Identifying Neighborhood Effects among Firms: Evidence from Location Lotteries of the Tokyo Tsukiji Fish Market*

Kentaro Nakajima† Hitotsubashi University

Kensuke Teshima‡ Doshisha University

Preliminary (This Version: May 2026)

Abstract: A growing literature identifies consumption externalities as a crucial mechanism of retail agglomeration, but their underlying structure remains poorly understood. In particular, little is known about how the composition of neighboring firms shapes these externalities, or about which firms benefit from them. Identifying these mechanisms is difficult because firms endogenously choose where to locate and whom to locate near. We overcome this challenge by studying intermediate wholesalers in the Tokyo Tsukiji Fish Market, where shop locations were periodically reassigned by relocation lotteries, generating quasi-random variation in the diversity and composition of neighboring firms. First, we confirm that these intermediate wholesalers' shop locations are indeed randomly distributed. Then, we find that the diversity of neighboring firms positively affects the performance of firms that are specialized in narrow product categories and whose main clients are sushi restaurants—buyers known to visit multiple wholesalers in a single trip. Furthermore, we find no effect from back neighbors—shops equally close but not sharing the flow of buyers through horizontal corridors. This provides direct evidence that the neighborhood effects operate through shopping behavior rather than knowledge spillovers or other proximity-based interactions. Overall, our analysis shows that the complementarity of products between specialized diverse stores is an important factor behind urban agglomeration.

Keywords: Neighborhood effects, Consumption externality, Lottery, Micro-foundation of agglomeration.

*We are indebted to Kyoko Fukuchi of Tsukiji Ginrinkai for her crucial support in the early stage of the project, without which this project would not have been possible. We thank many intermediate wholesalers in the Tsukiji market who tirelessly answered our questions. We thank Takatoshi Tabuchi and David Weinstein for their insightful discussions and encouragement. We are grateful to David Atkin, Andy Bernard, Don Davis, Jonathan Dingel, Michal Fabinger, Masahisa Fujita, Koichiro Ito, Amit Khandelwal, Daiji Kawaguchi, Ayako Kondo, Hideo Konishi, Hisaki Kono, Noriaki Matsushima, Kiminori Matsuyama, Yuhei Miyauchi, Takeshi Murooka, Ryo Nakajima, Stuart Rosenthal, Bryce Steinberg, Adam Storeygard, Jacques Thisse, Yosuke Yasuda and seminar/conference participants at various places for helpful discussions. We thank Francisco Carrera and Paola Ugalde for excellent research assistance. Both authors are affiliated with TDB-CAREE, Hitotsubashi University, which provided access to the Teikoku Databank (TDB) dataset. We thank Takeo Goto and Yoshiki Hiramine of TDB for their invaluable assistance with the data. We gratefully acknowledge financial support from the JSPS (Nos. 25220502, 26245037, 26285051, 15H03344, 22H00060, 26H01957).

†Institute of Innovation Research, Hitotsubashi University. 2-1 Naka, Kunitachi, Tokyo, 186-8603 Japan (E-mail: nakajima@iir.hit-u.ac.jp)

‡Faculty of Economics, Doshisha University, Imadegawa-Dori Karasuma-Higashiiru Kamigyo-ku, Kyoto 602-8580, Japan and Hitotsubashi University TDB-CAREE, 2-1 Naka, Kunitachi, Tokyo, 186-8603 Japan. (E-mail: kteshima@mail.doshisha.ac.jp)

1 Introduction

Access to a variety of goods and services is at the heart of urban attractiveness. Cities offer dense and diverse consumption opportunities, such as restaurants, retail stores, and other local services, and these consumption amenities are an important force behind urban agglomeration (Glaeser et al., 2001; Couture, 2015; Handbury and Weinstein, 2015; Leonardi and Moretti, 2023). Recent work further shows that consumer mobility and multi-purpose trips generate consumption externalities across nearby businesses, providing a microfoundation for the spatial concentration of retail and service activities (Koster et al., 2019; Miyauchi et al., 2025; Oh and Seo, 2023).

Yet, despite the growing evidence that consumption externalities are an important force behind the spatial concentration of retail and service activity, the mechanisms through which they operate remain poorly understood. Existing evidence tells us relatively little about how the composition of neighboring businesses shapes these externalities, including what forms of complementarity among neighboring products matter for firm performance. It is also unclear which firms benefit from such externalities. These questions are central because consumption externalities need not operate uniformly across all neighboring firms: they may depend on whether nearby products are complementary for consumers and whether the focal firm is positioned to benefit from consumer search, comparison shopping, or incidental purchases. Specialized sellers, whose niche products may be purchased as part of broader shopping trips, are a natural case in which such externalities may be especially important.

A simple descriptive pattern suggests that store specialization and neighborhood diversity are closely linked. Using store-level data from the Census of Commerce in Tokyo, Figure 1 plots each store's product variety—measured by the number of four-digit JSIC product categories the store sells—against characteristics of its surrounding retail environment at the zip-code level. Panel A uses the number of other retail stores in the same zip code as the horizontal axis and shows that stores in denser areas tend to sell fewer product categories. Panel B replaces the horizontal axis with the total number of distinct product categories represented across all other stores in the same zip code, and shows a similar negative relationship: stores surrounded by a wider variety of products tend to be more specialized themselves. One might expect diverse neighborhoods to be populated by large multi-product stores, yet Figure 1 shows the opposite: stores in more diverse areas tend to carry fewer product categories. This suggests that neighborhood diversity is driven at least in part by clusters of specialized sellers offering complementary products—a pattern consistent with the consumption-externality mechanism we study in this paper.

This pattern is consistent with specialized sellers benefiting from nearby variety, but it is only sugges-

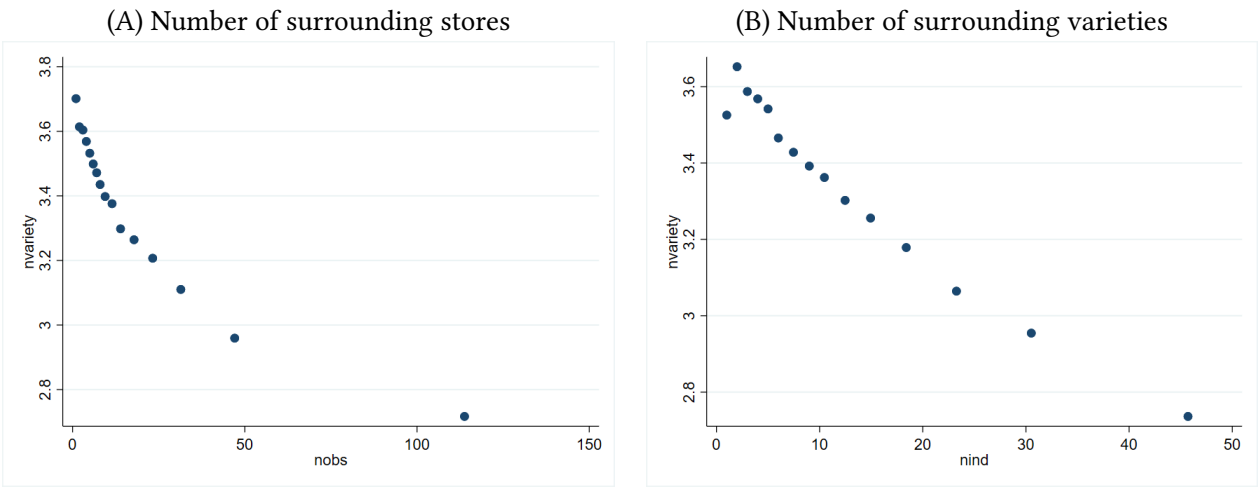


Figure 1: Specialization and the surrounding retail environment in Tokyo

Note: This figure uses store-level data for all retail establishments in Tokyo from the 2002 Census of Commerce. Areas are defined at the zip-code level. A store's product variety, shown on the vertical axis, is measured by the number of products the store sells, where products are defined at the four-digit JSIC level. Panel A plots store-level product variety against the number of surrounding retail stores in the same zip-code area. Panel B plots store-level product variety against the number of surrounding products in the same zip-code area, measured by the number of unique four-digit JSIC categories sold by other stores in the area. Each panel reports a binscatter, with observations grouped into bins along the horizontal axis.

tive: specialization and neighborhood diversity may be jointly shaped by local demand, sorting, or other unobserved features of retail locations. Moreover, spatial proximity may generate other forms of neighborhood effects, such as knowledge spillovers or advantages from unobserved location quality. Establishing whether neighborhood diversity itself improves firm performance therefore requires exogenous variation in the composition of neighboring firms.

We overcome this challenge by analyzing the neighborhood effect among fish intermediate wholesalers located in the Tokyo Tsukiji Fish Market and by exploiting a unique feature of their locations within the market; their locations are determined every 4-10 years by relocation lotteries. In our sample period, the Tokyo Tsukiji Fish Market hosts approximately one thousand intermediate fish wholesalers specializing in different types of fish.[1] The feature that the location of firms is randomly determined by the lotteries implies that the characteristics of the neighborhood are generated by chance, meaning that, within a market, there are areas with more diversity and areas with less of it exogenously. We also compile the exact location of every intermediate wholesaler within the market and thus those of their neighbors.[2]

[1]Intermediate wholesalers are intermediary firms that purchase fish from wholesale traders, who receive fish from outside the market and sell to restaurants and supermarkets. See Section 2 for more detail. We use the term "intermediate wholesalers" according to Bester (2004), a comprehensive anthropological study of the Tsukiji market. Throughout the paper, we also use "firms" to indicate the intermediate wholesalers.

[2]In this paper, we focus on one lottery cycle in 1990 and analyze the firm performance until 1995, which is the year that the next lottery occurred. We plan to extend our analysis to more lottery cycles.

The Tokyo Tsukiji Fish Market is particularly suitable for our purpose. In addition to the very rare opportunity to exploit a random source of variation in firm locations, there are two noteworthy features. First, the intermediate fish wholesalers are all firms in a very narrowly defined industry; however, they exhibit substantial heterogeneity along important dimensions, particularly the type of fish they sell, their operation scales, and how they are specialized in that type. Thanks to the heterogeneity in the type of fish, the lotteries generate exogenous variations in diversity of neighboring firms as well as other neighborhood-level characteristics such as fraction of firms of a particular type. We then can analyze what type of firms benefit from such neighborhood characteristics.

