

Discussion Papers
Department of Economics
University of Copenhagen

No. 08-21

Tax Salience, Voting, and Deliberation

Rupert Sausgruber,
Jean-Robert Tyran

Stu­di­estræde 6, DK-1455 Copenhagen K., Denmark
Tel.: +45 35 32 30 82 – Fax: +45 35 32 30 00
<http://www.econ.ku.dk>

ISSN: 1601-2461 (online)

Tax Salience, Voting, and Deliberation

Rupert Sausgruber and Jean-Robert Tyran^{*}

October 2008

Tax incentives can be more or less salient, i.e. noticeable or cognitively easy to process. Our hypothesis is that taxes on consumers are more salient to consumers than equivalent taxes on sellers because consumers underestimate the extent of tax shifting in the market. We show that tax salience biases consumers' voting on tax regimes, and that experience is an effective de-biasing mechanism in the experimental laboratory. Pre-vote deliberation makes initially held opinions more extreme rather than correct and does not eliminate the bias in the typical committee. Yet, if voters can discuss their experience with the tax regimes they are less likely to be biased.

Keywords: Tax salience, learning, deliberation, voting

JEL-codes: C92, H22, D72

^{*} Sausgruber: University of Innsbruck, Department of Public Economics, Universitätsstr. 15, A-6020 Innsbruck and University of Copenhagen, Department of Economics, Studiestræde 6, DK-1455 Copenhagen. Rupert.Sausgruber@uibk.ac.at
Tyran: University of Copenhagen, Department of Economics, Studiestræde 6, DK-1455 Copenhagen, and CEPR (London). Jean-Robert.Tyran@econ.ku.dk
We gratefully acknowledge financial support by the Austrian Science Fund (FWF), under Project No. P17029. We thank Alexander Sebald and Frederik Øvlisen for helpful comments.

1 Introduction

Tax incentives can be more or less salient, i.e. noticeable or cognitively easy to process. Salient taxes loom large in people's minds while taxes are not salient if tax incentives are more opaque and difficult to compute. Accumulating evidence suggests that misperception of tax incentives due to salience effects is common and can distort choices. Chetty, Looney and Kroft (2007) show in a field experiment that sales taxes that are explicitly stated on price tags in a grocery store (and are thus more salient) have larger effects on demand than if the tax is only added at the cashier. Finkelstein (2007) finds that electronic toll collection is less salient than when cash has to be handed at the toll station, and that reduced salience induces higher toll rates, and Blumkin, Ruffle and Ganun (2008) find salience effects in comparing consumption and income taxes in a laboratory experiment. Some evidence suggests misperception of estate tax rates (Slemrod 2006), and there is suggestive evidence of misperception of marginal tax incentives (Liebman and Zeckhauser 2004, Feldman and Katuscak 2007).

Our paper studies salience effects in voting on taxation. The focus of our paper is on how tax salience biases the choice of a tax regime by voting, rather than how individuals react to the salience of existing taxes, as in the literature cited above. We proceed in three steps.

In a first step, we show that salience effects bias voting on commodity taxes in a laboratory experiment. Our hypothesis is that taxes on consumers are more salient to consumers than equivalent taxes on sellers. We assume that consumers find taxing the "other side of the market" rather than taxing one's own side more attractive because consumers underestimate the extent of tax shifting in the market. In our experiment, participants earn incomes by trading a good in a market. After some periods of trading, consumers are told that market transactions need to be taxed in the future and that tax revenues are wasted, i.e. not returned to market participants. Only consumers can vote, and they vote on implementing a buyer tax or a seller tax. To make tax salience effects costly, consumers are presented with a choice between a buyer tax of t and a seller tax of more than t . According to standard economic theory, consumers correctly perceive that higher taxes reduce their market income independent of which side of the market is taxed and, therefore, vote for the low buyer tax. In contrast, salience effects induce consumers to vote for the less salient but more costly seller tax. In line with the salience hypothesis, we find that consumers vote for the seller tax in 2 out of 3 cases.

In a second step, we investigate if pre-vote deliberation mitigates tax salience effects. We compare voting when voters can freely discuss the two tax regimes with a treatment without this option. We show that deliberation is persuasive, i.e. that the quality of messages a voter receives affects his voting. Correct messages tend to improve voting while incorrect messages tend to promote voting for the less salient seller tax. We find that deliberation induces agreement, i.e. that opinions are more homogenous in the presence of deliberation. However, since opinions are divided at the beginning of the deliberation process, voters by chance receive more correct or incorrect messages. As a result, initially held opinions tend to become more extreme and voters agree on more or less arbitrary values (see Glaeser and Sunstein 2007 for a discussion). We therefore find that deliberation does not de-bias voting of inexperienced voters.

In a third step, we repeat the trading and voting sequence to study if experience mitigates salience effects. We find that voters who experienced the income-reducing effects of seller taxes tend to vote for the buyer tax in the sequel. Absent such experience, repeatedly discussing the tax proposal does not reduce salience effects. However, we find that deliberation and experience interact. If voters can discuss their adverse experience with costly seller taxes, they are more likely to vote for the buyer tax in consecutive votes.

Our results show that salience effects bias voting, that experience effectively de-biases voters but that deliberation *per se* does not improve voting in our laboratory environment. These results are surprising from the perspective of standard economic theory and call for taking misperception and learning into account. Below, we motivate some of our design choices and discuss the advantages of an experimental approach to tax salience. In addition, we discuss some related literature.

Our first main result is that salience effects bias individual voting and the choice of tax regimes. While the literature cited in the first paragraph suggests that salience effects can bias market choices, we are the first to show that salience effects can bias voting in a controlled environment. Our result is surprising given that the consequences of tax salience on voting are *a priori* indeterminate. On the one hand, salience effects are likely to be common among voters because incentives to exert cognitive effort are inherently weak in voting. Voters can free-ride on other voters' cognitive efforts because they are unlikely to be pivotal for the outcome. On the other hand, because individual biases aggregate in a non-linear fashion into committee choices, even a considerable share of voters prone to salience effects may not matter for voting outcomes, as long as biased voters are in minority.

While our demonstration of the existence of salience effects in a controlled environment is novel in our view, we believe that investigating the robustness of salience effects to de-biasing mechanisms is also important. We investigate the effects of pre-vote deliberation because of its potential to de-bias voters and because of its practical relevance.¹ We study deliberation in the guise of public, non-structured, free, and non-strategic discussion between voters by relating the content or “quality” of messages to voting.² Deliberation may improve voting outcomes in heterogeneous committees, in which some voters are prone to salience effects while others are not, if the “smart few” persuade at least a majority about the right course of action. In this case, deliberation de-biases committees even if a majority of voters initially falls prey to salience effects. However, deliberation may also induce “group polarization” which makes initially held opinions more extreme but not necessarily more correct. According to Glaeser and Sunstein (2007: 8) there is “pervasive evidence of group polarization on issues that bear directly on politics and political behavior”, and these facts “appear to cast doubt on the wisdom, and certainly the moderation, of crowds” (p. 2).

We study salience effects resulting from tax shifting, i.e. from a discrepancy between who nominally pays the tax and who economically bears its cost. While any undergraduate textbook in economics explains that it is matter of indifference whether a tax is levied on the buyers or the sellers in a market (the so-called Tax Liability Side Equivalence, or tax LSE for short), popular perception differs markedly from the textbook. For example, Graetz (2002: 270) claims that “many families underestimate their payroll tax burden because the employers’ share is hidden to employees”, and public debate in the media and the political arena is often much concerned with statutory incidence (see Ruffle 2004 for examples). However, it is important to note that such a concern for statutory incidence need not necessarily mean that consumers or voters are biased. In fact, statutory incidence does make a difference to economic incidence if prices are rigid and cannot freely adjust to changes in taxation (see Riedl and Tyran 2005 for a discussion). The relevance of price and wage rigidity is much debated in the field, but easy to control in the laboratory. An advantage of our

1 Deliberation is common in many types of committees. For example, trial jurors converse before casting their votes, hiring committees deliberate before making their final decisions, and top management teams hold meetings before determining their firm’s investment strategies (see Gerardi and Yariv 2007 for a discussion). Deliberation in the guise of political discussion is key to the political process. For example, Huckfeldt and Sprague (1987: 1197) note that “politics is a social activity imbedded within structured patterns of social interaction. Political information is conveyed not only through speeches and media reports but also through a variety of informal social mechanisms – political discussions on the job or on the street ... even casual remarks.”

