

NBER WORKING PAPER SERIES

INCOMPLETE DISCLOSURE:
EVIDENCE OF SIGNALING AND COUNTERSIGNALING

Benjamin B. Bederson
Ginger Zhe Jin
Phillip Leslie
Alexander J. Quinn
Ben Zou

Working Paper 22710
<http://www.nber.org/papers/w22710>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2016, Revised February 2018

We thank three anonymous referees. We are grateful to Yang Yue for excellent research assistance; to Elliot Anenberg, Liad Wagman, and attendees at the 2014 DC Industrial Organization Day and the 12th Annual International Conference of Industrial Organization for constructive comments; and to a number of Maricopa County government officials for patiently answering our questions over the phone. All remaining errors are ours. We are grateful to the Sloan Foundation for financial support. Part of the paper was revised when Jin is on leave at the Federal Trade Commission. The opinions expressed here are those of the authors and not necessarily those of the Federal Trade Commission, any of its Commissioners, or the National Bureau of Economic Research. All errors are ours.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Benjamin B. Bederson, Ginger Zhe Jin, Phillip Leslie, Alexander J. Quinn, and Ben Zou. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Incomplete Disclosure: Evidence of Signaling and Countersignaling
Benjamin B. Bederson, Ginger Zhe Jin, Phillip Leslie, Alexander J. Quinn, and Ben Zou
NBER Working Paper No. 22710
October 2016
JEL No. D82,H75,I18,L15,L81

ABSTRACT

In 2011, Maricopa County adopted voluntary restaurant hygiene grade cards (A, B, C, D). Using inspections results between 2007 and 2013, we show that only 58 percent of the subsequent inspections led to online grade posting. Although the disclosure rate in general declines with inspection outcome, higher-quality A restaurants are less likely to disclose than lower-quality As. After examining potential explanations, we believe the observed pattern is best explained by a mixture of signaling and countersignaling: the better A restaurants use nondisclosure as a countersignal, while worse As and better Bs use disclosure to stand out from the other restaurants.

Benjamin B. Bederson
Department of Computer Science
University of Maryland
College Park, MD 20742
bederson@cs.umd.edu

Alexander J. Quinn
School of Electrical & Computer Engineering
Purdue University
West Lafayette, IN 47907
aq@purdue.edu

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

Ben Zou
Michigan State University
Department of Economics
East Lansing, MI 48824
benzou@msu.edu

Phillip Leslie
Anderson School of Management
UCLA
Box 951481
Los Angeles, CA 90095
phleslie@amazon.com

JEL: D82, H75, I18, L15, L81

Keywords: quality disclosure, unraveling, voluntary disclosure, signaling, countersignaling, restaurant hygiene.

I. Introduction

This paper documents and explains an incomplete and non-monotonic disclosure of product quality. Theorists have argued that all firms should have incentives to voluntarily disclose their product quality if disclosure is costless and truthful via a third-party verification agency (Milgrom, 1981; Grossman, 1981; Jovanovic, 1982).¹ When disclosure involves a positive cost, disclosure incentive is predicted to increase with quality: all firms with quality above a certain threshold should choose to disclose; and all firms below the threshold should remain silent (Jovanovic, 1982). In contrast to these predictions, we observe an empirical setting where disclosure is neither complete nor monotonic.

The empirical setting is Maricopa County in Arizona. Every restaurant in Maricopa is subject to unannounced routine food safety inspections twice per year. In October 2011, Maricopa adopted a voluntary letter grading system. Under this new system, an inspector arrives for an unscheduled inspection and asks the restaurant manager whether she would like to receive a letter grade (*A*, *B*, *C*, or *D*) and allow the letter grade to be posted online after the inspection. If the answer is yes, the inspector will conduct the inspection, assign a letter grade according to the inspection outcome, post the letter grade on the county official website, and give the restaurant a physical report card with the letter grade. Whether, when and where to post the

¹This is because the incentive to disclose is driven by consumers holding the pessimistic belief that any non-disclosing firm must have the worst quality. Assuming that consumers are willing to pay for higher quality and there is no cost to disclose, this pessimistic belief alone will motivate all but the worst type of firms to disclose, in both monopoly and competitive markets.

physical grade card is up to the restaurant. What's special about Maricopa's grading system is that the detailed inspection results, along with the detailed metric for calculating letter grades, are always available on the county website. Therefore, although the disclosure is voluntary, the public can still impute the letter grade even if the restaurant chooses not to disclose. This provides a unique opportunity to study voluntary disclosure in an environment close to what classical disclosure theories have assumed.

Inconsistent with the theory, only 58 percent of restaurants in Maricopa chose to disclose, and this percentage declined slightly during the first 18 months after the introduction of the new grade card policy. As we expect, *A* restaurants are more likely to disclose than *Bs*, and *Bs* are more likely to disclose than *Cs* and *Ds*. Nevertheless, 49 percent of non-disclosing restaurants would have obtained a grade *A* if they had chosen to disclose. More specifically, if we plot disclosure rate against the total number of violations per inspection, the better *As* are less likely to disclose than worse *As*, although the better *Bs* are more likely to disclose than worse *Bs* (and the *Cs* and *Ds*).

The literature has documented incomplete disclosure in many empirical settings.² But few papers document a pattern of disclosure that is not only incomplete but also non-monotonic with quality (Harbaugh and To, 2016, is an exception). Consistently, many theories attempt to explain incomplete disclosure but still predict a monotonic relationship between disclosure decision and the underlying quality.³ Two notable exceptions are Grubb (2011) and Feltovich, Harbaugh and To

²For example, prior to the 1990 Nutrition Labeling and Education Act, many low-fat salad dressings had a nutrition label, but most of the higher fat dressings did not provide a label, and there were large fat content variations among the non-disclosing ones (Mathios, 2000). Similar incompleteness exists in the disclosure of SUV roll-over risk, financial information of public companies, hazardous substances in the workplace, toxic pollution, medical mistakes, and many other markets (Fung, Graham and Weil, 2007). Typically, disclosure is not complete until the government mandates, or threatens to mandate, disclosure.

³As summarized in Dranove and Jin (2010), theorists often attribute incomplete disclosure to

(2002, referred to as FHT hereafter) . In Grubb (2011), firms fear that today’s disclosure implies a liability to disclose in the future; hence even the highest quality firms may choose non-disclosure due to dynamic concerns. If such dynamic concerns are stronger for higher quality firms, we could observe the disclosure rate decreasing or becoming non-monotonic with quality. FHT (2002) shows that, when the quality information is coarse, sellers of best quality may use non-disclosure as a countersignal to distinguish themselves from eager-to-disclose medium-quality sellers. Following this logic, Harbaugh and To (2016) show an interesting example that economics faculty avoid using titles such as “Doctor” or “Professor” in voicemail greetings and course syllabi.

We test both explanations in our data, as well as the classical disclosure theory and other explanations such as grade uncertainty, consumer inattention, consumer prior knowledge, and competition. We find evidence for both signaling and countersignaling, but not for the dynamic concern as in Grubb (2011). More specifically, the disclosure rate declines monotonically across grades (from *A* to *D*), which is consistent with the classical disclosure theory (with positive and heterogeneous disclosure cost). But within *A* restaurants, those with better records on Yelp and better historical inspection results are less likely to disclose.⁴ Moreover, this pattern only holds for *A* restaurants, not for *B*, *C*, or *D* restaurants. We interpret this as evidence for countersignaling at the highest end of the quality spectrum. The evidence for signaling and countersignaling persists after we account for alternative explanations, although some alternative explanations do have power explaining

assumptions underlying the classical unraveling theory – for example, disclosure may not be costless (Jovanovic, 1982), consumers may not understand the disclosed information (Fishman and Hagerty, 2003), firms may not know the truth (Matthews and Postlewaite, 1985), strategic concerns may hinder disclosure in oligopoly (Board, 2009), and litigation risk may deter disclosure (Marinovic, Skrzypacz and Varas, 2015).

⁴Yelp.com is a popular business listing and consumer review website mainly for restaurants and other retail services.

some variations in disclosure.

The rest of the paper is organized as follows: Section 2 describes the background. Section 3 describes the data and the sample. Section 4 presents evidence for signaling and countersignaling and checks alternative explanations of the observed patterns. Section 5 concludes.

II. Grade Card (GC) Policy in Maricopa County

Maricopa County started to post detailed hygiene inspection records on the county's official website in 2007. Before October 2011, this process did not involve any restaurant decision: records of all routine and follow-up inspections were posted online for all restaurants. Offline, Maricopa followed a star-grading system in which restaurants with hygiene conditions in the top 10 percent were awarded gold stars, the next 10 percent were awarded silver stars, and the other restaurants received no star.⁵ If a restaurant qualified for a gold or silver star, the restaurant received a physical card with a star and could choose to post it anywhere in the restaurant.⁶

On October 14, 2011, Maricopa County adopted a voluntary letter grade system in which a restaurant is assigned a letter grade (*A, B, C* or *D*) at *each* routine inspection. The letter grade is determined by violations of the items specified in the inspection form, which are categorized into Priority (P), Priority Foundation (PF) and Core. P items are those that have a quantifiable measure for controlling hazards in cooking, reheating, cooling, or handwashing. A violation on a P item is a major violation that directly contributes to increasing the risk of food-borne

⁵We could not find any official documents on the exact percentages of restaurants eligible for gold and silver stars. The reported percentages are based on our conversation with an inspector in Maricopa.

⁶It is possible that gold and silver stars were reflected in the online database before the letter grade system. However, because the Maricopa website changed to a new database system in July 2011, we cannot observe online whether a restaurant had a gold or silver star in its historical records.

illness or injury. PF items are those that support, facilitate or enable one or more P items. PF items include those that require the purposeful incorporation of specific actions, equipment, or procedures by industry management, such as personnel training, infrastructure or necessary equipment, hazard analysis and critical control plans (HACCP), documentation or record keeping, and labeling. Core items comprise those items not designated as P or PF; these are usually related to general sanitation, operational controls, standard operating procedures, facilities or structures, equipment design, or general maintenance.

