

MPRA

Munich Personal RePEc Archive

Anticipated Monitoring, Inhibited Detection, and Diminished Deterrence

Makofske, Matthew

5 February 2024

Online at <https://mpra.ub.uni-muenchen.de/121173/>
MPRA Paper No. 121173, posted 12 Jun 2024 10:43 UTC

ANTICIPATED MONITORING, INHIBITED DETECTION, AND DIMINISHED DETERRENCE

Matthew Philip Makofske*

February 4, 2024

Abstract

Monitoring programs—by creating expected costs to regulatory violations—promote compliance through general deterrence, and are essential for regulating firms with potentially hazardous products and imperfectly observable compliance. Yet, evidence on how monitoring deployment affects perceived detection probabilities and—by extension—compliance, is sparse. Beginning in May 2020, pandemic-related protocols in Maricopa County, Arizona, required routine health inspections to occur by video-conference at food establishments with vulnerable populations (e.g., hospitals and nursing homes). Unlike conventional on-site inspections—which continued at most food establishments—these “virtual” inspections were scheduled in advance, and thus, easily anticipated. The virtual format also likely inhibits observation of some violations, further reducing detection probability. Tracking five violations that are detected by tests in both inspection formats, I find evidence of substantial anticipation-enabled detection avoidance. Comparing against contemporaneous on-site inspections, virtual inspections detect 53% fewer of these specific violations relative to pre-treatment levels, and that decrease reverses entirely when treated establishments are subsequently inspected on-site. Detected counts of all violations decrease 41% in virtual inspections. Consistent with general deterrence, this decrease is *more* than offset in establishments’ first post-treatment on-site inspections, where detected counts exceed the pre-treatment average by 28%. Across establishment types and compliance histories, deterrence-effect heterogeneity suggests a simple dynamic enforcement rule would better allocate existing inspection resources, and might meaningfully reduce social noncompliance costs.

*Department of Economics, Colgate University. Email: mmakofske@colgate.edu. I thank John Bowblis, Carolina Castilla, Rishi Sharma, and seminar participants at Colgate University and Washington State University for many helpful comments. Any remaining errors are mine.

1 Introduction

Programs of routine unannounced inspections are nearly universal in enforcing food-service hygiene and safety regulation. Yet, while entirely preventable, the Centers for Disease Control and Prevention (CDC) estimates that 48 million Americans contract a foodborne illness each year, with an annual economic burden estimated at 15.5 billion dollars (Hoffmann et al., 2015).¹ And from 2017 through 2019, Moritz et al. (2023) report that the CDC was *voluntarily* alerted to 800 foodborne-illness outbreaks involving retail food establishments, by 25 state and local health departments.

Periodic compliance monitoring creates expected costs for regulatory violations—the penalty if detected multiplied by the perceived detection probability—and promotes compliance through general deterrence (Becker, 1968). This enforcement approach has profound reach. Beyond food safety, it is also central to regulating—among other things—environmental quality, workplace hazards, international maritime practices, nursing-home standards, and licensed firearm dealers.

With monitoring resources efficiently deployed, a tradeoff exists between enforcement- and noncompliance costs—the sum of which is minimized at the social optimum. Yet, efficient (noncompliance-cost minimizing) deployment of monitoring resources is practically complex, and requires knowledge of: (i) how deployment affects actual, and perceived, detection probabilities; (ii) how perceived detection probabilities affect compliance; and (iii) potential heterogeneity in these effects across regulated entities.

Empirical evidence regarding these relationships is sparse and challenging to attain. Variation in monitoring frequency is potentially endogenous to compliance, and even if not, accounting for firms' perceptions is difficult.² Finally, even with an exogenous and perceived detection-probability shock, cleanly separating that shock's *deterrence effect*

¹CDC estimate [here](#); economic burden is estimated in 2013 USD.

²Across several industries, an initial literature (Gray and Deily, 1996; Laplante and Rilstone, 1996; Eckert, 2004; Telle, 2009) estimates inspection propensity as a function of firm observables, and generally finds positive relationships between predicted probabilities (which proxy for firm perceptions) and compliance. Gray and Shimshack (2011) review the challenges of accounting for perceptions of regulatory stringency and monitoring intensity.

from its opposing—and often simultaneous—*detection effect*, is seldom feasible.³ Exploiting a regulator’s pandemic-induced shift to remote inspections for *some* entities under their jurisdiction, I largely overcome these issues.

From the COVID-19 pandemic’s onset, the Maricopa County Environmental Services Department (MCESD) continued conducting routine health inspections on-site at most permitted food establishments. However, in late May 2020 they began conducting these inspections by video-conference for establishments serving especially vulnerable populations, such as hospitals, nursing homes, and assisted living facilities. These “virtual” inspections required advance scheduling with an establishment’s person-in-charge, and were thus, easily anticipated. Advance notice of inspections undermines a fundamental aspect of enforcement *via* deterrence—the continual threat of detection and punishment. By knowing in advance when detection will occur, establishments treated with virtual inspections can avoid punishment by correcting violations just prior, and upon recognizing this, will likely relax compliance effort. Moreover, the remote format likely inhibits inspector ability to observe some violations, *further* reducing their detection probability.

Using MCESD inspections spanning January 2018 to early August 2022, I leverage this sudden format adjustment as a policy experiment, and test multiple facets of the imperfect-monitoring paradigm. Concurrent on-site inspections at untreated establishments provide control for contemporaneous factors that may have affected compliance generally, and the sudden return of unannounced on-site inspections at treated establishments enables identification of a deterrence effect. In initial post-treatment on-site inspections, actual detection probabilities return to pre-treatment levels (removing any detection effect), but compliance efforts are still based on virtual-regime perceptions.

Initially, I track a subset of five MCESD codes where—regardless of inspection mode—compliance is checked through tests.⁴ Violations of these particular codes will isolate potential anticipation-enabled avoidance, because the virtual format doesn’t inhibit their

³E.g., following an exogenous and perceived detection-probability increase, fewer violations will be committed (the *deterrence effect*), but a greater share of committed violations will be detected (the *detection effect*).

⁴These tests involve demonstration of an appropriate holding temperature with a thermometer, or sufficient sanitizer concentration in cleaning solutions with pH test strips.

detection. Comparing against contemporaneous (same 14-day period) on-site inspections, and controlling for time-invariant establishment-specific differences, virtual inspections detect about 53% fewer of these “virtually demonstrable” violations. Consistent with last-minute and short-lived corrections, this decrease reverses entirely in subsequent on-site inspections. Notably, the decrease is largely evident in treated establishments’ first virtual inspections, suggesting fairly immediate detection avoidance.

While advance notice reduces detection probability on any violation capable of quick remedy, those five violations isolate anticipation’s effect because, even in virtual inspections, they *will* be detected if not corrected prior. Conversely, violations detected by visually observing premises are presumably less likely to be caught by virtual inspections, even when left uncorrected. Thus, I then expand focus to violations of any MCESD code, and use the return of unannounced on-site inspections at treated establishments to assess how overall compliance responds to the detection-probability shock.

Detected counts of all violations are 41% lower in virtual inspections, relative to the pre-treatment average. Notably, in establishments’ initial post-treatment on-site inspections—when their perceptions of detection probability are likely adapted to the virtual regime—that decrease is *more* than offset, yielding an estimated net increase that exceeds the pre-treatment average by 28%. Consistent with general deterrence, this suggests the detection-probability decrease caused a substantial decline in compliance effort.

