AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

AMERICAN ASSOCIATION OF WINE ECONOMISTS

AAWE WORKING PAPER No. 172 *Economics*

A TALE OF REPETITION: LESSONS FROM FLORIDA RESTAURANT INSPECTIONS

Ginger Zhe Jin and Jungmin Lee

Oct 2014

www.wine-economics.org

A TALE OF REPETITION: LESSONS FROM FLORIDA RESTAURANT INSPECTIONS

Ginger Zhe Jin and Jungmin Lee

ABSTRACT

We examine the role of repetition in government regulation. Using Florida restaurant inspection data from 2003 to 2010, we find that inspectors new to the inspected restaurant report 12.7-17.5% more violations than the second visit of a repeat inspector. This effect is even more pronounced if the previous inspector had inspected the restaurant more times. The difference between new and repeat inspectors is driven partly by inspector heterogeneity in inherent taste and stringency, and partly by new inspectors having fresher eyes in the first visit of a restaurant. These findings highlight the importance of inspector assignment in regulatory outcomes.

JEL No. D02, D81, D82, K32, L51

Ginger Zhe Jin, University of Maryland, Department of Economics,
3115F Tydings Hall, College Park, MD 20742-7211 and NBER, jin@econ.umd.edu

Jungmin Lee, Sogang University Seoul, Korea junglee@sogang.ac.kr

We have received constructive comments from Roger Betancourt, David Becker, Joerg Claussen, John Ham, Panle Jia, Terence Johnson, Phillip Leslie, Michael Luca, Louis Rossiter, John Rust, Andrew Foster, and seminar participants at Maryland, College of Williams and Mary, MIT, Johns Hopkins, New York University, Brown University and the 2011 IIOC conference. Jin acknowledges financial support from the Alfred P. Sloan Foundation (grant # 2011-10-14). Special thanks to the Florida Division of Hotels and Restaurants for providing us the data and patiently answering our questions. Yiyan (Echo) Liu provided excellent research assistance. Ginger Zhe Jin co-founded Hazel Analytics, a company that promotes the use of open government data and owns some equity in that company. The views expressed herein are those of the authors and do not

necessarily reflect the views of the National Bureau of Economic Research.

1 Introduction

From nuclear power to food safety, many regulations mandate that government employees inspect economic entities on a regular basis. Such an inspection introduces a classical double-moral-hazard problem: on the inspector side, government employed inspectors may not detect or report every violation as the principal desires; on the inspectee side, regulated firms may not comply with every rule set by the principal; their degree of compliance depends on the inspector's ability to detect violations and the subsequent punishment for reported violations. Many a suggestions have been made to alleviate the inspector moral hazard, including outcome-based contracts,¹ targeted auditing, reduction of information rents, high penalties for corrupt inspectors, or intentional selection of biased employees.² However, these "optimal" solutions – often made in a theoretical framework – are difficult to implement in reality because bureaucratic agencies are subject to rigid compensation schemes and limited resources.

This paper focuses on repetition in inspection programs, a tool commonly used in practice but rarely studied. In particular, for regular unannounced inspections, the program inspects the firm repeatedly and the inspector can be repeat or new to the firm. In addition, a typical inspection program may schedule follow-up visit(s) to ensure that violations detected in a regular inspection are corrected in a timely manner. Both types of repetition aim to enhance compliance, but follow-up visits often target a small fraction of firms with severe violations and only focus on the violations found in the last regular inspection. In comparison, regular unannounced inspections are applicable to all firms and all kinds of potential violations. For this reason, this paper will focus on repetition of regular unannounced inspections, leaving the economics of follow-up inspections to a separate paper.

Using a simple game-theoretical framework, we show that, in a regular unannounced inspection, both detection and compliance may differ according to whether the inspector is new or repeat. By definition, new inspectors have never inspected the firm before and therefore may find it more difficult to detect problems than repeat inspectors; however,

¹For example, the principal may set inspectors' compensation conditional on the reported violations or design a dynamic contract to prevent collusion between inspectors and the regulated (Tirole 1986, Martimort 1999).

²See Laffont and Tirole (1993), Mookherjee and Png (1989, 1995) for specific theories and Prendergast (1999) and Dixit (2002) for comprehensive reviews. Prendergast (2007) focuses on biased bureaucrats in particular.

repeat inspectors may slack over time and an on-going relationship may encourage the inspectee to learn and cater to the idiosyncratic taste of the inspector rather than comply with the regulation as a whole. The first visit to the firm may also equip new inspectors with “fresher eyes” and encourage them to examine the firm more thoroughly. Which effect dominates remains an empirical question.

We examine and measure these effects in a universe of restaurant hygiene inspections in Florida from July 2003 and March 2010. In particular, we find that inspector-restaurant relationship plays an important role in the outcome of regular inspections. Within repeat inspectors, one extra visit leads to merely a 0.7-1.9% reduction in reported violations. In comparison, not only do new inspectors report 12.7-17.5% more violations than the second visit of a repeat inspector, but this effect is also more pronounced if the previous inspector has had a longer relationship with the restaurant. Regarding the potential mechanisms behind these data patterns, we find that the difference between new and repeat inspectors is partly driven by inspector heterogeneity in inherent taste and stringency, and partly driven by new inspectors having a pair of fresher eyes in the first visit to a restaurant. These two effects suggest that both inspector heterogeneity and inspector rotation are important in determining the effectiveness of government inspection.

Our work contributes to several strands of the economics literature. First, our empirical analysis complements the large literature on principal-agent theory. Using a method articulated in our previous work (Jin and Lee forthcoming), we show that partial identification of detection and compliance can be achieved in administrative data without explicit random experiments.³ Second, we highlight the role of repetition in a typical inspection program. Theorists often worry that inspectors may be captured by a cozy relationship with the inspectee. This conjecture is unlikely to hold in our context because the probability of a fine is extremely low for a regular inspection (1.6%) and the fine amount (average \$861 if fined) is determined by a separate branch, not the inspector. Rather, we find that inspectors are inherently heterogeneous in taste and new inspectors tend to find more violations even after we control for inspectors’ inherent taste heterogeneity. Based on these two effects, we conduct counterfactuals to compare the benefits of inspector rotation ver-

³Levine, Toffel and Johnson (2012) utilized the randomness of workplace safety inspections to examine how inspections affected injury rates and other outcomes in California. Kleven et al. (2010) studied how tax payers responded to randomized tax auditing in Denmark. Slemrod and Yitzhaki (2002) provided a survey of taxpayer response to taxation.

sus the benefits of homogenizing inspector preferences. We believe our approach can have important implications beyond restaurant inspections, as inspector heterogeneity has been found in other inspection programs, but the previous literature does not differentiate the potential reasons behind the documented heterogeneity.⁴

The rest of the paper is organized as follows. Section 2 starts with a simple static model highlighting the game-theoretical interaction between inspector and restaurant. We then extend the model to emphasize the difference between new and repeat inspectors in regular inspections. Section 3 describes the data and background of Florida restaurant inspection, with an emphasis on the randomness of inspector assignment. Section 4 presents empirical estimates of new and repeat inspectors, and explicitly separates the influence of inspector heterogeneity from the fresh eye effect of new inspectors. These structural estimates allow us to conduct counterfactual simulations in Section 5, of how many more violations could have been detected if inspectors were assigned randomly or inspectors were trained to be homogeneous. A brief conclusion is offered in Section 6.

2 Model and Identification

This section starts with a stylized static model that incorporates restaurant effort, inspector effort, and inspector taste assuming perfect information. We then extend it to allow for the restaurant being uncertain about the identity of the next inspector. Under this uncertainty, we discuss several scenarios, clarify the extent to which detection and compliance can be partially identified under each scenario, and derive an econometric specification that encompasses all scenarios.

2.1 Benchmark Model with Perfect Information

Consider a regulatory regime of three parties – the principal (Florida Division of Hotels and Restaurants, DHR hereafter), inspectors (government employees), and clients (restaurants). The principal defines inspection criteria, inspection technology, inspector assignment, and inspector compensation. Each inspector earns a fixed wage. Assuming there are two categories of violations (e.g. critical and non-critical), the principal imposes a fine

⁴Evidence of inspector heterogeneity has been documented in nuclear plant inspections (Feinstein 1989), IRS tax auditing (Feinstein 1991), and FDA inspections of pharmaceutical manufacturing (Macher et al. 2011). See Dranove and Jin (2010) for a review of literature on the economics of certifiers.

structure $F(y) = \tau_1 y_1 + \tau_2 y_2$ where y_1 and y_2 denote the number of violations detected and τ_1 and τ_2 denote penalty rates on the two categories, reflecting the principal’s preferences.

The main task of an inspector is to visit a restaurant at an unannounced time, detect all the hygiene violations, and report them to the principal. Within the restaurant, the inspector has discretion as to how much effort to exert in detecting violations and how much information to report. In the eyes of the principal, hiding detected violations is equivalent to shirking on detection effort, so we do not distinguish between the two in the model.⁵ Rather, we consider every inspector honest and assume the cost of detection effort for inspection i is $C(e_i) = \theta_i e_i^2$, where θ_i is the parameter of detection cost.

Not only do inspectors differ in detection cost, but each inspector may also have her own interpretation of the regulation. Given the two categories of violations, we assume inspector i puts weight α_i on category 1 and weight $(1 - \alpha_i)$ on category 2. Accordingly, the inspector’s detection efforts are e_{i1} and e_{i2} for the two categories. Assumed to be between 0 and 1, e_{i1} and e_{i2} can be interpreted as the probability of detection for category 1 and 2. If true violations are $\tilde{y}_1$ and $\tilde{y}_2$, detected violations are $y_1 = \tilde{y}_1 e_{i1}$ and $y_2 = \tilde{y}_2 e_{i2}$. Here we do not allow inspectors to report non-existent violations (extortion) because an appeal procedure allows restaurants to contest any reported violation in Florida. Moreover, the expected fine is very low (\$14 per inspection) and the fine amount is not determined by the inspector.

For tractability, we assume effort costs on categories 1 and 2 are independent and both depend on the same cost parameter θ_i . In other words, θ_i denotes the overall stringency of i . If θ_i differs by category, it is observationally equivalent to the inspector putting different weights on different items. Empirically, the assumption of common θ_i across categories can be relaxed as we observe many inspectors and regulatory categories and category-specific cost of effort can be controlled for by inspector-category fixed effects.

The goal of regulation is enforcing food safety, which implies minimization of actual violations. Since we focus on the interaction between inspector and restaurant, we do not model the principal-inspector relationship explicitly. Rather, we assume the inspector, as an agent of the principal, derives negative utility from both detected and undetected violations. Because undetected violations may be ignored by the restaurant and pose a

⁵The incentive to hide perfectly-observable violations was the focus of many theories on inspector-firm collusion.

bigger public health risk, we assume the inspector is more concerned about undetected violations. In other words, the DHR and its employees would like to see zero violations if all violations could be detected. However, given the existence of violations, identifying them is better than leaving them undetected. In the model, we introduce $\lambda > 1$ as the disutility of an undetected violation relative to a detected violation. If $\lambda < 1$, the inspector will always prefer minimal effort and detect no violations. Note that λ reflects the inspector's preference, which may or may not coincide with that of the principal. In short, the inspector trades off her own preference for inspection outcomes versus her effort cost. This captures the fact that government inspectors are paid by fixed salary and their efforts are likely more motivated by intrinsic preferences than by monetary returns (Prendergast 2007).

