Having it Easy: Discrimination and Specialization in the Workplace

Sacha Kapoor and Arvind Magesan*

April 10, 2019

Abstract

Most studies analyzing worker adjustments to discrimination focus on sorting decisions, such as occupations workers pursue. We instead analyze on the job adjustments, focusing on the effects of discrimination by customers and managers. Specifically, we use observational and experimental data from large-scale restaurants to investigate the effects of the symmetry of waiter facial attributes on the trade offs they make, the extent to which the trade offs are shaped by consumer preference for symmetry, and whether the trade offs affect how they are treated by managers. We find customers have a preference for symmetry and that preferred workers consequently give them *less* personal attention. This allows preferred workers to serve more customers, generate more revenue for the firm, and more earnings for themselves. The large ensuing productivity and earnings differentials are facilitated largely by small differences in managerial treatment of preferred workers.

JEL: J31, J71, L80, M50

Keywords: Attractiveness, Task Specialization, Manager Discrimination, Customer Discrimination, Wage Differentials

Occupations in the modern economy typically require workers to carry out several tasks, where the allocation of effort across tasks is, at least to some degree, at the discretion of the worker. As such, workers who wish to maximize earnings should choose to specialize

*Kapoor: Erasmus University Rotterdam, Erasmus School of Economics H09-22, Burgemeester Oudlaan 50, 3062 PA Rotterdam, The Netherlands, kapoor@ese.eur.nl. Magesan: University of Calgary, Department of Economics, 2500 University Drive N.W., Calgary AB, T2N 1N4, Canada, anmagesa@ucalgary.ca. We thank Dwayne Benjamin, Gustavo Bobonis, Branko Boskovic, Laura Hering, Nicola Lacetera, Boris van Leeuwen, Robert McMillan, Daniel Parent, Aloysius Siow, Scott Taylor, Trevor Tombe, Jeroen van de Ven, and seminar participants at several universities for comments that greatly improved the paper. This research was supported by the Social Sciences and Humanities Research Council and the Canadian Labour Market and Skills Research Network. All omissions and errors are our own.

in tasks where they have a comparative advantage. In the classic Roy [1951] model,[1] for example, workers with the same skills carry out their tasks in the same way and earn the same amount. The theory is at odds with a large body of empirical work showing workers who have similar skills, but differ only in outward physical traits, have different earnings. A common reason given for the earnings disparities is the adjustments workers make in response to discrimination. For example, in response to anticipated employer discrimination, the trait may influence the types of jobs and occupations a worker pursues, thereby creating disparities in wages across otherwise identical workers. There are, however, few empirical studies examining whether and why earnings disparities emerge in the workplace itself. This is important because, even after workers have sorted into particular jobs and occupations, the same trait may continue to be a source of earnings disparity amongst workers.

In this view, the role of comparative advantage is not limited to sorting across jobs and occupations. The outward physical traits of workers may also shape comparative advantage in the workplace and, consequently, induce workers with similar skills to specialize in different tasks. But this then begs the question: how could outward physical traits alone be a source of comparative advantage at work? The answer may lie with the nature of work in the modern economy. Many jobs today, particularly those found in the ever-expanding services sector,[2] require interactions with customers. If customers buy more from workers who possess traits they prefer, then they shape the opportunity costs of these interactions and alter the tasks workers focus on. If managers are aware of the preference, and how workers respond to it, then the consumer preference can also generate disparities in how workers are treated by managers. All this can exacerbate productivity and earnings differences among workers.

We investigate the interplay between discrimination and comparative advantage in the workplace using detailed data from a couple of franchises of a North American restaurant chain. Specifically, we study the effects of the symmetry of waiter facial attributes - an outward, immutable, and subtle physical trait [Jones et al., 2017, Perry et al., 2015] - on how they are treated by customers and managers, on how they trade off the quality for the quantity of service they provide to customers, as well as the origins of the differential treatment

[1]For more recent formalizations of the Roy model see [Borjas, 1987] or [Heckman and Honore, 1990]. Many studies use the Roy model to examine the effects of comparative advantage on sorting decisions, such as those relating to the pursuit of higher education [Willis and Rosen, 1979], the choice of industry [Heckman and Sedlacek, 1985] or occupation [Miller, 1984], job assignments [Foster and Rosenzweig, 1996], union status [Lemieux, 1998], and the region to live in [Dahl, 2002].

[2]The rise of the service sector is one of the most stark economic developments of the last several decades. In the United States, for example, the service sector currently comprises 62.9 percent of U.S. GDP (CIA World Factbook).

by customers and managers.[3] We show customers have a preference for symmetry, that this induces preferred workers to pay less attention to individual customers and serve more of them, and to thus generate more revenue for the firm and earnings for themselves. Importantly, we show most of the earnings/revenue differences can be explained by allocations of work assignments by managers.

The venue has several useful features for understanding how discrimination shapes workplace specialization. First, the scale - thousands of consumers visit each franchise each week - introduces a trade off between personalized attention and customer volume. In restaurants of any scale, workers can sell - and earn - more by attending to customer wants. At large-scale restaurants, however, workers can also sell more by serving more customers. This is important because, without a tradeoff, it is impossible to study the role of comparative advantage in the workplace. Second, a quasi-random process pairs customers and workers, such that customers are exogenously assigned a bundle of worker traits, and workers are exogenously assigned a set of customers.[4] Our conclusions should therefore apply to customers who were assigned other workers, and to workers who were assigned a different set of customers. This feature ensures our samples are at least representative of the specific setting under study.

Facial symmetry is a useful trait in this context because there is a general agreement - across cultures and countries - that symmetric faces have greater appeal than asymmetric ones.[5] A general agreement reinforces an advantage of the exogenous matching process, in that it helps to further circumvent the need to control for customer attributes in our empirical analysis. If consumers were to differ in how they rank symmetry, then we would require information on customer attributes to distinguish the effects the trait has on worker behavior and performance. A general agreement on the appeal of symmetry also increases the chances that workers know whether customers prefer their endowment of symmetry. If there was some ambiguity about the value of the endowment to consumers, as is typically the case with other traits, then the responses we observe might be attributed to a lack of information on the part of workers rather than a consumer preference for the trait itself.

Coupling observational data with the (testable) assumption that matches are exogenous,

[3]While previous studies investigate the impact of cognitive and noncognitive skills on comparative advantage [Heckman, Stixrud, and Urzua, 2006], we investigate the impact of a trait that has no ostensible role for the skill and dexterity of workers.

[4][Persico, 2009] discusses the pitfalls when only one attribute, such as race, is exogenously assigned to a party with the opportunity to discriminate.

[5]Social psychologists argue the agreement partly originates in the fact that symmetry is a signal of the ability to reproduce [Thornhill and Gangestad, 1999]. Other characteristics that signal this ability include youth in females, and resource control (and the willingness to share) in males [Mulford et al., 1998].

we show workers with more symmetric faces pay less attention to customers and, in doing so, serve more customers. Symmetrical workers in turn earn 8.4 percent more than their more asymmetrical counterparts, and generate 12 percent more revenue for the firm. The findings suggest the natural advantage facial symmetry bestows allows workers to neglect individual customers and instead focus on generating customer volume. Using the fact that symmetry directly affects interactions with customers, but not the number served, we argue this finding is consistent with symmetry being a substitute for personal attention in consumer demand.

We consider an alternative test of substitutability that uses data generated by a field experiment. One possible criticism of using observational data is worker behavior may be explained by unobserved heterogeneity. The field experiment holds characteristics like ability fixed, and exogenously rewards workers with a bonus for customer volume. We show that, in theory, if symmetry is a substitute for personal attention, then workers with symmetric faces increase their customer volume by more than workers with less symmetric faces. We show further that this differential increase is borne out in the data. The differential increase in customer volume increases the earnings gap by another 13.5 percent, and the productivity gap by another 18.3 percent.

We use rich data on work assignments to show a large share of the productivity and earnings differentials between symmetrical and asymmetrical workers can be explained by small differences in how managers treat preferred workers. Our evidence implies further that managers treat preferred workers favorably because of attempts to exploit the productivity advantages the consumer preference bestows on symmetrical workers, and not because they themselves have a symmetry preference.

Finally, we use tip data to explore the origins of substitutability in consumer demand, focusing on whether it originates in a consumer preference for symmetrical faces. The tip data is useful in this regard because tips are paid after services are rendered. The information facial symmetry conveys should play no role in the tips a worker receives since, at this point of the meeting, there are no further product exchanges between worker and customer. Thus, if workers with symmetrical faces earn more in tips, we can conclude it is because consumers have a preference for symmetrical faces. We use this to present direct evidence of substitutability (**c.f.** Belot, Bhaskar, and van de Ven [2012]), to argue that the differential tip rates reflect a preference for symmetry, and investigating whether the preference reflects a taste for symmetry or a distaste for asymmetry.

Related Literature. This article contributes to a literature that examines the relation-

ship between facial attractiveness and labor market outcomes. The literature has shown robust positive correlations between facial attractiveness and earnings in large representative samples [Hamermesh and Biddle, 1994, Harper, 2000, Scholz and Sicinski, 2015]. Most of what we know about why there is this premium, or any other wage premium for outward physical characteristics, comes from studies of job advertisements [Kuhn and Shen, 2013], the capacity to negotiate for higher wages [Mobius and Rosenblat, 2006], hiring decisions [Goldin and Rouse, 2000], the choice of occupation ([Mocan and Tekin, 2010] and [Biddle and Hamermesh, 1998]), the skills or experiences workers acquire before entering the labor market ([Case and Paxson, 2008], [Persico, Postlewaite, and Silverman, 2004], and [Neal and Johnson, 1996]). Unlike these other studies, we use data from inside the workplace to investigate potential explanations for the wage premium from facial attractiveness. The workplace allows for a number of new insights, of which many are linked to the extraordinary quality of the data and the exogenous variation the field experiment generates.

This article contributes to a literature that examines the causes and economic consequences of customer discrimination against workers [Becker, 1957]. Previous literature has focused on customer discrimination on the basis of race or gender: Leonard, Levine, and Guiliano [2010] use administrative data from retail establishments to investigate the impact of consumer racial discrimination on overall sales, finding customers are fairly insensitive to employee racial composition; Holzer and Ihlanfeldt [1998] use surveys of large employers in select US cities to show the racial composition of customers affects the race of new hires; Neumark et al. [1996] conducts an audit study of Philadelphia restaurants and finds the sex composition of customers affects the gender of new hires.[6] To draw inferences about customer discrimination, these studies have historically relied on the observed choices of firms or workers, rather than that of customers, and on the assumption consumers prefer interacting with persons similar to themselves.

We differ from these other studies in that we examine customer discrimination based on facial attractiveness, and in that we rely on the observed and independent choices of customers (tips) to draw direct inferences about customer discrimination. In the latter regard, our paper complements [Bar and Zussman, 2017], which utilizes a direct survey-based approach to measure customer discrimination by Israelis against Arabs. Our data lets us go further because, in addition to observing independent choices by customers, we observe independent choices made by managers. This information facilitates an analysis

[6]There are a handful of older but notable studies in this literature [Borjas and Bronars, 1989, Nardinelli and Simon, 1990, e.g.]. See [Leonard, Levine, and Guiliano, 2010] and [Holzer and Ihlanfeldt, 1998] for details and further references.

of how customer and manager discrimination interact to generate workplace earnings and productivity differentials. From this angle, our analysis and findings complement recent empirical work on the deleterious effects of managerial bias for the productivity of workers from minority populations [Glover, Pallais, and Pariente, 2017].

1 Context

The context is large-scale franchise restaurants that have upwards of 3000 customers per week and that are only open for dinner. The franchises are part of a long chain of higher end big-box restaurants. The areas around the franchises are not easily walkable - customers cannot be solicited from the street. Our analysis focuses on the waiters of these restaurants.

Managers use a well-defined formula to match customers with workers. A new customer is either seated right away or informed of the current wait time for seating. If there is no wait, the customer is assigned on the basis of the worker's start time. The worker with the earliest start time gets the customer. If there is a wait, the customer is assigned to the first seat that comes available. In this case, matches implicitly depend on worker table assignments, as some tables turn over faster than others. The formula makes customer-worker matches quasi-random, i.e. random conditional on the start time and table assignment of the worker. These matches are exogenous therefore to the facial symmetry of the worker, as well as to the traits of customers. The feature lets us separate the direct role of physical appearance on worker production from its role decisions about who matches with whom.

Together, the scale and the matching procedure implies workers likely face the same customer on average. Since matches are exogenous to the facial symmetry of the worker, and the scale ensures the workers interact with many customers, differences in consumer preferences at the individual level are 'averaged' out at the daily level. That is, under the assumptions that the preference of one customer is unrelated to preferences of others and that workers draw customers from the same distribution,[7] the law of large numbers implies differences in the customers workers face average out over the course of a shift. Having workers face a representative consumer is useful because not having information on the individual traits of customers becomes even less important for the conclusions that we draw.

Although managers have no direct control over the matching process, they have a couple of informal instruments they can use to exploit customer discrimination against workers, as

[7] In this way, another nice feature of the setting is that patrons come from a relatively homogeneous population. The franchises are located in large suburbs of Toronto which consists mostly of married people of homogenous (European) origins. **Source:** Statistics Canada - 2006 Census of Population.

well as worker responses to it. Both instruments give managers opportunities to expose more customers to their preferred workers, and to take advantage of differences in the number of customers workers serve.

The first instrument is the assignment of start times. While workers have latitude over the days they work,[8] they have little latitude over when their shifts begin.[9] Instead, a single manager decides on start times, usually a couple of weeks in advance of a given work week. The manager can expose symmetrical workers to more customers, and give them more opportunities to serve lots of customers, by assigning them earlier start times. Early starts are better for exposure and customer volume because there are more consumers arriving before the rush begins (around 5pm) than after the rush ends (around 10:30pm).

The second instrument is table assignments.[10] Before a shift begins, the firm demarcates tables into groups of two to four (10-16 seats) and assigns the groups to workers. Table assignments are fixed. Workers cannot increase customer volume by taking on additional tables. Tables are also not shared. Once a worker is assigned a table, that worker serves all of the customers who get seated at the table. In order to increase the exposure of customers, the firm can assign workers with symmetric faces larger sections. The firm can assign these workers tables that turn over more quickly to take advantage of their responses to customers.