Individual-level responses to this shock provide insight on a fundamental dilemma: should firms with strong compliance records receive fewer inspections, so that severe violators can receive more? Deterrence-effect heterogeneity supports redirecting some routine inspections away from highly compliant establishments in lower risk classes, and toward establishments in the highest risk class where significant violations have been found. I show that a simple dynamic enforcement rule would achieve this targeted redirection with existing inspection resources, enhance general deterrence in the highest risk classification, and perhaps meaningfully reduce social noncompliance costs on net.

My findings build on a nascent literature utilizing field and natural experiments to empirically test enforcement *via* imperfect monitoring. In Florida food-service health in-

spections, following adoption of handheld devices which reminded inspectors of potential violations, [Jin and Lee \(2014\)](#) find an immediate 11% increase in detected violations; subsequent inspections suggest modest compliance-effort improvements in response. [Duflo et al. \(2018\)](#) study an experimental doubling of environmental-inspection frequency at Indian factories. Treated plants perceive elevated scrutiny, and are more frequently cited for violations, but no effect on average emissions is found. Most closely related to this work, two recent studies draw identifying variation in detection probability from the ability of some entities to anticipate monitoring in advance.

[Makofske \(2021\)](#) examines Las Vegas facilities housing multiple food-service establishments. At such facilities, inspectors often conduct many inspections during one visit, and establishments inspected second or later likely anticipate those inspections in advance. The study finds that, within establishment, detected noncompliance is significantly higher when inspected first—an effect driven by violations capable of quick remedy, suggesting anticipation-enabled avoidance—but is unable to test deterrence.⁵ [Zou \(2021\)](#) exploits every-sixth-day pollution monitoring under the Clean Air Act, which the US Environmental Protection Agency allows at some monitor sites. Near intermittent sites, [Zou \(2021\)](#) finds satellite pollution measures are 1.6% lower during monitor on-days than off-days, and that air-quality advisories are more likely during on-days, suggesting strategic responses by local governments. Following the retirement of some intermittent monitors, Zou finds that pollution levels significantly increase on what would have been on-days, and change little otherwise, consistent with deterrence.

[Makofske \(2021\)](#) and [Zou \(2021\)](#) use variation in anticipation ability—within-entity and across-entity, respectively—that is due to established institutional features and present throughout their samples. Here, firms with no prior anticipation ability acquire it from an abrupt and unforeseeable inspection-format change. The immediate response that ensues, shows how practices which inadvertently enable anticipation, even if short-lived, can meaningfully undermine enforcement. Further, observations before and after the

⁵Using Los Angeles County health inspections, [Makofske \(2019\)](#) compares detected noncompliance within establishment, across days when receiving the sole inspection, or one of many inspections, at a facility. Significantly more violations are detected on sole-inspection days, when anticipation is less likely.

virtual regime, enable comparisons across inspections where detection probabilities are similar, but perceived to be quite different. This yields an exceptionally clean test of deterrence, and—through firm-specific responses—an ability to assess potential policy improvements not possible in prior work. In evaluating response heterogeneity, I also contribute to recent work (Blundell, 2020; Blundell et al., 2020) on dynamic-enforcement mechanisms—individual-level expected-cost adjustments triggered by compliance history.

In the space remaining, I discuss details of the MCESD inspection program and the virtual regime begun in 2020. Next, I review the data and estimating sample, explain the methodology employed and test underlying assumptions. I then present estimates of anticipation-enabled avoidance and general deterrence. Finally, I examine deterrence-effect heterogeneity, discuss policy implications, and conclude.

2 Background

2.1 Maricopa County Inspection Program

The Maricopa County Environmental Services Department (MCESD) regulates and inspects food service and retail food establishments whom—per the MCESD—receive “required unscheduled food safety inspections.” MCESD issues 26 different food establishment permit types which, based on the nature of food and population typically served, are assigned risk classifications (from lowest risk to highest): *class 2*, *class 3*, *class 4*, and *class 5*.⁶ Establishments in these classes are prescribed 2, 2, 3, and 4 annual routine inspections, respectively.

Inspections check health code compliance and violations are specified—from most to least severe—as *priority*, *priority foundation*, and *core*.⁷ MCESD supplements inspections with ratings and disclosure. Inspection performances are graded: *A*, *B*, *C*, and *D*, according to the schedule in Appendix Figure A1. A peculiarity of this grading policy is that participation is voluntary. Prior to every inspection, the establishment’s person-in-

⁶*Class 1* applies only to Micromarket permits, none of which are in the primary estimating samples (see Section 3).

⁷Severity levels are not specific to the health code violated; i.e., a particular health code can be violated to each severity level.

charge chooses whether they will participate in the grading program for that inspection. If participation is elected, the grade—along with any cited violations—is shared on the county’s restaurant ratings [page](#); a grade card is also issued but display of the card is optional. If participation is declined, the inspection report with violations are posted online with “Not Participating” in place of a letter grade. The election is made before the inspection starts, and irreversible.

Despite the ability to preemptively opt out of grading, detected violations carry potential costs presumed sufficient to motivate avoidance. All inspection reports are published by Maricopa County in a searchable online database. For each establishment, an initial [page](#) provides the cited number of priority violations and hyperlinks to reports of all inspections from the last three years, regardless of grading participation. Inspection results are also incorporated into the consumer-review platform, Yelp. An establishment’s Yelp profile (e.g., [here](#)) shows their most recent inspection’s letter grade or “Not Participating” in the “Amenities and More” section, and a “Health Score” hyperlink leads to a [list](#) of *all* recent inspections with violation counts and descriptions.⁸

Detected violations carry other potential costs as well. MCESD inspectors have authority to suspend or revoke operating permits. Following routine compliance inspections, failure to correct any noted violation within the time limit given is cause for suspension of the permit.⁹ With priority and priority-foundation violations, if not immediately correctable, a re-inspection within 10 days to verify correction is required. Further, repeating the same priority violation in consecutive inspections requires an additional “Active Managerial Control Intervention plan” visit at the establishment, and a future priority violation of that particular code may result in permit suspension.

2.2 COVID-19 Pandemic and Virtual Inspections

On March 19, 2020, Arizona Governor Doug Ducey issued an executive order restricting restaurants in counties with confirmed cases of COVID-19 to offer food for dine-out only.

⁸In Louisville, KY, where mandatory on-site disclosure of a compliance score was already in place, [Makofske \(2020a\)](#) finds that independent restaurants improved average compliance when Yelp began publishing these scores.

⁹See Chapter 8.1 of the [Maricopa County Environmental Health Code](#).

On May 4, 2020, he issued executive orders providing guidance on re-opening of businesses during the COVID-19 pandemic, and allowing resumption of in-person dining on May 11.¹⁰ In a May 7, 2020 press conference, MCESD Director Darcy Kober explained that, throughout the pandemic, MCESD had continued conducting on-site inspection visits, as many establishments were providing dine-out service.¹¹ During that time, MCESD recorded many “ineffective visits”, where visited establishments were found to be temporarily closed. It’s noteworthy that MCESD continued visiting establishments without making status inquiries—it suggests reluctance to reveal an imminent inspection.

