
Masking Redistribution (or its Absence)

Jonathan Baron and Edward J. McCaffery

**USC CLEO Research Paper No. C04-1
USC Law and Economics Research Paper No. 04-5**

**CLEO RESEARCH PAPER SERIES
LAW & ECONOMICS RESEARCH
PAPER SERIES**

Sponsored by the John M. Olin Foundation

University of Southern California Law School
Los Angeles, CA 90089-0071

*This paper can be downloaded without charge from the Social Science Research Network
electronic library at <http://ssrn.com/abstract=528165>*

Masking Redistribution (or its Absence)

Jonathan Baron*
University of Pennsylvania

Edward J. McCaffery†
University of Southern California Law School
California Institute of Technology

July 21, 2004

Abstract

Research has shown that people vary widely in their support or opposition to progressive taxation. We argue here that the perception of progressiveness itself is affected by the nature of the tax system and by the way it is framed, or presented. Experiments conducted over the World-Wide Web and using within-subject design demonstrate that subjects suffer from a range of heuristics and biases in understanding and supporting progressive or redistributive taxation. After reviewing some prior results, we report four new studies. Two of them indicate that people do not sufficiently appreciate the reduction of progressiveness that results from the use of tax deductions to partly reimburse private expenditures. The other two indicate that people do not fully appreciate the reduction in progressiveness that results from cuts in government services.

Introduction

The two welfare theorems of neoclassical welfare economics suggest a two-part agenda for public finance, as well as for economic regulation more generally (Kaplou and Shavell 2000). First, policymakers should choose socio-economic policies that maximize social welfare, for which wealth-maximization typically serves as a suitable proxy; second, policymakers should redistribute the greater social wealth, through a tax and transfer system, to achieve a desired location on the pareto frontier. In public finance, this means that decisions about the proper role for the government, and the structure of that role — what

* baron@psych.upenn.edu. The research was supported by NSF grant SES 02-13409. The authors thank participants at a Behavioral Public Finance Conference held at the University of Michigan in April, 2004, especially Alan Lewis for his detail commentary

† emccaffe@hss.caltech.edu.

public goods ought to be provided, to what extent, and how financed — should be made in a strict welfare maximizing fashion, minimizing the deadweight losses or excess burdens of the fiscal intervention (Musgrave, 1959; Rosen, 2001).

Put in equivalent terms, any government fiscal action naturally has two effects, on the allocation of social resources and on their distribution. Logically, the welfarist view of public finance entails separating out these allocative and distributive dimensions, to achieve an optimal merger of efficiency and equity concerns, in the spirit, say, of Rawls's difference principle (Rawls, 1971). If a current government program is inefficiently designed, social welfare can be enhanced — and all individuals in society can stay at the same or a higher level of personal, welfare — by moving to a more efficient government (or private) provision, coupled with a redistribution of the greater wealth to meet the Paretian constraint.

So it is in logic. But there is a reason to believe that cognitive heuristics and biases get in the way of ordinary persons, such as citizens, understanding matters once two or more aspects of government action are separated; this is an instance of what we have elsewhere termed a *disaggregation bias* (McCaffery and Baron, 2003). We present a series of experiments in which we detail the ways in which the extent of redistribution that subjects are willing to support depends on the form of government provision of goods and services: for example, subjects are willing to support a higher level of redistribution or cross-subsidization through the provision of a good than they are when the redistributive function of government is broken off, or peeled away, and left to stand on its alone. This violates the spirit of the two welfare theorems by generating a conflict between equity and efficiency — between total social wealth and the well-being of society's members — that does not exist in pure, rational theory.

In the present chapter, we ask whether people understand, and take into account, the distributional effects of two kinds of policy changes: the use of tax deductions as a way of funding services, and the effect of privatization or elimination of government services and programs. Our main hypothesis is that, by failing to take distributional effects fully into account, people end up supporting policies that undercut their own values.

General Experimental Method

We followed the same general procedure in each of the studies reported below.

Subjects completed a questionnaire on the World-Wide Web. Subjects were paid \$3 or \$4 each. Subjects came to the studies through postings on various Web sites or Usenet news groups, or through prior participation in other studies. Subjects were paid by check (after some minimum amount was accumulated), and they had to register their address and (for U.S. residents) their Social Security number in order to get paid, but they identified themselves only with their email address after they registered, and the email address was stored separately from the data, to assure privacy and anonymity.

Consistent with standard methods in cognitive psychology, our experimental designs were all within-subject (Baron, 2000). Our subject pool is roughly representative of the

adult U.S. population in terms of income, age, and education (Babcock et al., 2003), but not in terms of sex, because (for unknown reasons) women predominate in our respondent pool.

In terms of precise methods, the studies were programmed in JavaScript so that one case was presented on each page, and subjects were required to answer all questions appropriately before proceeding. We recorded the time spent on each response, and we usually eliminated subjects who went noticeably faster than everyone else (outliers, typically 2–3%). Many of our experiments have had internal checks to assure that subjects understand the questions, and answer in the appropriate range; the first and third experiments reported below, for example, had such checks on understanding, the results of which we report.

Hypotheses

Research has long shown that subjects on average prefer at least moderately progressive tax rates (Lewis, 1978, Hite and Roberts, 1991, McCaffery and Baron 2003). Redistribution through a government typically comes bundled with other fiscal actions, such as the public provision of goods and services. A type of *isolation effect* seems to lie at the center of many cognitive heuristics and biases: subjects focus on one aspect of a choice set or problem, ignoring (or under-considering) others (Camerer, 2000, Read, Loewenstein & Rabin, 1999, McCaffery and Baron, 2003). Putting these facts together, we surmised that subjects would not always be consistently attentive to the redistributive aspects of a fiscal program.

In prior experiments we have found that subjects have preferences over the purely formal labeling of government charges — whether these are called “taxes” or something else, such as “payments” — that vary over the goods or services (such as education, social security, health care) being provided (McCaffery and Baron, forthcoming, and see also Eckel, Grossman & Johnston, 2004). In an instance of an isolation-type effect, we found that subjects’ preferences over the form of government fiscal mechanisms can trump their attitudes towards the redistributive effect of a tax or transfer program: that is, subjects would rather the government not use an income tax to raise money *or* to spend it, by means of allowing deductions through the tax (McCaffery and Baron, draft). But in the case of raising money, an income tax using progressive marginal rates is progressively redistributive (from rich to poor); in the case of spending money, such a tax is regressively redistributive (from poor to rich).

In a dramatic example of an isolation effect that we termed a *disaggregation bias*, subjects chose inconsistent tax systems when they were asked to design global tax systems as compared to when they were asked to design a part of the whole, with full degrees of freedom in the part under their control (McCaffery and Baron 2003). Once again we found that most subjects desired at least a moderately progressive tax system when asked to design a whole. But when half of the total revenue was collected by one tax, subjects underadjusted when choosing a tax system for the other half, although they were transparently able to effect the same overall level of redistribution or progressiveness as when they were choosing

a whole tax system. Subjects focused in on what they were being asked to choose, making it moderately progressive, although the sum tax system then varied. In one experiment, subjects ranked their preferences for three global tax systems as:

moderately progressive > steeply progressive > flat.

But when one-half of the tax was replaced by a flat-rate tax, and adjustments were made in the three taxes under comparison — such that the steeply progressive tax now had to feature a *negative* tax bracket at the lowest income range, to effect the same global distribution as before (as with the earned income tax credit under the current federal income tax (McCaffery 1999)) — subjects reversed their preferences, preferring:

moderately progressive > flat > steeply progressive.

It seems that the salience of the negative tax bracket changed preferences.

Based on these and related results, we expected that support for redistribution or the progressiveness of the tax system would vary with the system's form, contrary to the spirit of the two neoclassical welfare theorems. For example, subjects might consider a method of deductions from a progressive tax to be fair, until they consider the regressive nature of subtractions under a progressive marginal rate tax system (compare the *Schelling effect*, discussed in McCaffery and Baron, in press).¹ Subjects might also support more or less redistribution based on the nature and extent of the government's provision of goods and services. In sum, just as taxes can be "hidden," and different systems can evolve on account of this cognitive fact (McCaffery, 1994, McCaffery and Baron, draft), so too can progressiveness be "masked" or obscured from view by different fiscal mechanisms. We hypothesized that these "optics," as it were, would affect evaluation, leading to preference reversals and other inconsistencies.

