

FISHER COLLEGE OF BUSINESS
RESEARCH PAPER SERIES

Taxpayer Search for Information: Implications for Rational Attention

Jeffrey L. Hoopes, The Ohio State University

Daniel H. Reck, University of Michigan

Joel Slemrod, University of Michigan

Fisher College of Business RP 2014-02-004

Original: July 2014

This paper can be downloaded without charge from:

<http://ssrn.com/abstract=2504001>

An index to the research papers in the Fisher College of Business Research Paper Series is located at:

<http://www.ssrn.com/link/Fisher-College-of-Business.html>



THE OHIO STATE UNIVERSITY

FISHER COLLEGE OF BUSINESS

fisher.osu.edu

TAXPAYER SEARCH FOR INFORMATION: IMPLICATIONS FOR RATIONAL ATTENTION

Jeffrey L. Hoopes, Ohio State University, hoopes@fisher.osu.edu

Daniel H. Reck, University of Michigan, dreck@umich.edu

Joel Slemrod, University of Michigan, jslemrod@umich.edu

July 23, 2014

Forthcoming: American Economic Journal: Economic Policy

Abstract: We examine data on capital-gains-tax-related information search to determine when and how taxpayers acquire information. We find seasonal increases in information search around tax deadlines, suggesting that taxpayers seek information to comply with tax law. Positive correlations between stock market activity and search as well as year-end spikes in information search on capital losses when the market performs poorly suggest that taxpayers seek information for tax planning purposes. Policy changes and news events cause information search. These data suggest that taxpayers are not always fully informed, but that rational attention and exogenous shocks to tax salience drive taxpayer information search.

JEL Codes: D80, D83, H31, H24

*We wish to thank John Guyton and Ron Hodge of the Research, Analysis, and Statistics Division of the Internal Revenue Service, for helping us use the IRS aggregate administrative data. We thank the editor, Mark Duggan, and two anonymous referees, whose comments helped improve this manuscript. We are grateful for comments on an earlier draft received from Hal Varian, Raj Chetty, Alex Gelber, and participants at McMaster University, the University of Michigan Hosmer-Hall Luncheon series, the US Treasury Department Office of Tax Analysis, the University of Michigan Public Economic Summer Talks series, the University of Uppsala, and at the 2013 Public Economics United Kingdom Conference. The views expressed here are those of the authors alone, and do not reflect the views of the Internal Revenue Service.

Taxes, especially income taxes, can be complex and confusing. Despite a general awareness of this fact, the consequences of complexity and misunderstanding are not well-understood. Survey evidence suggests that many taxpayers do not understand basic tax concepts (Blendon et al, 2003), and the compliance cost of taxes, including learning enough about them to comply, is large (Slemrod, 1995; 2004). Given confusion surrounding tax incentives and tax law, many taxpayers might simply ignore or misperceive the incentives built into the tax code when they make decisions with tax consequences. Alternatively, taxpayers might collect information on the taxes that matter to them, and use this information to make tax-efficient decisions. Finally, taxpayers might learn just enough about tax policies to fill out their tax return and perhaps avoid an audit.

In this paper, we investigate these alternative views of taxpayer information search. We find evidence that is inconsistent with the notions that taxpayers are fully informed about the tax system, that they act in complete ignorance of the tax system, or that they gather information only for tax compliance purposes. The evidence we present suggests instead that at least some taxpayers employ rational attention to tax policies, in line with theories proposed by Sims (2003) and Reis (2006). In addition, we observe that exogenous shocks to tax salience from news events can substantially increase information search.

Modern technology has greatly expanded the accessibility of information. Any person with access to the Internet may, in a few minutes, learn at least something about the most obscure details of the tax code. Taxpayers undoubtedly do use these resources to seek information: people Google “tax” more often than they Google the names of public figures,¹ the IRS website has received on average 4.6 million visits *per day* since 2004, and the IRS call line has received on average 125 thousand calls per day since 1999. How tax knowledge matters hinges on 1)

¹ For example, Google searches for “tax” were more common than searches for “Barack Obama,” and were more common than “Obama” in all but the 2008 and 2012 election season.

how and when people seek out tax-related information, and 2) whether they change their behavior once they acquire it. In this paper, we address primarily the first of these questions, and provide some preliminary analysis of the second.

Our data consist of aggregate high-frequency time series on calls to the IRS toll-free phone number, aggregate visits to the IRS website, searches on the IRS website, measures of Google searches on tax-related terms, and views of tax-related web pages on Wikipedia. From these sources we collect data on information seeking regarding one specific tax-related topic: capital gains taxes. We select this topic because it is a perennially controversial policy issue, because data on the relevant taxed behavior, sales of capital assets, are available on a high-frequency basis, and because the American Taxpayer Relief Act of 2013 (ATRA) enacted a change in capital gains tax rates during our sample period.² We also note that capital gains taxes represent a non-trivial portion of total government revenues, so are interesting in their own right. In 2006, for example, capital gains realizations represented 5.96 percent of GDP, and taxes paid on those gains amounted to \$117 billion (Tax Foundation, 2010).

We study information-seeking around five different types of events: 1) time notches, 2) macroeconomic changes, 3) policy changes or the mention of potential policy changes, 4) filing deadlines or approaching filing deadlines, and 5) tax-related news events. Information search around the first three dimensions should be tied to the taxpayer's ability to make fully informed decisions affected by tax incentives. First, a particular date is often relevant for the incursion of tax liability or the tax efficiency of behavior. We call these *time notches*³ to indicate that tax liability can change abruptly, and non-incrementally, at certain dates, usually at year-ends. For example, the last date to affect one's capital gains tax liability in a

² We recognize that the type of people seeking information about capital gains tax are different from, say, those seeking information about the earned income tax credit. We expect that future work will examine whether the phenomena we document here also occur for other policies.

³ This terminology is taken from Slemrod (2013).

given year is December 31st. As the deadline approaches individuals may research the benefits from realizing a capital gain or loss in the current year as opposed to a future year. Second, macroeconomic changes may be associated with information search due to perceived changes in the importance of understanding the tax consequences of a particular behavior. For instance, asset owners may seek information about capital gains taxation as the stock market falls in order to understand the tax implications of their losses. Third, actual or potential changes in tax policy or tax enforcement may cause individuals to learn about the tax system, either to plan for the future or to make an informed voting choice. Taxpayers might also seek information on the tax system when completing their tax return in order to both fully comply with the law, as well as ensure they are using all the credits and deductions available. As such, the fourth dimension we study is filing deadlines and the onset of these deadlines, when individuals may learn about taxes due to approaching deadlines for filing tax returns. For the taxes we consider, the deadline affecting most taxpayers is in mid-April.

News events sometimes spark public interest in tax policy, and in this case individuals may search for tax information in order to develop an informed opinion on current events or out of curiosity. The fifth dimension we study consists of news stories such as the release of a public person's tax information. We can view these events as exogenous shocks to the salience of a particular aspect of tax policy. They have no direct bearing on an individual's tax burden, nor do they directly affect tax incentives. The search may, however, provide information that incidentally informs the searcher about her own tax situation or about the tax impacts her decisions. They may also affect other important decisions, such as voting behavior.

We observe strong seasonality in the search for information on capital gains taxes through all channels. Taxpayer information search increases substantially during the period commonly called "tax season," which runs from mid-January to

mid-April of each year. An even more pronounced spike in information search occurs very close to the filing deadline in mid-April. We also document the impact of several discrete events on taxpayer information search regarding capital gains taxes, through Google and Wikipedia. Presidential debates in which candidates discuss their proposals for capital gains taxes, the passage on January 2, 2013 of ATRA, the release of presidential candidate Mitt Romney's 2010 tax return in January of 2012, presidential elections, and policy changes all generate large and significant increases in taxpayer information search. In every case, these events cause a spike in taxpayer information search that fades within three to four days.

Next, we find that macroeconomic changes affect information search on capital gains taxes through Google and Wikipedia. We observe significantly elevated information search on days with large trading volume in the stock market. We also explore a measure of information search for personal investment advice generally, using Google searches. Daily searches for stock advice predict searches for capital gains tax on the same day, the previous day, and, in some specifications, one day in the future. This is the first evidence in the literature that some taxpayers investigate the tax consequences of an action while contemplating the action itself.

These patterns indicate that information search is consistent with a model in which individuals search when different events make understanding tax policy more important, i.e. when they increase the return to information search. However, these results are also consistent with individuals making choices and *then* researching the tax implications of those choices for their wealth. In an effort to provide evidence that more strongly suggests a causal relationship, whereby events cause individuals to search for information in order to make more informed decisions, we examine information search on capital losses. We document substantial elevation in information search related to capital *losses* at year-ends,

especially in 2008 (when capital losses from the stock market crash would have made taxpayers' more likely to be rebalancing their portfolios in order to take tax-efficient advantage of capital losses). Observing increases in information search just prior to the time notch is consistent with taxpayers seeking information on the tax system and using it to improve decision-making.

This paper provides the first-ever attempt at understanding how, and why, taxpayers search for information about government policy, specifically tax policy. We establish that taxpayers seek information in order to both comply with their tax obligation and to respond behaviorally to tax incentives. Taking taxpayer learning seriously has important implications for understanding the impact of tax policy and tax policy changes as well as for a full understanding of the nascent idea of tax salience.

2 Background

2.1 Public Understanding of Taxes

Political science research is fairly persuasive that voters know very little about the details of government. Delli Carpini and Keeter (1996), in a comprehensive survey of the political knowledge of voters covering several decades and hundreds of surveys, show that the majority of voters are ignorant of many key aspects of the U.S. political system. Surveys suggest that people are also largely ignorant of the tax system (Blendon et al (2003) summarize the results).

One specific setting where researchers have investigated understanding of the tax system is taxpayers' perceptions of their average and marginal tax rate. Brown (1968) and Fujii and Hawley (1988) find that individuals' self-reported marginal tax rate often differs from the true rate that can be estimated from their demographic characteristics. Research by de Bartoleme (1995) shows that, in a lab experiment, MBA students often confuse the average tax rate with the marginal tax rate when making investments in a taxable versus non-taxable project. Graham et al (2014) show that this misunderstanding of taxes extends to

corporate managers—only 13 percent of corporate tax executives surveyed responded that their firm used the theoretically correct marginal tax rate (as opposed to some other tax rate) in corporate decision making. Sheffrin (1994) and Liebman and Zeckhauser (2004) also find evidence that taxpayers systematically misunderstand income tax schedules.

Recently, the public economics literature has settled on the term “salience” to capture the extent to which tax aspects of the environment are noticed, and acted upon, by those affected. A key paper in the modern literature is Chetty, Looney, and Kroft (2009). Their findings come from an experiment at a large grocery retailer in California, at which prices inclusive of the 7.375 percent state sales tax were posted alongside the original pre-tax price over a three-week period for three product categories (cosmetics, hair care accessories, and deodorants). They estimate that the “tax treatment” reduced demand by 8 percent; given demand elasticities of 1 to 1.5 for the affected products, they conclude that most consumers do not take into account the sales tax revealed at the cash register. A crucial question for the policy implications of salience, relevant to our analysis, is whether taxpayers pay attention to tax incentives when more utility is at stake (Goldin, 2014; Reck, 2014).

2.2 Macroeconomics and finance

Many macroeconomists have considered how best to account for the inertia in observed economic behavior and to what extent imperfect information can account for it. For example, the models in Sims (2003) are motivated by the idea that information that is freely available to an individual may not be used, because of the individual’s limited information processing capacity. Alternatively, in Reis (2006) consumers rationally choose to only sporadically update their information and re-compute their optimal plans, while in between updating dates they remain inattentive. Both models imply that news disperses slowly throughout the population, so events have a gradual and delayed effect on behavior.

A recent literature in financial economics has taken advantage of newly available data and examined the demand by investors for information. Da, Engelberg, and Gao (2011) propose the Google Search Volume Index (GSVI) as a direct measure of investor attention. Drake, Roulstone and Thornock (2012) use the GSVI for public company ticker symbols to examine the timing and magnitude of Internet search around earnings announcements and the factors that influence Internet search. Among their results of interest is that investors extend more effort when the potential returns to search are higher.

2.3 Implications of theories of information acquisition

Although our analysis focuses on taxpayers' search for information about income taxes, it has more general implications regarding the information economic agents possess when making decisions (in our case, often very economically large decisions). Do individuals possess full information for important economic decisions? If not, is attention to important information rational? If attention is rational, how should we model information acquisition? In general, previous economic studies on attention have focused on the different predictions alternative assumptions make for behavior. However, every theory of behavior contains some assumption, implicit or explicit, about how individuals come to possess information. Our approach is to directly examine these auxiliary predictions of theoretical models, by documenting patterns in the timing of information search. Table 1 summarizes the predictions of the theories our analysis illuminates, as described in this section.

The null hypothesis in our analysis is that taxpayer information search is unresponsive to political or economic events. This hypothesis would be implied by the assumption of full information, wherein an individual always has all data necessary to make an optimal choice. Although frequently criticized and relaxed in the economics literature, this assumption is still common in public finance for modeling the response to tax incentives and the welfare cost of taxation.

Interestingly, the null hypothesis would also obtain in a model of full ignorance: if individuals “don’t know what they don’t know,” they will never search for information.⁴

In rational-expectations models like the one proposed by Muth (1961), individuals face uncertainty about the future, but know everything about the present and costlessly and instantaneously absorb all new information. A rational expectations framework might allow searches for tax information to respond instantaneously to unanticipated policy changes or shocks to the probability of a policy change, as these require that the individual absorb new information about policy. However, under rational expectations, information search related to one’s personal tax liability should not respond to economic or political events that change the importance of understanding tax policy but not tax policy itself. Information search should also not respond to the implementation of policy changes that were not surprises. Individuals should already possess all publicly available information on tax policy.

An attractive alternative framework is rational attention, under which individuals have limited capacity for processing new information and they allocate that resource optimally. Alternative assumptions for how limited attention may be rationally allocated are proposed by Sims (2003) and Reis (2006). In Sims’ (2003) approach, agents constantly update their beliefs based on their attention to different sources of information, while in Reis’ (2006) approach, updating beliefs is costly and, upon updating, the individual simultaneously plans future consumption and decides when to update again. In either model, information search should increase following events that increase the utility gain to understanding the incentives at play in making a decision (such as a market

⁴ From a Bayesian perspective, in order for individuals to have full information or full ignorance, they must have complete confidence in their own beliefs. Full information obtains when those beliefs turn out to be correct, full ignorance obtains when they are wrong. That the demand for information is zero in both cases follows from the nature of the dependence of the demand for information on prior beliefs documented in Keppo, Moscarini, and Smith (2008).

downturn or approaching time notch) and to events that introduce new information about incentives (such as policy changes). Vitrally important for the interpretation of our results is the fact that with a sufficiently relaxed constraint on attention, rational attention becomes behaviorally equivalent to rational expectations. Suppose individuals quickly obtain all the public information necessary to make an optimal choice. Then a rational expectations model with costless information acquisition—or even a full information model if there is little uncertainty—would accurately predict choices, despite a rational attention model performing better in predicting information search and choices jointly.

Attention is also the subject of a large literature in psychology. One useful distinction in this literature is between “exogenous,” or “stimulus-driven,” attention and “endogenous,” or “goal-directed,” attention (Theeuwes, 1994; Connor, Egeth, and Yantis, 2004). Endogenous attention is the same as rational attention: the individual voluntarily directs attention to meet a goal. Exogenous attention is driven, rather, by an external stimulus, such as increased salience. In this case, the individual might search for information regarding anything she sees in the news which she does not understand, regardless of her gain from the search. If attention to economic information is exogenously driven at times, information search should respond to news events that mention policy—because these are usually accompanied by news coverage that makes the policy more salient.

3 Capital Gains Taxation and Behavior

Capital gains generated from the sale of capital assets are subject to tax, but have received preferential treatment relative to labor income since 1921 in the United States. Income from the sale of capital assets is recognized in the year of sale, and the taxable income is equal to the sale price of the asset less its tax basis (the historical price plus any acquisition costs and improvements to the asset, less any accumulated depreciation). Capital gains and losses are divided into two categories, short-term and long-term. Long-term (short-term) capital gains and

losses arise from the sale of a capital asset that has been held for more than (less than) one year. Long-term capital gains receive a favorable tax treatment and currently face a maximum capital gains tax rate of 20 percent, along with a new 3.8 percent Net Investment Income Tax that is applicable to taxpayers in the highest tax bracket. Short-term gains are taxed as ordinary income. Gains and losses from short and long term assets are combined to determine a taxpayer's ultimate tax liability.

A rational investor should maximize after-tax utility, which would entail taking capital gains taxes into account in deciding whether to buy, sell, or hold a capital asset. Consider the purchase decision in isolation from the rest of an investor's portfolio. The taxation of realized capital gains reduces the expected after-tax rate of return to a capital investment, where the reduction depends on the expected appreciation, the expected holding period, and the likelihood that an asset with appreciation can be held until death. Because higher capital gains taxes reduce the attractiveness of assets expected to appreciate in value, one may expect that the level of asset prices would react negatively to unexpected news about tax increases (described as a "capitalization" effect).⁵

The decision of if and when to sell a capital asset should also be affected by the tax system. Certain provisions in the tax code (for example, a lower tax rate for long-term capital gains), may encourage taxpayers to postpone an asset sale in order to obtain the favorable tax treatment. Likewise, the annual nature of tax compliance may also create annual rebalancing of portfolios to achieve a favorable mix of capital losses and gains (such that the losses almost exactly cancel the gains), and may encourage asset sales.

Because capital gains taxes are triggered by asset sales that happen in year t , there is a planning deadline for capital asset sales at the end of the year. Therefore,

⁵ Of course, the effect of acquiring information about capital gains taxes depends on how it changes prior beliefs. One can imagine a potential investor being pleasantly surprised to learn about the preferential lower tax rate and step-up basis at death, or being discouraged upon learning that any tax at all applies upon the sale of appreciated assets.

information gathering in order to achieve the best after-tax outcome with regards to capital asset sales will have to happen by December 31 of year t . When capital gains tax rates increase between year t and year $t+1$, December 31 is also the last date to realize capital gains at the year- t tax rate. As such, information gathering in order to shift capital income through time in response to a rate change will also need to happen by December 31st of the end of the year.

Previous empirical evidence is largely consistent with taxpayers altering both the nature and the timing of transactions in order to achieve the greatest possible after-tax return on their capital investment. For example, as capital losses must generally be offset with capital gains, large downward movements in the stock market often leave investors with large unrealized capital losses that must be carried forward until years when these investors realize capital gains. This policy leads to the well-known strategy for minimizing capital gains tax liability of “loss harvesting,” selling capital assets with built-in losses to offset the gains realized during the year. For example, Ivković, Poterba and Weisbenner (2005) find evidence of tax-loss selling at the end of the year. Poterba and Weisbenner (2001) find that this is especially prevalent in years when changes in tax policy provided additional incentive to harvest losses at year-end. While tax-loss selling occurs in the corporate equity market, there is also evidence of such activity in municipal bond closed-end funds (Starks, Yong, and Zheng, 2006) and the market for long-term government and corporate bonds (Chang and Pinegar, 1986).

4 Data

Tax information to taxpayers is made available by the IRS, by accounting firms and other organizations in the business of providing information generally, such as Wikipedia. The IRS-provided information can be accessed through the Internet, via printed information booklets, and through toll-free numbers. Person-specific information can also be obtained by perusing one’s paycheck, one’s Form W-2, or one’s prior tax returns. Both information from the IRS website and

information from non-government organizations may be found quickly using Google searches. In all cases, the marginal cost of public information is only the time spent acquiring it, provided one has an Internet or phone connection. The cost of acquiring tax information decreased with the advent of the Internet, and search engines in particular. Long-run trends in information search may therefore vary due to both supply and demand for information. By focusing on high-frequency variation in information search, we isolate variation in the demand for information because the supply of information is unlikely to vary from day to day.

We examine measures of information search through many different channels. We first study data on tax-related information search from Google, accessed via Google Trends. We use Google data not only because Google makes its search data publicly available, but also because Google is the most widely used search engine throughout our sample period, capturing 66.9 percent of search volume in 2013 (comScore, 2013).⁶ One can use Google searches to find information from a variety of sources, including Wikipedia and the IRS. Using query data on Google searches, Google Trends provides a measure of the “propensity to search” for a given search query or set of queries. More specifically, an observation in the Google Trends data will be, for a given day and geographical region, the number of Google searches for the specified search terms divided by the total number of Google searches on any topic in the time and place.⁷ For our purposes, the search terms will be a broad set of capital-gains-tax-

⁶ Google’s share of internet searches has grown over the sample period. According to comScore (2005, 2007, 2009, 2013) Google’s share of online searches was 36 percent in 2005, slightly exceeding that of its closest competitor, Yahoo!. Google’s share grew rapidly until 2007, when it reached around 60 percent of the market. It has grown about 1 percent per year since then. Note that the Google’s scaling of the data and our removal of the long-run trend ensure that this does not jeopardize our interpretation of our results.

⁷ For more detail about Google Trends, see our appendix and <http://support.google.com/trends/?hl=en>. Unfortunately, Google does not provide their search data separately as a numerator (Capital Gains tax searches) and a denominator. While the scaling done by Google does help alleviate the concern we are capturing some temporal trend related merely to overall Internet usage, it raises the issue that we are unable to disentangle how much of our effect is coming from the numerator (Capital Gains tax searching) or the denominator (total Google searches). In the absence of data on total Google searches by day, which we are unable to attain, we are unable to investigate the separate numerator and denominator effect. However, the fact that our Wikipedia data are unscaled and that we obtain results that are materially the same, scaled or unscaled, suggests that the denominator effect in the Google data is not driving our results.

related search terms (listed in the Appendix), and the geographic region is the United States. After calculating the propensity to search, Google scales it from 0 to 100, where the number 100 represents the day with the highest search volume for this set of search terms in the entire sample period. Our sample period consists of January 1, 2004 (when data are available) through March 30, 2013.