Second, the Tsukiji market has several features that allow us to distinguish externality due to shopping from other forms of neighborhood or spillover effects. The fact that these firms are all located in the same marketplace means that there is no role for the consumption amenities and that there is no permanent vertical buyer-supplier relationship among intermediary wholesaler themselves.[3] Thus, these channels of neighborhood effects that could be potentially important elsewhere do not exist in this setting. Furthermore, the firms have two types of neighbors: those with whom they directly share buyer flow because they face each other across a corridor for buyers and those with whom they do not because they are back to back to each other without having a corridor for buyers in between. The effect of diversity due to shopping behavior should not exist from the second type of neighbors. Thus, finding neighborhood effects only from the first type of neighbors provides compelling evidence of complementarity due to shopping behavior. In short, the Tsukiji market provides an ideal setting for testing complementarity due to shopping externality.

We proceed with our analysis according to the following steps. First, we perform various types of analysis to show that the locations of the intermediate wholesalers are indeed randomly distributed by the location lotteries. Then, we find that the diversity of the neighboring firms positively affects the performance for specialized firms. We then explore the mechanisms behind these neighborhood effects. Furthermore, when we use information on wholesalers' main clients to identify firms that are more exposed to shopping-around buyers (Sushi restaurants), we find stronger effects of neighborhood diversity for these wholesalers. These results suggest that the neighborhood effects we estimate operate through buyers' shopping behavior: they are concentrated among firms for which product complementarity and shopping-around behavior of their customers should be most important.

The mechanism we empirically identify has been a central theoretical mechanism to explain retail clusters, marketplaces, and within-city agglomeration. A body of theory built on consumer search with

[3] A firm can buy fish from other firms if it encounters a shortage of fish on a particular day, or it can sell fish to other firms that run short of their fish, but these actions should cancel out over time.

heterogeneous preferences shows that clusters of specialized firms emerge because they reduce the cost of finding a product that matches each consumer's idiosyncratic demand (Wolinsky (1983); Konishi (2005)). The same logic implies that anchor stores or diverse assortments attract consumers to a location, generating demand spillovers to neighboring specialists who serve buyers whose specific needs are unmet by generalists (Konishi and Sandfort (2003); Rhodes and Zhou (2019)). In our setting, a diverse neighborhood of wholesale intermediaries collectively provides this one-stop convenience, and because the externality is not internalized within any single firm, it manifests as a measurable neighborhood effect. Despite the centrality of this mechanism in theory, few empirical studies have established its relevance owing to identification difficulties. We provide the first randomization-based evidence of this consumer-search-based shopping externality.[4]

Our paper also contributes to the recently growing literature on consumption externalities and the micro-geography of retail and service activity. Koster et al. (2019) show that pedestrian flows and multipurpose shopping trips generate shopping externalities across retail establishments on Dutch shopping streets. Miyauchi et al. (2025) use smartphone mobility data to show that consumer trips across urban destinations create demand linkages between locations within cities. Oh and Seo (2023) study the agglomeration of service establishments in Seoul and emphasize the role of consumer trips across nearby services. Closely related, Jardim (2015) estimates spillovers among neighborhood retailers and studies the agglomeration of the retail mix. Our analysis is in the same spirit, but focuses on the composition of neighboring sellers and on which firms benefit from the resulting consumption externalities.

Our paper is further related to work on the consumption value of urban variety and density. Couture (2015) studies the consumption benefits of urban density, while Schiff (2015) and Cosman (2014) analyze the value of diversity and density in restaurants and nightlife.[5] These papers quantify the benefits of dense and diverse consumption opportunities mainly from the consumer side, whereas our analysis studies how the same type of mechanism can affect the performance of the stores themselves. In addition, our paper contributes to the small and growing literature on neighborhood effects among firms in service sectors. Benmelech et al. (2018) and Shoag and Veuger (2018) find that retail firms' performance deteriorated when a nearby branch of a national retail chain disappeared because of the bankruptcy of the whole chain, and interpret this as evidence of neighborhood spillover effects.

This paper is organized as follows. Section 2 describes the background of the Tokyo Tsukiji Fish Market

[4] A related class of models derives agglomeration from consumers' love-of-variety preferences rather than search frictions. See Matsuyama (1992), Matsuyama (1995), Henkel et al. (2000), Arakawa (2006), and Tabuchi (2009), among others. In these models, the substitutability of heterogeneous products attracts spatially dispersed consumers to marketplaces and reduces their trip costs.

[5] Handbury and Weinstein (2015) show the gains from product variety at the city level.

and the details of the relocation lottery. Section 3 presents the data and empirical strategy. In Section 4, we provide evidence that the location assignment is indeed random. Section 5 reports the main results: Section 5.1 presents the baseline estimates and event study; Section 5.2 tests the shopping-externality mechanism using sushi-restaurant clients; and Section 5.3 presents robustness exercises, including an alternative diversity measure, full-sample results, and a placebo test using back neighbors. Section 6 concludes.

2 Background

2.1 Tsukiji Fish Market

The Tsukiji fish market, located in downtown Tokyo, was the largest wholesale fish and seafood market in the world.[6] In 1990, the period we focus on in this paper, the market was handling more than 700,000 metric tons per year (i.e., 2,000 tons daily) of more than 400 different types of seafood. There were more than 1,000 firms in the market, including wholesaler traders, intermediate wholesalers, and related equipment and machinery companies, which together employed more than 60,000 workers (Annual Report of Tokyo Central Wholesale Market).

The flow of fish from producers to consumers is summarized in Figure 2. Fish that arrived from all of the ports in Japan and from around world were first handled by seven large wholesale traders (i.e., "O-oroshi" in Japanese). Then, the wholesale traders sold the fish to intermediate wholesalers, either through auctions or through direct negotiations. Next, the intermediate wholesalers, at their shops within the marketplace, resold the fish to their buyers, which included supermarkets, local fish stores, or restaurants. Then, these buyers served their respective customers at their shops or restaurants outside the market, mostly in the Tokyo area. This paper focuses on the relationship between the intermediate wholesalers and buyers, as shown in the rectangular box in Figure 2.[7]

2.2 Types of Intermediate Wholesalers [Nakaoroshi]

Intermediate wholesalers can be classified according to several dimensions. First, they had different fish specialties. In the Tsukiji fish market, there were sixteen trade groups (i.e., "Gyokai") formed by firms of similar specializations. Intermediate wholesalers formed these trade groups to negotiate with other interested parties. Trade groups could be directly associated with specific fish, such as tuna, with fish caught from specific locations, or with the way the fish were handled. One exceptional trade group consisted

[6] The wholesale market functions of the Tsukiji Fish Market were relocated to the Toyosu Fish Market in October 2018.

[7] Some large supermarket chains could buy directly from wholesale traders if they obtained a license. They are shown as "Licensed Buyers" in the figure.

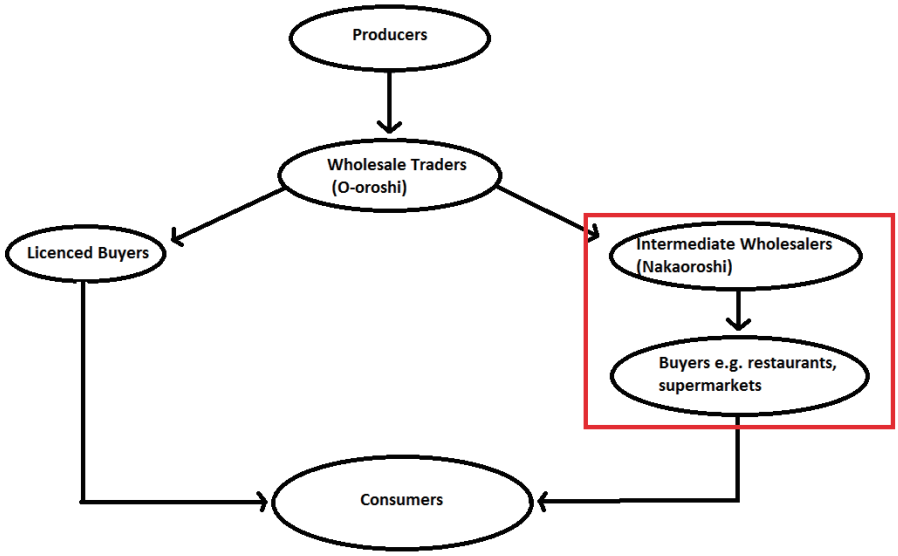


Figure 2: Flow of Fish

of intermediate wholesalers who specialized in high-quality fish for sushi and tempura served at sushi restaurants and high-end Japanese cuisine restaurants. As we mentioned above, firms in Tsukiji belonged to trade groups depending on the type of fish they handled. It was generally thought to be difficult for a business to change its fish category because of the extensive requirement for fish-specific skill, knowledge, and equipment.[8] Firms, however, could belong to more than one trade association. In this case, firms typically tended to belong to similar types of trade groups, such as shrimp and Japanese lobster. Table 1 shows the list of trade groups, the type of fish each trade group handled, and the number of firms and shops in 1990.[9]

Firms dealing with tuna had the largest share of firms and shops in Tsukiji. In second place were firms selling fish for sushi and tempura. The third and fourth largest shares were taken by firms that mainly sold their fish to supermarkets and small and mid-sized fish shops.[10] These four categories were the main

[8]Fish arriving from different areas were sold from the wholesale traders to intermediate wholesalers in different areas of the market in distinct ways. For example, different areas could have distinct ways of implementing auctions. This made it difficult for intermediate wholesalers, particularly small ones, to deal with a wide variety of fish, as they would have to send employees to distinct transaction places and learn the respective rules for these transactions. This is why some trade groups were based on where the fish were from.

[9]A firm could have multiple shop spaces. An explanation of the shop spaces is given in the next section.