For the restaurant, the benefits from cleaning up include reduced fines for detected violations and the reduced risk of bad publicity owing to foodborne illness outbreaks.⁶ To minimize both, the restaurant can exert efforts e_{r1} on category 1 and e_{r2} on category 2. Normalizing maximum violation (per category) as 1, we have the actual violations $\tilde{y}_1 = 1 - e_{r1}$ and $\tilde{y}_2 = 1 - e_{r2}$. Consequently, the detected violations are $y_1 = \tilde{y}_1 e_{i1} = (1 - e_{r1})e_{i1}$ and $y_2 = \tilde{y}_2 e_{i2} = (1 - e_{r2})e_{i2}$. For simplicity, we assume that the risk of bad publicity is a linear function of actual violations ($R_1 \cdot \tilde{y}_1 + R_2 \cdot \tilde{y}_2$), where R_1 and R_2 can be interpreted as the marginal expected penalty or reputational cost that consumers impose on restaurants with actual violations.

Assuming the cost of restaurant effort is strictly convex ($C(e_r) = \theta_r e_r^2$) and applies to both items independently, we can write the restaurant's problem as:

$$\min_{e_{r1}, e_{r2}} W_r = \tau_1(1 - e_{r1})e_{i1} + \tau_2(1 - e_{r2})e_{i2} + R_1(1 - e_{r1}) + R_2(1 - e_{r2}) + \theta_r e_{r1}^2 + \theta_r e_{r2}^2.$$

The inspector's problem can be written as:

$$\begin{aligned} \min_{e_{i1}, e_{i2}} W_i = & \alpha_i((1 - e_{r1})e_{i1} + \lambda(1 - e_{r1})(1 - e_{i1})) \\ & + (1 - \alpha_i)((1 - e_{r2})e_{i2} + \lambda(1 - e_{r2})(1 - e_{i2})) + \theta_i e_{i1}^2 + \theta_i e_{i2}^2. \end{aligned}$$

The timing of the game is as follows: at stage 0, the principal sets inspection criteria,

⁶Jin and Leslie (2003) show that restaurant revenue was insensitive to restaurant inspection outcomes before the introduction of restaurant hygiene report cards. As of 2011, Florida has no restaurant hygiene report card though inspection outcomes have been posted online since 2009. This change will be controlled for by year-month fixed effects throughout Florida.

inspector assignment, fine structure and inspector compensation. At stage 1, the restaurant chooses e_{r1} and e_{r2} . At stage 2, the inspector walks in and chooses detection effort e_{i1} and e_{i2} . At the end of stage 2, detected violations (y_1 and y_2) are reported to the principal.

In this subsection, to build up a benchmark model, we assume that all the cost parameters (θ_i, θ_r) and the inspector's emphasis on regulation (α_i) are common knowledge. However, the inspector does not know the restaurant's effort and the restaurant does not know the inspector's effort. Since no new information is generated between stages 1 and 2, the inspector-restaurant game is treated as a simultaneous game.

Figure 1 shows two reaction curves: the restaurant is more willing to clean up if it knows that the inspector exerts more effort (the restaurant's compliance curve), but the inspector will exert less effort if she knows that the restaurant has cleaned up (the inspector's detection curve). As the two curves intersect, we have a unique inner solution in equilibrium if $\theta_r > \frac{\max(R_1, R_2)}{2}$, $1 - \frac{4\theta_i\theta_r}{(2\theta_r - \tau_2 - R_2)(\lambda - 1)} < \alpha_i < \frac{4\theta_i\theta_r}{(2\theta_r - \tau_1 - R_1)(\lambda - 1)}$ ⁷:

$$\begin{aligned} e_{i1} &= \frac{(2\theta_r - R_1)\alpha_i(\lambda - 1)}{4\theta_i\theta_r + \tau_1\alpha_i(\lambda - 1)} & e_{i2} &= \frac{(2\theta_r - R_2)(1 - \alpha_i)(\lambda - 1)}{4\theta_i\theta_r + \tau_2(1 - \alpha_i)(\lambda - 1)} \\ e_{r1} &= \frac{\tau_1\alpha_i(\lambda - 1) + 2\theta_iR_1}{4\theta_i\theta_r + \tau_1\alpha_i(\lambda - 1)} & e_{r2} &= \frac{\tau_2(1 - \alpha_i)(\lambda - 1) + 2\theta_iR_2}{4\theta_i\theta_r + \tau_2(1 - \alpha_i)(\lambda - 1)}. \end{aligned}$$

Therefore, the equilibrium reported violations are as follows:

$$\begin{aligned} y_1 &= (1 - e_{r1})e_{i1} = \frac{2\theta_i\alpha_i(\lambda - 1)(2\theta_r - R_1)^2}{[4\theta_i\theta_r + \tau_1\alpha_i(\lambda - 1)]^2} \\ y_2 &= (1 - e_{r2})e_{i2} = \frac{2\theta_i(1 - \alpha_i)(\lambda - 1)(2\theta_r - R_2)^2}{[4\theta_i\theta_r + \tau_2(1 - \alpha_i)(\lambda - 1)]^2}. \end{aligned}$$

It is not difficult to show that (1) greater fines (τ) increase restaurant clean-up effort and decrease inspector detection effort; (2) restaurant clean-up cost (θ_r) decreases restaurant clean-up effort and increases inspector detection effort; (3) inspector detection cost (θ_i) decreases both inspector and restaurant efforts; and (4) inspector emphasis on category 1 (α_i) increases inspector and restaurant efforts in category 1 but decreases both efforts in category 2.

We face two fundamental identification problems if we want to use this framework to empirically identify detection from compliance: first, we observe only the intersection of

⁷These conditions imply that (1) the cost of restaurant effort must be sufficiently high so that the threat of bad publicity alone is not enough to guarantee maximum clean-up effort, and (2) the inspector puts relatively balanced weights on the two categories so that it is meaningful for the inspector to exert some but lower-than-maximum detection efforts to detect violations in both categories.

the two reaction curves. Interestingly, this problem resembles the typical identification problem in the supply and demand model where the difficulty can be resolved by using exogenous demand (supply) shifters to trace out the supply (demand) curve. However, identification is even harder in the inspection game because we only observe the product of non-compliance and detection ($\tilde{y}_j e_{ij}$), not the two separately. In other words, inspector heterogeneity (which shifts the detection curve) and restaurant heterogeneity (which shifts the compliance curve) cannot identify the two reaction curves. Similarly, exogenous policies that shift the inspector's detection curve or shift the restaurant's compliance curve cannot fully identify the two curves. Identifying compliance and detection separately is critical for conducting policy simulations.

Second, in the literature, researchers often regress detected violations on inspector fixed effects and interpret these fixed effects as inspector heterogeneity.⁸ Under the assumption of perfect information, our theory suggests that inspector fixed effects reflect not only inspector heterogeneity in overall stringency and taste, but also the differential compliance that restaurants adopt in response to inspector heterogeneity.

2.2 Extended model with uncertainty

Assuming perfect information, the benchmark model ignores two institutional features of a typical inspection program: first, regular inspections are designed to be random so that the restaurant may not perfectly predict the identity of the next inspector; second, the restaurant must choose the compliance effort before the inspector arrives. This implies that a restaurant's compliance is based on the restaurant's expectation instead of the actual identity of the next inspector.

These features imply that a risk-neutral restaurant will choose the optimal compliance effort in response to the *expected* inspector effort, $\tilde{e}_{i1}$ and $\tilde{e}_{i2}$ based on expected inspector taste ($\tilde{\alpha}_i$) and expected inspector effort cost ($\tilde{\theta}_i$).⁹ On the other hand, the inspector, knowing her own α_i and θ_i , will apply her best response to e_r as before. Mathematically, we have:

⁸See Feinstein (1989) and Macher et al. (2011) for example.

⁹For example, the restaurant may form an expectation about the next inspector's type in a naïve way: $\tilde{\alpha}_i = \alpha_{i-1} + u_i$ and $\tilde{\theta}_i = \theta_{i-1} + v_i$ where α_{i-1} and θ_{i-1} are the previous inspector's taste and effort cost parameter.

$$\begin{aligned}
e_{r1} &= \frac{\tau_1 \tilde{e}_{i1} + R_1}{2\theta_r} & e_{r2} &= \frac{\tau_2 \tilde{e}_{i2} + R_2}{2\theta_r} \\
e_{i1} &= \frac{\alpha_i(\lambda-1)(1-e_{r1})}{2\tilde{\theta}_i} & e_{i2} &= \frac{(1-\alpha_i)(\lambda-1)(1-e_{r2})}{2\tilde{\theta}_i}
\end{aligned}$$

Therefore, the equilibrium efforts are:

$$\begin{aligned}
e_{i1} &= \frac{(2\theta_r - R_1)\tilde{\alpha}_i(\lambda-1)}{4\tilde{\theta}_i\theta_r + \tau_1\tilde{\alpha}_i(\lambda-1)} & e_{i2} &= \frac{(2\theta_r - R_2)(1-\tilde{\alpha}_i)(\lambda-1)}{4\tilde{\theta}_i\theta_r + \tau_2(1-\tilde{\alpha}_i)(\lambda-1)} \\
e_{r1} &= \frac{\tau_1\tilde{\alpha}_i(\lambda-1) + 2\tilde{\theta}_i R_1}{4\tilde{\theta}_i\theta_r + \tau_1\tilde{\alpha}_i(\lambda-1)} & e_{r2} &= \frac{\tau_2(1-\tilde{\alpha}_i)(\lambda-1) + 2\tilde{\theta}_i R_2}{4\tilde{\theta}_i\theta_r + \tau_2(1-\tilde{\alpha}_i)(\lambda-1)}.
\end{aligned}$$

Comparing these solutions to those of perfect information, it is obvious that the difference is driven by the gap between an inspector's actual attributes (α_i and θ_i) and the restaurant's expectation ($\tilde{\alpha}_i$ and $\tilde{\theta}_i$). As shown in Figure 2, if the restaurant over-estimates the inspector's effort cost (i.e. under-estimates the inspector's overall stringency, $\tilde{\theta}_i < \theta_i$), the restaurant's expected inspector detection curve lies below the actual detection curve. Because the restaurant must choose the compliance effort before knowing the inspector's identity, its best choice is e_r^A . Knowing that the restaurant will choose e_r^A , the inspector will choose e_i^B according to her own stringency. In contrast, when the inspector is fully expected, the equilibrium point should be C.

Applying this framework to data entails several complications: the first is that inspectors may differ in inherent taste (α_i) and effort cost (θ_i), which by definition apply to all inspections that i conducts; second, on top of these inherent characteristics, an inspector's taste and effort cost may also vary by her relationship with the restaurant – the length of the relationship is zero if the inspector is new to the restaurant and it increases as the inspector returns repeatedly. In particular, the impact of the inspector-restaurant relationship on inspection outcomes depends on restaurant expectation of the next inspector. Below we discuss how these complications can be utilized to partially identify detection and compliance.

Inspector Heterogeneity

Consider Figure 2 as a scenario where the restaurant and an old inspector engage in a perfect-information equilibrium at point A. Suppose that in the next inspection, the restaurant expects the old inspector to return but a new inspector with different θ and α shows up instead. If the new inspector is more stringent (lower θ) or places more emphasis on category one (higher α), the outcome corresponds to point B. In the third inspection, the restaurant correctly expects the new inspector to return, which leads to another perfect-information equilibrium at point C.

