation.⁷ 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) 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

⁶See [Iachan et al. \(2021\)](#) for a further discussion of this point.

⁷See, 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. \(2024\)](#).

calibration.⁸

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 a fraction of households are stockholders with access to financial markets while remaining households are restricted from firm ownership (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).⁹

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. (2022). Consider an agent living for three periods, solving the following problem:

$$\begin{aligned} \max_{C_0; C_1; C_2} \quad & u(C_1) + u(C_2) + u(C_3) \\ \text{s.t:} \quad & C_0 + \frac{C_1}{R_1} + \frac{C_2}{R_1 R_2} = A_0; \end{aligned}$$

where A_0 is initial financial wealth and R is the (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 away from any portfolio choice (see, e.g., Iachan et al., 2021). Assume further that period utility is $u(C) = \frac{C^{1-\sigma} - 1}{1-\sigma}$ ($\ln(C)$ if $\sigma = 1$) where $\sigma > 0$ denotes the EIS.

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

⁸In 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.

⁹See 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 .

respect to R_2 are

$$\frac{\partial C_0}{\partial R_2} = -(\alpha - 1) \phi_0(A_0; R_1; R_2); \quad \frac{\partial C_1}{\partial R_2} = -(\alpha - 1) \phi_1(A_0; R_1; R_2); \quad \frac{\partial C_2}{\partial R_2} = \phi_2(A_0; R_1; R_2);$$

where ϕ_0 , ϕ_1 , and ϕ_2 are functions of A_0 , R_1 , and R_2 (Appendix A includes the details).

Hence, C_0 and C_1 are strictly decreasing in R_2 if $\alpha > 1$, while C_2 is always strictly increasing in R_2 . Intuitively, a higher $t = 2$ return induces the same $t = 0$ income and substitution effects as a contemporaneous increase in the interest rate. 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 reaction between the news shock and implementation of the reform can allow us to identify whether the EIS is greater or less than one. Note that while the model we present here is very stylized, these comparative statics also hold in richer models with multiple assets, returns risk, and labor income (Flynn et al., 2022).

3 Institutional Background

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 was taxed at very different rates. Wage earners faced a progressive income tax, with marginal tax rates from around 35% to 64.7%.¹⁰ Capital owners, on the other hand, faced a flat profit tax of 28%.¹¹ 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 (*NOU 2003:9, 2003*, p. 11).

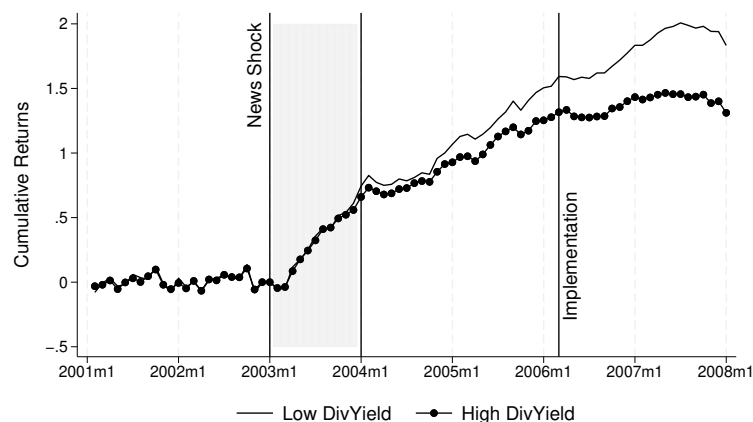
On February 6, 2003, the government-appointed commission published their findings (*NOU 2003:9, 2003*). Among the key recommendations was the implementation of a 28% dividend tax, raising the effective tax rate on firm owners from 28% to 48.2%. At the same time, the commission recommended reducing the 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*). 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

¹⁰This marginal tax rate included the employer's national insurance contributions (“arbeidsgiveravgift”).

¹¹Additionally, there was a wealth tax with a marginal tax rate of 1.1% of net (taxable) wealth above NOK580,000 during this period.



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 C presents details on how the portfolios are constructed.

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

returns in any given year.¹² 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.¹³ 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.¹⁴


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 2003-2004 news shock period 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-


¹²The 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.

¹³During the transition in 2005, personally held shares could be transferred to a holding company without triggering a capital gains tax.

¹⁴See Appendix B for further details on the tax reform.



(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.

median (high) and the other with below-median (low) dividend yield stocks, with a base month of January 2003.¹⁵ 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 2003 that the returns of the two portfolios began to diverge.¹⁶ Divergence accelerated noticeably during the transition period and after the implementation date. This finding suggests that the dividend tax reform was not anticipated by the market during the news shock period.

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 immediately clear

¹⁵See Appendix C for details on portfolio construction.

¹⁶Appendix C 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.

from the aggregate data whether spending 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 and loans to firms increased. Both saving in unlisted shares and corporate loans are examples of ways to shift the dividend tax burden across time.¹⁷ 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 2013,¹⁸ combined using unique personal identification numbers. Because Norway levies both a wealth tax and an income tax, the tax authorities collect information on household 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

¹⁷Alstadsæter and Fjærli (2009) document a similar pattern by looking at firm balance sheet data.