Experiment 1: Tax Deductions

The main purpose of this experiment was to ask whether people favor tax deductions in part because they do not know, or do not consider, the fact that deductions relatively benefit the rich when marginal tax rates are progressive. We asked this by looking at attitudes toward deductions before and after a de-biasing manipulation, in which we explained the effect of deductions on rich vs. poor and then tested understanding.

¹ Schelling (1981) reported a classroom observation. He asked his students, given that a tax system would have child bonuses added to it, whether these bonuses should be higher for the rich or the poor. Students overwhelmingly answered for the poor. Schelling then pointed out that child bonuses presume a childless default rule; given a baseline of having children, the economically equivalent policy would be to institute childless penalties. Should such penalties be higher for rich or poor? Students quickly reversed their preferences, choosing higher childless penalties for the rich. We confirmed what we call the Schelling effect experimentally, a combination of *penalty aversion* and what we call a *progressivity illusion*. (McCaffery and Baron, in press.)

A second purpose of the experiment was to examine some of the determinants of favoring government provision of a good. We examined attitudes toward progressiveness, and perceived self-interest.

Method

The questionnaire was completed by 104 subjects, ages 17–70 (median 34), of whom 72% were female. The introduction read:

Paying for goods and services

Goods and services can be provided in different ways. Here are the ways we ask about:

- Government pays, and the money comes from income taxes. The rich pay a higher rate of income tax than the poor.
- Users pay, and government reimburses them.
- Users pay from their own funds.
- Users pay, but the cost is adjusted (by law) based on users' different abilities to pay.

Note that when the government pays for something, the money must come from somewhere. Please assume that, in these cases, taxes increase to cover the cost, by some percent of the tax, for everyone.

Here is a list of the goods and services we will ask you about:

Primary education: Grades 1-8

Secondary education: Grades 9-12

Theft insurance/compensation: Insurance or compensation for replacement cost of stolen or vandalized items

Basic telephone service

Mail delivery

Health insurance: Including what basic health-policies now cover

Social security: Basic income after retirement

Disability insurance: Replaces some income if you have to stop working because of disability

Unemployment insurance: Provides time-limited salary after job loss

Natural disaster insurance: Replacement costs after damage by wind, flood, or earthquake

We are interested in your views about how these goods should be provided and paid for.

The 10 cases were presented in an order randomized for each subject, and then they were repeated in the same order with different questions following each one. Here are the questions (with their names and comments in brackets, not seen by the subjects), using health insurance as an example:

[First cycle]

The good or service in question is HEALTH INSURANCE

Rate each of the following ways of providing this good or service. (Remember that government funds come from income tax revenues.)

[Each item answered on a 9-point scale from "Awful" to "Excellent"]

Government pays.

Users pay, and government reimburses them.

Users pay, but they can deduct the cost from their income taxes.

Users pay. No tax deduction.

Users pay, but the cost is adjusted (by law) based on ability to pay.

Where you live now, how is this good or service provided?

Government pays.

[Not at all Partially Fully]

Users pay, and government reimburses them.

Users pay, but they can deduct the cost from their income taxes.

Users pay. No tax deduction.

Users pay, but the cost is adjusted (by law) based on ability to pay.

[Second cycle]

The good or service in question is HEALTH INSURANCE

A tax deduction saves more money for those with high incomes than for those with low incomes. The lowest earners pay no tax and thus save nothing from a deduction. Their tax cannot be reduced because there isn't any. But a person who pays 30% in tax will save 30% of the cost of anything that is tax deductible.

Assume that the rich pay a higher percent tax than the poor on their income. What will happen if the government allows people to deduct the cost of health insurance from their taxable income?

The rich will save more money than the poor, from the deduction.

The explanation says that the poor will save as much money as the rich, or more money.

The explanation says that the rich will save more, but that won't happen in this case.

I don't know.

Rate each of the following ways of providing this good or service. (You may change the ratings you made before.)

Government pays.

Users pay, and government reimburses them.

Users pay, but they can deduct the cost from their income taxes.

Users pay. No tax deduction.

Users pay, but the cost is adjusted (by law) based on ability to pay.

How should the cost of health insurance be distributed between those with high income (the 'rich') and those with low income (the 'poor')?

The rich should pay more and the poor should pay less.

The rich and poor should contribute equally.

The poor should pay more and the rich should pay less.

Who benefits more from health insurance?

The rich The poor Both equally

How are you and those you care about affected by whether users pay for health insurance or whether the government pays?

Better off if users pay.

Better off if the government pays, even if taxes were raised to cover the extra cost.

It doesn't matter.

How are you and those you care about affected by the relative cost of health insurance for the rich and the poor?

Better off if the rich paid more and the poor paid less.

Better off if the cost were the same for everyone.

Better off if the poor paid more and the rich paid less.

It doesn't matter.

Results

Factor analysis of the responses to the items about methods of payment (government pays, reimbursement, deduction, no deduction, ability to pay) together with the attitude questions (rich should pay more, rich benefit more, better for me if rich pay more, better for me if government pays) across subjects and across variables indicated two main factors, one associated with whether the government should pay. Of the payment items, "government should pay" and "no deduction" were the highest loadings on the government payment factor; on the graduation factor, the highest loadings were reimbursement, deduction, and ability to pay. (The attitude items had loadings completely consistent with this account.)

Because our main interest was in the perception of the effect of deductions on graduation, we scored the deduction item by subtracting the mean rating of the reimbursement and ability-to-pay items from the rating for the deduction item, using a nine-point scale. This insures that the attitude toward deductions will, insofar as possible, be measured by comparison to alternatives (thus removing the effects of tendency to rate all items highly), and also without contamination by attitude toward government payment.

Considering only trials on which the test question in the second cycle (“[w]hat will happen if the government allows people to deduct . . .”) was answered correctly (“rich will save more”), and on which the subject said that the rich *should* pay more — that is, only subjects who both understood the instruction and could be expected to have it move them in a less favorable direction — the mean deduction score was .26 on the first pass (where one unit is a step on the 9-point scale where 9 is “excellent”) and $-.09$ on the second ($t_{76} = -2.56$, $p = 0.0126$, across subjects). Each of the ten items, as shown in Figure 1, showed movement in the expected direction, that is, against deductions after the manipulation.

In sum, the subjects seemed to be unfamiliar with the idea that deductions help the rich more than the poor. Many of them could understand this fact after a simple explanation. When they understood, they liked deductions less, especially if they were already inclined to support redistribution. But the movement against deductions was slight, and 28% of the responses to the test question indicated continued misunderstanding (with 43% of the subjects making more than one error).

In general, differences in attitudes toward payment mechanisms as a function of subjects (means across items) or items (means across subjects) were predicted by the relevant questions about perceived self-interest (whether graduation is better and whether it is better if the government pays), by preferences about whether the rich should pay more, or by perception of whether the rich benefit more. Tables 1 and 2 show correlations for subjects and items, respectively, with the level required for uncorrected one-tailed significance. The tables use the second round of presentations, post debiasing, where subjects were more likely to understand the effects of graduation.

Of greatest interest is the result in Table 2, where subjects favored deductions on those items for which they felt that payment should be graduated, as indicated by the correlation of .85. (In fact, in a multiple regression of favoring deductions on attitude toward graduation and belief that the government should pay, only the former contributed significantly across items.) The same attitude toward graduation correlated with the other items that, in fact, helped the poor more than the rich (government pays, reimburse, and ability to pay), and it correlated negatively with the most obviously regressive proposal, no deduction. Subjects seemed to understand that with no deductions, the cost (in percent and hence arguably disutility terms) would fall more heavily on those least able to pay. But subjects apparently did not realize that this was also true for deductions, despite our informative manipulation (the effect of which was, as noted above, small). Indeed, given that the deductions occurred in a progressive marginal tax rate system, the users pay with deduction condition was *worse* for redistribution, given the lost or foregone government revenue going disproportionately to the rich — a subject tested further in the next experiment. Subjects did not understand this.