Because e_r^A is chosen by the restaurant's (*ex post* incorrect) expectation, the move from A to B is driven by the new inspector being inherently different from the old one. The move from B to C is a mixture of changes in both detection and compliance. After learning that the new inspector is more stringent (or pays more attention to category one), the restaurant will devote more effort to cleaning up; now that the restaurant has cleaned up more, the optimal detection effort will decline. It is tempting to argue that we can use the first arrival of the new inspector to identify changes in detection and the subsequent visit of this new inspector to identify an upper bound of compliance. However, this argument is incorrect because we have to identify each inspector's θ and α as well.

To fix our argument, suppose we have two identical restaurants Ra and Rc . By time 0, Ra has been inspected by inspector Adam and (incorrectly) expects Adam to come back; Rc has been inspected by inspector Calvin and expects Calvin to come back. Both have played their perfect-information equilibrium before time 0, so we denote their detected violations before time 0 as y_{Ra0} and y_{Rc0} . Graphically, these two numbers correspond to the detected violations at points A and C in Figure 2 ($y_{Ra0} = y^A$ and $y_{Rc0} = y^C$). Due to some exogenous reason (say Adam has suddenly retired), Calvin is assigned to inspect both restaurants from time 1 and on. For restaurant Rc , the equilibrium remains at C ($y_{Rc1} = y_{Rc0} = y^C$). For restaurant Ra , the surprising realization moves the equilibrium from A to B, so $y_{Ra1} = y^B$. At time 2, Calvin is expected to return to both restaurants and therefore both settle at C ($y_{Ra2} = y_{Rb2} = y^C$).

In the data, for restaurant r , inspector i at time t , we can run the regression:¹⁰

$$\log(y_{irt}) = \mu_i + \beta_{hetero} \cdot (\mu_i - \mu_{i-1})$$

where μ_i is the inspector fixed effect for the current inspector and μ_{i-1} is the inspector fixed effect for the previous inspector that inspected restaurant r in the last period. It can be shown that $\mu_{Adam} = \log(y^A)$, $\mu_{Calvin} = \log(y^C)$ and $\beta_{hetero} = \frac{\log(y^B) - \log(y^C)}{\log(y^C) - \log(y^A)}$. In other words, inspector fixed effects denote how different inspectors report different numbers of violations when they are fully expected¹¹, while β_{hetero} denotes the extra violations that Calvin will report after his first visit to restaurant Ra , as a percentage in addition to how

¹⁰Note that there is no error term in this equation because we have not allowed random factors in the theory. Empirically, in the presence of random factors, $\log(y_{irt})$ is replaced by $\log E(y_{irt})$.

¹¹As argued in Section 2.1, heterogeneity in inspector fixed effects reflects both detection heterogeneity and the corresponding differences in compliance.

the detected violations would have been changed should restaurant *Ra* know the inspector shift. This extra amount comes from *Ra* catering to Adam and will disappear when *Ra* adjusts its compliance effort to Calvin next time. In this sense, it represents an upper bound of the compliance response to inspector change from Adam to Calvin. The same logic can be adapted to a variation of Figure 2 if Calvin is less stringent or puts less emphasis on category one than Adam.

The Fresh-eye Effect of New Inspector

Now we turn to the case where all inspectors are inherently identical but a new inspector has a pair of “fresh eyes” in his first visit to a restaurant. As shown later, 27% of regular inspections in our final sample are conducted by a new inspector who did not visit the restaurant before. Moreover, within each restaurant, the probability of a new inspector coming in the next regular inspection increases almost linearly with the repetitiveness of the last inspector (labeled as *Lrepeat*) after the last inspector has visited the restaurant twice. This suggests that a restaurant may have anticipated a higher probability of encountering a new inspector as it develops a relationship with the current inspector.

Since the restaurant’s expectation on the probability of new inspector is different from the actual realization of whether the inspector is new or repeat, we can utilize this difference to partially identify detection and compliance. If a new inspector has fresher eyes and therefore is more stringent than repeat inspectors (the opposite can be shown symmetrically), Figure 3 shows four reaction functions of e_i with respect to e_r : the lowest one is for the repeat inspector, the highest one is for the new inspector, and the middle two capture the fact that the probability of a new inspector increases over time as *Lrepeat* increases from *Lrepeat*₁ to *Lrepeat*₂. Under *Lrepeat*₁, if the inspector turns out to be repeat, we observe y^A ; if the inspector turns out to be new, we observe y^B . Similarly under *Lrepeat*₂, we observe y^C when the inspector is repeat and y^D when the inspector is new.

If we fit points A, B, C and D into a regression framework:

$$\log(y) = \alpha + \beta_{new} \cdot NEW + \beta_{Lrpt} \cdot Lrepeat,$$

it is easy to show that $\beta_{new} = \frac{\log(y^B) - \log(y^A) + \log(y^D) - \log(y^C)}{2} = \frac{\log(e_i^B) - \log(e_i^A) + \log(e_i^D) - \log(e_i^C)}{2}$ and $-\beta_{Lrpt} = \frac{\log(y^A) - \log(y^C) + \log(y^B) - \log(y^D)}{2(Lrepeat_2 - Lrepeat_1)} > \frac{\log(1 - e_r^A) - \log(1 - e_r^C)}{(Lrepeat_2 - Lrepeat_1)}$. In other words, we can identify β_{new} as an average detection difference between new and repeat inspectors (β_{new});

and $-\beta_{Lrpt}$ as an upper bound of restaurant compliance in response to a growing fear of a new inspector as $Lrepeat$ increases.

Repeat Inspector Slacks Over Time

The fresh-eye effect of a new inspector explains why the same (repeat) inspector may report fewer violations over time: because the restaurant has a greater fear of a new inspector (who is more stringent due to the fresh-eye effect), the restaurant chooses to comply more. However, this is based on the assumption that a repeat inspector always has the same stringency and taste regardless of the length of her relationship with the restaurant. If a repeat inspector slacks over time (i.e. has higher θ thus become less stringent over time), he may report fewer violations over time but for a reason completely different from the fresh-eye effect of a new inspector.

For the sake of illustration, let us assume every inspector to be identical except that the degree of fresh eyes declines over time as the restaurant-inspector relationship is prolonged. This implies that, as shown in Figure 4, the detection curve of a new inspector lies above that of repeat inspector and this curve drops further as $Lrepeat$ increases from $Lrepeat_1$ to $Lrepeat_2$. For now suppose there is a constant probability of encountering a new inspector next time. Increased slackness of a repeat inspector over time implies that the restaurant's expected detection curve will drop as well from $Lrepeat_1$ to $Lrepeat_2$. As shown in Figure 4, this leads to less compliance by the restaurant. If the increased slack of detection dominates the decreased compliance, we may observe fewer violations over time within a restaurant-inspector relationship.

Econometric Specification

Above all, we offer three explanations as to why the detected violations may differ between new and repeat inspectors and why they may decline over time within the same restaurant-inspector pair: one reason is that the restaurant learns the inspector's taste and tailors its compliance effort accordingly; the second possibility is that repeat inspector slacks over time thus encourages less compliance; in the third explanation, new inspectors are known to have fresher eyes and the restaurant expects a greater probability of a new inspector as its relationship with the current inspector continues.

We can separate the first explanation from the rest because we observe multiple restaurants inspected by the same inspector, which allows us to identify each inspector's time-invariant stringency (θ_i) and categorical taste (α_i). As the restaurant caters to the taste

of the last inspector, the extra violations detected by a new inspector should increase with the stringency/taste difference between the new and repeat inspectors.

The other two explanations are not inspector-specific. They can be separated by comparing the effect of a new inspector coming at different phases of a restaurant-inspector relationship. If the main force is restaurant fear of a new inspector coming, Figure 3 shows that the increased detection due to the fresh-eye effect of a new inspector should be smaller for a longer relationship (i.e. $e_i^D - e_i^C < e_i^B - e_i^A$) because the restaurant should expect a higher probability of new inspector after a longer relationship with the current inspector and complies accordingly. In contrast, if the main force is the repeat inspector slacking over time, Figure 4 shows that restaurant compliance will be lower ($e_{r2} < e_{r1}$) for a longer relationship and the detection increase driven by the fresh-eye effect of the new inspector will be greater ($e_i^D - e_i^C > e_i^B - e_i^A$). Both lead to more extra violations being cited by the new inspector. This prediction is intuitive: a new inspector is a greater contrast to a slack-ing repeat inspector and therefore should uncover more problems after a repeat inspector has become more and more slack in her relationship with the restaurant.

Econometrically, let us denote y_{irct} as the number of observed violations in category c for restaurant r by inspector i at time t . Assuming y_{irct} follows a Poisson distribution with mean λ_{irct} , we propose the following specification:

$$\begin{aligned}
\log(\lambda_{irct}) &= \log(E(y_{irct})) \\
&= \underbrace{\beta_{new} \cdot NEW_{irt} + \beta_{Lrpt} \cdot Lrepeat_{i-1rt} + \beta_{newLrpt} \cdot NEW_{irt} \times Lrepeat_{i-1rt}}_{\text{New inspector's fresh eye effect or repeat inspector's shirking effect}} \\
&\quad + \underbrace{\mu_{ic} + \beta_{hetero}(\mu_{ic} - \mu_{i-1c})}_{\text{Inspector heterogeneity}} + \mu_{rc} + \mu_{ymc} + \zeta X_{rit}
\end{aligned}$$

where

- NEW_{irt} = A dummy equal to one if inspector i is new to r at t ;
- $Lrepeat_{i-1rt}$ = # of times the last inspector i_{-1} has visited r by t ;
- μ_{rc} = Restaurant-category fixed effects;
- μ_{ic} = Inspector-category fixed effects for current inspector i ;
- μ_{ymc} = Year-month-category fixed effects;
- μ_{i-1c} = Inspector-category fixed effects for the last inspector i_{-1} ;
- X_{irt} = Other restaurant-inspector observables such as restaurant age, inspector tenure, whether the inspection is during lunch hours, how many days since the last inspection of restaurant r , etc.

This specification includes a rich set of fixed effects. Restaurant-category fixed effects μ_{rc} capture restaurant r 's time-invariant difficulty (or willingness) to clean up category c . Inspector-category fixed effects μ_{ic} capture inspector i 's specific detection cost and relative taste in category c , and the corresponding compliance if the restaurant can perfectly expect that i is coming. These inspector fixed effects are identified from the average violations reported by the same inspector throughout all of her repeat inspections of restaurants. Any category-specific effort cost or taste change applicable to all inspectors and all restaurants is absorbed in year-month-category fixed effects μ_{ymc} .

All the key coefficients (β s) can be category-specific as well. Below we ignore category subscript for simplicity. The coefficient of β_{hetero} captures the extent to which inspection outcomes depend on the inherent stringency and taste heterogeneity between the current and previous inspectors. Note that β_{hetero} is not identified unless the current inspector is new and differs from the previous inspector. According to our theory, this coefficient should be zero if restaurants do not adjust their cleaning effort to suit the last inspector's stringency and taste. Since we control for the inspector-category fixed effect of the current inspector, the coefficient captures the extent to which a restaurant complies with the idiosyncratic stringency and taste of the last inspector rather than those of the current inspector.