¹⁸Our sample starts in 2000 because the accounting framework used by firms in Norway changed with the Accounting Act (LOV-1998-07-17-56) from January 1st, 1999. Because we need to observe the changes in paid-in-capital to impute consumption of firm owners in our treated group, the first observation of imputed spending for firm owners was in 2000. Moreover, our sample stops in 2013 because the government abolished the inheritance tax from 2014, resulting in a large increase in inter-vivo transfers among wealthier households and dynasties. Because these transfers no longer need to be reported for tax reasons, and these transfers were common among the wealthiest households, imputed spending for wealthy firm owners is imprecise after 2014.

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*. The main idea of imputed spending is to define spending 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 only defined at the household level. Moreover, we note that imputed spending is all spending by households on items that do not enter their balance sheet, including durable components such as housing refurbishment and household appliances, but also non-durable spending. Appendix D 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 controlling positions in the private business (above 50%) because these majority owners directly control the flow of resources from the firm to the owner. 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

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

where $\text{spending}_{i,t}^{\text{npbo}}$ is imputed spending for non-private business owners, defined in Appendix D. 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

$$\underbrace{\text{profits}_{i,t}}_{\{Z\}} - \underbrace{\left(\text{book value}_{i,t} - \text{capital gains}_{i,t} \right)}_{\text{dividends}_{i,t} + \text{retained earnings}_{i,t}} = \underbrace{\text{dividends}_{i,t}}_{\{Z\}} - \underbrace{\text{paid-in capital}_{i,t}}_{\text{paid-in capital}_{i,t}} \quad (2)$$

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

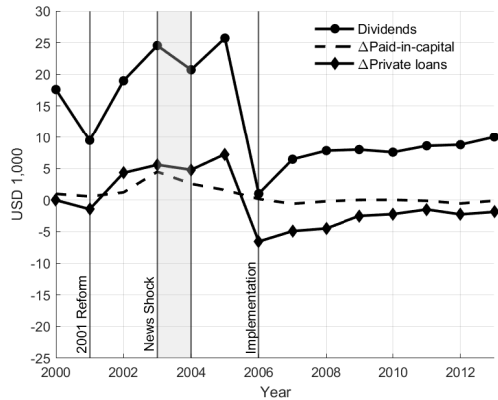
$$\text{spending}_{i,t} = \text{spending}_{i,t}^{\text{npbo}} + \text{dividends from the firm}_{i,t} - \text{paid-in capital in firm}_{i,t} \quad (3)$$

where “dividend from the firm” and “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.

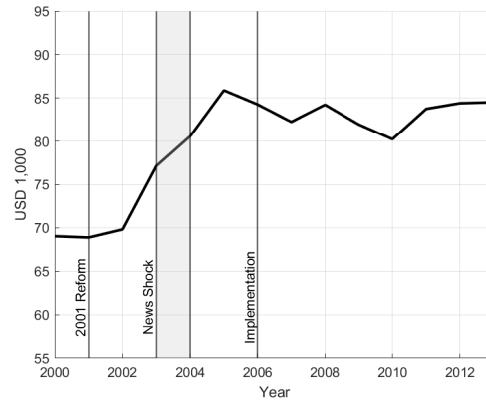
Our imputation procedure makes one additional underlying assumption: we assume that the ownership share is stable from one year to the next. This is necessary for two reasons. First, it allows us to use the definition of spending above without any adjustments for changes in ownership shares. Second, 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 of majority owners that mostly retain a stable ownership share from one year to the next.

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-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



(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

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 around \$25,000 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 businessperson 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.¹⁹

Second, our sample period starts in 2000, but the ownership registry starts in 2004. We therefore impute ownership shares in the four years between 2000 and 2003. We restrict attention to majority owners in 2004 (more than 50% 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.

¹⁹Because we focus on owners who exclusively hold majority ownership stakes, almost all private businesses in our sample are directly held.


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 may be that households respond to the dividend tax reform by moving wealth abroad to hide it. In that case, tax evasion will show up as a large spike in imputed spending as it involves assets disappearing from the household’s balance sheet. To limit the potential impact of such cases, we restrict attention to households whose imputed spending does not increase by more than 800% or decrease by more than 80% from one year to the next.²⁰

5 Empirical Setup

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 households exposed to the reform, the *treated* households, 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 and discuss threats to identification.

²⁰A concern is that the pattern we observe of heightened spending followed by a permanent decrease may be due to hiding wealth (resulting in a spending spike) and spending from this hidden wealth afterwards (resulting in a permanent reduction in spending). However, if these movements reflect somewhat realistic pricing of the hidden wealth, meaning that it should approximately equal the net present value of future spending, the increase in spending before the implementation is at least an order of magnitude too small to account for the observed permanent reduction in spending.



Notes: The figure shows the distribution of dividends from private business as a share of gross income among (majority) private business owners in 2000.

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

Treatment and control definitions. A household is considered 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 20%. 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 alluded to previously, we restrict attention to households holding only controlling ownership shares (more than 50%) in their private business ownership portfolio because we cannot impute spending of owners who hold minority ownership shares in private businesses. Figure 5 displays dividend income as a share of gross income among private business owners in our sample.²¹

The next step is to define a control group to serve as a counterfactual for the treatment group. The private business owners in the treatment group are wealthier than the general population. To define a relevant control group we therefore select households in the non-private business owner sample that hold more wealth than the median within their cohort.

²¹We choose the treatment threshold to ensure that dividends should make up a significant fraction of treated households income. In Section 7, we show that our results are robust to using different thresholds as well as to using a continuous treatment variable.