[10]Enkai and Kinkai mean "far coast" and "near coast," respectively. The former is outside the greater Tokyo bay area, while

Table 1: Overview of trade groups in 1990

Name	Product	Variable name	Number of firms	Number of shops
Omono	Tuna	Tuna	334	545
Tokushumono	Fish for sushi and tempura	Sushi	236	531
Enkai	Popular fish outside Tokyo	Enkai	144	224
Kinkai	Popular fish within Tokyo	Kinkai	109	174
Hokuyo	Salmon	Salmon	75	170
Ebi	Shrimp	Shrimp	68	134
Aimono	Half-processed fish	Aimono	63	108
Renseihin	Fish cake	Renseihin	61	113
Enkanmono	Dried fish	Enkan	55	102
Enyomono	Fish in Far Ocean	Enyo	27	94
Tako	Octopus	Octopus	22	29
Tansuigyo	Freshwater fish	Tansui	16	30
Togei	Whale	Whale	16	74
Tsukudawakai	Tsukudani	Tsukudawa	8	18
Iseebi	Japanese lobster	Iseebi	8	23
Tookakai	Shark	Shark	4	4

trade groups in the Tsukiji Fish Market. At the same time, we can observe a certain number of firms that handled extremely specialized seafood, such as shrimp, Japanese lobster, octopus, and shark.

Second, they were heterogeneous in the number of shops they had in the marketplace. Third, they differed in their degree of specialization and in the diversification of the types of fish they handled. Fourth, intermediate wholesalers sold either to large-scale supermarkets or to smaller restaurants. These factors were correlated with one another. Generally, intermediate wholesalers belonging to trade groups that were characterized by the location of fish were more likely to sell to large-scale customers and to handle many varieties of fish. On the other hand, intermediate wholesalers that belonged either to the trade group dealing with fish for sushi or to trade groups specializing in specific fish were more likely to sell to smaller-scale customers and to specialize in one or a few varieties of fish. Finally, small-scale intermediate wholesalers, typically those with only one shop, were likely to be more narrowly specialized or to operate in a niche segment, consistent with the argument by Holmes and Stevens (2014). The demand for specialized types of fish typically came from small-scale restaurants, while demand from large-scale supermarkets was less specialized.

2.3 Building for Intermediate Wholesalers

Figure 3 is a panoramic picture of the Tsukiji Fish Market. In the center of the picture, there is a quarter-

the latter is in the greater Tokyo bay area. These two trade groups merged into one group in the 2000s. The name for the new trade group is "Sengyo", meaning fresh fish. This generic name suggests that they are in the non-specialized segment.



Figure 3: Picture of Tsukiji Fish Market

circle-shaped building. The building housed intermediate wholesalers. Figure 4 shows the layout of the buildings in the market, the flow of fish, and the flow of buyers.[11] As we discussed, fish in Tsukiji were

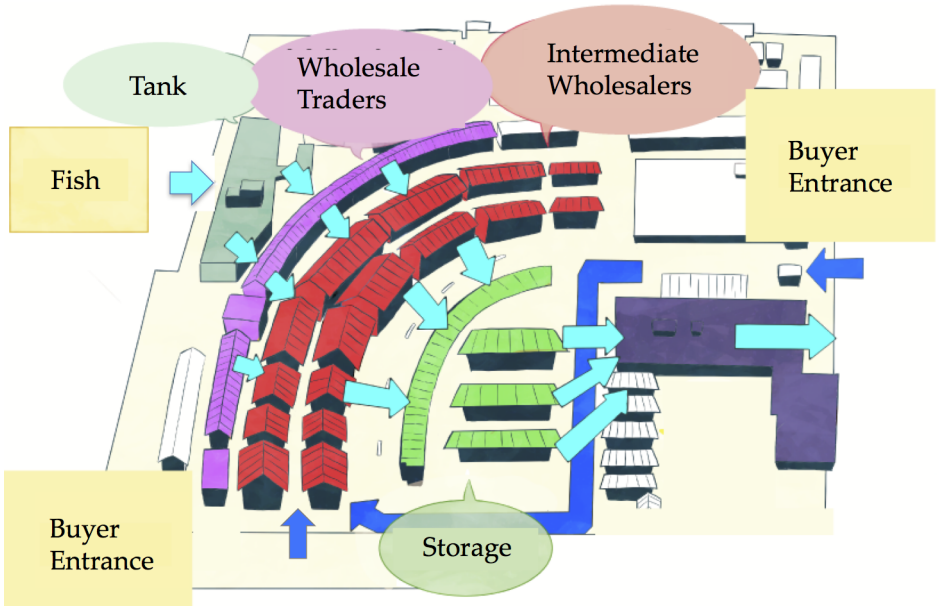


Figure 4: Buildings of Tsukiji Fish Market, Flow of Fish, and Buyers

first handled by seven large wholesale traders, who sold the fish to intermediate wholesalers at the light-colored (purple) curved buildings on the left side of the figure. The transactions between the wholesale traders and the intermediate wholesalers occurred at dawn. Then, intermediate wholesalers sold their fish to customers who came to the market early in the morning until around 9 a.m. The fan-shaped areas in

[11] There used to be railway tracks just outside the area for the wholesale traders. The market was built in this shape in order to handle efficiently fish carried by trains.

dark color (red) are the building housing intermediate wholesalers.[12] The fish flows are shown in the figure by light-colored (sky-blue) arrows. Customers of intermediate wholesalers came from two entrances of the market and entered the building of intermediate wholesalers from the bottom part of the figure. The buyer flows are shown in the figure by dark-colored (blue) arrows.

The detailed structure of the building housing intermediate wholesalers is shown in Figure 5. The

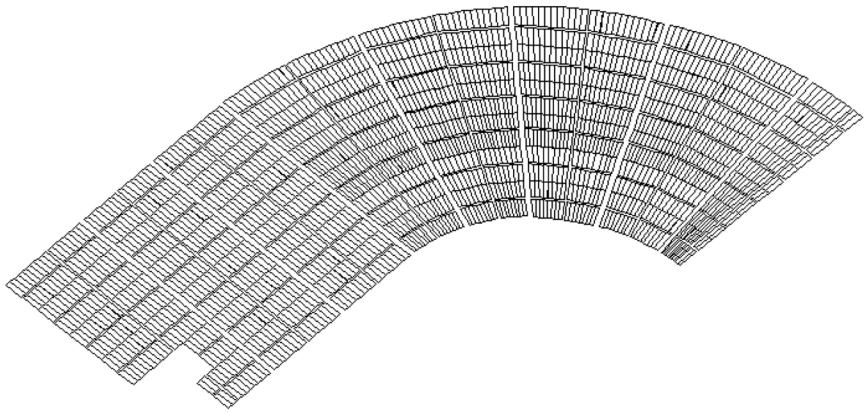


Figure 5: Building for Intermediate Wholesalers

width of the building was around 70 meters, and the length (walking through the center of the building from the left to the right ends) was around 300 meters. As the figure shows, the building consisted of 12 rows and 15 column blocks of stalls (i.e., shop spaces),[13] where each column block had eight to twelve stalls per row. Each small rectangle in the figure is one stall. There were approximately 1,700 stalls overall. The rows and column blocks were divided by vertical streets and horizontal corridors.[14] All of the horizontal corridors between rows had the same width, but the width of the streets between columns differed. There were seven large and eight small vertical streets.[15] Figure 6 shows the location of corridors and streets in the market. Figure 6a shows the horizontal corridors. Curves in black are the horizontal corridors. As the figure shows, there were six horizontal corridors. Figure 6b shows the location of small vertical streets in black. Finally, Figure 6c shows the location of large vertical streets in a similar way.

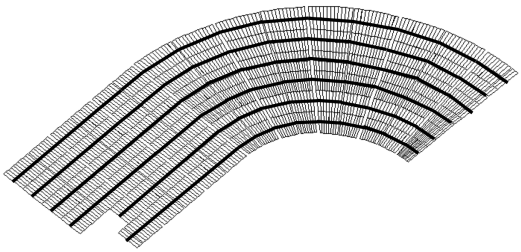
The characteristics of the stalls differed within the building. As we can see from Figure 5, there were stalls in the rectangular areas on the left and fan-shaped areas on the right. Stalls in the rectangular area were of the same size, but in the circular area, the stalls became larger from the inner to the outer parts (i.e.,

[12] In Figure 4, the buildings for the wholesale traders and intermediate wholesalers look separated, but they were in fact under one roof.

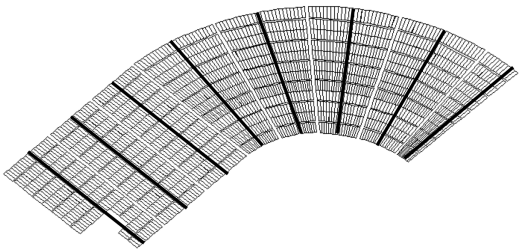
[13] We use stalls and shops interchangeably.

[14] We use "street" to indicate a vertical street, while we use "corridor" to indicate a horizontal corridor.

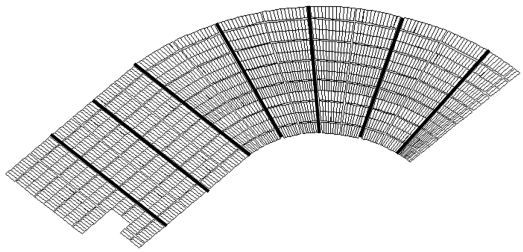
[15] They were literally called "large streets" and "small streets" in the market.



(a) Horizontal Corridors



(b) Vertical Streets (Small)



(c) Vertical Streets (Large)

Figure 6: Corridors and Streets in Tsukiji Fish Market

from the bottom part of the figure to the top part). On average, the area of one stall was 7.2 m^2. The main entrance of the building for buyers was located at the bottom left of the figure. Buyers walked through the building to shop for fish. Thus, the rectangular area and the stalls in the bottom part of the figure were attractive because they provided easy access for the customers. The stalls located on the corners of one of the large streets were also attractive to customers, since the other corridors and streets were quite narrow. As mentioned above, the attractiveness of stalls differed and greatly affected the performance of the firms. The intermediate wholesalers agreed that the stalls on the corner were good locations for them. Ensuring fairness for the firms in terms of their locations was one important reason why they continued to use the location lottery system for the long term.