The firm pays workers an hourly wage and lets them keep their tips. The hourly wage is the same for all workers, equalling the provincially-mandated minimum wage for servers in Ontario. The customer pays the worker a tip for their services. The tip is fully at the discretion of the customer. It is effectively proportional to the revenue from customer product purchases, where the proportion is the tip rate. The tip rate generally ranges from between 12 to 18 percent of the post-tax bill. Note that the tip comprises 76-77 percent of worker hourly earnings on average.

[8] Favorable treatment in the assignment of workers to shifts is unlikely in our setting. In order to minimize employee turnover, workers have latitude over the days they work. For shifts where there is a shortage of willing workers, workers take turns working the undesirable shifts.

[9] Start times are a useful feature of the context for other reasons as well. One is that they are staggered, ranging from 3:30pm until approximately 6:30pm, so that we observe different start times for workers working the same shift. Accordingly, we can account for differences in outcomes that are generated by differences in start times.

[10] Table assignments change from shift to shift and are at the full discretion of the manager-in-charge.

2 The Trade off at Work

In this section we construct a simple model to guide our empirical analysis. Formally, let $E(n, \mathcal{S})$ represent the total earnings of a worker with facial symmetry $\mathcal{S}$ who serves n customers. We view n as a choice of (average) personal attention, as it implies a personal attention of $q(n)$. q is a composite measure, including the effort workers put into generating revenue (*e.g.* discussing and recommending menu items), the prompt delivery of products, table maintenance (*e.g.* removing dirty dishes), and time spent with customers.[11] We have by definition:

$$E(n, \mathcal{S}) = e(q(n), \mathcal{S})n \tag{1}$$

where $e(q(n), \mathcal{S})$ represents tip earnings per customer.

We assume more personal attention increases earnings per customer, $\frac{\partial(e(q,\mathcal{S}))}{\partial q} > 0$ at a diminishing rate $\frac{\partial^2(e(q,\mathcal{S}))}{\partial q} \leq 0$, and the same for the facial symmetry of the worker, $\frac{\partial(e(q,\mathcal{S}))}{\partial \mathcal{S}} > 0$ and $\frac{\partial^2(e(q,\mathcal{S}))}{\partial \mathcal{S}^2} \leq 0$. The assumption that $\frac{\partial(e(q,\mathcal{S}))}{\partial q} > 0$ fits with economic theories that rationalize tipping [Azar, 2007], namely that tipping provides a means for rewarding service activities which are costly for the firm to monitor or measure.

$e(q(n), \mathcal{S})$ warrants additional comment. First, since tip earnings per customer equals customer expenditures on services rendered, the nature of the relationship between facial symmetry and personal attention in service demand determines the nature of their relationship in $e(q(n), \mathcal{S})$. Second, $e(q(n), \mathcal{S})$ is a reduced form representation of the tip rate times the revenue the worker generates. Since tip rates and revenue are likely non-decreasing in both personal attention and facial symmetry, this simplification has little bearing on the conclusions here. Third, social norms on tipping will likely be reflected in the intercept of a functional approximation of the relationship between tip rates and personal attention.

We assume that personal attention is decreasing in the number of customers a worker serves, $\frac{\partial q(n)}{\partial n} < 0$, an assumption whose sign we can verify empirically. We assume the decrease in personal attention happens at an increasing rate, $\frac{\partial q^2(n)}{\partial n^2} < 0$.

Our assumptions on the relationship between n and q are based on the realities of the setting. q is decreasing in n because workers work a fixed number of hours, and increasing

[11] We assume personal attention is implied by customer volume in order to obtain predictions about customer volume. We do this because we observe customer volume precisely in our data and because our measures of personal attention are fuzzy at best. As long as personal attention and customer volume move in opposite directions, however, the predictions of the model are unchanged.

the number of tables served in a fixed amount of time necessarily requires workers to spread themselves more thinly. The second assumption says the opportunity cost of volume is increasing in volume. This is analogous to the standard convex cost of effort assumption.

2.1. A First Test of Substitutability. We illustrate the trade off a worker with facial symmetry $\mathcal{S}$ faces when deciding on how many customers to serve and, by implication, the attention to provide each. The optimal choice $n^* = n(\mathcal{S})$ satisfies

$$e\big(q(n^*), \mathcal{S}\big) + n^* \frac{\partial e(q(n^*), \mathcal{S})}{\partial q} \frac{\partial q(n^*)}{\partial n} = 0 \tag{2}$$

The first term represents the gain in total earnings implied by a marginal increase in the number of customers served. The second term represents the loss in total earnings from the marginal decrease in personal attention, that itself follows from serving more customers. We can re-write Equation (2) as:

$$-\frac{\partial e\big(q(n^*), \mathcal{S}\big)}{\partial q} \frac{1}{e(q(n^*), \mathcal{S})} = \frac{1}{n^* \frac{\partial q(n^*)}{\partial n}} \tag{3}$$

Define $\epsilon_{e,n}(\mathcal{S}) \equiv -\frac{\partial e(q(n^*),\mathcal{S})}{\partial q} \frac{1}{e(q(n^*),\mathcal{S})}$ and $\epsilon_{q,n} \equiv \frac{1}{n^* \frac{\partial q}{\partial n}}$. A graphical depictions of $\epsilon_{e,n}(\mathcal{S})$ and $\epsilon_{n,t}$ is found in Figure 1. The figure illustrates the optimal solution for two types of workers, one with facial symmetry $\mathcal{S}''$ and the other with $\mathcal{S}'$, where $\mathcal{S}'' > \mathcal{S}'$, and where where symmetry substitutes for personal attention ($\frac{\partial^2 e(q(n),\mathcal{S})}{\partial q \partial \mathcal{S}} < 0$). The case when symmetry complements personal attention generates the opposite pattern for $n^*(\mathcal{S}') > n^*(\mathcal{S}'')$.

The model offers a lens for interpreting the data. If we observe workers with symmetric faces serving more customers while delivering less attention to each, this is evidence of facial symmetry and personal attention being substitutes in both earnings and consumer demand. On the other hand, if workers with more symmetric faces serve fewer customers while delivering more attention to each, this is evidence that facial symmetry and personal attention are complements in earnings, and therefore in consumer demand.[12]

[12] One can see this algebraically via the envelope theorem

$$\frac{\partial n^*}{\partial \mathcal{S}} \propto^+ \frac{\partial e}{\partial \mathcal{S}} + n^* \frac{\partial^2 e}{\partial q \partial \mathcal{S}} \frac{\partial q}{\partial n},$$

where $\propto^+$ means "positively proportional to".

Figure 1: Symmetry and Personal Attention are Substitutes

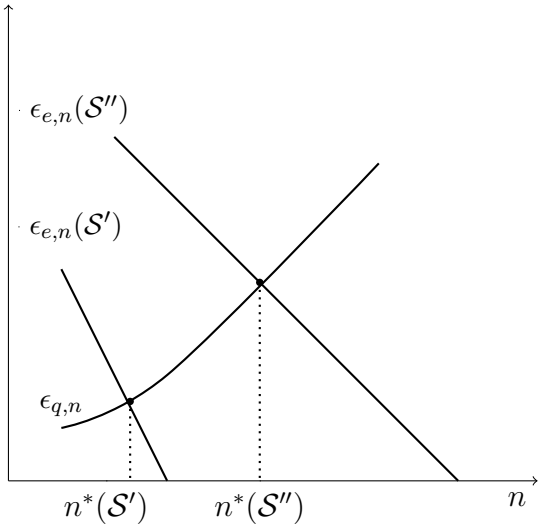

2.2. Empirical Content of Model. The model assumes symmetry has no direct impact on customer volume. A direct impact arises if workers solicit their own customers, either from the street or after the consumer arrives. In our setting, however, workers cannot solicit consumers from the street because the franchises are located in areas which are not easily walkable. Workers cannot solicit consumers after they arrive because of the quasi-random matching procedure. No direct impact on customer volume allows us to rule out the possibility that symmetry and customer volume are complements in generating personal attention.[13] This is useful because this kind of complementarity also generates the prediction that customer volume is increasing in symmetry.

Our test of substitutability is based on the proposition that if symmetry and personal attention are substitutes, then symmetrical workers should serve more customers. The limitation is that in the absence of functional form assumptions on the earnings function, the proposition is not an if and only if statement. As such, our prediction is also consistent with weak complementarities between symmetry and personal attention in consumer demand. For this reason, and because of any remaining doubts about the assumptions of our model, later we shall present a test for substitutability that only relies on the raw data alone. There we will show that the raw data also supports our claim that symmetry and attention are substitutes in consumer demand for services.

[13]In principle, facial symmetry could affect customer perceptions of how personal attention depends on customer volume. Namely, it could be that for workers with more symmetric faces, customers perceive a smaller reduction in personal attention when the worker serves more customers.

3 Data

Our empirical analysis rests on two samples of production data, from a couple of large franchises, belonging to the same North American corporation. The data includes information on worker earnings and production (revenue per shift), task performance, as well as rough proxies of the inputs they direct to each task. These outcomes all have natural counterparts in the model.

We first analyze observational data from a sample that runs from October of 2008 until May of 2009 and includes October 2009.[14] The sample has just under 40 workers and more than 3300 (worker-calendar date) observations. We then analyze data generated by the field experiment. The sample uses 65 workers, of which many come from the second franchise. This sample runs from October of 2008 until May of 2009, and from September of 2009 until May of 2010. It also has more than 3000 observations. As we proceed, we will explain the justification for each sample. For now, descriptive statistics are found in Table 1.[15]

Task performance is measured by customer volume and revenue per customer. We label the category 'task performance' because the firm uses these tasks, particularly revenue per customer, in worker performance assessments. Moreover, the tasks measure performance in the intermediate stage between the production of revenue and the actual inputs of workers. Actual inputs are measured by average price per item sold, number of items sold, and various measures of the time spent with customers. Collectively, these input measures provide a rough proxy for the attention workers give to customers. Accordingly, the last column of Table 1 reports conditional correlations between tip rates on the various input measures.

This tip data comes from the firm's point-of-sale software. The software tracks tips paid by credit or debit card but not by cash. The missing information is a minor obstacle. Almost 80 percent of bills are paid by credit or debit. Moreover, because the matching process is exogenous, workers are just as likely to serve customers who pay by credit as they are to serve customers who pay by cash.

3.1. Measuring Personal Attention. Most menu items come in several sizes. If the worker can sell the customer on the larger size, they will generate a higher price for each item they sell. Menu items are also often packaged together, so that buying the package is

[14] Data from summer months was not collected because we were told there are substantial slow downs in customer volume in these months. As a result, the opportunities for trading earnings per customer for the number of customers a worker serves are largely negated.

[15] One concern with our data relates to the sizes of our cross sections. While we refer to small sample issues throughout the paper, we deal with the issues specifically in Appendix A.2.

Table 1: **Summary Statistics for Worker Behavior, Performance, and Earnings.** The last column reports coefficients from a regression of tip rates on our measures of personal attention. The measures are standardized in the regressions. The coefficients therefore represent the effect of a one standard deviation increase in the measure of interest. Regressions include fixed effects for the calendar date and tables the worker is responsible for, the number of customers served, and the start time for the bill. Standard errors are clustered on the worker-calendar date combination and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

Variable	Description	Count	Mean	Standard Deviation	Effect on Tip Rate
Production					
Revenue	Revenue for the Shift	3770	970.57	401.10	
Task Performance					
Revenue From Each Customer	Revenue per Customer	3743	43.68	7.28	
Customers	Number of Customers	3770	22.53	9.65	
Inputs					
Average Price	Revenue per Item Sold	3770	5.97	0.81	-0.21*** (0.07)
Items Sold	Items Sold per Customer	3743	7.37	1.19	0.14** (0.06)
Time With Customer	Minutes Between First Order and Bill Settlement	3770	87.33	19.71	0.32*** (0.07)
Time to Linger	Minutes Between Last Dessert Order and Bill Settlement	3657	41.20	19.49	0.12* (0.07)
Time Between Customers	Minutes Between Bill Settlement and First Order on Next Bill	3205	23.84	14.39	0.00 (0.07)
Earnings					
Tip Rate		3715	14.06	2.32	
Hourly Wage (before April 2009)			7.60		
Hourly Wage (after)			8.25		
Hourly Tip Earnings		3346	26.09	23.88	
Hours of Work		3397	5.35	1.48	

less expensive than buying the items separately. If the worker can sell the customer on a bundle, they will sell more items to each customer they serve. To these ends, Average Price and Items Sold provide measures of efforts to upsell and bundle items respectively.

The remaining input variables describe how workers allocate their time. Time with Customer measures how much time workers spend on each bill. Time To Linger measures how much time they take to deliver the cheque. Finally, Time Between Customers measures the time it takes to reseat tables the worker is responsible for. All three variables are noisy by definition. The times with and between customers are noisy because we do not know when the customer first sat down. The time to linger is noisy because we do not know when the bill was delivered. Because of these data limitations, we use, as replacements, information on when the first order was taken and on when the bill was settled. Further to all this, time with customers and time to linger only provide rough proxies for the contact that actually takes place between the worker and customers.

The last column of Table 1 shows positive correlations of tip rates with Items Sold, Time with Customer, and Time to Linger, a negative correlation with Average Price, and no correlation with Time Between Customers. The positive correlations are consistent with customers rewarding bundling, attentiveness, and not being rushed. The negative correlation with Average Price is consistent with customers punishing workers for sales of more expensive items. No correlation with Time Between Customers is consistent with it only affecting customers though the time they spending waiting for a seat, and not necessarily through the attention they get from workers. It is important to keep in mind that these are correlations. They may simply reflect customer heterogeneity, as would be the case for example if heavy eaters pay higher tip rates.

3.2. Measuring Symmetry. We photographed workers in October 2009. The photographs were taken one half hour before the start of an already scheduled shift in order to ensure that workers were in uniform (cleaned and pressed white shirts, black pants, a tie) when the picture was taken. Interviews were mostly conducted at the same location to ensure that the photos were uniform in background lighting and background colors. To obtain measures of facial symmetry from the photographs, we use software that measures 8 geometric proportions at different points on the face. The software calculates the absolute difference between each of these proportions and the ideal.[16] The sum of these differences is then used to cal-

[16]The proportions are: 1. Distance between eyes/Nose width; 2. Head height/Face height; 3. Face height/(Face height - Chin height); 4. (Face height - Chin height)/Mouth width; 5. Face width in the mouth area/Nose length; 6. Head height/Face width in the eye area; 7. Face width in the eye area/(Face

culate the person's percentile in the population. In Online Appendix Figure A3 we present the photographs for two workers, one at the lower end of the distribution, the other at the top. According to the algorithm, the worker in the first photograph is in 42.5^{th} percentile and the worker in the second is in the 99.9^{th} percentile. A histogram of the full symmetry distribution can be found in Online Appendix Figure A1.