	Control		Treated	
	Mean	S.D.	Mean	S.D.
<i>Panel A: Household Characteristics</i>				
Age	45.49	11.14	50.71	8.71
<i>Panel B: Spending and Income Statement</i>				
Spending	43.68	38.45	86.07	145.04
Disposable income	39.38	18.22	94.67	142.8
Labor income	46.55	28.81	56.51	29.52
Transfers	8.84	11.67	4.49	8.39
Dividend income from private businesses	.	.	57.33	131.31
Taxes	15.07	10.98	26.75	22.68
<i>Panel C: Balance Sheet</i>				
Gross wealth	310.96	276.85	695.31	658.06
Housing wealth	283.62	269.38	548.34	543.89
Deposits	16.38	34.15	63.75	148.9
Public Stocks	0.32	20.02	8.22	56.1
Mutual Funds	2.21	9.22	11.92	36.84
Private Business Wealth	.	.	156.37	401.22
Net Wealth	410.09	294.83	715.72	663.84
Debt	59.42	60.75	71.04	117.09
<i>Panel D: Shares</i>				
Exposure to the reform (dividend share of gross income in %)	.	.	36.68	21.69
Number of individuals	1,320,970	.	2,959	.

Notes: We define *treated* as having, on average, more than 20% of gross income in the form of dividend income from private businesses in 2000 and 2002. Values in Panel B and C are in 1,000 dollars in 2011.

Table 1: Descriptive statistics in 2000.

Similar to the private business owners, these households are not likely to be financially constrained and thus can smooth consumption over time.

Additionally, we impose a few minor restrictions on our sample, which spans annual observations from 2000 to 2013. First, we concentrate on prime-age households between 25 and 65 years old in 2000. Second, we limit our analysis to households with disposable income and imputed spending greater than the base amount in the Norwegian social security system, approximately \$10,000. Third, as discussed above, we restrict attention to households above the median in the within-cohort wealth distribution in 2000 whose spending does not increase by 800% or decrease by 80% from one year to the next. Fourth, for the control group, we select households with no dividend income in the year 2000. Our final sample consists of 1,320,970 households in the control groups in 2000 and 2,959 in the

treatment group.


Table 1 presents descriptive statistics in 2000 by treatment status in our sample. Treated households spend more, earn more, are older and more educated, and are wealthier than the control group. 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 more than 35% of gross income.

The estimation framework. 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 rest of the population 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 year and estimate the following dynamic difference-in-differences model

$$c_{i;2000+h} - c_{i;2000} = \sum_{h=1}^H \beta_h D_{i;2000} \times \mathbb{1}_{2000+h} + \sum_{h=1}^H \gamma'_h \mathbf{X}_{i;2000} \times \mathbb{1}_{2000+h} + \epsilon_{i;h} \quad (4)$$

for $h = \{0; 1; 2; \dots; 13\}$, where $c_{i;2000+h}$ denotes log imputed spending for individual i in year $2000 + h$, $D_{i;2000} \in \{0; 1\}$ is our treatment variable for household i in year 2000, $\mathbb{1}_t$ 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 are clustered at the individual level to address that the error terms $\epsilon_{i;h}$ may potentially be correlated within households across time.

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.



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

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

Though not necessary for our identifying assumption, one weakness of our setup is the lack of balancing between the two groups. The treated group, being predominately private business owners, are older, have higher income, more spending, and are wealthier. 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 taxes as in [Thoresen et al. \(2010, Figure 1\)](#).²²

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,

²²Specifically, we choose the following three cut-offs in 2000 based on the 2006-NOK marginal tax groups: 394,000, 750,000, 936,800.

that would tend to bias our results. We therefore include a full set of age controls in our 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.

6 Empirical Results

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

Main results. Figure 7 displays the relative spending response to the dividend tax reform news shock and implementation of the treated households relative to the control group. Between the announcement and implementation of the reform, spending by treated households increased by around 5% relative to the control group. After the implementation of the reform in 2006, relative spending gradually decreased. In 2011-13, relative spending by treated households was about 8% lower than the control group. 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.

A gradual decline in relative spending characterizes the period between 2006 and 2011. According to the theoretical framework in Section 2, we expect a sharp decline in spending when the new dividend tax reform is in place. However, as discussed above, there was a delay between the reform's announcement and implementation, which allowed owners to engage in intertemporal shifting of the tax burden by taking out tax-free dividends and putting the resources back into the firm through private loans and paid-in capital. Hence, even though the reform was implemented in 2006, owners did not fully face the new dividend tax before they had depleted these vehicles of intertemporal tax shifting. The smooth decline from 2006 to 2011 is consistent with the gradual depletion of such




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.


vehicles of intertemporal tax shifting, where more and more owners are fully affected by the dividend tax reform as time passes. This discussion motivates our choice of introducing heterogeneous tax avoidance into the transition dynamics of our structural model in Section 7.

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 that concern, we control for a number of plausible confounders in our baseline specification, as explained above. Figure 9a displays the results for two alternative specifications, the first where we include fewer controls (only dummies for year and age) and a second where we include additional controls for the risky share of financial wealth and gross wealth. These additional controls account for the potential differential exposure of the treated and control groups to contemporaneous movements in financial markets during the same period. Adding these two additional controls only marginally affects our results.