Entering a given term into Google will frequently yield a Wikipedia page related to that term as one of the first suggested sites to visit. Our second source of information search data is page view data from Wikipedia, a free online encyclopedia edited by Internet users. We obtain the number of people who landed on the English language Wikipedia site, “Capital Gains Taxation in the United States” every hour. We use the summation of this hourly data over the 24 hours in a day to obtain daily data for January 1, 2008 through March 30, 2013.⁸ One disadvantage of the Wikipedia data is that it is available for a substantially shorter time period than is the Google data. On the upside, we are able to obtain the raw number of views of the webpage, which simplifies the interpretation of the size of changes in information search.⁹

Our third source of data regards aggregate calls made to the IRS’s toll free phone number, where taxpayers can call and speak to a representative from the IRS or listen to automated messages. In the course of a phone call with the IRS, some taxpayers will listen to prerecorded messages about various “tax topics.” We analyze daily phone calls that access IRS tax topic 409, “Capital Gains and Losses.”¹⁰ Taxpayers may reach the tax-topics function in a number of ways, including calling the tax-topics-specific phone number, or being referred to the

⁸ For a limited number of hours, Wikipedia data are unavailable. We scale the observation up by 24 divided by the number of hours present in a day when only some hours are missing (15 days), and we use the previous day’s value when the entire day is missing (128 days).

⁹ We have replicated the results in Section 6 and Section 7 using a measure of Wikipedia usage scaled in the same way as the Google data, using page views of the Wikipedia home page as a metric of overall Wikipedia usage. This replication suggests that scaling is not confounding our results.

¹⁰ Other example of tax topics are tax topic 304, “Extensions of Time to File Your Tax Return” or tax topic 151, “Your Appeal Rights.”

topic after a conversation with a representative or an interaction with the automated system. We use data from February 1, 2002 to March 31, 2012. We also obtain the total number of calls to the call line for individuals, which is not specific to capital gains.

The last data set we analyze consists of the aggregate daily number of visits¹¹ to the IRS’s webpage. The IRS maintains a website, IRS.gov, which hosts a vast amount of tax information. We obtain, directly from the IRS, the total number of page views and visits, by day, to any site hosted by IRS.gov. We are able to obtain this data for the entire time series for which the IRS has maintained the data, February 1, 2002 to March 31, 2012. Analysis of IRS.gov traffic carries the benefit of analyzing taxpayer information search about a very broad set of tax information.¹² This measure of information search is not limited to information search about capital gains taxes specifically.

5 Estimation Procedure

Before discussing data patterns, in this section we lay out the research design for our parametric analysis. Let I_{it} denote information search on date t from a source of information i , such as Google, Wikipedia, or the IRS call line. We wish to estimate the effect of several different events on information search through source i . The events we study may be either non-recurring, as in the case of a Presidential debate that mentions capital gains taxes, or may occur annually, as in the case of tax filing season. Non-recurring events are indexed by the subscript k , and annual events are denoted according to the day of the year on which they occur, denoted DoY . As we discuss in more detail later, the data on information search display marked seasonality at the yearly and weekly levels. In order to

¹¹ In contrast to page views, visits are a web analytic that count the number of page views from unique Internet Protocol (IP) addresses in a given time frame. As a result, a single viewer from the same computer may refresh the webpage multiple times, and each refresh will count as a new page view, but not a new visit. Wikipedia provides data only on page views.

¹² The phone call and website visit data was made available to us by the Research, Analysis and Statistics division of the IRS.

evaluate whether a particular (non-seasonal) event increased information search on date t , we must also properly specify the counterfactual level of search, conditional on the day of the week and day of the year of date t .¹³

Unlike for much analysis of time-series data, yearly seasonality is of intrinsic interest for our research questions. In particular, increases in information-seeking during tax season will affect the form of the yearly seasonality. Information searches occurring at the end of the year can be thought of as responses to a time notch. To address this concern, we employ a method that both estimates the seasonal patterns and allows us to perform classical statistical tests for whether information search is significantly higher on a particular day of the year than the average level of information search. Specifically, for each information search series i we estimate the function:

$$I_{it} = \sum_k [\beta_{k0}F_{kt} + \beta_{k1}F_{k(t-1)} + \dots + \beta_{k4}F_{k(t-4)}] + x_t' \gamma + f_i(DoY_t) + \delta_{DoW,i} + u_{it} \quad (1)$$

The term inside square brackets is a set of dummy variables: $F_{k(t-j)}$ equals 1 if event k occurred on date $(t-j)$ and zero otherwise. The four-day event window was selected because when large spikes in information occur in the data, search levels return to baseline within four days.¹⁴ We also include a vector of continuous, time-varying linear covariates x_t' —such as trading volume on the stock market—along with a non-parametric function in day-of-the-year $f_i(DoY_t)$, and a day-of-the-week fixed effect, $\delta_{DoW,i}$.

We estimate $f_i(DoY_t)$ using kernel-weighted local linear regression. This procedure estimates a weighted linear regression at each value of DoY , using data

¹³ While we explicitly control for these temporal trends in searching in our regression framework, as a result of how Google provides their Search Volume Index (SVI), Google has also helped alleviate some of these temporal issues. For example, people likely search in general more during Christmas break, so a concern may be that we are picking up Christmas break searching as part of our year-end search activity. However, as explained earlier, the SVI divides standardized total searches for the topic per day by total search volume on Google. As a result, factors such as mere increases in Internet usage likely do not explain the patterns in search behavior we see.

¹⁴ Figure A.4 in the Online Appendix contains a plot of the seasonally adjusted and detrended data around event dates, verifying that the impact of events tends to fade in four days.

from adjacent days (e.g. the j^{th} day before and the j^{th} day after DoY for some j).¹⁵ Dates further away from some DoY receive less weight in the estimation of the linear regression at DoY ; how the weights decay is governed by our choice of kernel function and bandwidth, discussed below. Our estimate will be consistent under the assumption that the value of $f_i(DoY_t)$ does not change *too* sharply from day to day.¹⁶ The strength of this estimator is that 1) it estimates a smooth trend in day-of-the-year, which allows us to control for day-of-the-year when estimating the effect of events, and 2) it increases the precision of the estimates of $f_i(DoY_t)$ relative to fixed effects. A potential downside to the estimator is that the assumption that $f_i(DoY_t)$ does not on average change too sharply from day to day may be wrong (especially around April 15th). The validity of this assumption is determined by the bandwidth of the estimator, and also by our choice of kernel density function. We discuss these choices in more detail in the online. Our preferred choices for kernel density function (Gaussian) and bandwidth (four days) produce estimates of seasonal trends that match the pattern implied by fixed-effects estimates. To avoid conflagration of seasonal effects with large outliers due to the events we document, we estimate all the components of equation (1) in a single regression. The appendix to this paper outlines the estimation procedure in more detail.

For a recurring annual event such as a filing deadline, we can examine the function $f_i(DoY_t)$ on days-of-year corresponding to the annual event to understand the effect of the event on information search. The comparison that seems most natural is to test the hypothesis that, on a particular day of the year, DoY' , inquiries are higher than the average over all days of the year, i.e. that $f_i(DoY') > E[f_i(DoY)]$.¹⁷

¹⁵ Leap days are dropped from the sample for simplicity, so that every year is 365 days long.

¹⁶ Alternatively, we could estimate $f_i(DoY_t)$ using day-of-the-year fixed effects. However, with only 5 to 9 years of data for each type of inquiry, these fixed effects would be imprecisely estimated.

¹⁷ This is analogous to a more conventional hypothesis test with fixed effects, such as the one we use to test for significant

Our estimate of the function $f_i(DoY)$ will inform us about the importance of tax season for information search, but it may not capture all variation due to compliance deadlines.¹⁸ For example, the mid-April filing deadline does not occur on April 15 of a given year if that date falls on a weekend or Emancipation Day. Instead, it can occur as late as April 18 in some years. We account for variation in compliance dates by year by including dummy variables for compliance events in addition to the locally linear function. Our estimate of the function $f_i(DoY)$ will estimate the average effect on April 15, but we add an additional dummy variable for the precise date of the filing deadline to capture variation that occurs specifically on the deadline each year. Doing so does not significantly change our picture of the importance of tax season for inquiries, but it does highlight the sharp spike that occurs exactly on the mid-April deadline each year. The inclusion of this dummy variable also helps mitigate the strength of the smoothness assumption in estimating $f_i(DoY)$, effectively imposing a bandwidth of zero on the date of the filing deadline. To further reflect that the smoothness assumptions should not apply when there are sharp changes on particular days of the year, we also include dummy variables for the following holidays: New Year’s Day, President’s Day, Martin Luther King, Jr. Day, Memorial Day, Independence Day, Labor Day, Thanksgiving, and Christmas.

6 Exploring the Raw Data

Figure 1 plots the evolution of three measures of information search on capital gains taxes over time. Calls to the IRS hotline inquiring about capital gains taxes occur almost exclusively during tax season, from mid-January to mid-April. The absolute volume of calls occurring during a tax season diminishes considerably

day-of-the-week effects. Rather than testing for significance relative to a “left-out” day of the year, however, we test for significance relative to the average over all days of the year.

¹⁸ For Google data, the interpretation of the estimates is somewhat complicated by the fact that the dependent variable is a measure of propensities to search. If there were yearly seasonality in overall Google searches, this would affect our estimate of $f_i(DoY)$ by decreasing it on days of the year when Google usage was highest. The fact that the pattern estimated by Google data and other measures of information search are similar largely alleviates this concern. The inclusion of day-of-the-week fixed effects also makes our results robust to weekly seasonality in Google searches.

throughout the time period, from 2002 to 2012. Wikipedia page views, in contrast, increase over time. These patterns are consistent with online information largely having supplanted telephone information as Internet access increased markedly over the sample period, as well as the fact that the increasing popularity of tax preparation software and online programs likely reduces the need for taxpayers to call the IRS regarding simple questions. The Google measure displays a slight downward trend.¹⁹

Visually apparent in each time series is a strong pattern of yearly seasonality. This pattern is most pronounced in the IRS call log series, but it is also present in Google searches and Wikipedia page views. Focusing on narrower time frames also reveals strong weekly seasonality in each time series.

To focus on high-frequency variation, we detrend the data using a Hodrick-Prescott filter.²⁰ To allow for visual comparison of the three measures of information search, we normalize each variable by dividing by its standard deviation.²¹ Figure 2 plots the estimated yearly seasonality in the standardized data.²² For each measure of information search, we observe a sizable and significant increase in search behavior during tax filing season, starting first in early January, as some taxpayers are likely filing early to get refunds as soon as possible, and others are receiving information returns (e.g., 1099-DIV, 1099-B) and pondering how they affect their capital gains tax liability. Then, we observe an even stronger spike in the immediate run-up to the April 15 filing deadline. Right after April 15, information search drops off sharply. The estimated seasonality in Figure 2 matches patterns in the timing of filing income tax returns

¹⁹ Recall that this does not mean that Google searches of capital gains taxes decreased over time, only that the share of Google searches that concerned capital gains taxes decreased over time.

²⁰ Whenever we use a Hodrick-Prescott filter, we use a smoothing parameter of 10^7 . This value was selected by trial and error, with the goal that the long-run trend captures long-run movements in the series but not variation due to yearly seasonality.

²¹ Because the volatility of the IRS call log series varies significantly over time, we standardize it by dividing by its standard deviation *by year*.

²² This estimate of seasonal effects comes from the regression in Section 7, and controls for event-driven outliers.

documented by Slemrod et al (1997). Clearly, the desire to comply with the tax law and take advantage of applicable tax credits and deductions as one fills out a tax return leads taxpayers to search for information.

In the next set of figures (which we analyze formally later in the paper), we focus more narrowly on variation in the standardized, detrended, and seasonally-adjusted series around three dates where we observe the largest spikes in information search.²³ Figure 3 focuses on the spike in October and early November 2008.²⁴ There are two obvious candidates for what triggered elevated information search in this time period: the presidential election of 2008, and the stock market crash. Because capital gains taxes were an issue on which candidates John McCain and Barack Obama differed substantially, the election could influence search behavior due to either the desire to make an informed choice about asset purchases or sales based on expected future tax policy, or due to the desire to understand the consequences for future tax policy of a President McCain or Obama. The market crash of 2008 resulted in large capital losses for many investors and, due to the extreme volatility of the market, potentially also large short-run gains. Our belief is that both the election and the crash played some role, although it is difficult to disentangle the effect of the two events. The largest swings in the stock market in this period (marked by vertical lines in Figure 3) are associated with modest elevation in search behavior, and information search surged strongly around the date of the presidential election. Interestingly, in this figure, we note no surge in search activity related to the IRS call line, suggesting that different information sources serve different purposes. Consistent with other evidence presented here, it appears that taxpayers mostly use the IRS call line for

²³ Seasonal adjustment consists of subtracting out estimated day-of-the-week effects, holiday effects, and the day-of-the-year effect. The full detrended and seasonally adjusted time series are provided in Online Appendix Figure A.1.

²⁴ If we estimate a simple regression of (1) including an indicator variable for whether the date was in September, October, or November of 2008, we find that Google searches were 1.06 standard deviations higher than the seasonal trend on average over the period and Wikipedia searches were elevated above the seasonal trend by an average of 0.501 standard deviations.

tax compliance purposes, as we do note a strong seasonality associated with tax season in the IRS call data. We present more evidence on how the crash might have affected information search by focusing more narrowly on capital losses in section 8, and Figure 6 provides further evidence on presidential elections.

Panel A of Figure 4 plots the same series in January and February of 2012. We attribute the observed surge in information search to the release of presidential candidate Mitt Romney's 2011 tax return on January 24. The release generated substantial news coverage, in part due to his low effective tax rate. As most news articles on the subject noted, much of Romney's income came from the realization of long-run capital gains taxes, taxed at 15 percent at that time. Another possible explanation for the spike in search behavior on this date is the State of the Union Address, which also occurred on January 24, 2012. In his speech, President Obama advocated taxing the wealthy at higher tax rates (supporting the "Buffett Rule"), but he mentioned neither capital gains taxes specifically nor capital income taxes more generally.²⁵ It is, however, possible that Obama's rhetoric on taxing the rich led the public to pay greater attention to the news about Mitt Romney's tax return, and thereby amplified its effect on capital-gains-tax-related information search.²⁶

Panel B of Figure 4 plots the data from detrended and seasonally adjusted data at the end of 2012 and the beginning of 2013. A presidential debate between Barack Obama and Mitt Romney appears to have sparked considerable information search. As in 2008, we also see elevated search immediately following the presidential election, perhaps as voters and investors researched what might happen to capital gains taxes in the aftermath of President Obama's

²⁵ The "Buffett Rule" is a tax plan proposed by President Obama, wherein individuals making over \$1.0 million in taxable income would be subject to a minimum average tax rate of 30 percent regardless of whether the income is composed of capital gains or ordinary income.

²⁶ Further insight into this issue can be gained by examining Wikipedia intra-day search activity. Search activity started rising dramatically mid-day on January 24th, 2012, consistent with Romney's tax return being at least partially responsible for this increase (as Obama's speech was not delivered until that night). This differs from the typical intra-day pattern of search activity, which generally shows more searching done after working hours.

re-election. Finally, the largest spike in information search in the 10-year period covered by our data—an increase of just over 4 standard deviations in Wikipedia page views and over 6 standard deviations of Google searches—occurred on January 2, 2013. We attribute this to the passage of the American Taxpayer Relief Act (ATRA) on that date. Again, we note the lack of any response in the IRS call data, consistent with taxpayers mostly using the IRS call line for tax compliance purposes. This bill temporarily resolved what was commonly called the “fiscal cliff” debate, and increased the top marginal tax rate on long-term capital gains from 15 percent to 20 percent.²⁷ Our evidence strongly suggests, therefore, that individuals search for information both in response to policy changes and in response to potential policy changes signaled by political events.

7 Regression Analysis of the Impact of Events

We start with the detrended and seasonally *unadjusted* time series. The regression procedure we use, outlined in Section 5 and in the appendix, explicitly controls for variation due to weekly and yearly seasonality. The notes to Table 2 describe the events we study, which are also discussed in the previous section.²⁸

Table 2 describes the estimated impact of events on information search through Google, Wikipedia, and the IRS web page. To examine statistical significance, we use a non-parametric permutations test based on the test proposed by Gelbach et al (2013) for single-firm event studies. This test is based on the comparison of the effects of the events we study with the estimated total effect of “placebo” events occurring on arbitrary days. The appendix describes this procedure in more detail. The resulting *p*-value corresponds to the probability

²⁷ As with many of our events, the focus of the ATRA was not the capital gains tax. Just as the presidential debates covered a broad spectrum of policy issues, the ATRA was an omnibus bill that combined many different tax and non-tax initiatives into a single piece of legislation. For a discussion of the effect of omnibus tax legislation, see Hoopes (2014).

²⁸ There is modestly elevated information search via Google, Wikipedia or both, on a few dates that we do not include in the analysis. These dates include January 18, 2004 (a State of the Union address delivered by George W. Bush advocating the extension of capital gains tax cuts), November 2, 2004 (re-election of George W. Bush), May 16, 2006 (the extension of the 2003 capital gains tax cuts), January 27, 2010 (a State of the Union address by Barack Obama advocating a cut in capital gains taxes for some taxpayers), and March 23, 2010 (passage of the Affordable Care Act, which included a “net investment income tax”). In each case, the response of searches is qualitatively similar—a spike in searches that fades in three to four days—but notably smaller than the events we do include in the formal analysis.

that increases in information search of the magnitude we observe would have occurred at random during the event window we specify. The results suggest that the events we study each have a large and significant impact on information search through Google and Wikipedia, but not through the IRS call line. When we estimate the overall impact of the event rather than examining a single-day impact, the release of Mitt Romney's tax return surpasses the passage of ATRA as the one-time event that generated the most taxpayer information search, through both Google and Wikipedia. This occurs because the effect of ATRA passage declined more rapidly, fading in two days instead of four. Notably, ATRA passage was anticipated in the days leading up to January 2. The bill passed Congress on January 1st, and the anticipated increase in the top capital gains rate, likely generated income shifting from 2013 to 2012, which may have also caused information search. From Figure 4, panel B, we can see that elevated information search occurred on the two days prior to ATRA passage, December 31 and January 1st. Adding these dates to the event window increases the estimated impact of ATRA by approximately 5 standard deviations for Google and Wikipedia, but changes little else.

8 Relating Tax Information Search to Stock Market Activity

An advantage of focusing on capital gains taxation is the availability of high-frequency data on sales of stock. Relating these data to data on information searches holds the promise of better understanding the causal connections between information search and capital-asset-related behavior. After all, we are interested not only in what causes people to search for information, but also in to what extent the acquisition of information affects behavior, in this case behavior related to the sale, purchase, and holding of capital assets.

8.1 The Lead-Lag Relationship to Volume, Volatility, and Market Return

Our first strategy is to examine the lead-lag relationship between measures of behavior and information searches. If searches lead behavior, then we have reason

to pursue the idea that the information obtained affected subsequent decisions. We investigate two data series, both of which represent general stock market activity.

The market-related measure we use is trading volume. We obtain the dollar value of shares traded from all publicly listed firms from the Center for Research in Security Prices (CRSP variable VOL), which we use as a measure of broad market activity.²⁹ We then detrend the measure using a Hodrick-Prescott filter, and include the log of daily trading volume as an independent variable in the regression described in equation (1). For our baseline regression, we include five days of leads and lags of log trading volume. These regressions control for the events we study and for weekly and yearly seasonality. Columns (1) through (3) of Table 3 report the coefficients on standardized data. See Online Appendix Figure A.2 for a graphical depiction of the lead-lag estimates.

When trading volume is high, individuals seek more information about capital gains taxes. The standard deviation of log trading volume is 0.2233, so the estimates in table 3 suggest that a one-standard-deviation increase in trading volume predicts a 0.064-standard-deviation increase in Google searches, and a 0.068 standard deviation increase in Wikipedia Page views. These effects are significantly different from zero at conventional significance thresholds ($p = 0.068$ or 0.038 for Google, $p = 0.055$ or 0.027 for Wikipedia, depending on whether one employs classical or Newey-West standard errors, respectively).³⁰ We find that searches on date t are not significantly related to trading volume on date $t+1$ or beyond. All of the association of market movement with information search occurs on the same date as the market movement, or one day before.³¹ As a

²⁹ Note that these transactions include many where the buyer and/or seller is not subject to capital gains taxes, may not be triggered by a human trader capable of information search or where the asset does not have an accrued gain that will be subject to taxation. These possibilities do not threaten our identification so long as the percentage of transactions that is not subject to capital gains taxes does not change substantially and systematically from day to day.

³⁰ This result suggests that a non-constant fraction of people trading consider the tax implications of the trade, as a constant fraction of traders would imply a coefficient of 1. It seems likely that on high-volume days, a smaller percentage of traders may be carefully researching the tax consequences of the trade, resulting in the fraction of tax-informed traders being non-constant across time.

