The SEC's Busted Randomized Experiment: What Can and Cannot Be Learned

Kate Litvak
Northwestern University, Pritzker School of Law

Bernard Black
Northwestern University, Pritzker School of Law and Kellogg School of Management, Department of Finance

(December 2017)

The SEC's Busted Randomized Experiment: What Can and Cannot Be Learned

Kate Litvak and Bernard Black*

Abstract

During 2005-2007, the US Securities and Exchange Commission ran a randomized experiment, in which it removed short sale restrictions for some "treated" firms, ostensibly chosen at random, and left them in place for others. The SEC experiment has been exploited by many finance and accounting scholars, who report evidence that removing short sale restrictions affected a wide range of financial outcomes, including market prices, firm investment strategy, accounting accruals, auditor behavior, innovation, and much more. We show that the SEC busted its own randomization experiment, in a way which undermines most prior studies. We discuss what one can and (often) cannot still learn from the SEC experiment, given how it was conducted. We also develop reasons why relaxing short-sale restrictions was unlikely to affect most studied outcomes, at most studied firms. We explain the need, when studying indirect effects of the SEC experiment, to account for nonrandom choices by short sellers on which firms to "target", and the nonrandom choices by firm managers and other market participants on how to respond to removal of short sale restrictions. We then revisit selected results from prior studies, and show that some are simply wrong and others are implausible.

---

* Litvak is Professor of Law, Northwestern University, Pritzker School of Law, email: k-litvak@northwestern.edu. Black is Nicholas D. Chabraja Professor at Northwestern University, Pritzker School of Law and Kellogg School of Management. Tel. 312-503-2784, email: bblack@northwestern.edu. We thank Hemang Desai, Mathew McCubbins, and participants in Conference on Empirical Legal Studies (2016), Burton Conference at Columbia Law School (2017), and workshops at City University Hong Kong and Tilburg Institute for Law and Economics and [*to come] for helpful comments and Woongsun Yoo for excellent research assistance. This work was supported by **[*to come]**.

Table of Contents

Contents

I. Introduction

In 2005-2007, the SEC ran a randomized trial – the only such trial in the SEC's history. The SEC randomly selected one-third of the Russell 3000 firms and exempted them from prior short sale restrictions, known as the "uptick rule." The exempted firms can be thought of as "treated" by the experiment, with the other 2,000 firms serving as "controls." The randomization process was conducted transparently. The list of treated and control firms was posted on the SEC website, encouraging research on the effects of the short-sale restrictions.

The SEC's own Office of Economic Analysis ("OEA") used the randomized trial to conclude that abolishing short-sale restrictions had only minor effects on share trading (OEA, 2007). The OEA found that short sale trading volume rose, but open short interest (the number of shares held in short positions at any point in time) did not. Share price volatility fell, but only for larger firms. Many other studies, often published in top journals (to date, JF twice, RFS twice, JAR, JAE) have relied on this randomized trial to study a wide variety of other financial market outcomes, including: share prices levels, share price reactions to the experiment announcement in 2004 and the experiment launch in 2005, abnormal accruals and other measures of earnings management, disclosure of bad news, disclosure opacity, capital investment and equity raising, the sensitivity of investment to share prices, audit fees, auditor behavior, and firm innovation**.**

However, the SEC, rather quietly, busted its own randomization. After the initial randomization, but before the treatment period began, the SEC took the randomly-selected "original control" firms; picked the largest one-third of these firms and suspended the uptick rule for trading in these firms' shares *after* regular trading hours. These firms became, in effect, "partly treated" – the uptick rule applied only during regular trading hours. The remaining original control firms, which we call "full controls", remained fully subject to the uptick rule.

That is, instead of having randomly-selected "treated" and "controls", we have randomly selected fully treated firms, but must compare these firms to non-randomly selected partly treated and full control firms. We cannot compare fully treated firms to full controls because the full controls are a non-random subset of the original controls. And we cannot compare fully treated firms to all original controls, because many of the original controls were partly treated.

The SEC did not keep the non-random division of the "original controls" into full controls and partly treated firms a secret. The decision rule the SEC said it followed (imperfectly, as shown below) was included in the 2004 SEC release announcing the experiment.[1] The supposed rule was: original controls in the Russell 1000 (roughly, the largest firms in the Russell 3000) were partly treated; while the rest of the original controls (Russell 2000 firms) remained as full controls. But apparently, researchers who studied the experiment did not closely read the SEC release.[2] The busted randomization went unnoticed by all prior researchers, *including* the SEC's own Office of Economic Analysis (the "OEA"). Indeed, Chester Spatt, the then-SEC Chief Economist who oversaw the experiment, told us that he was not aware that the SEC had created a partly treated group from the original controls. The best explanation we were able to find for why the SEC busted its own randomization is that "the lawyers did it," without telling the economists! Perhaps – as one of them later told us -- the lawyers believed that partial treatment didn't matter, and was similar to full control status.

The good news is that there are ways to deal with the busted randomization, given the mechanical way that it was busted. The partly treated firms are large and the full controls are small, with limited overlap between partly treated and full controls on size measures. In effect, the SEC conducted two separate randomized trials: For smaller firms, the SEC conducted the randomized trial everyone assumed took place: "fully treated" firms (no short sale restrictions) versus "full control" firms (regular restrictions) (below, "Small-Firm Experiment"). For larger firms, the SEC conducted a different trial: fully treated firms versus partly treated firms ("Large-Firm Experiment"). Those two experiments can be studied.[3]

The bad news is that the busted randomization invalidates most prior studies based on the SEC experiment. Only a few results, specific to small firms, reflect the intended comparison between fully treated and full controls. Results for the full sample would be valid only if one assumes that there is no difference between partial treatment and no treatment. There is no basis for this assumption, the busted randomization cannot supply such a basis, and we provide

[1] Securities Exchange Act Release 34-50104 (July 28, 2004).

[2] Neither did we, until we discovered the busted randomization in another way, and went looking for an explanation. Even then, it took months before we found the rule, in plain sight as it were, in the SEC release.

[3] All three groups of firms – fully treated, partly treated, and full controls -- overlap only within a narrow size band; too small to be usefully studied.

evidence below that it is likely incorrect. Thus, our first contribution is to document the experimental design that the SEC actually implemented, and consider what results can be obtained from the separate Large and Small Firm Experiments.

Our second contribution is to assess the institutional details behind short-selling, and how they affect what results one can plausibly expect from the short-sale experiments. One can crudely divide short-selling into two types: arbitrage and substantive. Arbitrage short-selling centrally involves rapid trading to profit from small departures of prices in two markets from parity. For example, stock index futures may become over- or under-priced relative to the underlying basket of stocks. If the futures are underpriced, arbitrageurs will seek to buy futures and sell short the equivalent shares – which implicates short sale restrictions. Arbitrage trades are usually short-term, should not produce large open short positions, and should have little or no effect on equilibrium share prices or on anything else that firm managers care about.

Substantive short-selling, in contrast, is driven by close investigation of firm fundamentals. Informed investors, often called "short sellers," seek to identify and "target" firms with overpriced shares, sell those firms' shares short, and profit from a subsequent price drop. Short-sellers typically hold positions for a substantial period. Their actions can lead to targeted firms having large open short positions. Open short interest has thus become a standard measure of substantive short-selling activity, and during this period was reported monthly by the stock exchanges. Substantive short-selling *could* affect share prices and thus firm behavior. It could, perhaps, also affect the behavior of other actors. But to assess any effect of the short-sale experiment on the behavior of firms or other actors, one must understand the causal chain. That chain starts with substantive short-selling, and operates largely through open short interest and the effects of short positions on share prices.[4]

These institutional features of short-selling imply that most of the indirect effects of the experiment, reported in the literature, are implausible for most of the studied firms. In brief, many larger firms were also traded on regional stock exchanges or electronic exchanges, most of which did not limit short sales. Removing short-sale restrictions on these firms' main trading platforms –the New York Stock Exchange ("NYSE") or the NASDAQ national market

[4] Some substantive short-sellers mount publicity campaigns, seeking to "talk down" the share prices of targeted firms. In this instances, open short-interest operates both as a direct marker of substantive short sale interest, and as a proxy for the other activities that short sellers may undertake.

(“NASDAQ”) – could facilitate arbitrage, by allowing short selling on more-liquid “main” markets, but should have had little effect on substantive short-selling.

Many firms also had publicly traded options, which let substantive short sellers engage in options trades equivalent to selling short shares, without restriction. Under the put-call parity theorem, dividends aside, the transactions [sell $\text{call}_{X,t}$, buy $\text{put}_{X,t}$] (with X = any exercise price and t = any expiration date), are an economic near-equivalent to selling short the underlying shares). For these firms too, removing short sale restrictions may have facilitated arbitrage, but should have had little effect on substantive short-selling.

Most Large-Firm Experiment firms had publicly traded options; many also traded on secondary exchanges. About half of the Small-Firm Experiment firms had public options. Thus, only for the smaller firms within the Small-Firm Experiment, without a secondary market or public options (call these firms the Plausible Subset”), would one expect relaxing the short-sale rule to meaningfully affect substantive short selling. No prior study discusses these institutional details, or assesses whether the effects they report are driven by this limited group of firms, although a few studies do report that some effects appear principally for smaller firms.

Even for the Plausible Subset of smaller firms, an effect on substantive short-selling is unclear. Substantive short-sellers are generally not in a rush to sell; they can often wait for the next uptick. Any effect of the uptick rule on substantive short-selling surely diminished as tick size shrank and index arbitrage trading grew (this arbitrage often generates upticks, divorced from firm fundamentals). A natural null hypothesis is that by 2005, the uptick rule affected arbitrage, but not substantive short-selling. This limited effect could explain why the experiment was politically possible – there was no political opposition from firm managers, because the short-sale rule no longer mattered much. The experiment attracted essentially no news coverage, which suggests unimportance.

Our third contribution is more subtle but likely more important for future efforts to exploit randomized trials in accounting and finance. We explain why, quite apart from the first two concerns, most prior papers used flawed research designs, which do not account for nonrandom compliance with treatment. The core problem is the following. Randomized assignment of firms to treatment or control, within each experiment, allowed traders to short treated firms more easily than (partial or full) control firms. The *only* research question that this randomization addresses directly is how removing these restrictions affects traders’ behavior.

This is a simple randomized experiment (well, two of them). One can estimate the average causal effect of each experiment on trader behavior by comparing mean outcomes for treated and control firms.[5] Outcome variables for "direct effects," plausibly due to increased arbitrage include: short sale volume, option trading volume, bid-asked spreads and other measures of liquidity, share price volatility, and shifts in short-selling among exchanges and trading platforms. One can see open short interest either as a direct effect on trading due to substantive short selling, or as an indirect effect, mediated by short-sellers' decisions on which firms to target; it also proxies we have for actual targeting.

Once we go beyond direct effects on trading, estimating causal effects becomes more challenging. For all other research questions – how do changes in substantive short-selling affect share prices, firm characteristics, management behavior, auditor behavior, etc. -- a simple comparison of means cannot identify an average causal effect. Substantive short-sellers will short-sell some treated firms more than others, based on many factors, some unobservable to researchers. To oversimplify, suppose that if short-sale restrictions are lifted, short-sellers will "target" some firms, beyond those that they would target with these restrictions, where "targeting" is a 0-1 variable. We discuss below how to operationalize this concept, based on open short interest.) In the language of randomized trials with two-sided noncompliance, firms that are targeted only without short-sale restrictions are "compliers." Firms that remain non-targeted are "noncompliers." Firms that would be targeted, with or without short-sale restrictions, are "always-takers; and one assumes that there are no "defiers" – firms that will be targeted only *with* short-sale restrictions (Angrist, Imbens, and Rubin, 1996). Compliance status is a nonrandom, partly observed firm characteristic (Frangakis and Rubin, 2002). Since only some firms that are offered treatment are actually treated (targeted), we have what is often called an "intent-to-treat" design.

---

[5] One could also add covariates and run regressions that control for what should be chance differences in these covariates between treated and control firms. Simple OLS, in which one regresses the outcome on a treatment dummy, will be unbiased for the average treatment effect (ATE), with or without covariates. Adding covariates can potentially improve precision, but omitting them (or including some relevant covariates but omitting others) should not cause bias.

One can also assess how the average treatment effect varies with firm characteristics – for example, is there a larger effect for firms without publicly traded options? But causal claims cannot go beyond average effects. If we observe, that the effect of treatment on outcome y varies with covariate x, we cannot causally ascribe that variation to x, because x may be associated with an unobserved covariate u, which is the true cause. For cross-sectional claims, we are back to association, not causation.

An illustrative parallel with medical trials: suppose researchers want to assess whether a particular medicine will improve children's health. They offer the medicine to randomly-selected children, to be administered by parents at home. Some parents in the treatment group give the medicine (compliers). Some do not, give the wrong dose, or stop giving the medicine because their children don't like it or react badly to it (noncompliers). Some parents in the treated and control groups would obtain the medicine whether treated or not (always takers). These behaviors reflect unobserved characteristics of children and parents. The only group for which one can estimate an average causal effect is for the compliers (known as a "local average treatment effect (LATE)"), because this is the only group whose behavior is changed by the treatment.

Even the LATE estimate requires further assumptions, which should always be stated explicitly and defended. In the medical example, one must assume that the treatment effect comes *only* from the medicine, not a placebo effect, or a change in parents' behaviors due to being in the treated or control group. Treated parents who administer the medicine can't relax other child-rearing behaviors. Treated parents who don't administer the medicine can't change their behavior either – for example, watching their children more carefully, to compensate for not giving them the medicine. Control parents can't change their behavior, to compensate for not having access to the medicine. And there cannot be interactions between treated and control children – as there would be for infectious disease.

These questions should be asked for any intent-to-treat design (Imbens and Rubin, 2015, chs. 23-24). We will call them "only through conditions." Econometricians often call them exclusion restrictions – one "excludes" – assumes away – some possible explanations for an observed intent-to-treat effect. These are substantive assumptions about the "science" of a particular experiment, which cannot be directly tested. They must be defended based on logic, substantive understanding of the world, and sometimes indirect checks, including (when possible) assessing the plausibility of the assumed causal sequence.

Note too that even if one finds an average benefit from the medicine for the whole sample (driven, as noted above, by the compliers), we cannot assume that this "intent-to-treat" estimate is "conservative" (biased toward zero) for all children. Suppose, for example, that some parents stop administering the medicine after observing an adverse reaction. If so, the medicine could cause net harm to noncompliers, and the intent-to-treat estimate would be an "upward" biased (away from zero) estimate of the average effect for all children.

The short-selling parallel is apparent. Relaxing short-sale restrictions allowed short-sellers to more easily target firms, but the short-sellers targeted only some treated firms for short selling. We do not know what would have happened to non-targeted firms if they had been targeted, and we cannot ignore this question by assuming that the intent-to-treat estimate is biased downwards. We need instead to use causal-inference methods for randomized trials with noncompliance.

Researchers have explored many indirect effects of the short-sale experiment. All involve substantive short-selling, rather than arbitrage. It is useful to loosely divide these indirect effects into second, third, and fourth-order-effects. Second-order effects involve the effect of short-selling (not regulatory permission to short-sell, but actual short-selling) on the market characteristics of a firm's shares. These questions are affected by short-sellers' nonrandom compliance, but in ways that can plausibly be addressed by standard methods, discussed above, drawn from the literature on randomized trials. This category of questions includes: the impact of short-selling on share prices, option prices, frequency of extreme price changes, and pricing efficiency (responsiveness to news).

Second-order effects (still assuming for simplicity that firms are either "targeted" by short-sellers or not) involve a classic "causal" instrumental variables (IV) setting, in which one uses IV technology to estimate a treatment effect for the compliers. The estimates in prior studies can be understood as "reduced form" or "intent to treat" estimates – averages of a treatment effect for the compliers, plus an assumed zero effect for always-takers and never-takers." To develop a causal IV estimate for the compliers, one needs to estimate how many firms are compliers, divide the intent-to-treat estimate by the fraction of compliers (this ratio is a "Wald estimate"), and then assess: (i) are any apparent results indeed found only, or more strongly, for firms in the treated group that are targeted by short-sellers (compliers and always takers); (ii) is the LATE estimate economically reasonable in magnitude, (iii) are the only through conditions plausible; and (iv) is there evidence to support the assumed causal channel? Yet no prior study goes beyond providing reduced form estimates, and treating these estimates as if they are averages for the entire sample. None of the researchers who study second (or higher) order effects state or defend the only through conditions. Yet, these conditions can be problematic, as we discuss below.

If an apparent intent-to-treat effect is found, we can use substantive knowledge about the causal chain to assess plausibility. A natural question: How many firms appear to be compliers -- targeted only if short-sale restrictions are relaxed? The best proxy for targeting is open short interest. One can ask: Is there an average (intent-to-treat) effect on open short interest? Is there an average effect for the Plausible Subset? Even if not, is there an effect on the *distribution* of open short interest across firms? If none of the above – as we find below -- this is strong evidence against the plausibility of second- and higher-order effects.

If there are few or no compliers, one would not expect an effect on share price, but one can also check for such an effect. Is there evidence of an effect on share price for all firms in the treated group? For the Plausible Subset? If no price effect is found – as we find below – this further diminishes the plausibility of other indirect effects.[6] In effect, we have an IV design with no first stage. Researchers have elided the substantive need for a first stage by providing only intent-to-treat estimates, without asking whether there is evidence for a first-stage.

Third-order-effects questions ask about the impact of changes in capital markets (not the impact of the short-selling itself) on firm behavior. These questions are affected by nonrandom compliance in complex ways, because they involve a series of endogenous choices by multiple parties: first, short-sellers nonrandomly decide which firms to target and how strongly; then, their targeting nonrandomly impacts aspects of the firm's securities (such as share price); then, firm managers and/or boards of directors nonrandomly respond to those changes by altering firm policies. The third-order effects that researchers have studies -- choices by firm managers in response to, or possibly in anticipation of short-sale targeting -- include capital and R&D investment; capital raising; mergers and acquisitions; the clarity, frequency, and other aspects of disclosures; executive compensation; abnormal accruals and other measures of earnings management, and so on. In our judgment, if second-order effects are implausible, third-order effects are still less likely.

