

Starving the beast: The psychology of budget deficits.

Jonathan Baron*

University of Pennsylvania

Edward J. McCaffery

University of Southern California Law School

California Institute of Technology

September 3, 2004

Abstract

Many opponents of big government favor a strategy of “starving the beast,” cutting taxes today with the expectation that spending cuts will follow tomorrow. Why might such a strategy work? Various heuristics and biases help to explain how it can. In two experiments conducted on the World Wide Web, subjects chose general levels of taxation and public spending from various hypothetical starting points. Subjects wanted to reduce both taxes and spending, preferring balanced budgets and even surpluses to deficits. When asked about *specific* spending cuts, however, subjects showed a marked reluctance to make cuts, leading to deficits. Subjects also showed an anchor and underadjustment bias, changing their responses in light of various baselines, and failing to completely close existing deficits. The “starving the beast” phenomenon, by pairing specific tax cuts with the general, abstract idea of spending cuts, can thus succeed in a population preferring fiscal balance. Once the deficit is created, it will likely persist, influencing future policy preferences.

*Email: baron@psych.upenn.edu. This work was supported by NSF grant SES 02-13409. We thank Judith Baron for some of the ideas.

Introduction and Hypotheses

Politicians, social scientists, and citizens disagree sharply about the appropriate size of government. The issue captures perhaps the major fault line between parties in two-party democracies. Some argue that big government is bad, but that people can be led to support it because they do not think about long-term issues and thus desire overly generous present programs (Buchanan & Tullock, 1962). Others argue that government is if anything too small, because of pressure for low taxes, which appeal to citizens on the basis of narrow — and myopic — self-interest.

A common element between the two extremes is that there is a disconnect between the present and the future: an at least implicit understanding that citizens will fail to integrate their beliefs and actions over time. Anti-government partisans fear that citizens will want programs now, neglecting their long-term costs, and then will be reluctant to cut these programs later: social security and medicare in the United States are leading case studies for such critics. Pro-government partisans fear that citizens will support tax cuts now, ignoring the long-term effects of any resulting deficit (or diminished surplus) on the ability of the government to continue to provide public goods and services in the future. Both of these sets of attitudes stand in stark contrast to the “rational choice” or “rational expectations” model of politics, where citizens properly integrate their actions over time. Thus, Barro (1974) has argued that government deficits may not even matter, because forward-looking citizens in an overlapping generations framework will rationally save today in anticipation of increased taxes tomorrow; conversely, surpluses today can lead to greater private debt in anticipation of lower taxes tomorrow.

Standard findings in cognitive psychology, most notably prospect theory and the endowment effect (Kahneman & Tversky 1979, Kahneman, Knetsch & Thaler 1991), support the popular understanding that timing matters. Once a government program is in place, it will become part of the status quo, and can be hard to cut. Thus the thumb is on the side of continued government growth. On the other hand, citizens are averse to taxes, a phenomenon that itself has cognitive

psychological dimensions. People react disproportionately to salient taxes and fail to consider the off-setting benefits of government programs (McCaffery & Baron, 2004); people are also likely to code a tax increase as a loss, making it hard ever to raise taxes.

A potentially psychologically-savvy political strategy used by those who favor smaller government has come to be called “starve the beast,” a term usually attributed to David Stockman, the budget director in U.S. President Ronald Reagan’s administration. The idea is to cut taxes before cutting spending, then use the resulting deficit as a political argument to reduce spending, or to reject new spending. Most commentators agree that this strategy has been used by both Reagan and the current U.S. president George W. Bush. In both cases, large deficits have resulted from fiscal policies. Although spending was not cut concurrently with taxes, government may have grown less than it would have without the tax cuts, because the baseline for future judgments was changed.

Can the starve the beast strategy work? If citizens were fully rational and consistent in the manner of Barro, such budgetary and political manipulations would not matter: private actions would perfectly counterbalance the public ones. Or citizens might be fully rational, anticipating the ultimate tax increases or spending cuts that deficits imply, but distrust their own discipline. In such a case, citizens would rationally want balanced budgets today as a self-commitment mechanism (Schelling, 1984), and politicians who cut taxes without reducing spending would lose public support. In these rational choice settings, either deficits do not matter or they cannot get started, and there is no room for the starve the beast technique to have real effects.