³¹ If we include only one lead and lag of log trading volume we estimate that the coefficient on the first lead and lag are

result, we are not able to determine from these data whether individuals seek information on capital gains taxes primarily before or after they make a decision regarding the sale of a capital asset. If they seek information before making a decision, they do so less than a day in advance, as far as these data are able to tell us.

We are sensitive to the possibility that some of the variation in these results may be driven by behavior during the 2008 stock market crash, a period of extremely high stock price volatility and trading volume. Columns (4) and (5) of Table 3 provide the estimates of the same regression specification, but limiting the inclusion of market variables to the period from September 2008 to February 2009. Columns (6) and (7) provide the estimates of the same regression specification, instead excluding this extraordinary period. For both Google and Wikipedia data, the estimated relationship between search and market activity during the extraordinary period is much larger, although imprecisely estimated. However, the estimated effects we have described survive when the extraordinary period is excluded from the sample.

8.2 Do taxpayers search for capital gains tax information before they trade?

In this section we pursue an alternative indicator of the taxpayer demand for capital gains tax information that is itself based on observed search volume. In particular, we use Google Trends to obtain a measure of the volume of searches for phrases related to personal investment advice such as “stock advice”, “should buy stock,” “should sell stock,” and “investment advice.”³² In so doing, we hope to learn more about the timing of the relationship between search for tax-related

slightly larger than in our preferred specification. Including up to 14 leads and lags does not change the estimates from our preferred specification. If we include one lag of log trading volume *without* contemporaneous log trading volume, the estimated coefficients are nearly identical to the ones in Table 3. If we include one lead of log trading volume *without* contemporaneous log trading volume, we do estimate a statistically significant relationship similar in magnitude to the contemporaneous effect. This last result is almost surely driven by autocorrelation in trading volume.

³² Unfortunately, Google Trends produces data only when there is sufficient search volume for a given search term. So, while more specific search terms would be useful (i.e., “how to short Facebook”, or, “should I sell my capital losses”), there is no available data on these searches.

information and decisions about whether to buy, sell, or hold assets. Note that, although Google searches for investment advice can lead one to websites that purport to provide such advice, searches for “capital gains tax” generally do not.

Table 4 shows the results of including the investment advice measure as an explanatory variable in a regression on searches for capital gains tax information, both in addition to stock trading volume and as a replacement for it.³³ Several results of interest emerge.

First, consider the results for Google searches. When investment advice search volume is included as an alternative to stock trading volume, the same positive contemporaneous association appears, and the effect of a one standard deviation change on capital gains tax related searches is 40 percent lower. In a “horse race” when both stock trading volume and investment-advice search volume are included as explanatory variables, unsurprisingly the statistical significance of the former declines because of the high correlation between the two variables. In the horse race, investment advice search volume wins, retaining its significance while stock trading volume loses its own significance. Strikingly, the one-day lag *and* *lead* values of stock investment advice search volume are both significantly positively associated with capital gains tax search volume. This estimated leading association is consistent with a story that taxpayers first recognize the need for information regarding buying or selling stock and, in the process, learn that relevant to this decision are the tax consequences.³⁴

8.3 Google Correlate

While the above mentioned stock-related terms that we selected (e.g., “stock advice”) were correlated with searches for capital gains taxes in

³³ Online Appendix Figure A.3 graphically depicts the lead-lag structure estimated when stock advice search volume is included in the absence of stock market trading volume.

³⁴ The key advantage of the stock investment advice search volume variable is that it likely captures the extent of taxpayer demand for information for which capital gains tax knowledge is crucial. One potential disadvantage is that in regression analysis one may pick up any shocks that affect all Google Searches. To that point note that in Table 5 the same pattern of results also applies when the volume of Wikipedia searches is the dependent variable, with one exception: the lead relationship, although positive, does not reach statistical significance.

predictable ways, Google, through its application Google Correlate, also has the ability to provide a list of search terms most highly correlated over time with any given user provided search term. Consistent with our other evidence that taxpayers simultaneously search for information about capital gains taxes and investing, among the top 20 search terms related to “capital gains tax” are “stock purchase,” “investing,” “td waterhouse,” “fidelity mutual funds,” and “mutual funds.”³⁵

8.4 Capital Losses

Because the capital gains tax rules related to the sales of assets with capital losses are especially important for many tax minimization strategies, such as loss harvesting, we also construct two measures of information search related specifically to capital losses. The first is weekly Google searches for the phrase “capital loss,” and the second is monthly searches related to capital losses using the search functionality on the IRS website. We obtain the latter measure directly from the IRS.³⁶ These time series are plotted in Figure 5.

As with searches for capital gains, much of the variation is seasonal: people tend to search for information on capital losses during tax season. There is also typically an increase in searches in December of each year, when some taxpayers “harvest” capital losses to reduce their tax liability. Of particular interest is the fact that searches for information on capital losses increased dramatically during October 2008, and surged even further in December of 2008. When the crash began in October, investors began to research the tax implications of the unrealized or realized losses they had sustained, perhaps evaluating the merits of pulling their wealth out of (or investing in) the declining stock market. Furthermore, for an investor who had lost money in the crash, harvesting capital

³⁵ <http://www.google.com/trends/correlate/search?e=capital+gains+tax&t=weekly&filter=capital+gains+tax&p=us#default.30> displays the full results. We thank Hal Varian for suggesting the use of Google Correlate.

³⁶ The IRS provided to us, at a monthly level, the most frequent search terms, and the number of searches for all terms with the word “capital”, and “loss”, in the search.

losses before the end of 2008 could reduce the taxpayer's tax liability substantially.³⁷ The heightened importance of loss harvesting from the stock market crash apparently caused even more information search at year-end.

For both Google and IRS searches, we can also see that searches for information on capital losses *during tax season* are higher after the 2008 crash than before or for three years after the crash. For several years after the crash investors realized capital losses with greater frequency than before the crash. At the same time, information search for tax compliance purposes also increased after 2008. This constitutes suggestive evidence of spillovers between macroeconomic changes and information search for the purpose of tax compliance.

9 Lessons from Aggregate Data on IRS Website and Call Line Usage

Figure 6 plots time series data on the use of the IRS website and call-line usage that are not specific to capital gains taxes. In order for their effects to be visible in broader measures of information search, events must significantly change the importance of understanding taxes at a given time, for a large number of taxpayers. Unsurprisingly, the IRS.gov domain and call line experience elevated traffic during tax season. Note that the amount of traffic experienced during tax season increases over time for the web page, probably due to increased use of the Internet by taxpayers. The call line, in contrast, decreased in usage from 1999 to 2008, and then experienced a resurgence. Together these results suggest that the Internet has not completely crowded out the use of the call line.

Most interestingly, we observe an abnormal surge in visits to IRS web pages and calls line traffic in May of 2008. Search volume usually drops sharply after the mid-April tax deadline, but in 2008 it remained high throughout the month of May. Among the top five searches on the IRS web page during this month were “stimulus,” “rebate,” and “stimulus check,” and many of the top

³⁷ This would naturally require that the investor had capital gains to offset with these losses.

pages viewed also dealt with the stimulus rebates. The tax rebates enacted by the Economic Stimulus Act of 2008 led millions of taxpayers to visit the IRS web page to investigate how the federal stimulus program affected them. From more detailed data on the use of the call line, we know that the IRS received over 50 million phone calls on stimulus checks, corresponding to at least 11 million unique taxpayers. This corroborates the evidence provided by Sahm, Shapiro, and Slemrod (2012) that these rebate checks were a salient form of economic stimulus with a relatively high (compared to a reduction in employer withholding rates) marginal propensity to consume. Using the intuition of rational attention, it makes more sense for individuals to seek information in May about their stimulus check if they intend to spend it compared to if they intend to save it.

10 Conclusions

It is well-established that in general taxpayers know little about the US income tax, and have systematic misperceptions. Given that acquiring information is costly, it may be optimal for individuals to learn only *if* the expected return is high enough and only *when* the information is most useful, known as rational attention. Because people are learning—and forgetting—things all the time, the process of net information acquisition is critical to a dynamic understanding of tax salience. Using newly available IRS administrative data and publicly available information on Google and Wikipedia searches, this paper establishes that people seek information about the US income tax in systematic ways that are consistent with the idea of rational attention. When policies change or seem likely to change, when filing deadlines or time notches loom, people turn to online resources like Google, Wikipedia and the IRS website, as well as traditional information resources like the IRS telephone hotline, to learn how the tax code affects them. In addition, people search for tax-related information when newsworthy events make taxes more salient, and in so doing they may incidentally obtain information relevant to their own decisions.

When policy or news events generate exogenous shocks to the demand for information, the responsiveness of information search to the event occurs remarkably quickly: search behavior usually spikes on the same day as the event, and falls back to baseline within three or four days. This timing pattern suggests that either 1) taxpayers acquire information about capital gains taxes much more quickly than envisioned in macroeconomic models of rational attention, or 2) the public's attention span is very short. In relation to the first point, we note that in the US, Google searches for “inflation”—the topic that most macroeconomists have in mind when thinking of rational attention—exhibit similar 3-4 day spikes following the release of inflation reports from the Bureau of Labor Statistics (BLS). These spikes are especially pronounced when BLS reports document large increases in the price level. A short public attention span could be caused by convex costs of attention to a particular topic, or it could mean that the discrete spikes in attention we observe are due to salience rather than rational attention.

We also present somewhat weaker evidence regarding to what extent acquiring information about taxes leads individuals to alter behavior. We show that individuals sought information about capital losses at year ends, especially after the stock market crash of 2008, since harvesting losses provides an opportunity to reduce one's tax liability. This finding is consistent with many investors not knowing the rules governing capital losses well enough to confidently apply them to make their behavior tax efficient in the wake of the crash. Information acquisition is a necessary component of the response to tax incentives, especially where more obscure details of the tax code are concerned.

Our attempt to learn about taxpayer information search is limited because we cannot account for all sources of information. For example, we cannot observe taxpayers obtaining information from paid professionals. About 60 percent of taxpayers do use a paid preparer when they file a tax return. Some of our measures may, however, capture financial advisors' information search. Further,

we cannot rule out “learning by doing,” whereby individuals learn about tax incentives while filling out a return and then improve their future decisions. Nor can we observe learning through social networks (as documented in Alstadsæter, Kopczuk, and Telle, 2010). As in all similar studies, we also cannot quantify the *amount* of learning that occurs via acquiring any one piece of information, especially because we do not observe which specific taxpayers are searching for information at a given time. Finally, taxpayers for whom capital gains taxes matter are generally wealthier than the typical taxpayer. These individuals may be more sophisticated in their use of information than the typical taxpayer, and many can afford to pay for high-quality information from financial advisors. For these reasons, future work should examine patterns in information search about other policies and, if possible, with individual-level micro data.

Appendix

Obtaining and Analyzing Google Trends Data

When querying Google Trends, the user provides 1) search terms, 2) a geographical window, and 3) a time range. As we are studying a tax issue within the United States, for all data used in this paper, the geographical range is specified to be the United States.

When the number of overall searches for a given term is too low, Google Trends will report an SVI of zero, or will report daily at a weekly or monthly level (as opposed to daily SVI data). We encounter this issue if we query searches for “capital gains tax” alone. As such, we use a set of search terms to maximize our sample period for which we are able to obtain daily data. The search terms we include are, according to Google Trends itself, highly correlated with searches for “capital gains tax.” The set of search terms is the following:

- Capital gains tax
- Capital gains tax rate
- Capital gains taxes
- Capital gains calculator
- Capital gains
- Capital gains rate

We have verified that 1) the daily time series of SVI for simply “capital gains tax” is virtually identical to the one from the broader set of search terms, but with fewer missing data, and 2) we obtain nearly identical results for event studies and market movement effects if we use simply searches for “capital gains tax” as our left-hand-side variable instead of the broader set of search terms, but with slightly larger standard errors (reflecting the decreased number of observations).

For the same reasons as above, we use multiple search terms related to stock advice. These are:

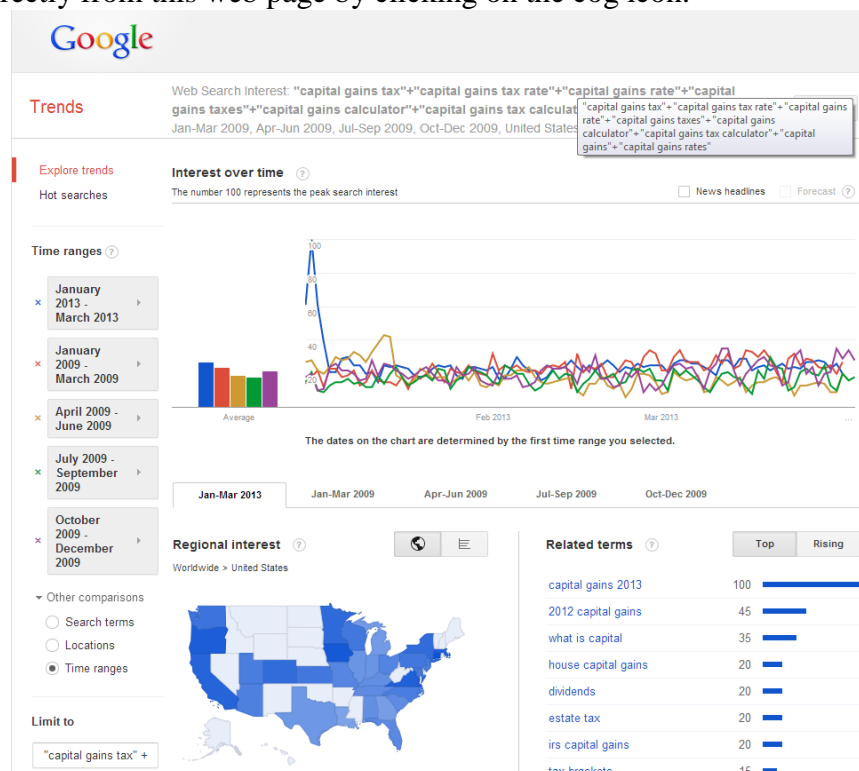
- Stock advice
- Stock market advice
- Stock tips
- Stocks to buy
- Stock to buy
- Stocks to sell
- Should buy stock
- Should sell stock
- Investment advice
- Investment tips

By default Google Trends provides weekly data when a user downloads a

time series longer than three months. In order to access daily search volume data, one must query Google Trends in three month intervals. Fortunately, one can query several (up to five) three-month periods at once.

In order to obtain daily data while maintaining proper scaling of the variable across the entire time series, however, one must query Google Trends very carefully. The data are scaled so that SVI takes the value of 100 on the day with the highest propensity to search out of any date range and search terms. For example, if January 2, 2013 were the day with the highest propensity to search for capital gains tax terms (as it is), SVI would equal 100 on that day if January 2, 2013 were in the period provided by the user. In order to obtain daily data that is properly scaled for the full sample period, one must first find the single day with the highest search volume, and then include a time period containing that day along with sets of other three-month periods.

To get a properly scaled daily time series, we therefore include the time period January 1, 2013-March 31, 2013 in every single one of our queries for capital gains tax SVI, along with other three-month periods, until we obtain data for our entire sample period. The figure below shows an example of what such a query would look like to pull daily data for the year 2009. One can then download the data directly from this web page by clicking on the cog icon.



Note: Google and the Google logo are registered trademarks of Google Inc., used with permission.

Estimation Procedure

Recall that the equation to be estimated is

$$I_{it} = \sum_k [\beta_{k0}F_{kt} + \beta_{k1}F_{k(t-1)} + \dots + \beta_{k4}F_{k(t-4)}] + x_t'\gamma + f_i(DoY_t) + \delta_{DoW,i} + u_{it} \quad (1)$$

In order to consistently estimate equation (1), which contains a non-linear function and a set of linear covariates, we use the double residual regression method suggested by Robinson (1988), and discussed in Hardle and Linton (1994). This first step of this three-step estimator consists of several non-parametric regressions of the following form:

$$E[X_t] = f_i(DoY_t),$$

where X_t is the dependent variable or one of the linear covariates, i.e. I_{it} , F_{kt} , x_t or $\delta_{DoW,i}$. We then obtain the residuals from this regression, which we denote with star superscripts. The residuals represent the component of I_{it} , F_{kt} , x_t or $\delta_{DoW,i}$ not correlated with the general within-year pattern represented by $f_i(DoY_t)$. The second step estimates the linear components of the model consistently using ordinary-least-squares regression on these residuals:

$$E[I_{it}^*] = \sum_k [\beta_{k0}F_{kt}^* + \beta_{k1}F_{k(t-1)}^* + \dots + \beta_{k4}F_{k(t-4)}^*] + x_t'^*\gamma + \delta_{DoW,i}^* \quad (2)$$

The third step regresses the residuals from the estimation of equation 2—denoted by I_{it}^{**} —non-parametrically on DoY_t to obtain a consistent estimate of the function $f_i(DoY_t)$.

$$E[I_{it}^{**}] = f_i(DoY_t) \quad (3)$$

The first and third steps involve non-parametric regression, which Robinson (1988) suggests implementing via kernel-weighted local polynomial regression. In this paper, we use local linear regressions with a Gaussian kernel density function and a bandwidth of four days.³⁸ This procedure estimates a weighted

³⁸ This bandwidth and kernel density function applies to non-parametric estimations from the third stage of the procedure, and nonparametric estimations involving continuous variables in the first stage. For discrete variables in the first stage, such as event dummy variables, we use a bandwidth of zero, to reflect that there should be no smoothing at this stage. Another complication is that traditional statistical packages and programs will implement the smoothing on a linear variable, while day-of-the-year is a cyclical variable. That is, January 1 and December 31 should be adjacent for smoothing purposes. Ignoring this problem results in a discontinuity in the seasonal pattern between these two days. We eliminate the discontinuity by estimating the seasonal pattern twice: once where the discontinuity is imposed at January 1, and a second time where the discontinuity is imposed at the 200th day of the year (July 19th). Then, we replace the 10 days around January 1 from the first estimation with these days from the second estimation.

OLS regression of the dependent variable on DoY at each value of DoY in the data, where the weights are determined by the kernel density function and the bandwidth. The bandwidth was selected to visually match the fixed-effects estimator of the function $f_i(DoY_t)$, but a data-driven choice of bandwidth—specifically selecting the bandwidth that minimizes the conditional weighted mean integrated squared error—yields nearly identical results. The estimates we present are also virtually unchanged by varying the degree of the local polynomial, the bandwidth, and/or our choice of kernel function, with the exception that a wider bandwidth results in a smoother function that no longer resembles the fixed-effects estimates and a narrower bandwidth results in a more jagged function.

Causal Inference for Event Studies

Here we describe the procedure used to obtain p-values for Table 2 in the body of the paper. The method we use is based on Gelbach et al (2013), who propose a similar estimator for causal inference over single-firm event studies. Most methods for robust causal inference in event studies consider the case where there are multiple firms affected by the same event, or the same type of event. This setting differs from ours because we have only one time series over which to examine the impact of events.

First, we estimate our statistical model using the procedure outlined in the previous section. After obtaining the residuals from this regression, we construct placebo event effects. Eligible placebo event dates consist of any date in the sample period that does not occur within five days of an event we study in Table 2, so that the five-day event windows never overlap. For every eligible placebo event date t , we construct a placebo event effect by adding the residual from the regression from t to $t+4$.

One can think of the distribution of placebo events as an approximation to the distribution of five-day effects in the absence of an event that significantly influences information search. The p-value reported in Table 2 is the fraction of placebo event effects that exceed the estimated effect of an event we specify, such as the release of Mitt Romney’s tax return. For example, the release of Mitt Romney’s 2011 tax return generated an estimated cumulative effect on Google searches of 7.17. Placebo event effects exceeded 7.17 in 0.3 percent of the time, so our p-value for this estimate is 0.003. The p-values we obtain using this method are slightly larger than p-values we would obtain from a classical F-test for the joint significance of coefficients for event days, which should be expected if there is some serial correlation in the data.

References

- Alstadsæter, A., W. Kopczuk, and K. Telle. (2010). Social Networks and Dynamics of Tax Avoidance: Evidence from a Well-defined Norwegian Tax Shelter. Working paper.
- Blendon, R.J., S.R. Pelletier, M.D. Rosenbaum, and M. Brodie. (2003). "Tax Uncertainty." *Brookings Review* (Summer): 28-31.
- Brown C.V. (1968). "Misconceptions about Income Tax and Incentives." *Scottish Journal of Political Economy* 16(2): 1-12.
- Chang, E. C., and J.M. Pinegar. (1986). "Return Seasonality and Tax-loss Selling in the Market for Long-term Government and Corporate Bonds." *Journal of Financial Economics* 17(2): 391-415.
- Chetty, R., A. Looney, and K. Kroft. (2009). "Salience and Taxation: Theory and Evidence." *American Economic Review* 99(4): 1145-1177.
- comScore. (2005). comScore Releases July 2005 U.S. Search Engine Rankings. <http://ir.comscore.com/releasedetail.cfm?ReleaseID=245083>. August 19, 2005.
- comScore. (2007). comScore Releases November U.S. Search Engine Rankings. <http://ir.comscore.com/releasedetail.cfm?ReleaseID=283126>. December 21, 2007.
- comScore. (2009). comScore Releases June 2009 U.S. Search Engine Rankings. <http://ir.comscore.com/releasedetail.cfm?ReleaseID=397028>. July 16, 2009.
- comScore. (2013). comScore Releases Oct. 2013 U.S. Search Engine Rankings. <http://ir.comscore.com/releases.cfm?Year=2013&ReleasesType=&PageNum=1>. November 13, 2013.
- Connor, C.E., H.E. Egeth, and S. Yantis. (2004). "Visual Attention: Bottom-Up versus Top-Down." *Current Biology* 14: R850-R852.
- Da, Z., J. Engelberg, and P. Gao. (2011). "In Search of Attention." *Journal of Finance* 66(5): 1461-1499.
- de Bartoleme, C.A.M. (1995). "Which Tax Rate Do People Use: Average or Marginal?" *Journal of Public Economics* 56(1): 79-96.
- Delli Carpini, M.X. and Keeter, S. (1996). *What Americans Know about Politics and Why It Matters*. New Haven, CT: Yale University Press.
- Drake, M.S., D.T. Roulstone, and J.R. Thornock. (2012). "Investor Information Demand: Evidence from Google Searches around Earnings Announcements." *Journal of Accounting Research* 50(4): 1001-1040.
- Fujii, E.T., and C.B. Hawley. (1988) "On the Accuracy of Tax Perceptions." *The Review of Economics and Statistics* 70(2): 344-347.