2 There is a considerable literature in experimental economics investigating how the option to communicate affects choices but most studies does not investigate the effects of the content of messages. Notable exceptions, and in this respect the closest match to our approach, are Brandts and Cooper (2007), Charness and Dufwenberg (2006) and Ellingsen and Johannesson (2004).

experimental approach is therefore that we can implement market institutions which are known to have flexible prices and quick equilibration, implying that tax LSE holds.³

Another advantage of an experimental approach to tax salience is its ability to control the environment, and in particular information conditions. Salience effects are hard to pin down in the field because inattention to in principle known information is often difficult to distinguish from sheer lack of information (see Chetty et al. 2007 for a notable exception).

Salience effects seem more plausible when tax schemes are complex because tax incentives are hard to compute in this case.⁴ Decision makers may find publicly available information hard to interpret in this case. While taxation is often complex in the field⁵, we believe that the empirical demonstration of salience effects is particularly convincing and striking when the environment is simple and relatively transparent. We therefore implement a simple commodity (per unit transaction) tax in a single market.

Tax salience is related to a literature going back to at least John St. Mill on “fiscal illusion” which discusses misperception of fiscal parameters more generally, including government spending and debt (e.g. Sausgruber and Tyran 2005, see Dollery and Worthington 2006 for a survey). Tax salience also relates to a literature on salience effects in other, non-fiscal, domains. Salience effects can bias consumers’ perception of the “net” or “final” price and make it less salient than the gross price. This may happen, in the diction of Gabaix and Laibson (2006), when the net price differs from the gross by price by some “shrouded attribute” (see Hossain and Morgan 2006, for a field experiment on the effect of shipping fees in eBay auctions, and DellaVigna 2007 for a review of this literature). Money illusion prevails if nominal, i.e. purely monetary, values are more salient than real values (Shafir, Diamond and Tversky, 1997). Similar to the distinction between net and gross prices, distinguishing real from nominal values requires cognitive effort because nominal values need to be deflated, i.e. corrected for inflation. Inattentive decision makers may therefore use nominal values as a proxy for real values (see e.g. Tyran 2007).

3 Tax LSE has been shown to hold for a broad range of market institutions, like double auction markets (Kachelmeier, Limberg, and Schadewald 1994), posted offer markets (Borck et al. 2002), pit markets (Ruffle 2004), and gift-exchange markets (Riedl and Tyran 2005).

4 McCaffery and Baron (2006: 107): “People decide complex matters – and tax raises a host of complex matters – by responding to the most salient or obvious aspect of a choice set or decision problem. They fail to take into account logically relevant information that is not immediately available to their mental models. In tax, this means, for example, that people ... prefer hidden taxes, however distorting to their own preferred allocation or distribution of resources, over more transparent levies.”

5 For example, Kotlikoff and Rapson (2006:1) argue that “thanks to the incredible complexity of the U.S. fiscal system, it’s impossible for anyone to understand her incentive to work, save, or contribute to retirement accounts absent highly advanced computer technology and software.” A similar argument is made by Liebman and Zeckhauser (2004) for tax and other (e.g. telephone) schedules.

2 Experimental Design

In our experiment, participants earn money by trading a hypothetical good in an initial market phase 0 and then vote on imposing a commodity tax to be levied on transactions in the following market phase. Participants are informed that markets need to be taxed and that all tax revenues are wasted, i.e. not returned to participants in any way and that only buyers can vote. In the main treatments, the choice is between a buyer tax of $t = 25$ and a seller tax τ of more than 25. The level of the seller tax is endogenously determined (see section 2.2 for details).

According to standard economic theory, higher taxes translate into lower market rents for consumers in a market with flexible prices (see Fullerton and Metcalf 2002). This relation holds irrespective of whether the tax is levied on the buyers or the sellers and follows from the so-called Tax Liability Side Equivalence (Tax LSE) which is a fundamental principle in public economics.⁶

Figure 1 illustrates the consequences of taxation in the two tax regimes. S_0 and D_0 show supply and demand prior to taxation and the pre-tax market equilibrium is at point A, at a quantity of $q_0 = 10$ units and a price of $p_0 = 100$. In the buyer tax regime (see left panel) demand shifts down to D_1 resulting in equilibrium B with $q_t = 5$ and $p_t = 90$.⁷ In the seller tax regime with $\tau > 25$, supply shifts up, resulting in the after-tax equilibrium C with the same quantity of $q_t = 5$ as in case of the buyer tax. However, with $\tau > 25$ prices increase by more than 25 points and consumer rent is smaller than with the buyer tax of $t = 25$. Hence, in competitive equilibrium $t = 25$ dominates $\tau > 25$ from the consumers' perspective.


Note that voters only need minimal information to correctly predict that buyer earnings are smaller with the seller tax. For example, voters do not need to know overall market parameters (e.g. demand and supply elasticities) or how much of the tax is shifted.⁸ All they need to know is that higher taxes reduce their incomes and that market prices respond quickly to shifts in demand and supply.

6 Blinder (1988: 12) describes tax LSE as the "...most basic theorem of public finance: the irrelevance of the side of the market on which a tax is levied."

7 Note that equilibrium prices fall by less than 25 points due to partial tax shifting. In Sausgruber and Tyran (2005) we investigated the more extreme cases of no or complete tax shifting.

8 The extent of tax shifting is market specific. Besley and Rosen (1999) analyze shifting of sales taxes in the US and "find a surprising variety of shifting patterns." (p. 158). A study for France (Carbonnier 2007) finds that a specific sales tax is shifted by 57% for new cars, and by 77% for housing repair services.

Figure 1: Induced supply and demand: buyer tax (left) and seller tax (right)



Our experiment has 3 treatments. The baseline treatment NO is as explained above. Treatment DELIB is the same as treatment NO, except that the voting stage is preceded by a deliberation stage and serves to isolate the effects of deliberation on tax salience. Treatment CON is the same as NO, except that the default tax is reversed. In CON, the choice is between a seller tax of 25 and a buyer tax of more than 25. Treatment CON serves to identify the bias and to discriminate the salience hypothesis from random or rational voting.

2.1 Details on the market mechanism

We use a uniform-price sealed bid/offer auction (a simplified version of Smith et al. 1982) with the supply and demand parameters shown in Figure 1 in all treatments. Using this market institution has three important advantages for our purposes. First, the market is known from previous experiments to quickly converge to equilibrium predictions implying that salience effects are costly in our design. Second, trading in our market is simple and is easy to explain to participants allowing participants to concentrate on the choice of the tax regime which is the main focus of our paper. Third, this market institution allows us to automate sellers which simplifies the analysis and enables us to isolate salience effects from fairness considerations.

The details of the market are as follows. A market has five buyers and automated sellers. Each buyer has privately known reservation values for two units $v_1 = 155$ and $v_2 = 95$. Buyers submit bids for each unit they want to buy. Buyers can bid at most their induced value per unit. The automated sellers have costs per unit ranging from 85 points to 100 points and

submit offers at true cost for each unit to sell. Buyers do not know the unit costs of the sellers, but do know that sellers submit offers at the true cost. We automate sellers to control for fairness considerations. Kerschbamer and Kirchsteiger (2000) argue that tax LSE may fail to hold if subjects perceive shifting the tax burden as unfair. Since we automate the sellers and since this is known to buyers, fairness considerations should not affect voting, allowing us to study salience effects in isolation.

Market outcomes are determined as follows: Bids are ordered from high to low, and offers from low to high. All bids exceeding offers are accepted, and the last accepted bid determines the uniform price for all units traded. A buyer's payoff per traded unit is the difference between the induced value and the market price or zero, if he does not trade. In our experiment, both buyer and seller taxes are transaction taxes, i.e. the tax only depends on the number of transactions, not on their value. We chose the tax structure deliberately simple to make tax incentives transparent. A tax on sellers adds to their cost. Since automated sellers bid their true (tax inclusive) cost, implementing a seller tax of τ shifts the supply function in figure 1 up by τ . Similarly, a tax on buyers shifts the demand function down by t .

Note that the market price is determined by the last accepted bid. Therefore, the bids of the (marginal) buyers affects market prices and buyers experience price fluctuations in the 10 periods of trading in phase 0. Phase 0 serves to familiarize participants with trading in the market and to make sure that participants experience that prices are flexible, which is a necessary condition for tax LSE to hold. Indeed, in phase 0 participants can see that market prices depend on participants' choices and that prices are flexible. In period 1 of phase 0, the variance of market prices is 5.3 points. Phase 0 also sets the stage for our analysis of voting behavior because markets converge well to equilibrium predictions. For example, market efficiency, defined as actual over equilibrium rents, averaged across all treatments is 98.4 percent in the last period of phase 0.

2.2 Details on voting

In the main treatments NO and DELIB, the choice is between a buyer tax of 25 and a seller tax of more than 25. Rather than having voters choose between two exogenous taxes, we let them choose between an exogenously given buyer tax of 25, and an endogenously determined seller tax. The value of the seller tax is determined by an incentive-compatible mechanism designed to elicit the maximum seller tax a buyer is willing to accept to avoid the buyer tax.