Table 2 summarizes the symmetry distribution as well as other personal information we collected. The facial symmetry of workers in our sample is well above the median in the population at large. Only 16 percent of the population have better looks than the average worker in our sample. In this sense, our empirical results are only generalizable to the population of faces if consumer preference for symmetry is linear in symmetry. If, as we assume in the model in the previous section, there are diminishing returns to symmetry, our estimates are in fact a lower bound on the true effect a one standard deviation increase has on customer and worker behavior.

Our primary interest, however, is in identifying the role of symmetry in worker-customer interactions for comparable jobs in the services sector. In many of these jobs, where workers regularly interact with customers, the distribution of facial symmetry is truncated or, at a minimum, skewed to the left. In this way, our results speak to our population of interest.

We focus on facial symmetry rather than facial attractiveness because we believe it broadens the applicability of our study. Specifically, it has a stronger link with other traits we often care about. Like these traits (race and gender for example), symmetry is basically immutable. It is difficult for workers to manipulate their endowment of symmetry. With attractiveness or, more specifically, measures of attractiveness, this is not necessarily the case. Workers can manipulate attractiveness with fresh haircuts, makeup, etc.

With that said, we had 80 evaluators rate worker facial attractiveness on a 5-point scale: significantly below average/below average/average/above average/very attractive. The evaluators inspected all the photos prior to assigning ratings. The rating is an assessment, therefore, of facial attractiveness relative to the other workers in our sample. We condition all our regressions on the average of the 80 assessments for each worker.[17]

3.3. Measuring Correlates of Symmetry. One challenge with understanding the effects and origin of a consumer preference for symmetry relates to the effects of other traits, such

height - Chin height); 8. Face height/Forehead height. The ideal is given by the golden section or divine proportion, which equals 1.61803398875. More information about the software can be found at http://www.facebeautyrank.com/index.html.

[17]The raw correlation between symmetry and attractiveness is 0.17. The low correlation is consistent with the mixed evidence on whether symmetry predicts attractiveness [Thornhill and Gangestad, 1999].

Table 2: **Attributes of Workers.** *** indicates $p < 0.01$, ** indicates $0.01 < p < 0.05$, and * indicates $p < 0.1$. The estimates in the second to last column are from multivariate OLS regressions of facial symmetry on the other worker attributes.

Category	Variable Name	Description	Count	Mean	Standard Deviation	Regression Coefficient	Standard Error
	Facial Symmetry	Facial Symmetry (Percentile)	38	84.16	19.94	-	-
Confidence	Expected Tip Rate		38	14.40	1.53	-2.90	2.42
	Expected Revenue Per Consumer		38	42.77	5.99	0.42	0.82
Social and Communication	Fun to Sit With	Number of coworkers reporting worker was someone they wanted to sit with at social gatherings	36	6.61	5.16	3.79	2.62
	Others Consider a Friend	Number of coworkers reporting worker was a friend	36	4.64	3.91	-4.59	3.46
	Friends at Work	Number of coworkers who are friends	38	4.47	2.70	-0.14	1.62
	Socializes with Coworkers	Times socialized with coworkers outside work last month	38	10.18	15.58	0.16	0.32
	Likes Working with Friends	Prefers working with friends	38	3.89	1.11	6.15*	3.44
Demographic and Employment	=1 if Male		38	0.55	0.50	-4.05	8.58
	=1 if Visibly Overweight		38	0.21	0.41	-5.79	12.44
	=1 if Still in School		37	0.35	0.48	-5.36	7.92
	Industry Experience (Months)		38	70.82	59.09	-0.02	0.11
	Older than 30		38	0.16	0.37	6.70	16.01
	Full Time		38	0.11	0.31	-18.59	21.74
	Days in Sample		38	63.09	51.17	-0.09	0.10

as the ability, joviality, friendliness, and confidence of the worker. These traits can confound the effects of symmetry. To address this challenge, we collected several additional pieces of information through interviews with workers. Statistics describing the information we collected are found in the remaining rows of Table 2.

Workers were asked about social relations with coworkers. They reported the coworkers they consider a friend, the number of times they socialized with coworkers outside work, and their preference for working with friends. Workers consider 4.47 coworkers a friend, socialized 10.18 times with coworkers outside work in the previous month, and have a preference for working with friends.[18] The variables give us a sense of the worker's sociability. They also capture biases such as overconfidence in self perceptions of sociability.

We use the interview questions to construct sociability measures based on coworker responses. As we asked workers to identify coworkers they consider a friend, we can also determine for each worker the number of coworkers who considers them a friend. The variables together measure (among other things) the sincerity of the worker. We also asked workers to identify coworkers they prefer to sit with at social gatherings. We use this information to determine for each worker the number of coworkers identifying them as 'fun to sit with'. The 'fun to sit with' variable measures more superficial aspects of sociability, such as the joviality of the worker. Table 2 shows that on average 4.64 coworkers report the worker is a friend and 6.61 report they are 'fun to sit with' at social gatherings.

In the interviews we solicited worker beliefs about the revenue they generate and the tip rates they receive.[19] The beliefs were used to derive worker expectations about revenue per customer and tip rates. Workers who believe they can upsell and deliver better than average attention should expect to sell more to each customer they serve. Expected revenue per customer thus measures worker confidence in their ability to convince customers to buy more than they otherwise would have. To the extent that convincing customers depends on the ability to communicate, it also measures the confidence of the worker in their communication skills. Workers who believe they can better personalize attention should expect higher tip rates from customers. While expected revenue per customer measure worker confidence in the ability to upsell and deliver personal attention, expected tip rates measure worker confidence in the ability to deliver personal attention. If, on the other hand, workers believe that beyond social norms tip rates are a matter of luck, then expected tip rates are a measure of worker optimism. The proxies are presented in rows 2 and 3 of Table 2. Row 2 gives the tip

[18] A five-point scale (where one is strongly disagree and five is strongly agree) was used to measure the preference for working with friends.

[19] The process for constructing beliefs follows from [Manski, 2004].

rate the worker expects to receive, while row 3 gives the revenue (per customer) the worker expects to make.[20] A comparison of the expectations in Table 2 with the measurements in Table 1 suggests workers forecast revenue and tips fairly accurately.

4 Baseline Results

The observational data alone allows for a simple, though indirect, test of the substitutability of symmetry for personal attention in generating earnings from customers. We estimate the relationship between the facial symmetry of the worker and how many customers they serve and compare the estimate with the relationship theory predicts. If the estimates show workers with symmetric faces serve more customers, we conclude symmetry and personal attention are substitutes in consumer preferences.

4.1. Empirical Specification. The baseline specification for our empirical analysis is:

$$y_{id} = \beta_{\mathcal{S}}\mathcal{S}_i + \mathbf{Z}_i\boldsymbol{\gamma} + \mathbf{X}_{id}\boldsymbol{\Gamma} + \gamma_d + \epsilon_{id}. \tag{4}$$

y_{id} measures outcomes for worker i on date d,[21] such as revenue, revenue per customer, customer volume, and tip rates. $\mathcal{S}_i$ is measures the facial symmetry of the worker. $\mathbf{Z}_i$ includes other important attributes of the worker such as rater assessments of their facial attractiveness, their confidence, social and communication skills, gender, and employment characteristics. The attributes help capture unobserved differences in worker ability that correlate with facial symmetry and outcomes observed in the data. They are particularly useful when workers with less symmetric faces are hired because they have personalities that compensate for the lack of symmetry.

$\mathbf{X}_{id}$ includes the an indicator of whether the worker often works on the day of the week that corresponds with date d, and a measure of their perception of how busy it usually is on that day. The fixed effects for the calendar date, γ_d, control for aggregate shocks to outcomes, such as daily changes in team composition or in the consumers that visit the firm. The random variable ϵ_{id} measures daily shocks to performance that are specific to the individual. We assume that $E[\epsilon_{id}|\mathcal{S}_i, \mathbf{Z}_i, \mathbf{X}_{id}, \gamma_d] = 0$. We cluster standard errors at the level of the worker to further account for correlations between facial symmetry and other

[20] The reported values for expected revenue and tip rates are averages of the revenue per customer and tip rates the worker expects to receive on a party of two and on a party of four.

[21] There is only one shift - the dinner shift - per date.

unobserved worker traits. Since we evaluate the effects of symmetry on several outcomes, we also report Bonferonni-Holm p-values for hypothesis testing with multiple outcomes.

Identification is facilitated by quasi-random matching of workers and customers, wherein well-defined rules imply customers cannot choose who serves them, workers cannot choose who they serve, and managers cannot, at least directly, choose who serves whom. Worker looks and other traits should therefore be exogenous to customer traits conditional on worker start times and table assignments.

There is a tension, however, between these rules and the incentives of managers and the firm more generally. If workers with symmetric faces generate value for the customer and more revenue for the firm, at the same cost, then the firm will have incentives to expose more customers to symmetrical workers and to take advantage of how many customers they serve. Alternatively, these incentives can arise because managers themselves simply have a preference for symmetric faces. We assess whether workers with symmetric faces receive more favorable work assignments, and why. We discuss the implications of favorable assignments for the empirical strategies used throughout the paper.

4.2. Quasi-Random Matching? The top panel of Table 3 reports estimates of regressions relating worker facial symmetry to various work assignments. Columns 1 through 6 describes work assignments decided upon before a shift begins. Columns 7 and 8 describe assignments that happen in real time, over the course of a shift.

The estimate in Column 1 of the top panel implies symmetrical workers are assigned favorable start times. All else equal, workers with symmetric faces start 28.8 ($p < 0.01$) minutes earlier than their more asymmetric coworkers. Early starts give workers with symmetric faces better opportunities to serve more customers.

The estimates in Columns 2 to 6 of the top panel imply symmetrical workers are also assigned favorable tables. Column 2 shows the firm assigns them 1.63 ($p < 0.05$) more booth seats. Booth seats are favorable because they offer customers more privacy than benches or chairs. Column 5 shows the firm assigns symmetrical workers 1.47 ($p < 0.05$) more seats in total. Column 6 suggests the firm assigns these workers tables that historically turn over more quickly, though the estimate is statistically insignificant at conventional levels.[22] Both

[22]The turnover rate for a table is the number of customers served divided by the number of seats. The rate for the worker's assigned section was computed as follows. First, we calculated one-year averages of the turnover rate for each table. Second, the one-year averages were used to rank tables by their turnover rate. The average ranking reduces concerns that rates reflect the productivity of workers typically assigned to the table, rather than productivity of the table itself. Third, for each shift, we computed an average ranking for the tables assigned to the worker. The average is the dependent variable in Column 6 of Table 3.

estimates imply workers with symmetric faces are given better opportunities to serve more customers.

The estimates in Columns 7 and 8 show symmetry has a statistically negligible effect on the unused capacity of the worker, as measured by the share of seats that are left empty (Column 7), and assignments of customers with a history of visiting the firm (Column 8). The estimates suggest assignments of customers to workers are based on factors other than worker facial symmetry, supporting our claim that matches are exogenous to symmetry. This is unsurprising in light of the incentives and constraints the firm faces. Selective assignments violate rules governing the process that matches customers with workers. Moreover, selective assignments are costly to the firm. They increase the time it takes to sort through line ups for seating, requiring the assignment of hundreds of customers to several workers over the span of a few hours, and by consequence increases the wait time for customers. Longer waits deter consumers from waiting and, in addition, lower perceptions of the overall service quality by customers who continue to wait. Although customers who wait might perceive better service quality because they are served by someone who is more aesthetically pleasing, they perceive worse service quality because they had to wait longer to interact with this worker.

In the bottom panel of Table 3 we assess whether symmetrical workers receive favorable assignments because they are more productive or because managers favor them. We assess this by showing how the effects of symmetry change when we account for the relative productivity of the worker. Our proxy for relative productivity is the cumulative revenue the worker generated, across the (previous) shifts they worked, divided by the average cumulative revenue of coworkers working the same shift. The basic idea is that if the effect disappears once we control for relative productivity, then symmetrical workers receive favorable assignments because they are more productive. If the effect persists then it is because managers themselves favor workers with symmetric faces.[23]

The bottom panel implies symmetrical workers receive favorable assignments mostly because they are more productive. The estimate in Column 1 shows that conditioning on relative productivity cuts the effect of symmetry by almost half. Instead of starting 28.8 minutes earlier, symmetrical workers start 15.0 ($p < 0.01$) minutes earlier. At the same time, the estimates in Columns 2 and 5 show that once we condition on relative productiv-

[23]The goal with including our proxy for relative productivity is to shed light on the path between facial symmetry and the allocation of work assignments by managers. We are not saying there is no causal effect of symmetry on work assignments. Rather, we just want to know if the effect operates through relative productivity. The distinction is important because our proxy (and others we considered) can inevitably be interpreted as outcomes of worker facial symmetry (*c.f.* Angrist and Pischke [2009]).

Table 3: **Exploiting Workers Who Have it Easy.** All regressions control for rater assessments of facial attractiveness, calendar date fixed effects, the demographic and employment traits of the worker, their social and communication skills, confidence, the number of sample days worked, an indicator for whether the worker usually works that day (of the week), the worker's ranking of how busy the day is (compared to other days of the week). The regressions in Columns 2 through 8 control for the start time of the worker. Columns 7 and 8 include fixed effects for the tables the worker is responsible for. Also note that the estimates in columns 7 and 8 provide tests of the (identifying) assumption that assignments of customers to workers are formulaic. Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$. Bonferonni-Holm p-values for hypothesis testing with multiple outcomes are in square parentheses.