Experiment 2: More about Deductions

The purpose of this experiment was to analyze further the attitude toward tax deductions. It addressed the following issues left open by the last experiment:

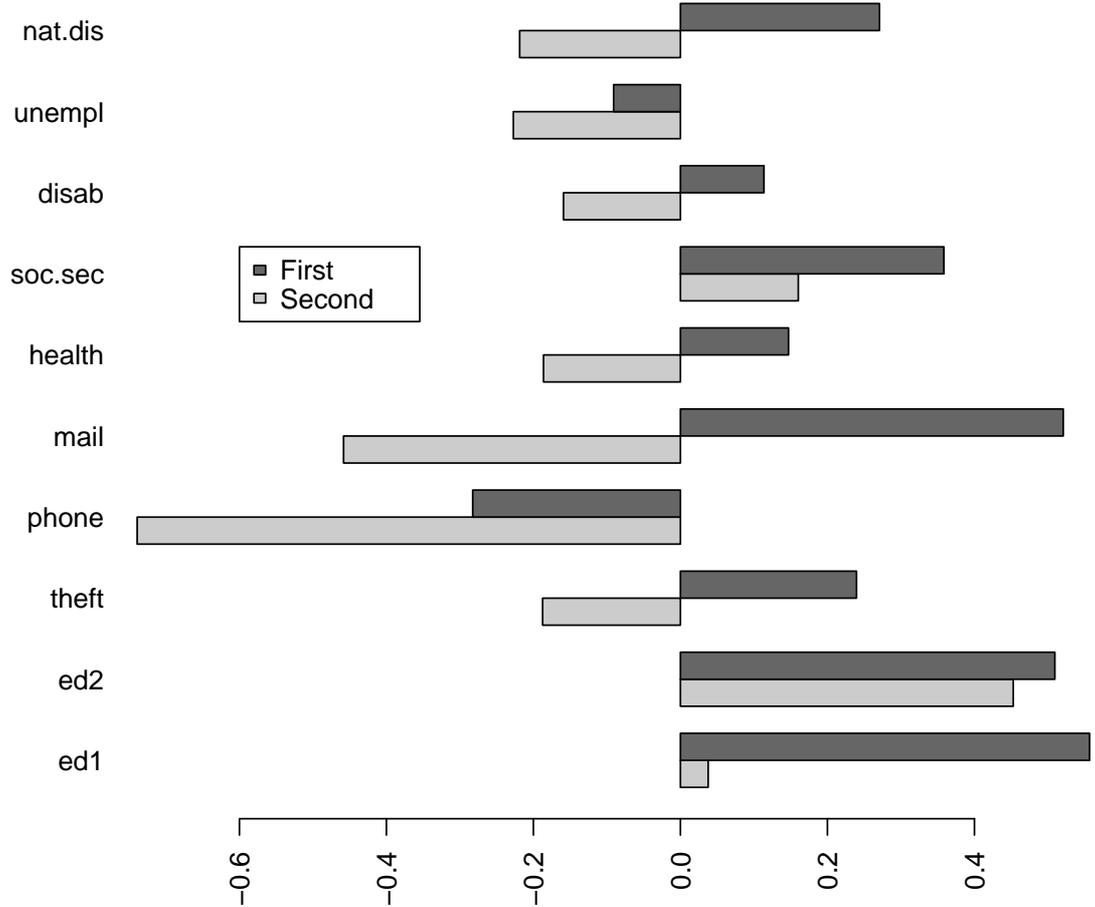


Figure 1: Mean ratings of deduction item minus mean of reimbursement and ability-to-pay items as a function of the item (on 9-point scale, for selected items).

Table 1: Correlations across the 104 subjects ($r > .16$ is “significant” at $p < .05$ one tailed)

	Government pays	Reimburse	Deduction	No deduction	Ability to pay
Should be graduated	0.18	-0.05	-0.10	-0.27	0.34
Rich benefit more	-0.11	-0.12	-0.14	-0.02	-0.02
Better if gov. pays	0.41	0.06	-0.23	-0.29	0.12
Better if rich pay more	0.01	0.09	0.05	-0.11	0.11

Table 2: Correlations across the 10 items ($r > .55$ is “significant” at $p < .05$ one tailed)

	Government pays	Reimburse	Deduction	No deduction	Ability to pay
Should be graduated	0.59	0.75	0.85	-0.78	0.85
Rich benefit more	-0.14	0.21	0.45	-0.23	0.40
Better if gov. pays	0.97	0.77	0.64	-0.91	0.59
Better if rich pay more	0.46	0.72	0.75	-0.69	0.88

- One source of support for deductions may have been the failure to notice their effect on the overall budget. Experiment 2 makes this effect explicit by showing it in a table. This is an experimental manipulation, so that we can assess the extent to which subjects fail to consider the cost of deductions, a kind of isolation effect.
- Experiment 1 explicitly said that taxes would have to be raised to cover the costs of any proposal, but subjects may have forgotten this, in part because it is unrealistic (taxes can be raised in the future, or spending cut, and so on). Subjects may have forgotten to consider tax increases relatively more in the case of deductions than in the case of subsidies. Experiment 2 drops this assumption.
- In Experiment 1, the attitude toward deductions was compared to several different alternatives. In Experiment 2, we used a single alternative, a rebate of 17% of the cost of the service. The rebate was more progressive than a deduction, assuming that it occurred within a progressive marginal rate tax system.

We asked a question about attitudes toward deduction four times: an initial baseline, a repetition of the baseline to control for order effects (Control), a condition showing the effect of a deduction on the government’s ability to spend (Loss, for lost revenues), and a condition showing both the loss and the effect on progressiveness (Progressiveness). The order of the last three conditions was randomized so that we could examine order effects.

Method

The experiment was completed by 135 subjects, ages 19–69 (median 38), 75% female. The introduction read:

Public vs. private goods

Some goods can be provided by government or by private companies. Here is a list of the ones we will ask you about. In some cases, government can provide a “basic” form of the good, and people who want more can buy more from private companies.

Basic housing

Basic food: Enough to live on

Trash collection

Tap water

Mail and basic telephone service

Basic health insurance: Including what basic health-policies now cover

Social security: Basic income after retirement

Disability insurance: Replaces a fixed amount of income if you have to quit work because of disability

Unemployment insurance: Provides time-limited salary, at a fixed rate, after job loss

Basic disaster insurance: Limited replacement costs after damage by wind, flood, terrorism, or earthquake

We are interested in why you favor government vs. private provision of these goods.

The issue here is how government should pay, if it pays at all. One way it can pay is by giving a **tax deduction**. This means that individual taxpayers pay for the good. Then the government allows the taxpayers to subtract the cost of the good from their taxable income. They do not pay tax on what they have paid for the good. They thus get a tax reduction in return for paying for the good.

Another way the government can help is simply to pay for the good, or some fraction of it.

The ten items were presented in four blocks. The cost of the various goods (used in questions presented below) was, respectively, \$1560, \$1040, \$520, \$260, \$520, \$1040, \$1040, \$260, \$520, and \$260. A random order of the cases was chosen for each subject and then used for each block. The questions for the first block were the same for all subjects. The order of the three remaining blocks was randomized for each subject. Here are the items, with names and comments in brackets, not seen by subjects, using tap water as an example.

[First set of questions: Baseline]

1. [Deduct]

Suppose that

- TAP WATER is now provided only by private companies.
- Individuals receive no help from government in paying for it.
- Current taxes now average 17%.

Which of these two proposals would you favor? (Click one.)

The government gives everyone a tax deduction for its cost.

The government gives everyone 17% of the cost.

These proposals are equal to me.

2. [GovHelp]

Under the same conditions, should the government help?

The government should not pay for tap water at all.

The government should pay some of the cost, but not all of it.

The government should pay the full cost.

3. [GovNow]

Which description is more accurate for you, where you live, now?

The government pays the full cost of tap water.

The government pays more than half of the cost, but not the full cost.

The government pays some of the cost, but less than half of it.

The government pays nothing.