Since there were approximately 1,700 stalls (i.e., shop spaces), which exceeded the number of firms, some firms occupied more than one shop space. Firms could extend their shops by buying the right of stall-holding from other firms. As we mentioned, the area of one stall was quite small, especially for high-performing firms, which had many customers and dealt with a high volume of fish. In general, high-performing firms tended to extend their shops to deal with a large amount of fish. When they occupied more than one shop space, they almost always occupied spaces that were horizontally adjacent to each

other because of the convenience for the business.[16] [17]

2.4 Buyer Trip Behavior

In this section, we explain background information and anecdotes about buyer trip behavior and product complementarity. There were two points worth mentioning that shaped buyers' trip behavior. First, there was uncertainty in the availability of fish in each store because the types of fish arriving at the market varied daily. Second, although the entire Tsukiji Fish Market was not very large, the corridors were narrow. Furthermore, the market was open only in the early morning (mainly 6 a.m.–9 a.m.), so navigating the narrow corridors in heavy congestion through the entire market was extremely costly for buyers with limited time.

These two factors led to the following buyers' trip behavior. First, buyers typically had one to five favorite sellers that they visited daily.[18] Buyers had these favorite sellers because these sellers treated them well to compensate for the uncertainty of fish availability. Second, buyers were simultaneously willing to try new sellers and eventually replace their existing favorite sellers with new ones. Third, because of all these factors, typical buyer trip behavior was to visit a few main sellers and then briefly look around the surrounding stores of the main sellers. This institutional setting suggests that buyers were likely to first visit their main sellers and then browse other shops along the same horizontal corridor, where facing shops were visible and easy to observe, while avoiding vertical streets where visibility was limited.

Under such trip behavior, we were able to collect ample anecdotes of complementarity from the diversity of neighboring stores. It was recognized as a realistic channel of neighborhood effects among intermediate wholesalers.

The term "Tsuidegai" was used in the Tsukiji fish market. This term refers to a mostly unplanned purchase of fish at one shop after a planned purchase of another type of fish nearby.[19] We had ample stories demonstrating "Tsuidegai" from our interviews with managers of intermediate wholesalers. One manager told us that shops selling fish for sushi could benefit from the presence of a high-performing shop specialized in tuna because a sushi restaurant chef typically bought tuna first and then looked for other

[16] Since the average size of a stall was only 7.2 m^2, the benefit of using connecting stalls without a wall was substantial.

[17] One may be concerned that it was difficult to buy the right of stall-holding from horizontally located firms. However, in the location lottery, a firm with multiple shops applied for one lottery for all of its shops and could thus have its shops next to one another. Therefore, a firm that planned to extend its shops did not need to buy the right of stall-holding from horizontally located firms if it bought the right of stall-holding just before the lottery. In fact, most of the trading of rights for stall-holding took place just before the location lottery.

[18] In fact, many sellers kept fish for their frequent buyers under their counters, hidden from wandering strangers, which made it even more difficult for new buyers to search.

[19] "Tsuide" literally means "on the occasion of" or "in addition to" in Japanese, and "gai" means "purchase."

fish.[20] Another manager who specialized in high-quality whitefish for sushi restaurants told us that his shop would benefit from the presence of other shops that specialized in other types of high-quality fish for sushi. On the other hand, a manager who operated a large-scale shop dealing with popular fish shared that "Tsuidegai" was not important for his business, as his customers were large-scale supermarkets. These episodes illustrate the consumer-search mechanism described in the theoretical literature: buyers visit a diverse area to reduce search costs, and specialized firms capture incidental demand from buyers whose primary purpose was to visit a different seller.

2.5 Relocation Lotteries

The market and the land were owned by the Tokyo Metropolitan government. The relocation lottery was introduced to mitigate inequality and unfairness associated with the location advantages described in the previous section. These relocations (i.e., Tempo Ido) have been major events, as they involve the whole re-installation of equipment for every shop. The relocations and associated lotteries took place every four years until 1990. After 1990, the timing became irregular because of debates on the relocation of the whole market. The lotteries were performed in 1990, 1995, and 2004.

Each lottery consisted of two parts: the preliminary and main lotteries. In the preliminary lottery, the market was divided into four blocks. Then, firms submitted their preferences to one another. If there were more firms than locations in these blocks, the lottery was implemented to determine who could enter the block. If a firm did not win a lottery for a block, it had to reapply for a different block where there was still a vacancy. Figure 7 shows the blocks for the preliminary lottery. Then, after the blocks the firms went to were determined, the main lottery was held for each block, and the exact location within the block was determined.

In the main lottery, a firm owner drew a ball printing a number by a lottery machine. The location was determined by the number printed on the ball in the manner shown in Figure 8. Namely, the location in each block was assigned to the first row from left to right and then to the second row from right to left, etc. For example, for the case of Block 1, colored white in Figure 8, the assignment of the shop locations started from a stall located at the left-end of the block highlighted by a circle. A firm that drew the number one in the lottery was assigned to the stall. If the firm had multiple rights of, for example, four shops, four contiguous stalls from the starting stall to the direction of the arrow were assigned to the firm. Then, the firm that drew the number two assigned their stalls next to those for the firm that drew the number one. Once this process reached the right-end, it turned to the left and was repeated until the all the shops in

[20] This anecdote may better correspond to what marketing studies call "anchor stores."

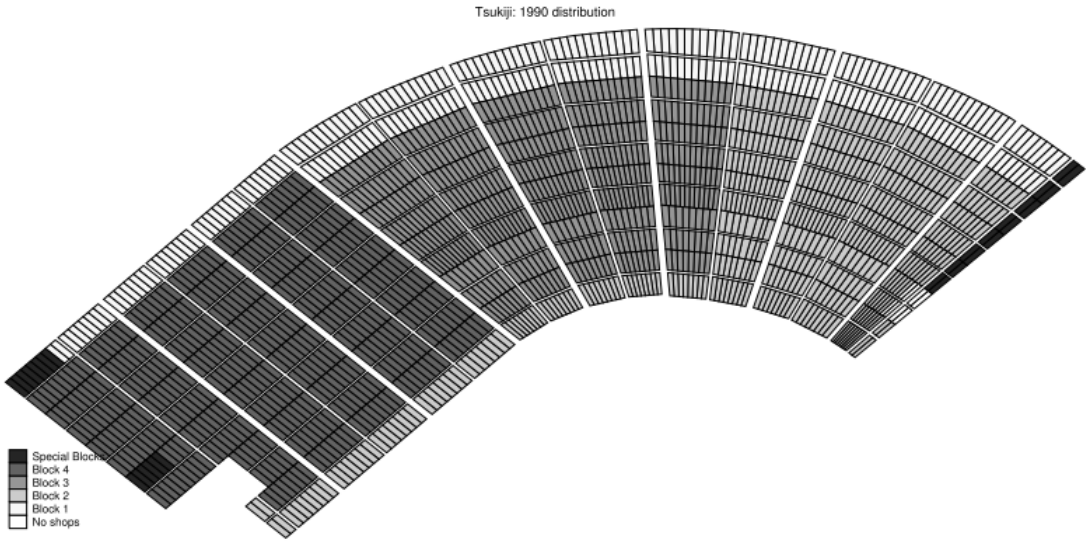


Figure 7: Blocks for the preliminary lottery

the block were assigned, as Figure 8 shows. The other blocks were assigned by the same procedure. The starting stalls are also highlighted by circles. This lottery system allowed a firm with multiple shops (i.e.,

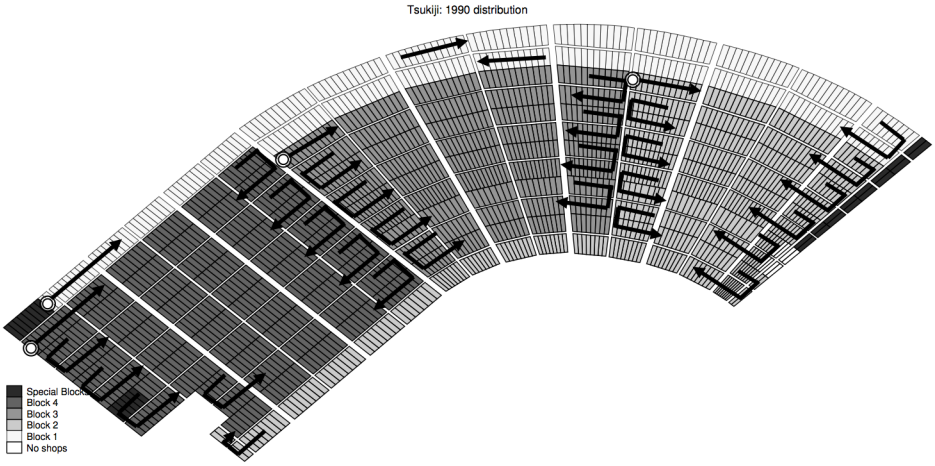


Figure 8: Assignment of shop location by main lottery

multi-shop firms) to place their shops next to one another. One feature of the lottery worth emphasizing is that the lottery number would not determine the exact location ex ante because of the assignment rule and because of the existence of multi-shop firms. The exact location of a firm depended on its lottery number and the number of shops that would have been occupied by the other firms that drew lottery numbers smaller than theirs. This feature was useful in preventing fraud, since there was no point in aiming for a particular lottery number. In addition, exchanges of shops after the lottery is not allowed.[21]

[21]For the 1995 lottery, we were able to obtain hand-written records of the lottery published internally a day after the lottery

There was another important feature of the lottery. The unit of application for the lottery was two shops. Single-shop firms as well as firms with odd-numbered store spaces had to form a pair. Furthermore, group application for more than one unit was allowed. According to Bester (2004), the average size of a group was 3.6 shops in the 1990 lottery.

This system of group applications could potentially pose a threat to our identification strategy, as there is a degree to which horizontal neighbors are self-selected. Unfortunately, we do not have information on joint applications. However, we find that when one space of a unit (two consecutive store spaces) is occupied by a firm with an odd number of store spaces, the other space is also occupied by another firm with an odd number of store spaces for more than 95 % of the cases where the second firm is different from the first one. This suggests that large-scale groups that involve more than two firms are rare.[22] Our analysis for the validity of randomization, presented in Section 4, suggests that group applications do not lead to a systematic violation of randomization. A suspicious reader might still wonder whether this could be because our data do not capture all of the relevant characteristics. However, whether neighbors faced each other across horizontal corridors was completely out of the control of the firms because the lottery numbers of the two shops facing each other were typically far apart, which is a point also mentioned by Bester (2004). We will explore this feature in one of the robustness check analysis presented in Section 5.3.

