Grade Inflation and Social Ties in Restaurant Hygiene Inspections*

Balázs Kovács
Yale University

David W. Lehman
University of Virginia

Glenn R. Carroll
Stanford University

April 2018

---

*We appreciate the cooperation of LAPHD in providing data and answering questions about the inspection process. We thank John DeFigueiredo, Daniel Ho, Greta Hsu, Özgecan Koçak and Jesper Sørensen for comments on earlier drafts. We also benefitted from comments at the Harvard-MIT Economic Sociology Seminar, the MEaN meetings at Yale University and the MacroOB lunch group at Stanford.

Grade Inflation and Social Ties in Restaurant Hygiene Inspections

Abstract

Discontinuities frequently appear in the distributions of scores underlying public grading schemes such as restaurant hygiene inspections. The typical economics-based study of such discontinuities generally sees them as rational responses to the scheme's inherent incentives. By contrast, we examine how evolving social relationships between inspectors and restaurants influence grades in a systematic manner. We study 512,313 inspections in Los Angeles County from 2000–2010 and find that the stronger the social tie between an inspector and a restaurant (i.e., duration of tie and number of repeated visits), the more likely the inspection score will be artificially inflated to the next highest grade category. The effect of tie strength is robust. When aggregated, such inflated grading produces discontinuities – kinks – in the distribution of scores around grade boundaries. We also find that restaurants with likely inflated grades subsequently receive more consumer complaints. These findings carry important implications for theory, practice, and policy.

Keywords: public grading schemes, social ties, kinks, targeted transparency, inspections

Introduction

Local and other governmental entities often use mandated public information disclosure as a means of achieving policy objectives (Weil et al. 2013). These efforts, often referred to as "targeted transparency," require organizations to disclose certain specific information to the public in standardized formats. The aim is to empower consumers to make informed choices and – presumably – to prompt organizations to take steps to mitigate risks and improve quality (Fung et al. 2007, Sunstein 1999). The use of information disclosure schemes has diffused to various domains including hospitals, food calorie posting, mutual funds, and motor vehicle emissions, among others.

Restaurant grading schemes have been referred to as "the poster child" for targeted transparency (Ho 2012). Restaurant hygiene inspections generally follow scripted guidelines for checking compliance with practices regarding issues such as food temperatures, clean counters, insect and rodent infestations, and employee handwashing. The typical inspection protocol begins with a perfect score (often 100) and deducts points for each violation observed. Many local polities collapse these numerical scores into letter grades resembling those assigned to students in schools (i.e., A, B, C, etc.) and require restaurants to display publicly the assigned grades for patrons to see. By some accounts, these schemes produce the desired effects: studies show higher grades to be associated with greater market demand (Jin and Leslie 2009), higher online consumer ratings (Lehman et al. 2014), and even reduced cases of foodborne illnesses (Simon et al. 2005, but see Ho et al. 2017). Indeed, a restaurant's grade appears to many diners to be a certification of food safety.

However, a curious discontinuity can frequently be seen in the distribution of numeric scores underlying such grading schemes. Scores just above grade cutoffs are disproportionately common whereas scores just below them are infrequent. Figure 1 shows this pattern for Los Angeles County for 2000-2010. The distribution exhibits discontinuities – or kinks – at the boundaries between each of the grade categories (i.e., 70, 80, and 90). As Ho (2012) demonstrates for New York, the kinking pattern emerges after restaurant grading schemes are put in place. Moreover, for places such as San Francisco that require only the release of numeric scores, the same discontinuity does not appear. The goal of this study is to understand better why these discontinuities appear.

Over the past decade, analysts have studied discontinuities in the distribution of various kinds of scores. The corresponding kink points were originally recognized in the context of taxation where tax brackets provided economists a way to study the implications of different incentive schemes (Lalumia 2008, Saez 2010). More recently, similar patterns have been reported in welfare programs (Camacho and Conover 2011), automobile fuel economy ratings (Sallee and Slemrod 2012), electricity pricing schedules and power consumption (Ito 2014), mortgage interest rates and demand (DeFusco and Paciorek 2017), marathon finishing times (Allen et al. 2017), among others.