Finally, fourth-order-effects questions involve the response of *other* market participants (auditors, analysts, other firms) to the changes in firm behavior caused by the changes in capital

[6] If a price effect were found, one could continue to explore the causal channel by assessing whether the price impact is principally for the compliers. If not, then once again, other indirect effects become less plausible.

market characteristics caused by targeted short selling caused by changes in short-sale rules. Fourt-order outcome variables that researchers have studied include auditor fees and the likelihood of earnings restatements. These outcomes are affected by nonrandom compliance at every step in the causal chain.[7] The further one moves from the original shock, the more suspect the ability of randomized trial methods to deal with endogeneity, and one is increasingly dependent on untestable assumptions about how the causal chain operates. Moreover, the longer the causal chain, the more important – for plausibility – to have a large shock at the beginning

Our fourth is to begin the large task of assessing which results from prior studies will survive the busting of the SEC's original experiment, and reassessment based on the substantive and causal inference factors noted above. A full assessment is a multi-paper project; we begin it here. We find evidence that:

> Neither experiment significantly affected open short interest. This is consistent with OEA (2007), but we provide separate findings for the two experiments, and also study the distribution of open short interest across firms.
>
> Neither experiment significantly affected share prices, either when then experiment was announced in late July 2004, when it began in May 2005, or when it ended in May 2007.[8]
>
> Results from the SEC's own study support an increase in short sale trading volume for the Small Firm Experiment; in contrast, there was no apparent effect on trading volume from the Large Firm Experiment.

These results are consistent with the Small Firm Experiment facilitating arbitrage trading, but not substantive trading and the Large Firm Experiment simply not mattering much. Given the lack of apparent effect of relaxing the uptick rule on open short interest or share price, most of the reported indirect effects of the short-sale experiment on firm outcomes are likely spurious. To be sure, "likely" is not the same thing as "definitely": replication is needed for one to be sure whether a particular result will survives.

We hope that this project can serve several goals. First, it should underscore the need to examine experiment design closely and assess possible design flaws, rather than assuming that a

---

[7] Some outcomes do not fit these categories neatly; for example, abnormal accruals depend on both firm decisions and auditor decisions, and boards of directors can be seen either as acting for the firm, or as responding to managers' actions.

[8] We view the evidence in Grullon, Michenaud and Weston (2015) of an effect on open short interest after the experiment was announced, but before it took effect, as implausible and spurious. We view their evidence in Grullon, Michenaud, and Weston (2015) of a price effect earlier in July 2004, *before* the treated and control groups were announced, as highly implausible and spurious.

supposedly randomized experiment was randomized in fact. One must also look closely at the science of the subject being studied, to assess which effects are plausible. And one must assess which statistical methods for randomized studies with partial compliance can be used, for which outcome variables. We hope that our stress on the need for greater care in carrying out and studying randomized trials will not discourage these trials, for which there is often no good substitute except guesswork (compare Romano, 2013). For us, the principal message from the SEC experiment is that relaxing the short-sale rule had little impact on substantive short-selling. This is an important policy message. In this sense, the experiment was a success – it led to relaxing a rule which served little purpose. We suspect that most studies reporting indirect effects will not withstand scrutiny, but only begin that scrutiny here.

This paper proceeds as follows. Part II summarizes the literature exploiting the SEC's randomized experiment. Part III explains the SEC randomized experiment, how the SEC busted its own randomization, discusses what little we could learn about how this happened, and shows the consequences of the busted randomization. In Part IV, we discuss whether the ostensible treatment, which affected only a firm's main exchange, in fact meaningfully expanded the opportunities for substantive short-selling, as opposed to arbitrage. We also discuss the third problem with prior studies: failure to address consider nonrandom compliance. In Part V, we show that there is no evidence that the Large Firm and Small-Firm Experiments affected either open short interest or share prices. Yet, at the same time, data from the SEC's own study supports different "direct" effects of the two experiments on trading markets.

The remainder of this project, begun but not yet completed, will be as follows. We will revisit core outcome variables studied by others, including the effect of removing the short-sale restriction on: (i) short-sale trading volume; (ii) daily price changes; (iii) open short interest; (iv) firm investment; and (v) abnormal accruals. We will choose prior studies to replicate emphasizing recent, highly placed research on indirect effects. A tentative initial list of papers that we plan to re-examine, data permitting:

DeAngelis, Grullon and Michenaud (RFS, forthcoming 2017)
Fang, Huang and Karpoff (JF 2016)
Grullon, Michenaud, and Weston (RFS 2015)
He and Tian (working paper 2016, presented at AFA 2016 annual meeting)
Hope, Hu and Zhao (JAE 2016)
Li and Zhang (JAR 2015)

We will first assess and document how we would answer the research questions in these prior studies, taking into account the research design concerns discussed above. We will then carry out our own analysis, using our own data. After that, we will seek to formally replicate the results in these prior studies. Where feasible, we will request the authors' own datasets and code, and compare their datasets to our own hand-cleaned dataset.

We are, at this point, intentionally agnostic about which specific results we will seek to re-examine, or how. In the spirit of "design before analysis," (Rubin, 2008), we do not want our own thinking about how *we* would answer the core research questions in these studies to be too strongly influenced by how the authors of these papers chose to proceed. We intend, in sequential order, to: (i) build and clean our own data; (ii) think carefully about how one might answer the research questions in these prior studies, in light of the data; (iii) write down the tests we plan to run; (iv) run those tests (and any follow-on tests that we consider useful, in light of the initial findings. Only then will we (v) replicate the specific results in prior studies, and assess the robustness of those results in light of the research design questions outlined above.

II: Prior Uses of the SEC Experiment

A number of studies have investigated the impact of the SEC experiment on treated firms. These are complex studies, often involving multiple outcomes (short sale volume, short sale interest, share returns, daily and intraday volatility, bid and offer depth, spreads, investment, abnormal accruals, and more). We summarize here selected results.

A. Direct Effects on Behavior or Traders and Second-Order Effects on Trading Markets

One category of papers studies the effect of the experiment on trading and basic measures of securities markets. These studies typically address both the direct effect of the randomized trial (the impact of the uptick rule on traders' choices of whether and where to short sell), and second-order effects (the impact of short-selling on price and related market characteristics). The direct-effects questions are vulnerable to our first concern (busted randomization), but the research questions often could reflect arbitrage, and thus are less sensitive to the second concern (the substantive importance of the uptick rule, given that substitutes were often available). Studies of share price and other second-order effects are vulnerable to all three of the concerned in the Introduction, including nonrandom noncompliance.

The SEC's own OEA (2007) investigated the effect of short-sale restrictions on the short sale volume, open short interest, option trading, share prices, liquidity, and volatility, market efficiency, and the frequency of extreme price changes. The OEA reports that abolition of the uptick rule increased short sale volume, especially for smaller stocks, but not open short interest; below, we confirm their finding of no effect on open short interest. The OEA also found evidence of lower daily volatility, mostly in larger stocks. Most other effects were small. This study, along with several academic studies with similar findings, led the SEC to abolish the uptick rule in 2007. We discuss selected results from the OEA study below.

Deither, Lee, and Werner (2009) study a six-month period from February to July, 2005 – the three months before the experiment began and the first three months in the experiment period. They ask, first, whether the exemption from the uptick rule affects short selling, liquidity, and volatility, and second, whether the increases in short selling affect market quality. They find mixed results: short sale volume increased (as the OEA also found), but there was no effect on returns or daily volatility. There was a small increase in spreads and intraday volatility for firms traded on the New York Stock Exchange (NYSE), but not for NASDAQ-traded firms. The modest results they found could largely reflect arbitrage activity.

Alexander and Peterson (2008) find that treated firms experience higher short sale volume, faster execution speed, smaller trade sizes, smaller ask depth, but similar spreads and liquidity. They find a drop in liquidity of NYSE stocks on several measures, but no effect on several measures of volatility. The authors also study the smallest 20% of the Russell 3000, and find no evidence that treated firms differ from control firms on measures of liquidity and volatility during periods when the overall market declines 1% or more.

Wu (2006) studies NYSE firms and finds increases in short sale volume, but only for small firms. She also finds declines in some liquidity measures, primarily for small firms; no changes in price volatility; and no changes in price efficiency.

Bai (2008) asks a different question. One commonly cited benefit of short selling is that it allows allowing market participants to quickly incorporate negative information into stock prices. One important source of new information is earnings announcements. Bai asks whether treated experienced different reactions to negative earning shocks, and finds no such differences. She also finds no differences in short sale volumes when negative earnings news is released.

Deng, Gao, and Kim (2017) report evidence that relaxing the short-sale rule reduced a measure of the risk of market-wide stock price crashes.

B. Second-Order Effects on Trading Markets

Grullon, Michenaud, and Weston (2015) study selected direct effects (open short interest), second-order effects (share price), and third order effects (investment in capital expenditures and R&D; asset growth; capital-raising). They find that: (i) open short interest in treated firms increased during the period between the announcement of the SEC experiment in 2004 and its launch in 2005; (ii) share prices of treated firms fell in the two weeks before the experiment was announced; and (iii) capital expenditures fell and asset growth slowed, especially for smaller firms. We reconsider their first two findings below, and find no support for them; we leave their other findings to future work.

C. Third-Order Effects on Firm Behavior

A number of other recent papers study how being treated affects firm behavior. Fang, Huang, and Karpoff (2016) ask whether short selling impacts firms' reporting behavior. They find that treated firms reduce several measures of earnings management (including abnormal accruals and probability of meeting or beating analyst consensus earnings forecast) during the treatment period. They also report a fourth-order result: Treated firms are more likely to be caught for financial misrepresentations occurring 2004 or earlier.[9]

Li and Zhang (2015) ask whether the increase in short-selling pressure caused by the SEC's experiment affected firms' voluntary disclosure choices. They find that treated firms reduced the precision of bad news forecasts and the readability of bad news annual reports. Chang, Lin, and Ma (2015) ask whether short-selling threat affects managers' mergers and acquisitions decisions. They find that treated firms reduced abnormal capital investment, consistent with this hypothesis.

Kim and Park (2015) study the impact of short selling on sensitivity of firm's investment to price. They find that treated firms experienced greater reduction in the sensitivity of investment to price than controls, and that the subsequent operating performance of treated firms is lower, which they attribute to lower usefulness of price signals to managerial decision-making.

---

[9] The logic here is to assume that the underlying probability of financial misreporting is the same for treated and control firms, and infer that the higher probability of detected misrepresentation flows from increased short selling.

He and Tian (2016) report that short-selling positively affects the quality of firms' investment in innovation, measured by patent counts and citations to those patents.

C. Fourth-Order Effects on Firm Behavior

Hope, Hu, and Zhao (2016) study the impact of short selling on auditor behavior, in particular, audit fees. Their proposed channel is that short-selling increases bankruptcy risk, which increases the risk of investor suits against auditors, which leads auditors to charge higher fees. They find that treated firms have higher audit fees, but only when they also have high abnormal accruals or default risk. They also report higher audit fees for firms whose managers have high levels of in-the-money options, under even more opaque reasoning for why this might flow from short selling.

Chang, Lin and Ma (2015) study the effects of short-selling on takeovers. They find evidence that greater opportunity for short selling leads to more hostile takeovers and higher acquirer returns. We classify this as a fourth-order paper because the authors find an effect on acquirers, not just targets.

DeAngelis, Grullon, and Michenaud (2015) ask whether boards of directors of firms subjected to short-selling respond by providing more compensation or employment protections to managers. They find that treated firms granted more stock options to top managers and adopted more antitakeover provisions. We treat this as a fourth-order paper because the implicit causal chain runs from short sale restrictions to short seller behavior to manager behavior to board response to managers.

III: The Busted SEC Experiment

A. Short Selling: Policy Considerations

Until July 2007, Rule 10a-1 imposed a restriction on short selling, known as the uptick rule. With limited exceptions, aa listed security could only be sold short if the prior transaction was at a price above the immediately preceding sale (plus tick), or at a price equal to the last sale price if that last price was higher than the last different price (zero-plus tick). The rule's original purpose, when it was adopted in 1938, was to prevent short sellers from accelerating downwards price spirals. By the early 2000s, the uptick rule was widely considered counterproductive because it prevented negative information from being quickly incorporated into market prices.

There was also increasing doubt as to whether the uptick rule still mattered much.[10] Tick size had moved from one-eighth of a dollar to one-sixteenth and then, in 2001, to a penny. Index arbitrage had grown into a major business, and would often generate plus ticks for reasons unrelated to firm fundamentals. There were also end-runs around the short-sale rule available, especially for larger firms. First, by 2005, about two-thirds of the Russell 3000 firms had publicly traded options. For these firms, short sellers could rely on the put-call parity theorem to short-sell a near-equivalent to shares through option trades. Dividends and differences in upfront cash aside, [short shares = short $call_{X,t}$, long $put_{X,t}$], where X is any exercise price and t is any expiration date. Second, many larger firms also traded on one or more regional exchanges (the Philadelphia, Boston, and National Stock Exchanges). The regional exchanges did not enforce the uptick rule. Thus, "substantive" short-sellers, who wanted to bet on a price decline, had little difficulty doing so. The rise of institutional ownership also meant that short-sellers usually had little difficulty borrowing the shares they needed to sell short.

For NASDAQ firms, additional factors weakened the effect of the short sale restrictions. The NASD used a bid-based equivalent of the NYSE uptick rule, which was considered to be weaker. Moreover, two electronic exchanges, ArcaEx and INET, had siphoned off around 40% of the trading volume in NASDAQ firms, and did not enforce the bid test.

Thus, by 2005, there was reason to believe that the uptick rule had little effect on substantive short-sellers, and thus on open short interest (the standard measure of the extent of substantive short selling). However, it did limit arbitrage, especially index arbitrage. Consider a typical index arbitrage strategy involving, say, the Standard & Poor's (S&P) 500 stock index. If futures on the S&P 500 index, or the index options equivalent to futures (the parity equation is [long index $future_{X,t}$ = long $call_{X,t}$, short $put_{X,t}$], became overpriced relative to the underlying shares, index arbitrageurs sell short the futures (or options equivalent) and buy the underlying shares. If futures on the S&P 500 index, or the options equivalent, became *under*priced relative to the underlying shares, index arbitrageurs will sell short the shares, and buy the futures (or options equivalent). This arbitrage keeps prices in all three markets (stock, options, and futures) tightly linked to each other.

[10] Our discussion here draws heavily on PowerPoint slides prepared by Hemang Desai, commenting on Fang, Huang and Karpoff (2016). We are grateful to Prof. Desai for sharing his slides with us.

There were two principal limits on index arbitrage. The first was transaction costs to buy or sell all of the component securities in an index of many companies. The second was the uptick rule. To short the S&P 500, an arbitrageur needed to satisfy the uptick rule for 500 separate securities, or else work around the rule. That index arbitrage was a vibrant business at the time confirms that work-arounds were available, even on the very short time-frames that are relevant for index arbitrage. Nonetheless, the work-arounds increased transaction costs, and hence the minimum spread between futures and shares needed to justify an arbitrage trade. Higher transaction costs also presumably reduced index arbitrage volume.

Similar considerations affected single-firm arbitrage, between different markets for shares, or between share and option markets. The uptick rule increased effective transaction costs, and thus reduced arbitrage trading volume. Thus, the uptick rule had a clear dampening effect on arbitrage, especially index arbitrage.

There could also be an interaction between arbitrage and substantive short-selling. If we put aside (illegal) naked short-selling, both draw on the same limited pool of shares available for loan. Arbitrageurs need an assured supply of borrowable shares; their demands could reduce the availability of shares for substantive short-selling.[11] This interaction, if important in practice, could provide a further reason why we don't find evidence of increased targeting of pilot firms.

B. The Intended Experiment: Data and Sample Cleaning

To investigate whether the uptick rule should be abolished, the SEC implemented its first – and to date only -- randomized experiment, known as the SHO Pilot Program (because the uptick rule was part of Regulation SHO). The experiment was announced on July 28, 2004, and ran from May 2, 2005 to July 3, 2007. The SEC's Office of Economic Analysis conducted its own study of the experiment. The results of this study, together with studies conducted by other parties, led the SEC to abolish the uptick rule in July 2007.

The SEC employed the following procedure for its experiment. It took the list of Russell-3000 firms as of sometime in 2004 -- we are not sure exactly when. It sorted the firms in the order of descending trading volume from June 2003 through May 2004, and picked every third

[11] We thank Charles Lee for suggesting this interaction.

firm for treatment. This is not strict random selection, but should be close to random in practice.[12]

1. Identifying Treated and Control Firms

To construct our own sample, we begin with a list of the Russell 3000 firms as of June 30, 2004.[13] This is a simple list of company names and ticker symbols. This may not be the exact list the SEC began with, but should be close; the SEC's actual list is not known.[14] We then match this list to Compustat. This is an intensive, often manual process. Many matches are clean (same ticker, and same or similar company name). But in about a quarter of the cases, we find similar company names with different tickers; or same tickers with different company names. We hand-match the remaining Russell 3000 firms to Compustat, relying on a combination of: (i) searching Compustat for a "Compustat company name" similar to the "Russell company name", and matching if the company is the same, even if the ticker is different; (ii) a web search for firm name changes. This produces 2,967 matches, out of 3,000 possible (a 99% match rate).

We then review the three firm lists available from the SEC: fully treated firms, which the SEC termed Category A; partly treated firms, which the SEC termed Category B, and firms that the SEC later excluded from these two categories.[15] We could not locate an SEC list of full control firms. The SEC did not explain why it excluded some firms from the experiment, but an internet search for events involving the excluded firms indicates that many exclusions reflect firms merging, going private, or going bankrupt.

The SEC lists of Category A, Category B, and excluded firms contain only company names and tickers. We can cleanly match some firms in the SEC's list to the merged Russell-Compustat list based on similar company name and a match on either the Russell or the Compustat ticker. Once again, the matching is imperfect: a nontrivial number of firms have an

---

[12] OEA (2007) explains: "Stocks were selected for this sample by sorting the 2004 Russell 3000 first by listing market and then by average daily dollar volume from June 2003 through May 2004, and then within each listing market, selecting every third company starting with the second.

[13] Source: Bloomberg maintains these historical lists of Russell 3000 firms.

[14] We obtained monthly lists of the Russell 3000 for this period; the SEC's actual division of firms does not precisely match any of them. See details below.

[15] Source: **[*to come].**

"SEC ticker" that does not match either the Russell or the Compustat ticker! Company names exhibit similar problems. We again engage in extensive hand matching, based on similar company name and similar ticker. We are able to match 2,890 firms across Russell 3000, Compustat, and the SEC. We then remove from the fully and partly treated groups the 99 firms that the SEC excluded from these groups. We cannot similarly exclude full control firms. This could introduce bias, but the bias should be minor given the small number of excluded firms. We confirm below covariate balance between the fully treated firms and the original controls.

The 3-way matching across Russell 3000, Compustat, and the SEC leaves a sample of 2,791 firms, matched across all three sources (1,272 full controls; 611 partly treated; 908 fully treated; overall 96% of the 2,901 possible firms). Of these firms, 2,639 have market capitalization in Compustat at year-end 2004 (1,219 full controls; 563 partly treated; 857 fully treated firms), of which 2,475 firms also have market capitalization at year-end 2005 ("2004-2005 Sample). This sample is substantially more complete than in most prior papers.[16]

We identify firms with publicly traded options using option-trading data from OptionMetrics. At May 2005, of the 2,791 firms in our sample, 1,720 had traded option (60%); this percentage rises to 63% if we focus on the 2,475 firms in the 2004- 2005 Sample (1,560 firms with traded options). The Plausible Subset of Small Experiment firms, within this 2004-2005 Sample but without public options includes 253 fully treated and 561 full control firms.