Finally, economists may wonder why the Tsukiji market has been implementing lotteries as their way of assigning store spaces to intermediate wholesalers. Different arrangements have indeed been discussed. It seems that assigning "good locations" to firms that have a higher willingness to pay is incompatible with the notion of fairness of firms in the market. Another interesting anecdote is that creating zones according to fish specialty was discussed but not adopted precisely because it would have forced buyers to walk around wider areas of the market if they want to buy several types of fish. This anecdote suggests that firms in this market recognize the value of diversity of shops in all parts of the market.[23]

3 Data and Empirical Strategy

3.1 Data

Our regression analysis combines two data sources. The first provides the location of each intermediate wholesaler within the Tsukiji market and the trade group affiliations of each firm. The second is a firm-

implementation and to compare the records to the actual location of shops. We did not find any evidence of shop space swapping.

[22] Anecdotes tell of enormous psychological pressure on a person who draws a lottery on behalf of other firms.

[23] Indeed, it is not theoretically straightforward to achieve optimal allocation when diversity is important or when externality is present. Simulating alternative mechanisms based on our estimates would be an interesting extension.

level panel data that provides our main outcome variables (sales and employment) and information on customer composition.

3.1.1 Tsukiji Directories and Industry Newspaper

To identify the exact location of every intermediate wholesaler within the market and to recover the institutional details of each location lottery, we use industry newspapers (*Nikkan Shokuryo Shimbun*) and the directories (*Kumiaiin Meibo*) compiled annually by the union (To-oroshi) and the newspaper company (*Nikkan Shokuryo Shimbun*). For the 1990 and 1995 lotteries, these sources provide each firm's stall-level location and its trade group affiliations, which allows us to identify the fish specialty of each firm. The industry newspaper additionally reports the lottery block of each firm in the 1990 preliminary lottery.

3.1.2 Teikoku Databank (TDB) Data

Our main outcome variables (sales and employment) and information on customer composition come from the Teikoku Databank (TDB) firm-level data set. We obtained the data from the Teikoku Databank Center for Advanced Empirical Research on Enterprise and Economy (TDB-CAREE) at Hitotsubashi University. Teikoku Databank is a major Japanese credit research company, and its data set covers most Japanese firms. The records contain basic firm-level information such as industry, the location of headquarters, the number of workers, and revenue, as well as up to five reported customers per firm. The data are organized as an annual panel covering the years from 1985 to 2012. For the analysis below, we use the period 1985–1995, which spans five years before the 1990 location lottery and the five years until the next lottery in 1995.

Two features of the TDB data are particularly useful for our purpose. First, the panel data begins in 1985, five years before the 1990 lottery, which allows us to estimate event-study specifications with a meaningful pre-period and to test the parallel-trends assumption directly. Second, the customer field, which reports up to five customers per firm, allows us to characterize each wholesaler's main client types and, in particular, to identify wholesalers whose primary customers are sushi restaurants. These restaurants are especially relevant for our analysis because they tend to shop across wholesalers and are therefore likely to generate the consumption externalities discussed in Section 5.2.

We link the TDB records to our directory-based data set of Tsukiji intermediate wholesalers using firm names, addresses, and telephone numbers.[24]

[24] Section 5 reports the resulting number of matched firms. The number of observations in our regressions in Section 5 is slightly smaller than the total number of matched firms because of missing values of the outcome variables in some years.

3.2 Empirical Strategy

Our empirical strategy is designed to test whether shopping externalities exist among firms in this market. The consumer-search theory of retail agglomeration provides the theoretical foundation for the form these externalities should take. Clusters of specialized firms reduce the cost of finding a product that matches each consumer's idiosyncratic demand (Wolinsky, 1983; Konishi, 2005), and diverse assortments—whether provided by anchor stores or by a concentration of heterogeneous firms—attract consumers to a location, generating demand spillovers to neighboring specialists who serve buyers whose specific needs are unmet by generalists (Konishi and Sandfort, 2003; Rhodes and Zhou, 2019). The institutional features described in Section 2 suggest that this mechanism operates through buyers' shopping-around behavior in the Tsukiji market. The theory thus predicts that if shopping externalities are present, neighborhood diversity should disproportionately benefit specialized firms.

We test this prediction by comparing changes in performance across specialized and non-specialized wholesalers randomly assigned to neighborhoods with different levels of diversity by the 1990 lottery, controlling for lottery block, trade group, own diversity, shop size, and physical location. In Section 5.2, we further examine this interpretation using information on customer composition.

To implement this test, we consider the neighborhood region as the unit of the areas in which neighborhood effects arise. We define a neighborhood region as a set of stalls consisting of two rows and column blocks of stalls between large vertical streets and stalls sharing one horizontal corridor. An example of a neighborhood region is shown in Figure 9a. The black-colored area is a neighborhood region. We assume that neighborhood effects on firm i's operate within the neighborhood region in which the firm is located.

Our hypothesis also predicts that the neighborhood diversity effect should be concentrated among firms whose product offerings make them more exposed to shopping-around behavior. We examine this implication using a firm-level specialization measure, Specialization_i, which captures the extent to which a wholesaler's business is concentrated in a narrow set of products. We define this measure in detail in the subsection below.

We test this prediction by estimating the following specification:

$$\Delta y_{igr} = \beta_0 + \beta_1 \text{Diversity}_r + \beta_2 \text{Specialization}_i + \beta_3 \big(\text{Diversity}_r \times \text{Specialization}_i\big) + \gamma X_{igr} + \lambda_g + \varepsilon_{igr},$$

where Δy_{igr} is the change in firm i's log sales or log employment between 1989 (one year before the 1990 lottery) and one of 1991, 1993, or 1995; Diversity_r is the number of unique trade groups represented by the other firms in the neighborhood region r; Specialization_i is a dummy indicating whether firm i is

specialized; X_{igr} is a vector of firm and location controls; and λ_g denotes trade group fixed effects.

Diversity$_r$ is measured using the shop locations determined by the 1990 lottery and is held fixed throughout the analysis. Because we do not observe the exact shop locations before 1990, we cannot construct pre-lottery neighborhood variables. However, the random assignment of locations established in Section 4 implies that Diversity$_r$ is uncorrelated with pre-determined firm characteristics conditional on lottery block and trade group fixed effects, so using the post-lottery neighborhood composition as a fixed regressor is appropriate. We verify this directly in the event-study specification below, where the pre-lottery coefficients are indistinguishable from zero.

Taking the first difference of the outcome between 1989 and a post-lottery year removes time-invariant firm-level heterogeneity. The coefficient of interest is β_3, which captures how the effect of randomly assigned neighborhood diversity differs between specialized and non-specialized firms. Because Diversity$_r$ is randomly assigned conditional on the lottery block, β_3 has a causal interpretation. The hypothesis predicts $\beta_3 > 0$. We focus on post-lottery outcomes through 1995 because 1995 is the year of the next location lottery, after which the neighborhood composition generated by the 1990 lottery is replaced by a new random reassignment.

The control vector X_{igr} includes a dummy indicating whether firm i occupies a corner stall on a large vertical street, a dummy for owning multiple shops, the number of trade group affiliations of firm i (own diversity), the 1990 lottery block fixed effects, and two coordinate variables capturing the horizontal and vertical position of the firm within the market. In all specifications, standard errors are clustered at the neighborhood region level. Below, we explain our measures of neighborhood variables and firm and location characteristics in greater detail.

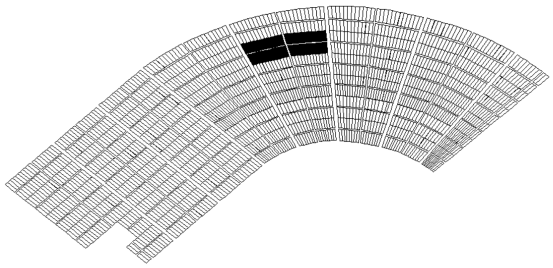
We use the total number of unique trade groups to which neighborhood firms belong as the neighborhood diversity variable.[25] As discussed in the introduction to this subsection, we consider the interaction of Diversity$_r$ with firm-level variables that, according to our hypothesis, should determine the strength of the neighborhood diversity effect. The main variable we consider captures supply-side specialization. We classify firm i as specialized, Specialization$_i = 1$, if it satisfies two conditions as of 1990: (i) it belongs to exactly one trade group, and (ii) that trade group is not Kinkai or Enkai.[26] The variable is constructed once, as of 1990, and is held fixed throughout the analysis.

The variation in neighborhood diversity generated by the lottery is illustrated in Figure 9. To build

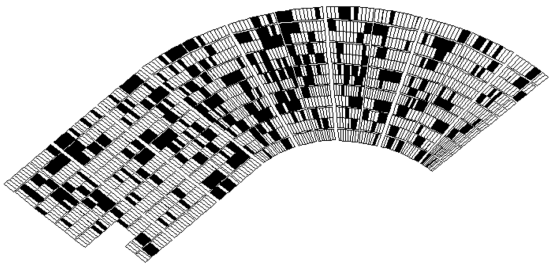
[25] In Section 5.3, we use an alternative measure of diversity.

[26] The Kinkai and Enkai trade groups deal with a wide variety of popular fish caught in nearby and distant coastal areas of Japan, respectively, and are widely understood within the market as non-specialized segments. As described in Section 2.2, intermediate wholesalers in these two groups typically supply supermarkets and other large-scale buyers.