Figure 1 shows how the numbers of P, PF, and Core violations are used to calculate the letter grade. A restaurant receives a letter *A* if the inspection finds no P or PF violations. If a restaurant does not qualify for an *A* but has no more than one P violation or two PF violations, it gets a *B*. A restaurant gets a *C* if it has no more than two P violations and three PF violations. Finally, a restaurant's grade drops one level (*e.g.*, *A* to *B*) if it has four or more Core violations. *D* is the lowest grade possible.

In this new system, posting of the grade card is voluntary both online and offline. The restaurant has to decide, *before* the inspection starts, whether to have the letter grade calculated and posted online. If it agrees to disclose, the letter grade will be posted online. The restaurant also receives a physical card with the letter grade for in-store posting, although the law does not specify where to post the card or whether it must be posted at all. Figure 2 exhibits several screenshots of online records after Maricopa adopted the grade card policy. Clicking on a restaurant's name brings up every inspection for that restaurant since July 2007. An inspection can be routine or follow-up: routine inspections are unscheduled, while follow-ups often focus on whether the restaurant has corrected the violations found in the last routine inspection. The choice of disclosure is only relevant for routine inspections. If the restaurant chooses to disclose the grade corresponding to a routine inspec-

tion, the letter grade is shown next to the inspection date and inspection purpose. If the restaurant chooses not to disclose, the record reports “not participating” in the place of a letter grade. This way, non-disclosure is highlighted and therefore distinguishable from not having been given a choice of disclosure. Regardless of the letter grade, clicking on an inspection record leads to detailed violations found during that inspection. If the inspection was conducted after the adoption of the grade card policy, each violation is labeled “P”, “PF”, or “C” in the middle of a large block of text describing the violation (the original text is not highlighted). Because the grading chart is posted on the same page as the detailed inspection result, an interested consumer still has all the elements needed to calculate the letter grade. Compared with disclosing restaurants, a non-disclosing restaurant simply hides its letter grade behind one more click and some calculations.

Because disclosing restaurants have the choice of posting the letter grade offline, our research assistant picked one random neighborhood in Maricopa and checked out restaurants there. Online Appendix Figure A presents two examples of physical card posting in that neighborhood. Of the 26 restaurants he checked, five did not participate in any disclosure (and therefore were labeled “not participating” in the online database), 14 participated in online disclosure but did not post the letter grade inside the restaurant, and the other seven disclosed both online and offline. All the offline disclosing restaurants had an *A* grade. Six of the 14 online-only disclosing restaurants had a *B* or *C* grade, and the remaining eight in this group had an *A* grade. This neighborhood is clearly not representative of the entire Maricopa County, but it highlights the facts that restaurants have discretion in both online and offline postings, and that their disclosure decision depends on factors beyond the underlying hygiene quality. The two examples in Appendix Figure A also highlight the fact that restaurants have discretion regarding where to post the card inside the restaurant if they choose to post offline. Given the difficulty of observing offline

posting status for every restaurant, our empirical analysis focuses on the decision of online posting only.

Maricopa's new grade system has received mixed reaction in the media. On the one hand, the letter grade system is more salient to consumers than the previous star system, although it is still difficult for consumers to infer the grade of non-disclosing restaurants; on the other hand, critics have expressed concern that the county's intention to maintain a friendly relationship with the industry may encourage dirty restaurants to opt out of posting and have little effect in reducing the public health risk of food-borne illnesses. We believe it is necessary to understand the incentives behind disclosure before drawing any welfare implications from the letter grade policy.

III. Data and Summary Statistics

III.A. Data

Our main data come from the public website of the Maricopa County government.⁷ We scraped the website in March 2013. Our dataset contains inspection results from July 2007 to March 2013 of 23,863 food installations, 19,719 of them with at least one inspection since the adoption of grade cards. For each inspection, we know the date of inspection, type of inspection (routine, follow-up, etc.), description of each violation item from the Food Code, and letter grade or "not participating" label after the adoption of grade cards. Because disclosure is relevant only for routine inspections, we exclude all follow-up inspections. In total, there are 211,627 inspection records, of which 146,498 were from before the grade card policy. Of

⁷All records of inspections since July 2007 can be found on Maricopa County's official website: <https://www.maricopa.gov/EnvSvc/OnlineApplication/EnvironmentalHealth/FoodInspections/Business>. The layout of the website has been changed slightly since we have extracted the data.

the grade card policy records, 37,600 participated in online disclosure, and 27,529 chose “not participating”.

Figure 2 shows one inspection record for Alexis Grill, a restaurant in Phoenix.⁸ It received a grade *B* in an inspection on August 8, 2013 and chose to disclose its letter grade online. Before that, it had five inspections for which it did not participate in grade posting. These were preceded by an inspection on January 31, 2012, the restaurant’s first inspection after Maricopa adopted the grade card system, at which it chose to disclose and got a *C*. The website also contains all inspections prior to the voluntary posting scheme, going back to July 2007, which are listed below. Details of each inspection are available with one click on the date. We show some details of the inspection conducted on August 8, 2013 (shown on the right of the page), and of the inspection conducted on February 16, 2010 (shown on the bottom of the page). Both records have violation #14, corresponding to the cleanliness of food-contact surface. However, the textual explanation of this violation is slightly different and the letter “P” is only added after grade cards were adopted. We can predict the letter grade for both disclosing and non-disclosing restaurants after the introduction of grade cards, but it is difficult to do so for pre-grade-card records. This is because the violations reported in pre-GC inspections do not contain the labels “P”, “PF”, and “C”. Moreover, each violation code may include multiple items in the county’s Food Code, and there is no guarantee that items under the same violation code are always classified in the same P, PF, or C category. In theory, we can use text matching to create a correspondence between pre- and post-grade card violation codes and then define grades before grade card, but we believe the outcome is likely to be noisy, as Maricopa revised its food code when it adopted the grade card system.

⁸The screenshot was taken when we visited the website in 2014. When we scraped the website in March 2013, the two latest inspections were not available.

Given this limitation and our interest in restaurants' disclosure decision, we focus on post-GC inspections. As shown in Online Appendix Figure B, the number of violations per inspection has declined steadily since 2007, from between 1.5 and 2 before grade cards to roughly 1 after grade cards were adopted. The decline is statistically significant, but it is unclear whether the drop is attributable to changes in Food Code, the grade card policy, or both. It does not seem to have any structural break after the grade card policy was adopted.

Table 1 provides a more detailed data summary. For pre-GC inspections, we report the number of total violations per inspection; after grade cards adoption, we report summary statistics on the number of total, P, PF, and Core violations by letter grade and disclosure status. According to Figure 1, P and PF violations are more serious than C violations: while one extra P or PF violation tends to trigger a downgrade in the letter system (the only exception is a *B* to *C* downgrade requires two extra PFs), it takes four extra Core violations to do the same. For simplicity, we construct a single weighted sum of violations per inspection after the adoption of grade cards, counting each P or PF violation as four and each Core violation as one. This continuous variable, referred to as *WSUMVIOL*, is more detailed than a letter grade and allows us to look into the heterogeneity within each letter grade. For example, the best *A* restaurant may have no violations at all; thus its *WSUMVIOL* is equal to 0, while the worst *A* restaurant can have up to three core violations, which implies a *WSUMVIOL* equal to 3. Similarly, the best *B* restaurant has either one P, or one PF violation, or four Core violations, which corresponds to a *WSUMVIOL* equal to 4, while the worst *B* restaurant can have a *WSUMVIOL* as high as 15. One shortcoming of *WSUMVIOL* is that *B* and *C* restaurants may overlap in the range between 8 and 15 because two PF and one P violations correspond to a *B* grade, while two P violations correspond to a *C* grade. However, *WSUMVIOL* is a monotonic proxy of quality within *A* and *B*. Given the fact that over 90 percent

of post-grade card inspections fall in either the *A* or *B* range, and these restaurants are the main focus of our discussion, we use WSUMVIOL to explore overall and within-letter variations.

III.B. Grade Distribution and Disclosure Patterns

Maricopa's grade card system offers a rare opportunity to study voluntary disclosure. In fact, several patterns in the raw data stand out, but they are only partially consistent with the classical disclosure theories.

The first data pattern is incomplete disclosure. Of all the routine inspections conducted after the adoption of the grade card scheme, only 58 percent led to on-line letter grade posting and this percentage declined slightly over time. Clearly, the disclosure rate is much lower than the prediction from the most classical unraveling theory that 100 percent will disclose (Milgrom, 1981; Grossman, 1981). One common explanation is heterogeneous disclosure cost (Jovanovic, 1982): though the disclosure decision does not involve any monetary cost, the restaurant owner/manager may not be on site at the inspection time, employees may have a hard time finding the owner/manager in time, or the owner/manager may find it mentally demanding to make any decision other than the default of non-disclosure. All of these can contribute to a positive disclosure cost and discourage disclosure. It is also possible that a restaurant cannot fully anticipate the inspection outcome beforehand. If the owner/manager is risk averse, the prospect of receiving an imperfect grade may discourage disclosure as well.

In light of the potential grade uncertainty, we examine the relative importance of across- and within-restaurant variations. To what extent do inspection outcomes vary over time within the same restaurant? How often do restaurants jump between grades?