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



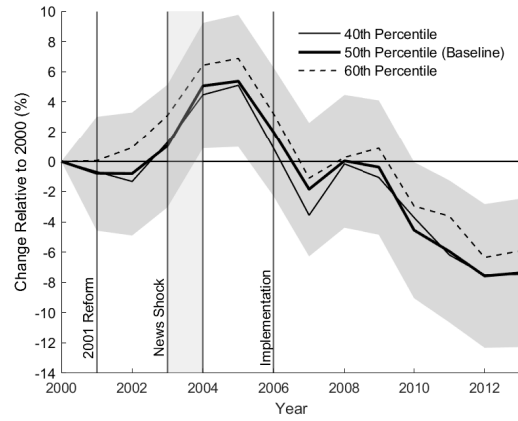
(a) Alternative sets of controls.



(b) Continuous treatment definition.



(c) Alternative treatment thresholds.



(d) Alternative sample restrictions.

Notes: The figures display the estimated coefficients of equation (4) with 95% confidence bands computed using standard errors clustered at the individual level. The continuous treatment coefficients in Panel (b) are estimated using equation (5). The continuous treatment coefficients are adjusted to be interpreted as the spending effect in % of a 50 pp. increase in the dividend share of gross income. Panel (d) shows results using alternative cut-offs in the (within-cohort) wealth distribution in the sample restriction.

Figure 8: Robustness of the empirical results.

specification we use is

$$C_{i;2000+h} - C_{i;2000} = \sum_{h=\underline{h}}^{\overline{h}} \alpha_h \text{exposure}_{i;2000} \times I_{2000+h} + \sum_{h=\underline{h}}^{\overline{h}} \beta'_h \mathbf{X}_{i;2000} \times I_{2000+h} + \epsilon_{i,h} \quad (5)$$

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 by 50 percentage points. Figure 9b 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 exposure from a zero dividend share to a 50% dividend share of gross income.

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 20% in 2000 and 2002. The choice of this threshold represents a trade-off between having a treatment group that is exposed and a treatment group that is relatively large. Figure 9c displays results when using some alternative thresholds (15% and 25%), showing that the results do not change materially.

Fourth, in our sample restrictions we focus on the relatively wealthy households in order to ensure that the control group is relevant for the counterfactual evolution of the treatment group. Thus, we restricted our sample to households in the top 50% of wealth within their cohort (i.e., above the median). Figure 9d therefore presents two alternative sample restrictions, at 40% and 60% with similar results as the baseline.

Taking stock. We find that news about a future permanent dividend tax reform caused a positive response in the relative spending of households with high dividend income intensity. Implementation of the reform, in turn, caused a persistent negative decline in the relative 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 also illustrated in Section 2 by our simple model. To pinpoint the exact value of the implied EIS and theoretically account for various partial and general equilibrium confounding channels, we now turn to our structural general equilibrium framework that we eventually bring to the data.

7 A Model of Capitalists and Workers

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

7.1 Preferences

There is a continuum of households of measure one. A fraction α of households are *workers* indexed by w . The remaining fraction $1 - \alpha$ 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. Capitalists can save and smooth consumption by purchasing claims on firms, which in turn produce capital. Workers, on the other hand, smooth consumption by accessing a risk-free bond and are restricted from firm ownership. This modelling approach is similar to the “stockholder and non-stockholder” setup of [Mankiw and Zeldes \(1991\)](#) and [Guvenen \(2009\)](#).

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

$$\max E_t \sum_{j=0}^{\infty} \beta^j U(C_{k;t+j}) \quad \text{s.t.} \quad S_{k;t+1}P_t + C_{k;t} \leq (1 - \tau_{k;t})N_k + S_{k;t}(D_t + P_t) + T_{k;t} \quad (6)$$

The period utility function is of the CRRA form: $U(C_{k;t}) = \frac{C_{k;t}^{1-\alpha}}{1-\alpha}$, if $\alpha \neq 1$, and $U(C_{k;t}) = \ln C_{k;t}$, if $\alpha = 1$. $\beta \in (0; 1)$ 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:

$$1 = E_t \beta \frac{C_{k;t+1}^{-\alpha}}{C_{k;t}^{-\alpha}} \frac{D_{t+1} + P_{t+1}}{P_t} \quad (7)$$

Of crucial interest to us is the parameter α , the EIS. Absence of arbitrage implies existence of a unique stochastic discount factor $\beta_{t;t+1}$ that prices all assets in the economy. $\beta_{t;t+1}$ is defined in terms of consumption of the capitalists, because they own all firms and are the marginal investor: $\beta_{t;t+1} \equiv \frac{C_{k;t+1}^{-\alpha}}{C_{k;t}^{-\alpha}}$.²³ Together with the usual transversality condition, by forward substitution the pricing equation for shares can be obtained as: $P_t =$

²³We 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 8 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 $\frac{C_{t+1}}{C_t}^{-\alpha}$. With recursive preferences

$E_t \sum_{j=0}^{\infty} \beta^j D_j$. One then obtains the intertemporal budget constraint: $E_t \sum_{j=0}^{\infty} \beta^j C_{k,t+j} \leq P_t + E_t \sum_{j=0}^{\infty} \beta^j (1 - \delta_{k,t+j}) N_k$.