2. Data on Covariates and Outcomes

The principal outcomes we examine here are open short interest, available from Compustat, and share returns, from the Center for Research on Security Prices (CRSP). We take financial variables from Compustat. We examine all covariates and outcomes for outliers and missing values. We address large outlier values for some covariates by winsorizing sales growth at 100%, R&D/sales at 0 and +1.00, and EBIT/assets and EBIT/sales at ±0.5, and trading volume at 1%/99%. We also replace missing values for R&D and advertising with zero, and negative values for sales with +1 (in units of $ millions).[17]

---

[16] For example, we have 857 fully treated firms, compared to 651 in Grullon, Michenaud and Weston (2015).

[17] We make these decisions in a design phase of the analysis, with outcomes hidden.

C. Summary Statistics for Intended Randomization

Table 1, Panel A contains summary statistics for firm characteristics for fully treated firms versus original controls. All variables are at year-end 2004, or for the year 2004, except for sales growth, which is percentage sales growth from 2003 to 2004. All variables are from Compustat, except trading volume, which is from CRSP. The table is based on the 2,639 firms in the sample with market capitalization in Compustat at year-end 2004. The fully treated and original control firms are similar in all observable characteristics: size, leverage, profitability, trading volume, the number of issued shares, and so on. The fully treated firms are slightly larger; this likely reflects the slightly nonrandom way the SEC chose the fully treated firms. From Table 1 and other unreported tests for covariate balance, we are satisfied that the fully treated firms were effectively randomly selected.

Figure 1, Panel A, provides a histogram of ln(market capitalization), winsorized at 1%/99%, for the two groups. There is no evidence of imbalance on this size measure.

D. The Actual (Busted) Experiment and News Coverage of the Experiment

The SEC's randomization procedure was compromised. After completing the original randomization, the SEC moved the larger original controls – those in the Russell 1000 (below, "R1000") -- into a separate category of "partly treated" firms, which it called "Category B Pilot Securities". These firms were subject to short-sale restrictions during the trading day, but exempt from these restrictions in after-hours trading, from 4:15pm ET until the opening of trading the next day. That is, they were regulated like full controls during the trading day and like fully treated firms in after-hours trading. The SEC designated the remaining, full control firms (from the Russell 2000, below "R2000") as "Category C".

The separation of original controls into partly treated and full controls was stated in the SEC's release announcing the experiment, published in the Federal Register. Nonetheless, the experiment as a whole attracted very little press attention, and the separation of original controls into two subgroups attracted no press attention at all. We searched for news stories and other information about the experiment, both during the launch period (July 2004 - May 2005) and the experiment period (May 2005-May 2007). The announcement of the experiment, in July 2004, attracted no news attention whatsoever. There was very limited news coverage of the experiment between then and the actual launch, but that was limited to technical explanations.

We found no evidence of political opposition from firm managers. The SEC's July 2004 announcement of the pilot was not covered in any of the standard business news sources (including the Dow Jones News Service, the Wall Street Journal, and the New York Times).[18] A New York Times story about the SEC's proposal to repeal the short-sale rule, from late 2006, explains:[19]

> You may not have read of this proposal. It was virtually ignored by the news media, and if any companies are upset about it, they have not made themselves known. A pilot program that exempted some companies from the so-called uptick rule starting in 2005 drew little attention.

This is the *only* New York Times story about the experiment we found. The only Wall Street Journal story we found was a short, technical story in April 2005, simply copying a Dow Jones Newswire announcement of the experiment launch.[20]

We found no news coverage of the separation of original controls into partly treated and full controls. That separation was mentioned in a few law firm newsletters commenting on the experiment, and nowhere else.[21] As noted above, all prior researchers, including the SEC's own OEA, ignored this separation, and the SEC's then Chief Economist was unaware of it.

In Figure 1, Panel B, we provide a histogram comparing the two subsamples that the SEC formed from the original controls: full controls and partly treated firms. The separation on size, measured by market capitalization, is nearly complete. This confirms that there was not a single treated group and a single control group, as all prior studies have assumed. Instead, there were three groups: fully treated firms, partly treated firms (drawn from original controls in the R1000), and full control firms (drawn from original controls in the R2000). Moreover, assignment to partial treatment was highly non- random; the partially treated firms are larger than the full controls.

---

[18] Appendix Table App-1 lists all of the business press stories we found. The first story was on 30 November 2004. The Dow Jones News Service published a short summary of a delay in launching the experiment, to let the exchanges make programming changes. This was *not* a separate story about the SHO Pilot; it was appended to a main story about another SEC rule.

[19] Floyd Norris, 70 Years Later, A Scapegoat Gets a Break, *New York Times* (Dec. 8, 2006).

[20] SEC Pilot Program To Halt `Uptick' Rule, *Wall Street Journal*, April 29, 2005.

[21] Appendix Table App-2 lists the law firm newsletters that we found. A couple of these newsletters were published in the legal press, in specialized journals that were unlikely to attract much investor attention.

This produces very large covariate imbalance if one compares all treated firms to either the partially treated group or the full control group. Table 1, Panel B provides a covariate balance table for fully treated and full control firms. Fully treated firms are significantly larger than full controls, by every measure. "Ratio measures," such as sales growth or advertising/sales are less dissimilar, but some are meaningfully different. For example, full control firms have higher short interest (as a percentage of outstanding shares), than fully treated firms. Table 1, Panel C compares fully treated and partly treated firms. The difference are again dramatic. Partly treated firms are much larger than fully treated firms. This time, short interest is significantly higher for the fully treated firms.

In short, the SEC's separation of intended controls into partly treated firms and full controls compromised its initial randomization. Only in a narrow range of *ln*(market cap) $\in$ [7.2, 7.8] is there significant overlap between partially treated and full control firms. Even there, the division of intended controls into partly treated and full controls is not random, and instead depends on whatever factors led the purveyors of the Russell indices to place some firms in the R1000, but not others.

There is no easy solution to this problem. We cannot compare the full set of treated firms to full controls, to partly treated firms, or to original controls. The nonrandom division of intended controls into partly treated and full control groups compromises both comparisons.

E. Partial Workaround: Separate Small-Firm and Large-Firm Experiments

There is, however, a serviceable workaround, that restores randomization. In effect, the SEC ran two different randomized experiments. For R2000 firms (almost all with *ln*(market cap) < 7.8), the SEC ran the experiment everyone thought it ran, which we term the Small-Firm Experiment, which compares fully treated to full control firms. For R1000 firms (almost all with *ln*(market cap) > 7.2), the SEC ran the Large-Firm Experiment, which compares partly treated to fully treated firms.

We reconstruct those two experiments as follows. First, we divide the sample into "pseudo" R1000" and R2000 subsets. This division is imperfect because the SEC's division of firms into these subsets was imperfect. The R1000 and R2000 membership changes monthly; the SEC release does not specify which specific list the SEC used. But *none* of the monthly lists for 2004 matches the SEC's division of original controls into partly treated and full controls, or comes particular close. For example, if we compare the 1,692 original controls (563 partly treated; 1,219 full controls) in the 2004 Sample to the Russell lists from June 2004, shortly before the July 28, 2004 experiment announcement date, there are:

517 partly treated firms in R1000 (this matches the rule the SEC claimed to follow)

1,175 full controls in R2000 (matching the SEC's asserted rule)

But there are also:

46 partly treated firms in R2000 (which violates the SEC's asserted rule)

44 full controls in R1000 (again violating the SEC's asserted rule)

No other plausible monthly Russell 3000 list does better. Given this data, we decided to:

(i) define the Large-Firm Experiment as comparing (a) all fully treated firms in the R1000 with ln(market cap) > 7.2 to (b) all partly treated firms in this size range. The resulting subsample contains 899 firms (295 fully treated; 604 partly treated);

(ii) define the Small-Firm Experiment as comparing (a) all fully treated firms in the R2000 with ln(market cap) < 7.8 to (b) all full control firms in this size range. The resulting subsample of the 2004 Sample contains 1,786 firms (569 fully treated; 1,217 full controls).[22]

(iii) treat the remaining xx firms as not within either experiment. These firms consist of: xx fully treated firms in the R1000 with *ln*(market cap) < 7.2; yy fully treated firms in the R2000 with ln(market cap) > 7.8; zz partly treated firms with *ln*(market cap) < 7.2; and zz full control firms with ln(market cap) > 7.8.

Table 2 provides separate covariance balance tables for the Small-Firm and Large-Firm Experiments. The Large-Firm Experiment shows good balance; only two of the *t*-statistics for means are above 2 and barely so; one of the p-values for medians is slightly below 0.05. There is mild evidence for the Small Firm Experiment that the fully treated firms were a bit larger than the full controls, but the differences are economically small and could plausibly arise by chance. Table 2 provides evidence that despite some research design fuzziness in the middle-size range with 7.2 < ln(market cap) < 7.8, our workaround recovers two separate apparently randomized experiments.

[22] We considered creating a third "middle" group, containing all firms with 7.2 < *ln*(market cap) < 7.8 – the range with meaningful overlap between the partly treated and full control groups. In this range, one could ignore then actual decision rule that the SEC imperfectly applied and use methods for pure observational studies to seek to imperfectly replicate the (assumed unobserved) rule governing assignment of original controls to partly treated versus full controls. For example, one could use financial characteristics to estimate the propensity of each original control firm to be assigned to the partly treated or full control group, and then use inverse propensity weights to estimate three treatment effects: fully treated versus full control; fully treated versus partly treated; and partly treated versus full control firms. However, the sample size in the overlap range is too small to make that a useful exercise.

F. Methods for Analyzing the Direct Effects of the SEC Experiment

If one puts recognizes that the SEC ran two experiments on different subsets of firms, studies each subset separately, and studies only the direct effects of the experiment on traders' behavior – including total trading volume, short-sale volume, monthly open short interest, bid-asked spreads, and the speed of price response to new information – there are several ways to estimate an unbiased causal effect. One is a simple comparison of means during the treatment period. The difference in means can be estimated by regressing the outcome on a treatment dummy, using either cross-sectional or panel data. The regression equation is:

$$y_{its} = \alpha_s + \lambda_s * \tau_{is} + \epsilon_{its} \qquad (1)$$

Here y is the outcome, τ is the treatment dummy, s indicates the subsample (large firms, small firms, or a further subset such as large NYSE firms or small NASDAQ firms), i indexes firms, and t indexes time.

A more careful analysis would use difference-in-differences (DiD) analysis with firm and period fixed effects, using panel data covering the pre-treatment and treatment periods (and potentially the post-treatment period as well). One would regress the outcome on an interaction between a treatment dummy and a treatment period dummy:

$$y_{its} = \alpha_s + f_i + g_t + (\lambda_s * \tau_{is} * P_t) + \epsilon_{its} \qquad (2)$$

Here, f_i and g_t are the firm and time fixed effects and P_t is the treatment period dummy (=1 during the treatment period, 0 before and after). The non-interacted treatment dummy and treatment period dummy are absorbed by the firm and period fixed effects. The outcomes of interest are likely to be persistent within firm, across time, so standard errors should be clustered on firm.

Below, we also present leads-and-lags graphs based on regressions in which we allow the treatment effect to vary by period, both before and after the start of the treatment period:

$$y_{its} = \alpha_s + f_i + g_t + + \sum_{k=-m}^{+n} (\beta^k * D_{it}^k) + \varepsilon_{its} \qquad (3)$$

Here, k indexes months relative to the start of the experiment in May 2005. $D_{it}^k = 0$ for control states for all t and k. For treated firms, $D_{it}^k = 1$ for the k^{th} month relative to May 2005, 0 otherwise. For example, $D_{it}^4 = 1$ four months after the experiment start, 0 otherwise. Therefore, β_1 provides the effect of the experiment in the first month after it starts, and β_{-1} is the estimated effect one month before it starts. We adjust the coefficients so that the coefficient just before the experiment is announced is zero: $\beta_{-11} \equiv 0$.

One could add to these regressions a vector of time-varying covariates $\mathbf{x}_{it}$ that are unlikely to be affected by treatment. For a randomized experiment, adding such "proper" covariates can increase precision without introducing bias.

IV. Complexity for Studing Indirect Effects

A. Indirect Effects of Short Selling on Firms and the Need for Causal IV Analysis

To study second-order effects, such as the impact of short selling on share price and other capital market characteristics, one must proceed more carefully. The logical progression of causation starts with direct effects. Any second order effect of relaxing the uptick rule on price is likely to be mediated by a change in open short interest; so are any third- and fourth-order effects. If open short interest does not appear to change – if there are no or very few apparent compliers, we should not expect indirect effects either. Thus, a careful research design should begin with a close look at open short interest. Below, we find no evidence for a change in open short interest.

If there is evidence of a significant number of compliers -- firms which are not only assigned to treatment, but are actually treated (targeted by substantive short-sellers), one can conduct further consistency checks:

> Any third- and fourth-order effects are likely to be mediated by an effect of targeting on share price. Does such an effect exist?
>
> We should expect a second-order effect on share price only for targeted firms, and can test whether that expectation is met. One can compare targeted firms in the control group (who are always-takers) to targeted firms in the treated group (a mix of compliers and always takers). In a DiD analysis of these firms, is there a relative change in share price for the targeted firms in the treated group?

Below, we find no evidence of a change in share prices.

Could there be other channels for a causal effect? One possibility is an effect thru managers' heads: Maybe managers of treated firms were afraid of bear raids, even though the raids never came. This is possible in the abstract, but is it plausible, and does it lead to testable predictions?

If the channel is not actual short selling, but only manager fear, we would expect that fear to be reflected in publicity about the proposed rule, and in manager opposition to the rule. As discussed above, the SHO Pilot was notable for the near-complete absence of publicity, and we know of no political opposition. We might also expect any effect to shrink over time, as managers realized the bears are not charging, with publicity about that too. That is testable, but none of the third-and fourth -order papers considers a manager fear channel.

A manager fear channel has further problems with testability. Managers' reactions would likely not be independent, yet independence is a core condition for valid causal inference – part of the "Stable Unit Value Treatment Assumption" (SUTVA) (Imbens and Rubin, 2015, ch. 1). And the policy response would likely be to educate managers about short-seller behavior, rather than an ineffective uptick rule. So the manager fear channel has major problems, and we are back to the natural channel: Relaxing uptick rule → more targeting → lower share price for targeted firms.

B. Causal IV Analysis

To estimate the effect of treatment on the compliers, the proper method is causal IV, using assignment to treatment (inclusion in "Category A") as an instrument for actual treatment (targeting by short-sellers). One uses standard two-stage least squares (2SLS) analysis. The first stage question is whether assignment to treatment predicts targeting. If not – as we find – there is no reason to expect a second-stage, and analysis should ordinarily stop here.

To be sure, it is mechanically possible to ignore the absence of a first-stage, and use DiD methods to compute an intent-to-treat effect, but it will rarely make sense to do so, as any results are likely to be spurious. This is especially so here because a natural second-order effect – the effect of targeting on share price – is also absent. This makes the third- and fourth-order effects studied in many papers still less plausible.

If one nonetheless wanted to proceed with an intent-to-treat analysis – we ourselves would not -- we believe it should be incumbent on researchers to justify why the outcome they study is likely to be affected by relaxing the short sale rule, despite the absence of evidence for a first stage. We also believe that one's priors should tilt against believing any effect that is found. How strongly? There is no clean answer, but one approach would derive from McCrary, Christensen and Fanelli (2016). They show that if publication bias results in only statistically significant results (2-tailed test, 5% significance level) being published, one should require a t-statistic of around 3, rather than 2, to in fact have 95% confidence that the result is real.

Concerns with specification searching or data mining would suggest an even more conservative approach.

C. Causal IV with Covariates

We have this far discussed Causal IV without covariates. If one wants to add covariates, or assess how second-order effects vary with observable firm characteristics, ordinary two-stage-least-squares (2SLS) analysis is insufficient, because it assumes -- implausibly -- that treatment effects are homogeneous, and that the proportion and nature of the compliers in the sample does not depend on the covariates. One can relax both assumptions and estimate a "local average response function", conditional on covariates (Abadie, 2003). This approach is rarely adopted in practice – the only example we know about is Angrist and Fernandez-Val (2013). The cost is additional complexity relative to 2SLS, and less precision in one's estimates, but the precision offered by standard 2SLS was spurious anyway. The papers studying the indirect effects of the SHO Pilot often include covariates in their intent-to-treat estimates; but none undertakes an explicit 2SLs analysis, let along the further step of estimating a local average response function.

D. Assessing Third- and Fourth-Order Effects

If one wants to study third- and forth-order effects, such as the impact of short selling on firm behavior and the behavior of other parties (like auditors), standard causal IV methods become problematic as well. Consider, for example, the effect of short selling on capital raising, studied by Grullon, Michenaud, and Weston (2015). Suppose there were a plausible number of compliers, sufficient to let researchers use assignment to full treatment as an IV for targeting of firms by substantive short sellers. Would that be enough? Or would one also need to show that targeting impacted share price?

If share prices fall, the causal story is easy to understand: (relax uptick rule → targeting of complier firms → lower share prices for these firms → reduced capital raising). In contrast, without a change in share prices, the logic for an effect on capital raising is much less clear. The firm's ability to raise capital would not be directly affected. Perhaps firm managers would respond to evidence of short-seller targeting by seeking to raise more capital and expand. One might then want to confirm that any observed effects of capital raising appear only, or more strongly, for treated firms that experience share price declines, relative to control firms with similar declines.

As another example: to study the impact of short-selling on auditor behavior, due to (assumed) higher litigation risk, as in Hope, Hu, and Zhao (2016). One would again want to assess the consistency of the assumed causal chain with the available evidence. For example: Does assignment to treatment predict litigation risk? If yes, does litigation risk rise principally for targeted firms? Principally for targeted firms which suffer share price declines?

E. Assessing SUTVA and the Only-Through Conditions

Even if all else were favorable for a study of indirect effects, one would still need to assess the only-through conditions for plausibility. Consider SUTVA first. It requires that short-sellers decision to target firm A is independent of their decision to target firm B. This implicitly requires that there be an unlimited supply of capital for short-selling. An alternative hypothesis is that, at least in the short- to medium-term, short-selling is a specialized niche, which limited available capital. If so, and if the uptick rule were a strong constraint on short-selling, then short-sellers would divert scarce capital to treated firms and away from control firms.

