Journal of Economic Behavior & Organization 124 (2016) 43–54

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

The articulation of government policy: Health insurance mandates versus taxes☆

Keith Marzilli Ericson[a,b,*], Judd B. Kessler[c,d]

[a] Boston University School of Management, United States
[b] National Bureau of Economic Research, United States
[c] University of Pennsylvania, United States
[d] Leonard Davis Institute Center for Health Incentives and Behavioral Economics, United States

ARTICLE INFO

Article history:
Received 24 March 2014
Received in revised form
24 September 2015
Accepted 27 September 2015
Available online 9 October 2015

Keywords:
Mandate
Tax
Framing
Articulation
Social norms
Health insurance

ABSTRACT

Can the articulation of government policy affect behavior? Participants in our experiment report their probability of purchasing health insurance under one of two financially equivalent policies: a government mandate to purchase insurance or a tax on the uninsured. During our one-year study frame, controversy arose over the Affordable Care Act's individual mandate. Pre-controversy, the mandate articulation increased purchase by 10.2 percentage points relative to the tax articulation (equivalent to a $1000 decrease in premiums). Post-controversy, the mandate was no more effective than the tax. We show that articulation affects behavior and should be considered when evaluating the efficacy of policy.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Governments aim to discourage or encourage certain behaviors of their citizens, and they employ a variety of policy levers to achieve this goal. Economists analyzing these policies typically assume that the efficacy of a policy is driven exclusively by the extrinsic incentives it creates.[1] For example, decision makers are assumed to care about the tax-inclusive price and about the size of a penalty along with the probability this penalty is imposed (Becker, 1968). This assumption has a natural appeal, since it allows the analysis of government policy to be simplified to easily observable incentives (e.g. financial penalties in the form of fines, taxes, or subsidies; and criminal penalties, such as forced community service or incarceration).

However, the way a policy is articulated to the public might also impact how individuals respond to it. Compliance with a given policy might depend not only on its explicit financial (or criminal) incentives, but also on how it is perceived by agents. For example, if the goal is to discourage a behavior, the government could prohibit the behavior and enforce the prohibition

☆ We thank Lucas Coffman, Andreas Fuster, Gianna Marzilli Ericson, Jim Rebitzer, and David Weil for helpful comments. We acknowledge funding from Boston University internal research funds and The Wharton School at the University of Pennsylvania internal research funds.

* Corresponding author at: Boston University School of Management, United States. Tel.: +1 617 353 4553.
E-mail address: kericson@bu.edu (K.M. Ericson).

[1] A notable exception is Auerbach et al. (2010), which partially motivates this study, as well as the literature on tax salience (e.g. Chetty et al., 2009; Finkelstein, 2009).

http://dx.doi.org/10.1016/j.jebo.2015.09.021
0167-2681/© 2015 Elsevier B.V. All rights reserved.

with a fine or, alternatively, tax the behavior without an explicit prohibition. Traditional models of tax policy would treat the tax and the fine as equivalent so long as the magnitude and probability of their being imposed were equal, but articulating the policy with a prohibition may emphasize a moral obligation or invoke different levels of moral suasion (see Liberman et al., 2004). In addition, the prohibition could carry different information about social norms or directly shape social norms regarding a behavior (Benabou and Tirole, 2011; Elster, 1989). As a result, a policy articulated as a prohibition and a fine might affect compliance differently than one articulated as a tax.

Importantly, how a policy is articulated is determined both by political actors and popular discourse, as each influence how individuals understand a policy's meaning, motivation, authority, and legitimacy. Policy makers act first by choosing, for example, whether a real-estate developer faces identical financial incentives articulated as either: a "requirement" to include affordable housing or else pay a *fine*, a "requirement" to include affordable housing with the option to pay a *fee* instead, or a "suggestion" to provide affordable housing paired with a discount off of higher base fees. Similarly, policy makers decide whether late tax filers are assessed a "penalty" from the IRS or receive a "loan" from the IRS with a high interest rate; and they decide whether to regulate parking with fees or fines. More broadly, policy makers can allocate the nominal assignment of tax liability to consumers or producers; while standard theory suggests that who you tax is irrelevant for tax incidence (which is simply determined by supply and demand elasticities), assigning a carbon tax to individuals as opposed to firms might affect behavior differently by highlighting the government's goal of dissuading individuals from polluting. After policy makers choose an articulation, popular discourse acts second: statements by political actors and commentators, as well as news media coverage, all influence the perceived meaning and legitimacy of a policy.

We study how the articulation of government policy affects behavior by analyzing the decision to purchase health insurance in the context of two financially equivalent policies: one that taxes individuals who do not purchase insurance, and another that mandates insurance purchase and fines those who do not comply. This setting has a particular appeal for the study of the articulation of government policy. An individual mandate to purchase insurance was a cornerstone of the 2010 Patient Protection and Affordable Care Act (PPACA). Our study takes place before the policy was implemented—during a time of policy uncertainty—and spans a period of active debate about the policy.

At first, government officials attempted to articulate the policy as a *mandate* to purchase insurance, with an associated fine for disobeying the mandate, rather than as a *tax* on remaining uninsured. The stated logic for employing this articulation was that a mandate would affect behavior beyond the fine's financial incentive in a way a tax would not, presumably because the mandate implies an obligation to comply with the law (Elmendorf, 2011; Auerbach et al., 2010).

During the months leading up to the U.S. Supreme Court's ruling on PPACA, political opponents and discussions in the popular press undermined the government's desired policy articulation. Positions articulated in the press suggested that the mandate had no particular moral suasion (i.e. it was unconstitutional). Throughout the paper we denote this period as a "controversy" over the policy, and we document its rise with the frequency of mandate-related articles in the press. While President Obama specifically argued that the mandate was not a tax[2] (Pear, 2010), Justice Robert's decisive opinion in the Supreme Court case on PPACA upheld the mandate as constitutional precisely because it could be re-articulated as a tax.[3]

We run an experiment to investigate whether the articulation of government policy affects behavioral intentions before and after the controversy over the mandate.[4] Before the controversy, individuals were more likely to purchase insurance when the policy was articulated as a mandate with a fine than when it was articulated as a tax. After the controversy, which undermined the legitimacy of the mandate and highlighted its equivalence to a tax, this effect was gone.