But we now have abundant evidence that ordinary people are not always rational, in the simple sense of having stable consistent preferences over choices and outcomes. We test for two hypotheses about why people might accept tax cuts in the absence of spending cuts.

One, people might think excessively or even exclusively about the short term. They neglect the fact that deficits must be covered in the future. More generally, they engage in a kind of optimism bias (Kahneman & Lovallo, 1993; Camerer & Lovallo, 1999): believing matters will all work out

in the end. In this case, they would favor budget deficits in the short term and respond differently when asked about the future than when asked about the present.

Two, people might think differently about tax cuts and spending cuts today because of their levels of salience or abstraction. Specifically, people are averse to taxes (Fennell & Fennell, 2003), and so support specific tax cuts. They actually favor fiscal discipline, that is, balanced budgets, but are thinking about spending cuts *in the abstract* when they support the tax cuts. Then, later, when the rubber hits the road, so to speak, the same people generally oppose cuts in *particular* government programs. This is analogous to the “identified victim” effect (Jenni & Loewenstein, 1997; Small & Loewenstein, 2003; Kogut & Ritov, 2004). Deficits result. We test this by asking about spending in the abstract and cuts in particular programs.

We also ask why budget deficits might be difficult to remove once they are created. We test the hypothesis that people anchor on current levels of taxation and government spending, so their preferred levels of each do not go far enough to remove the deficit, once a deficit exists.

General Experimental Method

We report experiments carried out on the World Wide Web. The subjects were drawn from a panel of about 1500 people who have done other experiments for pay (usually \$3 or \$4 each, for about 10–20 minutes). The subjects are not a random sample of any particular population, but they turn out to be typical of the U.S. adult population in age, income, and education (Babcock, et al., 2003), but not in gender, since most are women. Consistent with most cognitive psychological research, we were interested in within-subject results, that is, in the existence (or not) of inconsistencies in the judgments of individuals.

In Experiment 1, we presented people with hypothetical government budgets in which taxes and spending vary independently, leading to deficits, surpluses, or balanced budgets. We asked people for their preferences about taxes and spending. We compared their preferred levels to the starting

levels they were given. If subjects are optimistic, they will prefer higher spending and lower taxes, leading to deficits in all cases. We ask whether they adjust completely so as to maintain a constant balance and size of government, or, conversely, whether they under-adjust, in which case they will fail to correct surpluses and deficits. They will prefer levels of taxing and spending that fail to fully correct existing deficits or surpluses.

To further test the optimism-bias account, we asked subjects in Experiment 1 about their preferences in the short term and in the long term, on separate trials. They might be willing to tolerate deficits in the short term more than in the long term. This might also be true if they think about the pain of adjustment, or if people agree with the rational-expectations theory.

Experiment 2 tests the hypothesis that people prefer spending cuts in the abstract but not in particular. We told subjects, somewhat realistically, that only certain categories of the budget could be cut, and we asked how they would change the levels of spending in each category.

Experiment 1

Method

Seventy-five subjects completed the questionnaire. Their ages ranged from 20 to 71 (median 37); 29% were male. The questionnaire began with an introductory page, the text of which read:

This is about taxes and spending by the U.S. government. At any given time, taxes can be higher than spending, a surplus, or spending can be higher than taxes, a deficit.

When there is a surplus, the government can cut taxes, increase spending or reduce the debt. (The U.S. debt about 3 times the size of the annual Federal budget, so there is plenty of room to reduce it.) Reducing the debt saves money in the future, which can be used either for future tax cuts or spending increases.

When there is a deficit, the government can raise taxes, reduce spending or borrow

more. An increase in the debt means higher taxes or lower spending in the future.

In all of these cases, assume that the economy is performing normally. It does not need special stimulation or dampening.

Right now, a typical American household pays about 20% of income in federal, taxes combined (including payroll taxes). The rich pay somewhat more because the income tax is graduated (higher percent for higher income), but payroll taxes are actually lower for the rich, as a percent of income.

Some questions here ask you about where budget changes should be made. Increases or decreases in spending can occur in the following categories:

- Health care and public health
- Social security (pensions)
- Aid specifically for the poor
- Armed forces (including foreign military aid)
- Foreign aid for health programs, including AIDS treatment and prevention
- Scientific research, including health research

There are other categories, such as interest on the national debt.