In comparison, β_{new} captures how new and repeat inspectors report violations dif-

ferently, even if they have exactly the same intrinsic stringency and taste specific to a category. More specifically, β_{new} pools all the possibilities that a new inspector might be different from the previous inspector, for example, she might be less familiar with the kitchen, equipped with fresher eyes, or less willing to hide detected violations. We call the sum of all these possibilities the fresh-eye effect of new inspectors. β_{new} and β_{hetero} can be identified separately because the former compares new versus repeat inspectors no matter who the last inspector was, while the latter explores the identity and idiosyncratic taste/stringency of the last inspector.

As stated above, the effect of $Lrepeat$ is ambiguous: On the one hand, if a repeat inspector slacks over time, she encourages less compliance as she develops a relationship with a restaurant; on the other hand, if a longer relation with the last inspector implies a growing expectation of a new inspector’s arrival, the restaurant may comply more (or less) over time, depending on whether the new inspector is expected to be more (or less) stringent than the last inspector. In other words, $\beta_{Lrepeat}$ is a mixture of compliance, slack detection of a repeat inspector, and the restaurant’s expectation as to the stringency and taste of a new inspector. The coefficient on the interaction of NEW and $Lrepeat$ helps to distinguish these possibilities. If the effect from the gradual slacking of the repeat inspector dominates the effect from the growing fear of a new inspector who should be more stringent due to the fresh-eye effect, we predict $\beta_{newLrpt}$ to be positive.

Our model assumes that restaurants cannot perfectly predict the identity of the next inspector and therefore the time and information lag between restaurant choice of compliance and inspector choice of detection can be utilized to achieve partial identification. What if the expectation is perfect, for example, the previous inspector tells the restaurant in advance? This implies that we observe only points A and C in Figure 2 and their difference is driven by inspector variation as well as different compliance in response to inspector variation. In this setting, we can no longer (partially) separate detection from compliance, but we can still examine whether the difference between A and C is due to inspector heterogeneity or inspector-restaurant relationship, because the latter is restaurant specific while the former is not. After controlling for inspector heterogeneity, if we observe more violations reported at C (new inspector) than at A (repeat inspector), the extra violations can be interpreted as a lower bound of the detection difference between new and repeat inspectors due to the fresh-eye effect.

Another empirical implication from perfect foresight is that the restaurant will never adjust compliance efforts to suit the last inspector’s preference and then get caught by the next inspector with different preferences. This implies $\beta_{hetero} = 0$. Moreover, perfect foresight suggests that the restaurant will clean up according to the next inspector’s preference and her relationship with the restaurant. If this next inspector is new, restaurant effort will not depend on its relationship with the last inspector, which implies $\beta_{newLrpt} = 0$. In sum, statistical tests on $\beta_{hetero} = 0$ and $\beta_{newLrpt} = 0$ will help us infer whether restaurants can perfectly predict the identity of the next inspector.

To summarize:

- The distribution of inspector fixed effects (μ_{ic}) captures inspector heterogeneity in their inherent stringency and taste (as well as the corresponding compliance response if restaurants can predict perfectly who will come next time).
- We predict $\beta_{hetero} > 0$ if restaurants cannot perfectly predict the identity of the next inspector. Under that assumption, β_{hetero} represents an upper bound of compliance in response to inspector heterogeneity.
- We predict $\beta_{new} > 0$ if new inspectors have fresher eyes or are more reluctant to ignore observed violations. In contrast, we may have $\beta_{new} < 0$ if new inspectors are less familiar with the restaurant and thus incur a higher detection cost.
- We predict an ambiguous sign of β_{Lrpt} but if $\beta_{new} > 0$ we predict $\beta_{newLrpt} > 0$ if restaurants cannot perfectly predict the arrival of the new inspector and the increased slack of repeat inspector dominates the growing fear of a new inspector as *Lrepeat* increases.

Several econometric issues arise when we implement the above specification. First, Florida DHR classifies violations into 55 categories. To save space, we report the full-model estimation for three groups of categories separately, namely critical, non-critical and risk-factor violations. As shown below, results are qualitatively similar if we run the model for four frequently cited violation categories separately.

Second, following Hausman, Hall and Griliches (HHG 1984), we estimate the above model with maximum likelihood conditional on the sum of violations per restaurant throughout the whole sample period. This allows us to estimate the other coefficients without

estimating restaurant fixed effects for each run. While HHG (1984) derived the likelihood function from a Poisson model of count data, which assumes that the mean of the Poisson distribution is equal to its variance (so called equal-dispersion), Woodridge (1999) shows that the conditional maximum likelihood function proposed by HHG (1984) is robust to other distribution assumption and does not require equal-dispersion so long as the standard errors are adjusted. Therefore, for category c , we follow HHG (1984) to construct the conditional likelihood as shown below but compute robust standard errors according to Woodridge (1999):

$$\log(L_c) = \sum_r \sum_t \left[y_{irct} \log(\lambda_{irct}) - y_{rct} \log\left(\sum_s \lambda_{ircs}\right) \right].$$

The likelihood function does not sum over i because every t corresponds to an i in the raw data and therefore summing over t is implicitly summing over i . It is worth noting that although we circumvent the estimation of restaurant fixed effects, we will estimate inspector fixed effects explicitly for each group of categories. This enables us to quantify the nature of inspector heterogeneity. The heterogeneity that is common to all three groups of categories indicates how inspectors differ in their overall stringency, while the heterogeneity that varies across the three groups tells us how inspectors differ in their relative emphasis on critical, non-critical and risk-factor categories.

3 Data Summary

3.1 Data Description

Our sample is constructed from three administrative data sets collected by the DHR: (1) restaurant/food service inspection files, (2) license files, and (3) restaurant disciplinary activity reports. The data include all restaurant inspections in Florida from July 2003¹² to March 2010. There are two types of inspections: the first type is regular inspections conducted at an unannounced time, which Florida officials refer to as initial inspections. Depending on the results of a regular inspection, a callback inspection may follow to ensure compliance for a small proportion of restaurants. In the raw data, 81% are regular inspections and 19% are callbacks.

¹²July 2003 is the start of the 2003-2004 fiscal year.

All food establishments are required to be inspected twice per fiscal year by state laws and three times by administrative rules. However, due to labor shortage, the average number of regular inspections per restaurant per year is fewer than 2 except for the 2008-2009 fiscal year. About 20-40% of restaurants receive only one regular inspection a year.¹³ In this paper, we focus on regular inspections only.

The data on regular inspections are cleaned in several steps. Starting with 600,492 regular inspections in the raw data, we exclude the first six months of a restaurant since its first appearance in our data because we use these months to define history. If the first six months do not cover the restaurant’s first regular inspection in our data, we exclude the restaurant’s history up to its first regular inspection. We also exclude any inspection conducted before March 2004 because Florida reclassified some non-critical violations as risk-factors starting March 2004.¹⁴ As we apply restaurant fixed effects in all estimations, we also exclude the 11,819 restaurants that have only one inspection throughout the sample (1.97% of regular inspections).

The final sample includes 426,831 regular inspections, covering 60,976 unique restaurants and 358 individual inspectors.¹⁵ Each year there are around 220 active inspectors. Each inspector conducts, on average, more than 200 inspections per year although this number varies greatly across inspectors. Since there are 55 violation categories, the sample includes 23,475,705 category-by-inspection observations. Out of the 55 categories, 18 are critical violations, 26 are non-critical, and 11 are risk factors (See Appendix table A.1 for the exact content of each category). As shown in Table 1, an average inspection finds 7.9 violations, of which 1.62 are critical, 2.54 are risk factors, and 3.75 are non-critical. About 96% of regular inspections are routine ones, while 3.7% are initiated by complaints and 0.1% are licensing inspections.

According to our theory, the empirical identification of detection and compliance will explore several variations in inspector identity and restaurant-inspector relationship. Below we summarize these variations separately.

¹³The average number of regular inspections is 1.66 in FY 2003-04; 1.93 in FY 2004-2005; 1.67 in FY 2005-2006; 1.72 in FY 2006-2007; 1.85 in FY 2007-2008; 2.14 in FY 2008-2009. The proportion of restaurants that receive only one inspection is 50.6%; 22.4%; 39.9%; 26.2%; 15.2%, respectively.

¹⁴This reclassification requires inspectors to pay more attention to risk factors, although these factors are not denoted as critical violations on the inspection form.

¹⁵The original inspection files include 386 inspectors and 97,990 restaurants.

3.2 Summary of Inspector Heterogeneity

A crucial assumption underlying inspector heterogeneity is that inspectors differ significantly in stringency and taste. To get a sense of this, we regress the total number of violations per inspection on a full set of inspector dummies, controlling for fiscal year, month, and restaurant fixed effects. Such a regression yields an adjusted R-square of 0.514, which is higher than the adjusted R-square without inspector fixed effects (0.457). Based on this regression, Figure 5 plots the estimated inspector fixed effects. The range of these inspector fixed effects is huge, given the fact that the average number of violations is 7.9 with a standard deviation of 7.16. This finding is consistent with the findings that have been documented for nuclear inspectors, tax auditors, and pharmaceutical plant inspectors (Feinstein 1989, 1991; Macher et al. 2011).

To shed light on inspectors' category-specific tastes, we repeat the exercise at the category level with inspector-category fixed effects. Within each inspector, we code the category with the largest inspector-category fixed effect as the inspector's favorite category.¹⁶ Figure 6 plots the histogram of favorite categories across all 358 categories. This picture is dispersed, with relatively high frequencies in certain categories (notably, 2, 8, 14, 22, 32, 37 and 45). Appendix Table A1 presents the complete list of all 55 categories and highlights these high-frequency favorite categories.

3.3 Summary of New and Repeat Inspectors

We define an inspector i as “new” to a restaurant r if the observed inspection is the first inspection conducted by i at r during the whole sample period. It is possible that the new inspector actually inspected r prior to the sample period because our data are left-censored. To address this issue, we restrict the analysis sample to the inspections conducted at least 6 months after the restaurant's first appearance in our data. The sample also excludes the first regular inspection of a restaurant if it occurred after the first six-month history of that restaurant. As shown in Table 1, about 27% of regular inspections are conducted by new inspectors, and a restaurant has been inspected by the previous inspector on average 3.65 times.

By regulation, all inspectors are subject to standard training.¹⁷ Despite the univer-

¹⁶Before determining which category is an inspector's favorite, we test the statistical significance of each inspector-category fixed effect and exclude all that are insignificant from zero by 95% confidence.

¹⁷A newly hired inspector should receive at least 120 hours of training in her first year of employment.

sal training requirement, the first panel of Table 2 reveals significant difference between new and repeat inspectors. Pooling all categories, new inspectors on average find 9.36 violations, almost two additional violations (or 27% higher) than repeat inspectors. Such differences are found in all three types of violations. The rest of Table 2 further presents the new-repeat difference by the number of times that the previous inspector has inspected the restaurant ($Lrepeat$). The new-repeat difference remains highly significant, and this difference increases slightly with $Lrepeat$.

Figure 7 shows the average number of violations during the typical new-repeat history of a restaurant. It starts with the first new inspector for a restaurant, followed by the first, second, third, and fourth repetition of this inspector and then the next round for the next new inspector. Because not every restaurant has such a regular pattern, the subsample corresponding to each point of Figure 7 may differ. That being said, Figure 7 shows a striking pattern: a new inspector reports more violations; as the restaurant-inspector relationship is prolonged, the number of reported violations declines until the next new inspector.