To estimate our effects, we asked a sample of U.S. residents to report their probability of purchasing health insurance at two different prices ($3000/year and $2000/year) under one of two government policies: either (1) a mandate with a fine of $700 for not having insurance, or (2) an uninsurance tax of $700 that must be paid by anyone without insurance. The two policies are financially equivalent, since in each case the individual's wealth is reduced by $700 if insurance is not purchased (and by the insurance premium if it is purchased). By randomly assigning subjects to respond to one of the two policy articulations we can identify the *relative effectiveness* of the mandate articulation and the tax articulation. By randomizing different subjects into these articulations at different points in time, we can see if the *relative effectiveness* of the mandate changed over time. Section 2 explains how we identify the relative effect of the two articulations.

Participants' purchase intentions are hypothetical choices. Measuring actual choices is preferable but infeasible, as doing so would require policy makers to randomize how the actual policy is articulated to different individuals.[5] Fortunately, hypothetical choices serve our purpose quite well. Hypothetical choices are regularly used to provide valuable information

[2] President Obama may have had additional political reasons to make this argument beyond the mandate's effect on behavior.

[3] As Roberts wrote in the majority opinion: "it is reasonable to construe what Congress has done as increasing taxes on those who have a certain amount of income, but choose to go without health insurance. . ." (National Federation of Independent Business v. Sebelius 2012).

[4] We investigate behavioral intentions as opposed to public opinion (see Jacoby, 2000), as individuals may support or oppose a policy without it affecting their compliance with that policy.

[5] The ideal test would require variation in how a policy is articulated (i.e. what the government called the policy and how it was discussed in the news) among otherwise identical individuals, but this is infeasible. Direct evidence on the effect of insurance mandates (as opposed to other financially equivalent incentives) is lacking. Indirect evidence is sparse: before the passage of the 2010 federal health reform, Massachusetts was the only U.S. state to have a mandate to purchase individual insurance. Compliance with the Massachusetts mandate was high (Steinbrook, 2008; Gruber, 2011), but it is unknown whether the mandate was more effective than a similarly sized tax or subsidy would have been. Similarly, while some U.S. states have mandates to purchase auto insurance, it is unknown whether these mandates would affect behavior differently if articulated as taxes.

when actual choices are unavailable, most often in valuing environmental attributes such as pollution (Carson and Hanemann, 2005), but also in health insurance (Krueger and Kuziemko, 2013). Evidence shows that reported intentions and hypothetical choices predict actual choices in a variety of contexts (Ajzen, 1991; Camerer and Hogarth, 1999), including in Medicare Part D insurance purchase (Kesternich et al., 2013). Indeed, in our data, higher prices are associated with lower reported probability of purchase, and those who currently have insurance report significantly higher likelihood of purchase. Moreover, while the most common concern with hypothetical choices is inflated willingness-to-pay, our hypothesis and test are about *differences* between government policy articulations, rather than about absolute willingness-to-pay levels.

As noted above, we find a large effect of the mandate articulation in December 2011 and early March 2012, before the controversy and the associated debate about the insurance mandate leading up to the Supreme Court ruling. During this pre-controversy period, calling the policy a mandate rather than a tax increased insurance purchase intentions by 10.2 percentage points, similar in magnitude to the effect of decreasing the annual premium from $3000 to $2000.

We document an increase in news media mentions of the health insurance mandate beginning in late March 2012, with the public debate continuing through the Supreme Court ruling on PPACA and its individual mandate in June 2012. The controversy has a complex effect, including challenging the legitimacy of the mandate and highlighting the financial equivalence of the mandate and tax. Because our study analyzes the effect of government policy articulation before, during, and after the controversy (including before and after the Supreme Court ruling on PPACA), we are able to examine how the relative effectiveness of the mandate changed during this debate. After the controversy surrounding the mandate, the relative effectiveness of the mandate disappeared. These results demonstrate that there is (1) an articulation effect of government policy and (2) that it is malleable. The results suggest that policy makers who want to encourage or discourage a particular behavior have a lever to pull in the way they articulate policy, but that this framing is not entirely under the government's control.

Our finding that the articulation of policy affects behavior is similar in spirit to the hypothesis within Law and Economics that law has an "expressive function" (Sunstein, 1995). The "expressive function" hypothesis suggests that a law can affect behavior by being codified even without being enforced (Funk, 2007). Researchers have argued that the law has an expressive function because it provides information about what others think of a particular behavior (Dharmapala and McAdams, 2003) or because it otherwise alters norms (Sunstein, 1996). Here, we take this logic a step further and argue that even after a law has been codified, how it is articulated to the public can affect behavior. Moreover, we provide an empirical test of this effect.

While our results are related to "framing effects" (see e.g. Johnson et al., 1993), we intentionally use a broader concept and refer to them as *articulation effects*. Typically, framing manipulations present identical information in different ways: for instance, describing probabilistic outcomes with frequencies versus probabilities (e.g. 1 in 25 versus 4%; see Hoffrage et al., 2000), or survival rates versus mortality rates for a surgical procedure (McNeil et al., 1982). We want to distinguish articulation from framing, as various articulations of a given policy may also provide different information about intentions, beliefs, or norms. As noted above, articulations can vary in whether there is a moral component associated with the policy; setting a price does not necessarily attach a moral dimension to a decision, but setting a fine typically implies that (at least someone) thinks the behavior is "wrong."[6]

Finally, our work is distinct from previous research that has found that the salience of a government policy can alter its effects. For instance, Chetty et al. (2009) show that posting sales-tax inclusive prices lead people to respond to the tax-inclusive price, and Finkelstein (2009) shows that drivers who receive a monthly bill for tolls (EZ Pass) have worse recall of the toll amount.[7] In our experiment, the policy and its associated monetary incentives (a fine or tax for not purchasing insurance) are equally salient in both conditions.

2. Articulation and behavior

This section provides a simple model of how policy articulation affects behavior in order to show how our experimental design identifies the *relative effectiveness* of the mandate articulation and the tax articulation of a policy designed to encourage insurance purchase.