Table 2 shows the transition matrix of grades and disclosure decisions between a restaurant's previous and current inspections (we impute grades for non-disclosing restaurants). Because only a small fraction of restaurants receive *C* or *D*, from now on we group these two letter grades into one group, *C/D*. Disclosure decisions are persistent over time. If a restaurant chose to disclose in the previous inspection, the probability of disclosure in the current inspection is 0.81. If a restaurant chose not to disclose in the previous inspection, the probability of non-disclosure in the current inspection is 0.71. The persistence in the disclosure decision is at least partly driven by the persistence in inspection results: if a restaurant received an *A* in the previous inspection, the probability to getting an *A* in the current inspection is 0.73. If a restaurant received a *B* in the previous inspection, it has a probability of 0.47 of getting an *A* in this inspection and is a lot more likely to receive a *B* than a restaurant that received an *A* in the previous inspection (0.42 versus 0.23). In fact, restaurant fixed effects alone account for 63 percent of the total variation in letter grades and 72 percent of the total variation in disclosure decisions. Table 2 also shows that, although the disclosure decision is made before the inspection result is available, to some extent restaurants can predict what grade they will get and decide accordingly. For a restaurant that received an *A* and chose to disclose in the previous inspection, if it still gets an *A* in the current inspection, its probability of disclosure is 0.87 ($0.65/(0.65 + 0.091)$). In comparison, if it gets a *B* in the current inspection, its probability of disclosure drops to 0.73 ($0.164/(0.164 + 0.061)$). In short, the disclosure patterns described above seem more likely to be driven by across-restaurant variation than by changes within a restaurant. This is the second pattern from the raw data. This pattern is consistent with our assumption that restaurants can anticipate inspection outcomes.

Despite incomplete disclosure, the third data pattern is that disclosure rate varies monotonically *across* letter grades. According to Table 1, 66 percent of *A*, 49 per-

cent of *B*, and 30 percent of *C/D* choose to disclose. This pattern is consistent with the classical disclosure theory: when every restaurant faces some disclosure cost (which could be fixed or random), higher quality firms are more likely to disclose. It is also consistent with signaling, as disclosure alone is a positive signal to consumers when higher-quality firms have more incentives to disclose.

However, within-grade variations depict a rather different picture. As Table 1 shows, restaurants that *would* receive an *A* but choose not to disclose have even better inspection outcomes (WSUMVIOL= 0.31) than those that get an *A* and choose to disclose (WSUMVIOL= 0.34). This is inconsistent with the classical prediction that higher-quality firms should be more likely to disclose. More interestingly, this counterintuitive pattern only holds for *A* restaurants. For *B* and *C/D* restaurants, the disclosing ones on average have better inspection outcomes than the non-disclosing ones (in terms of WSUMVIOL scores, 1.9 versus 2.1 for *B*s, and 4.2 versus 4.6 for *C/D*s).

For a more direct comparison, Table 3 compares the characteristics of disclosing and non-disclosing restaurants within each grade. We focus on each restaurant's last available inspection following the adoption of grade cards. This way, our sample counts each restaurant equally no matter how frequently they have been inspected in the raw data.⁹ The first two columns in Panel A present the mean of inspection outcomes for disclosing *A*s and non-disclosing *A*s. Disclosing *A*s have worse inspection outcomes by almost all measures (except for the number of *P* and *PF* violations, which by definition, are zero for restaurants receiving grade *A*), and these differences are statistically significant (Column 3). This is in sharp contrast to the differences between disclosing *B*s and non-disclosing *B*s, as shown in Columns 4 to 6. The disclosing *B*s have better inspection outcomes than non-disclosing *B*s by most measures, and these differences are statistically significant.

⁹Inspection frequency could vary by restaurants.

One may argue that since the disclosure decision is made before the inspection is conducted, using the current inspection result can be misleading. Panel B replaces current inspection outcomes with the average inspection outcomes before the studied inspection, which are undoubtedly known to the restaurant before the current disclosure decision. The differences in average previous inspection outcomes between disclosing *As* and non-disclosing *As* are similar to those in panel A, as are the difference between disclosing *Bs* and non-disclosing *Bs*. Panel C compares the standard deviations of past inspection outcomes for disclosing and non-disclosing restaurants. Again, the disclosing *As* tend to have greater variation in their past inspection outcomes than non-disclosing *As*, but the disclosing *Bs* tend to have smaller variation in past inspection outcomes than the non-disclosing *Bs*. In short, the three panels of Table 3 reinforce the impression that disclosure patterns within *A* are opposite to the disclosure patterns within *B*.¹⁰

To further look into the heterogeneity within each letter grade, we define half-marks within each letter grade based on WSUMVIOL. For restaurants receiving *A*, we define “strong *A*”, denoted as *A+*, as restaurants with WSUMVIOL scores not higher than the median WSUMVIOL score of all inspections that disclose as *A*. In other words, disclosing or not, the *A+* restaurants are at least as good as the typical restaurant that displays grade *A*. Restaurants with an *A* or an imputed *A* that do not qualify as *A+* are defined as “weak *A*” restaurants and are denoted as *A-*. *B+* / *B-* and *CD+* / *CD-* restaurants are defined in similar ways.¹¹

Figure 3 plots the average disclosure rate for each half-marked letter grade. Consistent with the data summary tables, disclosure rate displays a salient “hump” shape: in general, disclosure rate declines as the inspection result gets worse, but

¹⁰Table 3 tests the mean differences. We also test whether the distributions of disclosing and non-disclosing restaurants within the same grade are the same. It follows the same structure as in Table 3. The mean tests and distribution tests reach similar conclusions.

¹¹The median WSUMVIOL score cutoff is 0 for *A+* / *A-*, 5 for *B+* / *B-*, and 13 for *CD+* / *CD-*.

it increases from $A+$ to $A-$.¹² Both the overall decline in disclosure rate and the hump among A restaurants are substantial in magnitude. Grade A restaurants have an average disclosure rate that is 17 percentage points higher than that of B restaurants (62.6 percent versus 45 percent), while the disclosure rate for C/D restaurants is another 18 percentage points lower. In contrast, the disclosure rate of $A+$ restaurants is 7 percentage points lower than that of $A-$ restaurants.¹³

To summarize, raw data display many interesting patterns: disclosure is incomplete, disclosure difference is largely driven by cross-restaurant variations, disclosure rate declines monotonically across grade, there is a non-monotonic “hump” in disclosure rate between $A+$, $A-$, and $B+$, and the difference between disclosing A s and non-disclosing A s contrasts with the difference between disclosing B s and non-disclosing B s. We now proceed to tie these patterns to possible explanations.

IV. Potential Explanations and Empirical Tests

There are at least six economic explanations that support some but not all of the documented data patterns.

- First, the classical disclosure theory with a positive disclosure and heterogeneous cost may explain the declining disclosure rate from A to C/D but not why $A+$ restaurants are less likely to disclose than $A-$ restaurants. We refer

¹²The unraveling theory predicts that restaurants receiving a letter grade of A or B will all choose to disclose, as they have an incentive to separate from restaurants with C/D grades. We do not observe complete disclosure for A or B restaurants. This is not inconsistent with the unraveling theory. The unraveling theory assumes perfect information to consumers, which is unlikely to hold in this setting given the contrived design of the grade posting system. Restaurants also face uncertainty as they have to commit to disclosure before inspection is conducted. Thus the overall declining disclosure rate is consistent with a generalized version of the unraveling theory.

¹³The overall 66 percent of disclosure rate of A restaurants, as cited in Table 1, is computed from all post-GC inspections. In Figure 3, the 60 percent versus 67 percent comparison between $A+$ and $A-$ restaurants is from restaurants’ last post-GC inspection. Among the last post-GC inspections, the average disclosure rate for all A restaurants is 62.6 percent.

to this classical theory as signaling.

- Second, the grade uncertainty facing risk-averse restaurants can explain why some restaurants choose non-disclosure even if the inspection outcome turns out to be good. However, according to panel C of Table 3, disclosing As face more uncertainty from past inspections than the non-disclosing As, but the opposite is true for disclosing and non-disclosing Bs.
- Similar difficulty applies to the third explanation – dynamic concerns. As argued in Grubb (2011), a firm with excellent product quality in period t may be reluctant to disclose its high quality for fear that such disclosure may commit the firm to future disclosure or raise more doubts when it does not disclose the same amount of information in the future. Following this dynamic concern, the fear of unfavorable results in the future should be higher for firms that face more uncertainty within *each* grade, which we know is true within *B* but not true within *A*.
- The fourth explanation is consumer attention. If consumers do not pay attention to grade cards at all, disclosure is equivalent to non-disclosure for all grades. More realistically, if consumers pay more attention to the grade of some restaurants, these restaurants may have more motivation to disclose their grades. Selective attention could explain why disclosure is incomplete, but it brings up the question of why *A+* restaurants receive less consumer attention than *A-* restaurants.
- Similarly, the fifth explanation lies in consumers' prior information: if consumers already know an *A* restaurant is as clean as *A*, the restaurant may not bother to disclose. Prior information that is non-linear by quality could explain the hump in disclosure rate. But it is unclear why consumer prior

information is particularly weak in $A-$, or why that prior information is weak enough to give $A-$ restaurants an extra incentive to disclose more than both $A+$ and $B+$.

- The sixth explanation is the countersignaling theory laid out by FHT (2002). As detailed below, countersignaling focuses more on the non-monotonic disclosure rate between $A+$, $A-$ and $B+$ than on the other parts of the grade distribution.

Because the countersignaling explanation is the most complicated and it does not rule out the other five explanations, we will organize our empirical test around signaling and countersignaling. In the meantime, we control for variables that measure grade uncertainty, dynamic concerns, consumer attention, and consumer prior information. This way, we allow multiple explanations to coexist but still look for evidence of countersignaling.

In the rest of this section, we first apply FHT's countersignaling theory to our setting and then define our econometric specification and key variables. Our results discussion will start with basic evidence of countersignaling and other explanations and end with extra evidence for countersignaling.