On May 29, 2020, MCESD began conducting what it called “virtual inspections” at establishments with populations highly vulnerable to COVID-19, such as nursing homes, assisted living facilities, and hospitals. Specifics of the virtual inspection program are detailed in an [award application](#) submitted by MCESD. Per that application, virtual inspections were pre-scheduled and establishments were instructed they would need a thermometer and flashlight. Establishments were required to demonstrate appropriate holding temperatures for potentially hazardous foods, and sanitizer concentration for cleaning solutions with pH test strips (which MCESD code requires establishments have at all times), checks normally conducted by inspectors.

3 Data

Maricopa County’s website maintains a [list](#) of hyperlinks to inspection-result [pages](#) for all permitted food establishments. Those pages contain dates and hyperlinks to [reports](#), for all inspections conducted within the last 3 years. Establishment-page and inspection-result links were first collected on June 5, 2022. For inspections prior to June 5, 2019, I collect reports from a separate report of weekly inspection [summaries](#). Overall, data were collected for all routine inspections from January 2, 2018 through August 2, 2022.

From inspection reports I collect the health codes and information provided on cited violations and all text in the “Inspection Comments” section. In those comments, virtual

¹⁰See [here](#).

¹¹Video of the press conference is available [here](#).

inspections typically include the tag: “VIRTUAL INSPECTION – COVID-19”.¹² A total of 2,237 inspections are tagged as virtual. However, among treated establishments there are 198 inspections that occur between the earliest and latest virtual inspections observed at the establishment but not tagged as virtual, which calls into question whether these may have been virtual inspections that were mistakenly not tagged. Appendix section A1.1 describes these observations and explains why nearly all are likely reported correctly. There are however 8 observations where the virtual inspection tag appears to have been left out in error; I code these as virtual inspections.

For estimation, I restrict the sample to establishments observed in at least two inspections before, and at least one inspection on or after, May 29, 2020. Subject to that restriction, my primary sample consists of treated establishments, and all untreated establishments with the same permit type as a treated establishment (see Appendix Table A1). All establishments—with the exception of 105 classified as Micromarkets—require 2 to 4 inspections annually. I exclude observations from Micromarkets, as they receive only 1 annual inspection, are cited for 0.071 violations per on-site inspection (compared with 0.828 among all others), and all but 1 are untreated. Finally, some establishments temporarily closed during the pandemic; I exclude observations from 243 establishments (2 treated, and 241 untreated) that went an entire calendar year without an inspection due to temporary closure. This leaves a primary estimating sample of 102,807 inspections, from 505 treated, and 9,663 untreated, establishments.

4 Methodology

A total of 52 different MCESD code violations are cited within the data, all of which presumably carry lower detection probability in virtual inspections. These detection-probability decreases have two potential sources. First, *inspection anticipation* enables avoidance—committed violations that would otherwise be detected in unannounced inspections, can be corrected before the virtual inspection begins. Second, detection of

¹²See [here](#), for example. Naturally—as all inspections prior to May 29, 2020 were conducted in person—inspection reports don’t explicitly indicate inspections conducted on site.

some violations may be subject to *format limitations*—inspector difficulty observing certain violations when not physically present. Initially, I seek to isolate changes in detected compliance attributable only to inspection anticipation.

To isolate an effect of anticipation, I track a subset of regulations: (i) “food-contact surfaces: cleaned and sanitized”, (ii) “proper cold holding temperatures”, (iii) “proper cooling methods used, adequate equipment for temperature control”, (iv) “proper cooling time and temperatures”, and (v) “proper hot holding temperatures”. As in on-site inspections, compliance with these regulations must be demonstrated during virtual inspections *via* thermometer and sanitizer-test-strip readings. As such, the remote format should not inhibit detection of these “virtually demonstrable” violations.

I estimate

$$y_{i,j}^d = \alpha_1 [(1 - Virtual_{i,j}) \times Post_{i,j}] + \alpha_2 Virtual_{i,j} + \mathbf{X}'_{i,j}\boldsymbol{\omega} + a_i + \epsilon_{i,j}, \quad (1)$$

where $y_{i,j}^d$ is the count of virtually demonstrable violations detected in inspection j of establishment i . $Virtual_{i,j}$ indicates that an inspection was virtual, and a_i is an establishment fixed effect. $Post_{i,j}$ equals one if inspection j of treated establishment i occurs on or after the date of their first virtual inspection, and 0 otherwise. In the primary sample, there are 1,055 on-site inspections of treated establishments, that occur after the establishment has received a virtual inspection(s). In such inspections, $[(1 - Virtual_{i,j}) \times Post_{i,j}] = 1$, which prevents $\hat{\alpha}_2$ from reflecting comparisons against post-treatment on-site inspections. In the full specification, vector $\mathbf{X}_{i,j}$ contains fixed effects for an inspection’s day of week, month of year, and 14-day period of the sample.

In estimating α_2 , observably similar and contemporaneous on-site inspections provide a counterfactual estimate for virtual inspections. This counterfactual estimate is valid if, absent the virtual-inspection regime, treated and untreated establishments would have exhibited a common trend in $y_{i,j}^d$ following May 28, 2020. To gauge the plausibility of that assumption, I test whether the two groups exhibit common trends prior to the

virtual-inspection period. Using inspections before May 29, 2020, I estimate

$$y_{i,j}^d = \gamma_1 (\text{Treated}_i \times \text{Trend}_{i,j}) + \gamma_2 \text{Trend}_{i,j} + \gamma_3 \text{Treated}_i + \mathbf{X}'_{i,j} \boldsymbol{\omega} + c_i + \epsilon_{i,j}. \quad (2)$$

$\text{Trend}_{i,j}$ is an inspection’s month of the sample. Under a null hypothesis of common trends prior to the virtual-inspection period, $\gamma_1 = 0$. Table 1 reports these estimates. In column (1), the vector of controls is empty. In column (2), 14-day period and establishment fixed effects are included. Both specifications estimate a very small difference in pre-period trends, with fairly precise null effects—in column (2), the 99-percent confidence interval on $\hat{\gamma}_1$ is $[-0.005, 0.004]$. Columns (3) and (4) report analogous estimates using detected count of all violations, $y_{i,j}$, as the dependent variable. The dependent variable in columns (5) and (6) is a severity-adjusted count of all violations, $y_{i,j}^a$, in which each core violation adds only 0.25.¹³ Appendix Table A2 reports these same estimates using a quarterly trend; all results are very similar.

To visualize the trend comparison, Figure 1 presents simple quarter-year averages of $y_{i,j}^d$ among untreated establishments (powder-blue diamonds), on-site inspections of treated establishments (solid red circles), and virtual inspections of treated establishments (hollow red circles). Prediction lines for each group are from the simple quarterly-trend estimates reported in column (1) of Table A2. Averages for both groups track closely prior to the virtual inspection period, after which there is a marked drop among treated establishments, but *only* in virtual inspections. When on-site inspections return, average y^d for treated establishments returns the levels predicted by the simple pre-period trend.

5 Results

5.1 Anticipation and Virtually Demonstrable Violations

Columns (1), (2), and (3) of Table 2 report estimates of equation (1). Standard errors, clustered multi-way on establishment and 14-day period, are reported in parentheses. In

¹³The inspection grade becomes B given one priority violation, one priority foundation violation, or four core violations, hence the weights of 1, 1, and 0.25.

column (1), 14-day-period fixed effects and an indicator for treated establishments are the only controls. Establishment fixed effects are added in column (2), and column (3) reports estimates under the full specification.