After this initial page, the subjects then saw a series of screens in which they were asked for their input, as explained below. A typical screen contained the following statement of the baseline condition, questions and options (with variables shown in brackets, [], not shown on the screens):

The current level of government spending per household comes to 15% [20%, 25%] of the typical household's income.

The current level of total taxes (all levels) comes to 15% [20% 25%] of the typical household's income.

What would you do about government spending in next year's budget? (Pick the closest.)

Cut government spending from 15% of income to 5%

Cut government spending from 15% of income to 10%

Leave government spending unchanged at 15% of income

Increase government spending from 15% of income to 20%

Increase government spending from 15% of income to 25%

What would you do about total taxes in next year's budget? (Pick the closest.)

Cut total taxes from 25% of income to 15%

Cut total taxes from 25% of income to 20%

Leave total taxes unchanged at 25% of income

Increase total taxes from 25% of income to 30%

Increase total taxes from 25% of income to 35%

The choices you have made so far imply that the government will have a . . . You can change them by clicking a different button.

The three levels of taxing and spending were combined to produce nine combinations. Each of these was presented in a short-term condition (as shown above) or in a long-term condition, where the questions were both, "What should the government aim for over the long term?" The resulting 18 cases were presented in a random order chosen for each subject.

At the end of each of the long-term or the short-term cases (depending on the subject's group, which was randomly chosen), the subject was asked, "What would you do to the spending levels in each of the following categories in this case? (Remember, these aren't the only categories.)" After each of the categories, the subject could choose "Decrease," "No change," or "Increase."




Figure 1: Preferred levels of taxation and spending, Experiment 1. (Diagonal lines represent no change from starting point.)

Results

Figure 1 shows subjects' preferred levels of taxation and spending as a function of the starting levels of each. Three features of the results are of interest.

First, subjects preferred lower taxes, reflecting a general tax aversion ($t_{74} = -4.58$, $p = 0.0000$). In the high (25%) and medium (20%) initial tax conditions, subjects lowered the tax rate. In the low (15%) initial tax condition, they supported a slight (non significant) tax increase, although it is worth noting that the introductory page had set a current condition default at 20%.

Second, however, subjects generally favored a surplus over a deficit. Preferred levels of taxation were higher than preferred levels of spending by an average of 1.3% ($t_{74} = 4.90$, $p = 0.0000$, across subjects). There was no significant difference between short-term and long-term. The optimism-bias hypothesis receives no support. Subjects do not want deficits: no subject showed a significant pro-deficit inclination by a within-subject t test across the 18 cases (with the p-level corrected for multiple tests using the step-down resampling procedure of Westfall and Young, 1993,

as implemented by Dudoit and Ge, 2003). The surpluses were created because the subjects cut spending, *here presented in an abstract, general way*, by more than they cut taxes.

Third, subjects adjusted their responses to the current balance of spending and taxation, although it was trivial not to do so — subjects easily could have maintained a constant level of tax and spending independent of the artificially set initial conditions. But they did not. Responses depended on both the starting levels of spending ($t_{74} = -3.69$, $p = 0.0004$; in a regression with taxation minus spending as a function of current spending, current taxation, and their interaction, with the coefficients tested across subjects) and taxation ($t_{74} = 3.56$, $p = 0.0007$, and no significant interaction between current taxation and spending). But subjects did not go far enough to maintain a constant level of taxes, spending, or the balance between them, showing an *anchor and under-adjustment* effect. The upshot was that their preferences led to significant surpluses when surpluses were already present, or even when the budget was balanced ($t_{74} = 5.05$, $p = 0.0000$), but, when deficits already existed, they were maintained ($t_{74} = -3.62$, $p = 0.0005$). The under-adjustment hypothesis is supported.

Experiment 1 thus revealed that subjects are generally tax averse but are *also* deficit averse. Given free rein, they generally support cutting taxes, but aim to balance the budget by cutting general levels of spending. They are not naively optimistic. But they are influenced by initial conditions, however thinly framed or presented.

Experiment 2

Experiment 2 was like Experiment 1, except that we removed the short-term condition, because we found no short term/long term divergence, and we added a new condition in which subjects made particular judgments about category spending. We attempted to approximate the major categories of spending in the U.S. federal budget. (Subjects could comment, and no subject commented that the numbers seemed unrealistic.) In this way, we test the identified-victim explanation, which is

that people oppose particular budget cuts, although they are happy with spending reductions in the abstract.