3.4 Control Variables

We control for a number of dynamic factors that may affect restaurant or inspector efforts. One factor is restaurant age; another is inspector tenure. With no access to the full employment record, we proxy inspector tenure by the number of regular inspections that an inspector has conducted in our data before a specific inspection. As shown in Table 1, the average inspector tenure is 1,535. In addition, 44% of inspections are conducted by inspectors with less than median tenure (“inexperienced”), and 1% are conducted by inspectors with a tenure less than 30 inspections. A third dynamic factor is the number of inspections that the inspector has done in a day. An inspector may become tired during the day and incur higher effort costs due to fatigue. On average, an inspector has completed 1.6 inspections before coming to the inspection under study and 28% of inspections are the first one conducted by that inspector in that day. Table 1 also reports that the average time span between this and the last inspection is 184 days, and 38% of the inspections occur during lunch time (12-2pm). These two variables may affect inspection outcomes because

Also existing inspection staff receive a minimum of 20 hours of training each year. Each inspector is checked by the FDA every three years to ensure compliance with national standards. Each inspector is required to pass a certified food manager examination every five years.

restaurants are likely to adjust cleaning efforts according to when the next inspection is expected to occur and most restaurants are busy at lunch time and probably pay less attention to food safety.

3.5 OLS Results and Inspector Assignment

Table 3 reports the OLS results, where the dependent variable is the number of detected violations and the key right hand variables are *New*, *Lrepeat*, and *New · Lrepeat*. We run the ordinary least square regression for critical, non-critical, and risk-factor violations separately. The control variables mentioned above are all included, in addition to restaurant fixed effects, year-month fixed effects, and inspector fixed effects. Consistent with the raw data summary, Table 3 suggests that inspections by new inspectors report more violations and this pattern is more conspicuous if the inspected restaurant has had a longer relationship with its last inspector. Also, within repeat inspections, the number of reported violations declines over time.

Our theory assumes that restaurants cannot predict for sure whether the next inspector will be new or not. This assumption could be violated if the DHR sticks to a predetermined inspector rotation schedule. To check this, Figure 8 plots the likelihood of a new inspector’s arrival as a function of the number of inspections that the last inspector has made for this restaurant (*Lrepeat*), after controlling for year, month and restaurant fixed effects. The curve increases steadily after *Lrepeat* = 3, suggesting that there is no obvious rotation in inspector assignment, although the DHR seems reluctant to change inspectors twice consecutively within a restaurant.

What factors drive the DHR to send a new inspector to a restaurant rather than using the previous inspector? We were told that inspectors are typically assigned territories near their residence in order to minimize transportation costs. We do observe most inspectors’ assignments clustered by no more than four zip codes ¹⁸. Further analysis suggests two other factors are important for inspector assignment, namely inspector retirement and new

¹⁸In particular, for each inspector-year, we list all the zip codes where an inspector conducted initial inspections and find that 80 to 90% of her assignments concentrate in four zipcodes. The average Herfindahl index of zip codes within each inspector in a given year is around 3,500. If we examine inspector assignment by zip-code, the Herfindahl index of inspectors within each zip-code is on average 7,000 to 9,000 each year. This suggests that a zip-code is typically served by only one or two inspectors. Moreover, comparing inspector assignment from one year to the next, we find that 57% of inspectors carry over at least 50% of her top-4 zip-code assignments to the next year. Should new inspectors be assigned mostly in an attempt to break a restaurant-inspector relationship, we should see weaker geographical concentration by inspectors and greater turnover of assignments between two consecutive years.

hiring. We define an inspector’s retirement date as the date of her last appearance in our inspection data. An inspector is counted as a new hire in the quarter of her first appearance in our data.

A more relevant concern is that the DHR may assign new inspectors according to a restaurant’s last inspection record. To check this, we use the inspection-level data and regress the dummy of new inspector on the restaurant’s total violations found in the previous inspection, in addition to restaurant fixed effects, year-month fixed effects, *Lrepeat*, and other attributes of the previous inspection. Results are reported in the first column of Table 4, showing that the total violations in the previous inspection has a small (-0.0014) but statistically significant effect on the propensity of new inspection. The second column of Table 4 breaks the total violations in the previous inspection into critical, risk-factor, and non-critical categories. To our surprise, the number of critical violations found in the previous inspection is positively correlated with having a new inspector next time, but the number of risk-factor or non-critical violations shows an opposite effect. We do not have a good explanation for this, but this pattern cannot explain the universal sign of new-repeat difference in our OLS regression (Table 3).

In the last two columns of Table 4, we include on the right hand side the proportion of retired inspectors and the proportion of new hires in the corresponding subdistrict and the previous quarter (there are 186 subdistricts in Florida). As we expect, both variables are highly correlated with whether a restaurant receives a new inspection. To the extent that inspector retirement and hiring are beyond the control of any single restaurant¹⁹, we believe the inspector assignment driven by subdivision-level retirement and hiring can be treated as exogenous.

In light of this, we use two dummy variables as the instrumental variables for *New*, indicating whether there was any retirement or new hire respectively in the corresponding subdistrict and the previous quarter. Their interactions with *Lrepeat* serve as instrumental variables for *New·Lrepeat*. The instrumental variable results are reported in the first panel of Table 5, for critical, risk-factor, and non-critical violations separately. The coefficients of *New* and *Lrepeat* are all of the same sign and the same significance level as in the OLS results (Table 3). These magnitudes are also similar to the OLS estimates, except

¹⁹The DHR faced severe labor shortages in the sample period (OPPAGA 2005; 2007). This implies that inspectors had a hard time meeting the inspection frequency as required by law, and hence had little room to conduct extra inspections.

for the *New* coefficient in the non-critical regression. Furthermore, the coefficients of $New \cdot Lrepeat$ are all positive and of similar magnitude as in the OLS; however, only the risk-factor regression retains the same statistical significance as in the OLS on this interaction term.

As another test of the potential endogeneity of *New*, the second panel of Table 5 uses propensity score matching (PSM), where the propensity score prediction is based on Column 1 of Table 4. We conduct the propensity score matching for each of the three violation categories, and by the exact value of *Lrepeat*. The reported estimate is for the coefficient of *New*, along with its 90% bootstrapped confidence interval. These PSM estimates are of similar magnitude as in the OLS results, and none of the confidence intervals includes zero.

Lastly, should new inspector assignment be targeted (and restaurants know the assignment rule), we shall observe a bigger new-repeat difference when the chance of getting a new inspector is low. To check this, we rerun the OLS regressions for a few subsamples, depending on whether there is high inspector employment turnover at the subdistrict level in the previous quarter, whether $Lrepeat \leq 3$, and whether the predicted probability of getting a new inspector is below 10%. As shown in Appendix Table A2, in each of these subsamples, we obtain statistically similar results as in the full sample OLS regression (Table 3). The coefficients of *New* are usually smaller than those of full sample OLS, suggesting that restaurants do not increase their compliance effort as the predicted probability of new inspector increases.

Above all, we conclude that inspector assignment is largely driven by factors beyond the control of individual restaurants, and the realization of inspector identity is hard to predict by individual restaurants. Given the similarity between our OLS and other estimations, we believe the OLS estimates based on the assumed exogeneity of new inspection reflect little bias. The rest of the paper will proceed as if new inspector assignment were exogenous, as neither IV nor PSM can be easily applied to the maximum likelihood estimation of our full model.

4 Results from the Full Model

Table 6 reports coefficient estimates from our full model. As discussed in Section 2.2, we estimate the specification for critical, risk-factor, and non-critical violations separately, so that inspectors are allowed to differ in both overall stringency and relative taste across the three types of violations. All estimations maximize the conditional likelihood with restaurant fixed effects. According to HHG (1984) and Wooldridge (1999), restaurant fixed effects will drop out of the likelihood function conditional on the total violations found at a restaurant. Hence, we only need to estimate μ_{ic} , β_{hetero} , β_{new} , β_{pda} , β_{Lpda1} , β_{Lpda0} , β_{new} , β_{Lrpt} , $\beta_{newLrpt}$ and ζ . Our (partial) identification of detection and compliance depends on three variations – inspector heterogeneity, new inspector assignment, and inspector-restaurant relationship – so we discuss the corresponding coefficients below.

4.1 Inspector Heterogeneity

We have two kinds of coefficients for inspector heterogeneity: one is inspector fixed effects for each type of violation (μ_{ic}), and the other is the coefficient β_{hetero} . According to our theory, μ_{ic} captures how inspector i differs in her stringency on category group c from the benchmark inspector (defined as the most frequent inspector in our sample), evaluated when the restaurant fully anticipates such stringency. In comparison, β_{hetero} captures an upper bound of how the restaurant has complied with the stringency of the last inspector when this inspector and the last inspector differ in μ_{ic} by one unit.

Figure 9 plots the kernel density of the estimated μ_{ic} for critical, risk-factor, and non-critical violations respectively. An estimate of $\mu_{ic} = 0.5$ should be interpreted as inspector i on average reports $\exp(0.5) - 1 = 65\%$ more violations than the benchmark inspector. According to this interpretation, Figure 9 shows enormous inspector heterogeneity in all three types of violations.²⁰

In the bottom panel of Table 6, we show that, within each inspector, the estimated μ_{ic} has a correlation coefficient of 0.73 between critical and risk-factor violations, 0.44 between critical and non-critical, and 0.47 between risk-factor and non-critical. This suggests that inspector heterogeneity is driven by differences in both overall stringency (applicable to all

²⁰After constructing robust standard errors according to Wooldridge (1999), we find that 22 inspectors have a stringency statistically different (at 95% confidence) from the reference inspector for critical violations, 8 inspectors are different for risk-factor violations, and 166 inspectors are different for non-critical violations.

three groups) and relative taste specific to each group of categories.

Despite the enormous variation in μ_{ic} , β_{hetero} is small. Our theory predicts $\beta_{hetero} > 0$ because, if the last inspector is more stringent than the current inspector ($\mu_{ic} - \mu_{i-1c} < 0$) and the restaurant anticipates that the last inspector will come back, it should comply more; such extra compliance will be reflected in fewer violations reported by the current and less-stringent inspector. This prediction holds for critical violations: for a new inspector who is 65% less stringent than the last (corresponding to $\mu_{ic} - \mu_{i-1c} = -0.5$), she will find 1.9% fewer violations in addition to what she would find if the restaurant fully anticipates her arrival. As we have argued above, this indicates rather low compliance in response to the high stringency of the last inspector. Less consistent with the theory, we find that β_{hetero} is indifferent from zero for risk-factor violations and significantly negative for non-critical violations. The latter could occur if the restaurant faces time or personnel constraints in compliance so that more compliance in critical violations implies less compliance in non-critical violations.²¹

4.2 New versus Repeat Inspectors

After controlling for inspector heterogeneity and corresponding compliance response, the coefficient of *NEW* captures the extent to which a new inspector reports more violations even if her μ_{ic} is the same as the last inspector. The full-model estimates suggest that a new inspector reports 17.5% more critical violations, 14.6% more risk-factor violations, and 12.7% more non-critical violations compared to a repeat inspector coming back for her second visit to the restaurant.²² This large effect, combined with its increase with the length of the relationship between the last inspector and the restaurant (as indicated by the significant positive coefficient of $\beta_{newLrpt}$), suggests that new inspectors may have significantly fresher eyes in their first visit of a restaurant. In comparison, the coefficient on the relationship alone, β_{Lrpt} , is small but negative, suggesting that among repeat inspectors one extra visit brings down the reported violations by only 0.7-1.8%. As elaborated in the theory, this number could reflect a mixture of gradual compliance to the last inspector's stringency, the gradual slacking of repeat inspectors over time, or more compliance with the growing fear of a new inspector coming next time. The positive and significant estimate of

²¹Our theory does not address such constraints.