- Gelbach, J.B., E. Helland, and J. Klick. (2013). "Valid Inference in Single-Firm, Single-Event Studies." *American Law and Economics Review* 15(2): 495-541.
- Goldin, J. (2014). "Optimal Tax Salience." Unpublished.
- Graham, J.R., M. Hanlon, T. Shevlin, and N. Shroff. (2014). "Tax Rates and Corporate Decision Making." Unpublished.
- Hardle, W. and O. Linton. (1994). "Applied Nonparametric Methods." In *Handbook of Econometrics Volume IV*, edited by R.F. Engle and D.L. McFadden, pp. 2295-2239. Elsevier.
- Hoopes, J. (2014). "Financial Accounting Consequences of Temporary Tax Law: Evidence from the R&D Tax Credit." Unpublished.
- Ivković, Z., J. Poterba, and S. Weisbenner (2005). "Tax-Motivated Trading by Individual Investors." *American Economic Review* 95(5): 1605-1630.
- Liebman, J.B. and R.J. Zeckhauser. (2004). "Schmeduling." Unpublished.
- Keppo, J., G. Moscarini, and L. Smith. (2008). "The Demand for Information: More Heat than Light." *Journal of Economic Theory* 138(1): 21-50.
- Muth, J. F. (1961). "Rational Expectations and the Theory of Price Movements." *Econometrica* 29(3): 315-335.
- Poterba, J.M. and S.J. Weisbenner. (2001). "Capital Gains Tax Rules, Tax Loss Trading, and Turn-of-the-Year Returns," *Journal of Finance* 56(1): 353-368.
- Reck, D.H. (2014). "Taxes and Mistakes: What's in a Sufficient Statistic?" Unpublished.
- Reis, R. (2006). "Inattentive Consumers." *Journal of Monetary Economics*, 53(8), 1761-1800.
- Robinson, P.M. (1998). "Root-N-Consistent Semiparametric Regression." *Econometrica* 56(4): 931-954.
- Sahm, C.R., M.D. Shapiro, and J.B. Slemrod. (2010). "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications." In *Tax Policy and the Economy, Volume 24*, edited by Jeffrey R. Brown, pp. 69-110. Chicago: Chicago University Press.
- Sheffrin, S.M. (1994). "Perceptions of Fairness in the Crucible of Tax Policy." In *Tax Progressivity and Income Inequality*, edited by Joel Slemrod, pp. 309-334. Cambridge: Cambridge University Press.
- Sims, C. A. (2003). "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3): 665-690.
- Slemrod, J. (1995). "The Compliance Cost of Income Taxation, and How to Reduce It," National Commission on Economic Growth and Tax Reform (Kemp Commission), Oral testimony given July 12, 1995.
- Slemrod, J., Christian, C., London, R., and Parker, J. A. (1997). "April 15 Syndrome." *Economic Inquiry* 35(4): 695-709.

- Slemrod, J. (2004). "New Estimates of the Compliance Cost of Income Taxation," Committee on Ways and Means, Subcommittee on Oversight, Hearing on Tax Simplification, invited written testimony given June 15, 2004.
- Slemrod, J. (2013). "Buenas Notches: Lines and Notches in Tax System Design." Unpublished.
- Starks, L.T., L. Yong, and L. Zheng. (2006). "Tax-Loss Selling and the January Effect: Evidence from Municipal Bond Closed-End Funds." *Journal of Finance* 61(6): 3049-3067.
- Tax Foundation. (2010). Federal Capital Gains Tax Collections, 1954-2009. <http://taxfoundation.org/article/federal-capital-gains-tax-collections-1954-2009>. Accessed 5/7/2014.
- Theeuwes, J. (1991). "Exogenous and Endogenous Control of Attention: The Effect of Visual Onsets and Offsets." *Perception & Psychophysics* 49(1): 83-90.

TABLE 1: TESTABLE IMPLICATIONS OF THEORIES OF INFORMATION AND ATTENTION

Table 1: Testable Implications of Theories of Information and Attention

Behavioral assumption	Should information search respond to...					
	time notches?	macroeconomic change?	enactment of policy changes?	implementation of policy changes?	filing deadlines?	news events?
Full information	No	No	No	No	No	No
Full ignorance	No	No	No	No	No	No
Rational expectations	No	No	Yes, instantly	No	No	No
Rational attention	Yes	Yes	Yes	Yes	Yes	Unclear*
Exogenous Attention (Salience)	No	No	Yes, while in the news	Yes, while in the news	Yes**	Yes
Summary of Our Findings	Yes: Dec 31 time notch, esp. capital losses (due to loss harvesting)	Yes: day-of stock market trading volume, response to stock market crash	Yes, in 3-4 day window: ATRA, presidential debates and elections	Yes: May 2008 stimulus	Yes: seasonality due to tax season	Yes, in 3-4 day window: release of Mitt Romney's tax returns

Notes: *News events unrelated to the other types of events listed here may inspire attention rationally due to the desire to be informed about current events, but they will not inspire rational attention for the purpose of improving economic decisions, such as capital gains realizations or tax compliance. ** Our answering yes is based on the observation that, as the deadline approaches, news stories and private discussions of tax payments are common.

TABLE 2: ESTIMATING THE IMPACT OF EVENTS ON
INFORMATION SEARCH FOR CAPITAL GAINS TAXES

	Google Searches	Wikipedia Page Views	Calls to IRS
Obama/McCain Debate	12.08 93.89 (0.001)	5.36 2662.51 (0.016)	-0.22 -16.17 (0.718)
Obama Elected	16.08 124.66 (<0.001)	5.87 2914.76 (0.012)	-0.16 -12.10 (0.655)
Mitt Romney's 2010 tax return released	20.72 161.00 (<0.001)	34.96 17351.70 (<0.001)	0.38 -559.08 (0.196)
Mitt Romney's 2011 tax return released	7.16 55.69 (0.003)	12.65 6277.82 (<0.001)	-0.34 -32.47 (0.810)
Obama/Romney Debate	9.39 72.99 (0.002)	9.88 4904.11 (0.002)	0.25 -17.69 (0.258)
Obama Re-Elected	8.94 71.02 (0.002)	14.56 7228.50 (<0.001)	-0.31 -14.27 (0.794)
American Taxpayer Relief Act signed	15.31 118.98 (<0.001)	9.62 4776.24 (0.003)	-6.17 -188.52 (0.989)
Sample Period	Jan 2, 2004— March 21, 2013	Dec 9, 2007— March 31, 2013	Jan 2, 2002— March 31, 2013
Number of Days	3364	1938	4100

Notes: The event dates listed as dependent variables represent, in order, Oct 15, 2008 (Presidential debate between Barack Obama and John McCain), November 4, 2008 (Barack Obama elected President), January 24, 2012 (Mitt Romney releases his 2010 tax return), September 21, 2012 (Mitt Romney releases his 2011 tax return), Oct 16, 2012 (Presidential debate between Barack Obama and Mitt Romney), Nov 6, 2012 (Barack Obama re-Elected), Jan 2, 2013 (American Taxpayer Relief Act signed into law, includes an increase in top capital gains tax rate). For each event, the top number reports the cumulative information search attributed to this event, in daily standard deviation units, added over the five-day event window. The second number reports the same estimate in the original units of the search volume measure, i.e. the Google Trends index or, Wikipedia page views, or number of calls to the IRS. The bottom number, in parenthesis, reports the p-value from a one-sided non-parametric permutations test of the hypothesis is that the event had positive impact on information search (refer to the appendix for details), i.e. that the variation in information searches over the event window is insignificant. The estimation controls for variation due to yearly and weekly seasonality.

TABLE 3. INFORMATION SEARCH AND STOCK MARKET ACTIVITY

Sample Period	Jan 2, 2004— Dec 29, 2012	Dec 9, 2007— Dec 29, 2012	Jan 7, 2002— Dec 29, 2012	Sep 1, 2008— Feb 28, 2009	Sep 1, 2008— Feb 28, 2009	Jan 2, 2004— Dec 29, 2012, except Sep 1, 2008— Feb 28, 2009	Dec 9, 2007— Dec 29, 2012, except Sep 1, 2008— Feb 28, 2009
Dependent Variable:	Google Searches (Standardized)	Wikipedia Page Views (Standardized)	Calls to IRS (Standardized)	Google Searches (Standardized)	Wikipedia Page Views (Standardized)	Google Searches (Standardized)	Wikipedia Page Views (Standardized)
Log Trading Volume	0.287 (0.156)* [0.138]**	0.306 (0.159)* [0.139]**	0.043 (0.075) [0.069]	0.763 (0.801) [0.789]	0.926 (0.470)* [0.352]**	0.216 (0.159) [0.128]*	0.385 (0.171)** [0.156]**
Log Trading Volume Lag 1	-0.051 (0.146) [0.120]	-0.072 (0.151) [0.100]	0.099 (0.070) [0.044]**	-0.097 (0.735) [0.599]	0.263 (0.415) [0.251]	-0.111 (0.149) [0.126]	-0.093 (0.162) [0.115]
Log Trading Volume Lead 1	0.062 (0.177) [0.160]	-0.023 (0.181) [0.120]	-0.038 (0.085) [0.071]	-0.083 (0.921) [0.727]	-0.123 (0.553) [0.353]	0.105 (0.179) [0.159]	-0.05 (0.193) [0.131]
Number of Days	2252	1266	2753	124	124	2128	1142


Notes: The dependent variable and log trading volume are de-trended prior to estimation, and the dependent variable is standardized by dividing by the standard deviation of the detrended data. We control for weekly and yearly seasonality and the events in Table 2 (except those events that occur outside the sample period, since market data are not available past December 29, 2012). Weekends and holidays are omitted, since stock market data are not generated on weekends. The regression included five market-dated leads and lags of each measure. We only report one lead and one lag for brevity and clarity (See Online Appendix Figure A.1). The coefficients are similar if we use 14 leads and lags instead of 5, and none of the results change substantially with the inclusion of lags of the dependent variable. Classical standard errors are provided in round parentheses below point estimates. Newey-West standard errors allowing for up to 5 (market-dated) lag orders of autocorrelation are reported in square brackets. Newey-West standard error estimates are unchanged to two significant digits with 7 lag orders instead of 5. * indicates $p < 0.10$, ** indicates $p < 0.05$, and *** indicates $p < 0.01$. The standard deviation of log trading volume is 0.223.

TABLE 4: SEARCHES FOR STOCK ADVICE AND SEARCHES FOR CAPITAL GAINS TAX

Sample Period	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Jan 2, 2004— Dec 29, 2012	Jan 2, 2004— Dec 29, 2012	Jan 7, 2004— Dec 29, 2012	Jan 7, 2004— Mar 21, 2013	Dec 9, 2007— Dec 29, 2012	Dec 9, 2007— Dec 29, 2012	Dec 9, 2007— Dec 29, 2012	Dec 9, 2007— Mar 31, 2013
Dependent Variable:	Google Searches (Standardized)				Wikipedia Page Views (Standardized)			
Log Trading Volume	0.287 (0.156)* [0.138]**	0.176 (0.155) [0.143]	0.130 (0.155) [0.140]		0.306 (0.159)* [0.139]**	0.162 (0.157) [0.140]	0.178 (0.157) [0.136]	
Log Trading Volume Lag 1	-0.051 (0.146) [0.120]	-0.096 (0.145) [0.117]	-0.06 (0.146) [0.119]		-0.072 (0.151) [0.100]	-0.141 (0.147) [0.101]	-0.137 (0.148) [0.101]	
Log Trading Volume Lead 1	0.062 (0.177) [0.160]	0.022 (0.175) [0.159]	0.002 (0.175) [0.158]		-0.023 (0.181) [0.120]	-0.056 (0.177) [0.118]	-0.098 (0.178) [0.119]	
Stock Advice Search		0.031 (0.004)*** [0.008]***	0.008 (0.005) [0.007]	0.009 (0.005)* [0.007]		0.030 (0.004)*** [0.008]***	0.016 (0.006)*** [0.005]***	0.017 (0.005)*** [0.005]***
Stock Advice Search Lag 1			0.017 (0.006)*** [0.007]**	0.017 (0.006)*** [0.007]**			0.017 (0.008)** [0.006]***	0.015 (0.008)* [0.006]**
Stock Advice Search Lead 1			0.009 (0.005)** [0.005]**	0.008 (0.005)* [0.004]*			0.005 (0.006) [0.004]	0.003 (0.006) [0.003]
Number of Days	2252	2239	2200	2211	1266	1266	1266	1272


Notes: The dependent variable, log trading volume, and stock advice search volume are de-trended prior to estimation, and the dependent variable is standardized by dividing by the standard deviation of the detrended data. We control for weekly and yearly seasonality and the events in Table 2 (except those events that occur outside the sample period, since market data are not available past December 29, 2012). Weekends and holidays are omitted, since stock market data are not generated on weekends. The regression included five market-dated leads and lags of trading volume and stock advice search volume. We only report one lead and one lag for brevity and clarity (See Online Appendix Figure A.3). Classical standard errors are provided in round parentheses below point estimates. Newey-West standard errors allowing for up to 5 (market-dated) lag orders of autocorrelation are reported in square brackets. Newey-West standard error estimates are unchanged to two significant digits with 7 lag orders instead of 5. * indicates $p < 0.10$, ** indicates $p < 0.05$, and *** indicates $p < 0.01$. The standard deviation of log trading volume is 0.223. The standard deviation of Stock advice search volume is 4.458.

FIGURE 1: PLOTS OF THE RAW DATA



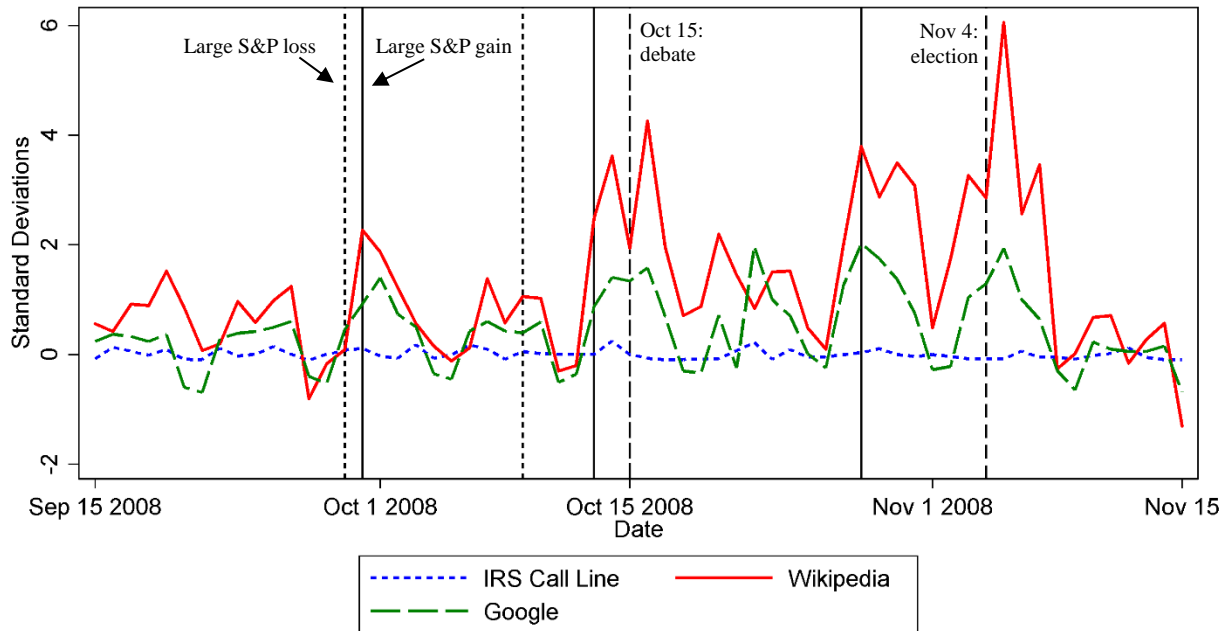
Notes: These graphs show the raw data for (in order) Google search volume, Wikipedia page views, and calls made to the IRS.

FIGURE 2: YEARLY SEASONALITY



Notes: Yearly seasonality is estimated using the smooth-fixed-effects method described in Section 5, with controls for the events in Table 2 and the market movements in Table 4. The thin dashed lines are the bounds of 95 percent confidence intervals.


FIGURE 3: CAPITAL GAINS TAX INFORMATION SEARCH IN OCTOBER 2008



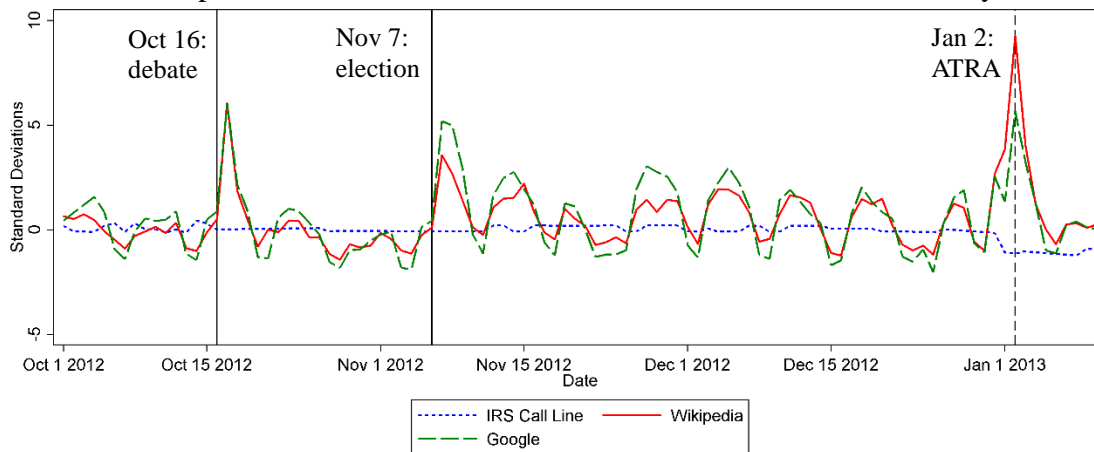
Notes: We plot the detrended, standardized, and seasonally adjusted data over a narrower time period to examine abnormal behavior in late 2008. The vertical dotted lines mark the worst three days of the stock market crash of 2008 according to the drop in the S&P 500 index: October 15 (-9.03 percent) , September 29 (-8.81 percent), and October 9 (-7.62 percent). The vertical solid lines mark the three largest gains in the S&P 500 during this volatile period: October 13 (11.58 percent), October 28 (10.79 percent) and September 30 (5.42 percent). October 15, 2008, marked with a vertical dashed line, was also the date of a Presidential Debate between Barack Obama and John McCain, in the course of which both candidates made proposals for changing capital gains tax rates. The other debates occurred on September 26 and October 7, and did not discuss capital gains taxes. The other vertical dashed line marks the 2008 Presidential election, November 4, 2008.

FIGURE 4: NARROW TIME WINDOWS OF SEARCH FOR TAX INFORMATION

Panel A. Capital Gains Tax Information Search in January 2012




Panel B. Capital Gains Tax Information Search in October 2012 – January 2013



Notes: In Panel A, we plot the detrended, standardized, and seasonally adjusted data (from Figure 4) over a narrower time period to examine the spike in January 2012. The vertical dashed line marks January 24, 2012, the date that Mitt Romney released his 2010 tax return. The spike a few days prior is coincident with a related news story wherein Mr. Romney told the press about his effective tax rate. In Panel B, the first solid vertical line marks the date of the Presidential debate on October 16, 2012, during which Barack Obama and Mitt Romney debated the merits of Romney's proposals for lowering capital gains tax rates. The second solid line marks the date of the Presidential election, November 7, 2013. The dashed vertical line marks the passage of the American Taxpayer Relief Act on January 2, 2013.


FIGURE 5: INFORMATION SEARCH ON CAPITAL LOSSES




Notes: The top panel plots weekly Google Trends data on searches for just “capital loss.” Google Trends data are missing for some dates in 2004 and 2005, when search volume was too low for Google to provide data. The second graph plots monthly searches in the search bar on the IRS home page for terms related to capital losses. The thick dotted vertical lines correspond to October 1 and December 2008, to delineate the period in which there was increased information search on capital losses due to the 2008 stock market crash. Solid vertical lines denote January 1st of each year, to highlight that a spike in search volume typically occurs at the very end of each year. Dashed vertical lines denote mid-April filing deadlines, to highlight the spike occurring in the run-up to the filing deadline.

FIGURE 6: OVERALL USE OF IRS INFORMATIONAL RESOURCES

Panel A: Visits and Page Views of all URLs on the IRS.gov Domain.