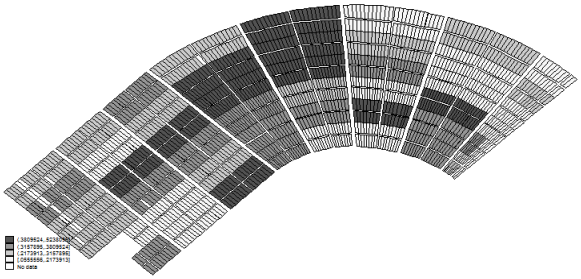
intuition, consider one trade group as an example. Figure 9b shows the location of tuna shops in black after the 1990 lottery. Because locations are randomly assigned, some neighborhood regions end up with a high concentration of tuna shops while others have few or none, as shown in Figure 9c. The same type of random variation arises for every other trade group. A neighborhood region where many different trade groups happen to be represented scores high on our diversity measure; one dominated by a single group scores low. Figure 9d shows the resulting diversity across neighborhood regions in 1990, with darker colors representing greater diversity.



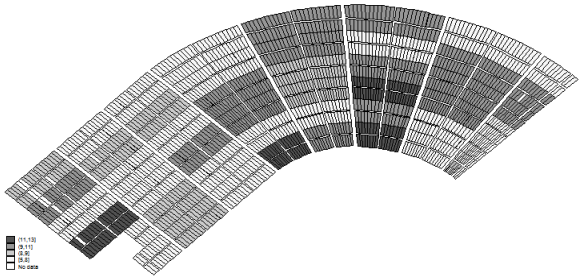
(a) Example of a neighborhood region



(b) Locations of Tuna shops in 1990



(c) Fraction of Tuna shops in neighborhood regions in 1990



(d) Diversity of neighborhood regions in 1990

Figure 9: Neighborhood region and neighborhood characteristics

As for the initial firm characteristics, X_{igr}, to control for firm and physical location characteristics, we include the following variables. We include a corner dummy indicating whether firm i is located on the corner of a large vertical street. We also include 1990 lottery block fixed effects as the locations are random with blocks. Furthermore, to control for local characteristics (e.g., areas nearby the entrance may attract more buyers), based on the neighborhood region, we construct two coordinate variables, one capturing the horizontal distance from the right end of the map of the market, another one capturing the vertical distance from the top end of the map of the market. To control for the size of firms, we include a multiple-shop dummy, as mentioned above. To control for specialization (or lack of it), we control the number of trade groups to which firm i belongs, which we call "own diversity".

In Table 2, we show the summary statistics of the key variables by specialized and non-specialized

firms and conduct a t-test for the difference in means. Non-specialized firms operate significantly more shops on average (1.90 vs. 1.66) and are more likely to be multi-shop firms (42% vs. 33%). The probability of owning a store on the corner is also higher for non-specialized firms (0.19 vs. 0.14), but this largely reflects the fact that they hold more shops: the more shops a firm owns, the greater the probability that at least one is located on a corner. The number of trade group affiliations is lower for specialized firms by construction, since specialization requires belonging to exactly one trade group. Importantly, the neighborhood-level diversity is nearly identical across the two groups, which is consistent with the random assignment of locations documented in the following section.

Table 2: Summary statistics

	Non-specialized	Specialized	Total
Number of Shops	1.90**	1.66**	1.73
	[0.09]	[0.06]	[0.05]
Dummy for Multiple Shops	0.42***	0.33***	0.36
	[0.03]	[0.02]	[0.01]
Dummy for Owning a Store on the Corner	0.19**	0.14**	0.15
	[0.02]	[0.01]	[0.01]
Number of Trade Group Affiliations (Own diversity)	1.58***	0.95***	1.13
	[0.04]	[0.01]	[0.02]
Number of Trade Groups in the Neighborhood Region	9.13	9.21	9.19
(Neighborhood diversity)	[0.10]	[0.07]	[0.06]

Notes: Standard errors of means in brackets. *, **, and *** indicate significance of the two-group t-test at the 10%, 5%, and 1% level, respectively. A firm is classified as specialized if it belongs to exactly one trade group, excluding the Kinkai and Enkai trade groups.

4 Validity of Randomization

The key assumption for our identification strategy is that the assignment of locations is random. In this section, we provide several types of test to confirm the assumption.

4.1 Balance Test

The causal interpretation of β_3 requires that Diversity_r is uncorrelated with firm characteristics—in particular those that determine Specialization_i—conditional on the lottery block. The summary statistics in Table 2 already show that the mean of neighborhood diversity is nearly identical for specialized and non-specialized firms (9.21 vs. 9.13). Here we test the assumption more systematically by regressing neighbor-

hood characteristics on firm i's attributes within lottery blocks.[27]

The results are shown in Table 3. In Column (1), we regress neighborhood diversity on firm characteristics, including trade group dummies that determine Specialization$_i$ and the number of trade group affiliations (own diversity). Since the number of firms in each trade group differs, even purely random assignment generates a mechanical correlation between firm i's trade group and the expected neighborhood composition. To absorb this, we include the market-level fraction of firms in firm i's trade group as a predicted neighborhood variable. None of the firm characteristics are significantly correlated with neighborhood diversity, with the exception of the Sushi trade group dummy, which is marginally significant at the 10 percent level. Columns (2) and (3) repeat the exercise using two other neighborhood characteristics—the fraction of same-trade-group firms and the fraction of multi-shop firms—as dependent variables, and again find no significant correlations. Overall, these results support the assumption that the lottery generates quasi-random variation in neighborhood composition.

4.2 Dartboard Approach

We check the validity of randomization in another way. If the shop location allocation is random, it should be indistinguishable from that generated by the "Dartboard approach". If shops are randomly assigned over the space in the market, there should be no agglomeration of a specific type of shop in the market. We show that shop locations in the market are randomly assigned by using the agglomeration index proposed in the literature on urban economics. Specifically, we use the bilateral distance approach proposed by Duranton and Overman (2005).[28] Duranton and Overman (2005) propose a way of testing the agglomeration of an industry. Intuitively, we first calculate the distribution of the bilateral distances between all the pairs of firms in the industry (in our case, trade group); then, we compare the obtained distribution with counterfactual distributions of bilateral distances generated by the random location assignment of firms.

Formally, let n be the number of shops belonging to the focusing trade group A, and we have $n(n-1)/2$ unique bilateral distances between shops in the trade group. Let d_{ij} be the linear distance between shops i and j.[29] The K-density estimator of bilateral distances at any distance d is

[27] Neighborhood variables can be correlated with firm characteristics across the lottery blocks because many firms can choose a block unless they did not get the desired block at the preliminary lottery.

[28] The Dartboard approach was initially developed by Ellison and Glaeser (1997). It calculates an agglomeration index using a fixed geographical unit. We calculated their agglomeration and coagglomeration indices using our neighborhood region as the unit. All the indices are close to zero. We report the results from the bilateral distance approach proposed by Duranton and Overman (2005) because this approach allows us to test the random allocation hypothesis as the null hypothesis and because the approach is the standard technique to test the concentration using micro-geographic data, as we do now.

[29] The distance here is calculated based on the unique x and y coordinates on stalls. From right to left, we assign an integer from one to a maximum number of stalls in rows as the x coordinate. From outside to inside, we similarly assign integers as the y coordinate for each stall. Based on these x and y coordinates, we calculate the Euclidean distance as the bilateral distance

Table 3: Balance test results

Dependent	(1) Neighborhood diversity	(2) Fraction of firms of the same trade group	(3) Fraction of multi-store firms
Dummy for multiple shops	-0.105	0.00493	-0.00739
	(0.112)	(0.00736)	(0.0104)
Corner	-0.0966	0.0115	-0.00529
	(0.0967)	(0.00786)	(0.00475)
Own diversity	0.661	0.0542	0.0257
	(0.594)	(0.0383)	(0.0361)
Enkanmono	-0.287	0.0259	0.00642
	(0.626)	(0.0380)	(0.0291)
Octopus	0.524	0.0155	-0.0341
	(0.651)	(0.0461)	(0.0465)
Aimono	-0.197	-0.00176	-0.0289
	(0.521)	(0.0393)	(0.0331)
Enkai	-0.264	0.00396	-0.0134
	(0.387)	(0.0289)	(0.0220)
Enyo	-0.0563	0.0372	-0.0188
	(0.976)	(0.0615)	(0.0465)
Kinkai	-0.673	-0.00344	0.00456
	(0.416)	(0.0300)	(0.0273)
Sushi	-0.371*	0.0159	-0.0109
	(0.215)	(0.0263)	(0.0118)
Renseihin	-0.0722	0.0105	-0.032
	(0.560)	(0.0366)	(0.0323)
Tansui	-0.151	0.0216	-0.0291
	(0.779)	(0.0521)	(0.0377)
Hokuyo	-0.435	0.0648	-0.0174
	(0.514)	(0.0398)	(0.0288)
Shrimp	-0.0785	0.0257	-0.00909
	(0.471)	(0.0407)	(0.0316)
Iseebi	-0.596	-0.026	-0.0457
	(0.849)	(0.0552)	(0.0412)
Predicted neighborhood variable	yes	yes	yes
Block dummy	yes	yes	yes
N	1064	1064	1064

Note: Standard errors are clustered at the region level and shown in parentheses. * indicates significance at the 10% level.

$$\hat{K}_A(d) = \frac{1}{n(n-1)h} \sum_{i=1}^{n} \sum_{j=1}^{n} f\left(\frac{d - d_{ij}}{h}\right),$$

where h is the bandwidth set as the optimal bandwidth as proposed by Silverman (1986), and f is the Gaussian kernel function.

Then, we consider the case where shops are randomly assigned in the market by simulation. In the simulation, we randomly assign n shop locations over all possible sites in the market and estimate the bilateral distance distribution under the simulated location distributions. By iterating this trial 1,000 times, we construct global confidence bands, i.e., an upper confidence band $K_A^U(d)$ and a lower confidence band $K_A^L(d)$, such that of the 1,000 randomly drawn K-densities, 95% lie below the upper band and 95% lie above the lower band over the entire distance range we focus on, $[0, d_{max}]$.[30] If $\hat{K}_A(d) > K_A^U(d)$ for at least one $d \in [0, d_{max}]$, trade group A is defined as globally localized at the 5% confidence level. On the other hand, if $\hat{K}_A(d) < K_A^L(d)$ for at least one d and trade group A is not defined as localized, we define trade group A as globally dispersed at the 5% level. In other words, if the estimated K-density lies between the upper and lower confidence bands, we cannot reject the null hypothesis that the shops in trade group A are randomly allocated at the 5% level.