The typical economics-based view of such discontinuities is one of rational behavior. Discontinuities in a rate scheme produce corresponding discontinuities in the behavioral patterns of agents. Standard, if nonetheless complex, applications of classic microeconomic (e.g., Burtless and Hausman 1978) and behavioral decision-making (e.g., Kahneman and Tversky 1979) theories can usually account for these behaviors. Economists use the term "bunching" to describe and analyze the observed patterns (see Kleven 2016); agents operating within the scheme behave in such a way that they "bunch" on one side or the other of a kink point, producing a discontinuity in the overall distribution. In short, a rate scheme may contain kinks and agents operating within that scheme exhibit bunching behavior.

These explanations do not, however, readily apply to grading schemes like those used for restaurant hygiene inspections. Here, it is the health inspector – not the restaurant's agents – that determines the score and assigned grade. Of course, agents still make choices; restaurateurs operate the eateries with varying degrees of attention to cleanliness. But, the regulatory official can – and often does – exercise discretion in assessing agents and assigning scores. Most research to date on discontinuities has not considered directly such active and ongoing involvement of officials. As such, it says little about how the social relationships between officials and agents might affect the grading process.

We adopt a different and novel approach to studying discontinuities in public grading schemes. Specifically, we consider the social relationships that likely develop between the regulatory officials who assign grades and the agents who are graded—and how these might affect scoring, especially at grade boundaries. We theorize about this relationship in terms of the strength of the social tie between the two parties. Granovetter (1973: 1361) offered an early definition of tie strength: "a combination of the amount of time, the emotional intensity, the intimacy (mutual confiding), and reciprocal services…" Various approaches have been adopted in empirical investigations (Blumstein and Kollock 1988, Marsden and Campbell 1984), but most rely on two observable aspects of the relationship: its duration and the number of interactions.

Various theoretical perspectives seem to point in different directions about how ongoing social relationships between regulatory officials and agents might affect grading processes and outcomes (Aven et al. 2018). On the one hand, macro-oriented theories of regulatory control suggest that stronger ties might be beneficial, reducing information asymmetry and allowing officials to conduct more thorough inspections (Baron and Besanko 1984, Macher et al. 2011). On the other hand, micro-oriented theories of auditor independence suggest that stronger ties might prove detrimental, resulting in biased or lenient assessments (Bazerman et al. 1997, Moore et al. 2006). We suggest that strong ties between officials and agents will prove most problematic in potentially impactful situations where minor score changes can lead to large differences in public perception and consumer behavior.

More specifically, we suggest that the effects of strong ties between officials and agents operating within the context of a public grading scheme will be particularly relevant when inspection scores are in the neighborhood of the grade category boundary. Diners and other relevant publics apparently find simple grades much more salient than scores and adjust their behavior accordingly (Ho 2012). Variations in scores in the middle of a grade category are thus not impactful; any minor adjustments up or down would not affect the final grade. Variations near grade category boundaries, however, are potentially more impactful; a minor difference might change the final grade and, therefore, carry significant consequences for the agent. In such cases, several scenarios are possible but a clearly plausible one sees officials as increasingly reluctant to assign the lower of two grades when the official and agent have a strong social relationship (Schelling 1968, Jenni and Loewenstein 1997). Accordingly, we expect that as the strength of the social tie between an official and agent increases, inspection scores will be more likely to be artificially inflated upwards to the next highest grade category. In addition, we expect that where artificially inflated grades appear, subsequent consumer complaints will be more likely.

Methods

Study Context and Data

The restaurant grading system in Los Angeles County was created nearly twenty years ago following a local investigative news report that exposed the squalid kitchen conditions of some of the city's trendiest – and presumably clean – eateries. The Los Angeles Department of Public Health (LADPH) subsequently created a system intended to increase and maintain compliance by restaurants: public health grades. Each inspection resulted not only in a numerical score but also a letter grade akin to school grades (i.e., A = 90-100 points, B = 80-89 points, C = 70-79 points). Each restaurant was required to post prominently near the front door of its facility a placard indicating the assigned grade. Details of the inspection process have changed in subtle ways over the years, but this system remains in place today. This particular study context was chosen because it is viewed by many as the vanguard of public grading schemes and has been modeled by many other health departments around the world.