Panel B: Calls to IRS Line for Individuals



Notes: The dotted vertical lines delineate the month of May, 2008. Dashed vertical lines correspond to the mid-April filing deadline for each year. Panel A: Data are unavailable from November 1, 2003 to February 29, 2004. Panel B: The phone number for this line is 1-800-829-1040, a number made available numerous places on the IRS website and on the instructions for individual tax returns (form 1040). Unique callers are identified by the telephone number from which the call originates.

TAXPAYER SEARCH FOR INFORMATION: IMPLICATIONS FOR RATIONAL ATTENTION

Jeffrey L. Hoopes, Ohio State University, hoopes@fisher.osu.edu

Daniel H. Reck, University of Michigan, dreck@umich.edu

Joel Slemrod, University of Michigan, jslemrod@umich.edu

July 23, 2014

Forthcoming: American Economic Journal: Economic Policy

Abstract: We examine data on capital-gains-tax-related information search to determine when and how taxpayers acquire information. We find seasonal increases in information search around tax deadlines, suggesting that taxpayers seek information to comply with tax law. Positive correlations between stock market activity and search as well as year-end spikes in information search on capital losses when the market performs poorly suggest that taxpayers seek information for tax planning purposes. Policy changes and news events cause information search. These data suggest that taxpayers are not always fully informed, but that rational attention and exogenous shocks to tax salience drive taxpayer information search.

JEL Codes: D80, D83, H31, H24

*We wish to thank John Guyton and Ron Hodge of the Research, Analysis, and Statistics Division of the Internal Revenue Service, for helping us use the IRS aggregate administrative data. We thank the editor, Mark Duggan, and two anonymous referees, whose comments helped improve this manuscript. We are grateful for comments on an earlier draft received from Hal Varian, Raj Chetty, Alex Gelber, and participants at McMaster University, the University of Michigan Hosmer-Hall Luncheon series, the US Treasury Department Office of Tax Analysis, the University of Michigan Public Economic Summer Talks series, the University of Uppsala, and at the 2013 Public Economics United Kingdom Conference. The views expressed here are those of the authors alone, and do not reflect the views of the Internal Revenue Service.

Taxes, especially income taxes, can be complex and confusing. Despite a general awareness of this fact, the consequences of complexity and misunderstanding are not well-understood. Survey evidence suggests that many taxpayers do not understand basic tax concepts (Blendon et al, 2003), and the compliance cost of taxes, including learning enough about them to comply, is large (Slemrod, 1995; 2004). Given confusion surrounding tax incentives and tax law, many taxpayers might simply ignore or misperceive the incentives built into the tax code when they make decisions with tax consequences. Alternatively, taxpayers might collect information on the taxes that matter to them, and use this information to make tax-efficient decisions. Finally, taxpayers might learn just enough about tax policies to fill out their tax return and perhaps avoid an audit.

In this paper, we investigate these alternative views of taxpayer information search. We find evidence that is inconsistent with the notions that taxpayers are fully informed about the tax system, that they act in complete ignorance of the tax system, or that they gather information only for tax compliance purposes. The evidence we present suggests instead that at least some taxpayers employ rational attention to tax policies, in line with theories proposed by Sims (2003) and Reis (2006). In addition, we observe that exogenous shocks to tax salience from news events can substantially increase information search.

Modern technology has greatly expanded the accessibility of information. Any person with access to the Internet may, in a few minutes, learn at least something about the most obscure details of the tax code. Taxpayers undoubtedly do use these resources to seek information: people Google “tax” more often than they Google the names of public figures,¹ the IRS website has received on average 4.6 million visits *per day* since 2004, and the IRS call line has received on average 125 thousand calls per day since 1999. How tax knowledge matters hinges on 1)

¹ For example, Google searches for “tax” were more common than searches for “Barack Obama,” and were more common than “Obama” in all but the 2008 and 2012 election season.

how and when people seek out tax-related information, and 2) whether they change their behavior once they acquire it. In this paper, we address primarily the first of these questions, and provide some preliminary analysis of the second.

Our data consist of aggregate high-frequency time series on calls to the IRS toll-free phone number, aggregate visits to the IRS website, searches on the IRS website, measures of Google searches on tax-related terms, and views of tax-related web pages on Wikipedia. From these sources we collect data on information seeking regarding one specific tax-related topic: capital gains taxes. We select this topic because it is a perennially controversial policy issue, because data on the relevant taxed behavior, sales of capital assets, are available on a high-frequency basis, and because the American Taxpayer Relief Act of 2013 (ATRA) enacted a change in capital gains tax rates during our sample period.² We also note that capital gains taxes represent a non-trivial portion of total government revenues, so are interesting in their own right. In 2006, for example, capital gains realizations represented 5.96 percent of GDP, and taxes paid on those gains amounted to \$117 billion (Tax Foundation, 2010).

We study information-seeking around five different types of events: 1) time notches, 2) macroeconomic changes, 3) policy changes or the mention of potential policy changes, 4) filing deadlines or approaching filing deadlines, and 5) tax-related news events. Information search around the first three dimensions should be tied to the taxpayer's ability to make fully informed decisions affected by tax incentives. First, a particular date is often relevant for the incursion of tax liability or the tax efficiency of behavior. We call these *time notches*³ to indicate that tax liability can change abruptly, and non-incrementally, at certain dates, usually at year-ends. For example, the last date to affect one's capital gains tax liability in a

² We recognize that the type of people seeking information about capital gains tax are different from, say, those seeking information about the earned income tax credit. We expect that future work will examine whether the phenomena we document here also occur for other policies.

³ This terminology is taken from Slemrod (2013).

given year is December 31st. As the deadline approaches individuals may research the benefits from realizing a capital gain or loss in the current year as opposed to a future year. Second, macroeconomic changes may be associated with information search due to perceived changes in the importance of understanding the tax consequences of a particular behavior. For instance, asset owners may seek information about capital gains taxation as the stock market falls in order to understand the tax implications of their losses. Third, actual or potential changes in tax policy or tax enforcement may cause individuals to learn about the tax system, either to plan for the future or to make an informed voting choice. Taxpayers might also seek information on the tax system when completing their tax return in order to both fully comply with the law, as well as ensure they are using all the credits and deductions available. As such, the fourth dimension we study is filing deadlines and the onset of these deadlines, when individuals may learn about taxes due to approaching deadlines for filing tax returns. For the taxes we consider, the deadline affecting most taxpayers is in mid-April.

News events sometimes spark public interest in tax policy, and in this case individuals may search for tax information in order to develop an informed opinion on current events or out of curiosity. The fifth dimension we study consists of news stories such as the release of a public person's tax information. We can view these events as exogenous shocks to the salience of a particular aspect of tax policy. They have no direct bearing on an individual's tax burden, nor do they directly affect tax incentives. The search may, however, provide information that incidentally informs the searcher about her own tax situation or about the tax impacts her decisions. They may also affect other important decisions, such as voting behavior.

We observe strong seasonality in the search for information on capital gains taxes through all channels. Taxpayer information search increases substantially during the period commonly called "tax season," which runs from mid-January to

mid-April of each year. An even more pronounced spike in information search occurs very close to the filing deadline in mid-April. We also document the impact of several discrete events on taxpayer information search regarding capital gains taxes, through Google and Wikipedia. Presidential debates in which candidates discuss their proposals for capital gains taxes, the passage on January 2, 2013 of ATRA, the release of presidential candidate Mitt Romney's 2010 tax return in January of 2012, presidential elections, and policy changes all generate large and significant increases in taxpayer information search. In every case, these events cause a spike in taxpayer information search that fades within three to four days.

Next, we find that macroeconomic changes affect information search on capital gains taxes through Google and Wikipedia. We observe significantly elevated information search on days with large trading volume in the stock market. We also explore a measure of information search for personal investment advice generally, using Google searches. Daily searches for stock advice predict searches for capital gains tax on the same day, the previous day, and, in some specifications, one day in the future. This is the first evidence in the literature that some taxpayers investigate the tax consequences of an action while contemplating the action itself.

These patterns indicate that information search is consistent with a model in which individuals search when different events make understanding tax policy more important, i.e. when they increase the return to information search. However, these results are also consistent with individuals making choices and *then* researching the tax implications of those choices for their wealth. In an effort to provide evidence that more strongly suggests a causal relationship, whereby events cause individuals to search for information in order to make more informed decisions, we examine information search on capital losses. We document substantial elevation in information search related to capital *losses* at year-ends,

especially in 2008 (when capital losses from the stock market crash would have made taxpayers' more likely to be rebalancing their portfolios in order to take tax-efficient advantage of capital losses). Observing increases in information search just prior to the time notch is consistent with taxpayers seeking information on the tax system and using it to improve decision-making.

This paper provides the first-ever attempt at understanding how, and why, taxpayers search for information about government policy, specifically tax policy. We establish that taxpayers seek information in order to both comply with their tax obligation and to respond behaviorally to tax incentives. Taking taxpayer learning seriously has important implications for understanding the impact of tax policy and tax policy changes as well as for a full understanding of the nascent idea of tax salience.

2 Background

2.1 Public Understanding of Taxes

Political science research is fairly persuasive that voters know very little about the details of government. Delli Carpini and Keeter (1996), in a comprehensive survey of the political knowledge of voters covering several decades and hundreds of surveys, show that the majority of voters are ignorant of many key aspects of the U.S. political system. Surveys suggest that people are also largely ignorant of the tax system (Blendon et al (2003) summarize the results).

One specific setting where researchers have investigated understanding of the tax system is taxpayers' perceptions of their average and marginal tax rate. Brown (1968) and Fujii and Hawley (1988) find that individuals' self-reported marginal tax rate often differs from the true rate that can be estimated from their demographic characteristics. Research by de Bartoleme (1995) shows that, in a lab experiment, MBA students often confuse the average tax rate with the marginal tax rate when making investments in a taxable versus non-taxable project. Graham et al (2014) show that this misunderstanding of taxes extends to

corporate managers—only 13 percent of corporate tax executives surveyed responded that their firm used the theoretically correct marginal tax rate (as opposed to some other tax rate) in corporate decision making. Sheffrin (1994) and Liebman and Zeckhauser (2004) also find evidence that taxpayers systematically misunderstand income tax schedules.

Recently, the public economics literature has settled on the term “salience” to capture the extent to which tax aspects of the environment are noticed, and acted upon, by those affected. A key paper in the modern literature is Chetty, Looney, and Kroft (2009). Their findings come from an experiment at a large grocery retailer in California, at which prices inclusive of the 7.375 percent state sales tax were posted alongside the original pre-tax price over a three-week period for three product categories (cosmetics, hair care accessories, and deodorants). They estimate that the “tax treatment” reduced demand by 8 percent; given demand elasticities of 1 to 1.5 for the affected products, they conclude that most consumers do not take into account the sales tax revealed at the cash register. A crucial question for the policy implications of salience, relevant to our analysis, is whether taxpayers pay attention to tax incentives when more utility is at stake (Goldin, 2014; Reck, 2014).

2.2 Macroeconomics and finance

Many macroeconomists have considered how best to account for the inertia in observed economic behavior and to what extent imperfect information can account for it. For example, the models in Sims (2003) are motivated by the idea that information that is freely available to an individual may not be used, because of the individual’s limited information processing capacity. Alternatively, in Reis (2006) consumers rationally choose to only sporadically update their information and re-compute their optimal plans, while in between updating dates they remain inattentive. Both models imply that news disperses slowly throughout the population, so events have a gradual and delayed effect on behavior.

A recent literature in financial economics has taken advantage of newly available data and examined the demand by investors for information. Da, Engelberg, and Gao (2011) propose the Google Search Volume Index (GSVI) as a direct measure of investor attention. Drake, Roulstone and Thornock (2012) use the GSVI for public company ticker symbols to examine the timing and magnitude of Internet search around earnings announcements and the factors that influence Internet search. Among their results of interest is that investors extend more effort when the potential returns to search are higher.

2.3 Implications of theories of information acquisition

Although our analysis focuses on taxpayers' search for information about income taxes, it has more general implications regarding the information economic agents possess when making decisions (in our case, often very economically large decisions). Do individuals possess full information for important economic decisions? If not, is attention to important information rational? If attention is rational, how should we model information acquisition? In general, previous economic studies on attention have focused on the different predictions alternative assumptions make for behavior. However, every theory of behavior contains some assumption, implicit or explicit, about how individuals come to possess information. Our approach is to directly examine these auxiliary predictions of theoretical models, by documenting patterns in the timing of information search. Table 1 summarizes the predictions of the theories our analysis illuminates, as described in this section.

The null hypothesis in our analysis is that taxpayer information search is unresponsive to political or economic events. This hypothesis would be implied by the assumption of full information, wherein an individual always has all data necessary to make an optimal choice. Although frequently criticized and relaxed in the economics literature, this assumption is still common in public finance for modeling the response to tax incentives and the welfare cost of taxation.

Interestingly, the null hypothesis would also obtain in a model of full ignorance: if individuals “don’t know what they don’t know,” they will never search for information.⁴

In rational-expectations models like the one proposed by Muth (1961), individuals face uncertainty about the future, but know everything about the present and costlessly and instantaneously absorb all new information. A rational expectations framework might allow searches for tax information to respond instantaneously to unanticipated policy changes or shocks to the probability of a policy change, as these require that the individual absorb new information about policy. However, under rational expectations, information search related to one’s personal tax liability should not respond to economic or political events that change the importance of understanding tax policy but not tax policy itself. Information search should also not respond to the implementation of policy changes that were not surprises. Individuals should already possess all publicly available information on tax policy.

An attractive alternative framework is rational attention, under which individuals have limited capacity for processing new information and they allocate that resource optimally. Alternative assumptions for how limited attention may be rationally allocated are proposed by Sims (2003) and Reis (2006). In Sims’ (2003) approach, agents constantly update their beliefs based on their attention to different sources of information, while in Reis’ (2006) approach, updating beliefs is costly and, upon updating, the individual simultaneously plans future consumption and decides when to update again. In either model, information search should increase following events that increase the utility gain to understanding the incentives at play in making a decision (such as a market

⁴ From a Bayesian perspective, in order for individuals to have full information or full ignorance, they must have complete confidence in their own beliefs. Full information obtains when those beliefs turn out to be correct, full ignorance obtains when they are wrong. That the demand for information is zero in both cases follows from the nature of the dependence of the demand for information on prior beliefs documented in Keppo, Moscarini, and Smith (2008).

downturn or approaching time notch) and to events that introduce new information about incentives (such as policy changes). Vitrally important for the interpretation of our results is the fact that with a sufficiently relaxed constraint on attention, rational attention becomes behaviorally equivalent to rational expectations. Suppose individuals quickly obtain all the public information necessary to make an optimal choice. Then a rational expectations model with costless information acquisition—or even a full information model if there is little uncertainty—would accurately predict choices, despite a rational attention model performing better in predicting information search and choices jointly.

Attention is also the subject of a large literature in psychology. One useful distinction in this literature is between “exogenous,” or “stimulus-driven,” attention and “endogenous,” or “goal-directed,” attention (Theeuwes, 1994; Connor, Egeth, and Yantis, 2004). Endogenous attention is the same as rational attention: the individual voluntarily directs attention to meet a goal. Exogenous attention is driven, rather, by an external stimulus, such as increased salience. In this case, the individual might search for information regarding anything she sees in the news which she does not understand, regardless of her gain from the search. If attention to economic information is exogenously driven at times, information search should respond to news events that mention policy—because these are usually accompanied by news coverage that makes the policy more salient.

3 Capital Gains Taxation and Behavior

Capital gains generated from the sale of capital assets are subject to tax, but have received preferential treatment relative to labor income since 1921 in the United States. Income from the sale of capital assets is recognized in the year of sale, and the taxable income is equal to the sale price of the asset less its tax basis (the historical price plus any acquisition costs and improvements to the asset, less any accumulated depreciation). Capital gains and losses are divided into two categories, short-term and long-term. Long-term (short-term) capital gains and

losses arise from the sale of a capital asset that has been held for more than (less than) one year. Long-term capital gains receive a favorable tax treatment and currently face a maximum capital gains tax rate of 20 percent, along with a new 3.8 percent Net Investment Income Tax that is applicable to taxpayers in the highest tax bracket. Short-term gains are taxed as ordinary income. Gains and losses from short and long term assets are combined to determine a taxpayer's ultimate tax liability.

A rational investor should maximize after-tax utility, which would entail taking capital gains taxes into account in deciding whether to buy, sell, or hold a capital asset. Consider the purchase decision in isolation from the rest of an investor's portfolio. The taxation of realized capital gains reduces the expected after-tax rate of return to a capital investment, where the reduction depends on the expected appreciation, the expected holding period, and the likelihood that an asset with appreciation can be held until death. Because higher capital gains taxes reduce the attractiveness of assets expected to appreciate in value, one may expect that the level of asset prices would react negatively to unexpected news about tax increases (described as a "capitalization" effect).⁵

The decision of if and when to sell a capital asset should also be affected by the tax system. Certain provisions in the tax code (for example, a lower tax rate for long-term capital gains), may encourage taxpayers to postpone an asset sale in order to obtain the favorable tax treatment. Likewise, the annual nature of tax compliance may also create annual rebalancing of portfolios to achieve a favorable mix of capital losses and gains (such that the losses almost exactly cancel the gains), and may encourage asset sales.

Because capital gains taxes are triggered by asset sales that happen in year t , there is a planning deadline for capital asset sales at the end of the year. Therefore,

⁵ Of course, the effect of acquiring information about capital gains taxes depends on how it changes prior beliefs. One can imagine a potential investor being pleasantly surprised to learn about the preferential lower tax rate and step-up basis at death, or being discouraged upon learning that any tax at all applies upon the sale of appreciated assets.

information gathering in order to achieve the best after-tax outcome with regards to capital asset sales will have to happen by December 31 of year t . When capital gains tax rates increase between year t and year $t+1$, December 31 is also the last date to realize capital gains at the year- t tax rate. As such, information gathering in order to shift capital income through time in response to a rate change will also need to happen by December 31st of the end of the year.

Previous empirical evidence is largely consistent with taxpayers altering both the nature and the timing of transactions in order to achieve the greatest possible after-tax return on their capital investment. For example, as capital losses must generally be offset with capital gains, large downward movements in the stock market often leave investors with large unrealized capital losses that must be carried forward until years when these investors realize capital gains. This policy leads to the well-known strategy for minimizing capital gains tax liability of “loss harvesting,” selling capital assets with built-in losses to offset the gains realized during the year. For example, Ivković, Poterba and Weisbenner (2005) find evidence of tax-loss selling at the end of the year. Poterba and Weisbenner (2001) find that this is especially prevalent in years when changes in tax policy provided additional incentive to harvest losses at year-end. While tax-loss selling occurs in the corporate equity market, there is also evidence of such activity in municipal bond closed-end funds (Starks, Yong, and Zheng, 2006) and the market for long-term government and corporate bonds (Chang and Pinegar, 1986).

4 Data

Tax information to taxpayers is made available by the IRS, by accounting firms and other organizations in the business of providing information generally, such as Wikipedia. The IRS-provided information can be accessed through the Internet, via printed information booklets, and through toll-free numbers. Person-specific information can also be obtained by perusing one’s paycheck, one’s Form W-2, or one’s prior tax returns. Both information from the IRS website and

information from non-government organizations may be found quickly using Google searches. In all cases, the marginal cost of public information is only the time spent acquiring it, provided one has an Internet or phone connection. The cost of acquiring tax information decreased with the advent of the Internet, and search engines in particular. Long-run trends in information search may therefore vary due to both supply and demand for information. By focusing on high-frequency variation in information search, we isolate variation in the demand for information because the supply of information is unlikely to vary from day to day.

We examine measures of information search through many different channels. We first study data on tax-related information search from Google, accessed via Google Trends. We use Google data not only because Google makes its search data publicly available, but also because Google is the most widely used search engine throughout our sample period, capturing 66.9 percent of search volume in 2013 (comScore, 2013).⁶ One can use Google searches to find information from a variety of sources, including Wikipedia and the IRS. Using query data on Google searches, Google Trends provides a measure of the “propensity to search” for a given search query or set of queries. More specifically, an observation in the Google Trends data will be, for a given day and geographical region, the number of Google searches for the specified search terms divided by the total number of Google searches on any topic in the time and place.⁷ For our purposes, the search terms will be a broad set of capital-gains-tax-

⁶ Google’s share of internet searches has grown over the sample period. According to comScore (2005, 2007, 2009, 2013) Google’s share of online searches was 36 percent in 2005, slightly exceeding that of its closest competitor, Yahoo!. Google’s share grew rapidly until 2007, when it reached around 60 percent of the market. It has grown about 1 percent per year since then. Note that the Google’s scaling of the data and our removal of the long-run trend ensure that this does not jeopardize our interpretation of our results.

⁷ For more detail about Google Trends, see our appendix and <http://support.google.com/trends/?hl=en>. Unfortunately, Google does not provide their search data separately as a numerator (Capital Gains tax searches) and a denominator. While the scaling done by Google does help alleviate the concern we are capturing some temporal trend related merely to overall Internet usage, it raises the issue that we are unable to disentangle how much of our effect is coming from the numerator (Capital Gains tax searching) or the denominator (total Google searches). In the absence of data on total Google searches by day, which we are unable to attain, we are unable to investigate the separate numerator and denominator effect. However, the fact that our Wikipedia data are unscaled and that we obtain results that are materially the same, scaled or unscaled, suggests that the denominator effect in the Google data is not driving our results.

related search terms (listed in the Appendix), and the geographic region is the United States. After calculating the propensity to search, Google scales it from 0 to 100, where the number 100 represents the day with the highest search volume for this set of search terms in the entire sample period. Our sample period consists of January 1, 2004 (when data are available) through March 30, 2013.