Neighborhood effects might arise not only within the same trade group but also across pairs of different trade groups. For example, as the anecdotal episodes mentioned, Sushi shops may benefit from Tuna shops being located nearby. If so, Sushi shops may choose their locations close to Tuna shops by applying the lottery jointly. To test the possibility of a concentration of shops between trade groups, we also conduct the test for coagglomeration proposed by Duranton and Overman (2005).

To test the coagglomerations, consider trade group A with n_A firms and trade group B with n_B firms. We calculate all the bilateral distances between firms in trade groups A and B. Thus, the K-density estimator is modified as follows,

$$\hat{K}_{AB}(d) = \frac{1}{n_A n_B h} \sum_{i=1}^{n_A} \sum_{j=1}^{n_B} f\left(\frac{d - d_{ij}}{h}\right).$$

Similar to the single trade group concentration, we can construct the confidence bands by counterfactual simulations that the locations of shops in trade groups A and B are randomly assigned.

As we mentioned, the location lotteries are conducted in two steps. In the first step, shops can choose

between shops.

[30] We set the threshold value of the maximum distance d_{max} as the median distance between all pairs of all industries, as suggested by Duranton and Overman (2005)

their preferable blocks. In this sense, the choice of blocks cannot be random. To address this issue, we conduct the analysis within the lottery block. That is, we restrict shops in a trade group in a block, and then we test their agglomeration. Thus, the results are obtained by the combination of trade groups and blocks.

As the results, on agglomeration, there are 31 testable combinations of trade groups and block.[31] There is no trade group significantly agglomerated at the five percent level. Similarly, on coagglomeration, there are 192 testable combinations of trade group and block, and there is no trade group significantly coagglomerated at the five percent level. Both in agglomeration and coagglomeration, there is no deviation from random assignment of locations.[32]

Furthermore, to check the overall tendency of the allocation of shops, we also conduct an analysis pooling the trade groups. That is, we calculate the bilateral distances between firms in each trade group, and then we pool all the calculated bilateral distances and estimate the K-density. The K-density for agglomeration is modified as follows.

$$\hat{K}_{\text{agg}}(d) = \frac{1}{h \sum_{A \in \mathfrak{A}} n_A (n_A - 1)} \sum_{A \in \mathfrak{A}} \left[\sum_{i=1}^{n_A} \sum_{j=1}^{n_A} f\left(\frac{d - d_{ij}}{h} \right) \right]$$

where $\mathfrak{A}$ is the set of trade groups. On the counterfactual simulation, we reshuffle all shops in all the trade groups simultaneously in each trial.

Similarly, we conduct an analysis for coagglomeration pooling the pairs of trade groups. In this case, we calculate bilateral distances between firms in each pair of trade groups. Then, we pool all the calculated bilateral distances and estimate the K-density as follows,

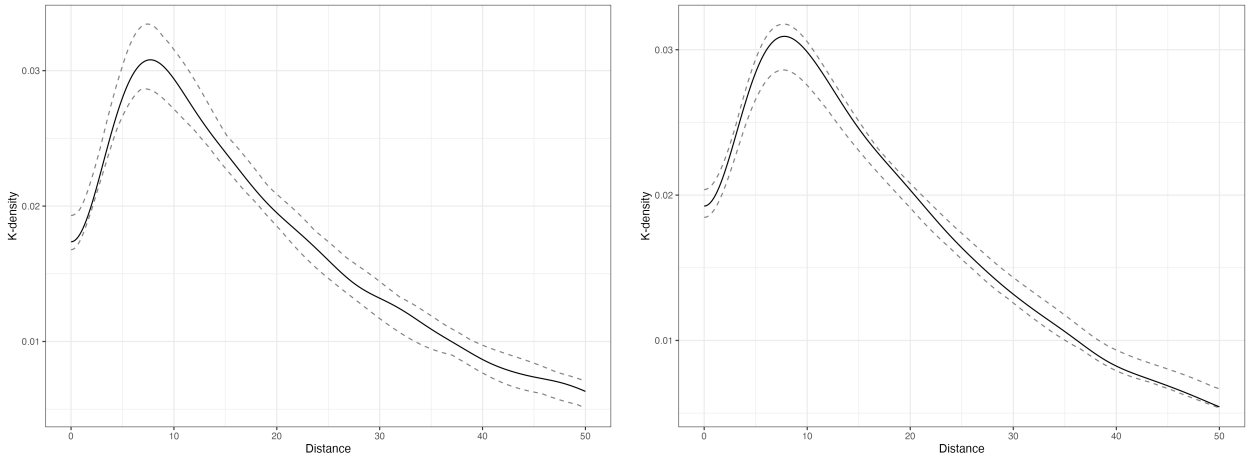
$$\hat{K}_{\text{coagg}}(d) = \frac{1}{h \sum_{\{A,B\} \in \mathfrak{A} \times \mathfrak{A}} n_A n_B} \sum_{\{A,B\} \in \mathfrak{A} \times \mathfrak{A}} \left[\sum_{i=1}^{n_A} \sum_{j=1}^{n_B} f\left(\frac{d - d_{ij}}{h} \right) \right]$$

The results are shown in Figure 10.

Figure 10a shows the result of agglomeration, and Figure 10b shows the result of coagglomeration. In both figures, the estimated K-densities are located inside the confidence bands. This implies that the actual location distribution of shops in the Tsukiji fish market does not differ from counterfactual random allocations.

[31] Some combinations have the number of observations too small for this test to be feasible.

[32] We conducted additional analysis where we run the regressions of a dummy indicating if there is a shop of trade group g within next K shops on the set of trade group dummies and block fixed effects. We find almost no statistically significant coefficients for any group g and for K between 1 and 10. The results are available upon request.



(a) Agglomeration (b) Coagglomeration

Figure 10: Results of pooled analysis

5 Results

5.1 Baseline: Specialization and Neighborhood Diversity

The hypothesis stated in Section 3.2 implies that neighborhood diversity should benefit specialized firms in particular. Firms dealing in a narrow line of fish are precisely the firms whose products are likely to be purchased as part of a shopping trip that involves complementary products from neighboring stores, and therefore the firms whose performance is most exposed to the shopping externality generated by the diversity of the surrounding neighborhood. We test this prediction by interacting neighborhood diversity with a measure of firm-level specialization, using the random variation in neighbors generated by the 1990 location lottery.

We estimate the baseline specification, in which we interact neighborhood diversity with the specialization indicator Specialization_i. The coefficient of interest is β_3, the differential effect of neighborhood diversity for specialized firms relative to non-specialized firms. We consider three horizons after the lottery (1991, 1993, and 1995) and stop in 1995, which is the year of the next location lottery.

In Panel A of Table 4, the coefficient on the interaction between specialization and neighborhood diversity is positive and statistically significant in all three columns. A specialized firm assigned by the 1990 lottery to a neighborhood region with one additional unique trade group enjoyed approximately a 5–6 percent higher growth in sales relative to a non-specialized firm in the same neighborhood, depending on the horizon. The magnitude is stable across the three horizons we consider, indicating that the differential gain accruing to specialized firms in diverse neighborhoods emerges quickly after the lottery and persists

Table 4: Baseline: Specialization and neighborhood diversity

	(1)	(2)	(3)
	(1989–1991)	(1989–1993)	(1989–1995)
Panel A: Δ log Sales			
Specialization $\times$ Neighborhood Diversity	0.0521*	0.0548**	0.0559*
	(0.0267)	(0.0260)	(0.0317)
Other covariates	✓	✓	✓
N	287	282	276
Panel B: Δ log Employment			
Specialization $\times$ Neighborhood Diversity	0.0232*	0.0390**	0.0374
	(0.0130)	(0.0181)	(0.0252)
Other covariates	✓	✓	✓
N	285	280	271

Note: Standard errors clustered at the neighborhood region level are in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively. Specialization is a dummy variable taking one if a firm belongs to exactly one trade group, excluding the Kinkai and Enkai trade groups. Neighborhood diversity is the number of unique trade groups represented by the other firms in the same neighborhood region. All specifications include the specialization dummy, neighborhood diversity, block fixed effects, trade group fixed effects, area coordinates, and the corner dummy.

until the next lottery, at which point a new random reassignment replaces the variation we exploit.

Panel B of Table 4 repeats the exercise using log employment as the outcome. The coefficient on the interaction is again positive in all three columns and statistically significant in the first two. The magnitude is somewhat smaller than for sales, which is consistent with the view that employment adjusts more slowly than sales but eventually moves in the same direction. Together, the two tables provide direct evidence that the neighborhood diversity effect operates primarily through specialized firms, exactly the pattern we would expect if the underlying mechanism is consumption externalities arising from the complementarity of specialized products.

It is also useful to put the magnitudes in context. Given that the standard deviation of neighborhood diversity is approximately two, one standard deviation increase in diversity is therefore associated with approximately 10 percent higher sales growth and 4–7 percent higher employment growth for specialized firms relative to non-specialized firms in the same neighborhood. Given that average sales growth over a comparable window is in the single digits, these effects are economically substantial.

To further examine the dynamics of the effect more flexibly, we estimate an event study version of the regression. Specifically, we estimate

$$y_{igrt} = \sum_{\tau \neq 1990} \delta_\tau \left(\text{Diversity}_r \times \text{Specialization}_i \times \mathbf{1}\{t = \tau\} \right) + \alpha_i + \eta_t + \gamma X_{igr,t} + \varepsilon_{igrt},$$

where y_{igrt} is the log of sales or employment in year t, α_i is a firm fixed effect, η_t is a year fixed effect, and 1990 is the omitted year. The coefficients δ_τ trace out the differential evolution of the outcome for specialized firms exposed to diverse neighborhoods, relative to non-specialized firms in the same neighborhoods.

As in the baseline specification, Diversity_r is measured once using the 1990 lottery outcomes and held fixed across all years; it does not vary over time. Because we do not observe pre-1990 shop locations, the pre-lottery coefficients $\hat{\delta}_\tau$ for $\tau < 1990$ test whether firms that were *subsequently* assigned to diverse neighborhoods were already on differential trajectories before the reassignment occurred. Under random assignment, we expect these pre-lottery coefficients to be zero.