	Assignments by Managers Before Shift Begins						After Shift Begins	
	Start Time	Tables					Empty	Repeat
	(Minutes)	Booth Seats	Bench Seats	Chair Seats	Total Seats	Historical Turn Rate	Seats (Share)	Customers (Share)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Facial	-28.8***	1.63**	-0.04	-0.13	1.47**	0.06	0.01	0.00
Symmetry	(0.11)	(0.77)	(0.17)	(0.14)	(0.71)	(0.04)	(0.02)	(0.01)
	[0.00]	[0.00]	[1.00]	[1.00]	[0.00]	[0.42]	[1.00]	[1.00]
Observations	3351	3351	3351	3351	3351	3343	3351	3351
R^2	0.51	0.29	0.18	0.12	0.28	0.54	0.39	0.26
Facial	-15.0***	0.88	-0.15	-0.18	0.55	0.05	0.01	0.00
Symmetry	(0.07)	(0.66)	(0.16)	(0.13)	(0.53)	(0.04)	(0.02)	(0.01)
	[0.00]	[0.14]	[0.44]	[0.51]	[0.37]	[0.51]	[1.00]	[1.00]
Historical Revenue	-40.2***	2.37***	0.36**	0.16	2.89***	0.04	-0.01	-0.01
Relative to Coworkers	(0.11)	(0.74)	(0.18)	(0.15)	(0.59)	(0.06)	(0.04)	(0.01)
Observations	3351	3351	3351	3351	3351	3343	3351	3351
R^2	0.53	0.30	0.18	0.12	0.30	0.54	0.39	0.26
Dep. Var. Mean	5:08pm	10.50	1.50	2.98	14.98	2.38	0.94	0.32
and Std. Dev.	(73 min)	(7.35)	(1.77)	(3.50)	(6.04)	(1.04)	(0.57)	(0.20)

ity, the effect of symmetry on table assignments disappears. Its disappearance is consistent with informal discussions we held with managers, wherein many stressed the importance of table assignments for maximizing revenue.

Two patterns imply managers themselves favor symmetrical workers, however. The first is that relative productivity negates the effects of symmetry on table or customer assignments. From the standpoint of discrimination by managers, the negation of these effects is intuitive. On top of being important instruments for maximizing revenue, it is more difficult for individual managers to exploit these assignments to discriminate against workers, as these assignments are usually done in consultation with other managers and support staff. The second pattern is the effect of facial symmetry only survives the inclusion of relative productivity when the dependent variable is the worker's start time. Unlike assignments of tables and customers, start times are decided in advance of the work week, each time by a single manager. In other words, start times are more susceptible to whims of individual managers. As such, if managers are going to discriminate, start times are the way to go.

The estimates in Table 3 have implications for empirical specifications based on observational data. Including controls for managerial assignments in our specifications increases the chances that isolate differences in how workers react to a consumer preference for symmetry. Excluding these controls allow us to include the influences of managers in our calculations, on top of the influences of worker reactions to customers, a potentially important source of discrimination in the workplace. Our preferred specification includes the influences of managers because, in our view, it best serves a major goal of the paper, namely to understand how having it easy shapes earnings and productivity differentials in the workplace. Nonetheless in most tables we will present specifications with and without these controls.

4.3. The Trade Offs Workers Make. Columns 1 through 5 of Table 4 report earnings per customer differentials for workers whose facial symmetry is one standard deviation above and below the average. Columns 6-10 do the same but for customer volume. For each dependent variable, as we move across the columns, we add controls in order to examine how effect sizes change when we account for fixed personality traits of the worker.

Column 6 shows workers with more symmetrical faces serve 2.88 ($p < 0.05$) more customers than workers with more asymmetrical faces. At the same time, symmetrical workers earn \$0.23 ($p < 0.01$) less for each. The trade off is a bit less pronounced once we account for social and communication skills. Columns 2 and 7 show workers with more symmetric faces serve 2.42 ($p < 0.01$) more customers while earning about the same (\$0.24 less) for

Table 4: **The Trade-Offs Workers Make.** The table reports outcome differentials for workers whose facial symmetry is one standard deviation above and below the average. Regressions include fixed effects for the calendar date, controls for the demographic and employment characteristics of the worker, an indicator for whether the worker usually works that day (of the week), and the worker's ranking of how busy the day is (compared to other days of the week). Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$. Bonferonni-Holm p-values for hypothesis testing with multiple outcomes are in square parentheses.

	Earnings Per Customer					Customers				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Facial	-0.23***	-0.24***	-0.24***	-0.23***	-0.19***	2.88**	2.42***	2.42***	3.15***	3.28***
Symmetry	(0.08)	(0.09)	(0.08)	(0.07)	(0.06)	(1.29)	(0.93)	(0.91)	(0.83)	(0.80)
	[0.00]	[0.00]	[0.00]	[0.01]	[0.02]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Fun to Sit With		-0.01	-0.01	-0.01	-0.01		0.35	0.34	0.21	0.24
		(0.02)	(0.03)	(0.03)	(0.02)		(0.27)	(0.28)	(0.24)	(0.24)
Others Consider		0.01	-0.00	0.00	0.01		0.27	0.29	0.54	0.53
a Friend		(0.03)	(0.03)	(0.03)	(0.03)		(0.36)	(0.38)	(0.34)	(0.34)
Friends at Work		0.01	0.01	0.01	-0.00		-0.46**	-0.46**	-0.46***	-0.45***
		(0.01)	(0.01)	(0.01)	(0.01)		(0.21)	(0.21)	(0.18)	(0.17)
Socializes with		0.00	0.00	0.00	0.00		-0.08**	-0.09**	-0.08***	-0.09***
Coworkers		(0.00)	(0.00)	(0.00)	(0.00)		(0.03)	(0.04)	(0.03)	(0.03)
Likes Working		-0.02	-0.01	-0.02	-0.05**		-0.48	-0.51	-0.79**	-0.77**
with Friends		(0.03)	(0.03)	(0.03)	(0.02)		(0.41)	(0.41)	(0.36)	(0.37)
Expected Revenue			0.01**	0.01*	0.01			-0.03	-0.11	-0.12
			(0.01)	(0.01)	(0.01)			(0.07)	(0.07)	(0.08)
Expected Tip Rate				0.02	0.01				0.88***	1.09***
				(0.03)	(0.03)				(0.31)	(0.29)
Facial					0.14***					0.05
Attractiveness					(0.04)					(0.49)
Observations	3491	3383	3383	3383	3319	3540	3417	3417	3417	3352
R^2	0.16	0.18	0.18	0.18	0.18	0.50	0.54	0.54	0.55	0.55

each. The change in effect size from Column 6 to 7 implies symmetrical workers are more social but that, in and of itself, sociability translates into fewer customers being served.

Columns 3, 4, 8 and 9 examine the role of worker expectations in our estimates.[24] As Column 9 shows, there is a large change in the estimated effect of symmetry once we account for the tip rates workers expect. The estimate implies symmetrical workers serve 3.15 more customers than their asymmetrical coworkers. The change in effect size suggests workers with symmetric faces have low expectations concerning the tip rates they earn, and that more customers are served by workers with high expectations.

Columns 5 and 10 examine whether and how the estimated symmetry coefficients change if we partial out rater assessments of facial attractiveness. Column 5 shows the earnings per customer differential decreases slightly to 19 ($p < 0.01$) cents, while Column 10 shows the customer volume differential increases slightly to 3.28 ($p < 0.01$) customers. The estimates imply the symmetry coefficients pick up something other than what is captured by the rater assessments. Altogether, the evidence, together with the model prediction, suggests workers can substitute the symmetry of their face for personal attention when generating earnings.

To this point, our conclusions about substitutability rest on the empirical content of the simple theory in Section 2. One concern with the theory is that it assumes personal attention is decreasing the number of customers served. This is a problem if workers who provide good service are also the ones who serve more customers. We examine the merits of this assumption empirically, by specifically studying the role facial symmetry plays in the delivery of personal attention. The last 5 columns of Table 5 reports estimates of its effect on measures of worker inputs. The bottom row reports estimates of the effect of the input on tip rates. This row provides some indication of whether whether we can equate more effort or time with better service.

Symmetrical workers sell more expensive items (Column 6), bundle fewer items (Column 7), spend less time with customers (Column 8), and the deliver the cheque more quickly (Column 9). The estimates suggest symmetrical workers are less attentive and are thus consistent with a negative relationship between the personal attention and customer volume of the worker. The estimate in Column 10 implies it takes 1.87 ($p < 0.01$) more minutes to reseat their tables. This suggests these workers are less diligent when it comes to keeping tables clean, consistent with their enjoying more leisure on the job.[25]

[24] We recognize that there is the potential for the so-called 'y on y' problem for the regression described in Column 4. Namely, the regressions include the revenue per customer and tip rates a worker expects to generate, of which the product strongly correlates with the dependent variable. We included the variables

Table 5: **Symmetry and Production.** The table reports the difference in outcomes for workers whose facial symmetry is one standard deviation above and below the average. Regressions control for rater assessments of facial attractiveness, the confidence of the worker, their social and communication skills, fixed effects for the calendar date, controls for the demographic and employment characteristics of the worker, an indicator for whether the worker usually works that day of the week, and the worker's ranking of how busy the day is (compared to other days of the week). Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$. Bonferonni-Holm p-values for hypothesis testing with multiple outcomes are in square parentheses.

	Overall Revenue (1)	Work Hours (2)	Overall Earnings (3)	Earnings Per Customer (4)	Customers (5)	Average Price (6)	Items Sold (7)	Time With Customers (8)	Time to Linger (9)	Time Between Customers (10)
					Panel A: Without Controls for Assignments by Managers					
Facial Symmetry	116.20***	0.09	15.16***	-0.17**	3.10***	0.06	-0.25**	-5.01***	-1.96**	1.87***
	(34.66)	(0.14)	(5.59)	(0.07)	(0.84)	(0.06)	(0.10)	(0.95)	(0.78)	(0.31)
	[0.00]	[0.25]	[0.00]	[0.13]	[0.00]	[0.28]	[0.00]	[0.00]	[0.13]	[0.09]
Observations	3352	3012	2982	3319	3352	3352	3346	3352	3282	2876
R^2	0.52	0.64	0.52	0.18	0.55	0.18	0.18	0.21	0.11	0.33
					Panel B: With Controls for Assignments by Managers					
Facial	40.21***	-0.20***	1.92	-0.10	1.13***	0.09	-0.21**	-4.12***	-2.09**	1.62***
	(13.04)	(0.07)	(2.00)	(0.07)	(0.25)	(0.06)	(0.10)	(1.16)	(0.91)	(0.29)
	[0.02]	[0.00]	[0.45]	[0.45]	[0.00]	[0.30]	[0.03]	[0.00]	[0.15]	[0.18]
Observations	3351	3012	2982	3318	3351	3351	3345	3351	3281	2876
R^2	0.71	0.76	0.69	0.21	0.73	0.21	0.20	0.30	0.13	0.37
Effect of Input on Tip Rate						-0.21***	0.14**	0.32***	0.12*	0.00
						(0.07)	(0.06)	(0.07)	(0.07)	(0.07)

4.4. Managers Make it Easier. The bottom panel of Table 5 assesses how accounting for favorable assignments by managers alters the estimates of the symmetry coefficients. It reports re-estimates of our main regressions where the regressions include controls for start times and table assignments.

Table 5 suggests managers set symmetrical workers up to do better than others. A comparison of the top and bottom panel implies managers have autonomy over the sensitivity of customer volume to the time and effort allocation of the worker. Specifically, the comparison reveals that assignments by managers have large impacts on the number of customers workers serve, and small impacts on allocations of time and effort. The estimate in Column 5 of the bottom panel shows symmetrical workers serve 1.13 ($p < 0.01$) more customers than their coworkers, whereas the estimate in Column 5 of the top panel shows they serve 3.10 ($p < 0.01$) customers. By contrast, Column 8 of the top panel shows symmetrical workers spend 5.01 ($p < 0.01$) fewer minutes with customers, whereas than their coworkers, whereas the estimate in Column 8 of the bottom panel equals 4.12 ($p < 0.01$) fewer minutes. The estimates in Columns 6,7,9, and 10 of the top and bottom panels are quite similar, on the other hand.

The first 3 columns of Panels A and B explore the implications of assignments by managers for the overall revenue, hours, and earnings of workers. The top panel shows symmetrical workers earn 15 ($p < 0.01$) dollars more than their coworkers, despite working the same hours. At the same time, these workers each produce 116 ($p < 0.01$) dollars more revenue for the firm. The estimate for earnings translates into a 8.4 percent difference, while the estimate for revenue translates into a difference of 12 percent (per worker). The bottom panel shows symmetrical workers earn 2 ($p > 0.1$) dollars more than their coworkers, while having shifts that are approximately 12 ($p < 0.01$) minutes shorter. They produce 40 ($p < 0.01$) dollars more revenue for the firm. The estimates suggest the earnings and productivity differentials are driven largely by the allocation of work assignments by managers.

5 Experimental Evidence of Substitutability

Our first test of substitutability assumes workers are fully characterized by the attributes we observe. In this section we derive an additional test, which allows for the possibility

because they have little impact on our estimates of the regressions.

[25] Online Appendix Tables A1 and A2 study the trade off in more depth, showing that the trade off between time with customers and customer volume is more severe for symmetrical workers on busy days, and for symmetrical workers who have more experience.

that workers differ along various unobserved dimensions such as ability. The test draws on a field experiment that exogenously rewards workers with bonuses (in addition to their tips and wages) for customer volume. Because the experiment rewards workers for serving more customers it exogenously increases the opportunity cost of personalized attention. When workers provide more attention, they now tradeoff the bonus as well as earnings from serving more customers. Using the theory we develop in Section 2, we show that if attention and symmetry are substitutes, then the response to the bonus is larger for workers with symmetric faces. They increase the number of customers they serve by more than their less symmetric coworkers. This theoretical result forms the foundation of the second test: If we observe the bonus elicits a larger increase in customer volume for workers with symmetric faces, then we can further infer that facial symmetry and attention are substitutes.[26]

5.1. Research Design. We briefly explain the field experiment here, focusing on design features which are most relevant for the present study. The original purpose of the field experiment was to study agency problems related to multitasking. Since the purpose was incidental to the facial symmetry of workers, conclusions regarding the effects of symmetry are somewhat immune to confounding factors, such as placebo effects, experimenter demand effects, and hawthorne effects. A detailed explanation of the design can be found in (redacted).

The experiment paid workers bonuses for customer volume on busy days in November, January, and May of the 2009-2010 season.[27] The experimental design yields three control groups: busy days from months in the 2009-2010 season where bonuses were off the table, all busy days in the treated restaurant from the 2008-2009 season, and all busy days from the control franchise.[28] There were no bonuses in any of the control groups. Workers were paid tips and hourly wages as usual.

[26]The experiment offers a second advantage. It allows us to investigate whether there is enough slack for workers to respond to favoritism by customers. If outcomes change during the experiment, then it would suggest that workers had slack in the first place.