Entering a given term into Google will frequently yield a Wikipedia page related to that term as one of the first suggested sites to visit. Our second source of information search data is page view data from Wikipedia, a free online encyclopedia edited by Internet users. We obtain the number of people who landed on the English language Wikipedia site, “Capital Gains Taxation in the United States” every hour. We use the summation of this hourly data over the 24 hours in a day to obtain daily data for January 1, 2008 through March 30, 2013.⁸ One disadvantage of the Wikipedia data is that it is available for a substantially shorter time period than is the Google data. On the upside, we are able to obtain the raw number of views of the webpage, which simplifies the interpretation of the size of changes in information search.⁹

Our third source of data regards aggregate calls made to the IRS’s toll free phone number, where taxpayers can call and speak to a representative from the IRS or listen to automated messages. In the course of a phone call with the IRS, some taxpayers will listen to prerecorded messages about various “tax topics.” We analyze daily phone calls that access IRS tax topic 409, “Capital Gains and Losses.”¹⁰ Taxpayers may reach the tax-topics function in a number of ways, including calling the tax-topics-specific phone number, or being referred to the

⁸ For a limited number of hours, Wikipedia data are unavailable. We scale the observation up by 24 divided by the number of hours present in a day when only some hours are missing (15 days), and we use the previous day’s value when the entire day is missing (128 days).

⁹ We have replicated the results in Section 6 and Section 7 using a measure of Wikipedia usage scaled in the same way as the Google data, using page views of the Wikipedia home page as a metric of overall Wikipedia usage. This replication suggests that scaling is not confounding our results.

¹⁰ Other example of tax topics are tax topic 304, “Extensions of Time to File Your Tax Return” or tax topic 151, “Your Appeal Rights.”

topic after a conversation with a representative or an interaction with the automated system. We use data from February 1, 2002 to March 31, 2012. We also obtain the total number of calls to the call line for individuals, which is not specific to capital gains.

The last data set we analyze consists of the aggregate daily number of visits¹¹ to the IRS’s webpage. The IRS maintains a website, IRS.gov, which hosts a vast amount of tax information. We obtain, directly from the IRS, the total number of page views and visits, by day, to any site hosted by IRS.gov. We are able to obtain this data for the entire time series for which the IRS has maintained the data, February 1, 2002 to March 31, 2012. Analysis of IRS.gov traffic carries the benefit of analyzing taxpayer information search about a very broad set of tax information.¹² This measure of information search is not limited to information search about capital gains taxes specifically.

5 Estimation Procedure

Before discussing data patterns, in this section we lay out the research design for our parametric analysis. Let I_{it} denote information search on date t from a source of information i , such as Google, Wikipedia, or the IRS call line. We wish to estimate the effect of several different events on information search through source i . The events we study may be either non-recurring, as in the case of a Presidential debate that mentions capital gains taxes, or may occur annually, as in the case of tax filing season. Non-recurring events are indexed by the subscript k , and annual events are denoted according to the day of the year on which they occur, denoted DoY . As we discuss in more detail later, the data on information search display marked seasonality at the yearly and weekly levels. In order to

¹¹ In contrast to page views, visits are a web analytic that count the number of page views from unique Internet Protocol (IP) addresses in a given time frame. As a result, a single viewer from the same computer may refresh the webpage multiple times, and each refresh will count as a new page view, but not a new visit. Wikipedia provides data only on page views.

¹² The phone call and website visit data was made available to us by the Research, Analysis and Statistics division of the IRS.

evaluate whether a particular (non-seasonal) event increased information search on date t , we must also properly specify the counterfactual level of search, conditional on the day of the week and day of the year of date t .¹³

Unlike for much analysis of time-series data, yearly seasonality is of intrinsic interest for our research questions. In particular, increases in information-seeking during tax season will affect the form of the yearly seasonality. Information searches occurring at the end of the year can be thought of as responses to a time notch. To address this concern, we employ a method that both estimates the seasonal patterns and allows us to perform classical statistical tests for whether information search is significantly higher on a particular day of the year than the average level of information search. Specifically, for each information search series i we estimate the function:

$$I_{it} = \sum_k [\beta_{k0}F_{kt} + \beta_{k1}F_{k(t-1)} + \dots + \beta_{k4}F_{k(t-4)}] + x_t' \gamma + f_i(DoY_t) + \delta_{DoW,i} + u_{it} \quad (1)$$

The term inside square brackets is a set of dummy variables: $F_{k(t-j)}$ equals 1 if event k occurred on date $(t-j)$ and zero otherwise. The four-day event window was selected because when large spikes in information occur in the data, search levels return to baseline within four days.¹⁴ We also include a vector of continuous, time-varying linear covariates x_t' —such as trading volume on the stock market—along with a non-parametric function in day-of-the-year $f_i(DoY_t)$, and a day-of-the-week fixed effect, $\delta_{DoW,i}$.

We estimate $f_i(DoY_t)$ using kernel-weighted local linear regression. This procedure estimates a weighted linear regression at each value of DoY , using data

¹³ While we explicitly control for these temporal trends in searching in our regression framework, as a result of how Google provides their Search Volume Index (SVI), Google has also helped alleviate some of these temporal issues. For example, people likely search in general more during Christmas break, so a concern may be that we are picking up Christmas break searching as part of our year-end search activity. However, as explained earlier, the SVI divides standardized total searches for the topic per day by total search volume on Google. As a result, factors such as mere increases in Internet usage likely do not explain the patterns in search behavior we see.

¹⁴ Figure A.4 in the Online Appendix contains a plot of the seasonally adjusted and detrended data around event dates, verifying that the impact of events tends to fade in four days.

from adjacent days (e.g. the j^{th} day before and the j^{th} day after DoY for some j).¹⁵ Dates further away from some DoY receive less weight in the estimation of the linear regression at DoY ; how the weights decay is governed by our choice of kernel function and bandwidth, discussed below. Our estimate will be consistent under the assumption that the value of $f_i(DoY_t)$ does not change *too* sharply from day to day.¹⁶ The strength of this estimator is that 1) it estimates a smooth trend in day-of-the-year, which allows us to control for day-of-the-year when estimating the effect of events, and 2) it increases the precision of the estimates of $f_i(DoY_t)$ relative to fixed effects. A potential downside to the estimator is that the assumption that $f_i(DoY_t)$ does not on average change too sharply from day to day may be wrong (especially around April 15th). The validity of this assumption is determined by the bandwidth of the estimator, and also by our choice of kernel density function. We discuss these choices in more detail in the online. Our preferred choices for kernel density function (Gaussian) and bandwidth (four days) produce estimates of seasonal trends that match the pattern implied by fixed-effects estimates. To avoid conflagration of seasonal effects with large outliers due to the events we document, we estimate all the components of equation (1) in a single regression. The appendix to this paper outlines the estimation procedure in more detail.

For a recurring annual event such as a filing deadline, we can examine the function $f_i(DoY_t)$ on days-of-year corresponding to the annual event to understand the effect of the event on information search. The comparison that seems most natural is to test the hypothesis that, on a particular day of the year, DoY' , inquiries are higher than the average over all days of the year, i.e. that $f_i(DoY') > E[f_i(DoY)]$.¹⁷

¹⁵ Leap days are dropped from the sample for simplicity, so that every year is 365 days long.

¹⁶ Alternatively, we could estimate $f_i(DoY_t)$ using day-of-the-year fixed effects. However, with only 5 to 9 years of data for each type of inquiry, these fixed effects would be imprecisely estimated.

¹⁷ This is analogous to a more conventional hypothesis test with fixed effects, such as the one we use to test for significant

Our estimate of the function $f_i(DoY)$ will inform us about the importance of tax season for information search, but it may not capture all variation due to compliance deadlines.¹⁸ For example, the mid-April filing deadline does not occur on April 15 of a given year if that date falls on a weekend or Emancipation Day. Instead, it can occur as late as April 18 in some years. We account for variation in compliance dates by year by including dummy variables for compliance events in addition to the locally linear function. Our estimate of the function $f_i(DoY)$ will estimate the average effect on April 15, but we add an additional dummy variable for the precise date of the filing deadline to capture variation that occurs specifically on the deadline each year. Doing so does not significantly change our picture of the importance of tax season for inquiries, but it does highlight the sharp spike that occurs exactly on the mid-April deadline each year. The inclusion of this dummy variable also helps mitigate the strength of the smoothness assumption in estimating $f_i(DoY)$, effectively imposing a bandwidth of zero on the date of the filing deadline. To further reflect that the smoothness assumptions should not apply when there are sharp changes on particular days of the year, we also include dummy variables for the following holidays: New Year’s Day, President’s Day, Martin Luther King, Jr. Day, Memorial Day, Independence Day, Labor Day, Thanksgiving, and Christmas.

6 Exploring the Raw Data

Figure 1 plots the evolution of three measures of information search on capital gains taxes over time. Calls to the IRS hotline inquiring about capital gains taxes occur almost exclusively during tax season, from mid-January to mid-April. The absolute volume of calls occurring during a tax season diminishes considerably

day-of-the-week effects. Rather than testing for significance relative to a “left-out” day of the year, however, we test for significance relative to the average over all days of the year.

¹⁸ For Google data, the interpretation of the estimates is somewhat complicated by the fact that the dependent variable is a measure of propensities to search. If there were yearly seasonality in overall Google searches, this would affect our estimate of $f_i(DoY)$ by decreasing it on days of the year when Google usage was highest. The fact that the pattern estimated by Google data and other measures of information search are similar largely alleviates this concern. The inclusion of day-of-the-week fixed effects also makes our results robust to weekly seasonality in Google searches.

throughout the time period, from 2002 to 2012. Wikipedia page views, in contrast, increase over time. These patterns are consistent with online information largely having supplanted telephone information as Internet access increased markedly over the sample period, as well as the fact that the increasing popularity of tax preparation software and online programs likely reduces the need for taxpayers to call the IRS regarding simple questions. The Google measure displays a slight downward trend.¹⁹

Visually apparent in each time series is a strong pattern of yearly seasonality. This pattern is most pronounced in the IRS call log series, but it is also present in Google searches and Wikipedia page views. Focusing on narrower time frames also reveals strong weekly seasonality in each time series.

To focus on high-frequency variation, we detrend the data using a Hodrick-Prescott filter.²⁰ To allow for visual comparison of the three measures of information search, we normalize each variable by dividing by its standard deviation.²¹ Figure 2 plots the estimated yearly seasonality in the standardized data.²² For each measure of information search, we observe a sizable and significant increase in search behavior during tax filing season, starting first in early January, as some taxpayers are likely filing early to get refunds as soon as possible, and others are receiving information returns (e.g., 1099-DIV, 1099-B) and pondering how they affect their capital gains tax liability. Then, we observe an even stronger spike in the immediate run-up to the April 15 filing deadline. Right after April 15, information search drops off sharply. The estimated seasonality in Figure 2 matches patterns in the timing of filing income tax returns

¹⁹ Recall that this does not mean that Google searches of capital gains taxes decreased over time, only that the share of Google searches that concerned capital gains taxes decreased over time.

²⁰ Whenever we use a Hodrick-Prescott filter, we use a smoothing parameter of 10^7 . This value was selected by trial and error, with the goal that the long-run trend captures long-run movements in the series but not variation due to yearly seasonality.

²¹ Because the volatility of the IRS call log series varies significantly over time, we standardize it by dividing by its standard deviation *by year*.

²² This estimate of seasonal effects comes from the regression in Section 7, and controls for event-driven outliers.

documented by Slemrod et al (1997). Clearly, the desire to comply with the tax law and take advantage of applicable tax credits and deductions as one fills out a tax return leads taxpayers to search for information.

In the next set of figures (which we analyze formally later in the paper), we focus more narrowly on variation in the standardized, detrended, and seasonally-adjusted series around three dates where we observe the largest spikes in information search.²³ Figure 3 focuses on the spike in October and early November 2008.²⁴ There are two obvious candidates for what triggered elevated information search in this time period: the presidential election of 2008, and the stock market crash. Because capital gains taxes were an issue on which candidates John McCain and Barack Obama differed substantially, the election could influence search behavior due to either the desire to make an informed choice about asset purchases or sales based on expected future tax policy, or due to the desire to understand the consequences for future tax policy of a President McCain or Obama. The market crash of 2008 resulted in large capital losses for many investors and, due to the extreme volatility of the market, potentially also large short-run gains. Our belief is that both the election and the crash played some role, although it is difficult to disentangle the effect of the two events. The largest swings in the stock market in this period (marked by vertical lines in Figure 3) are associated with modest elevation in search behavior, and information search surged strongly around the date of the presidential election. Interestingly, in this figure, we note no surge in search activity related to the IRS call line, suggesting that different information sources serve different purposes. Consistent with other evidence presented here, it appears that taxpayers mostly use the IRS call line for

²³ Seasonal adjustment consists of subtracting out estimated day-of-the-week effects, holiday effects, and the day-of-the-year effect. The full detrended and seasonally adjusted time series are provided in Online Appendix Figure A.1.

²⁴ If we estimate a simple regression of (1) including an indicator variable for whether the date was in September, October, or November of 2008, we find that Google searches were 1.06 standard deviations higher than the seasonal trend on average over the period and Wikipedia searches were elevated above the seasonal trend by an average of 0.501 standard deviations.

tax compliance purposes, as we do note a strong seasonality associated with tax season in the IRS call data. We present more evidence on how the crash might have affected information search by focusing more narrowly on capital losses in section 8, and Figure 6 provides further evidence on presidential elections.

Panel A of Figure 4 plots the same series in January and February of 2012. We attribute the observed surge in information search to the release of presidential candidate Mitt Romney's 2011 tax return on January 24. The release generated substantial news coverage, in part due to his low effective tax rate. As most news articles on the subject noted, much of Romney's income came from the realization of long-run capital gains taxes, taxed at 15 percent at that time. Another possible explanation for the spike in search behavior on this date is the State of the Union Address, which also occurred on January 24, 2012. In his speech, President Obama advocated taxing the wealthy at higher tax rates (supporting the "Buffett Rule"), but he mentioned neither capital gains taxes specifically nor capital income taxes more generally.²⁵ It is, however, possible that Obama's rhetoric on taxing the rich led the public to pay greater attention to the news about Mitt Romney's tax return, and thereby amplified its effect on capital-gains-tax-related information search.²⁶

Panel B of Figure 4 plots the data from detrended and seasonally adjusted data at the end of 2012 and the beginning of 2013. A presidential debate between Barack Obama and Mitt Romney appears to have sparked considerable information search. As in 2008, we also see elevated search immediately following the presidential election, perhaps as voters and investors researched what might happen to capital gains taxes in the aftermath of President Obama's

²⁵ The "Buffett Rule" is a tax plan proposed by President Obama, wherein individuals making over \$1.0 million in taxable income would be subject to a minimum average tax rate of 30 percent regardless of whether the income is composed of capital gains or ordinary income.

²⁶ Further insight into this issue can be gained by examining Wikipedia intra-day search activity. Search activity started rising dramatically mid-day on January 24th, 2012, consistent with Romney's tax return being at least partially responsible for this increase (as Obama's speech was not delivered until that night). This differs from the typical intra-day pattern of search activity, which generally shows more searching done after working hours.

re-election. Finally, the largest spike in information search in the 10-year period covered by our data—an increase of just over 4 standard deviations in Wikipedia page views and over 6 standard deviations of Google searches—occurred on January 2, 2013. We attribute this to the passage of the American Taxpayer Relief Act (ATRA) on that date. Again, we note the lack of any response in the IRS call data, consistent with taxpayers mostly using the IRS call line for tax compliance purposes. This bill temporarily resolved what was commonly called the “fiscal cliff” debate, and increased the top marginal tax rate on long-term capital gains from 15 percent to 20 percent.²⁷ Our evidence strongly suggests, therefore, that individuals search for information both in response to policy changes and in response to potential policy changes signaled by political events.

7 Regression Analysis of the Impact of Events

We start with the detrended and seasonally *unadjusted* time series. The regression procedure we use, outlined in Section 5 and in the appendix, explicitly controls for variation due to weekly and yearly seasonality. The notes to Table 2 describe the events we study, which are also discussed in the previous section.²⁸

Table 2 describes the estimated impact of events on information search through Google, Wikipedia, and the IRS web page. To examine statistical significance, we use a non-parametric permutations test based on the test proposed by Gelbach et al (2013) for single-firm event studies. This test is based on the comparison of the effects of the events we study with the estimated total effect of “placebo” events occurring on arbitrary days. The appendix describes this procedure in more detail. The resulting *p*-value corresponds to the probability

²⁷ As with many of our events, the focus of the ATRA was not the capital gains tax. Just as the presidential debates covered a broad spectrum of policy issues, the ATRA was an omnibus bill that combined many different tax and non-tax initiatives into a single piece of legislation. For a discussion of the effect of omnibus tax legislation, see Hoopes (2014).

²⁸ There is modestly elevated information search via Google, Wikipedia or both, on a few dates that we do not include in the analysis. These dates include January 18, 2004 (a State of the Union address delivered by George W. Bush advocating the extension of capital gains tax cuts), November 2, 2004 (re-election of George W. Bush), May 16, 2006 (the extension of the 2003 capital gains tax cuts), January 27, 2010 (a State of the Union address by Barack Obama advocating a cut in capital gains taxes for some taxpayers), and March 23, 2010 (passage of the Affordable Care Act, which included a “net investment income tax”). In each case, the response of searches is qualitatively similar—a spike in searches that fades in three to four days—but notably smaller than the events we do include in the formal analysis.

that increases in information search of the magnitude we observe would have occurred at random during the event window we specify. The results suggest that the events we study each have a large and significant impact on information search through Google and Wikipedia, but not through the IRS call line. When we estimate the overall impact of the event rather than examining a single-day impact, the release of Mitt Romney's tax return surpasses the passage of ATRA as the one-time event that generated the most taxpayer information search, through both Google and Wikipedia. This occurs because the effect of ATRA passage declined more rapidly, fading in two days instead of four. Notably, ATRA passage was anticipated in the days leading up to January 2. The bill passed Congress on January 1st, and the anticipated increase in the top capital gains rate, likely generated income shifting from 2013 to 2012, which may have also caused information search. From Figure 4, panel B, we can see that elevated information search occurred on the two days prior to ATRA passage, December 31 and January 1st. Adding these dates to the event window increases the estimated impact of ATRA by approximately 5 standard deviations for Google and Wikipedia, but changes little else.

8 Relating Tax Information Search to Stock Market Activity

An advantage of focusing on capital gains taxation is the availability of high-frequency data on sales of stock. Relating these data to data on information searches holds the promise of better understanding the causal connections between information search and capital-asset-related behavior. After all, we are interested not only in what causes people to search for information, but also in to what extent the acquisition of information affects behavior, in this case behavior related to the sale, purchase, and holding of capital assets.

8.1 The Lead-Lag Relationship to Volume, Volatility, and Market Return

Our first strategy is to examine the lead-lag relationship between measures of behavior and information searches. If searches lead behavior, then we have reason

to pursue the idea that the information obtained affected subsequent decisions. We investigate two data series, both of which represent general stock market activity.

The market-related measure we use is trading volume. We obtain the dollar value of shares traded from all publicly listed firms from the Center for Research in Security Prices (CRSP variable VOL), which we use as a measure of broad market activity.²⁹ We then detrend the measure using a Hodrick-Prescott filter, and include the log of daily trading volume as an independent variable in the regression described in equation (1). For our baseline regression, we include five days of leads and lags of log trading volume. These regressions control for the events we study and for weekly and yearly seasonality. Columns (1) through (3) of Table 3 report the coefficients on standardized data. See Online Appendix Figure A.2 for a graphical depiction of the lead-lag estimates.

When trading volume is high, individuals seek more information about capital gains taxes. The standard deviation of log trading volume is 0.2233, so the estimates in table 3 suggest that a one-standard-deviation increase in trading volume predicts a 0.064-standard-deviation increase in Google searches, and a 0.068 standard deviation increase in Wikipedia Page views. These effects are significantly different from zero at conventional significance thresholds ($p = 0.068$ or 0.038 for Google, $p = 0.055$ or 0.027 for Wikipedia, depending on whether one employs classical or Newey-West standard errors, respectively).³⁰ We find that searches on date t are not significantly related to trading volume on date $t+1$ or beyond. All of the association of market movement with information search occurs on the same date as the market movement, or one day before.³¹ As a

²⁹ Note that these transactions include many where the buyer and/or seller is not subject to capital gains taxes, may not be triggered by a human trader capable of information search or where the asset does not have an accrued gain that will be subject to taxation. These possibilities do not threaten our identification so long as the percentage of transactions that is not subject to capital gains taxes does not change substantially and systematically from day to day.

³⁰ This result suggests that a non-constant fraction of people trading consider the tax implications of the trade, as a constant fraction of traders would imply a coefficient of 1. It seems likely that on high-volume days, a smaller percentage of traders may be carefully researching the tax consequences of the trade, resulting in the fraction of tax-informed traders being non-constant across time.

³¹ If we include only one lead and lag of log trading volume we estimate that the coefficient on the first lead and lag are

result, we are not able to determine from these data whether individuals seek information on capital gains taxes primarily before or after they make a decision regarding the sale of a capital asset. If they seek information before making a decision, they do so less than a day in advance, as far as these data are able to tell us.