The purpose of endogenizing the seller tax is to induce voters to think about how much they prefer the seller tax over the buyer tax, and to provide them with incentives to state their true preference.

The details of voting are as follows. Voters are in committees of five. Only buyers can vote and voting is compulsory. Voters simultaneously choose integer seller tax rates τ_i between 0 and 50 points.⁹ A random draw τ from a uniform distribution with support $[25, 50]$ determines if the seller tax regime is implemented and at what rate. If the random draw is below the median vote, the seller tax regime is implemented at rate τ . If the random draw is above the median vote, the buyer tax regime is implemented at rate $t = 25$. Therefore, the implemented seller tax is above 25 by design if the seller tax regime prevails in the main treatments. If the choice of the median voter is 25 or less, the committee is said to have voted for the buyer tax regime and the buyers are taxed at $t = 25$ in the following trading phase in this case. If the median choice is above 25, the committee is said to have voted for the seller tax regime and sellers are taxed at the randomly determined rate τ if the median voter was willing to accept a seller tax of at least τ .¹⁰ Intuitively, the mechanism can be thought of as an incentive-compatible shortcut to majority voting. It is equivalent to the outcome of a vote between t and an exogenous τ .

The implemented seller tax is determined by a random draw to make voting incentive-compatible. Suppose a voter thinks as the standard economist who is indifferent between a buyer tax of 25 and a seller tax of 25. Such a voter casts a vote of 25 or less. Suppose a voter is prone to salience effects and believes that taxes on buyers fall to a greater extent on buyers than equivalent sellers taxes. That is, the voter finds the incidence of buyer taxes more salient than the incidence of seller taxes. Such a voter would be indifferent between the buyer tax of 25 and a seller tax of more than 25. Our mechanism provides incentives to cast a vote according to his indifference point rather than to strategically shade his vote.¹¹ The intuition why the mechanism is incentive compatible is that the vote τ_i only affects the probability that

9 Admitting votes for seller taxes below 25 avoids confound with unsystematic decision errors. To see why, assume that we had confined voting choices to the interval $[25, 50]$ and assume that voters state $\tau_i = 25 + \varepsilon$, where ε is a random component following some symmetric distribution. In this case, the median vote would suggest a bias of $\varepsilon/2$ despite errors being unsystematic.

10 To ease the understanding for experimental subjects, the instructions refer to the buyer tax regime as “Party A” and to the seller tax regime as “Party B”. See Appendix A3 for instructions.

11 To illustrate that incentives to vote strategically prevail absent the random draw, suppose a voter is indifferent between a buyer tax of 25 and a seller tax of 35, say. If the voter thinks that the seller tax is partly shifted to the buyers, he has an incentive to vote for the minimal seller tax guaranteeing that the seller tax regime is implemented (i.e. 26).

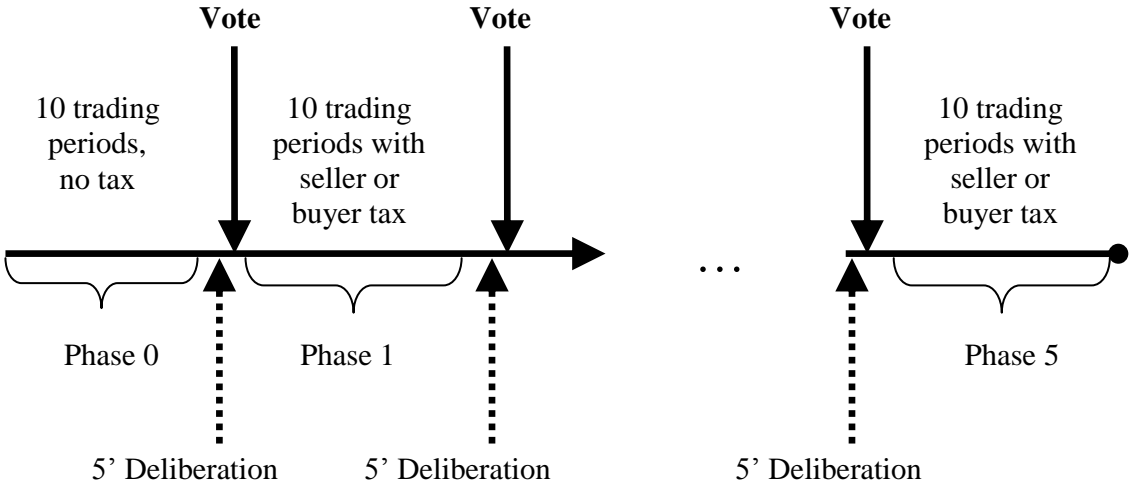
the seller tax is implemented at a subjectively preferred rate; it does not affect the distribution from which the seller tax rate is drawn (see Appendix A1 for a proof).¹²

The main advantage of this procedure is that it provides incentives to state the true preferences and that it is relatively simple to explain. The drawback is that observed outcomes (i.e. τ) are random and always below the median vote by design. We are therefore careful in distinguishing between individual votes, median (committee) votes and implemented taxes in our discussion of results.

2.3 Time line and feedback

To study the de-biasing effects of experience, we repeat the main phase 5 times (see figure 2). Each main phase consists of 2 stages in treatment NO: A voting stage and a market stage. The voting stage is explained both orally and in writing before the first vote (see Appendix A3 for instructions). Participants learn the outcome of the ballot (median tax vote and τ) at the end of the voting stage. The market stage has 10 periods with a buyer or a seller tax regime depending on the outcome of the voting stage.

Figure 2: Time line



12 The same intuition explains why an agent has an incentive to bid his true value in the second price auction. There, the agent chooses a bid which maximizes his expected payoff when the price (the second highest bid) is independent of the own bid.

2.4 Details on deliberation

Treatment DELIB is the same as treatment NO, except that a deliberation stage precedes voting. Deliberation is organized via computerized, written chat. Participants can send messages which are posted to all members of their committee at any time during 5 minutes. We ask subjects not to reveal their identity, induced values or personal information (participants have fixed ID numbers, see Appendix A2), and not to threaten or insult other subjects. The experimenter monitors participants' messages online to enforce compliance with these rules.

In our experiment, deliberation is public, non-structured, free, and non-strategic. It is public in the sense that voters post messages on a (within-committee) public chat board. It is non-structured in the sense messages are not constrained with respect to origin, sequence, frequency, and length and no priority is given to particular senders or messages. Deliberation is free as there is no cost of sending or receiving messages. Finally, deliberation is non-strategic in the sense that voters are in a common interest situation, i.e. face exactly the same incentives in voting (recall that only buyers can vote).

A total of 205 subjects participated in our experiment as follows: 75 subjects in treatment NO, 60 subjects in DELIB, and 70 in the control treatment CON. Participants were undergraduate students from various majors at the University of Innsbruck. An experimental session lasted approximately 2 hours. The average participant earned €24.6, including a show-up fee of €4. The experiments were run using the software z-Tree (Fischbacher 2007).

3 Results

Our main results are as follows. Section 3.1 shows that salience effects systematically bias voting absent deliberation and experience. Section 3.2 shows that participants actively discussed the tax regimes in DELIB. Voters tend to be influenced by what other say, i.e. we find that the quality of messages received is related to the recipient's voting. We find that deliberation induces committees to agree – but not necessarily on the income-maximizing action. Deliberation seems to make initially prevailing opinions more extreme rather than more correct. Overall, the positive and negative effects of deliberation cancel out, and deliberation does not eliminate salience effects, absent experience. Section 3.3 discusses the effects of learning from experience. We find that experiencing the adverse effects of less

salient but more costly taxes improves voting, and that deliberation promotes learning in the presence of clear evidence.

3.1 Salience effects among inexperienced voters absent deliberation

We now show that inexperienced voters were prone to tax salience effects absent deliberation. That is, we show that buyers systematically vote for the less salient but more costly seller tax in phase 1. The salience hypothesis is tested against two alternatives. First, the confusion hypothesis is that voters cast random votes. Second, the standard economics hypothesis is that voters do not vote for taxes above 25. To discriminate between these hypotheses, we compare voting in treatment NO and CON which both had no pre-vote deliberation. The two treatments differ with respect to the default tax. In NO, the default is a buyer tax of $t = 25$ and participants vote on a seller tax. In CON, the default is a *seller* tax of 25 and participants vote on a buyer tax.