The only-through condition for treated firms requires that there be no placebo effect – that relaxing the uptick rule affects firms only through its effect on short-selling. Above, we considered the possibility of an effect through managers' fear of bear raids, even if the raids never came. Any such fear, if it existed, would both likely not be independent across firms (and thus a SUTVA violation) and would be an anti-placebo effect – firm managers would feel worse, instead of the same, if offered the placebo of a relaxed short-sale rule without actual targeting.

Similarly, one must assume away – and argue for the validity of the assumption – any effect of being in the control group. For example, would managers of control firms feel relief – perhaps because they believe there is limited supply of short-seller attention, and behave differently because the uptick rule remains in place for them? We suggest here not that these assumptions necessarily fail, but that they are serious, substantive assumptions about how the business of substantive short-selling works, and how managers react to the treat of being targeted. Their plausibility needs to be assessed.

V: Reconsidering Prior Analyses

Does the busted randomization invalidate the prior studies discussed in Part II? If we allow for arbitrary treatment heterogeneity, the answer must often be yes. Any result that derives from both small and large firms must be re-evaluated, treating the two groups as subject to two different experiments. We begin that re-evaluation here by assessing the evidence for an effect on open short interest, or on share price.

We also consider whether there is an apparent difference in the effects of the Small and Large Firm Experiments. One cannot directly compare the two without assuming homogeneous treatment effects, for some range of firm sizes. Still, the difference in treatment effect estimates for the Small Firm and Large-Firm experiments, for firms near the R1000 cutoff, provides an estimate of the effect of partial treatment versus no treatment. If there were little difference in treatment effect estimates across a variety of outcome measures, we might place some credence in the SEC lawyers' belief that partial treatment was similar to no treatment, and therefore not important. This is not the case, however, as we show below.

A. Impact of Short Sale Restrictions on Open Short Interest

Most prior research, including the SEC's own study (OEA, 2007), found little effect of relaxing the short sale rule on open short interest. We assess here whether those results would change, if we study the Small Firm and Large-Firm Experiments separately. We investigate open short interest with care, because it is central to the validity of studies of indirect effects.

1. Graphical evidence

First, in Figure 2, we plot month-end open short interest, separately for treated and control firms. Panel A shows results for the Small-Firm Experiment, from July 2003 (about two years before the experiment began) through year-end 2007 (about 6 months after it ended). Vertical lines indicate the start and end of the experiment period; a dotted line shows when the SEC announced the experiment. There is no visual evidence of any impact of either the announcement of the experiment, in July 2004 (not that we would expect any), or the onset of the actual experiment in May 2005. Below, we confirm the absence of any significant average effect with regression analyses.

Panel B is similar, but is limited to the subsample of small NYSE firms. An effect might be more likely for these firms, because the uptick rule had more bite for NYSE firms than for NASDAQ firms. There is again no evidence of an increase in open short interest for treated firms. During most of the treatment period, treated firms generally have *lower* open short interest than control firms (opposite from predicted).

Panel C is again similar, but covers small NASDAQ firms. These firms tend to have somewhat higher short-interest percentages than the small NYSE firms. Looking at Panel C alone, one might be tempted to see evidence of a rise in open short interest starting around the onset of the experiment. But several factors counsel against a causal interpretation. First, there is a similar gap in 2003-2004, before the SEC experiment was even announced. Second, the relative rise reverses near the end of the experiment period. And third, we see an opposite drift for small NYSE firms.

In Panel D, we provide results for the Plausible Subset of firms without publicly traded options. **[* results to come]**

While trends in the Small Firm Experiment were similar for treated and control firms, the substantive short-selling world was changing rapidly during this time period. The average percentage short interest more than doubled from 4% in mid-2003 to 9% in late 2007. Average monthly short interest also strongly increased for large firms during this period, from around 3% to 5%. We have seen no discussion in the academic or practitioner literature for this great increase in short interest, especially for R2000 firms.

The remaining panels show results for the Large-Firm Experiment. Panel E shows results for all large firms. There is no difference between the graph lines for fully treated versus partly treated firms during the period from announcement to experiment onset, or from onset through the first half of the experiment period. There is a drift toward higher open short interest among treated firms in the second half of the experiment period. But the magnitude is comparable to the drifts we saw for small firm subsamples. As we show below, that drift is driven by the NASDAQ subsample, within the Large Firm Experiment.

Panel F shows results for the subsample of large NYSE firms. There is no evidence that experiment affected average open short interest.

Panel G shows results for the subsample of large NASDAQ firms. The sample of treated firms is small, only 44 firms. But it appears that there is an increase in open short interest for fully treated firms, starting in early 2006, about 6 months into the experiment period. Here too, however, there are other factors that counsel against a causal interpretation. First there is a similar gap that opens in the first half of 2004, before the experiment was announced, and then closes around the time the experiment begins. Second, the gap remains in place in 2007 after the experiment ends, and the uptick rule is repealed for all firms. Third, as we discuss above, large NASDAQ firms are the *least* plausible group for finding an effect of the uptick rule.

In Figure 3, we provide an alternate presentation of the same time trends, using leads and lags graphs. We present selected graphs, and relegate others to the Appendix.

Panel A presents results for the Small Firm Experiment; Panel B presents results for the Large-Firm Experiment. In both, the standard error bars are large, and there is no evidence for an effect of the SEC experiment.

Across both experiments and across subsamples, then, there is no solid graphical evidence that the experiment affected average open short interest. Overall, open short interest increases greatly, and the gap in average short interest fluctuates, and apparently random divergences between the treated and control groups can appear and disappear, but not at times that the timing and substance of the experiment can robustly predict.

2. DiD Regression Results

We confirm the absence of an average effect on open short interest in Table 3. We present DiD regressions, with firm and month fixed effects, for small firms in Panel A and for large firms in Panel B. The core variable of interest is the interaction between treatment dummy and treatment period dummy. In each panel, the odd-numbered regressions include no covariates, the even-numbered regressions add selected covariates that plausibly predict open short interest; results with a broader set of covariates are similar (see Appendix). For the Small-Firm Experiment, the coefficient on the interaction term is positive but small and statistically insignificant. The coefficient is larger for the NASDAQ subsample, but remains insignificant. **[*to come for Plausible Subset]**.

For the Large-Firm Experiment as a whole, the coefficient on the interaction term is positive but statistically insignificant for all large firms and for the NYSE subsample, but negative and insignificant for the NASDAQ subsample.

B. The Distribution of Open Short Interest

Even if the SEC experiment did not significantly affect *average* open short interest, there could still have been an increase in targeting of treated firms, which is buried in the "noise" of unrelated changes in open short interest at other firms. This might lead to treated firms being more likely to have high levels of open short interest.

In Figure 4, we check for that possibility by comparing the distribution of open short interest percentages between treated and control firms. We compute kernel density plots monthly over 2003-2007, and present representative plots in Figure 4.

Panel A provides a kernel density plot for the Small Firm Experiment for June 2004, before the experiment was announced. We would not expect any difference between the two groups, and do not find any. Panel B provides a similar plot as of March 2005, during the period between announcement and launch of the experiment. We have no reason to expect an expansion in targeted short selling, prior to the experiment start, and there is no evidence of any. Plots for other months during this period are similar.

Panel C provides a kernel density plot for the Small Firm Experiment in September 2005, after the experiment had been underway for five month -- long enough for any effect of the experiment on open interest to be likely to show up. An increase in the number of treated firms with “large” levels of open short interest, corresponding to firms being targeted by short-sellers, should be visible in this plot. There is, however, no evidence of a difference between the treated and control groups in the density of firms with high levels of open short interest. Plots for other months during the experiment period are similar. A Kolmogorov-Smirnov test cannot reject the null that the distributions of treated and control firms are the same. We cannot rule out the potential for relaxed short-sale restrictions to have facilitated bear raids on individual firms, but Figure 6 provides evidence that this must be quite infrequent, if it happens at all.

Figure 4, Panels D, E, and F provide kernel density plots for the Large Firm Experiment, and are otherwise similar to Panels A, B, and C. Plots for other months in each period are similar. There is no evidence of targeting of fully treated firms, versus partly treated firms.

In sum, there is no evidence that relaxing the uptick rule led to greater targeting of treated firms. The direct effect of the SEC experiment on open short interest, which in our view needs to be there for indirect effects to be plausible, is just not present. If there are any complier firms, they are so scarce that we cannot find them.

C. The Impact of Short-Selling Restrictions on Volume of Short Sale Trades

The SEC’s Office of Economic Analysis found that treated firms had higher short selling volume, but not higher outstanding short interest. We have confirmed the second finding, separately for the Small-Firm and Large-Firm Experiments.

For the first finding, on short-sale volume, the OEA helpfully reports results divided into market capitalization deciles. We focus here on the results for NYSE firms. Deciles 1-6 correspond to the firms that were subject to the Small-Firm Experiment, deciles 8-10 correspond to the firms that were subject to the Large-Firm Experiment, and decile 7 corresponds to the limited middle range, in which all three groups of firms exist in significant numbers.[23]

In Figure 5, we show the OEA's results for short sale volume, as a percentage of total trading volume in Figure 5. Vertical lines separate the three treatment regimes. The difference in treatment effects is readily apparent. The average coefficient for the Small-Firm Experiment, in deciles 2-6, is 3.64. Thus, the experiment apparently caused an increase in the ratio of short-sale trades to all trades of around 3.5%. However, the average effect for the Large-Firm Experiment, in deciles 8-10, is only 0.74. Thus, if we put aside the smallest decile, the estimated Small Firm treatment effect is five times as large as the Large Firm effect. If there are roughly homogeneous treatment effects across mid-sized and larger firms, this implies that, at least when it comes to the arbitrage trading that drives most short-sale volume, partial treatment is much closer to full treatment than to full control.[24]

This evidence suggests that, when it comes to the volume of arbitrage-related short-selling, the Large-Firm Experiment was scarcely an experiment at all. It compared two treatments that differed only slightly. We should not expect this weak experiment to tell us much about the importance of short-sale restrictions for large firms.

However, there was a much larger increase in short-sale trading for treated firms in the Small-Firm Experiment.

We caution that the small firm results do not imply that total short sale activity in fact increased for small firms. Perhaps it only migrated from regional exchanges to each firm's main exchange. To assess whether that happened, one would need to: (i) determine which small firms were also traded on regional exchanges; (ii) assess whether the treatment effects were similar for firms which were versus were not listed on regional exchanges, and (iii) for firms which were listed on regional exchanges, study short selling on those exchanges (assuming data exists).

---

[23] In our sample, there are only xx partly treated firms in decile 6 and only yy full control firms in decile 8.

[24] We do not present confidence intervals because the OEA does not report standard errors, only stars for significant results, and the stars are suspect because they are based on repeated daily regressions, without firm clusters, yet clustering on firm is clearly needed (Petersen, 2009).

D. Impact of Short-Selling Restrictions on Share Prices

Prior studies have found no effect of the short-sale experiment on average share prices (Diether, Lee, and Werner, 2009), extreme price changes (OEA, 2007), or price changes following negative earnings announcements (Bai, 2008). In Figure 6, we investigate whether the short-sale experiment affected average share price levels, considering the two experiments separately. The two panels of this figure present, for the Large-Firm and Small-Firm Experiments, average biweekly returns to the treated group, versus the control group, together with 95% confidence intervals (CIs).[25] We limit the sample to firms with share price > $5 at **[*date]**, in order to exclude "penny stocks," but obtain similar results without this exclusion. A vertical dotted line indicates when the SEC announced the experiment; solid lines show the start and end of the experiment period.[26]

If investors expected the lifting of short-sale restrictions to affect share prices once the experiment started, we might see a price response soon after the short-sale experiment, including the identities of treated firms, was announced in July 2004. Alternatively, if investors were unsure, when the experiment was announced, what effect the short-sale experiment would have on share prices, there might be an effect only when the experiment began in May 2005. Our prior, given reasons to doubt the effect of the uptick rule on substantive short-selling, and the lack of an apparent effect on open short interest, is to expect no significant price changes. If any were to occur, we have no prediction on whether they would be positive or negative.

In both panels, we see no evidence of a relative price change, when the experiment is announced, when it begins, or when it ends. Consider first the Small-Firm Experiment. A few biweekly relative returns are significant at the 5% level, but no more than would be expected from chance alone. None is especially large; indeed the largest biweekly return in magnitude is only 0.15% -- a tiny amount – and one that occurs after the experiment ends. Similarly, for the Large-Firm Experiment, there is no evidence of a share price reaction at the time of announcement, or the time of experiment onset.

---

[25] In unreported robustness checks, we obtain consistent results if we study either weekly or monthly returns.

[26] We use biweekly returns as a compromise between granularity and clutter. Results using weekly returns are similar to those we present. In unreported robustness checks, we investigated whether there was any effect of the short-sale experiment on share prices separately for NYSE and NASDAQ subsamples, within the Small-Firm and Large-Firm Experiments, and found no evidence of any effect. We also examined all firms, instead of only those with trading prices > $5; results were again similar.

Grullon, Michenaud and Weston (2015) report a price drop for small treated firms, during the period just *before* the SEC announced the experiment. And indeed, in Figure 4, Panel A, there is a statistically significant negative return to treated small firms just before the experiment is announced, with no similar significant return for large firms. Is this real?

We view the Grullon et al. result as highly implausible. Our explanation for why is a specific example of a core theme of this paper: There are often plausibility checks available for causal claims. Those checks should be searched for, and pursued when available.

First, the result is not very strong, as a statistical matter. The particular 2-week return that they focus on is negative and statistically significant for small firms, but: (i) the economic magnitude is tiny, at around 0.1%; (ii) a return of this magnitude could easily occur by chance; (iii) there are negative returns to the treatment group of similar magnitude for other 2-week periods, one in the middle of the experiment period in 2005, and one well after the experiment ends in 2007; and (iv) there is no similar negative return for large firms.

But the real problems with plausibility lie elsewhere. For the negative returns to small firms prior to the SEC announcement to be a real effect, rather than a random event:

(i) Someone needed to anticipate the SEC announcement, expect the treatment effect to be negative, and trade on this anticipation in large enough volume to move prices down and keep them down when the experiment was announced.

(ii) That someone would have to have *stolen* the exact SEC list of treated firms from the SEC. Even knowing in advance the rule the SEC would follow (one firm in three would be treated, with firms ranked by market capitalization as of a then-secret date) would not help, because the SEC exempted many firms from treatment for firm-specific reasons.[27]

(iii) That short selling must have somehow *not* been led to higher open short interest for treated firms on July 31, 2004, just after the experiment was announced (see Figure 2).

(iv) When the SEC released the actual list of treated firms, there was no significant price response. Thus, for the return two weeks earlier to be a true response to the future announcement, (a) the thief must have perfectly anticipated the actual market reaction by other investors, and (b) other investors must have realized that the effect of the experiment had already been impounded into share prices, through trading by an unknown thief, and therefore not traded further on their own views.

---

[27] The SEC often holds market-sensitive information; for example information about pending investigation of particular firms. We are unaware of instances of leakage or theft of this information.

All this is implausible, at the least. We're tempted to say, "impossible."

E. Impact of Short-Selling Restrictions on Daily Price Changes

The OEA also found that treated NYSE firms had smaller daily price changes than control firms, this time with the effect concentrated in larger firms. We present the OEA's results, by size decile, in Figure 7. We drop the lowest decile, which behaved oddly and very differently from other deciles across many of the SEC's measures. Deciles 2-6 correspond to the firms that were subject to the Small-firm experiment, deciles 8-10 correspond to the firms that were subject to the Large-firm experiment, and decile 7 corresponds to the middle range with all three groups. There some -- albeit less clear than in Figure 5 -- apparent difference in treatment effects between the two experiments is again readily apparent. The average effect for the Small-Firm Experiment, in deciles 2-6, is general small and averages +0.37. The average effect for the Large-Firm Experiment, in deciles 8-10, is -1.49. If we assume roughly homogeneous treatment effects across mid-sized and larger firms, one might find this result if relaxing short-sale restrictions fully did not affect daily price changes, but relaxing them only after hours pushed some trading and accompanying volatility into after-hours trading, leading to less volatility during the trading day. Here too, the Large-Firm and Small-Firm Experiments appear to have different treatment effects. Moreover, if our explanation for the Large Firm result is correct, then prior studies which find lower volatility during the trading day are suspect. There would be no true drop in share price volatility, only migration of volatility from within the trading day to after hours.

F. Prior Results Which Are Likely to Survive the "Busted Randomization" Problem

Our results above cast doubt on many prior studies of the SEC experiment, especially more recent studies, which focus on indirect effects. Thus, it is worth stating that some results seem likely to survive. One notable example is the Alexander and Peterson (2008) study of the effect of the experiment on treated firms in the *smallest* 20% of the Russell 3000 (by market capitalization), during market stress days, in which overall share price indices decline 1% or more. They find no effect on measures of liquidity, volatility or price efficiency. These results involve only firms within the Small-Firm Experiment, and thus are not by the busted randomization for R1000 firms. Moreover, many of the firms in the smallest 20% were in the Plausible Subset, for whom an effect of the short-sale experiment is more likely.

Some of their outcome variables, such as trading volume of trading and execution speed, involve direct effects, do not require causal IV analysis, and can stand as is. However, other outcome variables, those related to changes in market quality, involve second-order effects, and could usefully be re-examined using causal IV methods.

VI. Conclusion

The SEC's OEA designed and, during 2005-2007, ran the only randomized experiment in the SEC's history. The SEC removed short-sale restrictions from one-third of the firms in the Russell 3000 at random. The SEC's legal staff then busted the randomization, by making the larger control firms partly treated. This was not hidden, but the lawyers never told the economists (who didn't know about the busting until we told them), and the lawyers' action was obscure enough so that no prior researchers have noticed it. The actions by the SEC's legal staff created, in effect, two experiments: a Small-Firm Experiment for R2000 firms with *ln*(market cap) < 7.2, which compares fully treated to full control firms; and a Large-Firm Experiment for R1000 firms with ln(market cap) > 7.8, which compares fully treated to partly treated firms.

We study those two experiments, and find no evidence that either of them affected substantive short selling. Neither experiment affected either open short interest (the best available marker for substantive short selling, and neither had an apparent effect on share prices. The many papers which find indirect effects of an experiment which had such limited direct effects are thus suspect, but re-evaluating them is a task for a follow-up paper, now in progress.

References

Abadie, Alberto, 2003. Semiparametric Instrumental Variable Estimation of Treatment Response Models, 113 *Journal of Econometrics* 231-263.

Alexander, Gordon J. and Mark A. Peterson (2008), The Effect of Price Tests on Trader Behavior and Market Quality: An Analysis of Regulation SHO, 11 *Journal of Financial Markets* 84-111.

Angrist, Joshua D., and Ivan Fernandez-Val (2013), ExtrapoLATE-ing: External Validity and Overidentification in the LAE Framework, in Daron Acemoglu, Manuel Arellano, and Eddie Dekel eds., Advances in Economics and Econometrics vol. 3, Econometrics 401-433. Cambridge University Press.