IV.A. The Countersignaling Theory

According to FHT (2002), countersignaling may occur when a high-quality firm finds it desirable to signal high quality through non-disclosure. In a job interview example, FHT assume a setting with three types and two signals. Job candidates can have high, medium, or low ability. They may signal their hidden ability through an endogenous signal (sharing their GPAs) and an exogenous signal (a confidential recommendation letter). Both GPA and the letter are noisy but in a different way. On the one hand, high and medium types have good GPAs while the low type has

bad GPAs. Thus GPAs separate high and medium from low, but not high from medium. On the other hand, a recommendation letter is always good for high-ability candidates, always bad for low-ability candidates, but can be good or bad for the medium-ability candidates. Hence, the letter alone can separate high from low, but not necessarily high from medium or medium from low. These two signals are overlapped such that the combination of high GPA and good recommendation letter is still not enough to distinguish the high type from some lucky mediums.

FHT show that countersignaling can help to separate all three types, under some conditions. In particular, the high type may prefer not to mention their high GPAs, because this countersignaling action plus the favorable recommendation letter may help them stand out from the medium candidates that disclose high GPAs and have a good letter. In comparison, medium-ability high-GPA candidates cannot afford to hide their high GPAs because they do not know the exact content of the letter when they make the disclosure decision and high GPAs will clearly distinguish them from the low-ability candidates even if the letter turns out bad.

It is not difficult to translate this story into our setting, if we assume restaurants can anticipate the inspection outcomes perfectly and we only focus on the high end of the quality spectrum. The high, medium, and low types correspond to our $A+$, $A-$, and $B+$. The high and low GPAs correspond to our grade A and grade B , and disclosing GPA corresponds to disclosing A or B grade. From Yelp.com, we also obtain other measures of restaurant quality as of 2014, and for now, let us assume the restaurant's future Yelp review is the second signal, similar to the recommendation letter.

One key assumption of FHT is that the exogenous signal is noisy, has a random component (for at least the medium type), and its value is not known to the job candidate when he/she makes the disclosure decision. If this signal is known beforehand, the game boils down to the standard disclosure game conditional on the

realized signal. Since Yelp reviews are on-going, we believe it is reasonable to assume that the content of future Yelp reviews is random and $A+$ restaurants may have more confidence in obtaining favorable Yelp reviews than $A-$ restaurants. (This is actually an empirical question; we will test it later in the real data.) By this assumption, the countersignaling equilibrium could occur in our setting, where $A+$ are reluctant to display their A grade, $A-$ are eager to display their A grade, and $B+$ do not bother to display their B grade because $B+$ is the worst in our three-type world.

This stylized story does not explain why more than 40 percent of $B+$ restaurants do disclose their grade. This is because the above setting assumes away any types below $B+$. Adding in the lower types gives $B+$ restaurants an incentive to show that they are better than C or D . Whether there is another countersignaling incentive among B restaurants will depend on the nature of the exogenous signal and the number of types in the whole distribution (FHT 2002). Because our data summary only suggests countersignaling among A s, we believe it is reasonable to use FHT (2002) to focus on the high end only.

In short, the countersignaling theory of FHT could explain the hump in disclosure rate between $A+$, $A-$, and $B+$, but it alone does not explain the overall declining disclosure rate from A to B to C/D . FHT also implies that evidence of countersignaling can be found in the random nature of the exogenous signal and how the distribution of that signal overlaps with half grades. We will come back to test this in Section 4.4.

IV.B. Econometric Specification and Key Variables

To examine disclosure rate as a function of inspection results, we run a few versions of the following regression:

$$\begin{aligned}
 (1) \quad DIS_{it} = & \alpha_0 + \beta_{A+} \cdot \mathbf{1}\{\text{grade}=A+\}_{it} + \beta_{A-} \cdot \mathbf{1}\{\text{grade}=A-\}_{it} \\
 & + \beta_{B+} \cdot \mathbf{1}\{\text{grade}=B+\}_{it} + \beta_{B-} \cdot \mathbf{1}\{\text{grade}=B-\}_{it} \\
 & + \beta_{CD+} \cdot \mathbf{1}\{\text{grade}=CD+\}_{it} + \mathbf{X}_i \cdot \alpha_x + \mathbf{Z}_{it-1} \cdot \alpha_z \\
 & + \lambda_t + \varepsilon_{it}.
 \end{aligned}$$

The unit of observation is a restaurant’s last observed post-grade card inspection. We limit each restaurant i to only one observation in the regression because we know most of the variations in disclosure are driven by cross-restaurant rather than within-restaurant variations. Focusing on the last observed inspection allows the market to settle in a relatively stable situation after Maricopa County rolled out its grade card policy. That being said, we index t as the year-quarter of the studied inspection so that we can define that restaurant’s historical variables up to t .¹⁴

DIS_{it} is a dummy variable equal to 1 if the restaurant chooses to disclose. We include a set of dummy variables for half grades; for example, $\mathbf{1}\{\text{grade}=A+\}_{it}$ is an indicator variable that takes value 1 if the outcome of the inspection qualifies for the $A+$ category. $CD-$ is the left-out category; thus the coefficients associated with the half grade dummies are interpreted as the increased probability of disclosure for that grade relative to $CD-$ restaurants. \mathbf{X}_i indicates restaurant i ’s time-invariant characteristics such as chain status and characteristics of surrounding area.¹⁵ \mathbf{Z}_{it-1}

¹⁴We have tried the same regression in the full post-grade card sample, while allowing lagged variables to compute inspection history. Results are quantitatively similar to what is reported in this paper.

¹⁵In all regressions we include whether the restaurant is listed on Yelp.com, and whether it belongs to a restaurant chain. Results are the same without controlling for these covariates.

is a vector of restaurant i 's past inspection outcomes. λ_t is a full set of year-quarter dummies. ε_{it} is the error term. We are interested in β s and how they vary by half grades. In particular, we test the null hypotheses of $\beta_{A+} > \beta_{A-}$ and $\beta_{B+} > \beta_{B-}$. If signaling is the only story, neither null hypotheses would be rejected. If countersignaling prevails for A restaurants and signaling explains the disclosing behavior of other restaurants, we would expect $\beta_{A+} < \beta_{A-}$ but $\beta_{B+} > \beta_{B-}$. For easy coefficient interpretation across different versions, we report all regressions from the linear probability model. Results are similar when we repeat everything in probit or logit.

As stated before, we organize the empirical tests around countersignaling, but that does not rule out alternative explanations. Rather, we include variables suggested by the other explanations and let the data speak to their validity.

Consumer attention

One explanation for incomplete disclosure is *consumer inattention*. If consumers do not pay attention to the letter grade at all, disclosure is equivalent to non-disclosure for all grades. This hypothesis does not speak to why there is a hump in disclosure rate, but is still worth investigating. A plausible scenario is that consumer attention to grade cards varies over time and restaurants are more likely to disclose when consumers pay more attention to the grade card policy. If the timing of inspections and restaurant quality are somewhat correlated (for example, restaurants with poor hygiene may be inspected more frequently), consumer inattention could explain variations in disclosure rate.

From LexisNexis we searched for local news about Maricopa's grade card policy and could find a few news articles only around the time that Maricopa County introduced the policy (October 2011). Therefore, it is possible that consumer attention was heightened at the time of policy introduction. To account for overall changes in media coverage, we include a full set of year-quarter dummies in all

specifications of Equation 1.

On the other hand, the physical presence of grade cards (posted on the door, on the window, or inside a restaurant) may also remind consumers of the grade card policy. Even if the local news stopped covering grade cards soon after its introduction, consumer attention may linger over time and vary from one local area to another depending on the disclosure rate in each area. To account for the effects of history dependence, we classify restaurants into two groups: those that were examined in the first 10 percent of inspections after grade card policy adoption (which corresponds to the first three months after the introduction of grade cards) are referred to as the “first-batch”; the rest are “non-first-batch.” Among the 19,719 restaurants in our sample, about 30 percent are first-batch. In their last observed inspections, 58 percent of the first-batch restaurants chose to disclose, while only 52 percent of the non-first-batch restaurants disclosed. This difference could be driven by more consumer attention when first-batch restaurants were inspected for the first time after grade card adoption, combined with restaurants that disclosed before being more likely to disclose again. A dummy of first batch is included in all regressions.

Other restaurant information available to consumers

We use extra information extracted from Yelp.com, a popular online business listing and consumer review website. As shown in Kang et al. (2013), a Yelp review is correlated with restaurant inspection outcomes, but not all restaurants are reviewed by Yelp, and most Yelp reviews focus on restaurant food and service rather than hygiene. We obtain data on restaurants listed on Yelp from the 2014 Yelp Open Dataset Challenge.¹⁶ This dataset includes all restaurants in the Phoenix area that are listed on Yelp and had at least three reviews. We match restaurants in the inspec-

¹⁶The dataset is available at http://www.yelp.com/dataset_challenge/, last checked on June 5, 2015.

tion dataset with restaurants in the Yelp dataset by name and address. We were able to match 42 percent of restaurants in the inspection dataset. Among restaurants that can be matched with the Yelp data, we define three measures of quality. A restaurant is “popular” if it has more than 60 reviews (the median number of reviews in the matched sample). A restaurant has “good reviews” if it has 4 or more Yelp stars (ranging from 1 to 5 with half star increments). A restaurant is “pricey” if it has 3 or 4 dollar signs on Yelp (ranging from 1 to 4 dollar signs, with 1 dollar sign indicating that the average cost per person is below 10 dollars, 2 dollar signs for a range between 11 and 30 dollars, 3 dollar signs for a range between 31 and 60 dollars, and 4 dollar signs for above 61 dollars). These measures likely represent a better information source for consumers than letter grades. As discussed above, the online posting of actual hygiene inspection outcomes is long, requires multiple clicks, and can be confusing; in comparison, Yelp is a popular website among restaurant-goers, with easy-to-use information at one’s fingertip. On the other hand, there is no reason to ascertain that consumers will perfectly observe all the Yelp variables and draw perfect inference on restaurant quality.