[27]The experiment was run on busy days (Fridays and Saturdays, in the fall, winter, and spring) because, on these days, the trade offs workers make can result in substantial losses for the firm. On busy days, there are always line ups for seating. When workers focus on personal attention, it takes longer to seat customers. As a result of the long wait times, a lot of customers decide to go elsewhere.

[28]The second franchise is located about an hour away (by car) from the first. The franchises belong to the same corporation and have the same owners. They offer the identical products, charge identical prices, follow the same employment practices, including the rules that assign customers to workers. Even the physical layouts are similar. Furthermore, on average, workers at both franchises generate the same revenue per customer, and serve the same number of customers. The main difference between the franchises is that, as a whole, the treated franchise serves more customers.

The experiment paid workers a bonus when the number of customers they served exceeded an exogenously determined performance standard.[29] The bonus was proportional to the difference between the standard and the number of customers the worker served. The bonus rate (proportion) was chosen so that workers would earn between $20 and $30, or more than 10% of average daily earnings, for exceeding the standard by a standard deviation.[30]

The bonus rate was always the same for all workers. In November and January of the 2009-2010 season, the standard was also the same for all workers. In May of the same season, each worker was assigned their own standard.

Tailored standards help us equalize differences in the marginal incentives to earn the bonus. When the performance standard is the same for everyone, the bonus provides stronger marginal incentives for workers who normally serve lots of customers. For them, earning the bonus is easier. Part of this relates to the table assignments of workers, as workers are sometimes assigned tables that turn over more quickly. When they are, extra effort yields larger gains in the number of customers the worker serves. For this reason, we tailored the standards to the (already-determined) table assignment of the worker. It was set higher when the worker was assigned tables that were historically good for turnover. We set it lower when the tables were historically bad for turnover.

5.2. A Second Test of Substitutability. We revisit the model of Section 2 to illustrate why the performance incentive allows for another test of the substitutability of symmetry for personal attention. Worker earnings from the experiment can be written as:

$$E(s, n) = (e(q(n), \mathcal{S}) + \alpha)n \tag{5}$$

where $\alpha \geq 0$ captures the added incentive for customer volume. Workers with $\alpha = 0$ belong to the control group for the field experiment. The larger is α, the larger the opportunity cost of personal attention. As such, the experiment allows us to hold facial symmetry fixed and study how workers change their behavior in response to an exogenous endowment of α.

From Equation (5), it follows that a worker with symmetry $\mathcal{S}$ chooses $n^* = n(\mathcal{S}, \alpha)$ such

[29]The performance standard was calculated using the following steps: 1. compute long-run averages for the number of customers served, hours worked, and section size on high-demand days; 2. divide long-run customer turnover by long-run hours worked; 3. divide the number from the second step by long-run average section size. The steps yield a performance standard of 0.4 on Fridays and 0.41 on Saturdays.

[30]More specifically, workers were paid $3 for each tenth of a point above the performance standard.

that it satisfies:

$$\frac{\partial e(q(n^*),\mathcal{S})}{\partial q}\frac{q}{e(q(n^*),\mathcal{S})+\alpha} = -\frac{1}{\frac{n^*}{q}\frac{\partial q}{\partial n}} \tag{6}$$

Noting the analogy to the optimality condition presented in Equation (3) in Section 2, we define $\epsilon_{e,n}(\alpha) \equiv \frac{\partial e(q(n^*),\mathcal{S})}{\partial q}\frac{1}{e(q(n^*),\mathcal{S})+\alpha}$ and $\epsilon_{q,n} \equiv -\frac{1}{n^*\frac{\partial q}{\partial n}}$. Imposing substitutability between symmetry and attention allows for a simple prediction concerning how worker responses differ depending on the facial symmetry of the worker. Specifically, assume $e = q(n) + \rho S$, where $\rho > 0$ converts percentiles of facial symmetry into dollars per customer.[31][32] Also assume for purposes of tractability that $\frac{\partial^3 q}{\partial n^3} = 0$.[33]

The assumptions imply

$$\frac{\partial n^2}{\partial \alpha \partial \mathcal{S}} = \frac{3\frac{\partial^2 q}{\partial n^2}\rho}{(\frac{\partial^2 E(n,\alpha)}{\partial n^2})^3} > 0. \tag{7}$$

If facial symmetry substitutes for service quality then workers with more symmetric faces increase their customer volume by more than their less symmetric coworkers. If we therefore observe workers with symmetric faces increasing customer volume by more than their coworkers then we can infer that facial symmetry substitutes for service quality. In addition, the magnitude of the substitution parameter governs the strength of the differential response of workers with symmetric faces to the bonus. The easier it is for workers to substitute symmetry for service quality, the larger is the response for workers with symmetric faces.

5.3. Empirical Specification and Experimental Results. In order to evaluate differences in worker responses to bonuses, we assume outcomes for worker i, at franchise f, on calendar date d, are generated according to:

$$y_{ifd} = \alpha_i + (\beta_T + \beta_{\mathcal{S}}\mathcal{S}_i)I_f I_d + \gamma_d + \mathbf{X}_{ifd}\mathbf{\Gamma}_1 + \mathbf{Z}_{fd}\mathbf{\Gamma}_2 + \epsilon_{ifd}, \tag{8}$$

α_i captures time-invariant (ability) differences across workers. No worker works at both franchises. The fixed effects for workers of a franchise, in turn, add up to a fixed effect for

[31] This degree of substitutability is not necessary but is sufficient for the result.

[32] The parameter ρ therefore makes facial symmetry commensurate with the bonus rate.

[33] The prediction is unchanged if $\frac{\partial^3 q}{\partial n^3} \leq 0$.

the franchise. As before, S_i measures facial symmetry and γ_d is a fixed effect for the calendar date. I_f and I_d are binary variables indicating whether the franchise was treated and whether the date was treated. Since α_i encapsulates the influences of symmetry, it obviates the need to control for S_i alone.

Our interest is in the parameter $\beta_{\mathcal{S}}$, which measures the differential impact of the bonuses for more and less symmetric workers. When the dependent variable is the number of customers served, we infer symmetry substitutes for attention in consumer demand if $\beta_{\mathcal{S}} > 0$.

$\mathbf{Z}_{fd}$ includes a variable that indicates whether performance standards were tailored to individual workers. As mentioned earlier, the variable allows us to hold fixed the marginal incentives of the bonus. It allows us to condition estimates of $\beta_{\mathcal{S}}$ on differences in the sensitivity of customer volume to the efforts of the worker. $\mathbf{X}_{ifd}$ includes the number of days (in the sample) worked by the worker. The variable lets us condition estimates of $\beta_{\mathcal{S}}$ on the experience the worker has accumulated at the firm.

A major advantage of the experiment is it allows us to examine the effects of symmetry and, at the same time, control for unobserved correlates of symmetry. Specifically, the worker fixed effects, α_i, encapsulate influences of unobserved time-invariant correlates of symmetry, in addition to influences of symmetry itself.[34]

Estimates of Equation (8) are found in Table 6. Column 5 in the top panel shows symmetrical workers serve 4.52 more customers than their asymmetrical coworkers, where the estimate has $p < 0.10$ and $p = 0.01$ without and with the Bonferroni-Holm adjustment. Column 5 in the bottom panel shows symmetrical workers serve 2.96 more customers than their asymmetrical coworkers. The estimate has $p < 0.05$ and $p = 0.11$ without and with the Bonferroni-Holm adjustment. Both estimates are consistent with the theory, and specifically the prediction that the bonuses have larger effects on customer volumes for workers with symmetric faces. The estimates in Column 8 also supports the theory, showing larger reductions in the time symmetrical workers spend with customers. The top panel shows symmetrical workers spend 4.99 fewer minutes with customers when they could earn bonuses for serving more customers. The bottom panel shows a decrease of 4.36. Note, however, that the time with customer estimates are only statistically significant in the absence of the Bonferroni-Holm adjustment.

[34]Another advantage is the experiment allows us to use data from the control franchise to evaluate the effects of symmetry. This was not possible with the observational data because, due to various constraints, we were unable to collect several other pieces of personal information from employees of the control franchise. The experiment helps in this regard because we can use worker fixed effects to soak up the effects of the information we wanted to collect.

Table 6: **Performance Bonuses and the Trade Off at Work.** The table shows how treatment response (to bonuses for serving more customers) differs depending on whether the worker's facial symmetry is above or below average. The table excludes a coefficient for attractiveness alone because our regressions include fixed effects for the worker. The regressions control for differences in marginal incentives to earn the bonuses, and the number of days (in the sample) worked. Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$. Bonferonni-Holm p-values for hypothesis testing with multiple outcomes are in square parentheses.

	Overall Revenue (1)	Work Hours (2)	Overall Earnings (3)	Earnings Per Customer (4)	Customers (5)	Average Price (6)	Items Sold (7)	Time With Customers (8)	Time to Linger (9)	Time Between Customers (10)
					Panel A: Without Controls for Assignments by Managers					
Interaction of Symmetry and Bonus	212.18** (99.23) [0.00]	0.45** (0.20) [0.23]	29.55* (17.29) [0.06]	0.02 (0.22) [1.00]	4.52* (2.36) [0.01]	0.14 (0.14) [1.00]	0.08 (0.18) [1.00]	-4.99* (2.56) [0.94]	3.52 (3.05) [1.00]	-1.60 (1.42) [1.00]
Bonus	48.26 (38.75)	-0.18** (0.08)	14.95** (7.31)	-0.10 (0.14)	1.36 (0.94)	0.02 (0.09)	-0.10 (0.09)	-4.00** (1.59)	0.01 (1.88)	-0.16 (0.92)
Observations	3023	2854	2842	3007	3023	3011	3009	3011	2974	2938
R^2	0.15	0.20	0.22	0.07	0.20	0.07	0.07	0.03	0.02	0.09
					Panel B: With Controls for Assignments by Managers					
Interaction of Symmetry and Bonus	149.49** (59.45) [0.03]	0.25 (0.19) [1.00]	17.77 (12.15) [0.57]	0.03 (0.24) [1.00]	2.96** (1.38) [0.11]	0.17 (0.17) [1.00]	0.04 (0.16) [1.00]	-4.36* (2.31) [1.00]	2.81 (2.78) [1.00]	-1.95 (1.54) [1.00]
Bonus	53.04* (27.50)	-0.23*** (0.08)	13.57** (5.58)	-0.09 (0.13)	1.40** (0.59)	0.00 (0.09)	-0.08 (0.09)	-4.56*** (1.54)	-0.31 (2.02)	-0.13 (1.00)
Observations	3021	2854	2842	3005	3021	3009	3007	3009	2972	2936
R^2	0.52	0.39	0.51	0.11	0.56	0.09	0.11	0.18	0.05	0.15

The estimates of Columns 1 to 3 in the top panel show symmetrical workers earn 30 dollars more than their asymmetrical coworkers, work 27 minutes longer, and generate 212 dollars more in revenue for the firm. The earnings and revenue estimates are statistically significant with and without the Bonferroni-Holm adjustment. The analogous estimates in the bottom panel show symmetrical workers earn 18 dollars more than their asymmetrical coworkers, work 15 minutes longer, and generate 149 dollars more in revenue for the firm. The revenue differential has $p < 0.05$ with and without the Bonferroni-Holm adjustment. As with the observational data, the experimental estimates support the conclusion that managers set symmetrical workers up to do better than their more asymmetrical coworkers.

The results here ultimately imply that performance pay can play a central role in the relationship between discrimination and inequality in the workplace. If discrimination is a root cause of differences in productivity, and performance pay aims to better align earnings with productivity, then performance pay can amplify the effects of discrimination on inequality in the workplace (*cf.* [Lemieux, MacLeod, and Parent, 2009]).

6 The Origins of Having it Easy

We investigate the origins of the substitutability in consumer demand. We first examine the raw data for evidence that symmetry and personal attention are substitutes. In particular, we conduct a third test for substitutability, one that does not rely on the theory to interpret patterns in the data. Next we explain that substitutability arises because of a consumer preference for symmetry, rather than because symmetry bridges gaps in the information sets of consumers. We then assess whether the preference arises because of a taste for symmetry or because of a distaste for asymmetry. Finally, we evaluate whether consumers value proportionality or horizontal symmetry - the extent to which one side of the face mirrors the other.

6.1. Substitutability in the Raw Data. While we have provided complementary evidence that is consistent with the theory being correct, namely that personal attention and customer volume move in opposite directions, atheoretical evidence of substitutability makes our conclusions far more compelling. Our atheoretical test for substitutability boils down to evaluating the effects on tip rates of interactions between symmetry and personal attention. Specifically, we estimate variants of the specification:

$$\tau_{ibd} = \beta_{\mathcal{S}}\mathcal{S}_i + \beta_q q_{ibd} + \beta_{\mathcal{S}q}\mathcal{S}_i q_{ibd} + \mathbf{Z}_i\boldsymbol{\gamma} + \mathbf{X}_{ibd}\boldsymbol{\Gamma} + \gamma_d + \epsilon_{ibd} \tag{9}$$

where i indicates the worker, b the bill, and d the calendar date. τ_{ibd} is the tip rate, $\mathcal{S}_i$ is the facial symmetry of the worker, and q_{ibd} is service quality. As before, $\mathbf{Z}_i$ includes other important attributes of the worker such as rater assessments of their facial attractiveness, their confidence, social and communication skills, gender, and employment characteristics. $\mathbf{X}_{ibd}$ includes an indicator for whether the worker usually works that day (of the week), and the worker's ranking of how busy the day is (compared to other days of the week). It also includes the start time for the bill, table fixed effects, and number of customers on the bill. These variables help us control for heterogeneity in consumer preference.[35] In one of the specifications, $\mathbf{X}_{ibd}$ includes the service quality at neighbouring tables, for bills that end before b does. We do this to control for spillovers from the attention other customers are receiving. γ_d are fixed effects for the calendar date. The random variable ϵ_{ibd} measures bill-specific shocks relating to, for example, remaining differences in the budgets and preferences of consumers. We assume that $E[\epsilon_{ibd}|\mathcal{S}_i, q_{ibd}, \mathbf{Z}_i, \mathbf{X}_{ibd}, \gamma_d] = 0$.