In the standard model, individuals maximize utility subject to a budget constraint and do not care how a government policy is articulated. Let p be the price of doing the action favored by the policy (in our example, purchasing insurance). Let f be the financial consequence (e.g. a tax or fine), paid with probability 1, for failing to act as favored by the policy (here, remaining uninsured). Let an individual's wealth be given by w. We write utility as $U(w, i)$ where $i \in \{0, 1\}$ indicates whether the person has purchased insurance and the term i in the utility function represents the standard utility component of having insurance (e.g. as derived from protection against risk).

[6] Articulation can indeed include framing effects. For instance, evidence shows that individuals can be loss averse regarding items they own or expect to receive (Kahneman et al., 1990; Ericson and Fuster, 2011). A policy incentive can be presented as a fee plus a surcharge, or as a higher fee minus a discount. Loss aversion is not the focus of this paper, however, as we examine two different articulations in which the incentives are in the loss domain: fines and taxes.

[7] Similarly, attention may matter as in Lacetera et al. (2012).

In the standard model, an individual will take the desired action (e.g. buy insurance) if:

$$U(w-p, 1) > U(w-f, 0)$$

Since U is strictly increasing in wealth, more people purchase health insurance when $f > 0$ than when $f = 0$. This model does not distinguish between a fine and a tax, as only the financial consequences of failing to purchase insurance enter the utility function.

We augment the utility function above to include an additional moral or normative motive for responding to government policy (we assume for simplicity that this normative component of utility is separable from the rest of the utility function, but this is not essential). This normative component, denoted by $v(a, i)$, depends on the way the government policy is articulated, a, and whether the individual purchases insurance, i. Now, the individual purchases insurance whenever

$$v(a, 1) - v(a, 0) > U(w-f, 0) - U(w-p, 1)$$

Our experimental design holds constant the financial penalty f for remaining uninsured but changes the articulation a. The articulation a encompasses how individuals perceive the policy, and depends on both what the government calls the policy and how it is interpreted in public discourse. Whether policy called a "mandate" is perceived as legitimate (or, for example, constitutional) may affect whether a mandate carries normative weight. While the articulation a is the product of both what the policy is called and the broader political discourse, for simplicity in describing the experiment, we will simply refer to a as being an element of $\{mandate, tax\}$.

Since random assignment in our experiment will give us the same distribution of $U(w-f,0) - U(w-p,1)$ in both treatments, we will be able to identify the relationship between $v(mandate, 1) - v(mandate, 0)$ and $v(tax, 1) - v(tax, 0)$. If people are more likely to buy insurance under the *mandate* condition than under the tax condition at a point in time, then $v(\text{mandate}, 1) - v(mandate, 0) > v(tax, 1) - v(tax, 0)$, and we can conclude purchasing insurance increases utility derived from the normative component more in the mandate articulation than the tax articulation—that is, the mandate articulation is more effective than the tax articulation at that point in time.

3. Experimental design

In our experiment, all participants read a single policy vignette and decided how likely they would be to purchase health insurance. Participants were randomly assigned to a mandate or tax condition (a between-subject design).[8] Participants were told that the state of healthcare policy was in flux and asked to suppose that the government decided:

- *Mandate condition*: "to mandate everyone purchase insurance, or else pay a fine of $700 each year"
- *Uninsurance tax condition*: "to recommend that everyone purchase health insurance, and charge people without insurance an uninsurance tax of $700 each year"

Then, participants were asked what they would do if their current health insurance policy were no longer available and they were to become uninsured. Particularly, they were asked to choose between purchasing "coverage that is as good as the coverage that members of Congress get" at a market price of $3000 per year, or staying uninsured (i.e. this was the only insurance policy to which they had access). Participants indicated how likely they would be to purchase the policy on a 7 point scale, ranging from "almost certain to buy the policy (96–100% chance of buying the policy)" to "almost certain to stay uninsured (0–4% chance of buying the policy)"; see *Online Appendix* for the details of the question and response scale. After answering this question, participants were asked, on the same scale, how likely they would be to purchase the policy if it instead cost $2000 per year. They then saw a number of follow-up questions, which are described below and in *Online Appendix*.

By design, the mandate and tax conditions differ in two elements: the mandate condition has a mandate with a fine for violating it, while the tax condition has a recommendation with a tax for not purchasing insurance. We chose to have the tax condition contain a recommendation to purchase insurance so that both conditions conveyed the information that purchasing health insurance was a "good idea"; it is possible that a tax without a recommendation would be less effective. We choose to pair the mandate with a fine rather than a "penalty" or simply a "required payment" to make the violating the mandate similar to violating other laws (traffic, taxes, etc.).[9] Neither condition explicitly mentioned the probability that the fine/tax would be collected. Thus, the articulation of the policy could have affected participants' beliefs about resources devoted to enforcement, but given the phrasing of the question and pattern of results, it is natural to believe that the probability of paying the fine was similar (and near 1) in both cases. As we will show below, our effects cannot be explained by differential beliefs about enforcement. In addition, we did not mention or imply any additional penalty associated with the mandate (the enforcement mechanism for the mandate in the PPACA is done via the tax system).

[8] In the first and second wave, some participants were also assigned to a subsidy, or status quo condition. These conditions were discontinued as a result of a power calculation and because the status quo eventually became a mandate.

[9] "Civil penalty" and "civil fine" are often used as synonyms.

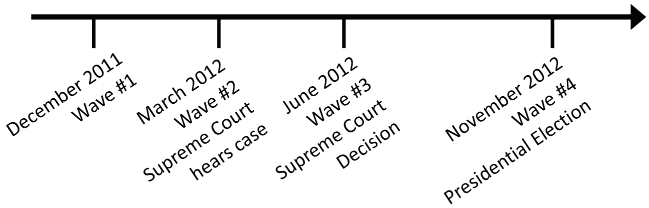

Fig. 1. Timeline of survey waves and major events.

Table 1
Subject demographics.