4. [DeductNow]

Which description is most accurate for you, where you live, now?

The government gives a tax deduction for some (or all) payments for tap water.

Employers pay for tap water and employees do not pay taxes on the cost.

Employers pay for tap water and employees do pay taxes on the cost.

The government pays some or all of the cost out of general revenue.

The government pays nothing.

5. [Progressive]

In general, how should people pay for this good?

The poor should pay more dollars than the rich.

The rich and the poor should pay the same number of dollars. (Thus, the poor pay a higher percent of their income than the rich.)

The rich and the poor should pay the same percent of income. (Thus, the rich pay more dollars than the poor.)

The rich should pay a higher percent of their income than the poor.

[Control condition]

1. [The question was identical to Deduct]

[Loss condition]

1. Suppose that

- TAP WATER is now provided only by private companies.
- The cost to each taxpayer is now \$1560.
- Individuals receive no help from government in paying for it.
- Current taxes now average 17%.
- If the government helps, it has less money per person to spend on other goods. This is shown in the row called 'Loss per person.'

Here is a table showing what happens if the government helps by giving a tax deduction for the cost (\$1560).

Average income of all taxpayers	\$56,000
Average income tax now	\$10,400
Tax deduction	\$264
Loss per person	\$264

If the government helps by giving everyone 17% of the cost, the table would be exactly the same.

Which of these two proposals would you now favor? (Click one.)

The government gives everyone a tax deduction for what they pay for it.

The government gives everyone 17% of the cost.

These proposals are equal to me.

[Progressiveness condition]

1. Suppose that [items were identical to loss condition]

Here is a table showing what happens to three taxpayer groups, if the government helps by giving a tax deduction for the cost (\$1560).

Group:	Lowest 1/3	Middle 1/3	Highest 1/3
Avg. Income	\$24,000	\$48,000	\$96,000
Income tax now	\$3,200	\$8,000	\$20,000
Tax deduction	\$208	\$260	\$325
Loss per person	\$264	\$264	\$264

Here is a table showing what happens to three taxpayer groups, if the government helps by giving everyone 17% of the cost.

Group:	Lowest 1/3	Middle 1/3	Highest 1/3
Avg. Income	\$24,000	\$48,000	\$96,000
Income tax now	\$3,200	\$8,000	\$20,000
17% subsidy	\$264	\$264	\$264
Tax deduction	\$208	\$260	\$325
Loss per person	\$264	\$264	\$264

Which of these two proposals would you now favor? (Click one.)

The government gives everyone a tax deduction for what they pay for it.

The government gives everyone 17% of the cost.
 These proposals are equal to me.

Results

The main issue concerns the Deduct question (preference for deduction vs. 17% subsidy) in the four conditions: Baseline, Control, Loss, and Progressiveness. We first examine the four conditions without respect to order of conditions. (As we discuss later, answers to the Deduct question depended on the order of the Control, Loss, and Progressiveness conditions. However, the answer to Deduct did not depend significantly on whether it was presented immediately after the baseline or last, with Baseline as a covariate.)

Figure 2 shows the proportion of pro-deduction, anti-deduction, and neutral responses in each condition. Of primary interest is the mean attitude toward Deductions, coding Pro as 1 and Anti as -1 . Subjects were, on the whole, less pro-deduction in the Progressiveness condition than the Control condition ($t_{134} = 2.86, p = 0.0049$, across subjects). They were also less pro-deduction in the Loss condition than in the Control condition ($t_{134} = 1.98, p = 0.0495$), but the Loss and Progressiveness condition did not differ significantly in their mean responses, although there were more neutral responses. (Nor did the mean difference between these conditions depend on the answer to Progressive.) It seems that attitude toward deductions changes the most when subjects are reminded of the effect of deductions on government spending. Adding information on the effect on progressivity accentuates the effect.

As we noted, we found no effect of order of testing on attitude toward deductions across the last three conditions. In particular, the Control condition did not change as a function of whether it came before or after the Loss and Progressiveness conditions. This result suggests that any debiasing effect of the Loss and Progressiveness conditions did not carry over to the Control condition when it followed them. However, the change in attitude from Baseline to Progressive was small (see Figure 2), so an overall test is not the most sensitive one. To look for transfer effects, we regressed the subject's mean score in the Control condition on the Baseline condition (to account for general individual differences), the position of the Control condition (before, between, or after the Loss and Progressive conditions, as an ordered variable), the mean attitude on the Progressive and Loss conditions combined, and the interaction between this mean and position. The interaction term was significant (standardized coefficient .24, $t_{130} = 4.09, p = .0001$), indicating that the Control question is more strongly related to the other two questions when the Control question comes last than when it comes first. It thus seems that there is some transfer from the Progressiveness and Loss questions. Because of this, the comparison of the Baseline condition and the Progressiveness condition might provide a better indication of the magnitude of the effect of the Progressiveness manipulation, a larger difference than that between Progressiveness and Control.

Notice in Figure 2 that the Loss condition led to more neutral responses. The wording of this item led subjects to think that the two options had identical effects. The effects

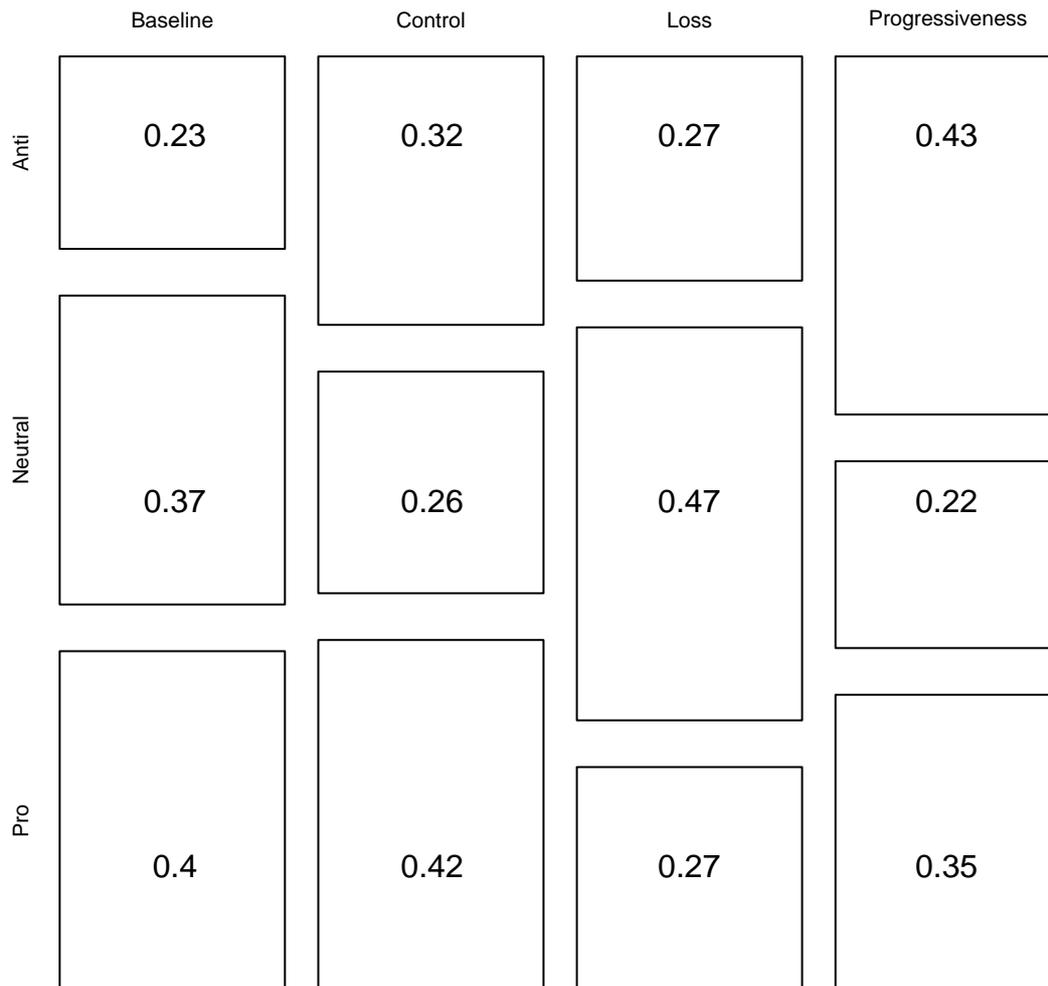


Figure 2: Proportion of responses to Deduct question by condition

were not in fact identical because they differed for different income groups, so our manipulation was misleading by virtue of being incomplete. This uncertainty, however, does not weaken the main conclusion: without prompting, people ignore the progressiveness effects of deductions. This conclusion is supported by the comparison of the Progressiveness and Control conditions.