The context has inherent properties which support the assumption. For one, neither workers, nor customers, nor managers have control over who matches with whom. The property alleviates concerns that an effect is an artefact of the quality of the match. For two, inputs are determined before tips are paid. The property somewhat alleviates concerns that tip rates affect our measures of personal attention, and not the other way around. For three, as noted elsewhere, we collected information on the tip rates workers expect to earn. The information allows us to condition our estimates on anticipatory behavior on the part of workers. Note that, even with these advantageous properties, the assumption $E[\epsilon_{ibd}|\mathcal{S}_i, q_{ibd}, \mathbf{Z}_i, \mathbf{X}_{ibd}, \gamma_d] = 0$ can fail if, for example, workers adjust their attention to unobserved characteristics such as the mood of the customer.

The unit of analysis warrants further comment. The unit of analysis is the bill, whereas for specifications we estimated elsewhere it was the worker-date combination. We have good reason for this. Our first tests required information on the performance of workers to draw inferences about substitutability in consumer demand. To this end, accurate measures of worker performance were important. We did not treat the bill as the unit of analysis in our

[35] We tried in other regressions to control for additional forms of heterogeneity, such as the means by which the bill was paid (by credit or debit card). The results of those regressions were the same as the results that we report here. We excluded the regressions from the paper because of space considerations and because of concerns that the payment represents a bad control.

first tests because worker performance is measured inaccurately at this level of aggregation. At this level, our performance measures reflect idiosyncrasies in whims of customers, as well as the performance of workers. We have better accuracy at the worker-date level because, together with the exogenous procedure that matches workers and customers, aggregation averages customer idiosyncrasies out of our measures of worker performance.

By contrast, the test here requires information on decisions by customers to draw inferences about whether symmetry substitutes for attention in consumer demand. To this end, accurate measures of customer decisions are important. In the bill data, the tip rate already measures a decision by the customer. This, in turn, makes aggregation unnecessary.

The conclusions of our test hinge on the sign for $\beta_{\mathcal{S}q}$. If large q is indicative of more personal attention, and $\beta_{\mathcal{S}q} < 0$, we conclude q substitutes for symmetry in generating tip rates. If large q is indicative of less personal attention, then $\beta_{\mathcal{S}q} > 0$ implies it substitutes for symmetry in tip rates. Average effects for the specification in Equation (9) are found in Table 7. The interaction effects measure the difference between workers whose facial symmetry is one standard deviation above and below the average. For the moment, the discussion focuses on the estimates in Column 4.

The interaction effect estimates for items sold, time with customers, and time to linger are statistically insignificant at conventional significance levels. The sole exception is the interaction effect for average price, which is marginally significant at the 10 percent level. Column 4 shows tip rates are lower when prices are high, but that the reduction is smaller for workers with symmetric faces. Specifically, for a below average worker, a one standard deviation increase in price is associated with a tip rate which is 0.37 percentage points lower. For above-average workers, the tip rate is only 0.15 percentage points lower. While this direct evidence suggests that symmetry substitutes for attention in consumer preference, it is important to keep in mind that the evidence is marginal from a statistical standpoint.

6.2. A Preference for Symmetry or Incomplete Information? Our goal here is to understand whether substitutability in consumer demand originates in a preference for symmetry or in the value of symmetry as a signal of other attributes customers care about. In a world where there are gaps in what they know about workers, customers can use the visible attributes of workers to draw inferences about invisible attributes. Customers in our case, for example, might use facial symmetry to infer the accuracy and reliability of advice from the worker. The advantage of tip data is that it allows us to sort through the origins of substitutability. It allow us to do so because tips are paid only after the exchange of prod-

Table 7: **Substitutabilty in the Raw Data.** Regressions include rater assessments of facial attractiveness, the start time of the bill, and the number of customers, the number of bills (in the sample) handled by the worker, fixed effects for the table and for the calendar date, controls for the demographic and employment characteristics of the worker, an indicator for whether the worker usually works that day (of the week), and the worker's ranking of how busy the day is (compared to other days of the week). Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Tip Rate				
	(1)	(2)	(3)	(4)	(5)
Facial Symmetry	-0.29***	-0.16*	-0.08	-0.07	-0.07
	(0.10)	(0.08)	(0.09)	(0.09)	(0.13)
Average Price			-0.25***	-0.26***	-0.23***
			(0.06)	(0.06)	(0.06)
Items Sold			0.11**	0.11**	0.15**
			(0.04)	(0.04)	(0.06)
Time With Customers			0.19***	0.19***	0.17***
			(0.05)	(0.05)	(0.06)
Time to Linger			-0.01	-0.01	0.03
			(0.04)	(0.04)	(0.04)
Interaction of Symmetry and					
Average Price				0.22*	0.22*
				(0.12)	(0.11)
Items Sold				-0.01	0.04
				(0.07)	(0.08)
Time With Customers				-0.02	-0.08
				(0.10)	(0.09)
Time to Linger				-0.04	-0.06
				(0.05)	(0.05)
Social Skills and Confidence		✓	✓	✓	✓
Service Quality at Neighbouring Tables					✓
Observations	17881	17299	16646	16646	10230
R^2	0.05	0.05	0.05	0.05	0.07

ucts (between the firm and customer) stops. Once this happens, there is no longer a need for customers to make use of the information symmetry conveys.[36] As a result, if workers with symmetric faces earn higher tip rates than their coworkers, despite giving customers

[36]Having tips paid after completion of product exchanges rules out statistical discrimination by customers possessing correct beliefs about the worker's product knowledge or trustworthiness. It also rules out statistical discrimination by customers possessing incorrect beliefs, such as considering workers with symmetric faces more competent than they are [Mobius and Rosenblat, 2006].

less attention, then it is likely the result of a preference for symmetric faces.[37] Such a finding would thus imply substitutability in consumer demand originates in a consumer preference for symmetry.

Evidence from the observational data is found in Table 7. Moving across the columns, we show how the coefficient on symmetry changes as we add controls for service quality and other factors affecting the relationship to tip rates. We also show how rewards for personal attention differ depending on worker facial symmetry.

Columns 1 to 3 of Table 7 imply customers are cognizant of less personal attention from symmetrical workers. Without controls (Columns 1 and 2), the tip rates for symmetrical workers (1σ above average) are 0.2 to 0.3 percentage points lower than for their asymmetrical coworkers (1σ below average). Once we control for correlates of personal attention, the effect of symmetry becomes statistically negligible.

Columns 4 to 5 of Table 7 imply consumers have a preference for symmetry, but that the preference manifests itself in differential rewards for personal attention. Both columns show the direct effect of symmetry is negligible. Both also show that, while they dislike prices, customers dislike them less when they interact with symmetrical workers.

To examine whether the experimental data implies a preference for symmetry, we consider how the effects of the bonuses on tip rates differs across workers with more or less symmetric faces. Having presented evidence that bonuses elicit less personal attention from symmetrical workers, if we observe no decrease in tip rates then we infer consumers have a preference for symmetric faces. The absence of a reduction in tip rates implies customers punish symmetrical workers less for poor performance.

Evidence from the experimental data is found in Table 8. Column 1 displays estimates where we exclude controls for the personal attention of the worker. While the estimate in Column 1 is supposed to capture the effects of the interaction between the bonus and attention, we evaluate in the remaining columns how the estimates change as we add controls for various aspects of service quality. The estimates in Table 8 show that symmetrical workers earn the same tip rate despite delivering less, in a comparative sense, personal attention to customers. As with the observational data, the experimental data implies consumers having a preference for symmetry.

[37] Most studies examining the sources of discrimination rely on tests that allow them to rule out discrimination based on incomplete information [Altonji and Pierret, 2001]. In this regard, several experimental studies find that outward traits matter less as agents accumulate information ([Todorov et al., 2005], [Andreoni and Petrie, 2008]). While the tests allow the authors to rule out incomplete information as a source for discrimination, they do not necessarily imply preferences is not the source.

Table 8: **Evidence of a Consumer Preference for Symmetry.** Regressions include fixed effects for the worker. They also include the start time of the bill, the number of customers on the bill, the number of bills (in the sample) handled by the worker, and fixed effects for the calendar date. Finally, the regressions control for differences in marginal incentives to earn the bonuses. Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Tip Rate				
	(1)	(2)	(3)	(4)	(5)
Interaction of Symmetry and Bonus	0.15	0.14	0.14	0.15	0.15
	(0.21)	(0.21)	(0.23)	(0.23)	(0.24)
Bonus	0.05	0.06	0.05	0.06	0.03
	(0.16)	(0.16)	(0.16)	(0.16)	(0.16)
Average Price on Items Sold		-0.33***	-0.18***	-0.21***	-0.21***
		(0.03)	(0.04)	(0.04)	(0.04)
Items Sold per Customer			0.19***	0.13***	0.14***
			(0.03)	(0.03)	(0.03)
Time (Minutes) With Customers				0.18***	0.19***
				(0.03)	(0.04)
Time (Minutes) to Linger					-0.00
					(0.02)
Observations	40306	40306	39360	39360	38487
R^2	0.04	0.04	0.04	0.04	0.04

6.3. The Good, the Bad, and the Average. The summary statistics for tip rates (Table 1) raise concerns about basing conclusions on the results reported in Table 7. Consistent with intuition, the summary statistics show the distribution for tip rates has light tails. Much of the mass is concentrated between tip rates of 9 and 19 percent. This fact alone raises concerns about using normal distributions to approximate distributions of our test statistics. The approximations are reasonable if the residuals from conditional tip rate regressions are normally distributed. Knowing this, we ran the residuals through several diagnostics tests. None implied normality. Furthermore, in conducting the tests, two properties stood out. First, a normal distribution has broader shoulders than the distribution for the residuals. In other words, the residuals have more concentration near the mean of the distribution. Second, a normal distribution has less mass in the tails.[38]

There are intuitive reasons for the shape of the distribution for the residuals. First, concentration at the mean is consistent with social norms (or rules of thumb) about tipping. Second, the masses in the tails are consistent with customers differing in their perceptions of these norms. Some customers view 15 percent as the norm, some view 10 as the norm, some view it as 20, etc. Moreover, some customers base the tip rate on the before-tax bill. Others base it on the after-tax bill. Both reasons, along with the fact that we have information from tens of thousands of bills, suggest we should examine the effects of symmetry at different points of the conditional tip rate distribution.[39]

Figure 2 plots results from quantile regression estimates of the specification in Column 4 of Table 7. The estimates come from regressions at the $5^{th}, 10^{th}, ..., 95^{th}$ quantiles of the tip rate distribution. We focus on estimates for the coefficients on price and its interaction with facial symmetry because of space considerations and because this was the only interaction with a statistically significant and robust effect. Figures 2(a) and 2(b) plot estimates for the coefficients on price and the interaction.

Figure 2(b) shows that good, bad, and average tippers all punish symmetrical workers less for selling high-priced items. The interaction effects of price and symmetry on tip rates range from about 0.06 percentage points (at the 65^{th} quantile) to 0.5 percentage points (at the 95^{th}). Figure 2(a) shows, on the other hand, the average (negative) effect of price largely originates in the lower and upper quantiles of the conditional tip rate distribution. Interestingly, price has a statistically insignificant effect on the tip rates of average tippers -

[38]The distribution for the residuals was also asymmetric. This is unsurprising because tip rates are bounded from below, but not from above.

[39]At the same time, such an analysis would allow us to speak to whether the results are driven by outliers in tip rates.

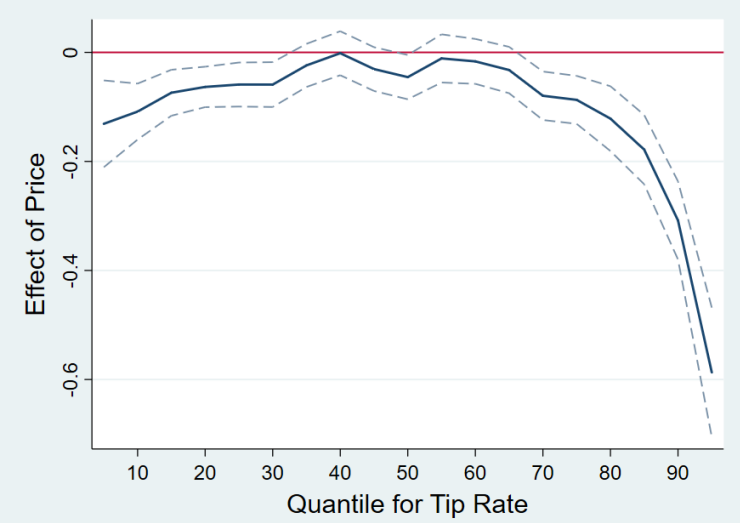

(a) Effect of average price.

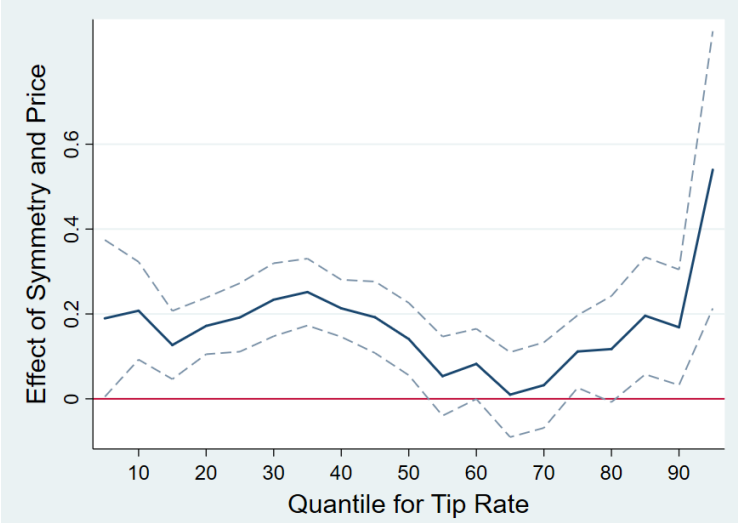

(b) Interaction effect of facial symmetry and average price.