Note that the Yelp variables are subject to multiple interpretations. They could be interpreted as a proxy for consumer prior knowledge before a restaurant decides whether to disclose the grade card. Under this interpretation, the disclosure decision should be conditional on the realized Yelp variables. Alternatively, the restaurant may have difficulty predicting future Yelp reviews, and different confidence in future Yelp reviews could create an incentive to countersignal. Whether we should interpret the Yelp variables as consumer prior knowledge or as an element in the countersignaling story is an empirical question.

Among all the potential hygiene violations, one may wonder whether some are more observable to consumers; bathroom cleanliness is an obvious example. To address this point, we classify 12 violations as consumer observable violations, all

of which are listed in Online Appendix Table B.¹⁷ As Table 3 shows, disclosing As have significantly more observable violations than non-disclosing As, but the opposite is true for B restaurants. If consumer observable violations constitute most of consumer prior knowledge, this does not explain why the comparison between disclosing and non-disclosing restaurants is reversed among As and Bs.

Grade Uncertainty and Risk Aversion

Because the restaurants do not know what their grade will be at the time they make disclosure decisions, one explanation for non-disclosure is *risk aversion* to potentially unfavorable results. We use the standard deviation of WSUMVIOL from past inspections as a proxy for the grade uncertainty facing a restaurant at the beginning of the current inspection. As Table 3 Panel C shows, disclosure decision is negatively correlated with the *mean* of a restaurant's historical WSUMVIOL score. This is not surprising because a restaurant consistently receiving a bad inspection outcome should be more reluctant to disclose, as disclosure is by and large a positive signal of underlying quality. However, if restaurants are risk averse to grade uncertainty, we should find that the *standard deviation* of past WSUMVIOL influences disclosure negatively, conditional on the same mean of past WSUMVIOL. For this reason, in some versions of the regression, we control for both the mean and the standard deviation of WSUMVIOL from past inspections.

Since restaurant quality is highly persistent over time, we would expect that restaurants that receive an A+ in the current inspection are likely to be the very high-quality restaurants in previous inspections as well. Knowing that, they are less likely to disclose in the current inspection if countersignaling is at play. Because the grade of the past inspection is known to the restaurant with certainty and presents no

¹⁷A violation is labeled as unobservable if the structure or procedure in question is behind the kitchen door. For example, insects and rodents are observable to the consumers, while proper cooking time and temperature are unlikely to be observable to consumers.

risk. If we observe a hump shape in *current* disclosure rate based on *past* inspection outcomes, we can argue that the hump shape cannot be driven by uncertainty and risk aversion. To test this hypothesis, we construct half-marked grades based on the second-to-the-last inspections and include them in the regression.

Dynamic Concerns

Dynamic concerns are related to risk aversion. As Grubb (2011) argues, even if an A restaurant knows that it will get an A for sure this time, its owner/manager may be reluctant to disclose the A grade because she is afraid that today's disclosure implies a commitment to disclose next time when the grade is below A. This concern should be more severe for restaurants facing more grade uncertainty, so the standard deviation of past WSUMVIOL can control for this story as well. More realistically, restaurants may choose not to disclose once they already have an A displayed from a previous inspection. Since A+ restaurants are more likely to have received an A in the past, this could explain why A+ restaurants are less likely to disclose later. To account for dynamic concerns, we first include the letter grade of the restaurant's previous inspection. We also include a dummy equal to one if the restaurant's most recent disclosure is A. These are the restaurants that already have an A displayed and are more likely to "play it safe" by not disclosing.

Competition

We also control for competition, although we have not mentioned it as an alternative explanation for disclosure. One may argue that a restaurant's disclosure decision is subject to competition from other restaurants nearby. We attempt to capture this by three variables: the first variable is the number of restaurants in the same ZIP code. Secondly, we use first-batch restaurants to define the fraction of competitors in the same ZIP code that are likely to be pushed to disclose. If consumer attention is heightened at the beginning of the grade card policy, which motivates first-batch restaurants to disclose more, the disclosure of first-batch restaurants may

raise consumer attention and make consumers more suspicious of a non-disclosing restaurant nearby, even if that restaurant itself is not in the first batch. The share of nearby restaurants in the first-batch intends to capture this spatial spillover effect. Thirdly, we compute the average WSUMVIOL of all restaurants in the same ZIP code. If disclosure (or non-disclosure) functions as a signal to stand out from competitors, it should depend on competitors' WSUMVIOL.

IV.C. Basic Evidence of Countersignaling and Other Explanations

Table 4 Column 1 reports the baseline version of Equation 1, including half grade dummies, year-quarter dummies, and dummy variables indicating whether it is a first-batch restaurant, whether it belongs to a restaurant chain, and whether it is listed on Yelp. Consistent with the classical disclosure theory, the overall disclosure rate declines as we move from better inspection outcomes to worse inspection outcomes. However, among *A* restaurants, the disclosure rate of *A+* restaurants is 5 percentage points lower than that of *A-* restaurants, and we can reject $\beta_{A+} > \beta_{A-}$ with more than 99 percent confidence. On the other hand, the disclosure rate for *B+* restaurants is 4.5 percentage points higher than that of the *B-* restaurants, and we cannot statistically reject $\beta_{B+} > \beta_{B-}$. Similarly, the disclosure rate of *CD+* restaurants is about 5 percentage points higher than that of *CD-* restaurants. Across grades, the disclosure rate of *B+* restaurants is 20 percentage points lower than that of *A-* restaurants, and the disclosure rate of *CD+* restaurants is 12 percentage points lower than that of *B-* restaurants. In other words, the disclosure pattern is consistent with signaling except for the very high end, where the non-linear pattern of *A+*, *A-* and *B+* is consistent with countersignaling.

Similar to what we have seen in the data summary, first-batch restaurants have

a disclosing probability roughly 6 percentage points higher than non-first-batch restaurants in their last observed inspections. This is consistent with first-batch restaurants receiving more consumer attention and therefore having more incentives to disclose.

Column 2 adds the mean and standard deviation of the restaurant's past WSUMVIOL. Consistent with the signaling story, the coefficient of the mean of past WSUMVIOL is negative, confirming that restaurants with better past scores are more likely to disclose. In contrast, although grade uncertainty and risk aversion would predict a negative coefficient on the standard deviation of past WSUMVIOL, that coefficient turns out to be slightly positive and not statistically significant at any conventional level, indicating that uncertainty about inspection result is not a good predictor of disclosure decisions; if anything, restaurants with a more uncertain history are slightly *more* likely to disclose. The hump-shaped disclosure pattern remains salient after controlling for these variables, while we can easily reject $\beta_{A+} > \beta_{A-}$ but cannot reject $\beta_{B+} > \beta_{B-}$.¹⁸

Column 3 includes the competition variables in addition to those already controlled in Column 2. Being in a ZIP code with a higher fraction of restaurants in the first batch is positively correlated with disclosure, while the number of nearby competitors is negatively correlated with disclosure. The coefficient associated with

¹⁸One may still be concerned that the volatility of the inspection outcomes works differently for *A*s and *B*s. Higher volatility in inspection outcomes for an *A* restaurant means more downward risk. Therefore, the restaurant will be reluctant to disclose. Higher volatility in inspection outcomes for a *B* restaurant could mean some upward risk, so a *B* restaurant may be willing to disclose in the hope of getting an *A*. We interact mean and standard deviation of past WSUMVIOL scores with whether the restaurant gets an *A* or a *B* in the current inspection. In the results not reported here, we show that while a higher mean WSUMVIOL score is negatively correlated with the probability of disclosure for both *A* and *B* restaurants, higher standard deviations of past WSUMVIOL scores are slightly positively correlated with disclosure decisions for *A* restaurants (and not statistically significant at conventional levels), and slightly negatively correlated with disclosure decisions for *B* restaurants (statistically significant at the 10 percent level). Furthermore, we also use the mean and standard deviation of getting an *A* as the alternative measure of past inspection outcomes, and results are similar.

mean WSUMVIOL score among the neighboring restaurants has a negative sign, which suggests that a higher average quality of competitors is associated with a higher probability of disclosure. Note that these variables do not attempt to identify any causal spillovers between competing restaurants. Rather, we use the competition variables to control for potential spillovers between competitors or whatever omitted variables may drive the interdependence of competitor decision. After these controls, the coefficients β s still exhibit a hump shape with statistical significance.

Column 4 accounts for dynamic concerns by including the letter grade of the restaurant's previous inspection. As expected, getting an *A* in the previous inspection is associated with a disclosure probability 10 percentage points higher than getting a *C* or *D*. In comparison, getting a *B* is associated with a disclosure probability 3 percentage points higher than getting a *C* or *D*. To better control for not only past inspection outcomes but also past disclosure decisions, Column 5 includes a dummy indicating whether the most recent disclosed grade (before the current inspection) is an *A*. Having an *A* displayed from previous inspections is actually *positively* correlated with the probability of disclosure in the current inspection. It is consistent with the signaling story where high-quality restaurants consistently have high grades and choose to disclose. It is inconsistent with the concern that once a restaurant gets an *A*, it will "play it safe" by choosing not to disclose in subsequent inspections. In both Columns 4 and 5, the hump shape of disclosure rate remains salient and statistically significant.¹⁹

Column 6 includes half-marked grades defined based on the second-to-the-last inspections. Since restaurant quality is persistent, *A+* restaurants in the previous

¹⁹In results not reported here, we also test dynamic concerns by including a dummy indicating whether the restaurant chose to disclose in the previous inspection, and the restaurant's previous disclosure decision interacted with the letter grade it received. The hump-shaped disclosure rate is robust to these additional tests, too, although because restaurants' inspection outcomes and disclosure decisions are highly serially correlated (see the transition matrix in Table 2), the hump-shaped pattern is only statistically significant at the 10 percent level.

inspections are likely to remain $A+$. If countersignaling is at play, they will choose non-disclosure in the current inspection and we should observe the hump-shaped disclosure rate among $A+$, $A-$, and $B+$ restaurants based on the past inspection outcomes. Column 6 shows exactly that: $A+$ restaurants in the previous inspection have a disclosure rate that is 2 percentage points lower than that of restaurants that received an $A-$ in the previous inspection, while restaurants that receive a $B+$ in the previous inspections have a disclosure rate that is 2 percentage points higher than that of $B-$ restaurants. We test the differences in these disclosure rates (ϕ 's) and show that the hump shape is statistically significant. Because the result of the previous inspection is known to the restaurant, this hump shape is not affected by potential uncertainty and risk-aversion concerns. At the same time, the hump shape between $A+$, $A-$, and $B+$ restaurants as defined by current inspection outcomes remain salient. Note that all columns of Table 4 have included a dummy for being listed on Yelp. Thus consumer prior knowledge about the restaurant, as embodied in the listing status on Yelp, does not explain away the hump shape of disclosure rate.