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

$$\max E_t \sum_{j=0}^{\infty} \beta^j U(C_{w,t+j}) \quad \text{s.t.} \quad C_{w,t} + \frac{B_{w,t+1}}{R_t^B} = (1 - \tau_{w,t}) N_w + B_{w,t} + T_{w,t}; \quad (8)$$

where $U(C_{w,t})$ is the same CRRA period utility function as that of capitalists and $T_{w,t}$ is the lump-sum tax or subsidy that is paid out to the workers. The Euler equation for the worker is:

$$1 = R_t^B E_t \left[\frac{C_{w,t+1}}{C_{w,t}} \right]^{-1} \beta; \quad (9)$$

Notice how through $\tau_{w,t}$ we allow the labor tax levied on workers to be different from $\tau_{k,t}$.

7.2 Technology

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

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

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 (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

$$\tau_{d,t} D_t = D_t(1 + \tau_{d,t}); \quad (11)$$

where $\tau_{d,t} \in [0; 1]$ is a parameter that governs (in equilibrium) the elasticity of the rate of

the SDF becomes $\frac{C_{t+1}}{C_t}^{-1} = \frac{V_{t+1}^{1-\gamma}}{E_t V_{t+1}^{1-\gamma}}$ with V_{t+1} the continuation value of the household's problem, γ the parameter that controls risk aversion, and $\beta = \frac{1}{1+\gamma}$. In the absence of uncertainty, the last term of the pricing kernel reduces to unity and we recover the CRRA case.

return on saving for capitalists to dividend tax rate shocks.²⁴ The parameter α 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, α 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, $\alpha = 1$ 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, $\alpha = 0$, 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 from the data.

The optimization problem of firms can be written recursively as:

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

where $V(K)$ is the market value of the firm and m' is the stochastic discount factor, equal to the stochastic discount factor of the capitalists β' . The first-order condition with respect to K' is:

$$E m' \frac{\partial D(D)}{\partial D(D')} [1 - \alpha + F_k(A; K'; N')] = 1: \quad (12)$$

| $\{ z \}$
Net Return on Investment R_t

7.3 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 is rebated back lump sum to the capitalist and any proportional labor tax gets rebated back to the respective agent:

$$\begin{aligned} T_{k;t} &= \tau_{d;t} D_t + \tau_{k;t} N_k \\ T_{w;t} &= \tau_{w;t} N_w \end{aligned} \quad (13)$$

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.

²⁴A 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.

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: $S_t = S_{t+1} = \frac{1}{1-\alpha}$. Labor market clearing in our environment is trivial and equates the weighted average of endowments to the production function input N : $N = \alpha N_w + (1 - \alpha)N_k$. Similarly, aggregate consumption is determined by: $C_t = \alpha C_{w,t} + (1 - \alpha)C_{k,t}$. 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 three policy instruments: the tax on dividends τ_d , the tax on capitalists' labor endowment τ_k , and the tax on workers' labor endowment τ_w . These instruments are assumed to follow the following exogenous stochastic processes:

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

where $\epsilon_{d,t-j}$, $\epsilon_{k,t-j}$, $\epsilon_{w,t-j}$ are drawn from $\mathcal{N}(0, 1)$.²⁵ The stochastic processes capture the expectation phenomenon of tax news shocks: an announcement at time $t-j$, captured by an innovation to $\epsilon_{d,t-j}$, represents a credible signal that the tax rate $\tau_{d,t}$ will change at t . Given the context of the Norwegian tax reform, we set the news lag to $j=3$ because the model is calibrated at an annual frequency.

Definition 1. A rational expectations general equilibrium, given tax policy innovation shocks $\{\epsilon_{d,t}, \epsilon_{k,t}, \epsilon_{w,t}\}$ 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, $m' = \frac{U_c(c'_k)}{U_c(c_k)}$, and all markets clear at any given time t .

7.4 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 search for tax avoidance opportunities, which are available at the beginning, i.e. upon arrival of the news, but disappear every period with an exogenous and i.i.d. probability $1 - \delta$. This represents the gradual closure of any tax loopholes by the

²⁵The 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.

relevant authorities. As a result, following a tax news shock, a fraction α of households is able to avoid the full tax burden while the remaining partition 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 α .²⁶ However, in the short run the incidence of dividend taxes is distributed heterogeneously.

It is useful to follow [Auclert et al. \(2021\)](#) and define several sequence-space objects of interest. First, define $\{k_s\}_{s=0}^{\infty}$ as the dividend tax *input sequence* that agents take as given. Denote with $C_{it} = C_{it}(\{k_s\}_{s=0}^{\infty})$ the spending *output function* for agent i , which stands either for the capitalist or the worker. 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 the household of type i . Each entry of \mathcal{J}_i^{HI} satisfies: $\mathcal{J}_{i,t;s}^{HI} = \frac{\partial C_{it}}{\partial k_s}$. 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 $\alpha = 0$.

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 the following way, focusing on the relevant case of $s > 0$ because the tax reform is announced in advance:

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

Finally, the differential response in spending between the capitalists and the workers can be readily computed as the column-difference between $\mathcal{J}_{k;t;s}$ and $\mathcal{J}_{w;t;s}$ for all t and a shock at some \hat{s} . Parameter α , 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.

8 Identifying the Elasticity of Intertemporal Substitution

Our modeling approach facilitates an exact comparison with our empirical strategy. First, capitalists and workers represent the treated and the control groups, respectively. Second, a dividend tax news shock in the model corresponds to the year 2003, and the actual rate hike to the 2006 implementation. 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

²⁶Our 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
	0.9987	Share of workers
	0.9800	Discount factor
	0.0750	Depreciation rate
	0.3300	Capital share
N	0.3000	Labor endowment
A	1.0000	Productivity
d	0.2800	St. dev., capitalist dividend tax news
l	0.1040	St. dev., capitalist labor tax news

Table 2: Model parameters fixed externally

describe which parameters we calibrate externally. Next, we show how we leverage our structural model to identify the elasticity of intertemporal substitution.