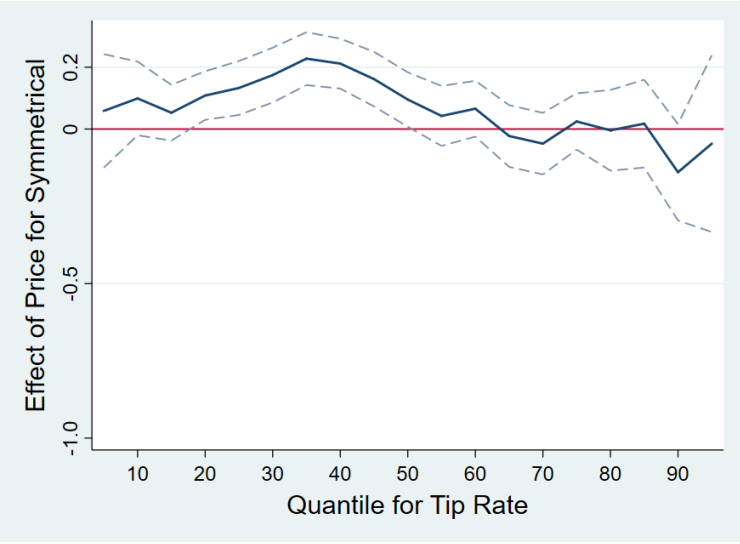

(c) Effect of Price for Symmetrical Workers.

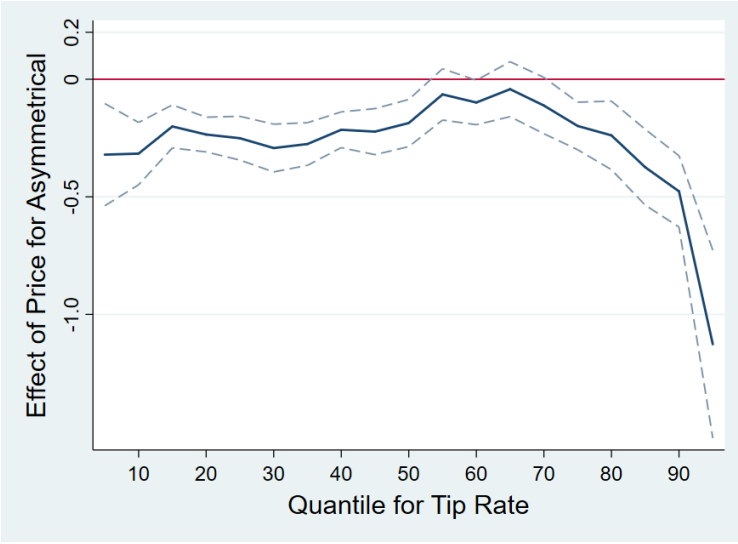

(d) Effect of Price for Asymmetrical Workers.

Figure 2: **Punishments for Sales of High-Priced Items.** Figure 2(b) plots the effects of the interaction between facial symmetry and average price on various quantiles (5,10,...,95) of the conditional tip rate distribution. Figure 2(a) plots the effects of average price alone. The solid lines are for coefficient estimates. The dashed lines are for the 90 percent confidence intervals. The figures show customers punish workers for sales of high-priced items, but that punishments are less severe for workers with symmetric faces.

ones between the 35^{th} and 65^{th} percentile of the conditional tip rate distribution.

6.4. A Taste for Symmetry or a Distaste for Asymmetry? Figures 2(c) and 2(d) let us study whether the differential effects of price on tip rates arise because customers like symmetry, or because they dislike asymmetry. Specifically, the figures plot the price effect for workers whose faces are one standard deviation above the average, and for workers whose faces are one standard deviation below.

Figures 2(c) and 2(d) show price and symmetry have a positive interaction effect for three reasons. First, all customers dislike asymmetry. Figure 2(d) shows that for asymmetrical workers, price has a negative effect at all quantiles of the conditional tip rate distribution. Second, some customers like symmetry. Figure 2(c) shows that for symmetrical workers, price has a positive effect between the 15^{th} and 45^{th} quantiles of the conditional tip rate distribution. Third, while customers in the top quartiles punish all workers for selling high-priced items, they punish asymmetrical workers more. Upwards of the 85^{th} quantile, price has a much stronger negative effect on the tip rates of asymmetrical workers than it does on the tip rates of their coworkers.[40]

7 Conclusion

The main takeaways from the present article are as follows. First, the evidence suggests consumers have preference for symmetric faces, as they are willing to trade off a bit less personal attention for service by a person with a more symmetric face. Second, this preference generates a comparative advantage for symmetrical workers. Since personal attention substitutes for symmetry in consumer preference, symmetrical workers can give customers less personal attention and instead focus on generating customer volume, allowing them to generate more revenue for the firm and more earnings for themselves. Third, the productivity gap encourages managers to treat preferred workers favorably, as managers place these workers in situations that take advantage of the additional revenue they generate. Our evidence shows that by doing this, managers widen productivity and earnings gaps among workers substantially.

The trade off that we analyze strongly resembles trade offs commonly found in the services sector, the largest and fastest growing economic sector in the industrialized world. The

[40] Our principal measure of symmetry captures deviations from the ideal face, as defined by the so-called golden ratio or divine proportion. In Appendix A.1, we show that our results are robust to using a different measure of symmetry, namely the extent to which one side of the face mirrors the other.

findings apply specifically to jobs where workers have contact with customers, and where there is a tension between the personal attention each customer receives and the number of customers the worker serves: residential real estate agents, salespersons at retailers such as Best Buy, Tax Agents at H&R Block, or the family doctor. The findings have less applicability to jobs where the tension is less pronounced: cashiers at supermarkets, for whom service quality (the speed of service) and customer volume move together; commercial real estate agents, attorneys at large law firms, celebrity sports agents, and surgeons, for whom customer volume is less important than the quality of the service they deliver to each.

References

Joseph G. Altonji and Charles R. Pierret. Employer learning and statistical discrimination. *The Quarterly Journal of Economics*, 116(1):313–350, February 2001.

James Andreoni and Ragan Petrie. Beauty, gender and stereotypes: Evidence from laboratory experiments. *Journal of Economic Psychology*, 29(1):73–93, 2008.

Joshua Angrist and Jorn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton, 2009.

Ofer Azar. Why pay extra? tipping and the importance of social norms and feelings in economic theory. *Journal of Behavioral and Experimental Economics (formerly The Journal of Socio-Economics)*, 36(2):250–265, 2007.

Revital Bar and Asaf Zussman. Customer discrimination: Evidence from israel. *Journal of Labor Economics*, 35(4):1031–1059, 2017.

Gary Becker. *The Economics of Discrimination*. University of Chicago Press, Chicago, 1957.

Michele Belot, V. Bhaskar, and Jeroen van de Ven. Beauty and the sources of discrimination. *Journal of Human Resources*, 47(3):851–872, 2012.

Jeff E. Biddle and Daniel S. Hamermesh. Beauty, productivity, and discrimination: Lawyers' looks and lucre. *Journal of Labor Economics*, 16(1):172–201, January 1998.

George J. Borjas. Self-selection and the earnings of immigrants. *The American Economic Review*, 77(4):531–553, September 1987.

George J Borjas and Stephen G. Bronars. Consumer discrimination and self-employment. *The Journal of Political Economy*, 97(3):581–605, June 1989.

Anne Case and Christina Paxson. Stature and status: Height, ability, and labor market outcomes. *Journal of Political Economy*, 116(3):499–532, 2008.

Gordon B. Dahl. Mobility and the return to education: Testing a roy model with multiple markets. *Econometrica*, 70(6):2367–2420, November 2002.

Andrew D. Foster and Mark R. Rosenzweig. Comparative advantage, information and the allocation of workers to tasks: Evidence from an agricultural labor market. *Review of Economic Studies*, 63:347–374, 1996.

Dylan Glover, Amanda Pallais, and William Pariente. Discrimination as a self-fulfilling prophecy: Evidence from french grocery stores. *Quarterly Journal of Economics*, 132(3): 1219–1260, 2017.

Claudia Goldin and Cecilia Rouse. Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review*, 90(4):715–741, November 2000.

Daniel S. Hamermesh and Jeff E. Biddle. Beauty and the labor market. *The American Economic Review*, 84(5):1174–1194, December 1994.

Barry Harper. Beauty, stature and the labour market: A british cohort study. *Oxford Bulletin of Economics and Statistics*, 62(s1):771–800, 2000.

James J. Heckman and Bo E. Honore. The empirical content of the roy model. *Econometrica*, 58(5):1121–1149, September 1990.

James J. Heckman and Guilherme Sedlacek. Heterogeneity, aggregation, and market wage functions: An empirical model of self-selection in the labor market. *The Journal of Political Economy*, 93(6):1077–1125, December 1985.

James J. Heckman, Jora Stixrud, and Sergio Urzua. The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics*, 24 (3):411–482, 2006.

Holzer and Ihlanfeldt. Customer discrimination and employment outcomes for minority workers. *The Quarterly Journal of Economics*, 113(3):835–867, 1998.

Kristen P. Jones, Dave F. Arena, Christine L. Nittrouer, Natalya M. Alonso, and Alex P. Lindsey. Subtle discrimination in the workplace: A vicious cycle. *Industrial and Organizational Psychology*, 10(1):51–76, March 2017.

Peter Kuhn and Kailing Shen. Gender discrimination in job ads: Evidence from china. *The Quarterly Journal of Economics*, 128(1):287–336, 2013.

Thomas Lemieux. Estimating the effects of unions on wage inequality in a panel data model with comparative advantage and nonrandom selection. *Journal of Labor Economics*, 16 (2):261–291, April 1998.

Thomas Lemieux, W. Bentley MacLeod, and Daniel Parent. Performance pay and wage inequality. *The Quarterly Journal of Economics*, CXXIV(1):1–49, 2009.

Jonathan S. Leonard, David I. Levine, and Laura Guiliano. Customer discrimination. *The Review of Economics and Statistics*, 92(3):670–678, August 2010.

Charles F. Manski. Measuring expectations. *Econometrica*, 72(5):1329–1376, September 2004.

Robert Miller. Job matching and occupational choice. *The Journal of Political Economy*, 92(6):1086–1120, December 1984.

Markus M. Mobius and Tanya S. Rosenblat. Why beauty matters. *American Economic Review*, 96(1):222–235, March 2006.

Naci Mocan and Erdal Tekin. Ugly criminals. *The Review of Economics and Statistics*, 92 (1):15–30, 2010.

Matthew Mulford, John Orbell, Catherine Shatto, and Jean Stockard. Physical attractiveness, opportunity, and success in everyday exchange. *The American Journal of Sociology*, 103(6):1565–1592, May 1998.

Clark Nardinelli and Curtis Simon. Customer racial discrimination in the market for memorabilia: The case of baseball. *The Quarterly Journal of Economics*, 105(3):575–595, 1990.

Derek A. Neal and William R. Johnson. The role of premarket factors in black-white wage differences. *The Journal of Political Economy*, 104(5):869–895, October 1996.

David Neumark, Roy J. Bank, and Kyle D. Van Nor. Sex discrimination in restaurant hiring: An audit study. *The Quarterly Journal of Economics*, 111(3):915–941, August 1996.

S.P. Perry, M.C. Murphy, and J.F. Dovidio. Modern prejudice: Subtle, but unconscious? the role of bias awareness in whites' perceptions of personal and other's biases. *Journal of Experimental Social Psychology*, 61:64–78, 2015.

Nicola Persico. Racial profiling? detecting bias using statistical evidence. *Annual Review of Economics*, 1:229–254, 2009.

Nicola Persico, Andrew Postlewaite, and Dan Silverman. The effect of adolescent experience on labor market outcomes: The case of height. *Journal of Political Economy*, 112(5): 1019–1053, October 2004.

A. D. Roy. Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 3(2): 135–146, June 1951.

A. J. Sankoh, M. F. Huque, and S. D. Dubey. Some comments on frequently used multiple endpoint adjustment methods in clinical trials. *Statistics in Medicine*, 16(22):2529–2542, 1997.

John Karl Scholz and Kamil Sicinski. Facial attractiveness and lifetime earnings: Evidence from a cohort study. *The Review of Economics and Statistics*, 97(1):14–28, 2015.

Randy Thornhill and Steven W. Gangestad. Facial attractiveness. *Trends in Cognitive Sciences*, 3(12):452–460, December 1999.

Alexander Todorov, Anesu N. Mandisodza, Amir Goren, and Crystal C. Hall. Inferences of competence from faces predict election outcomes. *Science*, 308(5728):1623–1626, June 2005.

Robert J. Willis and Sherwin Rosen. Education and self-selection. *The Journal of Political Economy*, 87(5):S7–S36, October 1979.

A For Online Publication

Figure A1: **Full Distribution for Primary Symmetry Measure.**

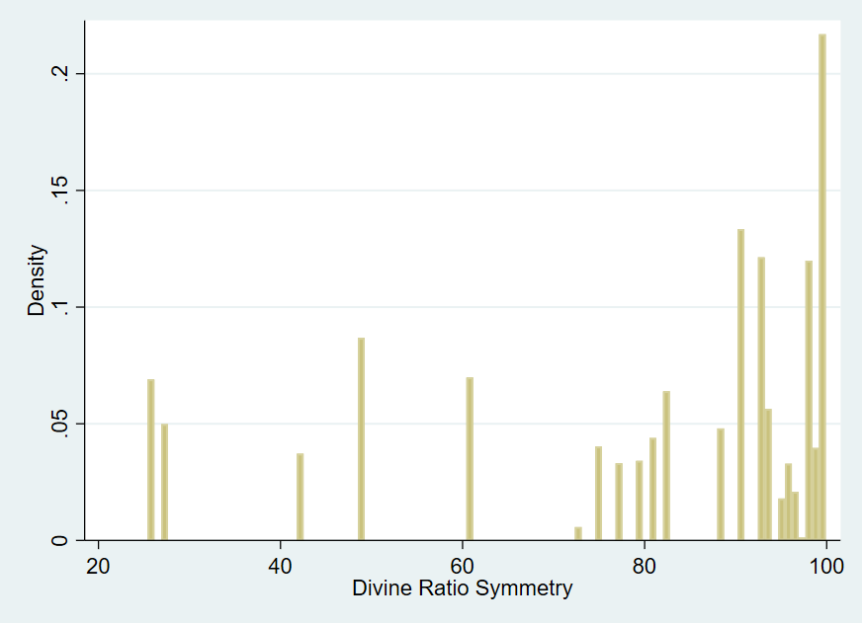

Table A1: **Opportunities for Serving More Customers.** The table reports the difference in outcomes for workers whose facial symmetry is one standard deviation above and below the average. Regressions control for rater assessments of facial attractiveness, the confidence of the worker, their social and communication skills, fixed effects for the calendar date,fixed effects for the calendar date, controls for the demographic and employment characteristics of the worker, an indicator for whether the worker usually works that day (of the week), and the worker's ranking of how busy the day is (compared to other days of the week). Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Slow Days			Busy Days		
	Customers	Time With Customers	Time Between Customers	Customers	Time With Customers	Time Between Customers
	(1)	(2)	(3)	(4)	(5)	(6)
Facial Symmetry	2.93***	-5.40***	1.96***	3.53***	-4.30***	1.02*
	(0.79)	(1.08)	(0.69)	(0.94)	(1.16)	(0.56)
Mean of Dependent Variable	19.50	87.21	28.42	27.31	87.50	18.15
Observations	2088	2088	1634	1264	1264	1242
R^2	0.49	0.24	0.26	0.46	0.18	0.14