The estimated effect of anticipation on detected compliance is substantial. Pre-treatment on-site inspections of treated establishments detect 0.240 demonstrable violations on average. Relative to that level, the full specification estimates a 53.4% relative decrease due to anticipation. Moreover, between pre- and post-treatment on-site inspections, the estimated difference in detected y^d is relatively small and statistically insignificant; the reduction observed in virtual inspections in no way persists when unscheduled on-site visits return.

Recall that $y_{i,j}^d$ tracks a subset of violations that are verifiably tested for in virtual inspections, meaning format limitations on detection ability are not likely driving these findings. Further, the 14-day-period fixed effects likely account for any general changes in compliance driven pandemic-related measures. A remaining alternative explanation however is that virtual inspections, because they required an establishment’s person-in-charge to assume a more active role, were educational and thereby caused hygiene improvements. The award application referenced in Section 2.2 suggests MCESD had hoped for this.¹⁴ If $\hat{\alpha}_2$ reflects an effect of learning through treatment: (i) these changes should be evident to some extent in subsequent on-site inspections (which estimates of α_1 contradict), and (ii) they should only manifest *after* an establishment receives a virtual inspection.

To assess whether the effect estimated by $\hat{\alpha}_2$ materializes after establishments’ first virtual inspections, I estimate

$$y_{i,j}^d = \beta_1 [(1 - Virtual_{i,j}) \times Post_{i,j}] + \beta_2 Virtual_{i,j} + \beta_3 (Virtual_{i,j} \times Post_{i,j-1}) + \mathbf{X}'_{i,j} \boldsymbol{\omega} + a_i + \epsilon_{i,j}, \quad (3)$$

¹⁴From that document: “An unexpected bonus of the virtual inspections has been the PIC being put in an active, hands-on role and learning from this. For example, the PIC must calibrate the food thermometer, verify the temperature of foods in hot-holding and/or cold-holding tables, open containers in the walk-in refrigerator and verify cold-holding temperatures, etc.”

where $Post_{i,j-1}$ is a one-inspection lag of $Post$. The interaction, $(Virtual_{i,j} \times Post_{i,j-1})$, equals 1 in all virtual inspections that come after an establishment’s first virtual inspection. If the effect estimated by equation (1) reflects better hygiene practices learned through virtual inspections, $\beta_2 = 0$.

Column (4) of Table 2 reports estimates of equation (3). The estimated decrease in establishments’ first virtual inspections ($\widehat{\beta}_2$) is substantial, and accounts for about 84.2% of the effect estimated among all virtual inspections. As an additional test, column (5) reports estimates of equation (1) under a restricted sample that ends following either: treated establishments’ first virtual inspections, or untreated establishments’ first inspections after May 28, 2020. This estimates an effect very similar to column (3), and further challenges the plausibility *any* learning effect in the initial α_2 estimates.

While the primary comparison group consists of untreated establishments with the same permit type as a treated establishment, estimates are robust to an expanded comparison group. Appendix Table A3 reports estimates analogous to Table 2, but with the comparison group expanded to include any permit type. Results are very similar.

Finally, recall that establishments irreversibly chose whether or not to participate in grade disclosure at the start of each inspection. Of the establishments in the primary sample: 996 (about 9.8%) never participate in disclosure; 3,340 (about 32.9%) always participate; and the remainder chose each option at least once.¹⁵ To assess whether participation decisions in virtual inspections are consistent with avoidance, I estimate equations (1) and (3), with $Disc_{i,j}$, a binary variable indicating establishment i chose disclosure participation in inspection j , as the outcome. These estimates are reported in Table 3, and suggest a statistically significant increase in disclosure propensity in virtual inspections. Across all virtual inspections, a 5.9% increase is estimated relative to the pre-treatment average of 0.780. While modest in magnitude, the direction of this change, and its partial reversal in post-treatment on-site inspections, are consistent with opportunistic use of anticipation ability.

¹⁵For comparison, from the grade program’s introduction in 2011, through 2013, [Bederson et al. \(2018\)](#) find that only 58% of establishments ever participate.

5.2 Testing Deterrence

The introduction of virtual inspections causes a sharp drop in detection probability at treated establishments. Deterrence theory suggests that treated establishments—conditional on recognizing this and expecting its continuation—will become less compliant. In initial post-treatment on-site inspections, while treated establishments’ compliance efforts likely reflect virtual-regime perceptions, actual detection probability returns to the pre-treatment level, thereby removing the detection effect and isolating any deterrence effect.

In assessing potential adjustments to detection probability, I use an inspection’s detected count of all violations, $y_{i,j}$, as well as the severity-adjusted count of all violations, $y_{i,j}^a$ (described in Section 4). Virtual inspections likely lowered detection probabilities for all health-code violations, hence the shift to these broader outcomes. Columns (3), (4), (5), and (6) of Table 1, suggest very similar pre-period trends in $y_{i,j}$ and $y_{i,j}^a$, between treated and untreated establishments.

I test deterrence by estimating equation (1), but using y and y^a as dependent variables. Any inspections of treated establishments after their initial post-treatment on-site inspections are excluded in estimation, as are all inspections from treated establishments not observed in a post-treatment on-site inspection. The coefficient of interest, $\hat{\alpha}_1$, estimates the difference in conditional expectation of y (or y^a) between treated establishments’ pre-treatment, and initial post-treatment, on-site inspections. If treated establishments don’t respond to the lower detection probability—or do respond, but anticipate the return of on-site visits and adjust back—then $\alpha_1 = 0$. Alternatively, if they respond in a manner consistent with general deterrence, and are caught unawares by the return of on-site inspections, $\alpha_1 > 0$.

These estimates are reported in Table 4. As expected, detected-violation counts are substantially lower in virtual inspections; with all controls included, relative to the pre-treatment average of 0.646, a 41.2% decrease in detected violations is estimated. Further, that decrease is more than offset by the return of unannounced on-site visits. Consistent with general deterrence, establishments’ initial post-treatment on-site inspections detect

violation counts that exceed the pre-treatment average by 28.2%. Columns (4) and (5) of Table 4 report similar estimates using the severity-adjusted count of violations, $y_{i,j}^a$ as the dependent variable. With all controls included, severity-adjusted violation counts in establishments' initial post-treatment on-site inspections are 20.5% higher than the pre-treatment average.¹⁶

5.3 Heterogeneity and Policy Implications

Here, as in many settings, uniform deployment of monitoring resources is presumably inefficient. Establishments differ in the social costs they pose through potential noncompliance, their propensities for noncompliance, and the sensitivity of their compliance effort to inspection frequency. MCESD's system for allocating inspections according to permit type, primarily accounts for differences along those first two dimensions. Exploiting the ability to assess heterogeneity along the third, I identify a simple dynamic-enforcement mechanism for improving inspection allocation, and enhancing general deterrence among high-risk-class establishments, with existing inspection resources.

Typically, dynamic-enforcement policies raise the expected costs of violations, at the individual level, following observed noncompliance. General deterrence is presumably the primary benefit of dynamic enforcement—the threat of greater penalties or scrutiny in the future, further deters violations at present. Triggering escalation under such policies may also induce specific deterrence.¹⁷ [Blundell \(2020\)](#) and [Blundell et al. \(2020\)](#) examine the EPA's approach to enforcing air quality regulation—where current noncompliance raises future fines—and find evidence supporting dynamic enforcement in that context. Here, I assess the potential gains from a policy of dynamically adjusting inspection frequency, and by extension, detection probability.