We are sensitive to the possibility that some of the variation in these results may be driven by behavior during the 2008 stock market crash, a period of extremely high stock price volatility and trading volume. Columns (4) and (5) of Table 3 provide the estimates of the same regression specification, but limiting the inclusion of market variables to the period from September 2008 to February 2009. Columns (6) and (7) provide the estimates of the same regression specification, instead excluding this extraordinary period. For both Google and Wikipedia data, the estimated relationship between search and market activity during the extraordinary period is much larger, although imprecisely estimated. However, the estimated effects we have described survive when the extraordinary period is excluded from the sample.

8.2 Do taxpayers search for capital gains tax information before they trade?

In this section we pursue an alternative indicator of the taxpayer demand for capital gains tax information that is itself based on observed search volume. In particular, we use Google Trends to obtain a measure of the volume of searches for phrases related to personal investment advice such as “stock advice”, “should buy stock,” “should sell stock,” and “investment advice.”³² In so doing, we hope to learn more about the timing of the relationship between search for tax-related

slightly larger than in our preferred specification. Including up to 14 leads and lags does not change the estimates from our preferred specification. If we include one lag of log trading volume *without* contemporaneous log trading volume, the estimated coefficients are nearly identical to the ones in Table 3. If we include one lead of log trading volume *without* contemporaneous log trading volume, we do estimate a statistically significant relationship similar in magnitude to the contemporaneous effect. This last result is almost surely driven by autocorrelation in trading volume.

³² Unfortunately, Google Trends produces data only when there is sufficient search volume for a given search term. So, while more specific search terms would be useful (i.e., “how to short Facebook”, or, “should I sell my capital losses”), there is no available data on these searches.

information and decisions about whether to buy, sell, or hold assets. Note that, although Google searches for investment advice can lead one to websites that purport to provide such advice, searches for “capital gains tax” generally do not.

Table 4 shows the results of including the investment advice measure as an explanatory variable in a regression on searches for capital gains tax information, both in addition to stock trading volume and as a replacement for it.³³ Several results of interest emerge.

First, consider the results for Google searches. When investment advice search volume is included as an alternative to stock trading volume, the same positive contemporaneous association appears, and the effect of a one standard deviation change on capital gains tax related searches is 40 percent lower. In a “horse race” when both stock trading volume and investment-advice search volume are included as explanatory variables, unsurprisingly the statistical significance of the former declines because of the high correlation between the two variables. In the horse race, investment advice search volume wins, retaining its significance while stock trading volume loses its own significance. Strikingly, the one-day lag *and* *lead* values of stock investment advice search volume are both significantly positively associated with capital gains tax search volume. This estimated leading association is consistent with a story that taxpayers first recognize the need for information regarding buying or selling stock and, in the process, learn that relevant to this decision are the tax consequences.³⁴

8.3 Google Correlate

While the above mentioned stock-related terms that we selected (e.g., “stock advice”) were correlated with searches for capital gains taxes in

³³ Online Appendix Figure A.3 graphically depicts the lead-lag structure estimated when stock advice search volume is included in the absence of stock market trading volume.

³⁴ The key advantage of the stock investment advice search volume variable is that it likely captures the extent of taxpayer demand for information for which capital gains tax knowledge is crucial. One potential disadvantage is that in regression analysis one may pick up any shocks that affect all Google Searches. To that point note that in Table 5 the same pattern of results also applies when the volume of Wikipedia searches is the dependent variable, with one exception: the lead relationship, although positive, does not reach statistical significance.

predictable ways, Google, through its application Google Correlate, also has the ability to provide a list of search terms most highly correlated over time with any given user provided search term. Consistent with our other evidence that taxpayers simultaneously search for information about capital gains taxes and investing, among the top 20 search terms related to “capital gains tax” are “stock purchase,” “investing,” “td waterhouse,” “fidelity mutual funds,” and “mutual funds.”³⁵

8.4 Capital Losses

Because the capital gains tax rules related to the sales of assets with capital losses are especially important for many tax minimization strategies, such as loss harvesting, we also construct two measures of information search related specifically to capital losses. The first is weekly Google searches for the phrase “capital loss,” and the second is monthly searches related to capital losses using the search functionality on the IRS website. We obtain the latter measure directly from the IRS.³⁶ These time series are plotted in Figure 5.

As with searches for capital gains, much of the variation is seasonal: people tend to search for information on capital losses during tax season. There is also typically an increase in searches in December of each year, when some taxpayers “harvest” capital losses to reduce their tax liability. Of particular interest is the fact that searches for information on capital losses increased dramatically during October 2008, and surged even further in December of 2008. When the crash began in October, investors began to research the tax implications of the unrealized or realized losses they had sustained, perhaps evaluating the merits of pulling their wealth out of (or investing in) the declining stock market. Furthermore, for an investor who had lost money in the crash, harvesting capital

³⁵ <http://www.google.com/trends/correlate/search?e=capital+gains+tax&t=weekly&filter=capital+gains+tax&p=us#default>, 30 displays the full results. We thank Hal Varian for suggesting the use of Google Correlate.

³⁶ The IRS provided to us, at a monthly level, the most frequent search terms, and the number of searches for all terms with the word “capital”, and “loss”, in the search.

losses before the end of 2008 could reduce the taxpayer's tax liability substantially.³⁷ The heightened importance of loss harvesting from the stock market crash apparently caused even more information search at year-end.

For both Google and IRS searches, we can also see that searches for information on capital losses *during tax season* are higher after the 2008 crash than before or for three years after the crash. For several years after the crash investors realized capital losses with greater frequency than before the crash. At the same time, information search for tax compliance purposes also increased after 2008. This constitutes suggestive evidence of spillovers between macroeconomic changes and information search for the purpose of tax compliance.

9 Lessons from Aggregate Data on IRS Website and Call Line Usage

Figure 6 plots time series data on the use of the IRS website and call-line usage that are not specific to capital gains taxes. In order for their effects to be visible in broader measures of information search, events must significantly change the importance of understanding taxes at a given time, for a large number of taxpayers. Unsurprisingly, the IRS.gov domain and call line experience elevated traffic during tax season. Note that the amount of traffic experienced during tax season increases over time for the web page, probably due to increased use of the Internet by taxpayers. The call line, in contrast, decreased in usage from 1999 to 2008, and then experienced a resurgence. Together these results suggest that the Internet has not completely crowded out the use of the call line.

Most interestingly, we observe an abnormal surge in visits to IRS web pages and calls line traffic in May of 2008. Search volume usually drops sharply after the mid-April tax deadline, but in 2008 it remained high throughout the month of May. Among the top five searches on the IRS web page during this month were “stimulus,” “rebate,” and “stimulus check,” and many of the top

³⁷ This would naturally require that the investor had capital gains to offset with these losses.

pages viewed also dealt with the stimulus rebates. The tax rebates enacted by the Economic Stimulus Act of 2008 led millions of taxpayers to visit the IRS web page to investigate how the federal stimulus program affected them. From more detailed data on the use of the call line, we know that the IRS received over 50 million phone calls on stimulus checks, corresponding to at least 11 million unique taxpayers. This corroborates the evidence provided by Sahm, Shapiro, and Slemrod (2012) that these rebate checks were a salient form of economic stimulus with a relatively high (compared to a reduction in employer withholding rates) marginal propensity to consume. Using the intuition of rational attention, it makes more sense for individuals to seek information in May about their stimulus check if they intend to spend it compared to if they intend to save it.

10 Conclusions

It is well-established that in general taxpayers know little about the US income tax, and have systematic misperceptions. Given that acquiring information is costly, it may be optimal for individuals to learn only *if* the expected return is high enough and only *when* the information is most useful, known as rational attention. Because people are learning—and forgetting—things all the time, the process of net information acquisition is critical to a dynamic understanding of tax salience. Using newly available IRS administrative data and publicly available information on Google and Wikipedia searches, this paper establishes that people seek information about the US income tax in systematic ways that are consistent with the idea of rational attention. When policies change or seem likely to change, when filing deadlines or time notches loom, people turn to online resources like Google, Wikipedia and the IRS website, as well as traditional information resources like the IRS telephone hotline, to learn how the tax code affects them. In addition, people search for tax-related information when newsworthy events make taxes more salient, and in so doing they may incidentally obtain information relevant to their own decisions.

When policy or news events generate exogenous shocks to the demand for information, the responsiveness of information search to the event occurs remarkably quickly: search behavior usually spikes on the same day as the event, and falls back to baseline within three or four days. This timing pattern suggests that either 1) taxpayers acquire information about capital gains taxes much more quickly than envisioned in macroeconomic models of rational attention, or 2) the public's attention span is very short. In relation to the first point, we note that in the US, Google searches for “inflation”—the topic that most macroeconomists have in mind when thinking of rational attention—exhibit similar 3-4 day spikes following the release of inflation reports from the Bureau of Labor Statistics (BLS). These spikes are especially pronounced when BLS reports document large increases in the price level. A short public attention span could be caused by convex costs of attention to a particular topic, or it could mean that the discrete spikes in attention we observe are due to salience rather than rational attention.

We also present somewhat weaker evidence regarding to what extent acquiring information about taxes leads individuals to alter behavior. We show that individuals sought information about capital losses at year ends, especially after the stock market crash of 2008, since harvesting losses provides an opportunity to reduce one's tax liability. This finding is consistent with many investors not knowing the rules governing capital losses well enough to confidently apply them to make their behavior tax efficient in the wake of the crash. Information acquisition is a necessary component of the response to tax incentives, especially where more obscure details of the tax code are concerned.

Our attempt to learn about taxpayer information search is limited because we cannot account for all sources of information. For example, we cannot observe taxpayers obtaining information from paid professionals. About 60 percent of taxpayers do use a paid preparer when they file a tax return. Some of our measures may, however, capture financial advisors' information search. Further,

we cannot rule out “learning by doing,” whereby individuals learn about tax incentives while filling out a return and then improve their future decisions. Nor can we observe learning through social networks (as documented in Alstadsæter, Kopczuk, and Telle, 2010). As in all similar studies, we also cannot quantify the *amount* of learning that occurs via acquiring any one piece of information, especially because we do not observe which specific taxpayers are searching for information at a given time. Finally, taxpayers for whom capital gains taxes matter are generally wealthier than the typical taxpayer. These individuals may be more sophisticated in their use of information than the typical taxpayer, and many can afford to pay for high-quality information from financial advisors. For these reasons, future work should examine patterns in information search about other policies and, if possible, with individual-level micro data.

Appendix

Obtaining and Analyzing Google Trends Data

When querying Google Trends, the user provides 1) search terms, 2) a geographical window, and 3) a time range. As we are studying a tax issue within the United States, for all data used in this paper, the geographical range is specified to be the United States.

When the number of overall searches for a given term is too low, Google Trends will report an SVI of zero, or will report daily at a weekly or monthly level (as opposed to daily SVI data). We encounter this issue if we query searches for “capital gains tax” alone. As such, we use a set of search terms to maximize our sample period for which we are able to obtain daily data. The search terms we include are, according to Google Trends itself, highly correlated with searches for “capital gains tax.” The set of search terms is the following:

- Capital gains tax
- Capital gains tax rate
- Capital gains taxes
- Capital gains calculator
- Capital gains
- Capital gains rate

We have verified that 1) the daily time series of SVI for simply “capital gains tax” is virtually identical to the one from the broader set of search terms, but with fewer missing data, and 2) we obtain nearly identical results for event studies and market movement effects if we use simply searches for “capital gains tax” as our left-hand-side variable instead of the broader set of search terms, but with slightly larger standard errors (reflecting the decreased number of observations).

For the same reasons as above, we use multiple search terms related to stock advice. These are:

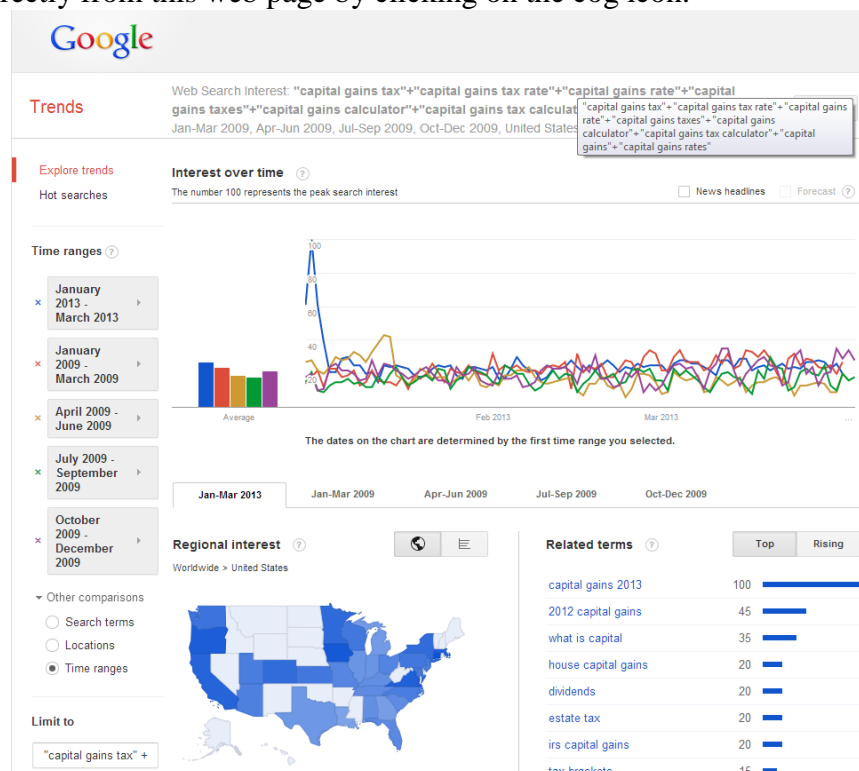
- Stock advice
- Stock market advice
- Stock tips
- Stocks to buy
- Stock to buy
- Stocks to sell
- Should buy stock
- Should sell stock
- Investment advice
- Investment tips

By default Google Trends provides weekly data when a user downloads a

time series longer than three months. In order to access daily search volume data, one must query Google Trends in three month intervals. Fortunately, one can query several (up to five) three-month periods at once.

In order to obtain daily data while maintaining proper scaling of the variable across the entire time series, however, one must query Google Trends very carefully. The data are scaled so that SVI takes the value of 100 on the day with the highest propensity to search out of any date range and search terms. For example, if January 2, 2013 were the day with the highest propensity to search for capital gains tax terms (as it is), SVI would equal 100 on that day if January 2, 2013 were in the period provided by the user. In order to obtain daily data that is properly scaled for the full sample period, one must first find the single day with the highest search volume, and then include a time period containing that day along with sets of other three-month periods.

To get a properly scaled daily time series, we therefore include the time period January 1, 2013-March 31, 2013 in every single one of our queries for capital gains tax SVI, along with other three-month periods, until we obtain data for our entire sample period. The figure below shows an example of what such a query would look like to pull daily data for the year 2009. One can then download the data directly from this web page by clicking on the cog icon.



Note: Google and the Google logo are registered trademarks of Google Inc., used with permission.

Estimation Procedure

Recall that the equation to be estimated is

$$I_{it} = \sum_k [\beta_{k0}F_{kt} + \beta_{k1}F_{k(t-1)} + \dots + \beta_{k4}F_{k(t-4)}] + x_t'\gamma + f_i(DoY_t) + \delta_{DoW,i} + u_{it} \quad (1)$$

In order to consistently estimate equation (1), which contains a non-linear function and a set of linear covariates, we use the double residual regression method suggested by Robinson (1988), and discussed in Hardle and Linton (1994). This first step of this three-step estimator consists of several non-parametric regressions of the following form:

$$E[X_t] = f_i(DoY_t),$$

where X_t is the dependent variable or one of the linear covariates, i.e. I_{it} , F_{kt} , x_t or $\delta_{DoW,i}$. We then obtain the residuals from this regression, which we denote with star superscripts. The residuals represent the component of I_{it} , F_{kt} , x_t or $\delta_{DoW,i}$ not correlated with the general within-year pattern represented by $f_i(DoY_t)$. The second step estimates the linear components of the model consistently using ordinary-least-squares regression on these residuals:

$$E[I_{it}^*] = \sum_k [\beta_{k0}F_{kt}^* + \beta_{k1}F_{k(t-1)}^* + \dots + \beta_{k4}F_{k(t-4)}^*] + x_t'^*\gamma + \delta_{DoW,i}^* \quad (2)$$

The third step regresses the residuals from the estimation of equation 2—denoted by I_{it}^{**} —non-parametrically on DoY_t to obtain a consistent estimate of the function $f_i(DoY_t)$.

$$E[I_{it}^{**}] = f_i(DoY_t) \quad (3)$$

The first and third steps involve non-parametric regression, which Robinson (1988) suggests implementing via kernel-weighted local polynomial regression. In this paper, we use local linear regressions with a Gaussian kernel density function and a bandwidth of four days.³⁸ This procedure estimates a weighted

³⁸ This bandwidth and kernel density function applies to non-parametric estimations from the third stage of the procedure, and nonparametric estimations involving continuous variables in the first stage. For discrete variables in the first stage, such as event dummy variables, we use a bandwidth of zero, to reflect that there should be no smoothing at this stage. Another complication is that traditional statistical packages and programs will implement the smoothing on a linear variable, while day-of-the-year is a cyclical variable. That is, January 1 and December 31 should be adjacent for smoothing purposes. Ignoring this problem results in a discontinuity in the seasonal pattern between these two days. We eliminate the discontinuity by estimating the seasonal pattern twice: once where the discontinuity is imposed at January 1, and a second time where the discontinuity is imposed at the 200th day of the year (July 19th). Then, we replace the 10 days around January 1 from the first estimation with these days from the second estimation.

OLS regression of the dependent variable on DoY at each value of DoY in the data, where the weights are determined by the kernel density function and the bandwidth. The bandwidth was selected to visually match the fixed-effects estimator of the function $f_i(DoY_t)$, but a data-driven choice of bandwidth—specifically selecting the bandwidth that minimizes the conditional weighted mean integrated squared error—yields nearly identical results. The estimates we present are also virtually unchanged by varying the degree of the local polynomial, the bandwidth, and/or our choice of kernel function, with the exception that a wider bandwidth results in a smoother function that no longer resembles the fixed-effects estimates and a narrower bandwidth results in a more jagged function.

Causal Inference for Event Studies

Here we describe the procedure used to obtain p-values for Table 2 in the body of the paper. The method we use is based on Gelbach et al (2013), who propose a similar estimator for causal inference over single-firm event studies. Most methods for robust causal inference in event studies consider the case where there are multiple firms affected by the same event, or the same type of event. This setting differs from ours because we have only one time series over which to examine the impact of events.

First, we estimate our statistical model using the procedure outlined in the previous section. After obtaining the residuals from this regression, we construct placebo event effects. Eligible placebo event dates consist of any date in the sample period that does not occur within five days of an event we study in Table 2, so that the five-day event windows never overlap. For every eligible placebo event date t , we construct a placebo event effect by adding the residual from the regression from t to $t+4$.

One can think of the distribution of placebo events as an approximation to the distribution of five-day effects in the absence of an event that significantly influences information search. The p-value reported in Table 2 is the fraction of placebo event effects that exceed the estimated effect of an event we specify, such as the release of Mitt Romney’s tax return. For example, the release of Mitt Romney’s 2011 tax return generated an estimated cumulative effect on Google searches of 7.17. Placebo event effects exceeded 7.17 in 0.3 percent of the time, so our p-value for this estimate is 0.003. The p-values we obtain using this method are slightly larger than p-values we would obtain from a classical F-test for the joint significance of coefficients for event days, which should be expected if there is some serial correlation in the data.

References

- Alstadsæter, A., W. Kopczuk, and K. Telle. (2010). Social Networks and Dynamics of Tax Avoidance: Evidence from a Well-defined Norwegian Tax Shelter. Working paper.
- Blendon, R.J., S.R. Pelletier, M.D. Rosenbaum, and M. Brodie. (2003). "Tax Uncertainty." *Brookings Review* (Summer): 28-31.
- Brown C.V. (1968). "Misconceptions about Income Tax and Incentives." *Scottish Journal of Political Economy* 16(2): 1-12.
- Chang, E. C., and J.M. Pinegar. (1986). "Return Seasonality and Tax-loss Selling in the Market for Long-term Government and Corporate Bonds." *Journal of Financial Economics* 17(2): 391-415.
- Chetty, R., A. Looney, and K. Kroft. (2009). "Salience and Taxation: Theory and Evidence." *American Economic Review* 99(4): 1145-1177.
- comScore. (2005). comScore Releases July 2005 U.S. Search Engine Rankings. <http://ir.comscore.com/releasedetail.cfm?ReleaseID=245083>. August 19, 2005.
- comScore. (2007). comScore Releases November U.S. Search Engine Rankings. <http://ir.comscore.com/releasedetail.cfm?ReleaseID=283126>. December 21, 2007.
- comScore. (2009). comScore Releases June 2009 U.S. Search Engine Rankings. <http://ir.comscore.com/releasedetail.cfm?ReleaseID=397028>. July 16, 2009.
- comScore. (2013). comScore Releases Oct. 2013 U.S. Search Engine Rankings. <http://ir.comscore.com/releases.cfm?Year=2013&ReleasesType=&PageNum=1>. November 13, 2013.
- Connor, C.E., H.E. Egeth, and S. Yantis. (2004). "Visual Attention: Bottom-Up versus Top-Down." *Current Biology* 14: R850-R852.
- Da, Z., J. Engelberg, and P. Gao. (2011). "In Search of Attention." *Journal of Finance* 66(5): 1461-1499.
- de Bartoleme, C.A.M. (1995). "Which Tax Rate Do People Use: Average or Marginal?" *Journal of Public Economics* 56(1): 79-96.
- Delli Carpini, M.X. and Keeter, S. (1996). *What Americans Know about Politics and Why It Matters*. New Haven, CT: Yale University Press.
- Drake, M.S., D.T. Roulstone, and J.R. Thornock. (2012). "Investor Information Demand: Evidence from Google Searches around Earnings Announcements." *Journal of Accounting Research* 50(4): 1001-1040.
- Fujii, E.T., and C.B. Hawley. (1988) "On the Accuracy of Tax Perceptions." *The Review of Economics and Statistics* 70(2): 344-347.