²²These percentages are computed by $(\exp(\beta_{new} + \beta_{newLrpt} \cdot 1) - 1) \cdot 100\%$ because the lowest value of *Lrepeat* is one.

$\beta_{newLrpt}$ (across all three groups of categories) suggests that slackness of repeat inspectors is probably the most likely among these possibilities.

One may argue that the large fresh-eye effect of new inspectors can be interpreted as restaurants catering to inspector heterogeneity within each of the three groups of violations. To address this possibility, we rerun the full model for four frequently cited categories separately (they are categories 2, 8, 22, and 23). In all four estimations, β_{new} is statistically significant from zero with 99% confidence. The coefficient magnitude implies that a new inspector, compared to a repeat inspector’s second visit, reports 16.69% more violations in category 2, 10.79% more in category 8, 3.80% more in category 22, and 8.99% more in category 23. In comparison, the fresh-eye effect derived from the group-wise model is 14.6% more for risk-factor violations (categories 2 and 8 are risk factors), and 12.7% more for non-critical violations (categories 22 and 23 are non-critical). These findings suggest that new inspectors demonstrate significant fresh-eye effects within each category, although some of the group-wise fresh-eye effects can be explained by inspector heterogeneity across specific categories within the same category group.

Coefficients of some control variables may suggest fresh eyes as well. To the extent that young inspectors (whose tenure, defined as the number of previous inspections done by the inspector, is no more than 30) probably have the freshest memory of the FDA training and are still learning where to pay more attention, they find 18.4-27.1% more violations across all three groups of categories. Short tenure (tenure less than median) also implies more detected violations, suggesting that the loss of fresh eyes is applicable not only to a long restaurant-inspector relationship but also to the long tenure of inspectors. The sensitivity to tenure is consistent with what Macher et al. (2011) found regarding the inspection of pharmaceutical manufacturing, although we do not have detailed training and tenure data as they do. Surprisingly, there is no obvious fresh-eye effect for the first inspection of the day, but inspectors do tend to find fewer violations throughout the day, probably due to fatigue. In addition, older restaurants tend to have more violations, as do inspections made during lunch time.

4.3 Model Fit and Counterfactual Simulations

In Table 6, we report the goodness of fit for each column. The calculation follows Cameron and Windmeijer (1996), which describes the log likelihood improvement from a constant-

only model to our model as a fraction of the log likelihood improvement from the constant-only model to the perfect fit using the raw dataset itself. By definition, it is bounded within 0 and 1. The goodness of fit measure is 0.45 for critical violations, 0.55 for risk factors, and 0.60 for non-critical violations, suggesting that our model fits the raw data reasonably well.

Table 6 also reports the comparison of our predicted violations versus the actual violations in both mean and standard deviation. The mean is literally zero for all three columns, while the standard deviation is between 1.5 and 2.7. Most of the seemingly large standard deviation is driven by the fact that our predicted violations are continuous but actual violations are integers. In light of this, Table 7 reports the discrepancy in predicted and actual violations by quartiles. Not surprisingly, the actual violations are more dispersed but their medians are close to the predicted violations, especially in risk-factor and non-critical categories (because these two groups are less censored at zero).

So far we have found evidence for both inspector heterogeneity and the fresh-eye effect of new inspectors. Which is greater in magnitude? To answer this question, we conduct three counterfactual simulations. The first simulation assumes that inspectors are assigned randomly within each district. Specifically, we compute the frequency of each inspector in the raw data and use this as the weight of random assignment for that inspector. This way, the number of assignments for every inspector is similar to that of raw data, but the assignment itself is random. To minimize simulation error, we simulate random assignment 100 times and compute the average predicted violations for each inspection. The simulation results are presented in Table 7. Many variables in the full model are related to inspector assignment, but the greatest change is that over 80% of random assignments involve new inspectors as compared to 27% in the existing assignments. Consequently, random assignments on average yield 11.35-17.57% more detected violations. If we decompose these effects by different parts of the full model, we find that most of the effects are driven by the large fresh-eye effect of new inspectors rather than inspector heterogeneity.

The second and third counterfactual simulations aim to compare the raw data with situations without any inspector heterogeneity. To do this, we keep the same inspector assignment as in the raw data but assume that every inspector is the same as either the average inspector or the most stringent inspector. Comparing the former to the raw data, inspector homogeneity leads to lower mean and lower dispersion in the simulated distri-

bution of detected violations. However, in the latter simulation, inspector homogeneity leads to higher mean and higher dispersion of simulated violations. This is because more stringency leads to more violations in every inspection, and this increase is greater in magnitude for dirtier restaurants due to the exponential functional form of the Poisson model. This also contributes to the increased dispersion of violations.

Overall, the counterfactual simulations reinforce the conclusion that both inspector heterogeneity and inspector-restaurant relationship contribute significantly to inspection outcomes. If the DHR wants inspectors to detect more violations, it could rotate the inspectors more often or train them to be more homogeneously stringent. A simple reduction of inspector heterogeneity will not do the trick, if all inspectors converge to be the average inspector rather than the most stringent inspector.

5 Conclusions

Government inspections often involve repeated interaction between inspectors and inspectees. In this paper, we use restaurant hygiene inspections as an example to show that inspector assignment and repetition can have significant impact on inspection outcomes. In particular, we find that new inspectors report 12.7-17.5% more violations than the second visit of a repeat inspector, and this effect is more pronounced if the previous inspector has had a longer relationship with the restaurant. The difference between new and repeat inspectors is attributed to two factors: (1) new inspectors tend to have fresher eyes in their first visit of a restaurant; and (2) inspectors differ greatly in stringency and taste, such inspector heterogeneity motivates restaurants to adjust their compliance effort according to the criteria of their previous inspectors. Both factors are found to be important in our data.

Our findings have important implications for the design of the inspection program. Counterfactual simulations suggest that detection can be further enhanced by a more frequent rotation of inspectors or greater efforts to ensure that inspectors are homogeneously stringent. More specifically, if one is willing to take the simulation numbers at their face value, random assignment of inspectors may report 1.086 more violations than the status quo assuming everything else is equal. The reality could be more complicated as restaurants are likely to adjust their expectations in response to more frequent rotation of inspectors,

which we cannot account for in a simple counterfactual simulation. Our results also suggest that inspector tenure and inspector training may affect inspection outcomes as well, a topic that is definitely worth further study.

References

- Cameron, A. Colin** and Frank A. G. Windmeijer. 1996. “R-Squared Measures for Count Data Regression Models with Applications to Health-Care Utilization.” *Journal of Business & Economic Statistics*, 14(2): 209-20.
- Dranove, David** and Ginger Zhe Jin. 2010. “Quality Disclosure and Certification: Theory and Practice.” *Journal of Economic Literature* 48(4): 935-63.
- Dixit, Avinash**. 2002. “Incentives and Organizations in the Public Section: An Interpretative Review.” *The Journal of Human Resources*, 37(4): 696-727.
- Feinstein, Jonathan**. 1989. “The Safety Regulation of U.S. Nuclear Power Plants: Violations, Inspections, and Abnormal Occurrences.” *Journal of Political Economy*, 97: 115-154.
- Feinstein, Jonathan**. 1991. “An Econometric Analysis of Income Tax Evasion and its Detection.” *RAND Journal of Economics*, 22(1): 14-35.
- Prentice Hall.
- Hausman, Jerry**, Bronwyn H. Hall, and Zvi Griliches. 1984. “Econometric Models for Count Data with an Application to the Patents-R & D Relationship.” *Econometrica*, 52(4): 909-938.
- Jin, Ginger Zhe** and Phillip Leslie (2003) “The Effects of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards,” *The Quarterly Journal of Economics*, May 2003, 118(2), 409-51.
- Jin, Ginger Zhe** and Jungmin Lee (2014) “Inspection technology, detection and compliance: evidence from Florida Restaurant inspections,” forthcoming *RAND Journal of Economics*.

- Kleven**, Henrik J., Martin B. Knudsen, Claus T. Kreiner, Soren Pedersen, and Emmanuel Saez (2010). "Unwilling or Unable to Cheat? Evidence from a Randomized Tax Audit Experiment in Denmark." *Econometrica* 79(3): 651-692.
- Laffont**, Jean-Jacques and Tirole, Jean. 1993. *A Theory of Incentives in Procurement and Regulation*. Cambridge, Mass.: MIT Press.
- Levine**, David I. ; Michael W. Toffel and Matthew S. Johnson (2012) "Randomized Government Safety Inspections Reduce Worker Injuries with No Detectable Job Loss." *Science*, 336 (6083), 18 May 2012, pp. 907-911.
- Macher**, Jeffery T., John W. Mayo, and Jack A. Nickerson. 2011. "Exploring the Information Asymmetry Gap: Evidence from FDA Regulation." *Journal of Law & Economics* 54(1): 25-54.
- Martimort**, David (1999). "The Life Cycle of Regulatory Agencies: Dynamic Capture and Transaction Costs." *The Review of Economic Studies* 66(4): 929-947.
- Mookherjee**, Dilip and Ivan Png (1989). "Optimal Auditing, Insurance and Redistribution." *The Quarterly Journal of Economics* 104(2): 399-415.
- Mookherjee**, Dilip and Ivan Png (1995). "Corruptible Law Enforcers: How Should They Be Compensated?" *The Economic Journal* 105(428): 145-159.
- Office** of Program Policy Analysis & Government Accountability. 2005. "Division of Hotels and Restaurants Improves Operations But Not Meeting Inspection Goals." Report No. 05-51, An Office of the Florida Legislature.
- Office** of Program Policy Analysis & Government Accountability. 2007. "Division of Hotels and Restaurants Improves Operations and Makes Progress in Meeting Inspection Goals." Report No. 07-41.
- Prendergast**, Canice. 1999. "The Provision of Incentives in Firms." *Journal of Economic Literature*, 37(1): 7-63.
- Prendergast**, Canice. 2007. "The Motivation and Bias of Bureaucrats." *American Economic Review*, 97(1): 180-196.

- Slemrod**, Joel and Shlomo Yitzhaki. 2002. "Tax Avoidance, Evasion and Administration" in the Handbook of Public Economics, No. 3, edited by Alan Auerbach and Martin Feldstein, Elsevier.
- Tirole**, Jean (1986) "Hierarchies and Bureaucracies: On the Role of Collusion in Organizations", Journal of Law, Economic and Organization, 2, 181-214.
- Woodridge**, Jeffery (1999) "Distribution-free Estimation of Some Nonlinear Panel Data Models." Journal of Econometrics 90: 77-97.

Figure 1. Equilibrium under Perfect information, Category 1

Figure 2. Overestimation of Inspector's Effort Cost ($\tilde{\theta}_i > \theta_i$), Category 1

Figure 3. Expectation of a New (More Stringent) Inspector. Category 1
 (Assume the effort cost of a repeat inspector remains the same, but the probability of a new inspector increases with the number of times the previous inspector has inspected the restaurant.)

Figure 4. Expectation of the Repeat Inspector's Slack, Category 1
 (Assume the effort cost of a repeat inspector increases with repetition, while the probability of a new inspector remains constant.)