If inspection costs are similar across establishments, then redirecting an inspection—away from an original designee, and toward a targeted establishment—is an improvement so long as it causes noncompliance costs to decrease by more at the target than they in-

¹⁶Appendix Table A4 reports the same estimates but with virtual inspections dropped in estimation. In all cases, estimates of α_1 are very similar or slightly larger than those in Table 4.

¹⁷[Makofske \(2020b\)](#) finds evidence of specific deterrence—future compliance improvements caused by receiving punishment—in Las Vegas, NV, food-service health inspections.

crease at the original designee. However, such improvements would be difficult to verify *ex post*; naturally, redirection makes compliance less observable at original designees. Alternatively, I evaluate the potential effects of redirecting inspections, by exploiting individual-level responses to the virtual regime’s detection-probability shock.

In the primary sample, 435 establishments are treated with virtual inspections *and* observed in at least one post-treatment on-site inspection. The distribution of these establishments across risk class (from *class 2* through *class 5*) is: 146, 23, 30, and 236. I evaluate establishment-specific deterrence effects using a simple difference—denoted \tilde{y}_i^a —between the severity-adjusted count detected in establishment *i*’s first post-treatment on-site inspection, and *i*’s average severity-adjusted count from pre-treatment inspections.

Deterrence effects are more likely among establishments where the potential social costs of noncompliance are generally higher. Among all 435 focal establishments, 30.5% of the *class-5* establishments, versus 13.6% of the sub-*class-5* establishments, have $\tilde{y}_i^a > 0$. Notably, lower risk-class establishments with perfect pre-treatment records, were overwhelmingly unresponsive to the detection-probability decrease. Among 108 sub-*class-5* establishments with perfect observed pre-treatment compliance, 88% also have $\tilde{y}_i^a = 0$ (even though *any* detected violation in their initial post-treatment on-site inspection would generate a positive difference). Redirecting some routine inspections away from this group, and toward targeted *class-5* establishments, might meaningfully reduce non-compliance costs on net. Further, publicly adopting a policy where additional inspections are triggered by specified inspection outcomes, will enhance general deterrence among the possible targets.

Among sampled risk classes, the social cost of any particular violation is presumed lowest in *class 2*. There are 79 treated *class-2* establishments with perfect observed compliance in four or more consecutive pre-treatment inspections. In their first post-treatment on-site inspections, 70 have perfect compliance, and average y_i^a is 0.104. Given the apparently low rate at which this group’s compliance efforts respond to the decline in detection probability, four consecutive perfect inspections seems a useful condition for redirecting inspections away from *class-2* establishments.

Among all (treated or untreated) *class-2* establishments observed in the data, 1,127 exhibit perfect compliance across four consecutive inspections from 2018 through 2019. Following four consecutive perfect inspections at a *class-2* establishment, if MCESD were willing to conduct one less routine inspection in the next year, it could free up considerable resources for inspecting establishments that may warrant greater scrutiny. Note that MCESD needn't commit to reducing inspections on this condition. Rather, they could privately use the condition to enable additional inspections as necessitated by a dynamic-enforcement rule.

Over the same period (2018 through 2019), *class-5* establishments had $y_i^a \geq 2$ in 434 inspections. In 420 observed and subsequent on-site inspections, $y_i^a \geq 2$ was repeated in 106 (25.24%). Consider a dynamic-enforcement policy requiring one additional routine inspection over the next year, for *class-5* establishments following $y_i^a \geq 2$. If this group's probability of repeating $y_i^a \geq 2$ in consecutive inspections were (figuring conservatively) 0.3, then the 434 instances observed over two years would ultimately require 620 additional inspections—well below the 1,127 inspections potentially freed up by highly compliant *class-2* establishments over the same period. Note also that with this policy in effect, $y_i^a \geq 2$ among *class-5* establishments would likely occur at a lower relative frequency due to general deterrence.

6 Concluding Remarks

General deterrence through imperfect monitoring is essential to enforcing a profound body of regulation. Yet, the theory of general deterrence is, by nature, difficult to empirically evaluate. Exploiting MCESD's temporary adoption of virtual compliance inspections among some establishments, I largely overcome the typical empirical obstacles.

I find that establishments exploit inspection anticipation to avoid detection of noncompliance. This contributes to recent work (Makofske, 2019, 2021; Zou, 2021) demonstrating the detrimental effect of anticipation ability on monitoring programs. Here, establishments with no prior history of anticipation ability suddenly acquire it; whereas prior

studies examine anticipation ability stemming from long-standing institutional practices. I find that avoidance behavior is immediate, suggesting that even sporadic provision of anticipation ability might significantly undermine enforcement.

I also find that compliance efforts respond to perceived detection probabilities in a manner consistent with general deterrence. In establishments' initial post-treatment on-site inspections, detected violations exceed pre-treatment levels by 28%. Moreover, considerable heterogeneity underlies this average effect. Notably, establishments that were highly compliant in observed pre-treatment inspections, and with permit types in lower risk classes, appear largely unresponsive to the reduction in detection probability and expected cost. Redirecting some inspections away from these establishments, and toward targets in the highest risk class, could significantly improve how inspections are allocated. Moreover, if targeting is explicitly tied to detected noncompliance, beyond improving inspection allocation, enhanced general deterrence should further reduce non-compliance costs. Existing MCESD inspection resources appear sufficient to comfortably accomplish this through a straightforward dynamic-enforcement policy.

Finally, note that MCESD was hardly alone in adopting virtual inspections; many agencies utilized the remote format during the COVID-19 pandemic, and some did so for all food establishments in their jurisdictions.¹⁸ This point is particularly important because presently—as with activities which rapidly migrated to remote format during the pandemic—debate exists over whether remote inspections should continue to some extent in food-safety regulation.¹⁹ While no doubt less costly, my results demonstrate that in this regulatory setting—or any where compliance status can change in the time between a virtual inspection's start and its requisite advance scheduling—remote inspections are a remarkably poor substitute for unannounced on-site visits.

¹⁸See <https://www.astho.org/topic/brief/virtual-food-safety-inspections-during-the-covid-19-pandemic/>.


¹⁹See, e.g., [here](#) or [here](#).

References

- Becker, G. (1968). Crime and punishment: An economic approach. *Journal of Political Economy* 76(2), 169–217.
- Bederson, B. B., G. Z. Jin, P. Leslie, A. J. Quinn, and B. Zou (2018). Incomplete disclosure: Evidence of signaling and countersignaling. *American Economic Journal: Microeconomics* 10(1), 41–66.
- Blundell, W. (2020). When threats become credible: A natural experiment of environmental enforcement from Florida. *Journal of Environmental Economics and Management* 101, 102288.
- Blundell, W., G. Gowrisankaran, and A. Langer (2020). Escalation of scrutiny: The gains from dynamic enforcement of environmental regulations. *American Economic Review* 110(8), 2558–2585.
- Duflo, E., M. Greenstone, R. Pande, and N. Ryan (2018). The value of regulatory discretion: Estimates from environmental inspections in India. *Econometrica* 86(6), 2123–2160.
- Eckert, H. (2004). Inspections, warnings, and compliance: The case of petroleum storage regulation. *Journal of Environmental Economics and Management* 47(2), 232–259.
- Gray, W. and M. E. Deily (1996). Compliance and enforcement: Air pollution regulation in the U.S. steel industry. *Journal of Environmental Economics and Management* 31(1), 96–111.
- Gray, W. B. and J. P. Shimshack (2011). The Effectiveness of Environmental Monitoring and Enforcement: A Review of the Empirical Evidence. *Review of Environmental Economics and Policy* 5(1), 3–24.
- Hoffmann, S., B. Macculloch, and M. Batz (2015). Economic burden of major foodborne illnesses acquired in the United States. Economic information bulletin number 140, United States Department of Agriculture Economic Research Service.