Finally, for robustness, we adopt an alternative definition of half-marked grades. In all the regressions presented in Table 4, we define “+” and “-” within each letter grade according to the disclosed outcomes of all restaurants in Maricopa. Maricopa County covers an area of over 9,000 square miles and has more than 4 million residents; arguably, a local restaurant may be more concerned with competitors down the street than with restaurants across the county. To address this concern, we re-define half marks within each letter grade relative to the median $WSUMVIOL$ score of the disclosed restaurants with the same letter grade in the same ZIP code, and re-run the analyses in Table 4. We cluster the standard errors at the ZIP code level as the key explanatory variables are now mechanically correlated within each ZIP code. We also include the restaurant’s consumer observable violations as an explanatory

variable, in case that enters consumer prior knowledge about the restaurant. The coefficient associated with it is economically small and not statistically significant. As Online Appendix Table A shows, results from this alternative definition are remarkable similar to those presented in Table 4, although statistical tests are relatively weaker because of error clustering.²⁰ Most tests reject the null hypothesis that $\beta_{A+} > \beta_{A-}$ at a statistical significance level between 0.05 and 0.1, while the null hypothesis of $\beta_{B+} > \beta_{B-}$ cannot be rejected.

IV.D. Further Evidence for Countersignaling

As FHT (2002) clarifies, countersignaling can only exist when there is another exogenous, noisy signal available in addition to the quality measure subject to disclosure. Moreover, both the exogenous signal and the to-be-disclosed measure must be coarse enough so that consumers cannot use just one of them to perfectly differentiate all levels of true quality.

For the Yelp variables to serve as the exogenous signal in countersignaling, it is important to check their correlation with grade cards. As Figure 4 shows, the Yelp variables are overall positively correlated with the vertical rank of letter grades overall, but their correlations are rather low and there is considerable overlap between any two of them.²¹ More specifically, Panel D of Table 3 presents the means of the Yelp variables by disclosure status and letter grade. Among *A* restaurants, the disclosing ones are less likely to be listed on Yelp than the non-disclosing ones. Conditional on being listed on Yelp, disclosing *A* restaurants are less likely to have good reviews or to be popular, although they are more likely to be pricey. In contrast, the Yelp variables are not significantly different between disclosing *B* restaurants and

²⁰If we use heteroskedasticity-robust standard errors instead, the null hypothesis that $\beta_{A+} > \beta_{A-}$ can be rejected at the 1 percent level.

²¹The correlation of half-marked grades (with *A+* denoted as 6 and *CD-* denoted as 1) with “popular”, “good review”, and “pricey” is 0.08, 0.17, and 0.03, respectively.

non-disclosing *B* restaurants. The *A* to *B* comparison is even more striking. When it comes to Yelp variables, disclosing *As* look much more like the *B* restaurants than like non-disclosing *As*. For example, roughly 22.5 percent of disclosing *As* are popular on Yelp. This percentage is much closer to that of disclosing *Bs* (22.6 percent) and non-disclosing *Bs* (25.7 percent) than to that of non-disclosing *As* (46 percent). Similar patterns occur for the other Yelp variables as well. All of the above is consistent with countersignaling at the high end of the quality spectrum.

A more explicit test of countersignaling calls for a direct comparison between “+” and “-” within letter grade. Following FHT (2002), countersignaling predicts that an *A+* restaurant is less likely to disclose if an extra and arguably better signal is available to show that it is of high quality. We test this hypothesis by taking Column 2 of Table 4 as the baseline and separately adding in each Yelp variable and its interaction with half grades. We test whether the coefficient associated with the interactive term with *A+* is smaller than that of the interactive term with *A-*.

Table 5 reports the results of these regressions. *A+* restaurants are much less likely to disclose than *A-* restaurants when they are listed on Yelp, or are popular on Yelp, or receive good reviews on Yelp. The differences in magnitudes of the coefficients are substantial and statistically significant. The only exception is for pricey restaurants, for which we do not observe a significant difference in disclosure rate between *A+* and *A-* restaurants. We postulate that it may be because that being pricey is not regarded as a good signal of quality.

Recall that countersignaling only predicts a lower disclosure rate for *A+* versus *A-*, not *B+* versus *B-*. This prediction is confirmed in Table 5: *B+* restaurants are more likely to disclose than *B-* restaurants, but overall their disclosure difference does not vary by the Yelp variables. In other words, both the differential effect of the Yelp variables on *A* restaurants and their lack of effect on *B* restaurants support countersignaling.

It is important to note that these findings cannot be explained by simply treating the Yelp variables as a proxy of consumer prior knowledge. Suppose that, before the disclosure decision, a restaurant knows that consumers already consider it to be good-quality because it has been listed on Yelp and received good reviews there. Given this prior belief, the restaurant should have less incentive to disclose the letter grade across the board. In other words, if favorable consumer prior belief discourages disclosure, it should hold for both *A* and *B* restaurants; however, we do not see this pattern for *B* restaurants. Thus consumer prior knowledge alone cannot explain the differential role of the Yelp variables for *A* and *B* restaurants. That being said, the countersignaling interpretation does not rule out the possibility that Yelp variables may incorporate some consumer prior knowledge. In fact, all regressions shown in Table 5 control for the Yelp variables before we interact them with half grades, and all columns in Table 4 include the dummy of being listed on Yelp, and the results essentially do not change if we include other Yelp variables. These controls do not explain away the hump shape of disclosure rate.

V. Conclusion

In contrast to the classical disclosure theory, we observe an empirical setting where disclosure is neither complete nor monotonic with quality. A closer look at the data suggests that the most likely explanation is a mixture of signaling and countersignaling.

In the post-GC inspections, disclosure rate drops steadily across letter grades, from 66 percent for *A* restaurants to 49 percent for *B* restaurants, and 30 percent for *C* and *D* restaurants. In other words, disclosure is by and large a positive signal to consumers over the whole distribution. However, at the high end of the grade spectrum, *A+* restaurants are less likely to disclose than *A-* restaurants. Focusing

on the last post-GC inspection of every restaurant, we find the disclosure rate of $A+$ restaurants is 7 percentage points lower than that of $A-$ restaurants. This is counter-intuitive, as $A+$ restaurants not only have fewer violations in the studied inspection, but also have fewer violations in previous inspections, have a smaller standard deviation in their inspection histories, are more likely to be listed in Yelp, and receive more and better reviews on Yelp. A likely explanation is countersignaling, where the very best restaurants have an incentive to use non-disclosure as a countersignal to distinguish themselves from eager-to-disclose medium-quality restaurants. This happens because the to-be-disclosed quality information (letter grade) is coarse and there exists another exogenous signal (the Yelp variables) that is noisy but overlaps with the distribution of letter grade. As a result, $A+$ restaurants expect sufficiently good Yelp reviews and can afford to use non-disclosure to signal their stellar quality, but $A-$ restaurants are eager to distinguish themselves from B s because their future Yelp reviews may turn out to be similar to those of B s.

While we are convinced that the non-monotonic pattern of disclosure is best explained by signaling and countersignaling, we also find significant influence from other factors. For example, restaurants that were examined sooner after the adoption of the grade card policy (so called first-batch) are more likely to disclose than other restaurants. This pattern persists over time, suggesting the important and long-lasting effect of consumer attention. There is also evidence that restaurants with more competitors in the same ZIP code are less likely to disclose, and restaurants with lower-quality neighbors are less likely to disclose, although these results should be interpreted as correlations rather than causal effects. Finally, there is some evidence that restaurants with a bigger standard deviation in historical violations are slightly more likely to disclose (although the difference is not statistically significant), and restaurants that already received an A in the last disclosed inspection are more likely to disclose. Both of these patterns are against the arguments of

uncertainty or dynamic concerns, which predict that risk-averse restaurants should be more reluctant to disclose when they face more uncertainty in inspection outcomes, and a restaurant that has already got an *A* in the window should play it safe by choosing nondisclosure this time. It seems that these predictions are dominated by the fundamental incentive to disclose high quality.

A remaining question is why, when we focus on the last observed post-GC inspection for each restaurant, only 67 percent of *A*– restaurants (and only 49 percent of *B* restaurants) choose to disclose. In theory, disclosure cost could explain any incomplete disclosure, but it is unclear what constitutes the disclosure cost in this particular context. Relatedly, one may wonder what information consumers derive from the market when over 40 percent of restaurants choose to be silent. To the extent that countersignaling exists, even non-disclosure contains extra information on top of what is already available to consumers. In this sense, disclosure rate – measured by the percentage of firms that choose to disclose – may not be the best measure of transparency, and 100 percent disclosure is not necessarily the most desirable metric to define the success of a disclosure policy. Non-disclosure, as a countersignal, may still add value for consumers if they can correctly understand its information content. That being said, whether consumers can comprehend a countersignal in reality remains unknown. On the one hand, experimental evidence presented in FHT (2002) suggests that senders (firms) are able to implement a countersignaling equilibrium in the lab. On the other hand, Jin, Luca and Martin (2015) show that, in a simple disclosure game (without countersignaling), senders are typically more sophisticated than receivers and receivers are not paying too much attention to what it means by non-disclosure. Consumer perception of non-disclosure warrants further study.