Experiment 3: Cutting Government Services

Experiments 1 and 2, on deductions, drew attention to the fact that government fiscal systems both provide goods—that is, spend—and raise money. All actions have benefits and burdens. There is no reason why the distributional effects of the two sides have to match. Thus in the Progressiveness condition in Experiment 2, subjects could see that the lost revenue from a tax deduction (resulting in reduced provision of public goods or services)—the burden—could be felt lump sum, while the increased private wealth effected by the deduction could fall disproportionately to the high income. Experiments 3 and 4 generalize this to the setting of the “privatization” of otherwise public goods.

When governments raise taxes by a progressive tax scheme and then pay for services that cost the same to rich and poor alike, the net effect is to redistribute income, a “cross-subsidy” through the provision of the good. The rich pay more, the poor less, both income classes benefiting the same. Government provision in such a case saves money for the poor, assuming that the poor would buy the services if the government did not provide them, and the government provides them with reasonable efficiency. This is a paradigm example of the “bundling” together of two distinct government actions, allocative (providing the good or service in the first place) and redistributive. Were government simply to “privatize” or otherwise cut government services, without continuing the redistribution effected through the tax and spend program, a greater burden would fall on those who are relatively poor. Logically, the government can continue to redistribute resources through the tax system with or without the provision of the good or service. But will subjects support a consistent level of redistribution, however, independent of government provision of goods?

To examine the effects of cutting services on support for the overall progressiveness of the public finance system, we asked subjects to imagine that their national government could provide five basic services, spending equal amounts on each: defense, education, health care, social security, and everything else. We presented 16 cases in which government provided all possible combinations of the first four. In each case, we asked the subjects to choose the fairest level of progressiveness, including the possible use of negative taxes. Consistent with our prior research on *disaggregation effects* (McCaffery and Baron 2003), we anticipated that subjects would not maintain the same level of redistribution — would not fully take into account or integrate the effects of the service cuts on household welfare — choosing less overall redistribution with fewer services.

Method

Seventy-eight subjects completed the questionnaire. Their ages ranged from 19 to 64 (median 40), and 18% were male.

The questionnaire began:

Tax and spend

We are interested in how you think government expenditures should affect the distribution of income taxes.

All governments spend money enforcing their own laws and regulations. Many other expenses are optional. We are concerned here with five categories of spending: national defense; basic health care; basic education; social security; and “all other functions” (law enforcement, etc.). For purposes of these questions, assume:

- The government spends the same amount on each of the five categories.
- The national government pays for all of these services, if any government does.
- The government contributes to a regional defense force (in the “other functions” category), so it does not absolutely need national defense.
- Otherwise, the country is like your own country in standard of living.
- The government pays \$2,000 each year for each good or service for each household.
- If the government does not provide health care, education, or social security, people must pay at least that much on their own.

Each screen shows several possible distributions of income tax among three groups of households: low, middle, and high pre-tax household income. (We exclude a small proportion of households that have little or no income, pay no tax, and receive some sort of assistance.) The three groups have the same proportion of all the pre-tax income. You see the average tax for each group. Here are some statistics for the three groups (approximately correct for the U.S.):

	Low	Middle	High
Percent of all income	33.3%	33.3%	33.3%
Percent of households in group	60%	30%	10%
Average annual household income	\$20,000	\$50,000	\$130,000

Government can re-distribute money from those with high incomes to those with low incomes. It can have a negative tax, in which it gives money to some people in order to accomplish this re-distribution. Negative taxes are shown in red.

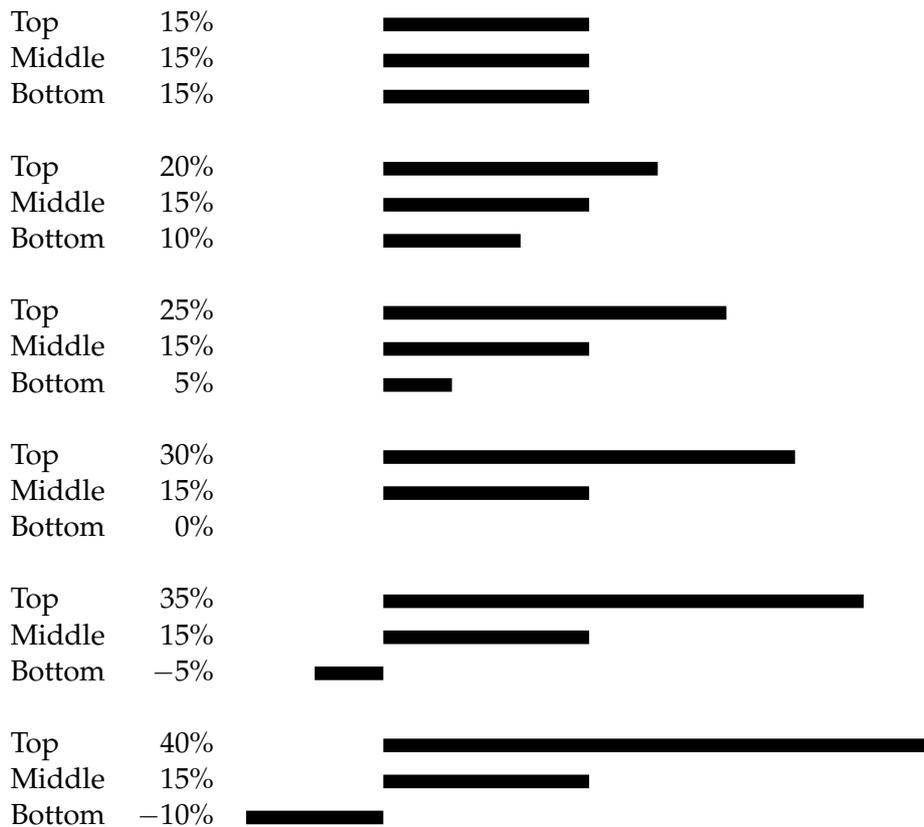
A typical item was:

In this case, suppose the government provides:
 NATIONAL DEFENSE
 BASIC EDUCATION
 ALL OTHER FUNCTIONS OF GOVERNMENT
 But the government does not provide:
 BASIC HEALTH CARE
 SOCIAL SECURITY

When the government does not provide a service that people must pay for themselves, the extra annual household cost is \$2,000, regardless of income. How much extra will each household have to pay per year, on the average, because of what the government does not provide in this case? (Pick the closest.)

\$0 \$2,000 \$4,000 \$6,000

Which of the following is the fairest distribution of income taxes among the three income groups, in this case? [A button was provided next to each of the six distributions.]



Reminder: The average annual pre-tax incomes of the three groups are, respectively, \$130,000, \$50,000, and \$20,000.

The first question, about the extra cost per household, was designed to encourage attention to this cost, but it could also serve as a check on understanding. The right answer was \$2,000 times the number of non-defense cuts.

The basic case (with no cuts, such that all five sets of services were provided by the government) had a 25% tax for each group in the least progressive (top) distribution. When cuts were made, the tax for each group was reduced by 5% of income for each cut. Thus, the same response would represent the same degree of progressiveness, regardless of the cuts, *if* each taxpayer would have to pay 5% of income for the missing service. But the cost in *dollars* of the services described is independent of income, so the poor would have to pay more, and the rich less, than implied by this assumption.

In the case of defense, a cut leads to an increase in income as a result of tax reduction with no additional costs. The average increase in income is 20% of the baseline tax paid, or 5% of income. We thus assume, for example, that taxes of 25%, 15% and 5% for three brackets, respectively, is just as progressive as taxes of 20%, 10%, and 0%, if the difference is a result of a cut in defense.