The system creates a situation whereby an ongoing social relationship likely develops between inspectors and restaurants. Each inspector is assigned to a district for a period of two to three years during which the inspector is given a caseload of particular restaurants; each restaurant is to be inspected two to three times per year depending on the type of restaurant. (The frequency varies some in practice). Inspections occur without advance warning except in the most general sense in that restaurateurs know that an inspector will visit every four to six months, but they do not know exactly when. After two- to three-year assignment period, the inspector is assigned a new caseload, typically in a new district.

Data came from inspections conducted in full-service restaurants between July 1, 2000, and June 30, 2010.[1] We retained in the dataset inspections of all restaurants that were inspected at least four times during the study period; more details below. This full dataset includes 512,313 inspections of 29,666 restaurants conducted by 606 inspectors.

Measures

The main outcome of interest was whether or not a restaurant received a grade that was likely inflated from the upper bound of one category (e.g., 89; grade B) to the lower bound of the next category (e.g., 90; grade A). Given the distribution of grades (A=80.27%; B=17.40%; C=2.08%; F=.24%), hypothesis tests were conducted on a subsample of observations that were near the A/B category boundary and therefore in question of being (or having been) inflated or not. This "risk set" was based on the actual score recorded at the end of the inspection and included scores ranging from 88 to 91 (i.e., two points on either side of the grade category boundary of interest.) Within this risk set, the grade inflation variable is equal to 1 if the actual score is 90 or 91 (i.e., it was inflated to the A grade category), and 0 if the actual score is 88 or 89 (i.e., it was not). See the Supplemental Material (SM) for robustness checks.

The primary independent variable measures the strength of the relationship between the inspector and the restaurant. Following prior research, we assess tie strength in two ways: (1) *inspector number of visits* as the number of inspections conducted by the inspector at the focal restaurant; and (2) *inspector duration* as the length of time in years since the first visit of the inspector to the focal restaurant.

Various control variables were included: (1) *restaurant size* as a factor variable based on the number of seats in the dining room (the omitted baseline is 0-10 seats); (2) *limited food prep* as a dichotomous variable, where 1 indicates that the restaurant does not prepare food multiple days in advance (i.e., generally a fast food restaurant; inspected twice per year) and 0 indicates that it does (i.e., inspected three times per year); (3) *inspector tenure* as the number of days since the inspector's first inspection in the database (left-censored at 2000); (4) *inspector experience* as the total number of inspections conducted by the inspector in the database (left-censored at 2000); (5) *inspector caseload today* as the count of inspections conducted on the same day by the same inspector as the focal inspection; (6) *previous grade* as a factor variable (the omitted baseline is an *F grade*); (7) *previously suspended* as a dichotomous variable, where 1 indicates that the previous inspection resulted in the restaurant having its health license temporarily suspended; and, (8) *type of inspection* as a factor variable (*routine + complaint* indicates that a consumer complaint was filed, and the inspection serves as a follow-up to that complaint and a routine inspection; *owner-initiated* indicates that the restaurateur requested an off-cycle inspection – typically due to dissatisfaction with the previously assigned grade– and has paid a fee to the LADPH for a

[1] Data was not available for inspections prior to July 1, 2000. Policy and data changes made it most appropriate to use as the end date June 30, 2010.

re-inspection; *department follow-up* indicates that the previous inspection was owner-initiated and, as such, a re-inspection is conducted 30-60 days later; the omitted baseline is a *routine* inspection that follows the standard four- to six-month cycle). Table S1 (SM) shows the descriptive statistics and correlations.