Age (mean)	30.2
Female	43.1%
College graduate or above	47.8%
Political affiliation	
Republican	16.1%
Democrat	43.6%
Independent	40.3%
Unemployed	21.9%
Has insurance	71.7%
Married	28.6%
Census region	
Midwest	24.2%
South	33.7%
Northeast	18.9%
West	22.7%
Purchase probability ($3000)	58.5%
Purchase probability ($2000)	66.8%
Neighbor's purchase probability	50.1%
Support for uninsurance (% of bills covered by charity/government)	53.9%
Socially appropriate to be uninsured (1–4), mean	2.4
N: Pre-controversy waves	263
N: Supreme court wave	784
N: Nov. election wave	623
Total *N*	1670

4. Data and waves

We conducted our experiment in four waves from December 2011 to November 2012, which spanned the controversy about the constitutionality of the individual mandate in the PPACA. See Fig. 1 for the timeline of events and when our survey waves were conducted.

We analyze results from 1670 participants recruited for our study from an online labor market. Each participant completed our study in one of the four waves; each wave had the same selection criteria and recruited participants to take a generic "survey". Recruitment used neutral language that did not mention healthcare to avoid selection into the experiment based on political beliefs.[10] Participation was limited to U.S. residents, and while not a representative sample of the U.S. population, this labor market (Amazon's Mechanical Turk) is regularly used for research studies (Horton et al., 2011; Paolacci et al., 2010). We excluded participants aged 65 and over a priori because they are covered by Medicare. Our participants vary significantly on educational background, age, employment status, and are geographically diverse, as reflected in Table 1.

The individual mandate in PPACA was the target of much discussion and controversy in 2012. In March 2012, the Supreme Court heard oral arguments regarding the constitutionality of elements of PPACA, including the individual mandate.[11] There was an increase in political discussion of the mandate in late March, leading up to and following the oral arguments. That discussion became less intense in April and May, but was followed by increased attention in June surrounding the release of the Court's ruling on PPACA. We will use this controversy to look for a changing effect of the policy articulation on insurance purchase behavior.

Fig. 2 provides a proxy for the intensity of political discussion regarding the mandate. It reports the number of articles published in U.S. newspapers by month that mention the terms: "health insurance mandate", "mandate AND ACA", and

[10] Participants on MTurk were recruited to "Participate in a brief survey: about 10 minutes," and paid 31 cents. The consent form (first page of the experiment) read: "This study examines decision-making. The results will be used to examine what people find important when making decisions."

[11] The minimum coverage provision was argued March 27th in *Department of Health and Human Services, et al., Petitioners v. Florida, et al.* However other aspects of the case were argued before and after this, and news media covered the issue in advance of the arguments, with some discussion in late February and the first two weeks of March.

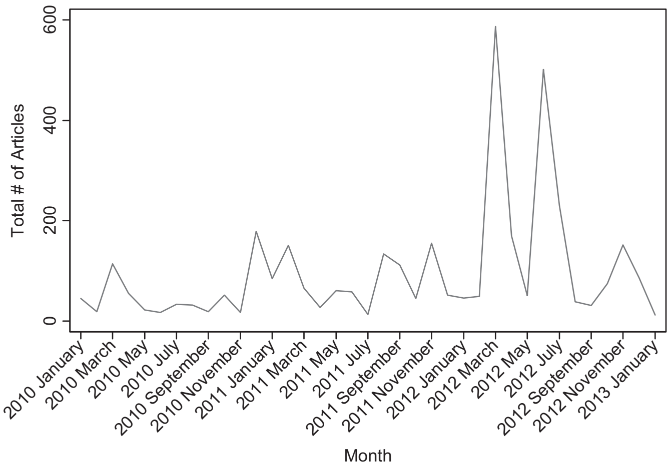

Fig. 2. Number of news articles relating to the mandate, by month. *Notes*: Plots the sum of three measures of news activity. Constructed from the number of news articles in U.S. Newspapers indexed in Factiva matching the terms "health insurance mandate", "mandate AND ACA", and "(individual mandate OR insurance mandate) AND unconstitutional".

"(individual mandate OR insurance mandate) AND unconstitutional".[12] The three individual measures are all highly correlated (pairwise correlations range from 0.63 to 0.81) and show the same pattern, so we simply display their sum.[13] The two peaks are occasioned by the oral argument (March 2012) and ruling (June 2012).[14]

We classify Waves 1 and 2 as pre-controversy and Waves 3 and 4 as post-controversy. The development of the controversy in March took place after most of our March wave was complete (see Fig. A1 in *Online Appendix*, which plots two series for March 2012: number of surveys completed by day and number of mandate-related news articles by day). However, our results are robust to alternative ways of defining the controversy period as described in Section 5 and displayed in Table 3.

5. Results

Our primary analysis uses participants' reported probabilities of purchase from the response scale. We set the probability of purchase to the midpoint of the probability range for each participant's selected option. We present means by condition and linear regression estimates that control for individual characteristics. With these methods, differences in choices between conditions can easily be interpreted as differences in reported probability of purchase. We also replicate our results using an ordered probit specification, which treats the probability of purchase as an ordinal variable and simply assumes that higher choices reflect higher likelihood of purchase. This specification is more robust in a statistical sense, but it does not have a natural interpretation in terms of purchase probability.

When we regress reported probability of purchase on insurance price, we use two observations per participant, since each participant was asked about their likelihood of buying insurance at a high price ($3000) and a low price ($2000). In those specifications, standard errors are clustered at the participant level to account for interdependence between the participant's two choices. We obtain similar results if we limit our analysis to an individual's first choice.

5.1. Mandate versus tax pre-controversy

We find a substantial and statistically significant effect of articulation on reported probability of insurance purchase in the pre-controversy period. Subjects who randomly receive the mandate articulation are significantly more likely to buy insurance. The increase due to the mandate articulation is comparable to a $1000 reduction in annual premiums.

Result 1. Before the controversy (i.e. in Waves 1 and 2), the mandate articulation generates more insurance purchase than a financially equivalent tax.

Table 2 analyzes results from all the waves in the experiment. Columns 1 and 2 show regressions on the reported probability of purchase. The coefficient on *Mandate (v. Tax)* displays the relative effectiveness of the mandate articulation, as

[12] Where "AND" and "OR" are the logical arguments commonly used in text-based search.

[13] Displaying the sum of the three measures double or triple counts articles that include more than one of our search terms. We think this is reasonable given that the individual mandate is more likely to be central to an article that involves multiple search terms. Regardless, the individual measures show a similar pattern.