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin (1996), Identification of Causal Effects Using Instrumental Variables, 91 *Journal of the American Statistical Association* 444-455.

Atanasov, Vladimir, and Bernard Black (2016), The Trouble with Instruments: Re-examining Shock-Based IV Designs, working paper, at http://ssrn.com/abstract=2417689.

Bai, Lynn (2008), The Uptick Rule of Short Sale Regulation—Can it Alleviate Downward Price Pressure from Negative Earnings Shocks, 5 Rutgers Business Law Journal 1-63.

Chang, Eric, Tse-Chun Lin, and Xiarong Ma (2015), Does Short-selling Threat Discipline Managers in Mergers and Acquisitions Decisions?, working paper.

De Angelis, David, Gustavo Grullon, and Sebastien Michenaud (2017), "The Effects of Short-Selling Threats on Incentive Contracts: Evidence from an Experiment", *Review of Financial Studies* (forthcoming), working paper at http://ssrn.com/abstract=2238236.

Deither, Karl, Kuan-Hui Lee, and Ingrid Werner (2009), It's SHO Time! Short-Sale Price-Tests and Market Quality, 64 *Journal of* Finance 37-73.

Deng, Xiaohu, Lei Gao, and Jeong-Bon Kim (2017), Short Selling and Stock Price Crash Risk: Evidence from a Natural Experiment, working paper, at http://ssrn.com/abstract=2782559.

Fang, Vivien W., Allen Huang, and Jonathan Karpoff (2016), "Short Selling and Earnings Management: A Controlled Experiment", 71 *Journal of Finance* 1251-1293.

Frangakis, Constantine E., and Donald B. Rubin (2002), *Principal Stratification in Causal Inference*, 58 Biometrics 21-29.

Grullon, Gustavo, Sebastien Michenaud, and James Weston (2015), The Real Effects of Short-Selling Constraints, 28 *Review of Financial Studies* 1737-1767.

Harmon, Dion, and Yaneer Bar-Yam (2008), Technical Report on SEC Uptick Repeal Pilot, New England Complex Systems Institute Technical Report 2008-11.

He, Jie (Jack), and Xuan Tian (2016), Do Short Sellers Exacerbate or Mitigate Managerial Myopia? Evidence from Patenting Activities, working paper, at http://ssrn.com/abstract=2380352.

Hope, Ole-Kristian, Danqi Hu, and Wuyang Zhao (2016), Third-Party Consequences of Short-Selling Threats: The Case of Auditor Behavior, xx *Journal of Accounting and Economics* yyy-zzz.

Imbens, Guido W., and Donald B. Rubin (2015), *An Introduction to Causal Inference in Statistics, Biomedical and Social Sciences*.

Kim, Pureum, and You-il (Chris) Park (2015), Regulatory Short Selling Constraints and the Sensitivity of Corporate Investment to Price, working paper

Li, Yinghua and Liandong Zhang (2015). "Short Selling Pressure, Stock Price Behavior, and Management Forecast Precision: Evidence from a Natural Experiment", 53 J*ournal of Accounting Research* 79-117.

McCrary, Justin, Garrett Christensen, and Daniele Fanelli (2016), Conservative Tests under Satisficing Models of Publication Bias, 11 *PLOS One* e0149590, doi: 10.1371/journal.pone.0149590**.**

Petersen, Mitchell A. (2009), Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches, 22 *Review of Financial Studies* 435-480.

Povel, Paul, and Michael Raith (2001). "Optimal Investment Under Financial Constraints: The Roles of Internal Funds and Asymmetric Information", University of Chicago mimeo.

Securities and Exchange Commission, Office of Economic Analysis (2007), Economic Analysis of the Short Sale Price Restrictions Under the Regulation SHO Pilot.

Romano, Roberta (2013), Regulating in the Dark, 1 *Journal of Financial Perspectives* 1-6.

Rubin, Donald B. (2008), *For Objective Causal Inference, Design Trumps Analysis*, 2 Annals of Applied Statistics 808-840.

Wu, Juan (Julie) (2006), "Uptick rule, short selling and price efficiency", working paper,

Figure 1, Panel A. Histogram: Fully Treated vs. Original Control Firms

Effect of SEC's initially intended randomization, in which it divided firms into Fully Treated and Original control groups. Histogram shows number of firms in each group with different levels of *ln*(market capitalization at year-end 2004). Market capitalization is winsorized at 1%/99%, corresponding to *ln*(market cap) of 2.97/11.34 (market cap of $19.6M/$84.2B). Sample is 1,783 original control firms and 857 fully treated Firms.

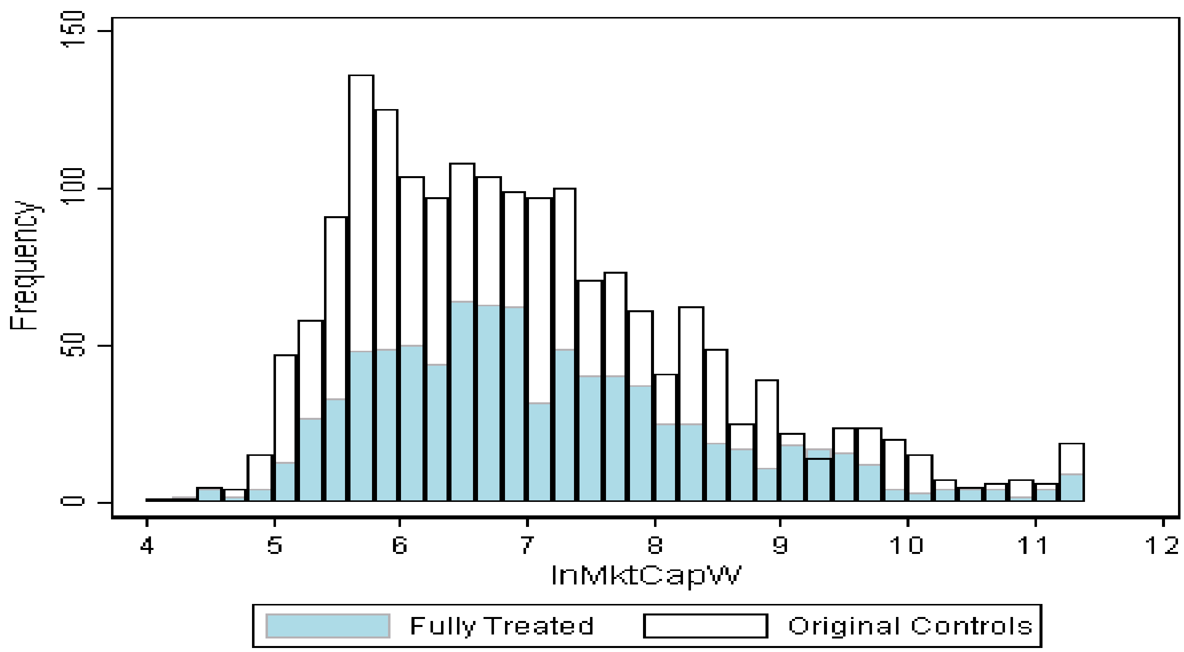

Figure 1, Panel B. Histogram: Partially Treated vs. Full Control Firms

Effect of SEC's division of original control firms into Partly Treated and Full Control firms. Histogram shows number of firms in each group with different levels of *ln*(market capitalization at year-end 2004). Market capitalization is winsorized at 1%/99%, corresponding to *ln*(market cap) of 2.97/11.34 (market cap .of $19.6M/$84.2B). Sample is 564 partly treated and 1,219 full control firms.

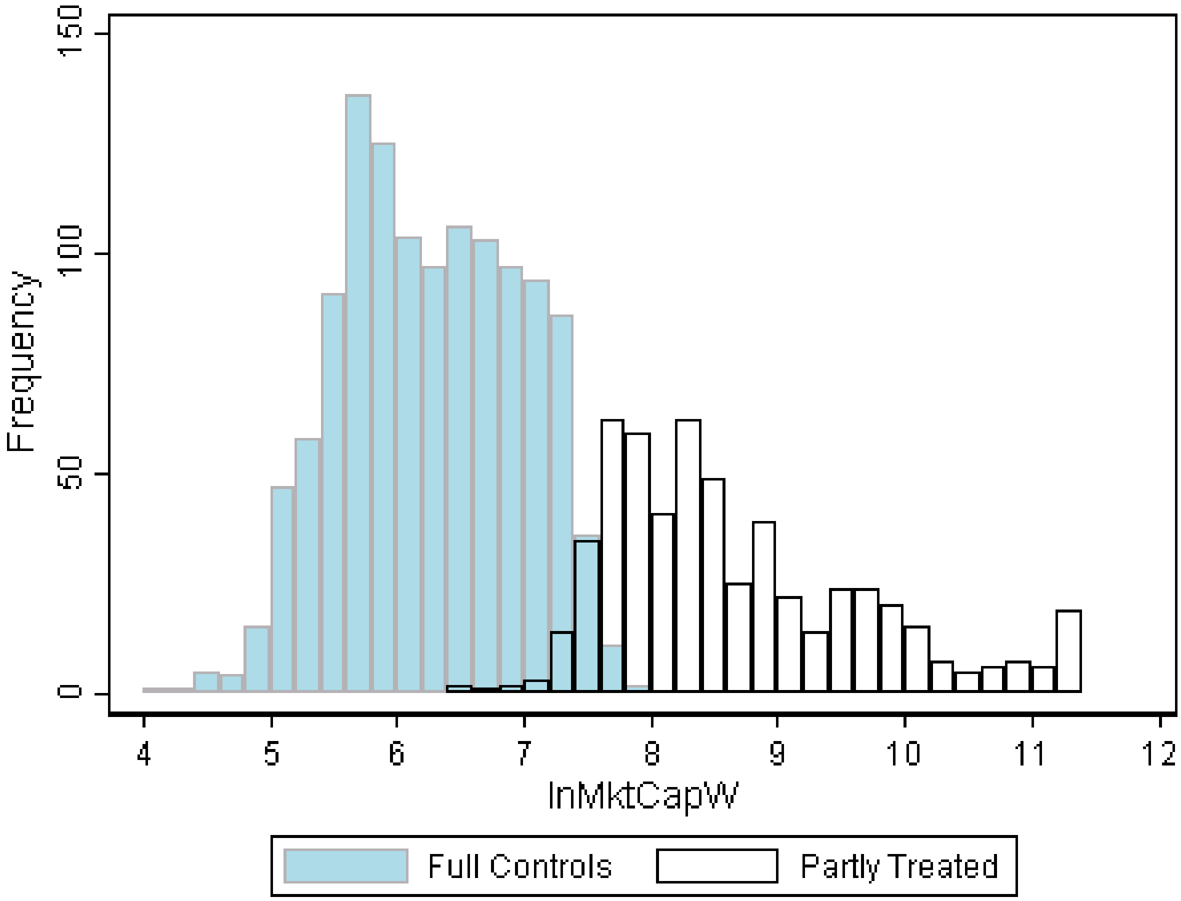

Figure 2. Month-End Open Short Interest

Mean month-end open short interest, as percentage of outstanding shares, over July 2003-December 2007. Short-sale experiment ran from May 2005 through July 2007. Dotted vertical line shows SEC announcement of the experiment; solid vertical lines separate pre-treatment, treatment, and post-treatment periods.

Panel A. Small-firm experiment. Sample is 483 fully treated and 1,041 full control firms with *ln*(market cap) < 7.8 at year-end 2004.

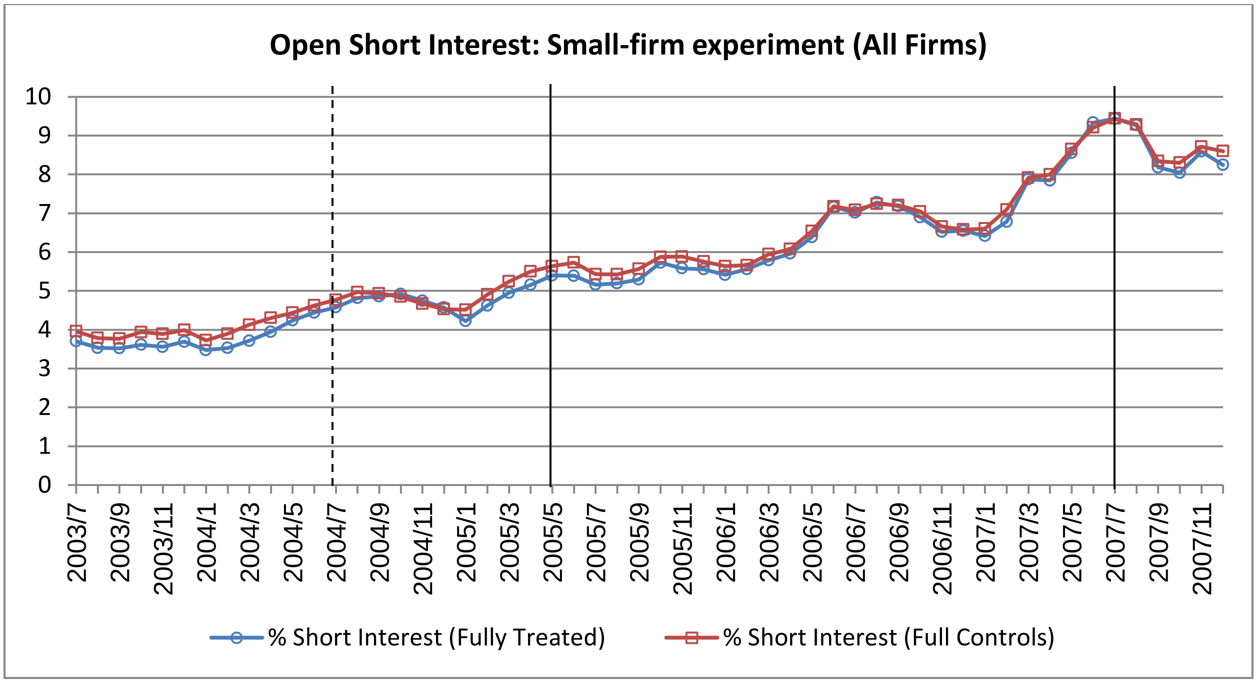

Panel B. Small-firm experiment: NYSE Subsample. Sample is 148 fully treated and 282 full control firms, listed on the NYSE with *ln*(market cap) < 7.8 at year-end 2004.

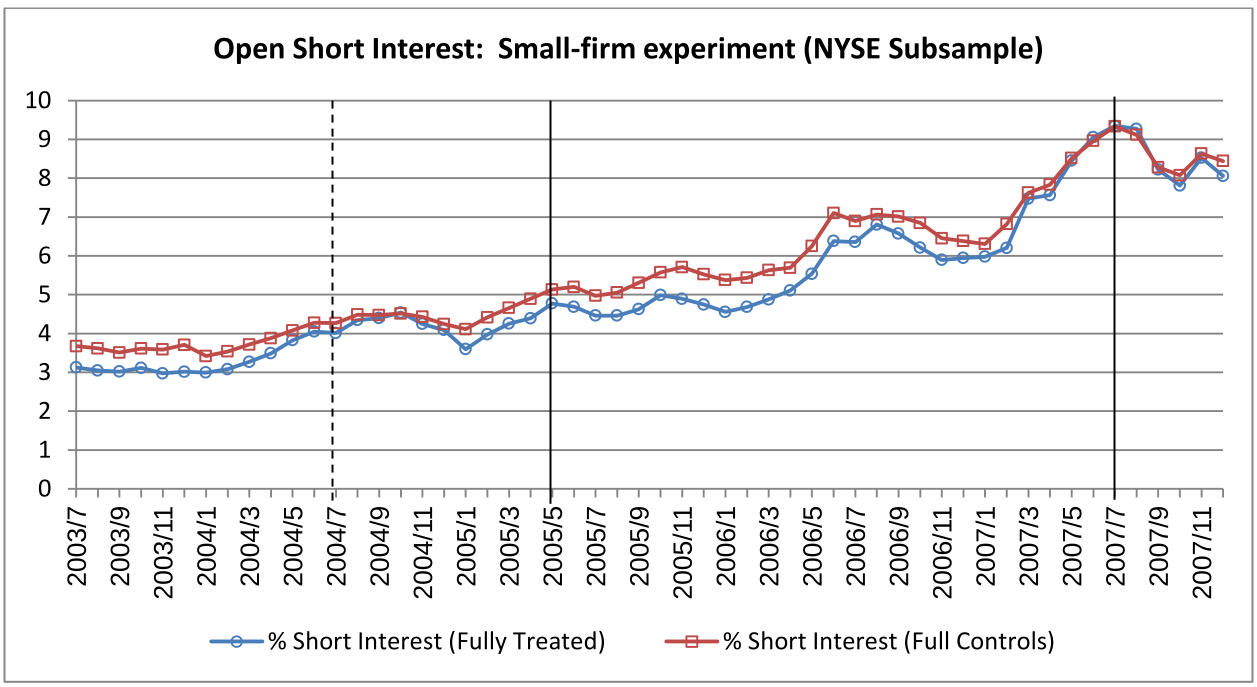

Panel C. Small-firm experiment: NASDAQ Subsample. Sample is 295 fully treated and 639 full control firms, traded on NASDAQ with *ln*(market cap) < 7.8 at year-end 2004

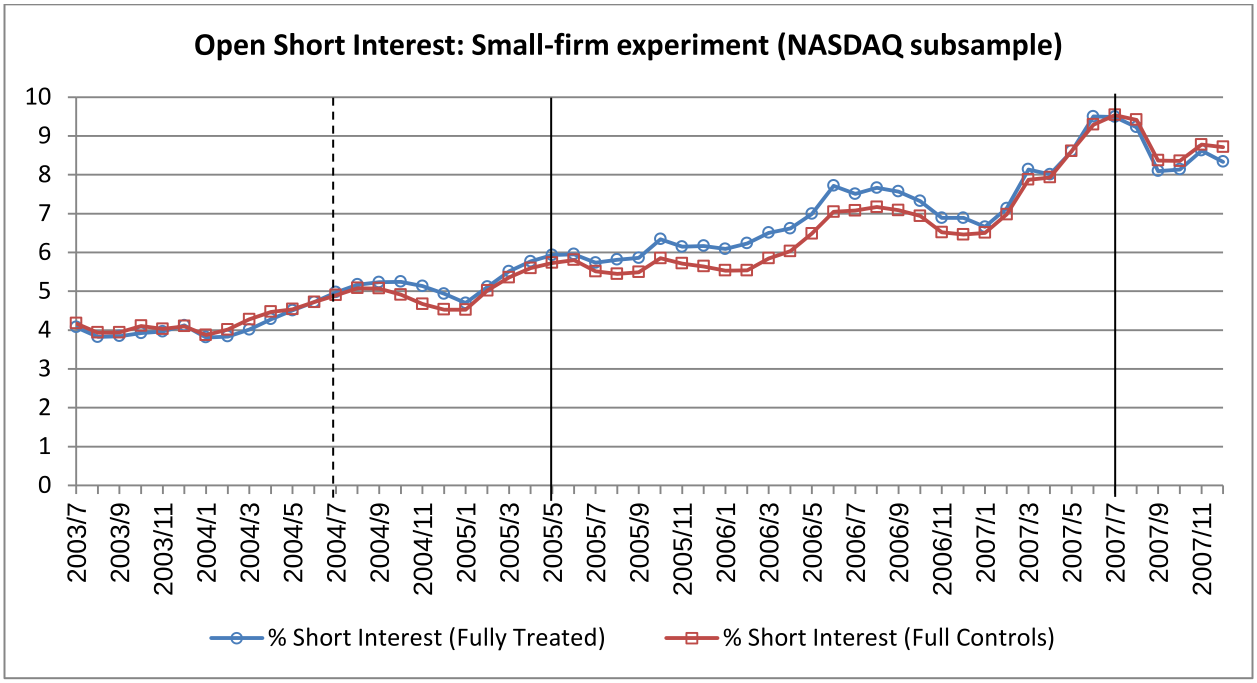

Panel D. Small-Firm Experiment. Plausible Subset, Sample is **xxx** fully treated and yyy full control firms, without publicly traded options as of May 2005, with *ln*(market cap) < 7.8 at year-end 2004.