Data Analysis

The analytical approach involved two steps. First, we calculated a predicted score for each inspection from previous inspections of the same restaurant (Table S2). This prediction model included the scores and inspection types for each of the three most recent inspections of the same restaurant. Year fixed effects were included because scores and grades increase over time as a general trend. Inspector fixed effects were also included because inspectors tend to vary in terms of overall leniency. The predicted scores explain approximately 50% of the variance in the observed scores. Because the prediction model includes the previous three scores, data used for the estimation models only includes restaurants that were inspected at least four times in the database and do not include the first three inspections for each restaurant. The estimation models use data from 426,763 inspections of 26,724 restaurants conducted by 557 inspectors. Second, controlling for the predicted score, we examined the effects of the inspector–restaurant relationship on the likelihood of grade inflation. We focused on a subsample of observations that were near the A/B category boundary and therefore in question of being (or having been) inflated or not. As noted above, this "risk set" was based on the actual score recorded at the end of the inspection and included scores ranging from 88 to 91. Using this risk set, we examined the effects of tie strength on the likelihood that the score assigned to a restaurant is inflated to the next highest grade category (i.e., it received an A rather than a B).

Results

Table 1 shows logistic regressions for the likelihood of grade inflation. Model 1 is a baseline model that includes control variables only. Models 2 and 3 show that, as hypothesized, a restaurant's score is more likely to be inflated upwards to the next highest grade category as the strength of the relationship between the inspector and the restaurant increases; this effect holds when tie strength is measured as the number of visits and as the duration of the relationship in years.

(Table 1 here)

Figure 2 plots the predicted probability of grade inflation as a function of the number of visits to the restaurant by the inspector. As shown, the probability of inflation increases approximately linearly with each visit by the same inspector to the restaurant.

(Figure 2 here)

Figure 3 plots the predicted probability that the hygiene score drops when the inspector assigned to a restaurant changes. On the set of 204,063 cases when a new inspector visited a restaurant, we

estimated how the length of the relationship between the *previous* inspector and the restaurant affects the difference between the new score by the new inspector and the last score by the previous inspector. In the model, we included inspector fixed effects as well as a factor variable for each number of visits conducted by the previous inspector prior to the change in caseload (Table S3). As expected, the longer the relationship between the previous inspector and the restaurant, the more the score will decrease upon arrival of a new inspector.

(Figure 3 here)

Placebo tests were conducted to examine whether the effects of tie strength were dependent on proximity to grade category boundaries. Models 2 and 3 were re-estimated with kink points in the middle of the A and B grade categories (Table S4), specifically at 96 (i.e., 94 and 95 vs. 96 and 97) and at 86 (i.e., 84 and 85 vs. 86 and 87). As expected, the effects of tie strength were not significant in these models. These non-results can be taken as evidence that tie strength does not produce leniency in general but, instead, inflationary behavior that occurs in the neighborhood of grade category boundaries.

Table 2 shows logistic regressions for the likelihood of a consumer complaint filed against the restaurant. The dependent variable is equal to 1 if the inspection is the result of a complaint and 0 if not. Models 4 and 5 are estimated using the full sample of inspections and control for the assigned hygiene score in the focal inspection.[2] As expected, a consumer complaint is more likely to occur the stronger the social tie between the inspector and the restaurateur. These findings suggest that diners are more likely to find fault in places with inflated grades, and that the observed discontinuities are not the result of learning or effort on behalf of the restaurant. Of course, these findings do not imply that outbreaks of foodborne illnesses or other such outcomes have occurred. Rather, these findings simply suggest that seemingly unhygienic practices are sometimes noticed by diners even if glossed over by inspectors.

(Table 2 here)

We checked the robustness of the findings in multiple ways. First, we considered alternative windows for the risk set in Models 2 and 3, as narrow as 89 to 90 and as wide as 85 to 94. Second, we used the predicted score rather than the actual score as the basis for the risk set; here, too, we considered various windows for the size of the risk set. Third, we examined the B/C grade boundary. Fourth, because the LADPH electronic archives only date back to the 2000 fiscal year, we considered if left-censoring may influence the results by estimating alternative models that restricted the data to 2001-2010, 2002-

[2] For most complaints, the inspector will conduct an investigation and a routine inspection at the same time. However, the LADPH database also includes some complaint-based inspections that are not gradable. Such instances occur when a complaint is filed but not enough time has passed since the previous inspection so as to justify the next inspection cycle. In such cases, the inspector will investigate the allegations and issue a report to the restaurant (including relevant fines or even closure of the restaurant, if appropriate). However, a full inspection is not conducted until the inspector returns for the next routine inspection. The dependent variable is coded as 1 in either case. Alternative models that excludes the non-gradable complaint investigations result in the same findings as what is reported here; findings are available from the authors upon request.