[14] Moreover, there is evidence that individuals were paying attention to the controversy. The Kaiser health tracking poll asked individuals how much they had seen in the news about the health reform law twice: in November 2011 and April 2012. Over that time, the fraction of people saying "a lot" rose from 18 to 36% and the fraction saying "a lot" or "some" rose from 41% to 60% (Kaiser Family Foundation, 2012). According to the same source, overall favorability ratings of the ACA were relatively stable during this time.

Table 2
The relative effectiveness of the mandate over time.

	Probability of purchase OLS		Ordered probit	
	(1)	(2)	(3)	(4)
Annual premium ($1000s)	−8.304***	−8.304***	−0.267**	−0.298***
	(0.374)	(0.375)	(0.012)	(0.014)
Mandate (v. Tax)	10.18**	9.910**	0.324**	0.350***
	(4.365)	(3.979)	(0.133)	(0.134)
+Mandate × [Wave 3 (June)]	−12.06**	−9.647**	−0.388**	−0.347**
	(4.950)	(4.518)	(0.152)	(0.153)
+Mandate × [Wave 4 (November)]	−11.28**	−12.54**	−0.346**	−0.430***
	(5.132)	(4.712)	(0.156)	(0.158)
Effect of survey wave	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes
R^2	0.03	0.20	N/A	N/A
N Participants	1670	1670	1670	1670
N Observations	3340	3340	3340	3340

Notes: Two observations per participant (purchase probabilities at different prices). Heteroskedasticity robust standard errors clustered at the participant level are included in parentheses. "Mandate v. Tax" reports the relative effect of the mandate in the pre-controversy waves. The interactions capture how that effect changes in later waves. Columns 1 and 3 report results without demographic controls. Columns 2 and 4 reports results with demographic controls (age, gender, indicators for current insurance status and source, marital status, number of children, educational attainment, employment status, risk aversion, political affiliation, survey wave, and region).
** $p < 0.05$.
*** $p < 0.01$.

compared to the tax articulation, in the first two waves of the experiment. In these waves, the mandate articulation increases insurance purchase by about 10 percentage points relative to the tax, a difference that is significant at $p < 0.05$. Comparing columns 1 and 2, we see that adding demographic controls improves precision but does not change the point estimate of the treatment effect.[15]

Because each person sees two different premiums, $3000 and $2000, we can estimate the effect of price on choice, as well as benchmark the articulation effect in dollar terms. Higher insurance premiums lead to significantly lower insurance purchase: a $1000 increase in annual premiums reduces insurance purchase by 8.3 percentage points. Dividing the coefficient on the mandate by the coefficient on price gives a dollar estimate of the mandate's articulation effect relative to the tax. Articulating the policy as a mandate rather than a tax has an estimated effect equivalent to a $1226 (column 1) or $1193 (column 2) decrease in price.

The ordered probit specifications in Table 2 treat the probability of purchase as an ordinal variable. The positive and statistically coefficient on mandate in columns 3 and 4 indicates that the mandate articulation induces more insurance purchase. In the ordered probit model, the coefficients represent changes in the probit index, with response cutpoints also estimated via maximum likelihood. A positive sign on *Mandate (v. Tax)* indicates that the model predicts that probability of purchase increases under the mandate. Comparing the coefficient on mandate to that on premiums confirms again that the effect of the mandate articulation is comparable $1000 change in premiums.

5.2. Mandate versus tax post-controversy

In response to the controversy that developed in March and April 2012, we ran additional waves of the experiment to identify how the relative effectiveness of the mandate and tax articulations were changing over this time period. We ran Wave 3 on two subsequent days, launching the evening of June 27 (the day before the Supreme Court ruling) and then the evening of June 28 (the day of the ruling).[16] Wave 4 began on October 31, 2012, ran through Election Day (November 6), and continued until November 8, 2012. While attitudes about health insurance may change over time, our design includes time fixed effects: our estimated effect of the mandate captures the relative effect of policy articulation on insurance purchase. Moreover, our results do not rely on participants' knowledge of the actual policy outside of the experiment—we delineate the consequences of both the mandate and tax policy articulations prominently in the experiment.

Result 2. After the political controversy over the mandate, there is no longer a difference between the mandate articulation and the tax articulation.

[15] We pool Waves 1 and 2 together, since they predated the measured controversy over the mandate. When we estimate the model in column 2 separately for each of these waves, we get similar results.
[16] We increased the sample size and ran on subsequent days to see if we could identify a discrete change in the articulation effect of the mandate. The secular trend in the articulation effect turned out to be much more important than the event of any given day.

Table 3
The effect of the mandate over time, robustness.

Specification	Mandate (v. tax) effect in pre-controversy period	Change in mandate effect in post-controversy period	Participants in sample
Full sample	9.90**	−10.92**	1670
	(3.98)	(4.30)	
Old (age > 26)	12.94***	−12.21***	865
	(4.66)	(5.26)	
Young (age ≤ 26)	4.30	−6.93	805
	(7.27)	(7.63)	
Women	10.82**	−10.35*	720
	(5.24)	(5.83)	
Men	9.08	−11.17*	950
	(6.20)	(6.54)	
Uninsured	14.88**	−16.34**	472
	(6.73)	(7.52)	
Has insurance	7.62	−8.09	1198
	(4.88)	(5.22)	
Married	22.32***	−21.55***	477
	(6.04)	(6.89)	
Unmarried	3.60	−5.26	1193
	(5.13)	(5.49)	
Employed full-time	11.92**	−15.92**	684
	(5.81)	(6.36)	
Not full-time	9.41*	−8.05	986
	(5.37)	(5.83)	
Exclude March	11.08**	−12.16**	1558
	(5.27)	(5.52)	
Post-controversy starts March 25	10.29**	−11.28**	1670
	(4.11)	(4.42)	
Post-controversy starts March 16	12.46***	−13.21***	1670
	(4.78)	(5.04)	

Notes: The regression specification parallels that of Table 2, Column 2, and includes annual premium, the same set of controls, and the main effect of survey wave. Each row reports the results of an OLS regression, with the dependent variable as probability of purchase. "Mandate (v. tax) effect in pre-controversy period" gives the coefficient describing the differential effect of the mandate in the pre-controversy period. "Change in mandate effect in post-controversy period" give the interaction between the mandate's effect and an indicator for being in Waves 3 and 4. As in Table 2, there are 2 observations per participant (purchase probabilities at different prices). Heteroskedasticity robust standard errors clustered at the participant level are included in parentheses.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