[*graph to come]

Panel E. Large-firm experiment. Sample is 229 fully treated and 445 partly treated firms with *ln*(market cap) > 7.8 at year-end 2004.

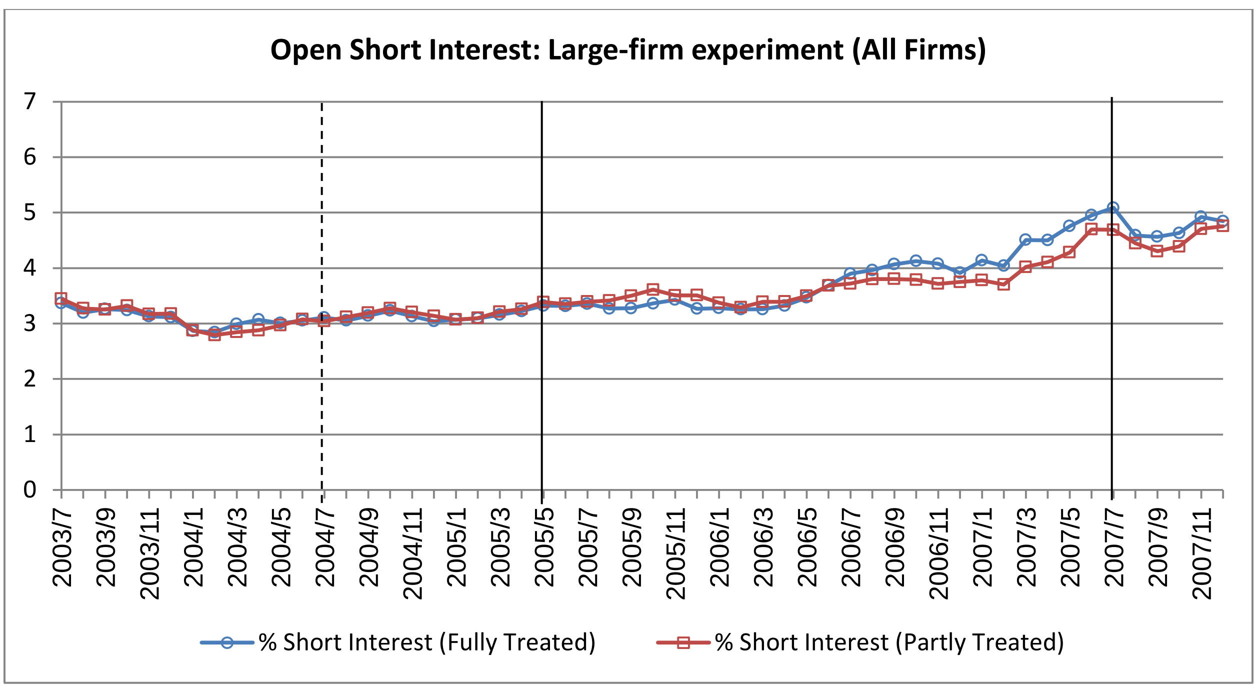

Panel F. Large-firm experiment: NYSE Subsample. Sample is 172 fully treated and 319 partly treated firms, listed on the NYSE with *ln*(market cap) > 7.8 at year-end 2004.

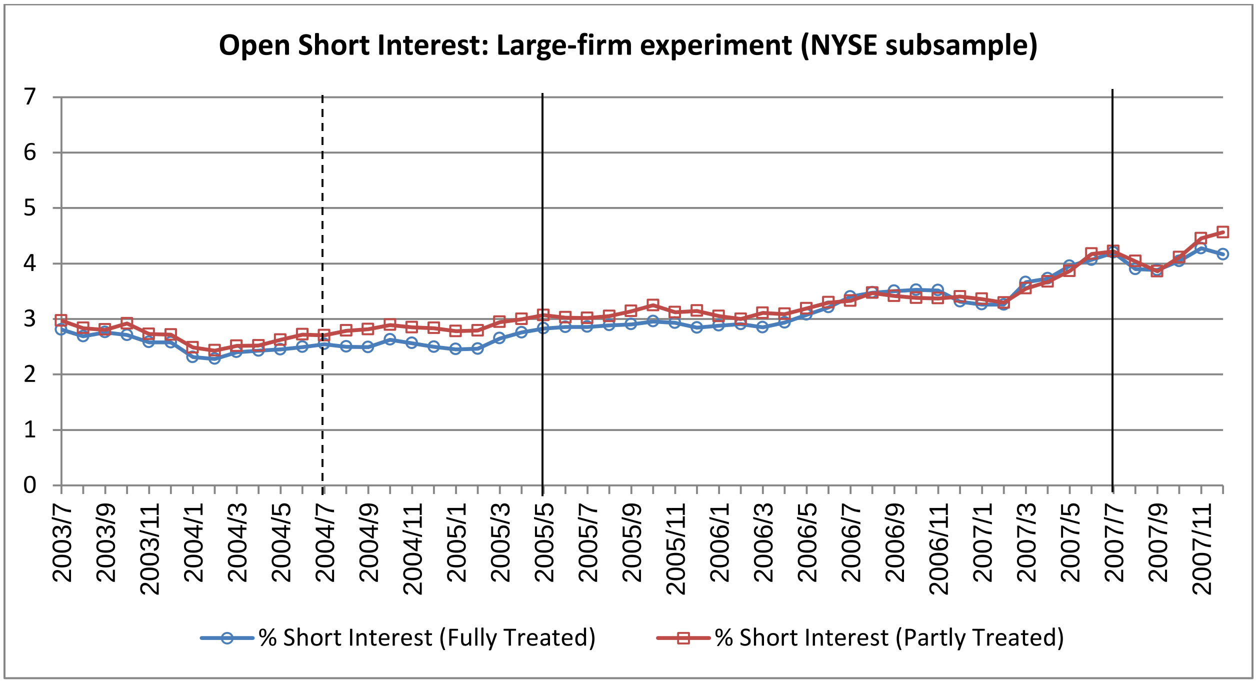

Panel G. Large-firm experiment: NASDAQ Subsample. Sample is 44 fully treated and 110 partly treated firms, traded on NASDAQ with *ln*(market cap) > 7.8 at year-end 2004.

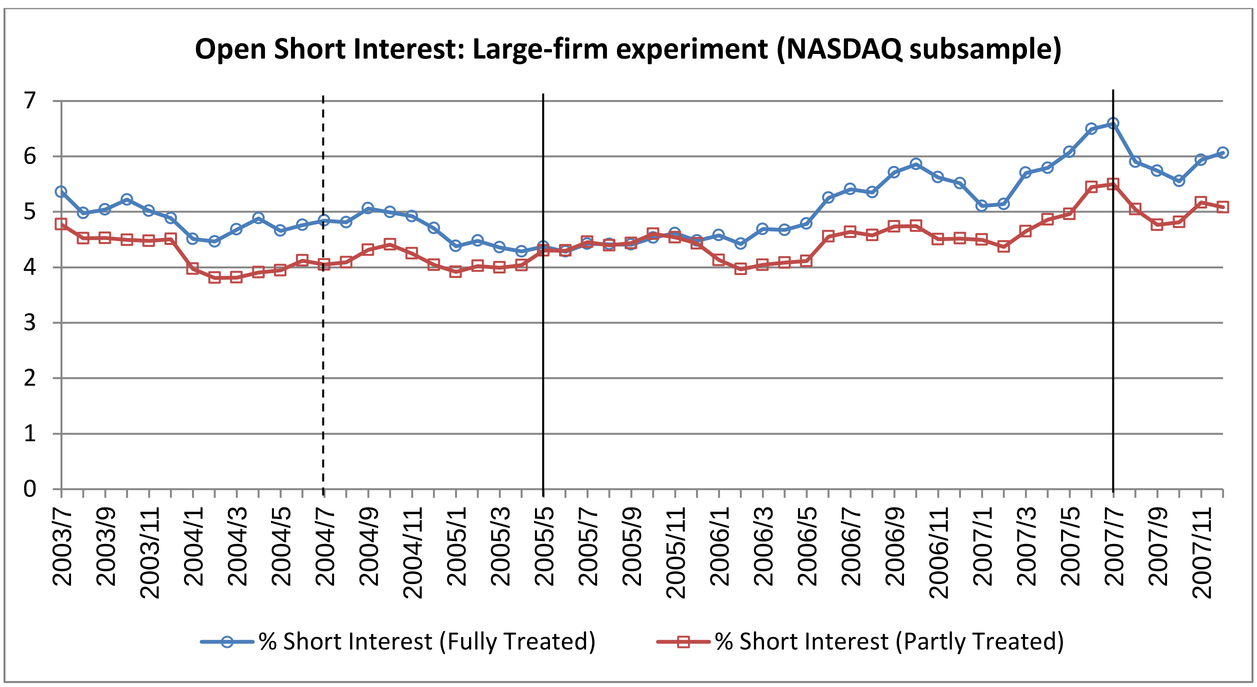

Figure 3. Leads-and-Lags Graphs

Leads-and-lags graphs of month-end open short interest for treated firms relative to control firms, as percentage of outstanding shares, over July 2003-December 2007. Experiment ran from May 2005 through July 2007. Dotted vertical line shows SEC announcement; solid vertical lines separate pre-treatment, treatment, and post-treatment periods.

Panel A. Small-firm experiment. Sample is 483 fully treated and 1,041 full control firms with *ln*(market cap) < 7.8 at year-end 2004.

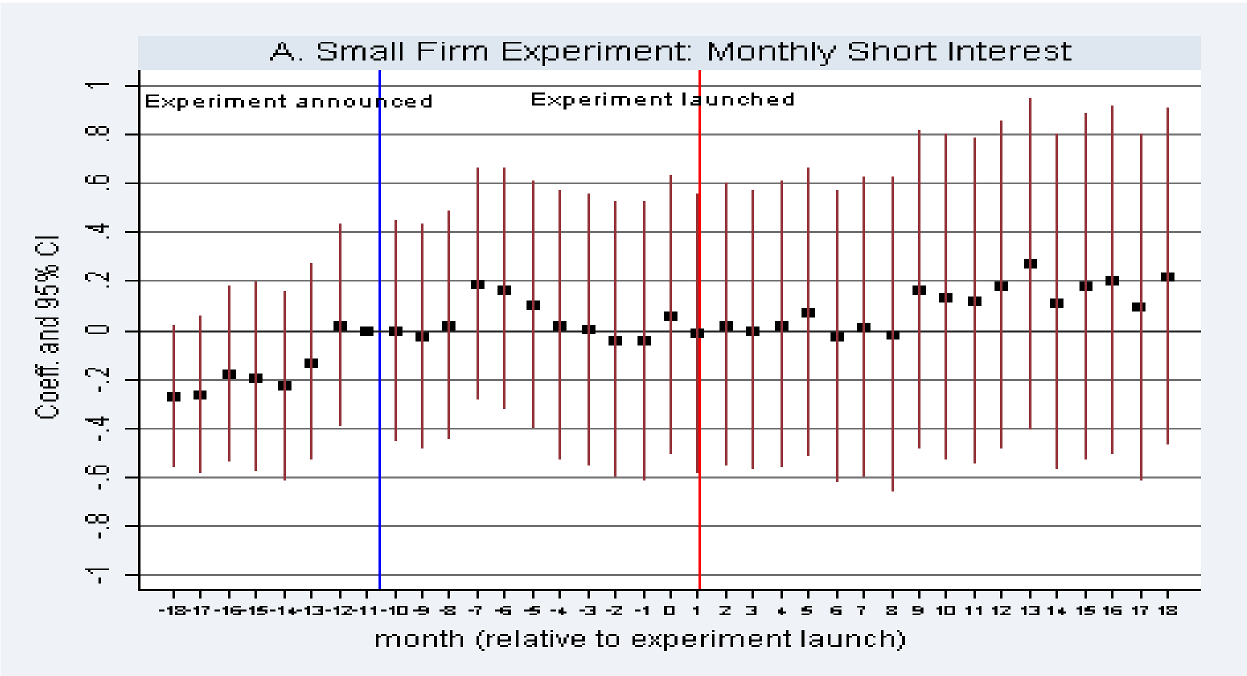

Panel E. Large-firm experiment. Sample is 229 fully treated and 445 partly treated firms with *ln*(market cap) > 7.8 at year-end 2004.

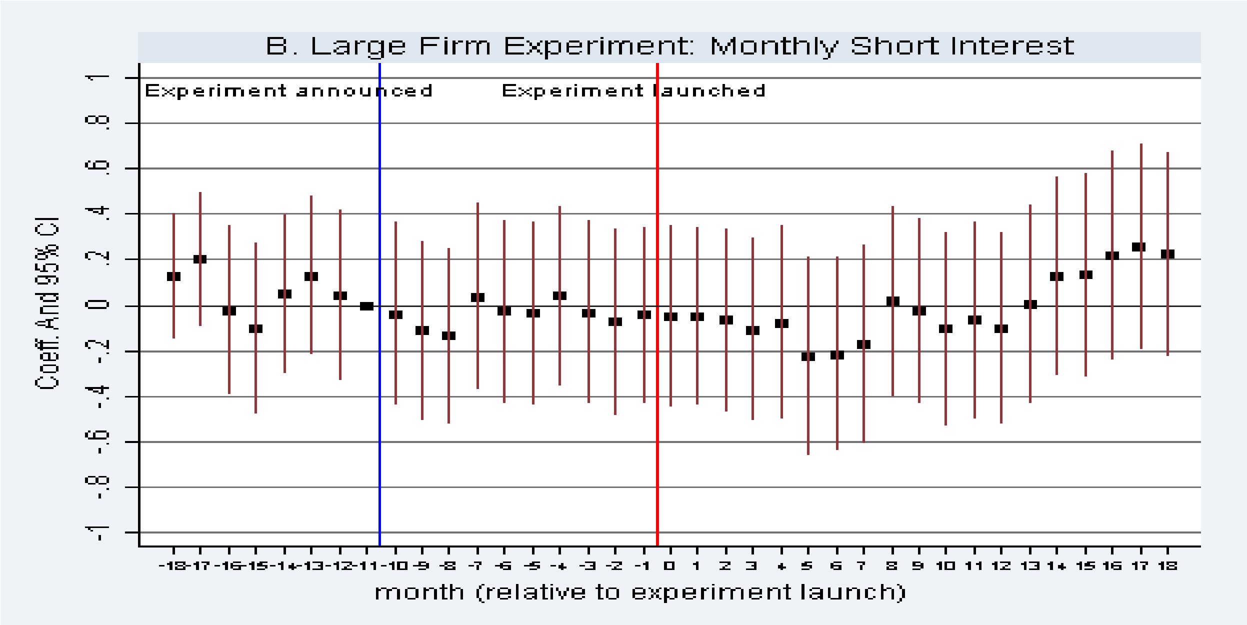

Figure 4. Distribution of Month-End Open Short Interest Across Firms

Panel A. Small Firm Experiment: Kernel Density, June 2004 (before experiment announcement)

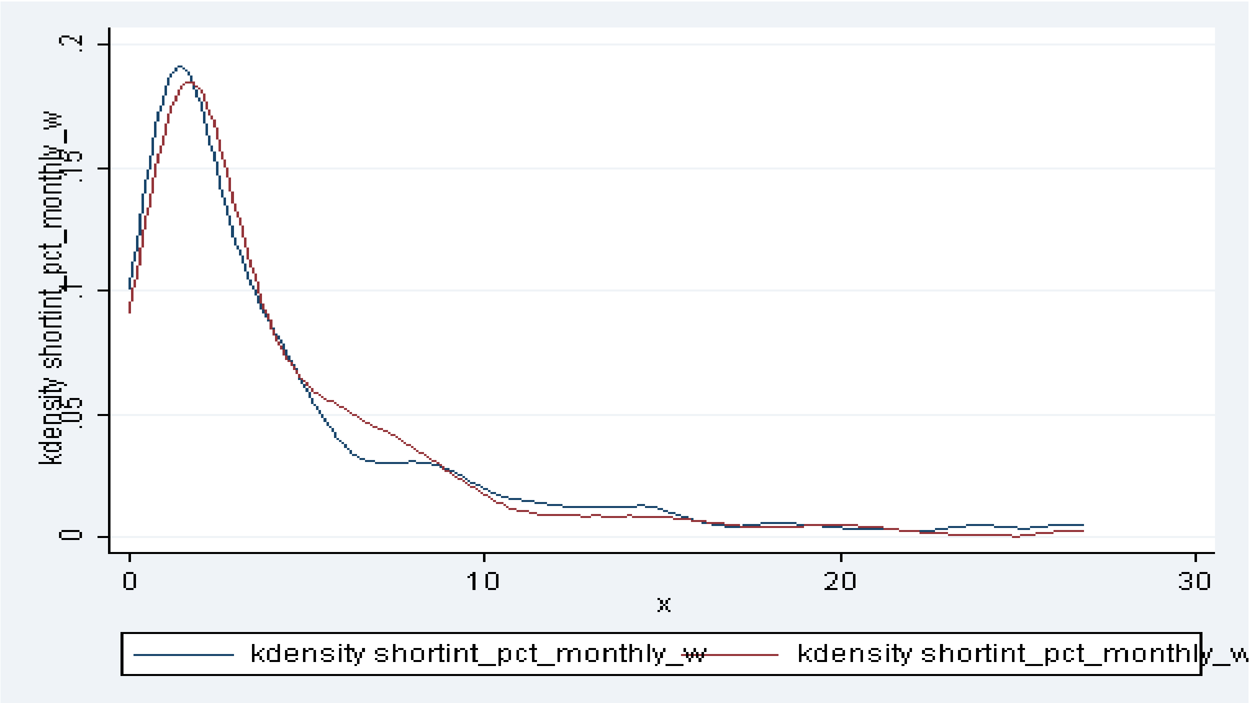

Panel B. Small Firm Experiment: Kernel Density, March 2005 (after announcement, before experiment starts)