The order of the 16 conditions (which differed according to what the government provided) was randomized for each subject.

Results

Six subjects always chose the least progressive distribution, which was equal percentage rates for all three groups, and 2 subjects always chose the most progressive. The mean choice was 3.42, on a 1–6 scale with 6 being the most progressive. This amounts to a difference (in absolute dollars) of 24.2% between the high and low income groups.

We calculated, for each subject, the mean effect of each cut on progressiveness, first ignoring the effect of cuts on out-of-pocket cost. The mean effects (in the change in percent difference between high and low groups) were, respectively, 1.1% for defense ($t_{77} = 1.70$, $p = 0.0931$), -0.1% for health care (n.s.), 0.4% for education (n.s.), and -0.4% for social security (n.s.), where a positive effect indicates less progressiveness with the cut than without it. The mean of these effects was not significantly positive, and the four services were not significantly different. Thus subjects maintained the same degree of progressiveness *without taking into account the effect of the cuts on out-of-pocket cost*. This is a clear example of the isolation effect.

But cuts do affect out-of-pocket costs, at least for three of the goods of interest: health care, education, and social security, both in the experiment and in the real world. A more relevant analysis of the data therefore includes the effects of these cuts in public services on net — after public tax and spending — household welfare. Do subjects use the tax system to compensate for the effects of public spending cuts? If so, they would *increase* the progressiveness of taxes for these three goods. An attempt to make such compensation

may account for the difference between the three goods and defense. But do subjects compensate enough? Or do they show an isolation effect or disaggregation bias: maintaining a constant level of redistribution or progressiveness in the tax system alone, but supporting less global redistribution on account of their failures to take into account (or take fully into account) changes on the expenditure side of the government ledger?

The information we provided allows us to estimate the effects of cuts on out-of-pocket costs as a percent of income. The \$2,000 figure we gave for the effect of each cut is, respectively, 10%, 4%, and 1.5% of the three income levels (\$20,000, \$50,000, \$130,000). To achieve a redistribution between the high and low groups of 10% vs. 1.5% would require almost two steps on our graduation scale. Thus, to be conservative, we assumed one step for inferential statistics. Of course, we could not expect subjects to make these precise calculations, but we do think that the numbers we explicitly provided were plausible, so that subjects might have made some attempt to compensate for cost effects even without calculating, and evidently some did so, given the difference with national defense.

All three of the cost-yielding cuts (health care, education, social security) yielded corrections far less than the one step we conservatively assumed would be the required (minimum $t_{=77} = 12.45$, $p = .0000$). Hence, it seems clear that, while some subjects attempted to offset the cost-increasing effects of cuts, on average the attempt falls far short of what is needed. (Note, however, that the response options we provided did not always allow the subjects to fully compensate for cuts. Still, only 12 subjects favored the maximum graduation response when all three cuts were made.)

Looking at the bottom line, Table 3 shows the mean response of subjects, using the same type of graph they saw, in the absence of any cuts and in the presence of three cuts. The lowest panel represents the results of including out-of-pocket costs. Table 3 gives an excellent look at the isolation effect or disaggregation bias, playing out in a unified tax and spending system. Subjects preferred at least moderate progressivity in the baseline, global condition, with government provision of all five sets of goods and services. With three major private-cost items removed from the mix of public goods, subjects continued to choose a tax system reflecting moderate progressivity, even accepting a negative tax bracket for the poor, as the second panel shows. But when realistic private replacement costs were built back in, showing a global tax and out-of-pocket effect, the overall system now looks regressive, in the bottom panel. Compared to the subjects' own chosen baseline, the bottom-line reflects a steep cut in costs (taxes plus out-of-pocket) for the upper income level, a slight drop for the middle income level, and a dramatic (230%) *rise* in effective burdens on the lowest income level. By focusing on the "optics" of taxes alone, or disproportionately, a preference reversal in the bottom-line effects — what really matters — followed.

Aversion to progressivity cannot explain the results, given that subjects (on average) consistently chose progressive taxes. Nor can ignorance of the financial effects of public spending cuts explain the results. Subjects made many errors on the test question about the extra cost per household caused by cuts. But the most common error seemed to be to simply count the number of cuts, including defense cuts, which (by specification) should

Table 3: Mean responses and inferred responses for presence and absence of health care, education, and social security, Experiment 3.

No cuts		
Top	33.6%	
Middle	22.5%	
Bottom	11.4%	
Three cuts, raw responses		
Top	18.9%	
Middle	7.5%	
Bottom	-3.9%	
Three cuts, responses plus out-of-pocket cost		
Top	23.4%	
Middle	19.5%	
Bottom	26.1%	

have had no effect on household spending. Yet 95% of the subjects gave the correct answer *or higher* for the effects of cuts on household finances. Moreover, the mean answer to the test question was 2.53 on a scale from 1 to 4, where the mean correct answer is 2.5. Subjects did not underestimate the effects of public spending cuts on net household costs. The results were essentially unchanged when we examined only the subjects who estimated cost correctly or overestimated it, on the average.

What *can* explain the results is the disaggregation bias or isolation effect. Subjects looked only (or primarily) at the tax system when adjusting it. They did not adequately factor in the effect of public spending cuts. The result is that effective progressivity decreased as the number of cuts increased — disappearing altogether with enough “downsizing” of government.

Experiment 4: More cuts

We carried out one final study, extending the last one. We made several changes:

- We eliminated defense, using only the three cuts that would require increased spending. This meant that we could present all possible combination in eight items. We presented these cases 4 times each, in a random order chosen for each subject, for a total of 32 items.
- In half of the items, we presented taxes and incomes as dollars rather than percent. This allowed us to replicate (with a different method) our finding that preferred graduation is higher when taxes are in percent (McCaffery & Baron, 2003).

- In half of the items, we presented a display indicating the extra costs required, in the same format as that used to display taxes. We thought that this might serve as a de-biasing method. (It did not.)
- We extended the list of response options but asked only the fairness question for each item.
 - We gave options in which the lowest tax was 0, to test the possibility that some people resist negative taxes.
 - We included two different “flat tax” options (constant percent, constant dollars). Subjects in Experiment 3 and other studies we have done often comment, after the study, that they favor a “flat tax,” but we have been unsure what they meant.
 - We included options that would look more evenly graduated when presented in dollars. We did this so that subjects would not reject a graduated option just because it did not look evenly spaced.
- We made other minor changes, mostly to make our calculations easier.

Method

Seventy-nine subjects completed the questionnaire. Their ages ranged from 24 to 58 (median 39), and 24% were male.

The instructions and questions were modified somewhat from Experiment 3. Of most importance, the introduction explained: “Taxes can be ‘flat,’ so that everyone pays the same percent, or the same dollars. Taxes can be ‘graduated,’ so that those with low incomes pay less and those with high incomes pay more.” It also stated, “If the government does not provide health care, education, or social security, each taxpayer must pay an average of \$2,500 per year to replace the service that government provided.” The mean incomes of the three groups were given as \$20,000, \$50,000, and \$125,000, so that they would differ by multiplicative constant.

We provided explanations of the response options and grouped them, so that a typical example (for the dollars version, with two cuts) was as follows:

Here is the extra cost that each of the three income groups must pay as a result of what the government does not provide.

Top	\$5,000	■
Middle	\$5,000	■
Bottom	\$5,000	■

Which of the following is the fairest distribution of income taxes among the three income groups, in this case? (All raise the same total amount.)