- Jin, G. Z. and J. Lee (2014). Inspection technology, detection, and compliance: evidence from Florida restaurant inspections. *RAND Journal of Economics* 45(4), 885–917.
- Laplante, B. and P. Rilstone (1996). Environmental inspections and emissions of the pulp and paper industry in Quebec. *Journal of Environmental Economics and Management* 31(1), 19–36.
- Makofske, M. P. (2019). Inspection regimes and regulatory compliance: How important is the element of surprise? *Economics Letters* 177(C), 30–34.
- Makofske, M. P. (2020a). The effect of information salience on product quality: Louisville restaurant hygiene and Yelp.com. *The Journal of Industrial Economics* 68(1), 52–92.
- Makofske, M. P. (2020b). Disclosure policies in inspection programs: The role of specific deterrence. *Economics Letters* 196(C), 109533.
- Makofske, M. P. (2021). Spoiled food and spoiled surprises: Inspection anticipation and regulatory compliance. *Journal of Economic Behavior and Organization* 190(C), 348–365.
- Moritz, E. D., S. D. Ebrahim-Zadeh, B. Wittry, M. M. Holst, B. Daise, A. Zern, T. Taylor, A. Kramer, and L. G. Brown (2023). Foodborne Illness Outbreaks at Retail Food Establishments—National Environmental Assessment Reporting System, 25 State and Local Health Departments, 2017-2019. *MMWR Surveillance Summaries* 72(6), 1.
- Telle, K. (2009). The threat of regulatory environmental inspection: Impact on plant performance. *Journal of Regulatory Economics* 35(2), 154–178.
- Zou, E. Y. (2021). Unwatched pollution: The effect of intermittent monitoring on air quality. *American Economic Review* 111(7), 2101–26.

Figure 1: INSPECTION FORMAT AND DETECTED VIOLATIONS



Average $y_{i,j}^d$ by quarter-year of sample. The “treated group” are establishments that received at least one virtual inspection. The “untreated group” are establishments with the same permit type as a treated establishment, that did not receive a virtual inspection. Prediction lines (navy for untreated, maroon for treated) are simple quarterly trend estimates from observations prior to May 29, 2020. Treated group averages from on-site inspections in 2020q3, 2020q4, and 2021q1 are suppressed due to few observations; 6, 5, and 14, respectively—by comparison there were 105 and 240 such inspections in 2021q2 and 2021q3.

Table 1: ASSESSING PRE-PERIOD TRENDS

| VARIABLE | (1) $y_{i,j}^d$ | (2) $y_{i,j}^d$ | (3) $y_{i,j}$ | (4) $y_{i,j}$ | (5) $y_{i,j}^a$ | (6) $y_{i,j}^a$ |
|-------------------------------|----------------------|--------------------|----------------------|------------------|----------------------|--------------------|
| <i>Trend</i> × <i>Treated</i> | 0.000 (0.002) | -0.001 (0.002) | 0.005 (0.004) | 0.003 (0.003) | 0.002 (0.003) | 0.001 (0.003) |
| <i>Trend</i> | -0.003*** (0.001) | 0.010 (0.011) | -0.010*** (0.002) | 0.012 (0.023) | -0.009*** (0.002) | 0.014 (0.019) |
| <i>Treated</i> | -0.055* (0.033) | | -0.430*** (0.070) | | -0.267*** (0.058) | |
| 14-day period FE | | ✓ | | ✓ | | ✓ |
| Establishment FE | | ✓ | | ✓ | | ✓ |
| R-squared | 0.002 | 0.384 | 0.006 | 0.514 | 0.006 | 0.482 |
| N | 52,006 | 52,006 | 51,696 | 51,696 | 51,696 | 51,696 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

OLS estimates of equation (2) from inspections prior to May 29, 2020. Standard errors, clustered two-way on establishment and 14-day period, are reported in parentheses. $y_{i,j}^d$ is an inspection's detected count of demonstrable violations, $y_{i,j}$ is an inspection's detected count of all violations. $y_{i,j}^a$ is a severity-adjusted count of all violations in which each core violation adds only 0.25. *Trend* is the month of sample and equals 1 in January 2018. Estimating sample in columns (3), (4), (5), and (6), excludes treated establishments that are not observed in a post-treatment on-site inspection.

Table 2: ANTICIPATION ABILITY AND DETECTED COMPLIANCE

| VARIABLE | (1) $y_{i,j}^d$ | (2) $y_{i,j}^d$ | (3) $y_{i,j}^d$ | (4) $y_{i,j}^d$ | (5) $y_{i,j}^d$ |
|--------------------------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| $(1 - Virtual) \times Post$ | -0.019 (0.023) | -0.014 (0.022) | -0.015 (0.022) | -0.014 (0.022) | |
| <i>Virtual</i> | -0.112*** (0.018) | -0.129*** (0.018) | -0.128*** (0.018) | -0.105*** (0.028) | -0.099*** (0.035) |
| <i>Virtual</i> \times $Post_{j-1}$ | | | | -0.030 (0.028) | |
| <i>Treated</i> | -0.063*** (0.017) | | | | |
| 14-day period FE | ✓ | ✓ | ✓ | ✓ | ✓ |
| Establishment FE | | ✓ | ✓ | ✓ | ✓ |
| Month-of-year FE | | | ✓ | ✓ | ✓ |
| Day-of-week FE | | | ✓ | ✓ | ✓ |
| R-squared | 0.013 | 0.292 | 0.293 | 0.293 | 0.355 |
| N | 102,807 | 102,807 | 102,807 | 102,807 | 62,239 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

OLS estimates of equations (1) and (3). Standard errors, clustered two-way on establishment and 14-day period, are reported in parentheses. $y_{i,j}^d$ is an inspection's detected count of demonstrable violations. $Post_{i,j-1}$ equals 1 in all inspections after an establishment's first virtual inspection, and 0 otherwise. Column (5) estimating sample: treated establishments dropped following first treated inspection; untreated establishments dropped following first inspection after May 28, 2020.

Table 3: ANTICIPATION ABILITY AND DISCLOSURE DECISIONS

| Variable | (1) $Disc_{i,j}$ | (2) $Disc_{i,j}$ | (3) $Disc_{i,j}$ | (4) $Disc_{i,j}$ |
|-----------------------------|---------------------|---------------------|---------------------|---------------------|
| $(1 - Virtual) \times Post$ | 0.028 (0.020) | 0.028 (0.020) | 0.028 (0.020) | 0.027 (0.020) |
| $Virtual$ | 0.047*** (0.017) | 0.046*** (0.017) | 0.020 (0.015) | 0.019 (0.015) |
| $Virtual \times Post_{j-1}$ | | | 0.035** (0.015) | 0.035** (0.015) |
| 14-day period FE | ✓ | ✓ | ✓ | ✓ |
| Establishment FE | ✓ | ✓ | ✓ | ✓ |
| Month-of-Year FE | | ✓ | | ✓ |
| Day-of-Week FE | | ✓ | | ✓ |
| R-squared | 0.572 | 0.572 | 0.572 | 0.573 |
| N | 102,807 | 102,807 | 102,807 | 102,807 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

OLS estimates. Standard errors, clustered two-way on establishment and 14-day period, are reported in parentheses. $Disc_{i,j}$ is a binary variable, indicating that establishment i participated in grading in inspection j . $Post_{i,j-1}$ equals 1 in all inspections after an establishment's first virtual inspection, and 0 otherwise.