In Table 2, the coefficient on *Mandate* × [*Wave* 3 (*June*)] gives the *change* in the mandate's relative effectiveness in Wave 3, compared to the pre-controversy waves; the coefficient in the row below gives the analogous comparison between Wave 4 and the pre-controversy waves. Post-controversy, the relative effectiveness of the mandate articulation disappears. The relative effectiveness of the mandate is significantly different pre- and post-controversy. Post-controversy, the point estimate on relative effect of the mandate is virtually zero in Wave 3 (10.18 − 12.06 = −1.88 in column 1; similarly for column 2) and in Wave 4 (10.18 − 11.28 = −1.10 in column 1; similarly for column 2) and we cannot reject the null hypothesis that the relative effectiveness of the mandate articulation is zero in Waves 3 and 4 (F-test: $p = 0.42$ and $p = 0.68$, respectively). Moreover, the estimated relative effect of the mandate is similar across the two post-controversy waves; the coefficients on the Waves 3 and 4 mandate interaction terms are not significantly different from each other ($p = 0.48$). The ordered probit results again confirm the OLS findings.

5.3. Robustness of results

While in Section 2 we showed that in each wave we are able to identify the relative effect of the mandate articulation and tax articulation on insurance purchase intentions, one might be concerned that demographic differences between participants in different waves might lead to different estimates of the articulation effect of the mandate.[17] To address this, we run specifications for a variety of subgroups. In Table 3, we present the relative effectiveness of the mandate pre-controversy and the post-controversy interaction (using one interaction for Waves 3 and 4 together) for a variety of subgroups. For every subgroup where we see a significant effect in the pre-controversy period, we find a similarly sized negative interaction in the post-controversy period.

[17] Similarly, if economic conditions were changing in a way that impacted sensitivity to the mandate only, it could affect our estimates. But evidence suggests that economic conditions were relatively constant; the unemployment rate was slowly and smoothly trending down during the period of our experiment. In addition, Table 3 shows we get the same effects when restricting our estimates to full-time workers in all waves.

While the demographics of the sample are similar across different waves, Table A1 in *Online Appendix* shows that there are some minor differences. Since Table 3 showed that the mandate's relative effect disappeared for all subgroups, changes in demographic composition do not explain our results. However, to further rule out any sample composition story, we create a synthetic post-controversy sample that matches our pre-controversy sample more closely using nearest neighbor matching (Blackwell et al., 2009). Table A1 shows how closely our nearest neighbor sample matches the pre-controversy sample. Table A2 in *Online Appendix* verifies that results on the mandate's relative effectiveness and its change over time are the same with this sample as in the main specification reported in Table 2.

Another dimension on which we show the robustness of our results is with the start of the post-controversy period. The initial spike in press about the PPACA mandate happened in March 2012, a month in which we also collect pre-controversy survey responses. Fig. A1 in the *Online Appendix* shows the news activity measure that comprises Fig. 2 broken down by day of March 2012. Looking at this figure suggests two alternative ways in which the start of the post-controversy period might be defined: the beginning of the press build up in early March or the major spike in articles in late March. Table 3 shows that our results hold in additional specifications that: (i) exclude March entirely from the dataset, (ii) move the first post-controversy date to March 25th 2012 (the day before the big spike in press), or (iii) move the first post-controversy date to March 16th 2012 (the day before the press begins to build up). All of these specifications show a significant relative effect of the mandate articulation pre-controversy and a negative interaction post-controversy.

5.4. Channels through which mandate articulation might have an effect

We have shown that the mandate articulation motivates a higher likelihood of insurance purchase than a financially equivalent tax before the controversy and that this relative effectiveness disappears post-controversy. What might be driving this effect? We argue that the relative effectiveness pre-controversy results from a moral suasion or a normative motive (i.e. a perceived individual obligation to comply with the law). Here, however, we discuss some other ways in which the mandate articulation may have affected behavior before the controversy and assess the evidence for these stories.

One alternative explanation is that the mandate affects behavior through social channels. For instance, the social pressure to have insurance that an individual feels from other people might differ under the mandate versus tax articulation, and the level of support an uninsured individual receives might change across these articulations.

To examine these social channels, each participant was asked to answer three additional questions, under the same mandate or tax articulation, after giving their purchase intentions:

1. Suppose an individual in your neighborhood was uninsured, but was given the opportunity to get themselves coverage by purchasing the same health insurance policy just described at the cost of $3000 per year. How likely do you think they would be to purchase this policy versus staying uninsured? (Same response scale as own purchase question).
2. Suppose someone in your community of average income was offered health insurance but chose not to buy it, despite the government's recommendation. After showing symptoms of weight-loss, nausea, abdominal pain, they were diagnosed with pancreatic cancer and needed expensive treatment to stay alive. Because they were uninsured, they might not be able to pay for this care. How much support should this person get from charity care and/or government safety net programs, such as Medicaid? (6 categories of response, from "A very generous amount of support: 81–100% of medical bills," to "No Support: 0% of medical bills").
3. Suppose someone in your community chose not to buy health insurance, despite the government's recommendation. How would you evaluate their decision not to buy health insurance? (Response options: "Very socially inappropriate," "Somewhat socially inappropriate," "Somewhat socially appropriate," "Very socially appropriate").

Because these channels might be affected by the articulation of the policy, we include them in Table 4 as control variables in a model of insurance purchase estimating an effect of the mandate articulation.

We find that the relative effectiveness of the mandate articulation is not substantially impacted when we control for these three measures of social incentives to be insured. The pre-controversy effect of the mandate, relative to the tax, drops slightly, from about 10 percentage points before controlling for these channels (Table 2) to about 8 percentage points with these controls (Table 4). To the extent that these questions capture the social incentives to get insured, the results indicate that the mandate articulation had an additional effect beyond the social effect, which we interpret as moral suasion. This moral suasion effect disappears in the post-controversy waves.