According to the confusion hypothesis, votes are randomly distributed in both NO and in CON. Since the default tax is located in the center of the choice range (from 0 to 50), the confusion hypothesis predicts a symmetric distribution of votes, i.e. equal shares of votes above and below 25 in both NO and CON. The standard economics hypothesis predicts 0 percent votes above 25 in both treatments. In contrast, the tax salience hypothesis predicts a bias, i.e. an asymmetric distribution of votes in the two treatments. According to the salience hypothesis, voters prefer to avoid the more salient buyer tax in both cases. In particular, voters are predicted to vote for high (seller) taxes to avoid the buyer tax in NO, but to vote for low (buyer) taxes to avoid the buyer tax in CON. Thus, the salience hypothesis predicts a higher share of votes above 25 in NO than in CON.

Figure 3 shows the distribution of tax votes of inexperienced voters in treatments NO and CON. In NO, the distribution of votes is clearly inconsistent with the standard economics hypothesis since 57 percent of votes are for taxes above 25. A comparison with CON suggests that a considerable share of these votes can be explained by tax salience effects since only 24 percent of the votes are for taxes above 25 in CON.

Of the three hypotheses under consideration, only tax salience can account for the observed bias, i.e. the asymmetry in voting in the two treatments. Figure 3 shows that tax votes above 25 are more than twice as common in NO than in CON (57 vs. 24 percent). The

average tax vote¹³ is significantly different ($p = 0.000$, Mann-Whitney (MW) test) and the distribution of tax choices is significantly different across treatments ($p = 0.002$, Kolmogorov-Smirnov (KS) test). We conclude that voters are heterogeneous with respect to how receptive they are to salience effects. In particular, we find that about a third (= 57% - 24%) of tax votes in NO can be attributed to systematic tax salience effects.

Figure 3: Distribution of tax votes of inexperienced voters (Phase 1)



How do individual tax votes translate into committee choices? In CON, only 1 out of 14 committees (7.1 percent) has a median tax vote > 25 and the median tax vote averaged across all committees is 25.1 points, i.e. very close to the rational choice benchmark of 25. In contrast, in NO, two thirds (10 out of 15) of all committees have a median tax vote > 25 and the median vote averaged over all markets is 31.5 points. The share of committees accepting taxes above 25 is significantly higher in NO than CON (χ^2 -test, $p = 0.001$), and the median tax votes are significantly higher in NO than CON (MW, $p = 0.000$). We conclude that tax salience significantly biased voting outcomes at the committee level.

Committee votes are linked to market outcomes through the incentive-compatible voting mechanism which involves a random draw to induce voters to reveal their true preference. While there was one committee that voted for taxes above 25 in CON, no committee experienced such a tax in phase 1. In NO, 10 committees voted for taxes above 25 and 3 markets experienced such a tax (with an average tax rate of 31.3).

13 In calculating average tax vote we count all tax votes below 25 as equal to 25. The reason is that stating a tax vote below 25 rather than exactly 25 does not affect which tax regime is implemented.

How do committee votes translate into market outcomes? As predicted by standard theory, higher taxes reduce market incomes for buyers. In particular, buyers in markets with taxes above 25 in NO earn less than with taxes of 25 in NO (25.9 vs. 32.6, $p = 0.061$) and in CON (25.9 vs. 32.2, $p = 0.058$). Also in line with standard theory, we find that buyer and seller taxes at a given rate have the same market consequences, i.e. tax LSE holds. In particular, we find that buyer earnings in markets with a tax of 25 do not differ between the NO and CON (32.6 vs. 32.2, $p = 0.897$). We conclude that our market institution generates outcomes predicted for perfectly competitive markets, a finding that is in line with results from numerous experiments. For the sake of brevity, we therefore abstain from providing a detailed discussion of how taxation affects prices and quantities and focus on the effects of deliberation instead.

3.2. Deliberation and salience effects

Does pre-vote deliberation mitigate the bias resulting from tax salience? To investigate, we compare treatments NO and DELIB. Treatment DELIB is the same as NO, except that voters are given the option to communicate during 5 minutes by sending each other messages before the vote. Recall that deliberation is public, non-structured, free, and non-strategic. Deliberation is also voluntary in the sense that voters cannot be forced to send messages or to pay attention to messages they receive. Yet, voters used the option to deliberate extensively. In the deliberation stage of phase 1, the average number of messages sent per voter is 6.5, and the median is 5.0.

To investigate the effect of the content or “quality” of messages, we classify them as follows. We call a message “related” if it directly refers to the tax proposal and “unrelated” otherwise. We find that 90.4 percent of all messages were “related” in phase 1 indicating that messages were not just casual conversation but involved problem-oriented deliberation. We call a related message “correct” if it is compatible with the standard economics prediction or proposes an improvement of the committee decision, and “incorrect” otherwise. Finally, we call a related message “none” if it is neither correct nor incorrect (see Appendix A2 for an example of our classification).

In all committees in DELIB, at least one correct message is sent. Therefore, if there is something like a “heureka effect”, according to which some voters are inattentive but recognize the optimal action if told, we would expect all committees in DELIB to vote in line with standard economics. As shown later in more detail, this is not the case. In fact, in all

committees also at least one incorrect message is sent which means that all voters are exposed to some correct and some incorrect messages in varying proportions. In fact, the average shares of correct and incorrect messages per committee are not significantly different according to a Wilcoxon signed-rank test (22.3 vs. 32.5 percent, $p = 0.308$).

Deliberation and persuasion

How is voting affected by deliberation? We proceed in two steps to answer this question. First, we investigate how what a voter says is affected by what other say. Second, we analyze how the votes depend on what other say.

How are the messages a voter sends affected by the messages he receives? To investigate, we run a probit regression where the quality of the last message voter i sends is regressed on the quality of the first message i sends and the quality of the messages i receives. The first coefficient in table 1 shows that those who send correct messages right away are more likely to end up sending correct messages than those who get it wrong in the beginning of the deliberation stage. In particular, a voter is 39 percent more likely to send a correct message at the end of the deliberation stage if his first message was already correct than if it was incorrect. The next two coefficients show that incorrect messages received from others significantly deteriorate the quality of messages sent at the end of the discussion stage, but correct messages do not.

Table 1: Quality of messages sent by quality of messages received

Dependent variable is 1 if last message of voter i is correct, and 0 else	
First message correct (dummy)	0.385*** (0.126)
Incorrect messages from others	-0.042** (0.019)
Correct messages from others	0.058 (0.048)
Observations	52
Pseudo R-squared	0.239
Wald $\chi^2(3)$	11.05**

Notes: Probit regression, coefficients show marginal effects. Robust standard errors (clustered on independent groups) in parentheses.

* significant at 10%; ** significant at 5%; *** significant at 1%

How does the quality of messages correlate with the quality of committee decisions? We find that these variables are strongly related. The more correct messages are exchanged in a committee, the better the committee decision. In fact, the larger the number of correct messages is, the lower the committee’s tax vote (Spearman's rho = -0.673, $p = 0.016$), and vice versa for incorrect messages (Spearman's rho = 0.856, $p = 0.000$). While these correlations are impressively strong, they need to be interpreted with care. On the one hand, correct messages may tend to reduce, and incorrect messages tend to enforce tax salience through persuasion. On the other hand, voters who are prone to salience effects may also be prone to sending more incorrect messages and vice versa. In this case, the correlation should not be interpreted as a causal effect of deliberation on tax choices but as a selection effect.

As we show next, both types of effects are present. We find that “persistent senders” vote in line with what they say, and that “fickle senders”, i.e. those who change their expressed views vote in line with what other voters say. To demonstrate that this is the case, we group persistent senders into “correct senders” (12/60 subjects send at least one correct and no incorrect message), and analogously defined “incorrect senders” (15/60 subjects), respectively. 87 percent of incorrect senders vote for taxes above 25, but only 17 percent of the correct senders do. In contrast, fickle senders’ votes are affected by what consistent senders say.

Table 2: Voting for the taxes above 25 by number of consistent senders

		# of correct senders		<i>sum</i>
		0	≥ 1	
# of incorrect senders	0	5/5	1/5	6/10
	≥ 1	13/15	5/8	18/23
<i>sum</i>		18/20	6/13	24/33

Notes: The first number is the number of fickle voters voting for taxes above 25 in DELIB and the second number is the number of voters in a category. For example, 5/5 in the upper left cells means that 5 fickle voters are in committees without correct or incorrect senders and that all vote for taxes above 25.

Table 2 shows how voting of fickle senders depends on the number of correct and incorrect senders. A comparison of the sums in the columns of the table measures the effect of having at least one correct sender in a group. Correct senders improve voting of fickle senders. While 90 percent (= 18/20) of the fickle senders in a committee without correct senders vote for taxes above 25, only 54 percent (= 6/13) do so in the presence of at least one

correct sender. The comparison of the rows in table 2 shows the effect of having incorrect senders in the group. We find that incorrect senders tend to deteriorate voting of fickle senders (60% absent, 78% in the presence of at least one incorrect sender).