Flat tax in dollars (everyone pays the same number of dollars):

Top	\$5,769	████████
Middle	\$5,769	████████
Bottom	\$5,769	████████

Flat tax in percent (everyone pays the same percent):

Top	\$18,750	████████████████████
Middle	\$7,500	██████████
Bottom	\$3,000	████

Graduated, varying in amount of graduation and position of middle group:

Top	\$25,000	████████████████████████
Middle	\$7,500	██████████
Bottom	\$2,000	██

Top	\$31,250	████████████████████████████
Middle	\$7,500	██████████
Bottom	\$1,000	█

Top	\$27,500	████████████████████████
Middle	\$10,500	██████████
Bottom	\$4,000	████

Top	\$37,500	████████████████████████████████
Middle	\$7,500	██████████
Bottom	\$0	

Top	\$33,750	████████████████████████████
Middle	\$10,500	██████████
Bottom	-\$600	█

Top	\$43,750	████████████████████████████████████
Middle	\$7,500	██████████
Bottom	-\$1,000	█

Top	\$40,000	████████████████████████████████
Middle	\$10,500	██████████
Bottom	-\$1,600	█

Graduated, lowest group pays zero. (Note: These may be identical to distributions listed above. If so, they are equivalent answers.)

Top	\$37,500	
Middle	\$7,500	
Bottom	\$0	
Top	\$31,875	
Middle	\$9,750	
Bottom	\$0	

When the display of amounts was not presented, subjects were simply reminded of the three average income levels. In the percent condition, the amounts were identical but expressed as a percent of income.

Results

Graduation: sensitivity to units and cuts

First consider the results concerning graduation, which replicated and extended the results of Experiment 3. We measured graduation as the difference in percent between the high and low income groups for the chosen option. Graduation was higher when the units were percent (21.9%) than when they were dollars (18.2%; $t_{78} = 4.42$, $p = 0.0000$); we confirmed the *metric effect* that we had found earlier (McCaffery & Baron, 2003 and forthcoming). The display of costs had no effect on graduation and did not interact with dollars vs. percent.

As the number of cuts increased, graduation *decreased* ($t_{78} = 3.47$, $p = 0.0009$, for the slope of the function relating graduation to number of cuts). The means of graduation were 21.2%, 21.2%, 19.0%, and 16.0%, respectively, for 0, 1, 2 or 3 cuts. This effect resulted from the fact that graduation of the last two options (those with the low group getting exactly zero tax) is lower when the number of cuts is greater. Subjects who chose those options apparently caused this effect. (We discuss this later.) When these options (and the first option, a flat dollar tax, which would have the reverse effect) were assigned a constant graduation level, graduation showed a small and non-significant increase with more cuts.

In sum, the results were consistent with those of Experiment 3, in which subjects essentially ignored the distributional effects of the cuts. Table 4 shows the mean responses for the conditions shown in Table 3, that is, the percent conditions. Once again, the implied or bottom-line tax and transfer rate is much more regressive with greater cuts, and the effect is just as great with taxes expressed in dollars.

For three pairs of response options, subjects could choose between distributions that were evenly spaced by percent and distributions in which the tax rate was higher for the middle-income group (and lower for the other two groups), giving an impression closer to equal spacing in the dollar-units condition. On the whole, subjects preferred the lower rate for the middle group, choosing it in 29% of all items, vs. 13% for the higher rate ($t_{78} = 4.78$, $p = 0.0000$, across subjects). This result suggests a “soak the rich” (or a “help the middle”?) attitude, rather than a “help the poor” (at the expense of the middle?) one. However, preference for the higher rate was greater in the dollar-units condition, where it appeared to be more equidistant from the rate for the other two groups: in dollar-units, low middle-

Table 4: Mean responses and inferred responses for presence and absence of health care, education, and social security, Experiment 4.

No cuts		
Top	36.7%	
Middle	25.0%	
Bottom	13.3%	
Three cuts, raw responses		
Top	18.5%	
Middle	10%	
Bottom	1.5%	
Three cuts, responses plus out-of-pocket cost		
Top	24.5%	
Middle	25.0%	
Bottom	39.0%	

group rate 21% and high rate 17%; in percent units, low 36% and high 9%; $t_{78} = 5.59$, $p = 0.0000$, for the effect of units). This is another appearance of the metric effect.

Strategies: flat tax; graduation; zero

We can examine the extent to which subjects followed various consistent strategies: flat percentage; flat dollars; graduation without adjustment for costs; graduation with compensation; the maximum possible graduation; negative taxes; and zero for the lowest bracket. The flat dollar option was chosen in 3.6% of the cases, while the flat percent option was chosen in 26.9%. It appears that subjects who favor flat taxes favor a flat percentage tax. The 3.6% could result from errors, although a few people might actually favor this. Henceforth we combine these two categories.

The last two options, which fix the lowest bracket at 0%, allow us to assess opposition to negative taxes. Subjects who generally favor graduation but oppose negative taxation will choose the "zero option" over one of the more graduated options. When there were no cuts, the zero option was the most graduated of all.

Table 5 shows the use of each strategy, including those who use it consistently and avoid it consistently, and the mean percent use for those who use it sometimes. (The negative tax is based on the 7/8 of the cases in which it was possible.) The category called "Maximum grad., no negative" is the most graduated option other than any option that involves a negative tax. The last two strategies were simply the proportion of subjects who increased graduation with more cuts and those who did not do so; this classification overlaps with the others, as it is not based on individual responses but on the overall pattern.

To determine whether subjects try to avoid negative taxes, we calculated for each sub-

Strategy	Percent who never choose	Percent who always choose	Mean % for the rest
Flat tax	39.0	12.7	30.0
Zero for low income	26.6	7.6	30.0
Negative tax	51.9	0	14.2
Maximum graduation	31.6	2.5	26.7
Maximum grad., no negative	26.6	2.5	24.7
Graduation with adjustment	N.A.	N.A.	30.4
Graduation, no adjustment	N.A.	N.A.	69.6

Table 5: Strategy use, Experiment 4.

ject the mean graduation for cases with no cuts, 1 cut, 2 cuts, and 3 cuts. We could thus examine graduation as a function of the number of cuts. Then we selected for further examination those subjects with a maximum graduation (across these four categories) exceeding 30%. These 34 subjects are the ones most likely to choose negative taxes with more than one cut, since a graduation rate of 30% would imply negative taxes in this case. If they tried to avoid negative taxes, however, their choices would become less graduated as the number of cuts increased. Specifically, the maximum graduation possible without negative taxes with 0, 1, 2, or 3 cuts, respectively, is 50%, 40%, 30%, and 20%. If the subject does not try to avoid negative taxes, the maximum possible graduation is 50%, 40%, 40%, and 40%.

Figure 3 shows the results for these subjects, one line per subject. The solid lines represent the subjects whose graduation decreased most sharply as the number of cuts increased. The circles represent the maximum graduation rates without negative taxes; the diamonds, the maximum possible graduation rate among the response options. The group of solid lines that fall near the circles are from subjects who approximated the strategy of maximizing graduation while avoiding negative taxes. Note that many of the other subjects, the dashed lines, had the opposite pattern, increasing graduation as the number of cuts increased (to the left), thus compensating for the effect of the cuts, but this required negative taxes, which resulted in graduation rates higher than the circles.

To test for an effect of avoiding negative taxes, we correlated the slope of the functions shown in Figure 3 (graduation as a function of number of cuts), with each subject's maximum graduation, across all subjects (excluding the leftmost point because the maximum possible graduation was higher than the other three points). Subjects with higher maximum graduation rates had a more negative slope (that is, high on the left, low on the right, in Figure 3); the correlation was $-.48$ ($p = .0000$), using all subjects. This is the opposite from what we would expect if subjects compensated for the effect of the cuts. In sum, part of the reason that people do not compensate for the effects of cuts is that they want to avoid negative taxes.