The channel controls are included linearly in column 1 of Table 4, allowing us to see how responses to these questions are related to insurance purchase probability. The responses are sensible: own purchase is positively correlated with estimates of neighbor's purchase, and negatively correlated with judgments of how socially appropriate it would be to remain uninsured (we do not find an association between "Deserving of Support if Uninsured" and probability of purchase). Column 2 includes

Table 4
The effect of the mandate, controlling for social incentives.

	Probability of purchase OLS	
	(1)	(2)
Annual premium ($1000s)	−8.304***	−8.304***
	(0.375)	(0.376)
Mandate (v. Tax)	7.847**	7.679**
	(3.251)	(3.223)
+Mandate × [Wave 3 (June)]	−6.716*	−6.668*
	(3.709)	(3.678)
+Mandate × [Wave 4 (November)]	−10.158***	−10.305***
	(3.735)	(3.701)
Probability neighbor would purchase	0.496***	As Categories
	(0.0239)	
Deserving of support if uninsured (% of medical bills)	0.0308	
	(0.0248)	
Social appropriateness of uninsurance (1–4)	−12.40***	
	(0.885)	
Effect of survey wave	Yes	Yes
Controls	Yes	Yes
R^2	0.473	0.488
N Participants	1670	1670
N Observations	3340	3340

Notes: Two observations per participant (purchase probabilities at different prices). Heteroskedasticity robust standard errors clustered at the participant level are included in parentheses. "Mandate (v. Tax)" reports the effect of the mandate in the pre-controversy waves. The interactions capture how that effect changes in later waves. Controls include age, gender, indicators for current insurance status and source, marital status, number of children, educational attainment, employment status, risk aversion, political affiliation, survey wave, and region.

* $p < 0.1$.
** $p < 0.05$.
*** $p < 0.01$.

controls for each category of response to the three social questions, as the linearity assumption may be inappropriate. The estimated effect of the mandate relative to the tax is unchanged.[18]

Another potential alternative explanation is that participants in the pre-controversy waves thought it would be easier to evade paying the tax than the fine (giving the mandate additional effectiveness) and that they were dissuaded of this in later waves. Again, we have multiple reasons to think that evasion beliefs are not driving the effect. First, this explanation would require participants to have ignored the explicit statement in the wording of the question that they must buy insurance or pay the tax/fine. Second, the financial consequences of the decision were integrated into a table detailing the costs faced under each policy at the end of the question about insurance purchase (see Experimental Materials in *Online Appendix*), making clear that they were being asked about an environment where they were expected to pay the tax or the fine. Third, the magnitude of the mandate's relative effectiveness (more than $1000 in premiums) is larger than the effect of being able to certainly evade the tax versus pay the fine for sure.

A related issue is why the articulation effect of the mandate dissipated in response to the political discourse between March and June. We have two, potentially complementary, hypotheses. First, the political discussion may have changed individuals' beliefs about the mandate's legitimacy and moral claim. A related version of this first hypothesis is that the controversy gave individuals "moral wiggle room" (i.e. an excuse) to avoid complying with the mandate, even if they did not themselves dispute its legitimacy.[19] Second, the increased attention to the issue may have made the mandate and the tax equivalent in the minds of individuals—in fact, advocates for the mandate's constitutionality frequently pointed out this financial equivalence.

One potential alternative explanation is that the relative effectiveness of the mandate might have changed over time because more information was available to participants about the Affordable Care Act that was unrelated to the mandate's legitimacy or making explicit the equivalence of the two policies. However, we do not believe changes in other information are driving our effect. Survey evidence suggests that general knowledge about the Affordable Care Act was not changing over time. The Kaiser Family Foundation's Health Tracking Polls show that measures of knowledge were stable throughout our

[18] Table A3 in the *Online Appendix* directly examines the effect of the mandate on these social measures, and how the mandate's effect on these measures varies by wave. The results show that the mandate articulation has a small and insignificant impact on answers to these questions. Moreover, the effect of the mandate articulation on these answers does not vary much by wave.

[19] See Dana et al. (2007) for an experimental demonstration that giving individuals moral wiggle room reduces pro-social behavior.

study period (they have four measures collected between April 2010 and March 2012).[20] Moreover, these polls show collaborating evidence that the normative element of the mandate was changing over time: when poll participants in Dec. 2011 were informed that a case challenging the constitutionality of the mandate was being heard at the Supreme Court, support for the mandate as a public policy dropped 10 percentage points. Thus, the evidence indicates that rather than changing information over time, the mandate's legitimacy was challenged by public debate, reducing its relative effectiveness.

One final issue to address is how our results might be affected by social desirability bias in which experimental participants distort their answers in a way they believe is socially desirable. For example, if it were socially desirable to comply with the mandate, participants may have been more likely to *say* they would purchase insurance under the mandate than under the tax, but purchase behavior would not respond accordingly. If this were the case, our estimates of relative effectiveness of the mandate would be inflated.

Note that the social desirability story relies on our story—people must believe it is more important to comply with the mandate than the tax in the pre-controversy waves for a social desirability bias to exist. However, we do not think our evidence is consistent with a social desirability bias. We directly elicit social appropriateness in one of our questions, as described in the previous section. If social desirability bias were playing a large part in our results, this question should pick it up, but Table 4 shows that the relative effectiveness of the mandate is not substantially changed when controlling for such social channels. Moreover, Table A3 in *Online Appendix* shows that the mandate did not affect how participants would judge the social appropriateness of being uninsured, which is inconsistent with a story in which social desirability bias led to differential lying under the mandate policy. Finally, it is worth emphasizing that as with almost all experimental papers our interest is not in estimating a precise magnitude (see Kessler and Vesterlund, 2015); rather, it is in identifying whether the efficacy of a government policy can be affected by articulation and whether this articulation effect can be undermined.

6. Discussion

The experimental results presented here demonstrate that the way a government policy is articulated can alter how it affects behavior. Before the controversy over the PPACA mandate, individuals are particularly inclined to obey a mandate to purchase health insurance as compared to a tax with the same financial consequences. Our data suggest that the relative effectiveness of the mandate is driven by moral suasion rather than perceptions of the social appropriateness of remaining uninsured. After the controversy about the individual mandate, we observe that the differential effect of the mandate disappears; the mandate and the tax then encourage insurance purchase at the same rate. This result demonstrates that the articulation of government policy can be influenced by political discourse. Our results give guidance to policymakers attempting to advance policy prescriptions and who may articulate their policy in different ways. The results suggest that the articulation of policy can substitute for financial incentives. Before the controversy, a larger tax would have been needed to achieve the same effect on behavior as a mandate with an accompanying fine. Individuals respond to these articulations in significant ways. Consequently, the effectiveness of a particular government policy depends not only on its financial incentives but also on how it is articulated, which is at least partially within the policy maker's control.