Figure 5. The Estimated Probability of New Inspector Arrival by the Number of Inspections of the Previous Inspector

Notes: The linear probability model for a new inspector's arrival is estimated by using the specification of Column (1) of Table 4. The probability is relative to that when the number of inspections by the previous inspector is one. The estimates when the number of inspections by the previous inspector is greater than 25 are omitted.

Figure 6. Average Violations by Inspection History

Notes: This graph tracks each restaurant's inspection history since the date of its first "new" regular inspection as observed in the sample where an inspection is counted as "new" if the inspector making this inspection has never inspected this restaurant in our data of regular inspections.

Figure 7. Inspector Heterogeneity: Individual Fixed Effects

Notes: Inspector fixed effects (FE) are estimated by regressing total violations on inspection month and fiscal year FE as well as restaurant FE. There are 358 inspectors in total. By adding inspector FE, the R squared increases from 0.034 to 0.131. The omitted inspector's ID = 59. The most outlying inspector (ID = 49233) has only 12 inspections in the final sample for all different restaurants. The average number of inspections per inspector is 1,192, the median 952 and the maximum 4,275.

Figure 8. Inspector Heterogeneity by Category

Notes: Inspector-category fixed effects are estimated by regressing category-specific violations on restaurant fixed effects, month and fiscal year fixed effects. The graph shows the frequency of inspectors in each category for which they are most likely to detect violations. There are 358 unique inspectors in the sample. See Appendix A for categories.

Figure 9. Kernel Densities of Inspector Fixed Effects Estimated from the Full Model

Notes: We use the most frequent inspector in our data as the benchmark. All the inspector fixed effects are relative to this single inspector.

Table 1. Variables and Summary Statistics

	Mean	Std.dev.	Min	Max
Number of critical violations	1.62	1.93	0	33
Number of risk factors	2.54	2.84	0	44
Number of non-critical violations	3.75	3.96	0	62
Number of all violations	7.90	7.16	0	111
New inspector	0.27	0.44	0	1
Number of inspections by previous inspector	3.65	3.07	1	38
PDA inspection	0.89	0.31	0	1
Number of previous PDA inspections	5.06	3.87	0	41
Whether the last inspection has used PDA	0.92	0.28	0	1
Restaurant age (years)*	4.10	2.74	0	14.5
Missing age	0.24	0.42	0	1
Complaint inspection	0.04	0.19	0	1
Licensing inspection	0.001	0.03	0	1
Inspector tenure (# of inspections done before t)	1535	1114	0	5791
Inexperienced inspector (tenure less than the median)	0.44	0.5	0	1
Novice inspector (tenure $\leq$ 30)	0.01	0.11	0	1
First inspection of the day**	0.28	0.45	0	1
Number of previous inspections in the same day**	1.6	1.53	0	35
Missing inspection time	0.1	0.3	0	1
Time span from the last regular inspection (in days)	184	92	1	2004
Lunch time	0.38	0.49	0	1
Total # of restaurants	60,976			
Total # of inspectors	358			
Total # of inspections	426,831			
Total # of category-inspection (55 categories per inspection)	23,475,705			

* N = 326,461 obs. with age not missing. ** N = 384,198 obs. with inspection time not missing.

Table 2. Number of Violations: New vs. Repeat Inspectors

	(1) Repeat Inspector	(2) New Inspector	(2) – (1) New – Repeat (%)
A. All inspections			
All violations	7.38 (6.63)	9.36 (8.27)	27%***
Critical violations	1.47 (1.79)	2.01 (2.24)	37%***
Risk factors	2.36 (2.66)	3.05 (3.23)	29%***
Non-critical violations	3.54 (3.74)	4.31 (4.45)	22%***
B. By number of inspections by previous inspector			
B.1 # of inspections by previous inspector = 1 to 3			
Critical violations	1.49 (1.81)	1.96 (2.19)	32%***
Risk factors	2.32 (2.68)	2.91 (3.13)	25%***
Non-critical violations	3.44 (3.77)	4.14 (4.32)	20%***
B.2 # of inspections by previous inspector = 4 to 6			
Critical violations	1.43 (1.73)	2.13 (2.32)	49%***
Risk factors	2.35 (2.63)	3.35 (3.40)	43%***
Non-critical violations	3.57 (3.70)	4.63 (4.63)	30%***
B.3 # of inspections by previous inspector = 7 to 9			
Critical violations	1.48 (1.76)	2.25 (2.42)	52%***
Risk factors	2.47 (2.66)	3.54 (3.53)	43%***
Non-critical violations	3.79 (3.72)	5.03 (4.96)	33%***
B.4 # of inspections by previous inspector = 10 or more			
Critical violations	1.49 (1.77)	2.20 (2.45)	48%***
Risk factors	2.57 (2.61)	4.00 (3.81)	56%***
Non-critical violations	3.94 (3.65)	5.42 (5.15)	38%***

Notes: Standard deviations are presented in parentheses. *** p<0.01.

Table 3. OLS Results with Restaurant, Inspector, Year-Month Fixed Effects

	(1)	(2)	(3)
Violation type	Critical	Risk factor	Non-critical
Sample average of dependent variable	1.62	2.54	3.75
New inspector	0.2657*** (0.0109)	0.3140*** (0.0148)	0.3930*** (0.0199)
# visits by the last inspector	-0.0128*** (0.0018)	-0.0457*** (0.0026)	-0.0186*** (0.0035)
New inspector × # visits by the last inspector	0.0152*** (0.0030)	0.0467*** (0.0042)	0.0243*** (0.0055)
Restaurant age	0.0113*** (0.0041)	-0.0067 (0.0063)	0.0372*** (0.0088)
Missing age	0.4948 (0.6044)	-0.1683 (0.4612)	-0.2660 (0.8647)
Complaint inspection	-0.3553*** (0.0160)	-0.4231*** (0.0219)	-0.7132*** (0.0320)
Licensing inspection	-0.4352*** (0.0917)	-0.6340*** (0.1222)	-1.2830*** (0.1557)
Short tenure	0.0959*** (0.0097)	-0.0268** (0.0132)	0.1443*** (0.0185)
Young inspector	0.4867*** (0.0366)	0.4658*** (0.0465)	1.0624*** (0.0681)
Fatigue	-0.0529*** (0.0024)	-0.0755*** (0.0033)	-0.1125*** (0.0044)
First of the Day	0.0344*** (0.0089)	0.0398*** (0.0121)	0.0787*** (0.0164)
Missing fatigue	-0.1123*** (0.0177)	-0.2238*** (0.0240)	-0.2904*** (0.0327)
PDA	0.1127*** (0.0210)	0.3996*** (0.0292)	0.5333*** (0.0398)
# of previous PDA inspections	-0.0864*** (0.0050)	-0.0422*** (0.0075)	-0.0105 (0.0104)
# of previous PDA inspections × PDA	-0.0233*** (0.0035)	-0.0379*** (0.0054)	-0.0854*** (0.0074)
Time span	-0.0770*** (0.0059)	-0.1229*** (0.0078)	-0.1019*** (0.0106)
Lunch	0.0401*** (0.0064)	0.0485*** (0.0088)	0.0228** (0.0115)
Restaurant FE	Yes	Yes	Yes
Inspector FE	Yes	Yes	Yes
Month-by-Year FE	Yes	Yes	Yes
Adj. R squared	0.1306	0.1579	0.1494
Number of observations	426,831	426,831	426,831

Notes: Robust standard errors are presented in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4. The Probability of New Inspector Arrival

	(1)	(2)	(3)	(4)
<i>Previous inspection's characteristics</i>				
Total violations	-0.0014*** (0.0002)		-0.0015*** (0.0002)	
Critical violations		0.0019*** (0.0005)		0.0016*** (0.0005)
Risk factor violations		-0.0012*** (0.0004)		-0.0012*** (0.0004)
Noncritical violations		-0.0030*** (0.0003)		-0.0030*** (0.0003)
Retired inspectors in previous quarter			1.0815*** (0.0212)	1.0805*** (0.0212)
New hires in previous quarter			0.3643*** (0.0169)	0.3632*** (0.0169)
Restaurant age	-0.0279*** (0.0011)	-0.0278*** (0.0011)	-0.0275*** (0.0011)	-0.0274*** (0.0011)
Missing age	0.0900 (0.1002)	0.0885 (0.1003)	0.0930 (0.1000)	0.0916 (0.1001)
Complaint inspection	0.0205*** (0.0041)	0.0207*** (0.0041)	0.0211*** (0.0041)	0.0214*** (0.0041)
Licensing inspection	-0.0057 (0.0278)	-0.0060 (0.0278)	-0.0091 (0.0282)	-0.0093 (0.0282)
Short tenure	-0.0666*** (0.0022)	-0.0665*** (0.0022)	-0.0639*** (0.0022)	-0.0639*** (0.0022)
Young inspector	-0.0215*** (0.0080)	-0.0219*** (0.0080)	-0.0320*** (0.0079)	-0.0323*** (0.0079)
FE of # visits by the last inspector	Yes	Yes	Yes	Yes
Restaurant FE	Yes	Yes	Yes	Yes
Month-by-Year FE	Yes	Yes	Yes	Yes
R squared	0.2732	0.2733	0.2837	0.2838
Number of observations	365,855	365,855	365,304	365,304

Notes: Linear probability models are estimated. Robust standard errors are presented in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 5. Instrumental Variable and Propensity Score Matching Results by Violation Type

Violation type	(1) Critical	(2) Risk factor	(3) Non-critical
<i>A. Instrumental variable results</i>			
New inspector	0.2751*** (0.0795)	0.3124*** (0.1052)	1.1706*** (0.1431)
# visits by the last inspector	-0.0067 (0.0050)	-0.0606*** (0.0069)	-0.0481*** (0.0089)
New inspector × # visits by the last inspector	0.0024 (0.0182)	0.0665*** (0.0249)	0.0049 (0.0319)
Control variables	Yes	Yes	Yes
Restaurant FE	Yes	Yes	Yes
Inspector FE	Yes	Yes	Yes
Month-by-Year FE	Yes	Yes	Yes
Adj. R squared	0.0654	0.0833	0.0485
Number of observations	418,038	418,038	418,038
<i>B. Propensity score matching results</i>			
# visits by the last inspector = 1	0.2174 [0.1713, 0.2440]	0.3410 [0.2767, 0.3846]	0.4650 [0.3758, 0.5310]
# visits by the last inspector = 2	0.4227 [0.3623, 0.4576]	0.5609 [0.4878, 0.6300]	0.6125 [0.5182, 0.7296]
# visits by the last inspector = 3	0.5307 [0.4615, 0.5970]	0.6461 [0.5414, 0.7541]	0.7264 [0.5723, 0.8687]
# visits by the last inspector = 4	0.5756 [0.4905, 0.7169]	0.7096 [0.5704, 0.8725]	0.5305 [0.2545, 0.7417]
# visits by the last inspector = 5	0.6400 [0.4284, 0.7908]	0.6622 [0.4698, 1.0024]	1.0307 [0.7300, 1.3983]
# visits by the last inspector = 6	0.4746 [0.1638, 0.8142]	1.4873 [0.8850, 1.9788]	1.0424 [0.4304, 1.7941]
# visits by the last inspector = 7	0.6429 [-0.0956, 1.3846]	1.6964 [0.5686, 3.2679]	1.3214 [0.1373, 3.2083]

Notes: Instrumental variables are a dummy variable for any retired inspectors in the subdistrict in the previous quarter, a dummy for any new hires in the subdistrict in the previous quarter, and their interactions with (# visits by the last inspector). For IV, robust standard errors are presented in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Cragg-Donald Wald F statistic is 1,396. For PSM, 90% bootstrapped confidence intervals are presented in squared brackets. The propensity score is estimated based on Column (1) of Table 4.