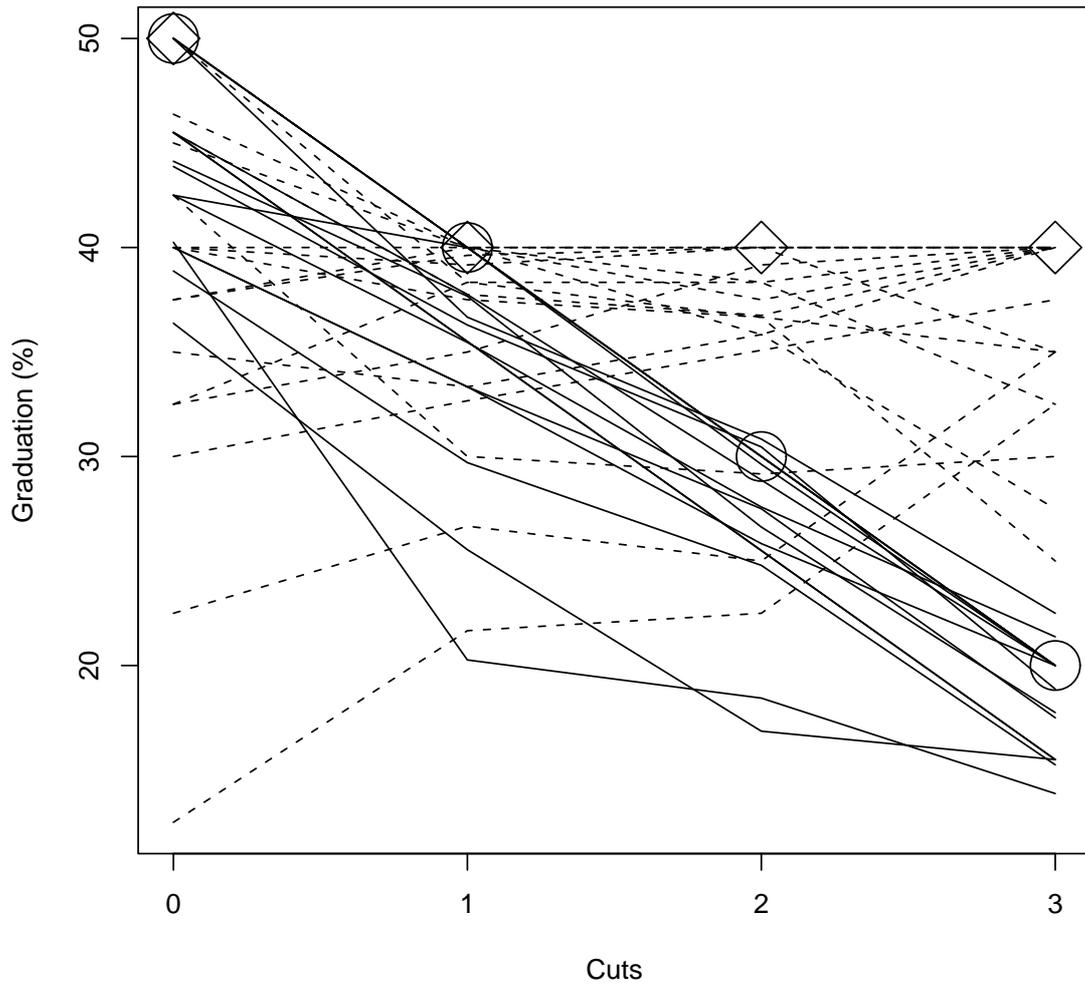


Figure 3: Graduation as a function of number of cuts for subjects with a maximum graduation greater than 30%. (The circles represents the highest graduation possible without negative taxes, and the diamonds represent the most graduated response including negative taxes.)

Conclusion

Subjects who favor at least moderate progressivity in the allocation of public benefits and burdens — which seems to be the norm across time and cultures (Lewis 1978, Hite and Roberts 1991, McCaffery and Baron 2003) — can be easily fooled by the form of public finance mechanisms.

We found for example that subjects think too little about the effects of tax deductions on progressiveness. When they are shown these effects, they are less likely to support the use of tax deductions, but only moderately so. Worse, they seem to prefer a program of private provision with tax deductions to a system of private provision without tax deductions, even when the fact of a progressive marginal rate tax system means that the deduction element, standing alone, subtracts from the redistributive effect of government action.

In the next two experiments, we found that people are also insufficiently aware of the effects of privatization or (in the case of defense) abandonment of government services on progressiveness. Most people want some progressiveness in the tax system: even those who favor a “flat tax” (in comments after Experiment 3) seem to mean a flat percentage rather than a flat dollar amount. Yet their failure to think through the effects of changes on progressiveness leads them to support proposals that undercut their own values.

The same mechanisms could of course work in the other direction. People who do not favor progressive taxation could fail to see the effect of changes that increase progressiveness, such as increased government funding of health care. In all cases, however, the puzzle is the same: that the *form* of the provision affects preference reversals in attitudes about the *substance* of government fiscal actions.

We have focused here on the conflict between subjects’ judgments and their own apparent values. But distributional effects are also real. Optimal taxation models typically take into account the declining marginal utility of money as well as the incentive effects of taxation (Mirrlees, 1971). Under most welfarist models, some redistribution is beneficial, and policymakers are increasingly looking to public finance mechanisms to provide it (Kaplow and Shavell, 2000). The kind of reasoning we have found here can pit equity (or redistributive) effects against efficiency (or allocative) ones, and can lead to truly harmful effects on the poor, beyond the conflict with attitudes.

The common psychological mechanism is one we have explored elsewhere (McCaffery and Baron, 2003 and in press): a disaggregation bias or isolation effect. In a nutshell, people do not think ahead, they do not integrate logically connected matters to form consistent judgments as to bottom-line effects and realities. In studies of the effects of marriage, for example, we have found that people focus on the question they are answering, without thinking of effects of their answer on other things they care about, such as the interactive effects of marriage bonuses, penalties, and tax graduation. Likewise, people prefer hidden taxes such as business taxes in large part because they do not think sufficiently about the necessity for the money to come from somewhere, and about where it will come from. Here we found that people under-account for the effects of government cuts on household finances when making judgements about the appropriate remaining tax system.

The problem of failing to think more than one step ahead — failing to put two and two together — is not of course limited to taxes and public policy. But especially given the absence of any natural corrective to these tendencies in the nonmarket sectors of public finance, they are likely, if left unchecked, to wreck considerable havoc there.

References

- Babcock, L., Gelfand M, Small D, Stayn H. (2003). The propensity to initiate negotiations: toward a broader understanding of negotiation behavior. Manuscript, Carnegie-Mellon University.
- Baron, J. (2000). *Thinking and deciding* (3d ed.). New York: Cambridge University Press.
- Camerer, C. F. (2000). Prospect theory in the wild: Evidence from the field. In Kahneman, D., & Tversky, A., eds. *Choices, values, and frames* New York: Cambridge University Press (and Russell Sage Foundation).
- Eckel, C., Grossman, P., & Johnston, M.R. (2004). An experimental test of the crowding-out hypothesis. Work in progress on file with authors.
- Hite, P.A. and Roberts, M.L. (1991). An experimental investigation of taxpayer judgements on rate structure in the individual income tax system. *Journal of the American Taxation Association*, 13, 47-63.
- Kaplow, Louis and Shavell, Steven. (2000). *Fairness v. welfare*. Cambridge, Mass: Harvard University Press
- Lewis, A. (1978). Perceptions of tax rates. *British Tax Review*, 358–66.
- McCaffery, E.J. (1994). Cognitive theory and tax. *UCLA Law Review*, 41, 1861-1947.
- McCaffery, E. J. (1999). The burdens of benefits. *Villanova Law Review*, 44, 445–493.
- McCaffery, E.J., and Baron, J. (2003). The Humpty Dumpty blues: disaggregation bias in the evaluation of tax systems. *Organizational Behavior and Human Decision Processes*, 91, 230-242.
- McCaffery, E.J. and Baron, J. (in press). Framing the family: evaluation of tax policies involving household composition. *Journal of Economic Psychology*.
- McCaffery, E.J. and Baron, J. (forthcoming). Heuristics and biases in thinking about tax. *National Tax Association Papers & Proceedings*.
- McCaffery, E.J. and Baron, J. (draft). Thinking about tax.
- Mirrlees, James A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38: 175-208.
- Musgrave, Richard. (1959). *The theory of public finance*. New York: McGraw Hill.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, Mass: Harvard University Press.
- Read, D., Loewenstein, G., & Rabin, M. (1999). Choice bracketing. *Journal of Risk and Uncertainty*, 19, 171-197.
- Rosen, H. (2001). *Public Finance*(6th ed.). Chicago: McGraw Hill/Irwin.
- Schelling, T. C. (1981). Economic reasoning and the ethics of policy. *Public Interest*, 63, 37–61.