Conflicts of interest

The authors declare no conflicts of interest.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2015.09.021.

References

Auerbach, D., et al., 2010. Will Health Insurance Mandates Increase Coverage? Synthesizing Perspectives from Health, Tax, and Behavioral Economics. Congressional Budget Office Working Paper Number 2010-05.
Ajzen, I., 1991. The theory of planned behavior. Organ. Behav. Hum. Decis. Process. 50 (2), 179–211.
Becker, G., 1968. Crime and punishment: an economic approach. J. Polit. Econ. 72 (2), 169–217.
Benabou, R., Tirole, J., 2011. Laws and Norms. NBER Working Paper Number 17579.
Blackwell, et al., 2009. Coarsened exact matching in Stata. Stata J. 9 (4), 524–546.
Camerer, C.F., Hogarth, R.M., 1999. The effects of financial incentives in experiments: a review and capital-labor-production framework. J. Risk Uncertain. 19 (1), 7–42.
Carson, R.T., Hanemann, W.M., 2005. Contingent Valuation. Handbook of Environmental Economics, vol. 2., pp. 821–936.
Chetty, R., Looney, A., Kroft, K., 2009. Salience and taxation: theory and evidence. Am. Econ. Rev. 99 (4), 1145–1177.
Dana, J., Weber, R.A., Kuang, J.X., 2007. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. Econ. Theory 33 (1), 67–80.

[20] One question in the poll asked whether participants had enough information to understand the health reform law. The fraction saying they did not have enough information was 56%, 52% 55%, and 59% in April 2010, March 2011, November 2011, and March 2012 (these are the months in which the question was asked). There is no upward trend in participants having knowledge about the health care law (see "March 2012 Toplines" report available at kff.org).

Dharmapala, D., McAdams, R.H., 2003. The Condorcet jury theorem and the expressive function of law: a theory on informative law. Am. Law Econ. Rev. 5 (1), 1–31.
Elmendorf, D., 2011. CBO's Analysis of the Major Health Care Legislation Enacted in March 2010. Congressional Budget Office.
Elster, J., 1989. Social norms and economic theory. J. Econ. Perspect. 3 (4), 99–117.
Ericson, K.M.M., Fuster, A., 2011. Expectations as endowments: evidence on reference-dependent preferences from exchange and valuation experiments. Q. J. Econ. 126 (4), 1879–1907.
Finkelstein, A., 2009. E-ZTAX: tax salience and tax rates. Q. J. Econ. 124 (3), 969–1010.
Funk, P., 2007. Is there an expressive function of law? An empirical analysis of voting laws with symbolic fines. Am. Law Econ. Rev. 9 (1), 135–159.
Gruber, J., 2011. Massachusetts points the way to successful health care reform. J. Policy Anal. Manag. 30 (1), 184–192.
Hoffrage, U., et al., 2000. Communicating statistical information. Science 290 (5500), 2261–2262.
Horton, J., Rand, D., Zeckhauser, R., 2011. The online laboratory: conducting experiments in a real labor market. Exp. Econ. 14 (3), 399–425.
Jacoby, W.G., 2000. Issue framing and public opinion on government spending. Am. J. Polit. Sci. 44 (4), 750–767.
Johnson, E.J., Hershey, J., Meszaros, J., Kunreuther, H., 1993. Framing, probability distortions, and insurance decisions. J. Risk Uncertain. 7 (1), 35–51.
Kahneman, D., Knetsch, J.L., Thaler, R.H., 1990. Experimental tests of the endowment effect and the Coase theorem. J. Polit. Econ. 98 (6), 1325–1348.
Kaiser Family Foundation, 2012. Kaiser Health Tracking Poll: April 2012, Topline. Publication Number 8302, http://www.kff.org/kaiserpolls/8302.cfm.
Kessler, J.B., Vesterlund, L., 2015. The external validity of laboratory experiments: the misleading emphasis on quantitative effects. In: Guillaume, R.F., Schotter, A. (Eds.), Handbook of Experimental Economic Methodology. Oxford University Press.
Kesternich, I., Heiss, F., McFadden, D., Winter, J., 2013. Suit the action to the word, the word to the action: hypothetical choices and real decisions in Medicare Part D. J. Health Econ. 32 (6), 1313–1324.
Krueger, A., Kuziemko, I., 2013. The demand for health insurance among uninsured Americans: results of a survey experiment and implications for policy. J. Health Econ. 32 (5), 780–793.
Lacetera, N., Pope, D.G., Sydnor, J.R., 2012. Heuristic thinking and limited attention in the car market. Am. Econ. Rev. 102 (5), 2206–2236.
Liberman, V., Samuels, S.M., Ross, L., 2004. The name of the game: predictive power of reputations versus situational labels in determining prisoner's dilemma game moves. Pers. Soc. Psychol. Bull. 30 (9), 1175–1185.
McNeil, B., et al., 1982. On the elicitation of preferences for alternative therapies. N. Engl. J. Med. 306, 1259–1262.
National Federation of Independent Business et al. v. Sebelius, Secretary of Health and Human Services et al., 567 U.S. 58 (2012).
Paolacci, G., Chandler, J., Ipeirotis, P., 2010. Running experiments on Amazon mechanical turk. Judgm. Decis. Mak. 5, 411–419.
Pear, R., 2010, July. Changing Tune, Administration Defends Insurance Mandate as a Tax. New York Times.
Steinbrook, R., 2008. Health care reform in Massachusetts—expanding coverage, escalating costs. New Engl. J. Med. 358, 2757–2760.
Sunstein, C.R., 1995. On the expressive function of law. Univ. Pa. Law Rev. 144, 2021–2053.
Sunstein, C.R., 1996. Social norms and social rules. Columbia Law Rev. 96 (4), 903–968.