Table 2 therefore suggests that fickle senders are persuaded by messages from both correct and incorrect senders. These results are confirmed by a Probit regression using all subjects with voting for taxes above 25 as the dependent variable. We find that (controlling for the number of messages) a one-percent increase in the share of correct messages sent by others within the committee reduces the probability to vote for taxes above 25 by 0.29 percent ($p = 0.006$) while a one-percent increase in the share of incorrect messages increases this probability by 0.15 percent ($p = 0.025$).

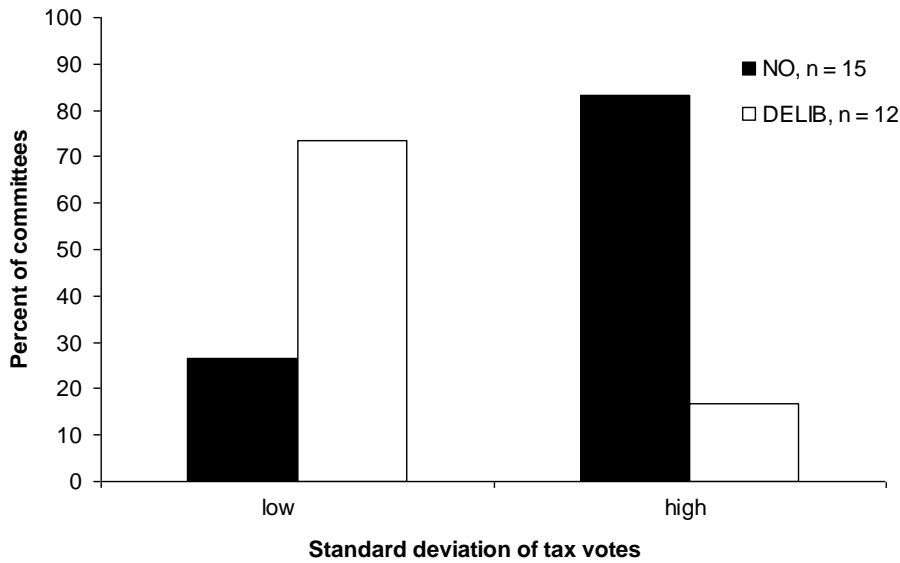
Deliberation and group polarization

The analysis above suggests that messages are persuasive, i.e. that voting is systematically affected by the quality of messages received. We show next that deliberation induces voters to agree – but not necessarily on the right solution. Instead, deliberation tends to make initially held opinions more homogenous and more extreme – a process called group polarization (e.g. Isenberg 1986). Committees that happen to exchange many incorrect messages tend to vote for taxes above 25 and vice versa for correct messages. As a result, we find that the option of pre-vote deliberation does not mitigate the illusion among inexperienced voters.

Figure 4 shows that deliberation induced committees to agree. Agreement is measured by the within-committee standard deviation of individual tax votes. We group committees into low and high agreement depending on whether a committee has a standard deviation above or below the median committee in NO and DELIB. The figure shows that 83 percent of committees have a high level of agreement in DELIB, but only 27 percent of the committees in NO do (χ^2 -test, $p = 0.005$).

The average within-committee standard deviation is almost twice as large in NO than in DELIB (14.8 vs. 7.9, $p = 0.007$, two-sided MW). While deliberation induces voters to agree, they do not systematically agree on the more salient but less costly tax. In fact, the correlation between committee votes and within-committee standard deviation is insignificant (Spearman's rho = -0.318, $p = 0.313$).

Figure 4: Deliberation induces agreement
(within-committee standard deviation of tax votes, phase 1)



While deliberation reduces the within-committee standard deviation of taxes votes, it increases its across-committee standard deviation. It is more than twice as large in DELIB (11.8) than in NO (5.8). If we consider a session an independent observation (we had 3 sessions per treatment), we find that across-committee standard deviations were significantly higher in DELIB than in NO ($p = 0.063$, one-sided Wilcoxon rank-sum test). Thus, deliberation seems to induce a “group polarization” effect.

Figure 5: Distribution of tax votes with and without deliberation (Phase 1)

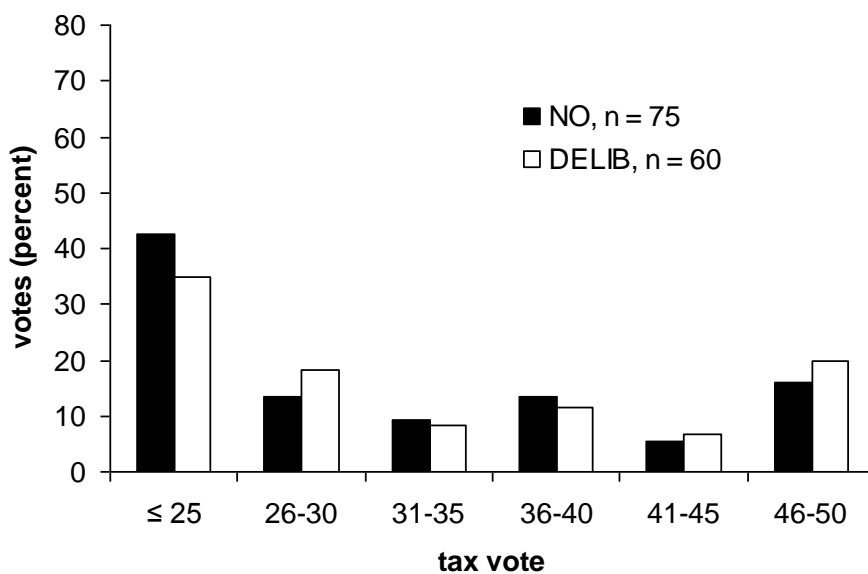


Figure 5 shows that pre-vote deliberation did not affect the distribution of individual tax votes. In DELIB, 65 percent of tax votes are above 25, and the average tax vote is 34.1 points. In NO, 57 percent of tax votes are above 25 and the average tax vote is 31.5. Neither the average tax vote ($p = 0.696$, MW) nor the distribution of tax votes above 25 is significantly different in DELIB and NO ($p = 0.839$, KS).

How are committee votes affected by the option to deliberate? The short answer is that they are not because the individual-level effects of deliberation tend to cancel out. In NO, two thirds of all markets (10 out of 15) have median tax votes above 25, and the median tax vote averaged across all markets is 31.5 points. The results for DELIB are very similar. In DELIB, again two thirds of markets (8 out of 12) have a median tax vote above 25, and the median tax vote averaged across all markets is 33.3 points ($p = 0.709$, MW).

3.3 Learning and tax salience

We now discuss how deliberation and learning from experience interact to reduce tax salience. Uninformative experience prevails if the market was not subject to a tax above 25 in a previous phase. In this case, voters do not learn the income-reducing effects of voting for high taxes but voting may still improve due to repetition and reconsideration. *Informative* experience prevails if the committee vote was for a tax above 25 and the committee experienced a tax above 25 in the market stage in a previous phase. In this case, voters experience losses in the market stage.

Our results show that informative experience massively reduces salience effects. We show that deliberation *per se* does not reduce tax salience even with repetition, but induces voting for the more salient but less costly buyer tax in the presence of easy-to-interpret informative experience.

Two facts set the stage for our analysis of learning and deliberation. First, deliberation can in principle affect outcomes after phase 1 because participants continue to actively discuss and because opinions are still divided after the first referendum. To illustrate the fact that participants continue to deliberate after phase 1, we note that the average number of messages a subject sends is not significantly different in phase 5 from phase 1 (5.7 vs. 6.7, $p = 0.530$, Wilcoxon sign-rank test).

As in phase 1, the quality of messages exchanged correlates well with the quality of committee decisions in the remaining phases. We find that the more correct messages are

exchanged, the better the committee choice averaged over phases 2 to 5 (Spearman's rho = -0.724, $p = 0.007$), while the reverse is true for incorrect messages (Spearman's rho = 0.718, $p = 0.008$).

Table 3 summarizes our main results with respect to tax salience, learning and deliberation. Regression (1) shows, first, that voters in treatment CON are significantly less likely (by 22.8 percent) to choose a seller tax above 25 than voters in treatment NO. This estimate suggests that the pronounced effect of tax salience discussed in figure 3 for inexperienced voters is also present in later phases. Second, deliberation does not improve voting over all phases. This result is in line with our finding for inexperienced voters (see figure 5). Third, there is considerable persistence in voting. A voter who voted for a tax above 25 in one phase is 53.2 percent more likely to choose a tax above 25 in the consecutive phase. Fourth, voters who experienced a tax above 25 in the market are much less likely (by 38.0 percent) to vote for taxes above 25 in later phases, indicating that voters learn to overcome salience effects in the presence of clear and immediate payoff feedback. Fifth, the variable Trend indicates that mere repetition does not significantly improve voting.