Method

Experiment 2's introductory page read much like that of Experiment 1, with the following significant change:

Some questions here ask you about where budget changes should be made. Increases or decreases in spending can occur in the following categories (with very rough estimates of current spending in percent):

Spending category	Current spending
Health care and public health	25%
Social security (pensions)	20%
Aid specifically for the poor	15%
Armed forces (including foreign military aid)	25%
Foreign aid for health programs, including AIDS treatment and prevention	2%
Scientific research, including health research	5%

Other categories, such as interest on the national debt, cannot be changed. Thus, any spending changes must be in the categories listed.

The items for 9 of the 18 trials were like the following, with current government spending levels of 16%, 20%, or 24% of income crossed with the same three levels of taxation (leading to surpluses or deficits of up to 8%). Here is an example with the spending of 24% and taxes of 20%, a 4% deficit.

The current level of government spending per household comes to 24% of the typical household's income. The current level of taxes comes to 20% of the typical household's income.

What should the government aim for over the long term? (Pick the closest to what you think.)

Cut taxes from 20% of income to 10%.

Cut taxes from 20% of income to 15%.

Leave taxes unchanged at 20% of income.

Increase taxes from 20% of income to 25%.

Increase taxes from 20% of income to 30%.

What should the government aim for over the long term? (Pick the closest to what you think.)

Cut government spending from 24% of income to 14% Cut government spending from 24% of income to 19% Leave government spending unchanged at 24% of income. Increase government spending from 24% of income to 29% Increase government spending from 24% of income to 34%

On the other 9 of the 18 trials, we replaced the spending question (the second one) with specific questions about spending in each of the six categories. As in the health-care example below, each category was listed in terms of its overall proportion of government spending and its proportion of average income.

What would you do to the spending levels in each of the following categories in this case? Remember, these are the only categories that allow changes in spending. (Pick the closest to what you think.)

Health care and public health (25% of government spending, 6% of income)

Cut spending from 6% of income to 0.

Cut spending from 6% of income to 3%.

Leave spending unchanged.

Increase spending from 6% of income to 9%.

Increase spending from 6% of income to 12%.

When government spending was 24%, the possible range of spending for the direct spending question was 14% to 34%, and the possible range for the category-spending question was 1.2% to 44.8%, given that it was possible to double all programs listed or cut them to zero. Thus, in principle subjects had more room for spending adjustments in the category-spending question.

The 18 cases were presented in a random order chosen for each subject.

The 87 subjects ranged in age from 20 to 68 (median 39), and 24% were male.

Results

Figure 2 shows the mean judgments for the four conditions. Tax1 and Spend1 are from the trials with the direct questions about spending as well as taxation. Tax2 is the taxation question from the trials with category-spending questions, and Spend2 is the level of spending inferred from the answers to the category-spending questions. The three bars in each group are for the three different starting points of taxation or spending. The lines at the top represent the responses that would be expected if the average judgments indicated no change from the given levels.

In Tax1 and Spend1, subjects want less spending and less taxation on the whole, especially when the level of each is high. As in Experiment 1, there is some attempt to adjust toward a constant level, but not enough to remove the influence of the starting point (which would make the bars all the same height within each group of three). Because of this under-adjustment, all deficits and surpluses remained incompletely corrected. On the whole, however, subjects favored neither surpluses nor deficits, although they did favor reductions in both spending and taxation.




Figure 2: Levels of taxation and spending implied by judgments, Experiment 2. (Lines represent no change from starting point.)

(All positive results described in this paragraph were highly significant by t tests across subjects; $p < .0001$.)

Tax1 and Tax2 did not differ significantly; the tax question came first, and, evidently, subjects did not try to adjust their answer for their response to the other questions. This is reminiscent of the disaggregation bias described by McCaffery and Baron (2003), in which subjects do not take advantage of the possibility of correcting a problem in one quantity (spending, here) by manipulating another quantity (taxes). Spend1 and Spend2, however, differed significantly ($t_{86} = 5.44$, $p = 0.0000$). Although subjects did adjust Spend2 somewhat by reducing spending more when initial spending was higher ($t_{86} = -3.31$, $p = 0.0014$), the amount of adjustment (change in from the starting point in Figure 2) was 7% of the amount in Spend1. Moreover, the Spend2 and Tax2 judgments together implied higher deficits than the starting point on the average ($t_{86} = 5.78$, $p = 0.0000$). Subjects wanted to cut taxes ($t_{86} = -6.00$, $p = 0.0000$) but did not want to change spending significantly, when they were faced with questions by category.