8.1 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 share of workers in the economy is $n = 0.9987$. This value corresponds to the share of the control group in our empirical analysis (see Table 1). 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 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.

8.2 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, e.g., 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 τ , the EIS σ , and heterogeneous tax incidence η . 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 $\hat{\tau}$: the value of



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 9: Model Meets Data

the EIS that corresponds to the global minimal IRF matching error. Further details on the computational procedure are provided in Appendix E.

8.3 Main Result

Figure 9 shows the best-fitting model-generated consumption response of capitalists less the consumption response of workers following the news shocks $\{d_{t,i}, k_t\}$ in $t = 2003$ and implementation in $t = 2006$. The implied $\bar{\gamma}$ is 2.06, 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 2.06 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 $\bar{\alpha}$ and $\bar{\beta}$ are 0.02 and 0.90, 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.

8.4 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 to be wide. To account for the coefficient uncertainty for the implied EIS, we find the implied EIS using 10,000 independent draws from the estimated spending response coefficients. The average EIS is 2.2, and the 95% confidence band is [1:35; 3:08].²⁷ We repeat the same bootstrapping procedure for every extension and sensitivity test below.

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: $W_t = A(1 - \alpha)K_t^\alpha N_t^{1-\alpha}$.

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: $(1 - \alpha_x)W_t = N_{x,t}$. Both types have the same Frisch elasticity $\frac{1}{\alpha_x}$. The labor

²⁷The average estimate turns out to be slightly greater than the baseline of 2.06 due to some (mild) non-linearities.


disutility parameter α 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: $(1 - \alpha)W_x C_{x,t}^{-1-\alpha} = N_{x,t}$. We fix the EIS of workers to unity and allow α to be determined by the impulse response matching exercise as before. Thus, this particular extension allows for the EIS to be heterogeneous across the two population types. Labor disutility is again re-adjusted such that hours equal 0.3 in the steady state.


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 $G_t = d_{t,t}D_t + k_{t,t}N_{k,t} + w_{t,t}N_{w,t}$. 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: $\frac{1}{R_t^F}B_{t+1} + d_{t,t}D_t + k_{t,t}N_{k,t} + w_{t,t}N_{w,t} = B_t$. In equilibrium, the risk-free rate is pinned down by the stochastic discount factor: $R_t^F = \frac{1}{E_t m_{t,t+1}}$.

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 τ to 0. The stochastic discount factor becomes: $m_{t,t+1} = \frac{C_{k,t+1} - C_{k,t}}{C_{k,t} - C_{k,t-1}}^{-1-\alpha}$. The habit parameter α replaces τ as the determinant of how slow the post-announcement transition is. Now, we do not parameterize α 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 α , β , and now α .

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



(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 and 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, low share of workers, low productivity, and high discount factor.

Figure 10: Model Robustness

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 1.01 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 . We now set β to a higher value of 0.996, thus targeting a low implied risk-free rate of 0.4% (p.a.), which loosely corresponds to the Norwegian low-interest rate environment of 2016-2022.

Figure 10 reports the results in two stages. First, in Panel (a) we plot relative spending responses implied by all 10 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 β 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 and is qualitatively unchanged from our baseline outcome: it is roughly 2 on average with a robust lower bound of unity. The only exception is the case of habit formation instead of heterogeneous tax incidence, which produces irregularly noisy bands in general but is still significantly larger than unity if we use 90% confidence bands.

Additional results. We present and discuss four supplementary results in Appendix F. First, we show how the model-implied spending response would look if we set the EIS (β) 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 $\tau = 0$ 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.

9 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.

Beyond the implications for macroeconomics and finance, our results also show that dividend taxation can successfully reduce income and, more importantly, *consumption* inequality. Nevertheless, there may have been some negative implications for aggregate demand via the negative spending response of the more exposed households. This aggregate demand channel of capital income taxation has been relatively under-explored in contrast to the neoclassical cost-of-capital channel. In light of major capital taxation reforms in the U.S. and other countries since the early 2000s and the evidence of a strong spending response that we show in this paper, investigating the aggregate demand channel of capital income taxation appears to be a priority for future research on this topic.

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,” *American 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 Joel 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.
- 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.
- Flynn, Joel P, Lawrence DW Schmidt, and Alexis Akira Toda**, "Robust Comparative Statics for the Elasticity of Intertemporal Substitution," *arXiv preprint arXiv:2201.10673*, 2022.
- 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," *Quarterly 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.
- 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 non-stockholders,” *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**, “Bad news and robust comparative statics for the elasticity of intertemporal substitution,” *Available at SSRN 2544688*, 2019.
- Stortingsmelding nr. 29(2003-2004)**, *Stortingsmelding nr. 29 (2003-2004) Om skattereform 2004*.
- Straub, Ludwig 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

March 2024

Contents

A Derivations for the Illustrative Model	1
B Institutional Background	2
C Stock Market Analysis Details	5
D Imputed Spending without Private Businesses	7
E Numerical Details for the Structural Model	9
F Additional Model Results	12