Table 3: Voting, deliberation and experience

Dependent variable is 1 if voter i votes for tax above 25, and 0 else		
	(1)	(2)
CON	-0.228*** (0.081)	-0.192** (0.078)
Deliberation (DELIB)	0.008 (0.081)	0.200 (0.132)
Vote _{i} previous phase > 25 (Vote _{i})	0.532*** (0.032)	0.535*** (0.031)
Informative experience (Info Exp)	-0.380*** (0.080)	-0.317*** (0.091)
Deliberation \times Info Exp		-0.248* (0.141)
Trend (Phase)	-0.012 (0.031)	-0.012 (0.029)
Observations	820	820
Pseudo R-squared	0.24	0.25
Wald $\chi^2(5,6)$	239.71	285.13

Notes: Probit regression, using all treatments and all phases. Coefficients show marginal effects. Robust standard errors (clustered on independent groups) in parentheses.

* significant at 10%; ** significant at 5%; *** significant at 1%

In regression (2), we add a term showing that deliberation significantly interacts with informative experience. Deliberation reduces the probability to vote for taxes above 25 by

24.8 percent if a voter made an informative experience in a previous phase. Thus, deliberation has no effect absent informative experience, but has a beneficial effect in the presence of informative experience.

Overall, our conclusions from table 2 are well in line with the discussion in sections 3.1. and 3.2, and the results seem robust.¹⁴ However, the regression results should be interpreted with due care. For example, selection effects are likely to occur because committees making particularly bad choices in one phase are more likely to learn the adverse effects of high taxes than those voting for moderate taxes. The next section provides some additional details on the learning process which, we believe, support our interpretation of the main effects given above.

Details on deliberation and learning

The first two coefficients in table 3 suggest that tax salience is strong (there is a pronounced difference between NO and CON) and that deliberation makes no difference for voting behavior (no difference between NO and DELIB) throughout all phases.

To illustrate the second point, we note that voting is not significantly different with and without deliberation both at the individual and the committee level. In particular, the average vote over all phases is 33.1 in DELIB and 32.8 in NO ($p = 0.941$, MW). The average committee vote is 30.8 points in NO and 32.0 points in DELIB ($p = 0.526$, MW). In addition, we find that convergence to the standard economics prediction is slow and similar in both NO and DELIB. For example, the share of individual votes for taxes above 25 falls in NO from 57 to 40 percent and DELIB from 65 to 42 percent from phase 1 to 5. The share of groups with committee votes above 25 falls from 67 to 52 percent in NO and from 67 to 42 percent in DELIB. In both treatments, the share of markets experiencing the seller tax regime falls by about a third (from 25 to 13 percent in NO and from 25 to 17 percent in DELIB).

To provide more details on the interpretation of the first estimate, and given the discussion in the previous paragraph showing that NO and DELIB have similar patterns, we merge NO and DELIB and compare it to CON. Recall that according to the salience hypothesis, NO and DELIB induce voting for high taxes but CON does not. Over all phases, NO and DELIB have significantly higher committee votes (NO and DELIB: 31.3, CON: 26.7, $p = 0.001$, MW), higher tax revenues (NO and DELIB: 127.5, CON: 121.3, $p = 0.041$, MW), and lower efficiency (NO and DELIB: 94.4%, CON: 95.8%, $p = 0.099$, MW) than CON.

14 Results are robust to alternative model specifications. For example, we have estimated a Poisson-model for count data, and a double-censored (at 25 and 50) Tobit-model. We find no difference for the interaction effects when considering the critique in Ai and Norton (2003).

The third estimate in table 3 suggests that salience effects are persistent absent informative experience. For example, only one third of subjects (= 10/31) who voted for seller taxes above 25 in one phase, but do not experience the adverse effects of a tax above 25 in phases 2-5 reduce their tax votes in the following phase. In other words, mere repetition and the opportunity to reconsider did not reduce tax salience.

The fourth estimate suggests that informative experience is powerful in de-biasing voters. Almost all committees make an informative experience at some point during the experiment. In particular, only one of the 27 committees in NO and DELIB never experiences taxes above 25. However, most markets experience a tax above 25 only once (67 percent in both NO and DELIB), indicating that they managed to avoid the inefficient tax after having experienced its adverse effects. How do committees react to informative experience? In phases 2-5, there are 33 instances of informative experience in NO and DELIB. In 94 percent (= 31/33) of these cases, the committee vote responds in the right direction. This finding not only shows that voters managed to respond rationally in the face of easy-to-interpret evidence but also that our experiment provided sufficient incentives for learning given informative feedback.

Finally, we find evidence that experience affects the quality of messages in DELIB. Without informative experience, the share of correct messages falls and the share of incorrect messages tends to increase, while with informative experience, the share of correct messages increases and the share of incorrect messages falls.¹⁵ Thus, deliberation in the presence of informative experience improves the quality of messages and votes (see interaction term in table 2). In later phases of the experiment committee polarization effects cede to exist. This is not surprising as subjects become increasingly experienced. For instance, in phase 5, average within-committee standard deviation is 4.5 points in NO and 6.1 points in DELIB ($p = 0.237$).

15 For example, a comparison of the first and the second vote reveals significant effects in a difference-in-difference test. Without informative experience, the share of correct messages falls by 18.2 percentage points, but increases by 20.8 percentage points with informative experience ($p = 0.021$, MW). For the share of incorrect messages, the differences are -1.0 and -43.6 percentage points ($p = 0.013$, MW). However, experience (whether informative or not) does not affect the absolute number of messages sent.

4 Concluding remarks

A striking result of our experiment is that deliberation does not prevent tax salience effects. Deliberation among inexperienced subjects causes voters to agree on arbitrary, not necessarily correct, values. Apparently, the smart voters are not able to convince sufficiently many other voters in the committee and the voters sending misleading messages are just as convincing as those sending correct messages. However, experience is a powerful argument to convince others. If subjects have experienced the income-reducing effects of voting for the less salient but more costly tax, deliberation improves the committee choices.

Our result that learning mitigates tax salience is perhaps less surprising in the light of a voluminous literature showing that cognitive biases can be considerably reduced by providing subjects with clear and immediate payoff feedback (e.g. Friedman 1998). However, in interpreting our result, one should keep in mind that learning is straightforward in our experiment while feedback is unlikely to be as direct and informative in natural environments.

In our main treatments, the less salient tax was inefficient because it was higher by design than the more salient buyer tax. The reason for this design choice is that we wanted to create a tension between salience and efficiency. This tension makes the demonstration of salience effects more convincing because voters have an incentive to overcome salience effects. However, the efficiency of non-salient taxes is debated. From the perspective of optimal taxation theory, non-salient taxes might be more efficient. The idea is that if a given tax revenue has to be raised, it should be raised in the least distorting manner, and that would be the case when individuals under-react to a change in taxation, i.e. with non-salient taxes (e.g. Chetty et al. 2007). On the other hand, the use of non-salient taxes could allow budget-maximizing politicians and bureaucrats to increase tax revenues beyond what citizens deem appropriate. Some authors have suggested that policy makers may deliberately choose non-salient taxes to increase tax revenues (e.g. Slemrod and Krishna 2003, Finkelstein 2007).

We believe that tax salience is an important determinant of individual reactions to existing tax incentives in the market place and we have shown that the acceptance of tax regimes in the political arena can depend on tax salience. We therefore hope our contribution motivates further investigations into the emerging field of behavioral public economics.