Figure 3 shows category-spending changes as a function of category, both for the actual changes, calculated on the assumptions given to the subjects, and changes under an “equal” condition in which each of the six categories was assumed to be a sixth of the spending listed (92% of all spending). It is apparent that subjects were willing to cut some spending, but their favorite target for cuts was foreign aid, which amounted to a small proportion of the budget. However, despite the strong desire to cut foreign aid, analysis of the data on the assumption that all categories were equal in spending (as shown in Figure 3) does not change the main result. Spending cuts were (non-significantly) less in this analysis (taking into account the size of the budget) than in the original analysis. Thus, it appears that a primary source of the reluctance to cut particular categories is the identification of the categories.

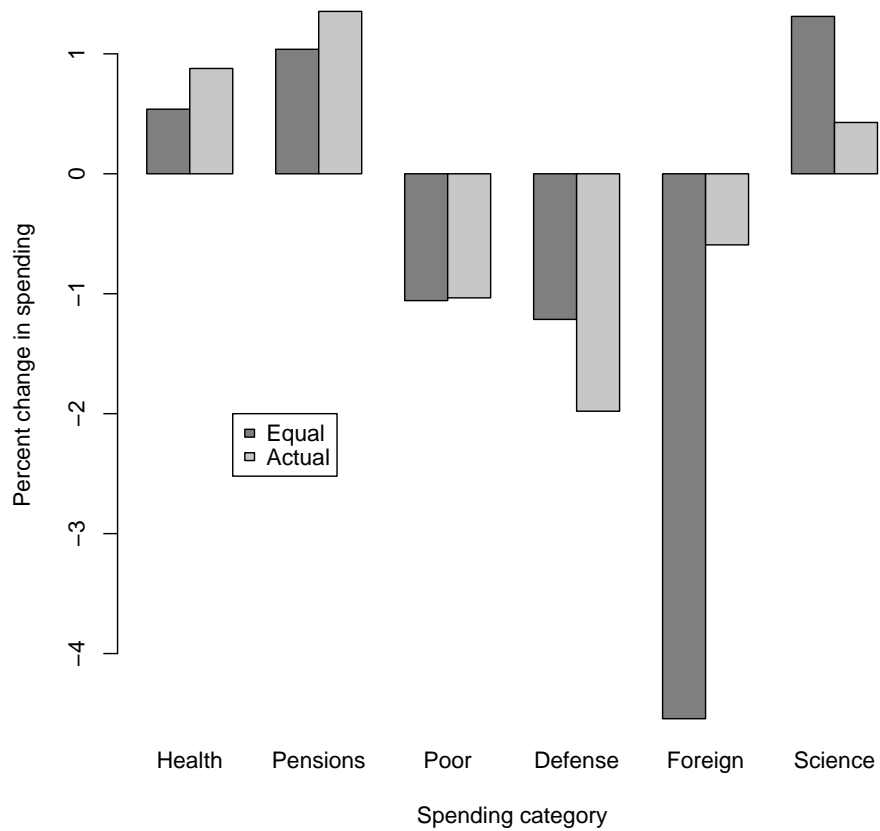


Figure 3: Category-spending changes, in percent of spending, calculated as if all categories were equal parts of the budget, or the actual percents given to the subjects.

Discussion

Our research shows how the “starve the beast” strategy might succeed, at least in part. It also shows why it is critical for opponents of big government to try it. The key to the technique’s success is to match specific tax cuts today, which subjects will support, with the abstract, general idea of spending cuts, which subjects will also support today. If tax cuts today must be matched by specific spending cuts today, then the opposition to both specific cuts and deficits is likely to preserve the status quo. On the other hand, if the tax and spending decisions can be separated in time and (logical) space, then the specificity of the spending cuts can recede, and a *disaggregation bias* effect can take hold (McCaffery & Baron, 2003). Subjects will focus on the tax cuts alone, where a generic tax aversion will lead them to support cuts. A budget deficit results. Once this deficit is created, the preference for fiscal prudence causes people to want to raise taxes and cut spending. But these desires are not strong enough to reduce the deficit to zero, even when people are asked about the “long run.”