2010, and 2003-2010. Fifth, we estimated the models in a linear regression framework with inspector fixed effects. Sixth, in order to check for potential issues with the skewness of the tie strength variables, we re-estimated the models using logged versions of the tie strength variables. Seventh, we re-estimated the models using two alternative model specifications for the prediction model, by including restaurant fixed effects and dummy variables for the letter grades of the previous three inspections. In each case, the findings generally agree with what is presented above (Tables S5-S11). Finally, we investigated whether the effect of tie strength differs by restaurant size and by whether the inspector interacts with the owner as opposed to a manager or other employee; no significant differences were detected (Table S12).

Discussion

This study examined how restaurant hygiene inspection scores and grades are systematically shaped by evolving social relationships between inspectors and restaurants. The findings show that inspection scores just below grade category boundaries are more likely to be inflated to the extent that the social tie between the inspector and the restaurant is stronger. When aggregated, this behavior produces discontinuities – or kinks – in the distribution of scores occurring around the grade boundaries. Importantly, the behavior generates public costs: restaurants strongly tied to inspectors are subsequently more likely to receive a consumer complaint filed with the health department.

These findings offer a new perspective on discontinuities at kink points. With rare exception (c.f., Diamond and Persson 2016), prior research has assumed that such discontinuities are the result of direct choices by agents (Kleven 2016). While we agree that agents responding to incentive schemes in intendedly rational ways often produce discontinuities, we also suggest that these patterns may be the result of other processes, including social ones. Differences in possible managerial remedies follow: rather than tinker with the incentives embedded in a rate scheme, problematic social relationships might be corrected by changing the assignments and staffing patterns of regulatory officials.

Targeted transparency initiatives appear in various domains of social life (Weil et al. 2013). From these and other findings, we see a possible trade-off inherent in public grading schemes. On the one hand, introducing categorical grades simplifies communication and comparison of entities in a format that is easier to comprehend than numerical scores, thus facilitating the usage and impact of information by consumers. On the other hand, categorical boundaries may introduce unintended consequences such as grade inflation. Along with others (e.g., Ho et al. 2017), we suggest that the verdict is still out on whether the benefits outweigh the potential drawbacks.

Of course, the obvious theoretical question the findings raise concerns how and why do strong social ties generate grade inflation. Unfortunately, the data we currently have do not allow us to identify exactly the specific underlying mechanism(s) through which inflation operates. Instead, we theorized at a more macroscopic level that inspectors will exercise discretion and leniency in the proximity of grade

category boundaries, an expectation consistent with most plausible mechanisms and not yet documented until this study. Accordingly, we regard the findings linking tie strength and grade inflation as an important first step in uncovering mechanisms, one that demonstrates there is something social worth studying here in greater depth. We imagine that any compelling test among possible mechanisms might require an experimental design as well as archival analysis. In any event, while the problem may be difficult, it merits further attention given its theoretical intrigue and its obvious policy importance.

References

Allen EJ, Dechow PM, Pope DG, Wu G (2017) Reference-dependent preferences: Evidence from marathon runners. *Mgmt. Sci.* 63: 1657-1672.

Aven B, Morse L, Iorio A. (2018) The valley of trust: The effect of relational strength on monitoring quality. Working paper, Carnegie Mellon University.

Baron DP, Besanko D (1984) Regulation, asymmetric information, and auditing. *Rand J. Econ.* 15: 447-470.

Bazerman MH, Morgan KP, Loewenstein GF (1997) The impossibility of auditor independence. *Sloan Mgmt. Rev.* 38: 89-94.

Blumstein P, Kollock P (988) Personal relationships. *Ann. Rev. Sociology* 14: 467-490.