References

- Ai, C. and Norton, E.C. (2003): Interaction Terms in Logit and Probit Models. *Economics Letters* 80(1): 123-9.
- Besley, T. and Rosen, H.S. (1999): Sales Tax and Prices: An Empirical Analysis. *National Tax Journal* 52(2): 157-78.
- Blinder, A.S. (1988): The Challenge of High Unemployment. *American Economic Review* 78, 1-15.
- Blumkin, T., Ruffle, B.J. and Ganun, Y. (2008): Are Income and Consumption Taxes Ever Really Equivalent? Evidence from a Real-Effort Experiment. Working paper 6479, University of Munich.
- Borck, R., Engelmann, D. Müller, W. and Normann, H.-T. (2002): Tax Liability Side Equivalence in an Experimental Posted Offer Market. *Southern Economic Journal* 68(3): 672-82.
- Brandts, J. and Cooper, D.J. (2007): It's What You Say, Not What You Pay: An Experimental Study of Manager-Employee Relationships in Overcoming Coordination Failure. *Journal of the European Economic Association* 5(6): 1223-68.
- Carbonnier, C. (2007): Who Pays Sales Taxes? Evidence from French VAT Reforms, 1987-99. *Journal of Public Economics* 91(6): 1219-29.
- Charness, G. and Dufwenberg, M. (2006): Promises and Partnerships. *Econometrica* 74: 1579-601.
- Chetty, R., Looney, A. and Kroft, K. (2007): Salience and Taxation: Theory and Evidence. NBER working paper no. 13330, August 2007.
- DellaVigna, S. (2007): Psychology and Economics: Evidence from the Field. Working paper UC Berkeley.
- Dollery, B.E. and Worthington, A.C. (2006): The Empirical Analysis of Fiscal illusion. *Journal of Economic Surveys* 10(3): 261-97.
- Ellingsen, T. and Johannesson, M. (2004): Promises, Threats and Fairness. *Economic Journal* 114: 397-420.
- Feldman, N.E. and Katuscak, P. (2007): Effects of Predictable Tax Liability Variation on Household Labor Income. Working paper Ben Gurion University.
- Finkelstein, A. (2007): E-Z tax. Tax Salience and Tax Rates. Working paper MIT, December 2007.
- Fischbacher, U. (2007), z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics* 10(2): 171-8.
- Friedman, D. (1998): Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly. *American Economic Review* 88(4): 933-46.
- Fullerton, D. and Metcalf, G.E. (2002): Tax Incidence. In: A.J. Auerbach and M. Feldstein (Eds.): *Handbook of Public Economics* Vol. 4. Amsterdam: Elsevier: 1787-872.
- Gabaix, X. and Laibson, D. (2006): Shrouded Attributes and Information Suppression in Competitive Markets. *Quarterly Journal of Economics* 121: 505-40.

- Gerardi, D. and Yariv, L. (2007): Deliberative Voting. *Journal of Economic Theory* 134: 317-38.
- Glaeser, E.L. and Sunstein, C.R. (2007): Extremism and Social Learning. NBER working paper 13687, December 2007.
- Graetz, M.J. (2002): 100 Million Unnecessary Returns: A Fresh Start for the U.S. Tax System. *Yale Law Journal* 112: 261-310.
- Hossain, T. and Morgan, J. (2006): "...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *Advances in Economic Analysis and Policy* 6(2): Article 3.
- Huckfeldt, R. and Sprague, J. (1987): Networks in Context: The Social Flow of Political Information. *American Political Science Review* 81(4): 1197-216.
- Isenberg, D.J. (1986): Group Polarization. *Journal of Personality and Social Psychology* 50(6): 1141-51.
- Kachelmeier, S., Limberg, S. and Schadewald, M. (1994): Experimental Evidence of Market Reactions to New Consumption Taxes. *Contemporary Accounting Research* 10(2): 505-46.
- Kerschbamer, R. and Kirchsteiger, G. (2000): Theoretically Robust but Empirically Invalid? An Experimental Investigation into Tax Equivalence. *Economic Theory* 16(3): 719-34.
- Kotlikoff, L.J. and Rapson, D. (2006): Does it Pay, at the Margin, to Work and Save? – Measuring Effective Marginal Taxes on Americans' Labor Supply and Saving. NBER working paper 12533.
- Liebman, J.B. and Zeckhauser, R.J. (2004): Schmeduling. WP Harvard University.
- McCaffery, E.J. and Baron, J. (2006): Thinking about Tax. *Psychology, Public Policy, and Law* 12(1): 106–35.
- Riedl, A. and Tyran J.-R. (2005): Tax Liability Side Equivalence in Gift-Exchange Labor Markets. *Journal of Public Economics* 89(11): 2369-82.
- Ruffle, B. (2004): Tax and Subsidy Incidence Equivalence Theories: Experimental Evidence from Competitive Markets. *Journal of Public Economics* 89(8), 1519-42.
- Sausgruber, R. and Tyran, J.-R. (2005): Testing the Mill Hypothesis of Fiscal Illusion. *Public Choice* 122: 39-68.
- Shafir, E., Diamond, P.A. and Tversky, A. (1997): Money Illusion. *Quarterly Journal of Economics* 112(2): 341-74.
- Slemrod, J. (2006): The Role of Policy Misconceptions in Support for Regressive Tax Reform, *National Tax Journal* 59(1): 57-75.
- Slemrod, J. and Krishna, A. (2003): Behavioral Public Finance: Tax Design as Price Presentation. *International Tax and Public Finance* 10(2): 189-203.
- Smith, V.L., Williams, A.W. Bratton, W.K. and Vannoni, M.G. (1982): Competitive Market Institutions: Double Auction versus Sealed Bid-Offer Auctions. *American Economic Review* 72(3): 58-77.
- Tyran, J.-R. (2007): Money Illusion and the Market. *Science* 314: 1042-3.

Appendices

Appendix A1: Proof of the incentive-compatibility of the voting mechanism

Assume that a buyer maximizes expected utility and attaches positive probability to being decisive. A transaction tax τ_i with $i = B, S$ shall be imposed. The index B indicates that the tax is levied on the buyers, index S denotes that the tax is levied on the sellers. The rate of the buyer tax is a constant, $\tau_B = \bar{\tau}_B$. The rate of the seller tax is variable, $\tau_S \in [\bar{\tau}_B, \tau_{MAX}]$. Let $U_B(\tau_i)$ be the buyer's net benefit under tax τ_i . Call $\bar{\tau}_S$ the rate of the seller tax that yields the same net benefit for the buyer as the buyer tax at rate $\bar{\tau}_B$. Assume that

$$\begin{aligned} U_B(\tau_S) &> U_B(\bar{\tau}_B) \quad \forall \tau_S \in [\bar{\tau}_B, \bar{\tau}_S[, \\ U_B(\tau_S) &= U_B(\bar{\tau}_B) \quad \forall \tau_S = \bar{\tau}_S, \text{ and} \\ U_B(\tau_S) &< U_B(\bar{\tau}_B) \quad \forall \tau_S \in]\bar{\tau}_S, \tau_{MAX}]. \end{aligned}$$

The tax is determined by the following mechanism: The buyer announces a rate for the seller tax τ'_S . A uniformly distributed random variable z is drawn from the interval $[\bar{\tau}_B, \tau_{MAX}]$. The implemented tax is

$$\begin{aligned} \tau_B &\quad \forall z > \tau'_S \\ \tau'_S &= z \quad \forall z \leq \tau'_S \end{aligned}$$

Claim: it is optimal for the buyer to announce $\tau'_S = \bar{\tau}_S$.

Proof: With every choice of τ'_S the buyer determines a lottery with two prizes. One prize is $U_B(\bar{\tau}_B)$, the other $U_B(\tau_S^Z)$. The lotteries fall into two classes. The first is with $\tau'_S \leq \bar{\tau}_S$ and has $U_B(\tau_S^Z) \geq U_B(\bar{\tau}_B)$. As long as $\tau'_S \leq \bar{\tau}_S$, the probability of getting the higher (or equal) prize $U_B(\tau_S^Z)$ is monotonically increasing in τ'_S . Therefore, the lottery with $\tau'_S = \bar{\tau}_S$ first order stochastically dominates all other lotteries in this class. The second class of lotteries is with $\tau'_S \geq \bar{\tau}_S$ and has $U_B(\bar{\tau}_B) \geq U_B(\tau_S^Z)$. As long as $\tau'_S \geq \bar{\tau}_S$, the probability of getting higher (or equal) prize $U_B(\bar{\tau}_B)$ is monotonically decreasing in τ'_S . Therefore, the lottery with $\tau'_S = \bar{\tau}_S$ first order stochastically dominates all other lotteries in this class. QED.

Appendix A2: Example of chat and classification

Subject and Message	Related	Code
Subject 2: <i>So?</i>	0	--
Subject 5: <i>What do you propose?</i>	1	None
Subject 3: <i>I propose to choose a high tax since we are the buyers.</i>	1	Incorrect
Subject 1: <i>Well, it depends on how the price is determined.</i>	1	None
Subject 2: <i>I suppose the tax will add to the price.</i>	1	Correct
Subject 5: <i>In any case, we shall all choose the same tax.</i>	1	None
Subject 1: <i>If we pay the tax anyway, why then not chose a low one?</i>	1	Correct
Subject 4: <i>All below 25?</i>	1	Correct
Subject 2: <i>I am for it.</i>	1	None
Subject 3: <i>Me too.</i>	1	None
....		

Appendix A3: Instructions for the auction (translated from German.)

General Instructions for Participants

You are now participating in an economics experiment which funded by the Austrian Science Fund. The purpose of the experiment is to analyze decision making in markets. You are paid Euro 4 for showing up on time. If you carefully read the instructions and follow the rules you can earn additional money. The €4 and all additional amounts earned during the experiment will be paid to you in cash immediately after the experiment. You can earn points in the experiment. These points will be converted to Euros according to the following exchange rate: Points 100 = €0.8. During the experiment we ask you do not to speak to other participants. If you have a question, please ask us. We will gladly answer your questions in private. It is very important that you follow this rule. Otherwise the results of the experiment have no value from a scientific perspective.