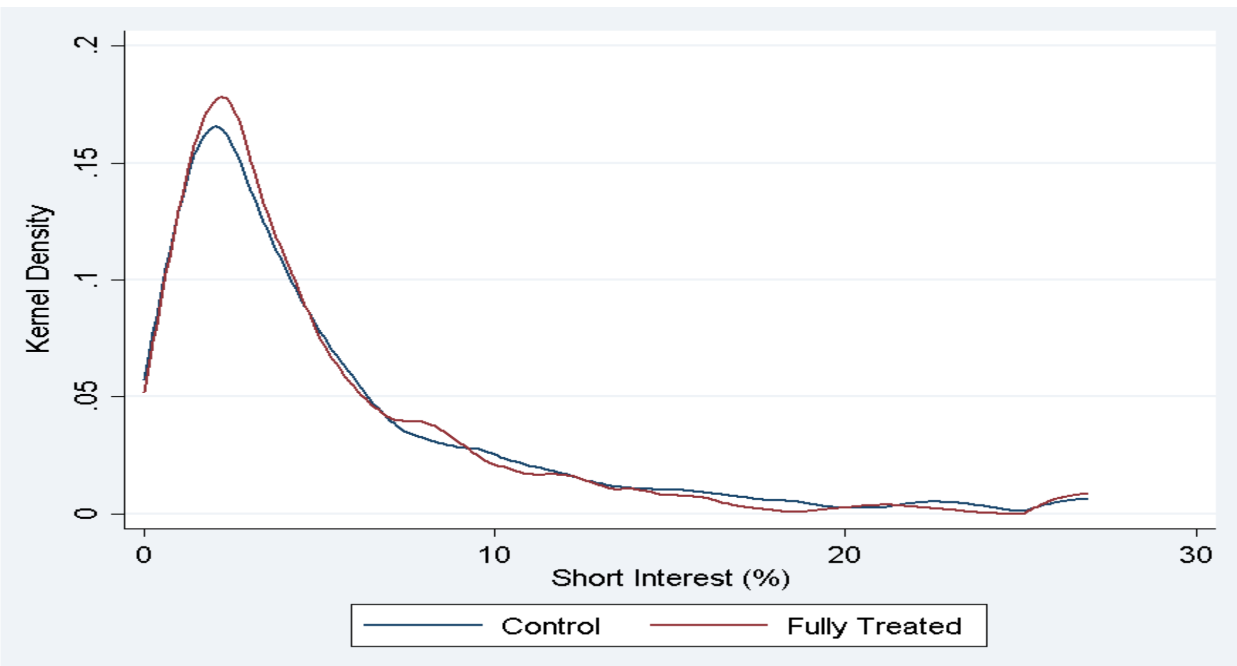

Panel C. Small Firm Experiment: Kernel Density, September 2005 (during experiment)

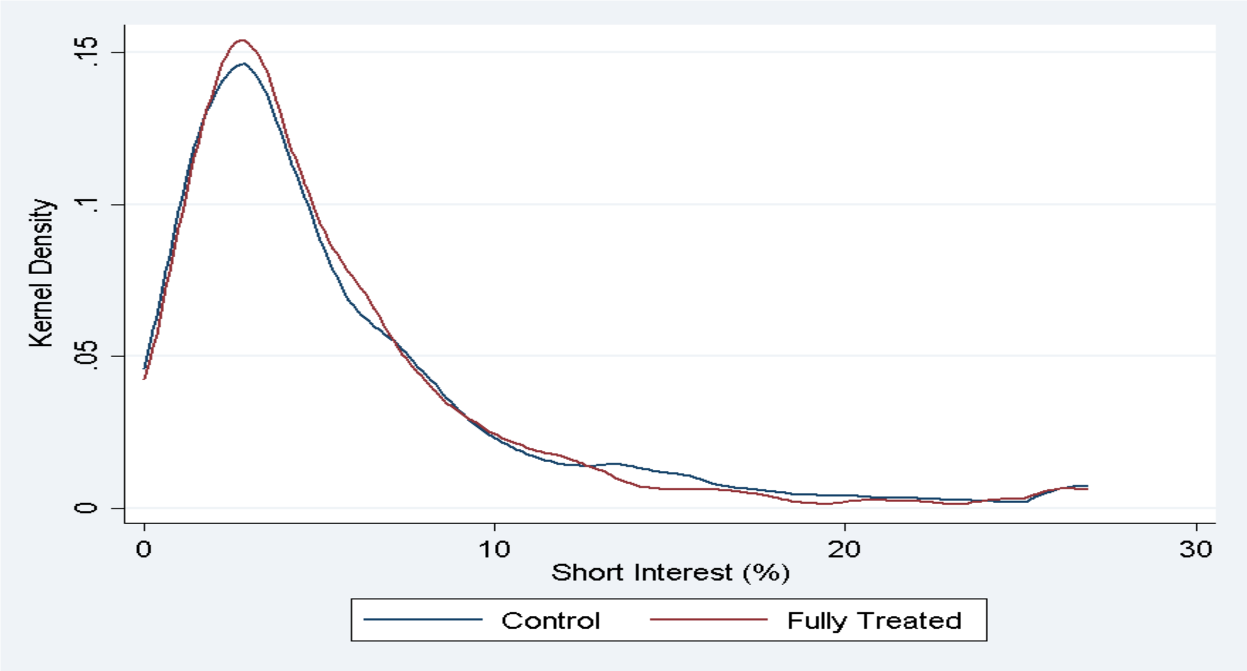

Panel D. Large Firm Experiment: Kernel Density, June 2004 (before experiment announcement)

[*graph to be added]

Panel E. Large Firm Experiment: Kernel Density, March 2005 (after announcement, before experiment starts)

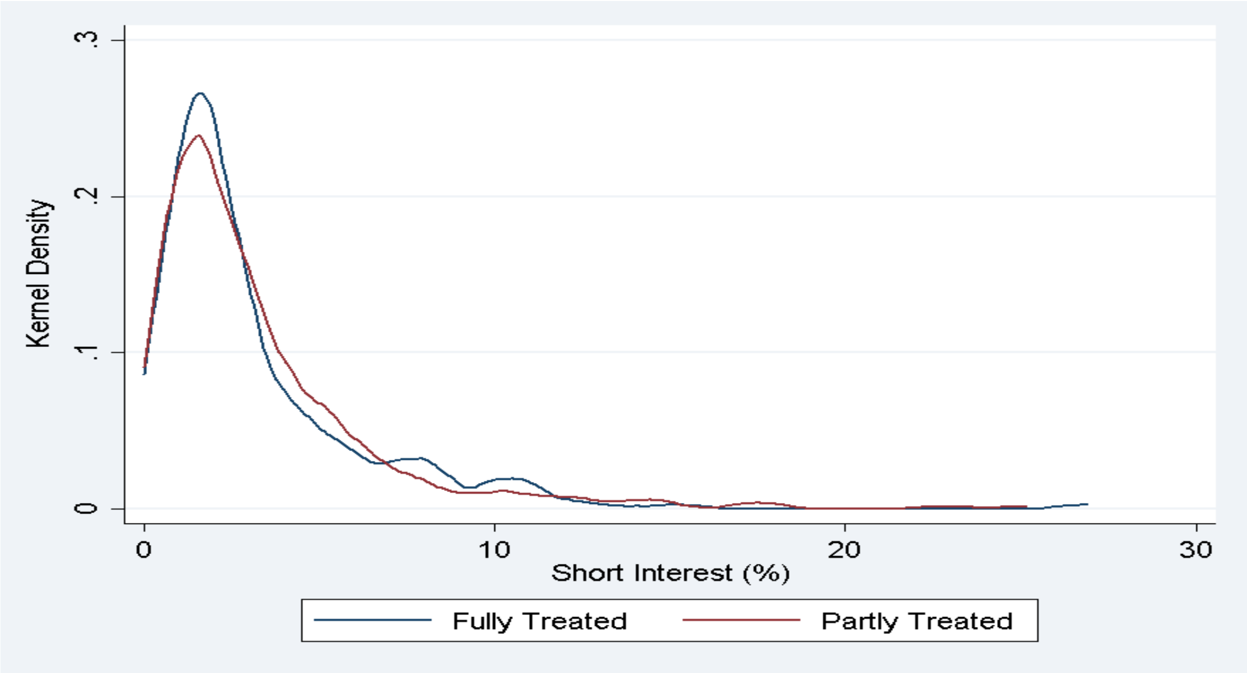

Panel F. Large Firm Experiment: Kernel Density, September 2005 (during experiment)

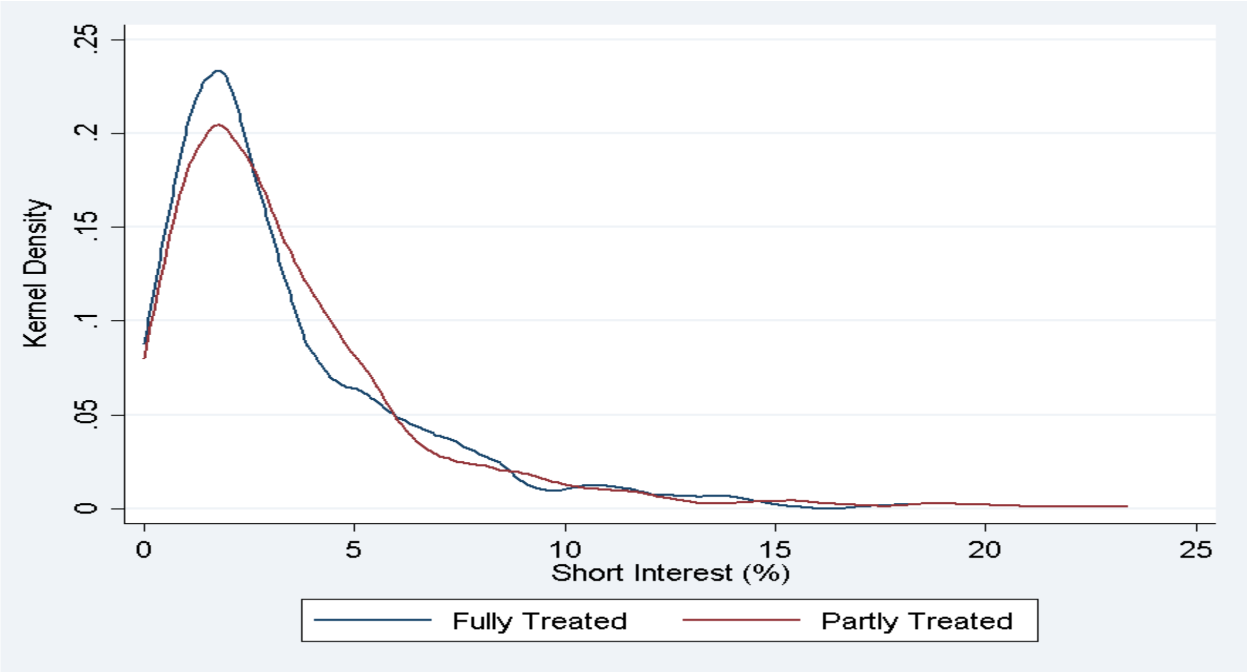

Figure 5. Short-Sale Restrictions and Short-Sale Trading Volume

Figure is based on OEA (2007), Table 14, and shows coefficients on full treatment dummy for size deciles (based on market capitalization) for NYSE firms. Treatment effect is for Small-firm experiment for deciles 1-6; Large-firm experiment for deciles 8-10, and in decile 7 is for fully treated firms vs. mix of full controls and partly treated. Vertical lines divide the treatment regimes. Coefficients are from daily regressions of:

$$\text{Short Selling Volume}_i = \alpha + \sum_{k=1}^{9} \beta_k \text{ Pilot}_i \times \text{Decile}_{ki} + \beta_{10} \text{ Pilot}_i + \beta_{11} \text{ Pre-Expt Short Sale Volume}_i$$

Coefficients for deciles 1-9 = $\beta_{10} + \beta_k$. Short selling volume = number of shares sold short/total trading volume. Regressions are estimated on each day during the Pilot Period (during the experiment).

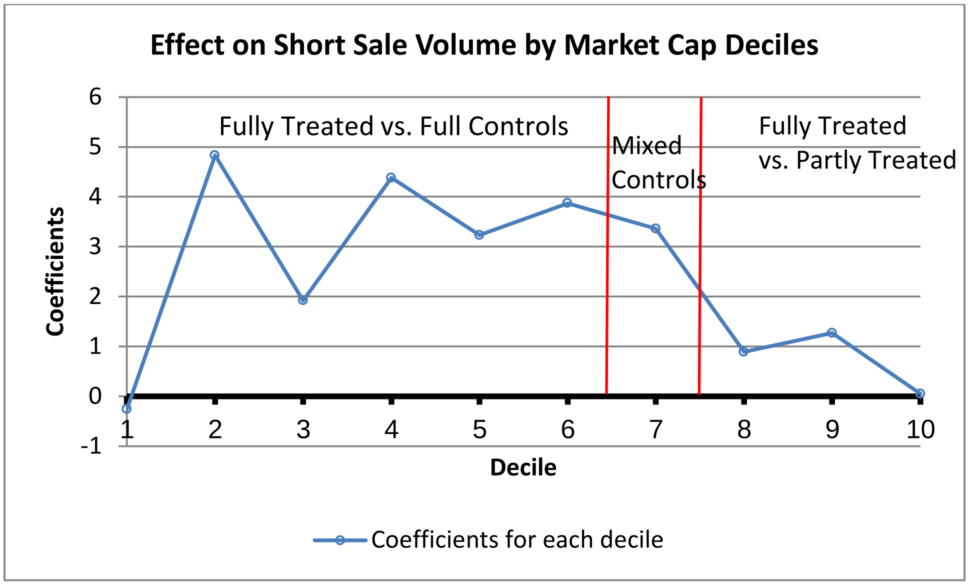

Figure 6. Effect of Short-Sale Experiment on Share Prices

Panel A. Small Firm Experiment: Return Difference

Biweekly average relative share returns for the Small-Firm Experiment, for fully treated versus full control firms, over 2003-2007. Sample size at year-end 2004 is xxx treated firms; yyy control firms. Firms in each group are equally weighted.

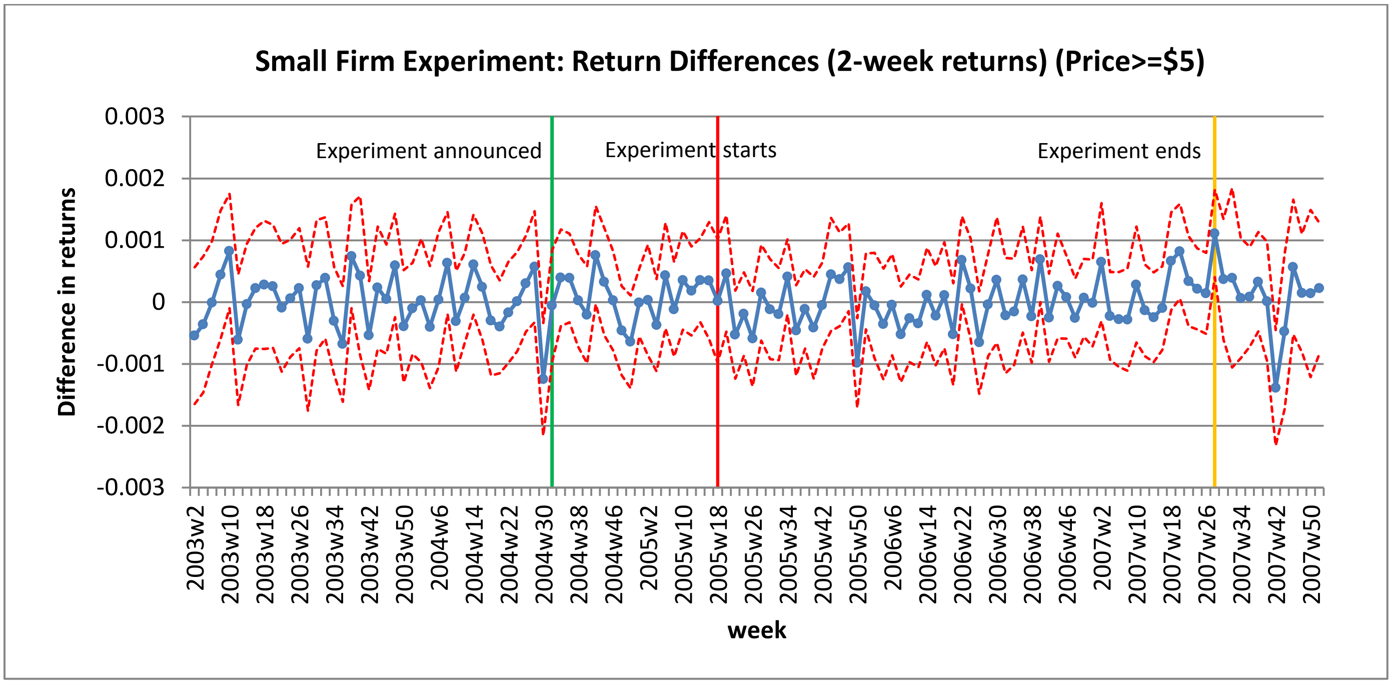

Panel B. Large Firm Experiment: Return Difference

Similar to Panel A, except biweekly returns are for the Large-Firm Experiment, for fully treated firms versus partly treated firms. Sample size at year-end 2004 is xxx treated firms; yyy control firms.