Burtless G, Hausman JA (1978) The effect of taxation on labor supply: Evaluating the Gary negative income tax experiment. *J. Political Econ.* 86: 1103-1130.

Camacho A, Conover E (2011) Manipulation of social program eligibility. *American Econ. J.: Econ. Policy* 3: 41-65.

DeFusco AA, Paciorek A (2017) The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit. *American Econ. J.: Econ. Policy* 9: 210-240.

Diamond R, Persson P (2016) The long-term consequences of teacher discretion in grading of high-stakes tests. Working paper, Stanford University.

Fung A, Graham M, Weil D (1999) *Full Disclosure: The Perils and Promise of Transparency*. Cambridge University Press.

Granovetter MS (1973) The strength of weak ties. *American J. Sociology* 78: 1360-1380.

Ho DE (2012) Fudging the nudge: Information disclosure and restaurant grading. *Yale Law J.* 122: 574-689.

Ho DE, Ashwood ZC, Handan-Nader C (2017) The false promise of simple information disclosure: New evidence on restaurant hygiene grading. Working paper, Stanford University.

Ito K (2014) Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Econ. Rev.* 104: 537-563.

Jenni K, Loewenstein G (1997) Explaining the identifiable victim effect. *J. Risk and Uncertainty* 14: 235-257.

Jin GZ, Leslie P (2009) Reputational incentives for restaurant hygiene. *American Econ. J: Microeconomics* 1: 237-267.

Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47: 263-291.

Kleven HJ (2016) Bunching. *Annual Rev. Econ.* 8: 435-464.

LaLumia S (2008) The effects of joint taxation of married couples on labor supply and non-wage income. *J. Public Econ.* 92: 1698-1719.

Lehman DW, Kovács B, Carroll GR (2014) Conflicting social codes and organizations: Hygiene and authenticity in consumer evaluations of restaurants. *Mgmt. Sci.* 60: 2602-2617.

Macher JT, Mayo JW, Nickerson JA (2011) Regulator heterogeneity and endogenous efforts to close the information asymmetry gap. *J. Law and Econ.* 54: 25-54.

Marsden PV, Campbell KE (1984) Measuring tie strength. *Social Forces* 63: 482-501.

Moore DA, Tetlock PE, Tanlu L, Bazerman MH (2006) Conflicts of interest and the case of auditor independence: Moral seduction and strategic issue cycling. *Acad. Mgmt. Rev.* 31: 10-29.

Pierce L, Snyder J (2012) Discretion and manipulation by experts: Evidence from a vehicle emissions policy change. *B.E. J. Econ. Analysis & Policy* 12: 1-30.

Purdum TS (1988) 15 more arrested in restaurant inspection bribes. *New York Times*, July 8.

Saez E (2010) Do taxpayers bunch at kink points? *American Econ. J.: Econ. Policy* 2: 180–212.

Sallee JM, Slemrod J (2012) Car notches: Strategic automaker responses to fuel economy policy. *J. Public Econ.* 96: 981-999.

Schelling TC (1968) The life you save might be your own. In S Chase (ed.) *Problems in Public Expenditure Analysis*. The Brookings Institute.

Simon PA et al. (2005) Impact of restaurant hygiene grade cards on foodborne-disease hospitalizations in Los Angeles County. *J. Environ. Health* 67: 32.

Sunstein C (1999) Informational regulation and informational standing: *Akins* and beyond. *U. Penn. Law Rev.* 147: 613-675.

Weil D, Graham M, Fung A (2013) Targeting transparency. *Science* 340: 1410-1411.

Table 1
Effects of Inspector-Restaurant Tie Strength on the Likelihood of Grade Inflation