Table 4: TESTING DETERRENCE

| VARIABLE | (1) $y_{i,j}$ | (2) $y_{i,j}$ | (3) $y_{i,j}$ | (4) $y_{i,j}^a$ | (5) $y_{i,j}^a$ |
|-----------------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| $(1 - Virtual) \times Post$ | 0.126* (0.064) | 0.181*** (0.050) | 0.182*** (0.051) | 0.100** (0.041) | 0.101** (0.041) |
| <i>Virtual</i> | -0.247*** (0.041) | -0.270*** (0.045) | -0.266*** (0.045) | -0.197*** (0.038) | -0.195*** (0.038) |
| <i>Treated</i> | -0.372*** (0.042) | | | | |
| 14-day period FE | ✓ | ✓ | ✓ | ✓ | ✓ |
| Establishment FE | | ✓ | ✓ | ✓ | ✓ |
| Month-of-year FE | | | ✓ | | ✓ |
| Day-of-week FE | | | ✓ | | ✓ |
| R-squared | 0.021 | 0.420 | 0.420 | 0.388 | 0.388 |
| N | 101,365 | 101,365 | 101,365 | 101,365 | 101,365 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

OLS estimates of equation (1). Standard errors, clustered two-way on establishment and 14-day period, are reported in parentheses. $y_{i,j}$ is an inspection's detected count of all violations. $y_{i,j}^a$ is a severity-adjusted count of all violations in which each core violation adds only 0.25. $Post_{i,j-1}$ equals 1 in all inspections after an establishment's first virtual inspection, and 0 otherwise. Estimating sample: for untreated establishments, all inspections; for treated establishments, all inspections prior to, and including, their first post-treatment on-site inspection. Treated establishments with no observed post-treatment on-site inspections are excluded.

A1 Appendix

A1.1 Notes on Data Cleaning

Among treated establishments, there are 190 inspections that are not tagged as virtual in the inspection comments, but that occur after the earliest, and before the latest virtual inspection observed at that establishment. A question arises: were any of these untagged inspections actually virtual? Naturally—as all inspections prior to May 2020 were in-person—inspection reports don’t explicitly indicate inspections conducted on-site. Here, I explain why most these inspections are very likely coded correctly, and only a small number are truly questionable.

Figure A2 shows the distribution of these “flagged” inspections by month of sample. About 93.2% of the inspections in question occur during or after May 2021, by which point COVID-19 vaccines had been widely available in Maricopa County.²⁰ Moreover, the decline in frequency of these inspections beginning November 2021 coincides with the rise of delta-variant infections and the eventual emergence of the omicron variant. This suggests the flagged inspections from May 2021 onward likely did occur on-site, and the subsequent return to virtual inspections at these establishments was driven by the rise of new and more transmissible variants. Thus, I treat all 177 flagged inspections from May 2021 onward as on-site inspections.

However, the 7 flagged inspections in September and December 2020 all occurred at facilities with highly susceptible populations, at a time where all other such facilities were receiving virtual inspections. Moreover, these are all cases where the establishments’ prior and next inspections are indicated as virtual. I treat these 7 inspections as virtual. Of the remaining 6 flagged inspections—those made in 2021 but before May—all but 1 begin a sequence of multiple inspections not tagged as virtual at respective establishments. I treat those 5 as correctly reported on-site inspections, and the other inspection as virtual inspection that was mistakenly not tagged.

²⁰See <https://www.maricopa.gov/5671/Public-Vaccine-Data>.

Figure A1: MCESD GRADING SCALE

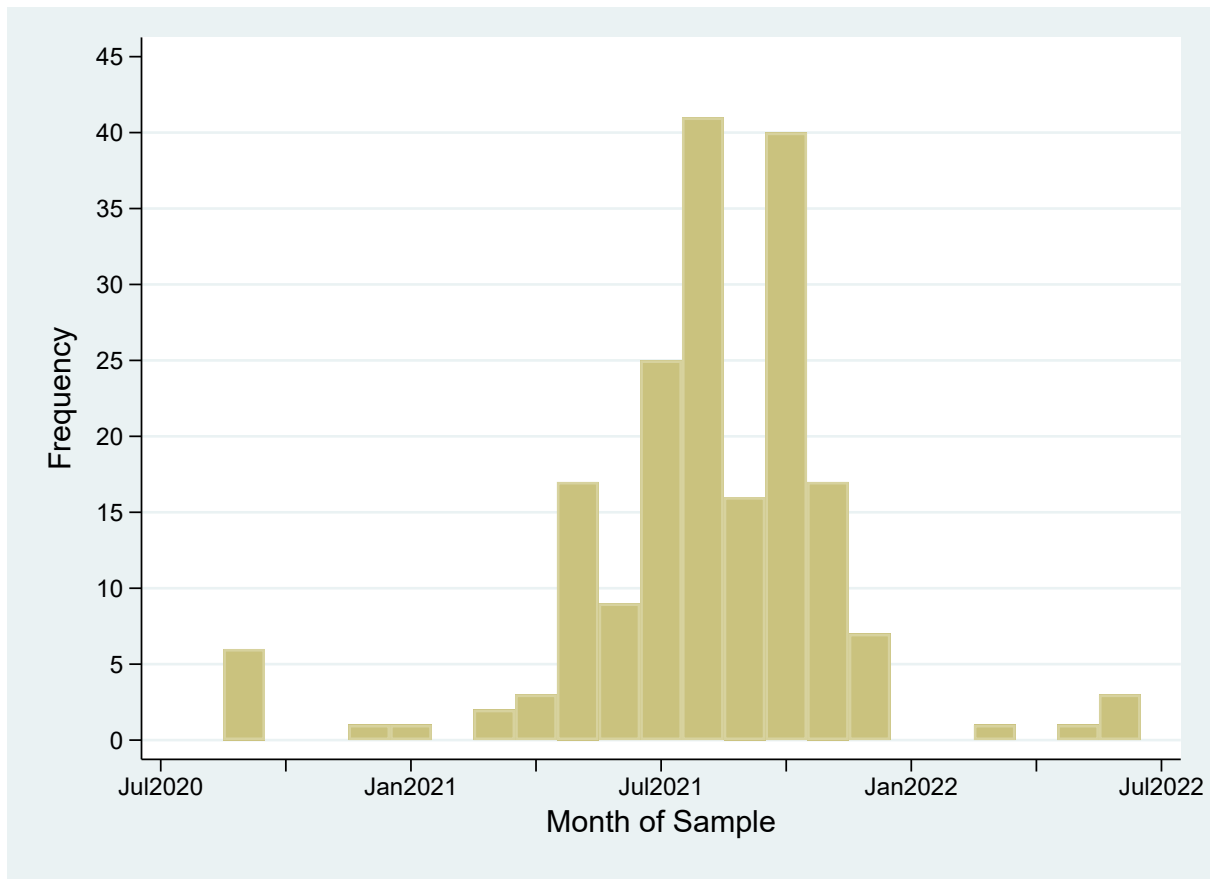
Grading System

| | | Priority Foundation Violations | | | | |
|---------------------|-----------------------|--------------------------------|---|---|---|----|
| | | 0 | 1 | 2 | 3 | 4+ |
| Priority Violations | Count of Violation(s) | | | | | |
| | 0 | A | B | B | C | D |
| | 1 | B | B | B | C | D |
| | 2 | C | C | C | C | D |
| 3+ | D | D | D | D | D | |

- Four or more Core violations drops one grade level
- Any legal action results in a D

Retrieved from <https://envapp.maricopa.gov/Images/JPG/GradingSystem.jpg> on June 12, 2022.