References

- Board, Oliver.** 2009. "Competition and Disclosure." *Journal of Industrial Economics*, 57(1): 197–213.
- Dranove, David, and Ginger Zhe Jin.** 2010. "Quality Disclosure and Certification: Theory and Practice." *Journal of Economic Literature*, 48(4): 935–963.
- Feltovich, Nick, Richmond Harbaugh, and Ted To.** 2002. "Too Cool for School? Signalling and Countersignalling." *RAND Journal of Economics*, 33(4): 630–650.
- Fishman, Michael J, and Kathleen M Hagerty.** 2003. "Mandatory Versus Voluntary Disclosure in Markets with Informed and Uninformed Customers." *Journal of Law, Economics, and organization*, 19(1): 45–63.
- Fung, Archon, Mary Graham, and David Weil.** 2007. *Full Disclosure: The Perils and Promise of Transparency*. Cambridge University Press.
- Grossman, Sanford J.** 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *The Journal of Law & Economics*, 24(3): 461–483.
- Grubb, Michael D.** 2011. "Developing a Reputation for Reticence." *Journal of Economics & Management Strategy*, 20(1): 225–268.
- Harbaugh, Richmond, and Theodore To.** 2016. "False Modesty: When Disclosing Good News Looks Bad."
- Jin, Ginger Zhe, Michael Luca, and Daniel Martin.** 2015. "Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure." National Bureau of Economic Research 21099.
- Jovanovic, Boyan.** 1982. "Truthful Disclosure of Information." *The Bell Journal of Economics*, 36–44.
- Kang, Jun Seok, Polina Kuznetsova, Michael Luca, and Yejin Choi.** 2013.

“Where Not to Eat? Improving Public Policy by Predicting Hygiene Inspections Using Online Reviews.” *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Marinovic, Iván, Andrzej Skrzypacz, and Felipe Varas. 2015. “Dynamic Certification and Reputation for Quality.”

Mathios, Alan. 2000. “The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market.” *Journal of Law and Economics*, 43(2): 651–77.

Matthews, Steven, and Andrew Postlewaite. 1985. “Quality Testing and Disclosure.” *The RAND Journal of Economics*, 328–340.

Milgrom, Paul R. 1981. “Good News and Bad News: Representation Theorems and Applications.” *The Bell Journal of Economics*, 380–391.

Figures and Tables

Figure 1: Grading System

Grading System

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

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

Source: Maricopa County government website.
<https://www.maricopa.gov/EnvSvc/EnvHealth/PermitScoring.aspx>. Last visited in August, 2016.

Figure 2: Example of the Online Database

FD-33452	AJ's Food and Grill	2209 N 99th Ave	Phoenix	
FD-14971	Alexis Grill	3550 N Central Ave Suite 120	Phoenix	85012
FD-00464	Alexis Grill	3550 N Central Ave Suite 120	Phoenix	85012

Alexis Grill - 3550 N Central Ave Suite 120 - Permit ID: FD-00464

Violations ID links to additional information.


Total Records Retrieved [16]

Inspected Date	Purpose	Grade	Violation ID	Violation Description	Violation Comments	Correct By
8/27/2013	Re-Inspection					
8/8/2013	Routine Inspection	B	14	Food-contact surfaces, cleaned & sanitized	SECOND OCCURRENCE; CONSECUTIVE 4-501.14, P: Manual and Mechanical Warewashing Equipment, Chemical Sanitization-Temperature, pH, Concentration and Hardness - Chlorine dishwasher behind bar had a concentration of 0ppm. Line had air bubbles and manager primed machine so that proper amounts of sanitizer were being dispensed. All chlorine sanitizers must be maintained between 50-100 ppm at all times.	Corrected At Time Of Inspection
4/12/2013	Routine Inspection	Not Participating				
2/11/2013	Routine Inspection	Not Participating				
11/26/2012	Routine Inspection	Not Participating				
7/26/2012	Routine Inspection	Not Participating				
5/18/2012	Routine Inspection	Not Participating				
1/31/2012	Routine Inspection	C	23	Consumer advisory provided for raw or undercooked foods	3-503.11 PFC: Consumption of Animal Foods that are Raw, Undercooked, or Not Otherwise Processed to Eliminate Pathogens - Establishment serves undercooked items. Items on menu are asterisked, and reminder is on bottom, but menus are missing disclosure statement that 'these items may be served raw/undercooked' or 'these items may contain raw/undercooked items.' Manager will modify menus to properly display consumer advisory.	Correct Prior To Reinspection
6/15/2011	Routine Inspection					
2/11/2011	Routine Inspection					
12/3/2010	Re-Inspection					
11/30/2010	Advisory					
11/16/2010	Complaint Inspection					
10/5/2010	Routine Inspection		31	Proper cooling methods used; adequate equipment for temperature control	3-501.15 (A), PFC: Cooling Methods - Cut tomatoes in prep unit were holding at 53°F. Manager said tomatoes were just cut, but were not cooled to 41°F prior to placing on line. All potentially hazardous foods	Corrected At Time Of Inspection
6/8/2010	Advisory					
2/16/2010	Routine Inspection					

Violations ID links to information about violations. Total Records Retrieved [2]

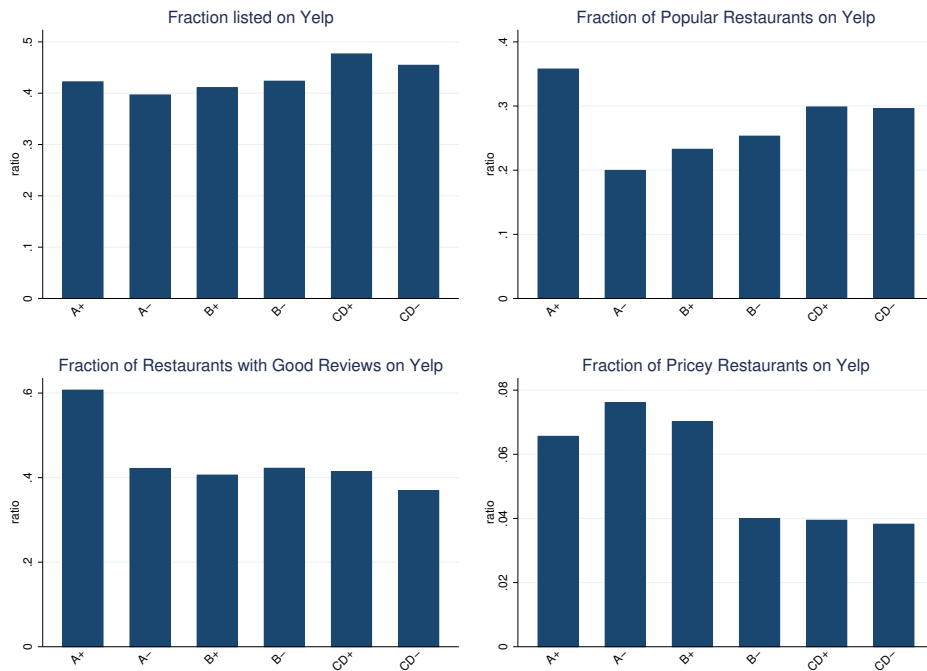
Violation ID	Violation Description	Violation Comments	Correct By
14	Food-contact surfaces, cleaned & sanitized	CHEMICAL SANITIZER (CHLORINE) IN THE 3 COMPARTMENT SINK FOUND AT 0PPM. EMPLOYEE REMADE SANITIZER AT TIME OF INSPECTION TO 100PPM.	
26	Toxic substances properly identified, stored, and used	SOUP FOUND COOLING IN A CONTAINER LABELED (MANUFACTURER'S LABEL) HAND SOAP. SEVERAL OTHER	

Figure 3: Disclosure by Grade



Note: The sample includes the last observed inspection in the post-GC period for each restaurant. The height of the bar indicates the proportion of restaurants that choose to disclose in each half-mark grade. The x-axis shows grade letters in half-marks. Both the mean test and the Kolmogorov-Smirnov distribution test show that the disclosure rate for A+ is significantly lower than that for A-, while the disclosure rates for B+ and CD+ are significantly higher than those for B- and CD-, respectively.

Figure 4: Correlation with Other Signals



Note: The sample includes last observed post-GC inspections. Each graph plots the average value of another signal of restaurant quality, against grades in half marks. Except for the top-left graph, the sample is conditional on being listed on Yelp.

Table 1: Summary Statistics

	N	total vios	WSUMVIOL	P	PF	Core
Before Disclosure	146498	1.763				
Disclosed	37600	0.917	2.293	0.242	0.217	0.458
A	26465	0.343	0.343	0.000	0.000	0.343
B	9420	1.940	5.828	0.629	0.667	0.644
C/D	1715	4.157	12.981	1.842	1.099	1.216
Non-disclosed	27529	1.563	4.474	0.519	0.451	0.592
A	13625	0.312	0.312	0.000	0.000	0.312
B	9938	2.054	6.213	0.618	0.768	0.668
C/D	3966	4.627	14.415	2.057	1.206	1.364

Note: Each observation is an inspection. Letter grades for non-disclosing restaurants are imputed.