- Gelbach, J.B., E. Helland, and J. Klick. (2013). "Valid Inference in Single-Firm, Single-Event Studies." *American Law and Economics Review* 15(2): 495-541.
- Goldin, J. (2014). "Optimal Tax Salience." Unpublished.
- Graham, J.R., M. Hanlon, T. Shevlin, and N. Shroff. (2014). "Tax Rates and Corporate Decision Making." Unpublished.
- Hardle, W. and O. Linton. (1994). "Applied Nonparametric Methods." In *Handbook of Econometrics Volume IV*, edited by R.F. Engle and D.L. McFadden, pp. 2295-2239. Elsevier.
- Hoopes, J. (2014). "Financial Accounting Consequences of Temporary Tax Law: Evidence from the R&D Tax Credit." Unpublished.
- Ivković, Z., J. Poterba, and S. Weisbenner (2005). "Tax-Motivated Trading by Individual Investors." *American Economic Review* 95(5): 1605-1630.
- Liebman, J.B. and R.J. Zeckhauser. (2004). "Schmeduling." Unpublished.
- Keppo, J., G. Moscarini, and L. Smith. (2008). "The Demand for Information: More Heat than Light." *Journal of Economic Theory* 138(1): 21-50.
- Muth, J. F. (1961). "Rational Expectations and the Theory of Price Movements." *Econometrica* 29(3): 315-335.
- Poterba, J.M. and S.J. Weisbenner. (2001). "Capital Gains Tax Rules, Tax Loss Trading, and Turn-of-the-Year Returns," *Journal of Finance* 56(1): 353-368.
- Reck, D.H. (2014). "Taxes and Mistakes: What's in a Sufficient Statistic?" Unpublished.
- Reis, R. (2006). "Inattentive Consumers." *Journal of Monetary Economics*, 53(8), 1761-1800.
- Robinson, P.M. (1998). "Root-N-Consistent Semiparametric Regression." *Econometrica* 56(4): 931-954.
- Sahm, C.R., M.D. Shapiro, and J.B. Slemrod. (2010). "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications." In *Tax Policy and the Economy, Volume 24*, edited by Jeffrey R. Brown, pp. 69-110. Chicago: Chicago University Press.
- Sheffrin, S.M. (1994). "Perceptions of Fairness in the Crucible of Tax Policy." In *Tax Progressivity and Income Inequality*, edited by Joel Slemrod, pp. 309-334. Cambridge: Cambridge University Press.
- Sims, C. A. (2003). "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3): 665-690.
- Slemrod, J. (1995). "The Compliance Cost of Income Taxation, and How to Reduce It," National Commission on Economic Growth and Tax Reform (Kemp Commission), Oral testimony given July 12, 1995.
- Slemrod, J., Christian, C., London, R., and Parker, J. A. (1997). "April 15 Syndrome." *Economic Inquiry* 35(4): 695-709.

- Slemrod, J. (2004). "New Estimates of the Compliance Cost of Income Taxation," Committee on Ways and Means, Subcommittee on Oversight, Hearing on Tax Simplification, invited written testimony given June 15, 2004.
- Slemrod, J. (2013). "Buenas Notches: Lines and Notches in Tax System Design." Unpublished.
- Starks, L.T., L. Yong, and L. Zheng. (2006). "Tax-Loss Selling and the January Effect: Evidence from Municipal Bond Closed-End Funds." *Journal of Finance* 61(6): 3049-3067.
- Tax Foundation. (2010). Federal Capital Gains Tax Collections, 1954-2009. <http://taxfoundation.org/article/federal-capital-gains-tax-collections-1954-2009>. Accessed 5/7/2014.
- Theeuwes, J. (1991). "Exogenous and Endogenous Control of Attention: The Effect of Visual Onsets and Offsets." *Perception & Psychophysics* 49(1): 83-90.

TABLE 1: TESTABLE IMPLICATIONS OF THEORIES OF INFORMATION AND ATTENTION

Table 1: Testable Implications of Theories of Information and Attention

Behavioral assumption	Should information search respond to...					
	time notches?	macroeconomic change?	enactment of policy changes?	implementation of policy changes?	filing deadlines?	news events?
Full information	No	No	No	No	No	No
Full ignorance	No	No	No	No	No	No
Rational expectations	No	No	Yes, instantly	No	No	No
Rational attention	Yes	Yes	Yes	Yes	Yes	Unclear*
Exogenous Attention (Salience)	No	No	Yes, while in the news	Yes, while in the news	Yes**	Yes
Summary of Our Findings	Yes: Dec 31 time notch, esp. capital losses (due to loss harvesting)	Yes: day-of stock market trading volume, response to stock market crash	Yes, in 3-4 day window: ATRA, presidential debates and elections	Yes: May 2008 stimulus	Yes: seasonality due to tax season	Yes, in 3-4 day window: release of Mitt Romney's tax returns

Notes: *News events unrelated to the other types of events listed here may inspire attention rationally due to the desire to be informed about current events, but they will not inspire rational attention for the purpose of improving economic decisions, such as capital gains realizations or tax compliance. ** Our answering yes is based on the observation that, as the deadline approaches, news stories and private discussions of tax payments are common.

TABLE 2: ESTIMATING THE IMPACT OF EVENTS ON
INFORMATION SEARCH FOR CAPITAL GAINS TAXES

	Google Searches	Wikipedia Page Views	Calls to IRS
Obama/McCain Debate	12.08 93.89 (0.001)	5.36 2662.51 (0.016)	-0.22 -16.17 (0.718)
Obama Elected	16.08 124.66 (<0.001)	5.87 2914.76 (0.012)	-0.16 -12.10 (0.655)
Mitt Romney's 2010 tax return released	20.72 161.00 (<0.001)	34.96 17351.70 (<0.001)	0.38 -559.08 (0.196)
Mitt Romney's 2011 tax return released	7.16 55.69 (0.003)	12.65 6277.82 (<0.001)	-0.34 -32.47 (0.810)
Obama/Romney Debate	9.39 72.99 (0.002)	9.88 4904.11 (0.002)	0.25 -17.69 (0.258)
Obama Re-Elected	8.94 71.02 (0.002)	14.56 7228.50 (<0.001)	-0.31 -14.27 (0.794)
American Taxpayer Relief Act signed	15.31 118.98 (<0.001)	9.62 4776.24 (0.003)	-6.17 -188.52 (0.989)
Sample Period	Jan 2, 2004— March 21, 2013	Dec 9, 2007— March 31, 2013	Jan 2, 2002— March 31, 2013
Number of Days	3364	1938	4100

Notes: The event dates listed as dependent variables represent, in order, Oct 15, 2008 (Presidential debate between Barack Obama and John McCain), November 4, 2008 (Barack Obama elected President), January 24, 2012 (Mitt Romney releases his 2010 tax return), September 21, 2012 (Mitt Romney releases his 2011 tax return), Oct 16, 2012 (Presidential debate between Barack Obama and Mitt Romney), Nov 6, 2012 (Barack Obama re-Elected), Jan 2, 2013 (American Taxpayer Relief Act signed into law, includes an increase in top capital gains tax rate). For each event, the top number reports the cumulative information search attributed to this event, in daily standard deviation units, added over the five-day event window. The second number reports the same estimate in the original units of the search volume measure, i.e. the Google Trends index or, Wikipedia page views, or number of calls to the IRS. The bottom number, in parenthesis, reports the p-value from a one-sided non-parametric permutations test of the hypothesis is that the event had positive impact on information search (refer to the appendix for details), i.e. that the variation in information searches over the event window is insignificant. The estimation controls for variation due to yearly and weekly seasonality.

TABLE 3. INFORMATION SEARCH AND STOCK MARKET ACTIVITY

Sample Period	Jan 2, 2004— Dec 29, 2012	Dec 9, 2007— Dec 29, 2012	Jan 7, 2002— Dec 29, 2012	Sep 1, 2008— Feb 28, 2009	Sep 1, 2008— Feb 28, 2009	Jan 2, 2004— Dec 29, 2012, except Sep 1, 2008— Feb 28, 2009	Dec 9, 2007— Dec 29, 2012, except Sep 1, 2008— Feb 28, 2009
Dependent Variable:	Google Searches (Standardized)	Wikipedia Page Views (Standardized)	Calls to IRS (Standardized)	Google Searches (Standardized)	Wikipedia Page Views (Standardized)	Google Searches (Standardized)	Wikipedia Page Views (Standardized)
Log Trading Volume	0.287 (0.156)* [0.138]**	0.306 (0.159)* [0.139]**	0.043 (0.075) [0.069]	0.763 (0.801) [0.789]	0.926 (0.470)* [0.352]**	0.216 (0.159) [0.128]*	0.385 (0.171)** [0.156]**
Log Trading Volume Lag 1	-0.051 (0.146) [0.120]	-0.072 (0.151) [0.100]	0.099 (0.070) [0.044]**	-0.097 (0.735) [0.599]	0.263 (0.415) [0.251]	-0.111 (0.149) [0.126]	-0.093 (0.162) [0.115]
Log Trading Volume Lead 1	0.062 (0.177) [0.160]	-0.023 (0.181) [0.120]	-0.038 (0.085) [0.071]	-0.083 (0.921) [0.727]	-0.123 (0.553) [0.353]	0.105 (0.179) [0.159]	-0.05 (0.193) [0.131]
Number of Days	2252	1266	2753	124	124	2128	1142


Notes: The dependent variable and log trading volume are de-trended prior to estimation, and the dependent variable is standardized by dividing by the standard deviation of the detrended data. We control for weekly and yearly seasonality and the events in Table 2 (except those events that occur outside the sample period, since market data are not available past December 29, 2012). Weekends and holidays are omitted, since stock market data are not generated on weekends. The regression included five market-dated leads and lags of each measure. We only report one lead and one lag for brevity and clarity (See Online Appendix Figure A.1). The coefficients are similar if we use 14 leads and lags instead of 5, and none of the results change substantially with the inclusion of lags of the dependent variable. Classical standard errors are provided in round parentheses below point estimates. Newey-West standard errors allowing for up to 5 (market-dated) lag orders of autocorrelation are reported in square brackets. Newey-West standard error estimates are unchanged to two significant digits with 7 lag orders instead of 5. * indicates $p < 0.10$, ** indicates $p < 0.05$, and *** indicates $p < 0.01$. The standard deviation of log trading volume is 0.223.

TABLE 4: SEARCHES FOR STOCK ADVICE AND SEARCHES FOR CAPITAL GAINS TAX

Sample Period	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Jan 2, 2004— Dec 29, 2012	Jan 2, 2004— Dec 29, 2012	Jan 7, 2004— Dec 29, 2012	Jan 7, 2004— Mar 21, 2013	Dec 9, 2007— Dec 29, 2012	Dec 9, 2007— Dec 29, 2012	Dec 9, 2007— Dec 29, 2012	Dec 9, 2007— Mar 31, 2013
Dependent Variable:	Google Searches (Standardized)				Wikipedia Page Views (Standardized)			
Log Trading Volume	0.287 (0.156)* [0.138]**	0.176 (0.155) [0.143]	0.130 (0.155) [0.140]		0.306 (0.159)* [0.139]**	0.162 (0.157) [0.140]	0.178 (0.157) [0.136]	
Log Trading Volume Lag 1	-0.051 (0.146) [0.120]	-0.096 (0.145) [0.117]	-0.06 (0.146) [0.119]		-0.072 (0.151) [0.100]	-0.141 (0.147) [0.101]	-0.137 (0.148) [0.101]	
Log Trading Volume Lead 1	0.062 (0.177) [0.160]	0.022 (0.175) [0.159]	0.002 (0.175) [0.158]		-0.023 (0.181) [0.120]	-0.056 (0.177) [0.118]	-0.098 (0.178) [0.119]	
Stock Advice Search		0.031 (0.004)*** [0.008]***	0.008 (0.005) [0.007]	0.009 (0.005)* [0.007]		0.030 (0.004)*** [0.008]***	0.016 (0.006)*** [0.005]***	0.017 (0.005)*** [0.005]***
Stock Advice Search Lag 1			0.017 (0.006)*** [0.007]**	0.017 (0.006)*** [0.007]**			0.017 (0.008)** [0.006]***	0.015 (0.008)* [0.006]**
Stock Advice Search Lead 1			0.009 (0.005)** [0.005]**	0.008 (0.005)* [0.004]*			0.005 (0.006) [0.004]	0.003 (0.006) [0.003]
Number of Days	2252	2239	2200	2211	1266	1266	1266	1272


Notes: The dependent variable, log trading volume, and stock advice search volume are de-trended prior to estimation, and the dependent variable is standardized by dividing by the standard deviation of the detrended data. We control for weekly and yearly seasonality and the events in Table 2 (except those events that occur outside the sample period, since market data are not available past December 29, 2012). Weekends and holidays are omitted, since stock market data are not generated on weekends. The regression included five market-dated leads and lags of trading volume and stock advice search volume. We only report one lead and one lag for brevity and clarity (See Online Appendix Figure A.3). Classical standard errors are provided in round parentheses below point estimates. Newey-West standard errors allowing for up to 5 (market-dated) lag orders of autocorrelation are reported in square brackets. Newey-West standard error estimates are unchanged to two significant digits with 7 lag orders instead of 5. * indicates $p < 0.10$, ** indicates $p < 0.05$, and *** indicates $p < 0.01$. The standard deviation of log trading volume is 0.223. The standard deviation of Stock advice search volume is 4.458.

FIGURE 1: PLOTS OF THE RAW DATA



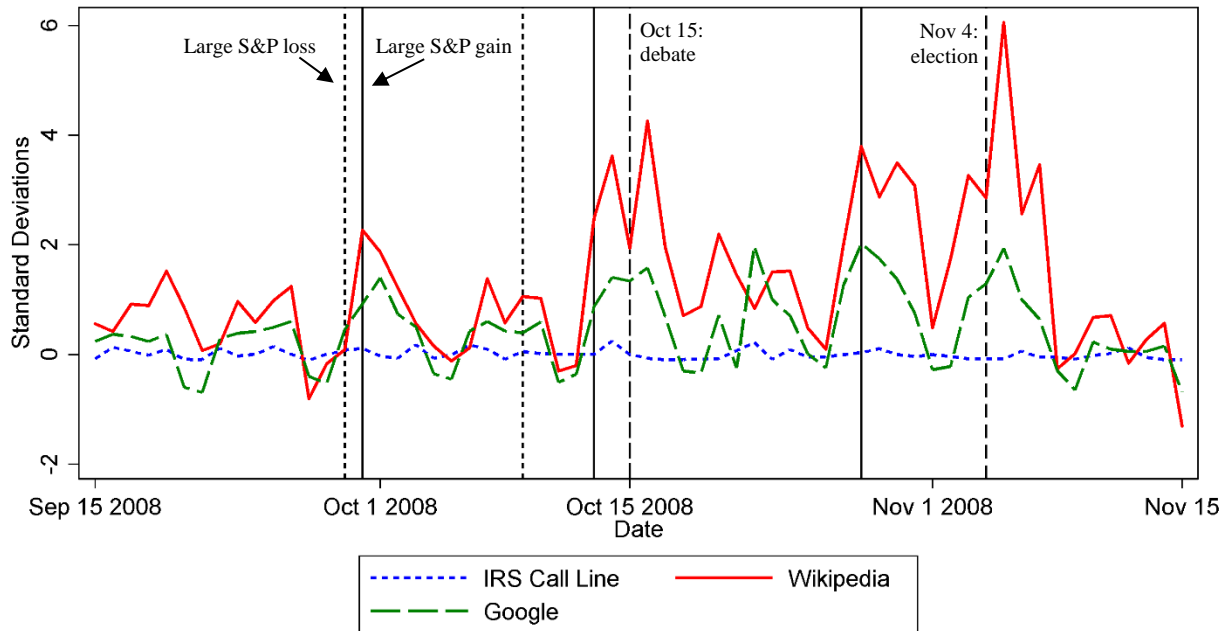
Notes: These graphs show the raw data for (in order) Google search volume, Wikipedia page views, and calls made to the IRS.

FIGURE 2: YEARLY SEASONALITY



Notes: Yearly seasonality is estimated using the smooth-fixed-effects method described in Section 5, with controls for the events in Table 2 and the market movements in Table 4. The thin dashed lines are the bounds of 95 percent confidence intervals.


FIGURE 3: CAPITAL GAINS TAX INFORMATION SEARCH IN OCTOBER 2008



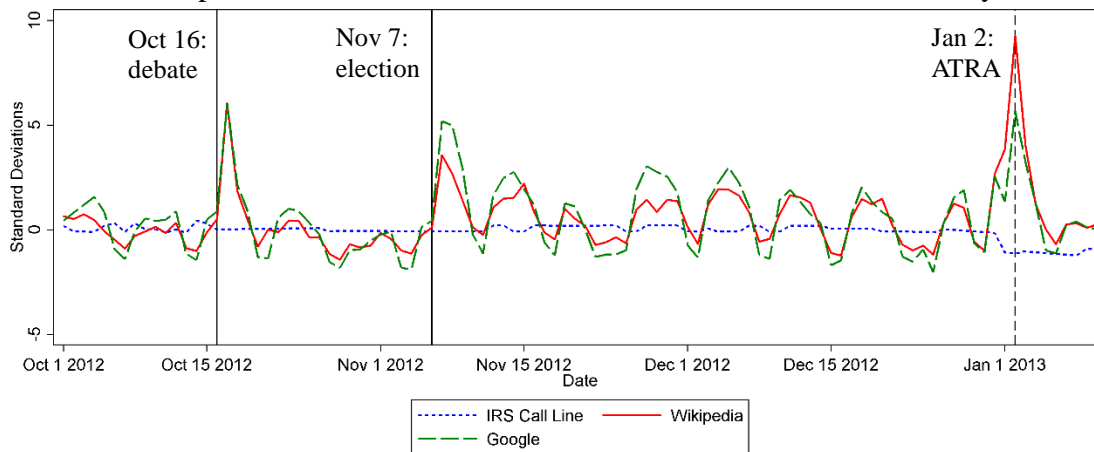
Notes: We plot the detrended, standardized, and seasonally adjusted data over a narrower time period to examine abnormal behavior in late 2008. The vertical dotted lines mark the worst three days of the stock market crash of 2008 according to the drop in the S&P 500 index: October 15 (-9.03 percent) , September 29 (-8.81 percent), and October 9 (-7.62 percent). The vertical solid lines mark the three largest gains in the S&P 500 during this volatile period: October 13 (11.58 percent), October 28 (10.79 percent) and September 30 (5.42 percent). October 15, 2008, marked with a vertical dashed line, was also the date of a Presidential Debate between Barack Obama and John McCain, in the course of which both candidates made proposals for changing capital gains tax rates. The other debates occurred on September 26 and October 7, and did not discuss capital gains taxes. The other vertical dashed line marks the 2008 Presidential election, November 4, 2008.

FIGURE 4: NARROW TIME WINDOWS OF SEARCH FOR TAX INFORMATION

Panel A. Capital Gains Tax Information Search in January 2012




Panel B. Capital Gains Tax Information Search in October 2012 – January 2013



Notes: In Panel A, we plot the detrended, standardized, and seasonally adjusted data (from Figure 4) over a narrower time period to examine the spike in January 2012. The vertical dashed line marks January 24, 2012, the date that Mitt Romney released his 2010 tax return. The spike a few days prior is coincident with a related news story wherein Mr. Romney told the press about his effective tax rate. In Panel B, the first solid vertical line marks the date of the Presidential debate on October 16, 2012, during which Barack Obama and Mitt Romney debated the merits of Romney's proposals for lowering capital gains tax rates. The second solid line marks the date of the Presidential election, November 7, 2013. The dashed vertical line marks the passage of the American Taxpayer Relief Act on January 2, 2013.


FIGURE 5: INFORMATION SEARCH ON CAPITAL LOSSES




Notes: The top panel plots weekly Google Trends data on searches for just “capital loss.” Google Trends data are missing for some dates in 2004 and 2005, when search volume was too low for Google to provide data. The second graph plots monthly searches in the search bar on the IRS home page for terms related to capital losses. The thick dotted vertical lines correspond to October 1 and December 2008, to delineate the period in which there was increased information search on capital losses due to the 2008 stock market crash. Solid vertical lines denote January 1st of each year, to highlight that a spike in search volume typically occurs at the very end of each year. Dashed vertical lines denote mid-April filing deadlines, to highlight the spike occurring in the run-up to the filing deadline.

FIGURE 6: OVERALL USE OF IRS INFORMATIONAL RESOURCES

Panel A: Visits and Page Views of all URLs on the IRS.gov Domain.



Panel B: Calls to IRS Line for Individuals



Notes: The dotted vertical lines delineate the month of May, 2008. Dashed vertical lines correspond to the mid-April filing deadline for each year. Panel A: Data are unavailable from November 1, 2003 to February 29, 2004. Panel B: The phone number for this line is 1-800-829-1040, a number made available numerous places on the IRS website and on the instructions for individual tax returns (form 1040). Unique callers are identified by the telephone number from which the call originates.