Figure A2: FREQUENCY OF FLAGGED INSPECTIONS



Frequency distribution of the 198 inspections that are not indicated as being virtual, but that occur in between the establishment's earliest and latest observed virtual inspections.

Table A1: ESTABLISHMENT TYPES

| PERMIT TYPE | Treated with Virtual Inspection | |
|--------------------------------|---------------------------------|-----|
| | No | Yes |
| | NUMBER OF ESTABLISHMENTS | |
| Adult Daycare | 2 | 2 |
| Assisted Living | 0 | 129 |
| Bakery | 252 | 0 |
| Boarding Home | 26 | 0 |
| Bottled Water & Beverage Plant | 28 | 0 |
| Damaged Foods | 4 | 0 |
| Daycare Food Service | 203 | 9 |
| Eating & Drinking | 7,980 | 116 |
| Food Bank | 29 | 1 |
| Food Catering | 363 | 6 |
| Food Jobber | 175 | 0 |
| Food Processor | 345 | 4 |
| Hospital Food Service | 1 | 38 |
| Ice Manufacturing | 3 | 0 |
| Jail Food Service | 2 | 0 |
| Meat Market | 322 | 0 |
| Micromarket | 104 | 1 |
| Nursing Home | 0 | 57 |
| Refrigeration Warehouse | 3 | 0 |
| Retail Food Establishment | 1,371 | 3 |
| School Food Service | 288 | 0 |
| Senior Food Service | 1 | 1 |
| Service Kitchen | 149 | 164 |

Count of different permit types among establishments in the raw data.

Table A2: ASSESSING COMMON TRENDS ASSUMPTION

| Variable | (1) $y_{i,j}^d$ | (2) $y_{i,j}^d$ | (3) $y_{i,j}$ | (4) $y_{i,j}$ | (5) $y_{i,j}^a$ | (6) $y_{i,j}^a$ |
|---|----------------------|--------------------|----------------------|-------------------|----------------------|--------------------|
| <i>Quarterly Trend</i> × <i>Treated</i> | -0.001 (0.005) | -0.002 (0.005) | 0.013 (0.011) | 0.009 (0.010) | 0.005 (0.009) | 0.004 (0.009) |
| <i>Quarterly Trend</i> | -0.010*** (0.003) | 0.009 (0.020) | -0.032*** (0.007) | -0.019 (0.046) | -0.027*** (0.005) | -0.006 (0.039) |
| <i>Treated</i> | -0.047 (0.034) | | -0.422*** (0.075) | | -0.260*** (0.061) | |
| 14-day period FE | | ✓ | | ✓ | | ✓ |
| Establishment FE | | ✓ | | ✓ | | ✓ |
| R-squared | 0.002 | 0.384 | 0.006 | 0.514 | 0.006 | 0.482 |
| N | 52,006 | 52,006 | 51,696 | 51,696 | 51,696 | 51,696 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

OLS estimates from inspections prior to May 29, 2020. Standard errors, clustered two-way on establishment and 14-day period, are reported in parentheses. Estimating sample in columns (3), (4), (5), and (6), excludes treated establishments that are not observed in a post-treatment on-site inspection. *Quarterly Trend* is the quarter-year of the sample, equal to 1 for January-March 2018. $y_{i,j}^d$ is an inspection's detected count of demonstrable violations. $y_{i,j}$ is an inspection's detected count of all violations.

Table A3: ROBUSTNESS TO EXPANDED COMPARISON GROUP

| Variable | (1) $y_{i,j}^d$ | (2) $y_{i,j}^d$ | (3) $y_{i,j}^d$ | (4) $y_{i,j}^d$ | (5) $y_{i,j}^d$ |
|--------------------------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| $(1 - Virtual) \times Post$ | -0.020 (0.023) | -0.015 (0.022) | -0.015 (0.022) | -0.015 (0.022) | |
| <i>Virtual</i> | -0.116*** (0.018) | -0.133*** (0.018) | -0.133*** (0.018) | -0.111*** (0.028) | -0.105*** (0.034) |
| <i>Virtual</i> \times $Post_{j-1}$ | | | | -0.029 (0.028) | |
| <i>Treated</i> | -0.043** (0.017) | | | | |
| 14-day period FE | ✓ | ✓ | ✓ | ✓ | ✓ |
| Establishment FE | | ✓ | ✓ | ✓ | ✓ |
| Month-of-year FE | | | ✓ | ✓ | ✓ |
| Day-of-week FE | | | ✓ | ✓ | ✓ |
| R-squared | 0.012 | 0.296 | 0.296 | 0.296 | 0.358 |
| N | 112,541 | 112,541 | 112,541 | 112,541 | 68,311 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

OLS estimates from expanded sample including establishments of any type, with at least two inspection before, and at least one inspection on or after May 29, 2020. Standard errors, clustered two-way on establishment and 14-day period, are reported in parentheses. $y_{i,j}^d$ is an inspection's detected count of demonstrable violations. $Post_{i,j-1}$ equals 1 in all inspections after an establishment's first virtual inspection, and 0 otherwise.

Table A4: TESTING DETERRENCE: VIRTUAL INSPECTIONS EXCLUDED

| Variable | (1) $y_{i,j}$ | (2) $y_{i,j}$ | (3) $y_{i,j}$ | (4) $y_{i,j}^a$ | (5) $y_{i,j}^a$ |
|-----------------------------|----------------------|---------------------|---------------------|---------------------|---------------------|
| $(1 - Virtual) \times Post$ | 0.125* (0.064) | 0.201*** (0.050) | 0.202*** (0.051) | 0.114*** (0.039) | 0.115*** (0.039) |
| <i>Treated</i> | -0.372*** (0.042) | | | | |
| 14-day period FE | ✓ | ✓ | ✓ | ✓ | ✓ |
| Establishment FE | | ✓ | ✓ | ✓ | ✓ |
| Month-of-year FE | | | ✓ | | ✓ |
| Day-of-week FE | | | ✓ | | ✓ |
| R-squared | 0.016 | 0.419 | 0.420 | 0.388 | 0.388 |
| N | 99,695 | 99,695 | 99,695 | 99,695 | 99,695 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

OLS estimates of equation (3), but with all virtual inspections excluded in estimation. Standard errors, clustered two-way on establishment and 14-day period, are reported in parentheses. $y_{i,j}$ is an inspection's detected count of all violations. $y_{i,j}^a$ is a severity-adjusted count of all violations in which each core violation adds only 0.25.