Table 2: Transition Matrix

previous inspection	current inspection									
	A		B		C/D					
	dis	non-dis	dis	non-dis	dis	non-dis	dis	non-dis	dis	non-dis
dis	0.807	0.193								
non-dis	0.290	0.710								
A			0.734		0.231			0.035		
dis			0.650	0.091		0.164	0.061		0.018	0.015
non-dis			0.244	0.471		0.052	0.193		0.005	0.034
B			0.469		0.415			0.116		
dis			0.426	0.070		0.284	0.115		0.053	0.052
non-dis			0.185	0.258		0.086	0.344		0.015	0.111
C/D			0.285		0.442			0.273		
dis			0.272	0.049		0.293	0.140		0.110	0.136
non-dis			0.134	0.135		0.100	0.346		0.031	0.254

Note: Letter grades are as disclosed or are imputed. Each cell indicates the probability of transition between two consecutive inspections. Each row sums up to one.

Table 3: Characteristics of Disclosing and Non-Disclosing Restaurants

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A	A or imputed A			B or imputed B		
Average of current inspections	dis	non-dis	(1)-(2)	dis	non-dis	(4)-(5)
WSUMVIOL	0.319	0.269	0.050***	5.923	6.268	-0.345***
priority vios	0.000	0.000	0.000	0.617	0.589	0.028**
priority foundation vios	0.000	0.000	0.000	0.706	0.786	-0.080***
core vios	0.319	0.269	0.050***	0.631	0.770	-0.139***
consumer observable vios	0.175	0.157	0.018**	0.617	0.684	-0.067***
all vios	0.319	0.269	0.050***	1.954	2.145	-0.191***
risk factor vios	0.041	0.036	0.005	1.205	1.249	-0.044**
retail practice vios	0.277	0.233	0.044***	0.749	0.896	-0.147***
Panel B	A or imputed A			B or imputed B		
Average of past inspections	dis	non-dis	(1)-(2)	dis	non-dis	(4)-(5)
WSUMVIOL	0.750	0.716	0.034	2.526	3.122	-0.596***
priority vios	0.075	0.069	0.006*	0.264	0.316	-0.052***
priority foundation vios	0.068	0.070	-0.002	0.284	0.361	-0.077***
core vios	0.177	0.157	0.020***	0.335	0.413	-0.078***
consumer observable vios	0.414	0.367	0.047***	0.682	0.791	-0.109***
all vios	0.320	0.296	0.024***	0.883	1.091	-0.208***
risk factor vios	0.467	0.424	0.043***	0.992	1.114	-0.122***
retail practice vios	0.572	0.543	0.029***	0.919	1.042	-0.123***
Panel C	A or imputed A			B or imputed B		
stdev of past inspections	dis	non-dis	(1)-(2)	dis	non-dis	(4)-(5)
WSUMVIOL	1.418	1.379	0.039	3.169	3.532	-0.363***
priority vios	0.171	0.159	0.012**	0.418	0.454	-0.036***
priority foundation vios	0.159	0.167	-0.008	0.428	0.490	-0.062***
core vios	0.322	0.286	0.036***	0.523	0.587	-0.064***
consumer observable vios	0.563	0.505	0.058***	0.752	0.798	-0.046***
all vios	0.548	0.514	0.034***	1.067	1.199	-0.132***
risk factor vios	0.619	0.586	0.033***	0.901	0.936	-0.035***
retail practice vios	0.720	0.707	0.013	0.944	0.982	-0.038***
Panel D	A or imputed A			B or imputed B		
Mean of other signals	dis	non-dis	(1)-(2)	dis	non-dis	(4)-(5)
on Yelp	0.384	0.469	-0.085***	0.428	0.409	0.019
good reviews on Yelp	0.491	0.672	-0.181***	0.408	0.419	-0.011
pricey on Yelp	0.213	0.172	0.041***	0.201	0.192	0.009
popular on Yelp	0.225	0.460	-0.235***	0.226	0.257	-0.031*

Note: The sample includes the last observed post-GC inspections. Panel A shows the mean characteristics for restaurants receiving an A or an imputed A (Columns 1 to 3) or restaurants receiving a B or an imputed B (Columns 4 to 6). Panel B shows the average of the mean of past inspections of these restaurants. Panel C shows the standard deviation of the past inspections. Panel D shows the mean values of other signals. The last three rows are conditional on the restaurant being listed on Yelp. Column 3 shows the mean difference between the value in Column 1 and that in Column 2. Column 6 shows the mean difference between the value in Column 4 and that in Column 5. Asterisks indicate statistical significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Disclosure by Imputed Grade

<i>devar</i> : disclosure in current inspection	(1)	(2)	(3)	(4)	(5)	(6)
A+ in current inspection (β_{A+})	0.371*** (0.017)	0.295*** (0.023)	0.294*** (0.023)	0.307*** (0.023)	0.347*** (0.023)	0.324*** (0.020)
A- in current inspection (β_{A-})	0.421*** (0.019)	0.341*** (0.023)	0.320*** (0.024)	0.356*** (0.023)	0.380*** (0.024)	0.371*** (0.021)
B+ in current inspection (β_{B+})	0.221*** (0.019)	0.150*** (0.023)	0.139*** (0.023)	0.165*** (0.023)	0.214*** (0.023)	0.179*** (0.021)
B- in current inspection (β_{B-})	0.176*** (0.019)	0.123*** (0.022)	0.114*** (0.022)	0.135*** (0.022)	0.169*** (0.023)	0.147*** (0.021)
C/D+ in current inspection	0.053** (0.024)	0.014 (0.026)	0.010 (0.026)	0.028 (0.026)	0.080*** (0.027)	0.038 (0.025)
first-batch	0.056*** (0.008)	0.041*** (0.008)	0.037*** (0.008)	0.042*** (0.008)	-0.005 (0.007)	0.045*** (0.008)
mean WSUMVIOL of own past inspections		-0.014*** (0.002)	-0.010*** (0.002)	-0.010*** (0.002)	0.002 (0.002)	
stdev WSUMVIOL of own past inspections		0.003 (0.002)	-0.001 (0.002)	0.008*** (0.002)	0.014*** (0.002)	
mean WSUMVIOL of ZIP restaurants			-0.034*** (0.004)			
fraction of ZIP restaurants in first batch			0.044 (0.039)			
# of restaurants in ZIP ($\times 1,000$)			-0.285*** (0.012)			
A in previous inspection				0.101*** (0.017)		
B in previous inspection				0.033** (0.016)		
latest disclosing grade is A					0.413*** (0.007)	
A+ in previous inspection (ϕ_{A+})						0.122*** (0.021)
A- in previous inspection (ϕ_{A-})						0.142*** (0.022)
B+ in previous inspection (ϕ_{B+})						0.064*** (0.022)
B- in previous inspection (ϕ_{B-})						0.042* (0.022)
C/D+ in previous inspection						0.021 (0.027)
year-season FE	X	X	X	X	X	X
<i>N</i>	19719	18482	18482	18431	17056	18431
<i>R</i> ²	0.065	0.064	0.093	0.068	0.217	0.067
$\beta_{A+} > \beta_{A-}$ (<i>p</i> -value)	0.000	0.000	0.006	0.000	0.000	0.000
$\beta_{B+} > \beta_{B-}$ (<i>p</i> -value)	0.999	0.972	0.961	0.982	0.999	0.990
$\phi_{A+} > \phi_{A-}$ (<i>p</i> -value)						0.027
$\phi_{B+} > \phi_{B-}$ (<i>p</i> -value)						0.942

Note: The sample includes last observed inspections in the post-GC period. All columns also include a dummy indicating whether the restaurant is listed on Yelp, and a dummy indicating whether it belongs to a restaurant chain. Robust standard errors are reported in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Extra Information and Disclosure

<i>depar</i> = disclosure	(1)	(2)	(3)	(4)
<i>extra information</i>	on Yelp	popular	good reviews	pricey
A+	0.304*** (0.022)	0.328*** (0.027)	0.371*** (0.029)	0.290*** (0.025)
A-	0.304*** (0.023)	0.358*** (0.029)	0.389*** (0.032)	0.394*** (0.026)
B+	0.106*** (0.023)	0.193*** (0.029)	0.198*** (0.031)	0.211*** (0.025)
B-	0.112*** (0.023)	0.111*** (0.029)	0.156*** (0.032)	0.131*** (0.025)
<i>extra info</i>	-0.040 (0.025)	-0.070* (0.037)	0.021 (0.037)	-0.073 (0.088)
A+ × <i>extra info</i> (γ_{A+})	-0.047* (0.027)	-0.148*** (0.041)	-0.161*** (0.040)	0.021 (0.094)
A- × <i>extra info</i> (γ_{A-})	0.068** (0.031)	0.123** (0.051)	-0.004 (0.047)	0.001 (0.105)
B+ × <i>extra info</i> (γ_{B+})	0.082*** (0.031)	0.026 (0.049)	0.009 (0.046)	-0.051 (0.103)
B- × <i>extra info</i> (γ_{B-})	0.006 (0.033)	0.042 (0.052)	-0.076 (0.049)	-0.062 (0.117)
other covariates	X	X	X	X
year-season FE	X	X	X	X
<i>N</i>	18482	7718	7718	7718
<i>R</i> ²	0.067	0.081	0.071	0.064
<i>p</i> -values of testing				
$\gamma_{A+} > \gamma_{A-}$	0.000	0.000	0.000	0.619
$\gamma_{B+} > \gamma_{B-}$	0.997	0.370	0.976	0.549

Note: Other covariates include whether the restaurant is in the first batch, whether it belongs to a restaurant chain, and the mean and standard deviation of past grading. “on Yelp” is a dummy indicating whether the restaurant is listed on Yelp.com. Columns 2 to 4 include only restaurants that are listed on Yelp. A restaurant is “popular” if it has more than 60 reviews (the number of reviews for the median restaurant in the sample). A restaurant has “good reviews” on Yelp if it has 4 or more Yelp stars (ranging from 1 to 5 with 0.5 increments). A restaurant is “pricey” if it has 3 or 4 dollar signs on Yelp (ranging from 1 dollar sign to 4 dollar signs). Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.