You are now participating in a market experiment. In the market, you can buy units of a hypothetical commodity. You earn money by trading. How much you earn depends on your and the decisions of others. The experiment consists of two practice periods followed by a number of trading periods. In the practice periods you do not earn money; but you should take these periods seriously since you will gain valuable experience for the paid trading periods.

Detailed Instructions for Buyers

In this experiment each participant is a buyer. You can buy units from automated sellers. The sellers will sell to you according to the rules described below. In your market there are 5 buyers who can buy units from sellers in each of the trading periods.

What participants can do: As a buyer you state at which price you would buy a unit. We call this your “bid”. You can buy two units at most. You can submit separate a bid for each of these two units.

How the market works: At the end of each trading period, the bids you and other buyers in your market have submitted are collected and ranked from high to low. The highest bid is ranked above the 2nd highest bid. The 2nd highest bid is ranked above the 3rd highest bid, and so on. If two or more bids are equal, ranks will be randomly assigned by the computer.

Sellers bear costs for each unit they sell. These costs are ranked from low to high. This means that the lowest cost is ranked above the 2nd lowest cost. The 2nd lowest cost is ranked above the 3rd lowest cost, and so on. Buyers buy units with bids above costs in this order. Bids that are below costs for a unit are rejected.

Example: Assume we collect four bids in a market period. The highest bid is 145, the 2nd highest bid is 130, the 3rd highest bid is 110, and the 4th highest bid is 90. The lowest cost is 60, the 2nd lowest cost is 80, the 3rd lowest cost is 95, and the 4th lowest cost is 105.

	1.	2.	3.	4.
Bids	145	130	110*	90
Costs	60	80	95	105

In the example the first 3 bids are above the costs of the sellers. Therefore, the 3 units sell to the buyers who have made these bids. The 4th bid at 90 is rejected. The buyer who submitted this bid does not buy the unit.

Please note: Buyers do not know the bids of others, nor do they know the costs of units for sellers.

How your earnings are computed: Each unit has a value for you. We call this “your value”. You will learn your value in the experiment. You only know your value; you do not know the values of other buyers nor do other buyers know your value. If you buy a unit you have to pay a price. Each buyer pays the same price per unit. This uniform market price is equal to the last accepted bid in the order explained above. In our example, 3 units have been bought. The bid of the last accepted unit is 110 (market with *). The uniform market price is therefore 110. Please note once again that the first two units of the order are also bought at a price of 110.

Your earnings per unit bought are calculated as follows: Earning = value minus price. Note that if you buy a unit you will pay less than your bid unless your bid is exactly equal to the market price. In our example, suppose that you submitted the bid of 130 for a unit which has a value of 155. This bid is above the market price of 110. Since you buy this unit at a market price of 110 your earnings will be $155 - 110 = 45$.

If you do not buy a unit, its value expires. On the other hand, you also do not pay a price. Your earnings are therefore zero.

Subjects’ original instructions contained figures showing how the computer screen would look like during the experiment

How is trading presented on the computer screen? In each trading period a Decision Screen appears (Figure 1). At the end of each period an Outcome Screen appears (Figure 2). After 15 trading periods a History of Results appears (Figure 3). The numbers in the figures of the instructions serve illustrative purposes only. Actual numbers may be different.

In the top area of the Decision Screen you see on the left the number of the current trading period (here: 2) and the total amount of trading periods (here: 15). Each trading period ends after a time limit. The remaining time within a period is shown in the top area on the right (here: 19 Seconds). When the experiment starts the time available for trading is generous but it will be continuously reduced in later periods.

In the first row you see your value for your first unit. In this example the buyer has a value of 155. The input field below your value serves to enter your bid. To enter a bid, click on the field labeled ‘Your Bid’ and type in a number. To submit that bid, click on the ‘Submit’ button. In the second row everything is repeated for your 2nd unit. Here, the buyer values the 2nd unit at 105 points.

Rules for bidding: An important rule for trading is to “trade at no loss”. Therefore, you may not submit bids above your values. In the example shown in Figure 1, this buyer’s bid must not be above 155 for the 1st and

105 for the 2nd unit. A second rule is that your bid for the 1st unit may not be above the bid of the 2nd unit. If you violate these rules, a message box appears. The message disappears if you press the 'OK' button.

The Outcome Screen (Figure 2) appears at the end of a trading period. The upper part of this screen looks the same as the Decision Screen. In the table below you find your value, your bid, the price, for each of the two units. In the last row of the table shows your earnings for the current period. In this example, the earnings from buying the 1st unit are 45. This is the difference between the value of this unit and the market price. The earnings from the 2nd unit are zero: the bid of 90 has been rejected because it was below the cost of the unit.

Below the table you see the total number of units traded in the market (Market Quantity). The next line shows how many units you have bought.

Finally you see Your Period Earnings. This is the sum of your earnings from buying units in the market.

After a total of 10 trading periods the History of Results (Figure 3) appears. Here, the results from each period are summarized. "Your Total Earnings over the last 10 periods" shows the sum of your earnings as sum over the last 10 periods.

Instructions on the tax proposal [*the wording in parenthesis has been used in the control treatment, CON*]

You and the 4 other buyers in your market now have to decide between two parties A and B. A tax at a rate of at least 25 points must be imposed.

Party A proposes that buyers [*sellers*] in each of the following 10 periods pay a tax equal to 25 ($t_A=25$) points for each unit they buy [*sell*].

Party B proposes that sellers [*buyers*] in each of the following 10 periods pay a tax greater than 25 ($t_B>25$) points for each unit they sell [*buy*].

The tax rate t_B on sellers [t_A on buyers]

At the time of your choice the tax rate t_B proposed by party B [*t_A proposed by party A*] is yet not known. This tax will be announced only after you have decided between the two parties A and B. The tax per unit sold [*bought*] is an integer number randomly determined between 25 and 50. All numbers are equally likely.

Decision rule

You and the 4 other buyers in your market propose a tax rate per unit. The tax rates proposed by all 5 buyers are ranked from low to high. The 3rd number in this order is called the median.

- If the median is smaller than the tax rate proposed by party B ($\text{median} < t_B$), party A wins and the tax is $t_A=25$ points: Buyers then pay a tax of 25 points on every unit they buy for the next 10 periods. [*If the median is smaller than the tax rate proposed by party A ($\text{median} < t_A$), party B wins and the tax is $t_B=25$ points: Sellers then pay a tax of 25 points on every unit they sell for the next 10 periods.*]
- If the median is greater or equal than the tax rate proposed by party B ($\text{median} \geq t_B$), party B wins and the tax is t_B points: Sellers then pay a tax of t_B points on every unit they sell for the next 10 periods. [*If the median is greater or equal than the tax rate proposed by party A ($\text{median} \geq t_A$), party A wins and the tax is t_A points: Buyers then pay a tax of t_A points on every unit they buy for the next 10 periods.*]

Please note: If you propose a tax smaller or equal than 25 you decide in favor of party A (= against party B) [*in favor of party B (= against party A)*]. The higher the median of the values proposed by you and the other 4 buyers, the more probable it is that party B [*party A*] wins.

Example 1: You and the other 4 buyers in your market propose the tax rates 35, 22, 39, 26, 37. After respective ranking, we get:

- 1: 22
- 2: 26
- 3: 35 (= median)
- 4: 37
- 5: 39

Suppose that the random tax of party B [*party A*] is determined as $t_B=32$ [$t_A=32$]. The median (35) is above $t_B=32$ [$t_A=32$]. Therefore, party B [*party A*] wins and sellers pay a tax of $t_B=32$ [$t_A=32$] points for each unit they sell.

Example 2: You and the other 4 buyers in your market propose the tax rates 35, 3, 17, 30, 37. After respective ranking, we get:

- 1: 3
- 2: 17
- 3: 30 (= median)
- 4: 35
- 5: 37

Suppose that the random tax of party B [*party A*] is determined at $t_B=38$ [$t_A=38$]. The median (30) is below $t_B=38$ [$t_A=38$]. Therefore, party A [*party B*] wins and buyers pay a tax of $t_A=25$ [$t_B=25$] points for each unit they buy [*sell*].

In the treatment with deliberation the following paragraph was added.

Discussion via chat

You now have the opportunity to discuss the above proposal with the other 4 buyers in your market. Discussion will be exclusively in written form via a chat on your computer. You have 5 minutes to discuss. Please, do not discuss anything apart from the proposal. For example, you may try to convince other buyers in your market about the tax rate they should choose. Please do not reveal your true identity and do not threaten or insult others. Otherwise the results of this experiment are of no value from a scientific perspective.