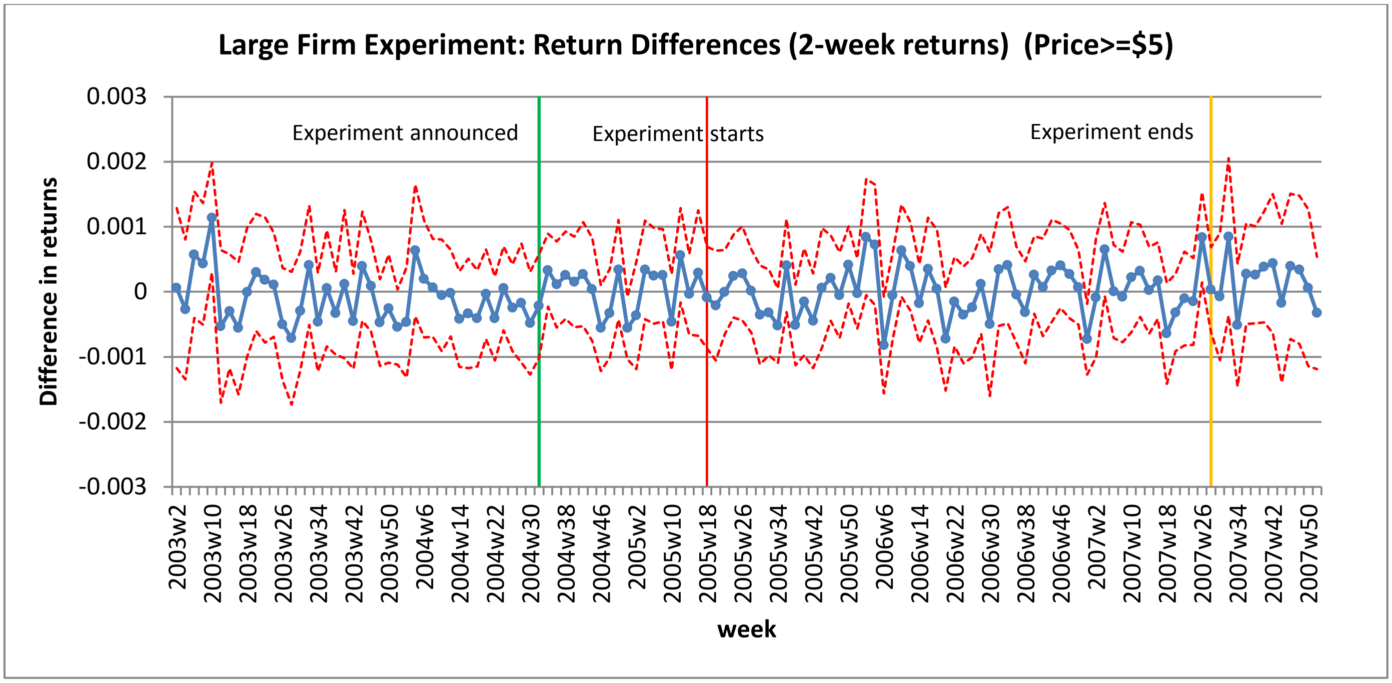
Large Firm Experiment: Return Differences (2-week returns) (Price>=$5)
Experiment announced
Experiment starts
Experiment ends
Difference in returns
0.003
0.002
0.001
0
-0.001
-0.002
-0.003
2003w2
2003w10
2003w18
2003w26
2003w34
2003w42
2003w50
2004w6
2004w14
2004w22
2004w30
2004w38
2004w46
2005w2
2005w10
2005w18
2005w26
2005w34
2005w42
2005w50
2006w6
2006w14
2006w22
2006w30
2006w38
2006w46
2007w2
2007w10
2007w18
2007w26
2007w34
2007w42
2007w50
week

Figure 7. Short-Sale Restrictions and Price Volatility

Figure is based on OEA (2007), Table 20. It shows coefficients on full treatment dummy for size deciles (based on market capitalization) for NYSE firms. Treatment effect is for Small-firm experiment for deciles 2-6; for Large-firm experiment for deciles 8-10, and in decile 7 is for fully treated vs. mix of full controls and partly treated. Vertical lines divide the three treatment regimes. Coefficients are from daily regressions of:

$$\text{Variance (Daily Returns)}_i = \alpha + \sum_{k=1}^{9} \beta_k \,\text{Pilot}_i \times \text{Decile}_{ki} + \beta_{10}\,\text{Pilot}_i + \beta_{11}\,\text{Pre-Expt Variance}_i$$

Coefficients for deciles 2-9 are sum of $\beta_{10} + \beta_k$. Short selling volume = number of shares sold short/total trading volume. Decile 1 is omitted. Regressions are estimated on each day during the Pilot Period (during the experiment).

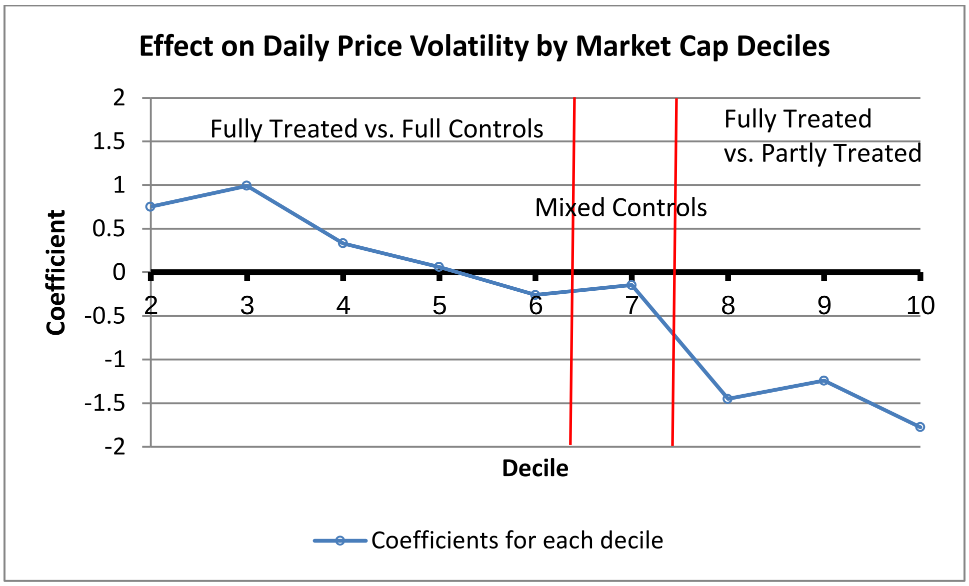

Table 1. Covariate Balance for Intended and Actual Experiments

Table shows sample means for 2004 (or as of end of 2004) for indicated groups. "Treated" firms are Russell-3000 firms designated as "Category A Pilot Securities" by the SEC. Partly treated firms are those designated as "Category B Pilot Securities by SEC. "Full controls" are all other Russell-3000 firms. "Original controls" are "full control" and "partly treated" firms. Sample is 2,640 firms with market capitalization data on Compustat at year-end 2004: 1,219 full controls, 564 partly treated firms (together, 1,783 original controls), and 857 fully treated firms with market capitalization data in Compustat at 31 Dec. 2004. Short interest is average for 2004 of month-end values, as % of outstanding shares. Normalized difference is defined as $ND_j = (\overline{x}_{jt} - \overline{x}_{jc}) / [(s_{jt}^2 + s_{jc}^2)/2]^{1/2}$ (see Imbens and Rubin, 2015). Final column shows two-sample *t*-statistics for difference in means. Dollar amounts in \$ millions; share amounts in millions; employees in thousands. *, **, *** indicates significant differences at 10%, 5%, and 1% levels, respectively. Significant differences, at 5% level or better, in **boldface**.

Panel A. **Covariate balance for original randomization**: Sample means for 857 fully treated to 1,783 original control firms.

Firm Characteristic	Fully Treated	Original Controls	Norm. Diff.	t-stat
ln(Assets)	7.07	6.99		1.07
ln(Sales)	6.46	6.36		1.21
ln(Market Cap)	7.14	7.03		1.82*
ln(Book Value of Debt)	6.27	6.22		0.65
Tobin's *q*	2.27	2.33		0.64
Short Interest (% of Shares)	4.00	4.10		0.52
Sales Growth (%)	1.21	1.22		0.35
PPE/Assets	0.22	0.21		0.94
Capex/Assets	0.04	0.04		0.38
R&D/Sales	1.61	1.63		0.02
EBIT/Assets	0.07	0.05		**2.75*** **
EBIT/Sales	-0.86	-1.57		0.85
Advertising/Sales	0.03	0.03		0.75
Debt/Equity	0.39	0.40		0.20
Debt/Assets	0.22	0.21		1.25
Market/Book Ratio	3.96	3.83		0.15
Market/Sales Ratio	14.96	53.43		0.90
Trading Volume (shares/year-end outstanding)	160,440	163,514		0.27

Panel B. **Covariate balance for fully treated vs. full control firms.** Comparison of 857 fully treated to 1,219 full control firms.

Firm Characteristic	Fully Treated	Full Controls	Norm. Diff.	t-stat
ln(Assets)	7.07	6.23		**13.33*** **
ln(Sales)	6.46	5.58		**11.79*** **
ln(Market Cap)	7.14	6.27		**18.80*** **
ln(Book Value of Debt)	6.27	5.39		**10.70*** **
Tobin's *q*	2.27	2.38		1.13
Short Interest (% of Shares)	4.00	4.55		**2.50** **
Sales Growth (%)	1.21	1.25		0.88
PPE/Assets	0.22	0.20		**2.22** **
Capex/Assets	0.04	0.04		0.24
R&D/Sales	1.61	2.32		0.63

Firm Characteristic	Fully Treated	Full Controls	Norm. Diff.	t-stat
EBIT/Assets	0.07	0.03		**5.18*** **
EBIT/Sales	-0.86	-1.96		1.23
Advertising/Sales	0.03	0.04		0.97
Debt/Equity	0.39	0.41		0.24
Debt/Assets	0.22	0.19		**2.56** **
Market/Book Ratio	3.96	3.07		1.68*
Market/Sales Ratio	14.96	46.46		0.85
Trading Volume (shares)	160,440	78,341		**9.01*** **

Panel C. Covariate balance for fully treated vs. partly treated firms. Comparison of 857 fully treated to 564 partly treated firms.

Firm Characteristic	Fully Treated	Partly treated	Norm. Diff.	t-stat
ln(Assets)	7.07	8.63		**18.97*** **
ln(Sales)	6.46	8.06		**17.77*** **
ln(Market Cap)	7.14	8.68		**22.46*** **
ln(Book Value of Debt)	6.27	8.01		**16.95*** **
Tobin's q	2.27	2.20		0.77
Short Interest (% of Shares Outstanding)	4.00	3.13		**4.01*** **
Sales Growth (%)	1.21	1.16		0.80
PPE/Assets	0.22	0.24		1.68*
Capex/Assets	0.04	0.04		0.62
R&D/Sales	1.61	0.08		1.42
EBIT/Assets	0.07	0.10		**4.66*** **
EBIT/Sales	-0.86	-0.75		0.11
Advertising/Sales	0.03	0.03		0.43
Debt/Equity	0.39	0.40		0.05
Debt/Assets	0.22	0.24		1.72*
Market/Book Ratio	3.96	5.43		1.01
Market/Sales Ratio	14.96	68.08		0.99
Trading Volume (shares)	160,440	350,671		**11.06*** **

Table 2. Covariate Balance for Small Firm and Large-Firm Experiments

Table shows means and medians for 2004 (or at end of 2004) for indicated groups, two-sample *t*-test for difference in means, and p-value for χ^2 test for difference in medians. *, **, *** indicates significance differences at 10%, 5%, and 1% levels, respectively. Significant differences, at 5% level or better, in **boldface.**

Panel A. Small-Firm Experiment. Comparison of 1,217 fully treated versus 569 full control firms with *ln*(market cap) < 7.8 at year-end 2004.

	Full Controls		Fully Treated		t-test for	*p*-value for
	Mean	Median	Mean	Median	means	medians
ln(Assets)	6.23	6.20	6.33	6.31	1.59	**0.048****
ln(Sales)	5.57	5.72	5.69	5.80	1.45	0.378
ln(Market Cap)	6.27	6.25	6.36	6.43	**2.65*****	**0.004*****
ln(Debt)	5.39	5.42	5.47	5.64	0.97	0.060
Tobin's q	2.37	1.74	2.26	1.74	0.97	0.942
Short Interest (% of Shares Outstanding)	4.53	2.73	4.38	2.76	0.55	0.944
Sales Growth (%)	1.25	1.13	1.22	1.14	0.48	0.786
PPE/Assets	0.20	0.11	0.20	0.12	0.12	0.259
Capex/Assets	0.04	0.02	0.04	0.02	0.85	0.827
R&D/Sales	2.31	0.07	2.32	0.05	0.01	0.252
EBIT/Assets	0.03	0.06	0.05	0.07	**2.34****	0.873
EBIT/Sales	-1.95	0.08	-1.39	0.08	0.51	0.873
Advertising/Sales	0.04	0.01	0.02	0.01	0.99	0.486
Debt/Equity	0.40	0.11	0.39	0.13	0.12	0.365
Debt/Assets	0.19	0.14	0.21	0.16	1.55	0.240
Market/Book Ratio	3.08	2.39	3.31	2.23	0.45	0.102
Market/Sales Ratio	46.26	1.76	21.26	1.65	0.55	0.569
Trading Volume (million shares)	78,270	37,590	69,547	39,568	1.31	0.606

Panel B. Large-Firm Experiment. Comparison of 295 fully treated versus 604 partly treated "control" firms with *ln*(market cap) > 7.2 at year-end 2004.

	Partly Treated		Fully Treated		t-test for	*p*-value for
	Mean	Median	Mean	Median	means	medians
ln(Assets)	8.75	8.49	8.65	8.50	0.92	0.925
ln(Sales)	8.09	8.03	8.02	7.93	0.63	0.538
ln(Market Cap)	8.71	8.41	8.72	8.46	0.17	0.460
ln(Debt)	8.15	7.94	8.00	7.95	1.18	0.925
Tobin's q	2.16	1.71	2.31	1.83	1.34	0.112
Short Interest (% of Shares Outstanding)	3.06	2.28	3.05	2.01	0.08	0.112
Sales Growth (%)	1.16	1.12	1.17	1.13	0.82	0.177
PPE/Assets	0.23	0.16	0.26	0.20	**2.05****	0.060
Capex/Assets	0.04	0.03	0.05	0.03	1.73*	0.216
R&D/Sales	0.08	0.03	0.09	0.03	0.20	0.814
EBIT/Assets	0.10	0.09	0.11	0.10	**2.05****	0.114
EBIT/Sales	-0.75	0.13	0.18	0.15	0.75	**0.042****
Advertising/Sales	0.03	0.02	0.03	0.02	1.66*	0.847
Debt/Equity	0.46	0.18	0.42	0.17	0.41	0.705
Debt/Assets	0.24	0.21	0.24	0.22	0.38	0.282
Market/Book Ratio	5.37	2.86	5.14	2.92	0.10	0.598
Market/Sales Ratio	67.91	1.68	2.93	1.93	0.71	0.200
Trading Volume (million shares)	337,445	198,914	344,125	191,206	0.25	0.770

Table 3: DiD Regressions: Month-end Open Short Interest

Firm and month fixed effects regressions, over July 2003-July 2007 of month-end short interest (as % of outstanding shares) on treatment period dummy (=1 for May 2005 on, fully treated dummy (drops out with firm FE), interaction between treatment period and fully treated dummies, and indicated covariates, with firm and month fixed effects. *t*-statistics, with standard errors clustered on firm, in parentheses. *, **, *** indicates statistical significance at the 10%, 5%, and 1% levels, respectively. Significant results, at 5% or better, in **boldface** (suppressed for constant term).

Panel A. Small-Firm Experiment. Sample is 1,217 fully treated firms and 569 full control firms, with *ln*(market cap) < 7.8 at year-end 2004.

	Month-end open short interest (% of outstanding shares)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All Small Firms		**NYSE Subsample**		**NASDAQ Subsample**		**Plausible Subset**	
Treatment period dummy	**5.430***	**4.761***	**5.557***	**6.003***	**5.112***	**3.946***		
	(27.53)	**(19.56)**	**(17.29)**	**(15.85)**	**(19.87)**	**(12.20)**		
Treatment period dummy * Fully treated dummy	0.144	0.123	0.006	0.005	0.378	0.326		
	(0.72)	(0.59)	(0.02)	(0.02)	(1.30)	(1.12)		
ln(Market Cap)		-0.348*		**-0.994**		0.179		
		(-1.73)		**(-2.35)**		(0.76)		
ln(Assets)		**1.444***		-0.230		**1.510***		
		(3.26)		(-0.22)		**(2.95)**		
ln(Sales)		0.024		0.534		0.395		
		(0.09)		(1.26)		(1.17)		
ln(Debt)		0.299		0.148		0.340		
		(1.21)		(0.23)		(1.19)		
Tobin's *q*		-0.020		0.062		-0.128*		
		(-0.29)		(0.26)		(-1.86)		
Sales Growth (%)		0.007		-0.007		-0.021		
		(0.20)		(-0.35)		(-0.38)		
PPE/Assets		**4.050***		-0.437		**7.515***		
		(3.31)		(-0.27)		**(3.82)**		
Constant	3.753	-5.134	3.390***	6.992**		-9.778		
	(37.54)	(-2.94)	(23.18)	(2.30)		(-4.60)		
Firm and month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No of Firms	1,776	1,615	581	524	1,018	945		
Observations	72,098	65,373	24,378	22,095	40,742	37,447		
R^2	0.715	0.721	0.712	0722	0.708	0.715		

Panel B. Large-firm experiment. Sample is 295 fully treated firms and 604 partly treated "control" firms, with ln(market cap) > 7.2 at year-end 2004.

	Month-end open short interest (% of outstanding shares)					
	(1)	(2)	(3)	(4)	(5)	(6)
	All Large Firms		**NYSE Subsample**		**NASDAQ Subsample**	
Treatment period dummy	**1.375***** **(7.54)**	**1.780***** **(8.09)**	**1.232***** **(6.49)**	**1.914***** **(6.61)**	**0.894**** **(2.55)**	**0.984**** **(2.00)**
Treatment period dummy * Fully treated dummy	0.127 (0.73)	0.071 (0.40)	0.117 (0.63)	0.095 (0.49)	-0.170 (-0.40)	-0.216 (-0.50)
ln(Market Cap)		**-1.912***** **(-5.61)**		**-1.757***** **(-4.01)**		**-1.691***** **(-2.75)**
ln(Assets)		-0.295 (-0.49)		-1.135 (-1.23)		0.512 (0.55)
ln(Sales)		0.230 (1.28)		-0.339 (-0.48)		0.432 (1.52)
ln(Debt)		0.362 (0.92)		**1.431***** **(2.82)**		-0.199 (-0.39)
Tobin's q		0.206 (1.31)		0.173 (0.63)		0.165 (0.94)
Sales Growth (%)		0.000 (1.16)		-0.203 (-0.70)		-0.000 (-0.47)
PPE/Assets		1.454 (0.86)		-1.913 (-1.27)		**10.278***** **(2.64)**
Constant	3.418 (42.96)	16.736 (6.07)	2.917 (33.77)	19.595 (4.65)	4.902*** (27.06)	10.819** (2.17)
Firm and month FE	Y	Y	Y	Y	Y	Y
No. of Firms	889	749	634	522	213	194
Observations	36,388	30,157	25,922	21,376	8,965	7,588
R^2	0.683	0.720	0.671	0.697	0.748	0.758