A Derivations for the Illustrative Model

The agent's problem is

$$\begin{aligned} & \max_{C_0; C_1; C_2} u(C_1) + u(C_2) + u(C_3) \\ & s.t: C_0 + \frac{C_1}{R_1} + \frac{C_2}{R_1 R_2} = A_0; \end{aligned}$$

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

$$C_0 = \frac{A_0}{1 + R_1^{-1} + R_1^{-1} R_2^{-1}}; \quad C_1 = \frac{R_1 A_0}{1 + R_1^{-1} + R_1^{-1} R_2^{-1}}; \quad C_2 = \frac{R_1 R_2 A_0}{1 + R_1^{-1} + R_1^{-1} R_2^{-1}};$$

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

$$\begin{aligned} \frac{\partial C_0}{\partial R_2} &= -(-1) \frac{R_1^{-1} R_2^{-2} A_0}{1 + R_1^{-1} + R_1^{-1} R_2^{-1}} \equiv -(-1)_0(A_0; R_1; R_2) \\ \frac{\partial C_1}{\partial R_2} &= -(-1) \frac{R_1^{-1} R_2^{-2} A_0}{1 + R_1^{-1} + R_1^{-1} R_2^{-1}} \equiv -(-1)_1(A_0; R_1; R_2) \\ \frac{\partial C_2}{\partial R_2} &= \frac{R_1 R_2^{-1} A_0 (1 + R_1^{-1}) + R_1^2 R_2^{-2} A_0}{1 + R_1^{-1} + R_1^{-1} R_2^{-1}} \equiv {}_2(A_0; R_1; R_2) \end{aligned}$$

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

B Institutional Background

This appendix presents additional details on the dividend tax reform of 2006 and leans on [Sørensen \(2005\)](#), [Alstadsæter and Fjærli \(2009\)](#), [Thoresen et al. \(2012\)](#), and [Alstadsæter et al. \(2014\)](#).

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.¹ 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^c , which is 28%. The dividend tax above the rate of return allowance (RAA) is 28%. The marginal dividend tax rate τ^d is thus

$$\tau_t^d = \underbrace{0}_{\text{Corporate tax}} + \underbrace{(1 - 0.28)}_{\text{Net of corporate tax}} \times \underbrace{0.28}_{\text{Dividend tax rate over RRA}} \approx 0.482$$

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.

¹During the transition in 2005, personally held shares could be transferred to a holding company without triggering a capital gains tax.

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 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 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.

C Stock Market Analysis Details



Notes: This figure shows the cumulative returns of two-way sorted portfolios where the second dimension of sorting is the market beta in panels (a) and (b) and market capitalization in panels (c) and (d), respectively, all relative to January 2003. Appendix C presents further details on how the portfolios are constructed.

Figure A.1: Two-Way Sorted Cumulative Portfolio Returns

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 1NOK or greater than 1000NOK. This amounts to roughly the bottom 1% and the top 99% 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. Fourth, we partition all stocks in a high and low dividend yield portfolio based on dividend yield being above or below the median. Fifth, we compute portfolio-specific cumulative returns, relative to January 2003.

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). Figures A.1a and A.1b show 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. Figures A.1c and A.1d report 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.

D 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.²

The relevant measure of saving is the sum of active depositing or withdrawing of resources from all asset classes. The main challenge is that the tax authorities only report total valuations within asset classes at the end of the year, and changes in these values could be due to either saving or capital gains. Our approach to computing saving within each asset class differs depending on data availability. For nominal assets, such as debt and deposits, saving during the year is the change between end-of-year and beginning-of-year values. For assets with market prices, we either impute a measure of capital gains or use transaction data directly. For housing, the main asset for almost all households, we directly observe housing transactions from the transaction registry, allowing us to observe the relevant measure of saving. For stocks, we compute capital gains on household stocks using information from the stock ownership register after 2005. In particular, we observe a household's ownership of stocks in a specific company at the end of each year and combine this information with the evolution of that company's stock price during the year to compute capital gains. Before 2005, we only observe total wealth in stocks (no information on household 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 off. 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 when we ignore private businesses, $spending^{npbo}$, is thus computed

²An 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.

as



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

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

E 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) groups as given, e.g., as in Figure 7. 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 α and the EIS, η . The grid for α is agnostic, ranging from 0.01 to 0.99. The grid for η 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 $\{\alpha; \eta\}$ pair, i.e., for every point on a two-dimensional grid. The grid comprises 50 nodes in each direction. We thus solve the model 2,500 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 π_d and a one-standard deviation negative shock to π_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 τ , ranging from 0.01 to 0.99 using 50 non-linearly spaced nodes. For each value of τ on this grid, we compute and store a new \mathcal{J} . This completes the first step of our approach.



Notes: Panel (a) shows a heatmap of IRF matching errors produced by the calibration procedure over the three-dimensional grid $\{ \theta ; \kappa ; \psi \}$ with tax avoidance intensity and reform pass-through on the horizontal and vertical axes, respectively. Colder colors correspond to lower mean squared errors. Panel (b) shows the heatmap of the implied EIS values. Warmer colors correspond to greater values.