There is inconsistency here, but it is not the perhaps more usual optimism bias, ruled out in Experiment 1. Rather it is a failure to properly anticipate the depth of the difficulty in making *specific* cuts tomorrow — the depth, that is, of the endowment effect. At a high level of generality, the starve-the-beast strategy works by pairing a specific (salient) tax cut with an abstract (non-salient) set of spending cuts.

This conceptualization suggests two broad ways for governments to avoid deficits. One is to keep everything abstract: to pass laws, as in the form of constitutional restrictions, about balanced budgets. Experiment 1 shows that in the abstract, subjects indeed support fiscal balance. Many state governments in the U.S. are indeed required to have balanced budgets each year, and the U.S. government has occasionally tried to bind itself in advance by various budgetary rules.

An alternative takes the opposite: to make everything concrete and specific. We could break taxes down into categories earmarked for particular services, as in the case of the various wage

taxes in the U.S. If citizens come to think of each tax as linked with a particular service, they may be less willing to cut taxes (Garrett, 1998); indeed, this could explain why the social security and medicare tax in the United States, now the largest tax for *most* American taxpayers, is also the one major federal tax that has *never* been cut (McCaffery & Baron, 2003). This alternative would probably lead to a larger, more active government than the first method (binding in the abstract).

It may also be possible simply to confront people with the conflict in their opinions, as a device of argumentation. Advocates of larger government are often tempted to answer their opponents, who want to cut both taxes and spending, by saying, “OK. Where? Exactly what do you want to cut?” The usual answer, “government waste,” — an abstraction — may stop working after a while, given that practically every single politician elected to public office has been against waste. But then that would be rational, as politics seldom is.

References

- Barro, R. (1974). Are government bonds net wealth? *Journal of Political Economy* 82: 1095-1117.
- Buchanan, J. M., & Tullock, G. (1962). *The calculus of consent: Logical foundations of constitutional democracy*. Ann Arbor: University of Michigan Press.
- Camerer, C. F., & Lovallo, D. (1999). Overconfidence and excess entry: An experimental analysis. *American Economic Review*, 89, 306–318.
- Dudoit, S., & Ge, Y. (2003). Bioconductor R packages for multiple hypothesis testing: multtest. <http://www.bioconductor.org>.
- Fennell, C. C., & Fennell, L. A. (2003). Fear and greed in tax policy: A qualitative research agenda. *Journal of Law and Policy*, 13, 75–138.
- Garrett, E. (1998). Harnessing politics: the dynamics of offset requirements in the tax legislative process. *University of Chicago Law Review*, 65: 501-569.
- Jenni, K. E., & Loewenstein, G. (1997). Explaining the “identifiable victim effect.” *Journal of Risk*

and Uncertainty, 14, 235–257.

- Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science*, 39, 17–31.
- Kahneman, D., and Tversky, A. (1979). "Prospect theory: An analysis of decision under risk." *Econometrica* 47(2): 263-291.
- Kahneman, D., Knetsch, J.L., and Thaler, R.H. (1991). Anomalies: The endowment effect, loss aversion, and the status quo bias. *Journal of Economic Perspectives*, 5:1, 193-206.
- Kogut, T., & Ritov, I. (2004). The "identified victim" effect: an identified group, or just a single individual? Hebrew University of Jerusalem.
- 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., & Baron, J. Thinking about tax. USC CLEO Research Paper No. C04-10; USC Law School, Olin Research Paper No. 04-13. <http://ssrn.com/abstract=567767>
- Schelling, T. (1984). "The intimate contest for self-command," *Choice and Consequence* Cambridge: Harvard University Press.
- Small, D. A. & Loewenstein, G. (2003). Helping the victim or helping a victim: Altruism and Identifiability. *Journal of Risk and Uncertainty*, 26, 5–16.
- Westfall, P. H., & Young, S. S. (1993). *Resampling-based multiple testing: Examples and methods for p-value adjustment*. New York: John Wiley & Sons, 1993.