	Model 1	Model 2	Model 3
Inspector No. Visits to Restaurant		0.0638***	
		(0.0114)	
Inspector Experience (Total Inspections)		0.0002***	
		(0.0000)	
Inspector Duration w/ Restaurant (in Years)			0.0543**
			(0.0222)
Inspector Tenure (in Years)			0.0699***
			(0.0149)
Predicted Score	0.1255***	0.1187***	0.1186***
	(0.0089)	(0.0090)	(0.0090)
Restaurant Size			
11-30 Seats	0.2057***	0.2068***	0.2065***
	(0.0414)	(0.0413)	(0.0412)
31-60 Seats	0.3484***	0.3462***	0.3461***
	(0.0446)	(0.0444)	(0.0444)
61-100 Seats	0.5153***	0.5143***	0.5135***
	(0.0560)	(0.0559)	(0.0557)
101-150 Seats	0.5594***	0.5601***	0.5582***
	(0.0711)	(0.0717)	(0.0713)
151-200 Seats	0.7841***	0.7799***	0.7800***
	(0.0981)	(0.0983)	(0.0982)
201-400 Seats	0.8218***	0.8200***	0.8193***
	(0.1021)	(0.1025)	(0.1023)
400+ Seats	0.9916***	0.9902***	0.9969***
	(0.2093)	(0.2092)	(0.2093)
Limited Food Preparation (0/1)	-0.1949***	-0.1691***	-0.1918***
	(0.0400)	(0.0397)	(0.0399)
Inspector Caseload Today	0.0479***	0.0482***	0.0524***
	(0.0122)	(0.0121)	(0.0121)
Previously Suspended (0/1)	0.0873	0.0515	0.0594
	(0.0839)	(0.0842)	(0.0842)
Previous Grade			
C	-0.2971	-0.3125	-0.3003
	(0.2627)	(0.2635)	(0.2625)
B	-0.0041	-0.0114	0.0040
	(0.2644)	(0.2664)	(0.2647)
A	1.2884***	1.2872***	1.3018***
	(0.2673)	(0.2694)	(0.2676)
Type of Inspection			
Routine + Complaint	-0.3098***	-0.3395***	-0.3303***
	(0.0685)	(0.0684)	(0.0681)
Owner-Initiated	1.3229***	1.2952***	1.3273***
	(0.0811)	(0.0808)	(0.0809)
Department Follow-Up	0.4816***	0.4264***	0.4811***
	(0.0775)	(0.0779)	(0.0777)
Constant	-9.7464***	-9.4340***	-9.3297***
	(0.7949)	(0.7954)	(0.7954)
No. Observations	125,811	125,811	125,811
No. Inspectors	532	532	532
Log-likelihood	-22276	-22181	-22229

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$

Notes: Logistic regression with inspector random effects; standard errors (in parentheses) are clustered on the inspectors.

Table 2
Effects of Inspector-Restaurant Tie Strength on the Likelihood of Diner Complaints

	Model 4	Model 5
Inspector No. Visits to Restaurant	0.0107**	
	(0.0046)	
Inspector Duration w/ Restaurant (in Years)		0.0342***
		(0.0080)
Assigned Score (Previous Inspection)	-0.0035**	-0.0036**
	(0.0016)	(0.0016)
Restaurant Size		
11-30 Seats	0.2070***	0.2073***
	(0.0322)	(0.0322)
31-60 Seats	0.5299***	0.5296***
	(0.0336)	(0.0336)
61-100 Seats	0.8205***	0.8197***
	(0.0360)	(0.0360)
101-150 Seats	0.9641***	0.9630***
	(0.0425)	(0.0425)
151-200 Seats	1.0293***	1.0281***
	(0.0515)	(0.0515)
201-400 Seats	1.2151***	1.2140***
	(0.0505)	(0.0505)
400+ Seats	0.7602***	0.7597***
	(0.0954)	(0.0954)
Limited Food Preparation (0/1)	0.3759***	0.3727***
	(0.0236)	(0.0235)
Type of Inspection		
Routine + Complaint	0.3916***	0.3908***
	(0.0328)	(0.0328)
Owner-Initiated	-0.9311***	-0.9225***
	(0.0564)	(0.0564)
Department Follow-Up	0.4316***	0.4432***
	(0.0332)	(0.0331)
Constant	-3.7423***	-3.7261***
	(0.1458)	(0.1459)
No. Observations	426,763	426,763
No. Inspectors	26,724	26,724
Log-likelihood	-70079	-70073

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$

Notes: Logistic regression with restaurant random effects; standard errors (in parentheses) are clustered on the restaurants.