Figure A.2: Impulse Response Matching Results

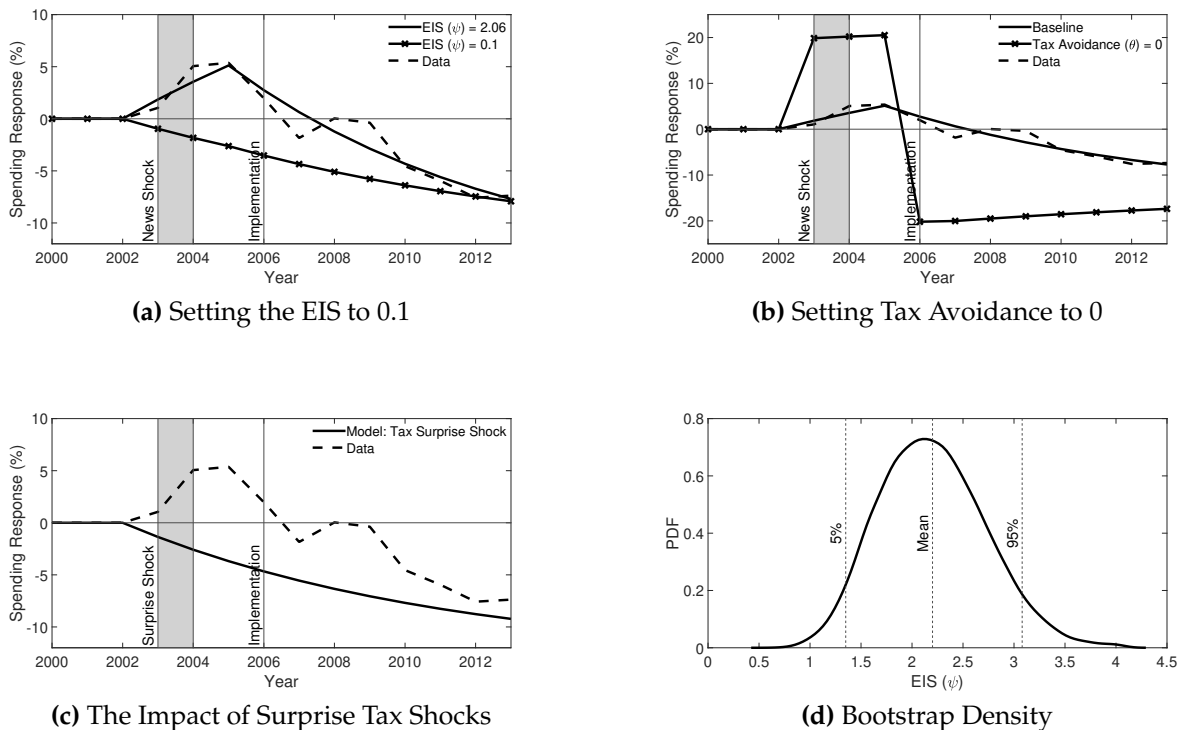
In the second step of our procedure, we locate the point on the grid that minimizes the distance between the empirical and the model-implied spending differentials. In other words, we identify the values of parameters that “match” the empirical impulse responses as closely as feasible. Our candidate model-based sequences are stored in a $11 \times 50 \times 50 \times 50$ array, corresponding to the three-dimensional grid $\{ \theta ; \kappa ; \psi \}$ with eleven rows (years). Our target is the full empirical relative spending sequence over 2003-2013, 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 $\mathcal{E}_z = E^h \left[\mathcal{J}_z - \mathcal{J}^{DATA} \right]^2$.

Finally, we identify the index \bar{z} that corresponds to the triple $\{\bar{\alpha}; \bar{\beta}; \bar{\gamma}\}$ 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{z} , 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. In Panel (a), we present a heatmap with α and β 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 of the graph 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. Panel (b) shows the respective values of the implied EIS. We notice that the northeastern region of this graph corresponds to relatively high EIS values, generally greater than 1. The patch corresponding to the global minimum matching error is consistent with an EIS in the region of 1.5 to 2.5, with a baseline estimate of 2.06.

F Additional Model Results



Notes: Panels (a) and (b) plot relative spending responses implied by alternative parameterizations in which we set $\theta = 0.1$ 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

In this section, we present four additional model results that complement the main text. First, we test whether our IRF matching exercise is quantitatively meaningful, i.e. whether picking a wrong model with high matching errors indeed produces empirically inconsistent results. Specifically, we counterfactually set the EIS (ψ) to 0.1, a very different value from what our calibration suggests. Panel (a) of Figure A.3 shows the result. Note that this parameter choice yields a relative spending response that is way off the data; it corresponds to a mean squared error that is at least an order of magnitude above the minimum. In addition, such a low EIS completely fails to match the front-loading of spending during the implementation period. Second, Panel (b) presents relative spending responses from a parameterization that sets θ to 0, a case that corresponds to no tax avoidance in the short run. This showcases that the tax avoidance friction merely affects the smoothness of the short-run response to tax news shocks and not the long-run impact.

In Panel (c) of Figure A.3, for completeness, we also present the relative spending response to a *surprise* shock to taxes, i.e., instantaneous changes in d_t and k_t *without a pre-announcement*. The plot demonstrates that in order to match the front-loading of spending over the 2003-2006 period, a *news* shock is essential. Finally, Panel (d) of Figure A.3 reports the density of EIS estimates as produced by the model bootstrap procedure with 10,000 draws. The mean, 5th percentile, and 95th percentile values are 2.2, 1.35, and 3.08, 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.
- 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.
- Thoresen, Thor O, Erlend E Bø, Erik Fjærli, and Elin Halvorsen**, “A suggestion for evaluating the redistributive effects of tax changes: With an application to the 2006 Norwegian tax reform,” *Public Finance Review*, 2012, 40 (3), 303–338.