Table A2: **Trade offs by Experience.** The table reports the difference in outcomes for workers whose facial symmetry is one standard deviation above and below the average. Days in sample refers to the number of days the worker worked in our sample. 50 is the median days in sample. Regressions control for rater assessments of facial attractiveness, the confidence of the worker, their social and communication skills, fixed effects for the calendar date,fixed effects for the calendar date, controls for the demographic and employment characteristics of the worker, an indicator for whether the worker usually works that day (of the week), and the worker's ranking of how busy the day is (compared to other days of the week). Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Days in Sample < 50			Days in Sample ≥ 50		
	Customers	Time With Customers	Time Between Customers	Customers	Time With Customers	Time Between Customers
	(1)	(2)	(3)	(4)	(5)	(6)
Facial	2.84***	-6.74***	3.15***	3.05***	-2.76***	1.01**
Symmetry	(1.01)	(1.81)	(1.09)	(0.83)	(1.00)	(0.42)
Observations	1182	1182	984	2170	2170	1892
R^2	0.63	0.34	0.49	0.59	0.27	0.37

A.1. Reflections on Symmetry. We consider more precisely what it is that customers value when it comes to the facial features of the worker. Our principal measure of symmetry captures deviations from the ideal face, as defined by the so-called golden ratio or divine proportion. In doing so it aggregates, into a composite measure, reflection symmetry - the extent to which one side of the face mirrors the other - with how well-proportioned the face is (in the classic sense). In this section, we isolate the role of reflection symmetry. We do this both for the purposes of robustness and because many consider reflection symmetry as a barometer for gauging the attractiveness of a person. It measures attractiveness because, as social psychologists postulate, symmetry of two-sided traits carries information about the ability to resist parasites (see, for example, [Grammer and Thornhill, 1994]).[41]

As our principal measure of symmetry subsumes reflection symmetry, we expect a positive correlation between the measures. We illustrate the correlation in Figure A2, which plots the relationship between the measures, after conditioning on controls for rater assessments of facial attractiveness, worker confidence, their social and communication skills, and demographic and employment characteristics.[42] Our estimates of the correlation imply that

[41] To calculate the measure for a given worker, we identify the center of each pupil. We then draw a vertical line through the midpoint of the pupils. Using the midpoint, we calculate left and right deviations at various points of the face. For example, we calculate the outer points of the mouth and, for each point, measure the deviation (in pixels) from the midpoint of our line. This calculation alone results in two measurements. We do this for five more times, using different points of the face. We then add up the absolute value of the twelve deviations from the midpoint. The calculation yields a value that indicates the extent of asymmetries in the face of a worker. For the purposes of our regression analysis, we standardize the measure by subtracting the mean and dividing by the standard deviation. We multiply the standardized value by negative one in order to simplify comparisons with our principal measure of symmetry.

[42] Specifically, we obtained the residuals from separate OLS regressions of our principal measure and

Figure A2: **Reflection Symmetry.** The figure plots residuals from regressions of each symmetry measure on controls for rater assessments of facial attractiveness, the social skills of the worker, their demographic and employment characteristics, an indicator for whether the worker usually works that day (of the week), and the worker's ranking of how busy the day is (compared to other days of the week). The y-axis presents the primary measure considered in the paper. The x-axis presents the measure we use to check the robustness of the results.

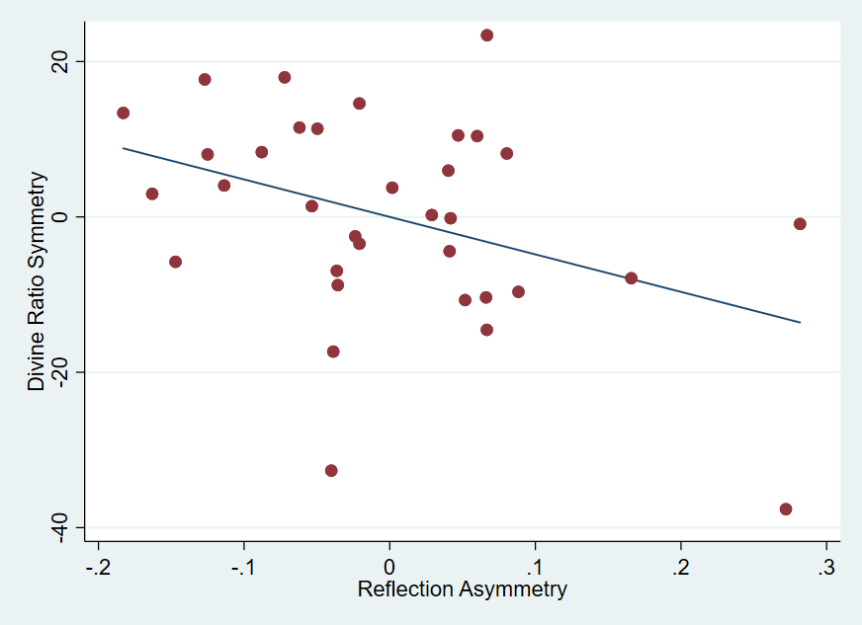

a one unit increase in reflection asymmetry translates into a 40.53 percentile reduction in our principal measure, where the estimate is statistically significant at the ten percent level.

In order to evaluate precisely the facial features customers value, we take the following approach. In a first step, we partial reflection symmetry out of our principal measure of symmetry. We regress our principal measure on reflection symmetry and obtain the residuals. The residuals are by definition the part of our measure that is not in the reflective symmetry measure, thereby providing a net measure of the proportionality of the face. We then re-estimate our primary specifications, replacing our principal measure with the residuals from the first stage regression.[43] The results of our approach are found in Table A3.

The results in Table A3 show that our conclusions thus far are robust to the measure of symmetry. Columns 1-2 show workers with well-proportioned faces serve more customers at the expense of their earnings per customer. The remaining columns illustrate that these workers are the ones who deliver less personal attention.

reflection symmetry on the controls. Figure A2 plots the residuals from these regressions.

[43]The conclusions are the same if the regressions use the reflective measure instead of the principal measure.

Table A3: **Reflection Symmetry and Allocations of Effort and Time.** The table reports the difference in time and effort allocations for workers whose facial symmetry (net of reflection symmetry) is one standard deviation above and below the average. Regressions control for rater assessments of facial attractiveness, the confidence of the worker, their social and communication skills, fixed effects for the calendar date, controls for the demographic and employment characteristics of the worker, an indicator for whether the worker usually works that day of the week, and the worker's ranking of how busy the day is (compared to other days of the week). Standard errors are clustered on the worker and are in parentheses. *** for $p < 0.01$, ** for $0.01 < p < 0.05$, and * for $p < 0.1$.

	Earnings From Each Customer	Customers	Average Price on Items Sold	Items Sold to Each Customer	Time With Customer	Time to Linger	Time Between Customers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Net Facial	-0.14*	3.32***	0.06	-0.22**	-4.82***	-1.91**	1.94***
Symmetry	(0.08)	(0.87)	(0.06)	(0.10)	(0.90)	(0.80)	(0.31)
Observations	3319	3352	3352	3346	3352	3282	2876
R^2	0.18	0.55	0.18	0.18	0.21	0.11	0.33

A.2. Sample Size. As noted in [Bloom et al., 2013], three broad issues arise with smaller cross sections (usually much smaller than ours). The first issue relates to statistical power, whether the cross section is large enough to detect significant effects. The second issue is about whether conventional statistical tests are reliable given the size of our cross section. The third issue pertains to whether our sample is representative of a broader population of interest. We will discuss these issues one by one.

Statistical Power. The data has features that help with concerns about power. First, there is little to no measurement error in the production data, as it comes from the point-of-sale software the firm uses. Second, most of our analysis aggregates the raw transactions data, either to the bill level or the shift level. By doing so, we average out some of the idiosyncracies in our performance outcomes that arise because of, for example, idiosyncracies in the preferences of consumers. Third, in addition to having large T, the production data has low autocorrelation. For instance, regressing the revenue of the worker on (just) their revenue from their last shift yields an estimated coefficient of 0.25. The autocorrelation is low because, as is common in retailing, outcomes are volatile. Intuitively, with low autocorrelation, each realization offers new information about the production of the worker. In essence we get several 'independent' observations for each worker. This in turn increases our ability to detect small effects [McKenzie, 2012].

Table A4: **Bootstrapped-based Inferences.** The table reports various p-values for Wald statistics that test for an effect of facial symmetry on the number of customers served or the earnings from each. Columns 1 and 3 report p-values of the Wald statistics for the coefficient estimates in Columns 5 and 10 of Table 4. In that table, the estimated effect of facial symmetry on earnings per customer is -0.19, with a Wald test statistic of -2.98, and the estimated effect on customer volume is 3.28. Its Wald test statistic is 4.09.

	Earnings Per Customer		Customers	
	Cluster Robust	Wild Cluster Bootstrap-t	Cluster Robust	Wild Cluster Bootstrap-t
	(1)	(2)	(3)	(4)
p-value of Wald Statistic	0.003	0.076	0.000	0.043

Statistical Inference. Our main statistical tests are based on the cluster-robust variance estimator and the assumption that each worker is their own cluster. As such, in conducting the tests, we allow for within-worker correlation across the shifts that they work. While the cluster-robust variance estimator allows for within-worker correlation, tests based on it rely on a large-cluster (normal) approximation of the true distribution for the test statistic [Cameron, Gelbach, and Miller, 2012]. However, if the number of clusters is not large, then the large-cluster distribution can be a poor approximation for the finite-cluster distribution of the test statistic [Cameron, Gelbach, and Miller, 2012]. Accordingly, we might infer an effect of symmetry where, in fact, there is none.

There are a couple of reasons why this is not a major issue with our study. First, Monte Carlos in [Cameron, Gelbach, and Miller, 2012] show that cluster-robust standard errors do well with fewer clusters (workers) than we have. Specifically, they show that rejection rates become fairly accurate when the number of clusters (workers) increases from 25 to 30. Second, we evaluated the robustness of our results by following recommendations in [Roodman et al., 2019] and [Mackinnon and Webb, 2017] on how to conduct cluster-robust statistical tests when there are few clusters (and when cluster sizes vary wildly). In particular, we use bootstrap-t procedures to derive bootstrap p-values for our test statistics. The procedures yield an asymptotic refinement over the large-cluster approximation, in the sense that the procedures yield distributions which converge more quickly to the true distribution of the test-statistic. In Table A4, we show the alternative procedures leave the conclusions from Table 4 unchanged relative to conventional significance levels.[44]

[44] [Ibragimov and Muller, 2010] propose an alternative strategy when the number of clusters is not large. In our context, their strategy boils down to estimating a separate regression for each worker and then averaging

Representativeness of the Sample. In this subsection, we give reasons for why the sample is at least representative of the specific setting we study. We also discuss whether the sample allows us to draw conclusions about a broader population of interest.

Having an exogenous matching procedure makes it more likely that the sample from the observational data alone is representative of the specific setting we study. As we explain in the paper, the firm uses well-defined procedures to match workers with customers. The procedure makes matches exogenous. Customers are exogenously assigned a bundle of worker traits and workers are exogenously assigned a set of customers. In turn, our conclusions should apply to customers who were assigned other workers, and to workers who were assigned a different set of customers.

Our field experiment further increases the chances that our findings apply to other samples drawn from our setting. The field experiment holds worker characteristics fixed, and exogenously rewards workers with a bonus for serving more customers. It allows us to consider the effects of symmetry while including worker fixed effects in our specifications. Worker fixed effects help with the criticism of that our effects are driven by unobserved correlates of symmetry, such as the persuasive capacity of the worker.

The applicability of study is even broader than this. The findings should apply to other jobs where workers also have personal contact with customers, and also face trade offs between the number of customers they serve and the service quality they deliver to each. This includes jobs where the trade off is of primary importance, such as sales jobs in Best Buy or in residential real estate. Albeit to a lesser extent, this also includes more complex jobs, such as the family doctor, where our trade off is but one of several trade offs workers make. Furthermore, within these jobs, our analysis applies to other traits workers might possess. If other traits also substitute for service personal attention in consumer demand then, all else equal, workers who possess the traits should serve more customers than their coworkers.[45]

the results. We can not use their strategy with the observational data because facial symmetry is part of the intercept in the time series for each worker. We do not use the strategy with the experimental data because the strategy requires a more restrictive specification. Specifically, the strategy does not allow us to include time dummies in the specification for each worker.

[45]The 'all else equal' part of this statement is important. Unlike facial symmetry, other traits will usually have a direct effect on physical productivity. For example, the gender of workers in our context probably has a direct effect on the number of customers they serve, in addition to a direct effect on customer perceptions of the service quality they deliver. Thus, while the mechanism we describe can exist with other traits, it is important to separate their effect on physical productivity from their effect on customer perceptions of service quality.

Appendix References

[A1] Nicholas Bloom, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. Does management matter? evidence from india. *The Quarterly Journal of Economics*, 128: 1–51, February 2013.

[A2] A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427, 2012.

[A3] Karl Grammer and Randy Thornhill. Human (homo sapiens) facial attractiveness and sexual selection: The role of symmetry and averageness. *Journal of Comparative Psychology*, 108 (3):235–242, 1994.

[A4] Rustan Ibragimov and Ulrich K. Muller. t-statistic based correlation and heterogeneity robust inference. *Journal of Business and Economic Statistics*, 28:453–468, 2010.

[A5] James G. Mackinnon and Matthew D. Webb. Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, 32(2):233–254, 2017.

[A6] David McKenzie. Beyond baseline and follow-up: The case for more t in experiments. *Journal of Development Economics*, 99:210–221, 2012.

[A7] David Roodman, James G. Mackinnon, Morten Ørregaard Nielsen, and Matthew D. Webb. Fast and wild: Bootstrap inference in stata using boottest. *Queen's Economics Department Working Paper*, (1406), January 2019.