Table 6. Structural Estimation Results

Dependent variable	(1) Critical	(2) Risk factor	(3) Non-critical
Sample average of dependent variable	1.62	2.54	3.75
New inspector	0.1508*** (0.0075)	0.1224*** (0.0043)	0.1163*** (0.0059)
# visits by the last inspector	-0.0180*** (0.0012)	-0.0189*** (0.0007)	-0.0068*** (0.0011)
New inspector × # visits by the last inspector	0.0102*** (0.0018)	0.0141*** (0.0011)	0.0030** (0.0015)
Restaurant age	0.0074** (0.0032)	-0.0059*** (0.0020)	0.0149*** (0.0026)
Missing age	0.3614 (0.3212)	0.1532 (0.1206)	0.1273 (0.2570)
Complaint inspection	-0.2017*** (0.0096)	-0.1491*** (0.0053)	-0.1531*** (0.0077)
Licensing inspection	-0.3082*** (0.0625)	-0.2978*** (0.0479)	-0.4065*** (0.0572)
Less than median experience	0.0593*** (0.0080)	-0.0082* (0.0048)	0.0437*** (0.0065)
Novice inspector	0.2377*** (0.0162)	0.1686*** (0.0098)	0.2396*** (0.0137)
# inspections before the current inspection	-0.0391*** (0.0019)	-0.0358*** (0.0011)	-0.0362*** (0.0015)
First inspection of the day	0.0030 (0.0057)	0.0046 (0.0034)	0.0085* (0.0045)
Missing inspection time	-0.0830*** (0.0113)	-0.1191*** (0.0074)	-0.0938*** (0.0102)
Time span	-0.0626*** (0.0037)	-0.0523*** (0.0022)	-0.0365*** (0.0030)
Lunch	0.0250*** (0.0041)	0.0211*** (0.0025)	0.0081** (0.0032)
Inspector heterogeneity	0.0379*** (0.0142)	0.0009 (0.0114)	-0.0547* (0.0309)
PDA=1	0.1809*** (0.0208)	0.2908*** (0.0150)	0.2915 (0.0275)
# of previous PDA=1 * PDA=1	-0.0396*** (0.0108)	-0.0483*** (0.0061)	-0.0304*** (0.0084)
# of previous PDA=2 * PDA=1	-0.0633*** (0.0114)	-0.0887*** (0.0067)	-0.0343*** (0.0090)
# of previous PDA>=3 * PDA=1	-0.1069*** (0.0121)	-0.1363*** (0.0074)	-0.0492*** (0.0097)

# of previous PDA=1 * PDA=0	-0.0006 (0.0216)	0.0554*** (0.0156)	0.0973*** (0.0273)
# of previous PDA=2 * PDA=0	0.0023 (0.0226)	0.0286* (0.0165)	0.1163*** (0.0281)
# of previous PDA>=3 * PDA=0	0.0679*** (0.0211)	0.0644*** (0.0155)	0.2275*** (0.0059)
Restaurant FE	Yes	Yes	Yes
Inspector FE	Yes	Yes	Yes
Month-by-Year FE	Yes	Yes	Yes
Log Likelihood	-1,405,660	-2,219,280	-3,350,400
Number of observations	426,831	426,831	426,831

. Note: *** p<0.01, ** p<0.05, * p<0.1.

<u>Correlations between Inspector FEs</u>			
Risk factor	0.727		
Non-critical	0.442	0.474	
Goodness of fit (pseudo R-squared)	0.451	0.553	0.605
Mean of (Predicted Y - Y)	0	0	0
Std. dev. of (Predicted Y - Y)	1.468	1.974	2.672

Notes: The three columns are estimated separately. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. All columns control for whether or not this and previous inspections are paperless.

Table 7. Counterfactual Simulations

Simulations by Category and Scenario	Mean	Std. dev.	1 st quartile	2 nd quartile	3 rd quartile
Critical					
Raw data	1.616	1.931	0.000	1.000	2.000
Predicted	1.616	1.334	0.694	1.293	2.161
Random weighted assignment	1.900	1.388	0.927	1.608	2.542
Every inspector same as average	1.486	1.179	0.674	1.215	1.989
Every inspector same as the most stringent	5.024	3.987	2.279	4.106	6.725
Risk factor					
Raw data	2.542	2.839	0.000	2.000	4.000
Predicted	2.542	2.175	1.013	1.975	3.459
Random weighted assignment	2.919	2.174	1.324	2.442	4.036
Every inspector same as average	2.527	2.178	1.020	1.963	3.428
Every inspector same as the most stringent	8.808	7.591	3.554	6.840	11.948
Non-critical					
Raw data	3.745	3.958	1.000	3.000	5.000
Predicted	3.745	3.090	1.535	2.998	5.117
Random weighted assignment	4.170	3.011	1.958	3.583	5.752
Every inspector same as average	3.498	2.737	1.539	2.911	4.789
Every inspector same as the most stringent	9.384	7.342	4.129	7.809	12.847

Notes: Based on the coefficients estimated in the full model (Table 6).

Appendix Table A1. Inspection Categories

Category Number	Details	Type
1	Approved source	R
2	Original container: properly labeled, date marking, consumer advisory	R
3	Food Out of Temperature	R
4	Facilities to maintain product temperature	C
5	Thermometers provided and conspicuously placed	C
6	Potentially hazardous food properly thawed	C
7	Unwrapped or potentially hazardous food not re-served	R
8	Food protection, cross-contamination	R
9	Foods handled with minimum contact	R
10	In use food dispensing utensils properly stored	N
11	Personnel with infections restricted	R
12	Hands washed and clean, good hygienic practices, eating/drinking/smoking	R
13	Clean clothes, hair restraints	N
14	Food contact surfaces designed, constructed, maintained, installed, located	N
15	Non-food contact surfaces designed, constructed, maintained, installed, located	N
16	Dishwashing facilities designed, constructed, operated	C
17	Thermometers, gauges, test kits provided	C
18	Pre-flushed, scraped, soaked	N
19	Wash, rinse water clean, proper temperature	N
20	Sanitizing concentration or temperature	C
21	Wiping cloths clean, used properly, stored	N
22	Food contact surfaces of equipment and utensils clean	N
23	Non-food contact surfaces clean	N
24	Storage/handling of clean equipment, utensils	N
25	Single service items properly stored, handled, dispensed	N
26	Single service articles not re-used	N
27	Water source safe, hot and cold under pressure	C
28	Sewage and wastewater disposed properly	C
29	Plumbing installed and maintained	N
30	Cross-connection, back siphonage, backflow	C
31	Toilet and hand-washing facilities, number, convenient, designed, installed	C
32	Restrooms with self-closing doors, fixtures operate properly, facility clean, supplied with hand-soap, disposable towels or hand drying devices, tissue, covered waste receptacles	R
33	Containers covered, adequate number, insect and rodent proof, emptied at proper intervals, clean	N
34	Outside storage area clean, enclosure properly constructed	N
35	Presence of insects/rodents. Animals prohibited. Outer openings protected from insects, rodent proof	C
36	Floors properly constructed, clean, drained, coved	N
37	Walls, ceilings, and attached equipment, constructed, clean	N

38	Lighting provided as required. Fixtures shielded	N
39	Rooms and equipment - vented as required	N
40	Employee lockers provided and used, clean	N
41	Toxic items properly stored, labeled and used properly	R
42	Premises maintained, free of litter, unnecessary articles. Cleaning and maintenance equipment properly stored. Kitchen restricted to authorized personnel	N
43	Complete separation from living/sleeping area, laundry	N
44	Clean and soiled linen segregated and properly stored	N
45	Fire extinguishers - proper and sufficient	C
46	Exiting system - adequate, good repair	C
47	Electrical wiring - adequate, good repair	C
48	Gas appliances - properly installed, maintained	C
49	Flammable/combustible materials - properly stored	C
50	Current license properly displayed	C
51	Other conditions sanitary and safe operation	N
52	False/misleading statements published or advertised relating to food/beverage	N
53	Food management certification valid / Employee training verification	R
54	Florida Clean Indoor Air Act	N
55	Automatic Gratuity Notice	N

Notes: Those categories where a relatively large number of inspectors (at least 10 inspectors) are concentrated in Figure 7 are in bold. In the third column, C represents critical violations, R risk factors, and N non-critical violations. There are 17 critical violation categories, 11 risk factor categories, and 27 non-critical categories. The last column presents the number of subcategories per classification in each category. For example, for category 1, there are two classification codes, A and B. Under A, there are 16 subcategories and under B there are 26. There used to be 3 more categories, which were eliminated later: category 56 Copy of Chapter 509, Florida Statutes, 57 Hospitality Education Program Information provided, and 58 Smoke Free. On the paper inspection form, categories are divided by the classification code. All the information contained in this table can be downloaded from the website of the Division of Hotels and Restaurants of Florida Department of Business & Professional Regulation (www.myfloridalicense.com/dbpr/hr/index.html).

Appendix Table A2. Robustness Checks

Violation type	(1) Critical	(2) Risk factor	(3) Non-critical
<i>A. High inspector turnover at previous quarter</i>			
New inspector	0.1480*** (0.0406)	0.2265*** (0.0524)	0.3208*** (0.0657)
# visits by the last inspector	-0.0592*** (0.0081)	-0.0320*** (0.0107)	-0.0320** (0.0136)
New inspector ×# visits by the last inspector	0.0529*** (0.0120)	0.0307* (0.0159)	-0.0010 (0.0183)
R squared	0.6984	0.7192	0.7451
Number of observations	65,616	65,616	65,616
<i>B. # visits by the last inspector ≤ 3</i>			
New inspector	0.1850*** (0.0198)	0.2593*** (0.0264)	0.3491*** (0.0355)
# visits by the last inspector	-0.0735*** (0.0052)	-0.1335*** (0.0071)	-0.0726*** (0.0096)
New inspector ×# visits by the last inspector	0.0557*** (0.0107)	0.0841*** (0.0143)	0.0558*** (0.0192)
R squared	0.4935	0.5728	0.6030
Number of observations	262,333	262,333	262,333
<i>C. Prob(new inspector) < 10%</i>			
New inspector	0.2140*** (0.0662)	0.2084** (0.0992)	0.4848*** (0.1336)
# visits by the last inspector	-0.0261*** (0.0039)	-0.0421*** (0.0050)	-0.0236*** (0.0069)
New inspector ×# visits by the last inspector	0.0279** (0.0132)	0.0538** (0.0221)	0.0387 (0.0263)
R squared	0.5538	0.6513	0.6808
Number of observations	112,412	112,412	112,412

Notes: All regressions include all control variables and fixed effects of our full specification. Robust standard errors are presented in parentheses. Panel A restricts the sample to those where more than 10% of inspectors in the same subdistrict retired or more than 10% of inspectors in the same subdistrict were newly hired in the previous quarter. For Panel C, the probability of new inspector's arrival is predicted from Column (1) of Table 4. *** significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.