The specification also includes two-way interactions $\text{Diversity}_r \times \mathbf{1}\{t = \tau\}$ and $\text{Specialization}_i \times \mathbf{1}\{t = \tau\}$ within $X_{igr,t}$, so that δ_τ captures the differential effect of diversity for specialized firms net of any level effect of diversity or specialization in year τ. Because firm fixed effects absorb all time-invariant controls (including Specialization_i, the corner dummy, coordinates, and own diversity), the remaining controls in $X_{igr,t}$ consist of these two-way interactions.

Figures 11 and 12 display the estimated coefficients $\hat{\delta}_\tau$ together with their 95 percent confidence intervals for log sales and log employment, respectively.

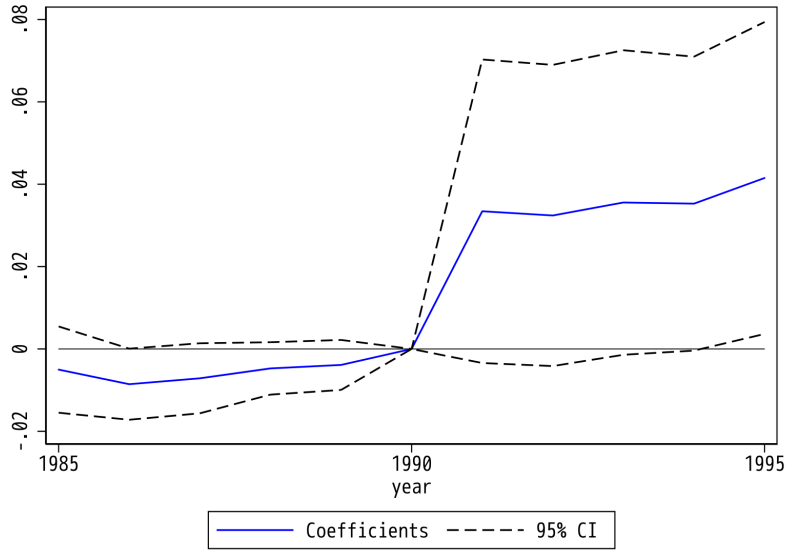


Figure 11: Event study: Specialization × neighborhood diversity, log sales

Note: The figure plots the coefficients $\hat{\delta}_\tau$ from the event study specification described in the text, with 95 percent confidence intervals based on standard errors clustered at the neighborhood region level. The vertical line indicates the timing of the 1990 lottery.

Two features of the figures are noteworthy. First, the pre-lottery coefficients are small and statistically indistinguishable from zero, supporting the parallel trends assumption: specialized firms that ended up

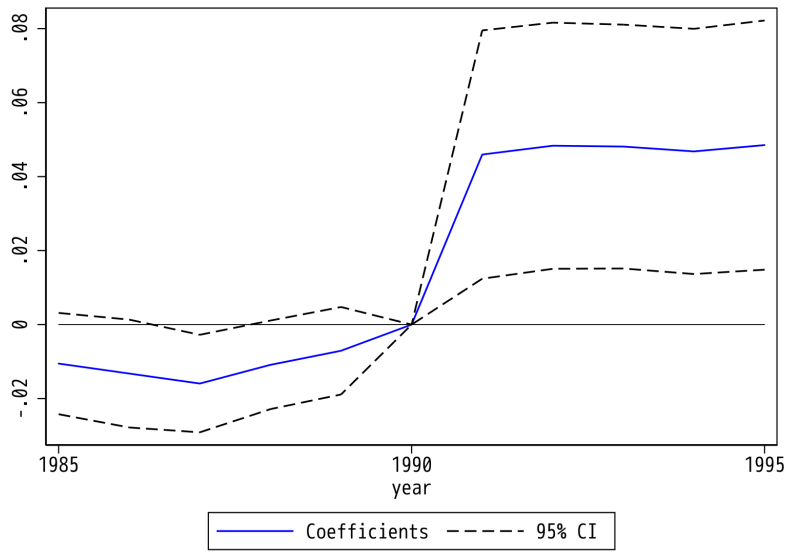


Figure 12: Event study: Specialization × neighborhood diversity, log employment

Note: See Figure 11. The dependent variable is the log of employment.

in diverse neighborhoods were not on a different trajectory before the lottery. Given the random assignment of locations established in Section 4, this is what we should expect, and the event study provides direct visual confirmation. Second, the post-lottery coefficients become positive and statistically significant within a year or two of the lottery and remain so until the next lottery. The pattern is consistent across the sales and employment outcomes. This dynamic pattern strengthens our causal interpretation by showing that the differential gain accruing to specialized firms in diverse neighborhoods materializes after, and not before, the lottery that randomly reassigned their neighbors.

5.2 Mechanism: Evidence from Customer Composition

The results in Section 5.1 establish that specialized wholesalers benefit disproportionately from being assigned to neighborhoods with diverse trade groups. This pattern is consistent with consumption externalities arising from buyers' shopping behavior, but it is not, on its own, decisive: one might worry that the specialization variable picks up some unobserved firm characteristic correlated with the responsiveness to local conditions for reasons unrelated to shopping. In this section, we provide an independent test of the shopping-channel interpretation by exploiting heterogeneity in the customer composition of the wholesalers.

Identifying buyers prone to shopping around

As discussed in Section 2.4, certain buyers are far more prone to "shopping around" multiple wholesalers within a single trip than others. Sushi restaurant chefs are the prototypical example. Because the daily availability of high-quality fish varies and because preparing a sushi menu requires several different fish species of high quality, sushi chefs typically visit multiple wholesalers in one morning to assemble their daily assortment. In contrast, buyers from supermarkets and other large-scale customers tend to procure standardized fish from a small number of dedicated suppliers and have less incentive to walk around the market. Under the hypothesis that the neighborhood diversity effect operates through shopping externalities, wholesalers whose main clients are sushi restaurants should benefit more from neighborhood diversity than wholesalers serving other client types.

To operationalize this idea, we use customer information from the TDB dataset. For each firm in the TDB data, the records include up to five reported customers. The customer field in TDB does not require respondents to list customers by formal corporate name. We define the dummy variable $SushiMain_i$ to take one if the first reported customer in the TDB data refers to a sushi restaurant—either as a generic mention of "sushi restaurants" or as the name of an identifiable sushi restaurant—and zero otherwise. Crucially, we take this customer field from the 1993 wave of the TDB data, which is the earliest year for which this information is available.[33]

Specification and results

We estimate the analog of the baseline specification, replacing the specialization indicator with the sushi-main-client indicator:

$$\Delta y_{igr} = \beta_0 + \beta_1 \text{Diversity}_r + \beta_2 \text{SushiMain}_i + \beta_3 \left(\text{Diversity}_r \times \text{SushiMain}_i\right) + \gamma X_{igr} + \lambda_g + \varepsilon_{igr},$$

where the controls are as in Section 5.1. The coefficient of interest, β_3, captures the differential effect of neighborhood diversity for firms whose main clients are sushi restaurants. The shopping-channel hypothesis predicts $\beta_3 > 0$.

The customer-composition variable is constructed from data independent of the specialization measure and is not mechanically related to either neighborhood diversity or the trade group structure. So this analysis provides a complementary test of the shopping mechanism: while the specialization analysis identifies firms whose *products* are likely to be subject to shopping-around behavior, the customer-composition

[33] The type of fish the firms deal with is a highly persistent, structural feature of their business in this setting.

analysis identifies firms whose *customers* are likely to engage in such behavior.

Table 5: Sushi main client and neighborhood diversity

	(1)	(2)	(3)
	(1989–1991)	(1989–1993)	(1989–1995)
Panel A: Δ log Sales			
Sushi Main Client × Neighborhood Diversity	0.0588*	0.0624*	0.0645
	(0.0313)	(0.0344)	(0.0401)
Other covariates	✓	✓	✓
N	281	282	276
Panel B: Δ log Employment			
Sushi Main Client × Neighborhood Diversity	0.0613**	0.0632**	0.0470
	(0.0244)	(0.0271)	(0.0300)
Other covariates	✓	✓	✓
N	279	280	271

Note: Standard errors clustered at the neighborhood region level are in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively. Sushi Main Client is a dummy variable taking one if the first customer reported by the firm in the Teikoku Databank (TDB) data refers to a sushi restaurant, either as a generic mention (e.g., "sushi restaurants") or as the name of an identifiable sushi restaurant. All specifications include the Sushi Main Client dummy, neighborhood diversity, block fixed effects, trade group fixed effects, area coordinates, and the corner dummy.

Table 5 reports the results for log sales and log employment, respectively. The pattern in Panels A and B of Table 5 closely mirrors the baseline pattern in Table 4. The interaction between SushiMain_i and neighborhood diversity is positive across both outcomes and all horizons, and is statistically significant at the 10 percent level or better in four of the six columns. The magnitudes are similar to, and if anything somewhat larger than, those obtained for the specialization interaction.

Two remarks are in order. First, taken together with the baseline results, these estimates suggest that the firms that benefit most from neighborhood diversity are precisely those most exposed to shopping-around behavior, whether identified on the supply side (what a firm sells) or on the demand side (whom a firm sells to). This evidence makes it difficult to reconcile our findings with mechanisms other than consumption externalities arising from shopping behavior. For example, a story based on knowledge spillovers across firms would not naturally predict that the effect is concentrated among firms serving customers known to walk around the market. Second, the customer-composition test is informative beyond the specialization test because being a specialized firm and serving sushi restaurants are conceptually distinct: some specialized firms sell mainly to processors, while some non-specialized firms serve sushi restaurants for a subset of their products. The fact that both interactions are positive and significant indicates that the mechanism operates through the demand-side shopping behavior of customers, regardless of the precise

supply-side specialization status.

5.3 Robustness

5.3.1 Alternative Measure of Neighborhood Diversity

Our baseline results use the number of unique trade groups in the neighborhood region as the measure of diversity. To confirm that the findings are not driven by this particular choice, we re-estimate the baseline specification replacing the trade-group count with the Herfindahl–Hirschman Index (HHI) of the distribution of shops across trade groups within the neighborhood region. The HHI is an inverse measure of diversity: a higher HHI indicates a more concentrated, less diverse neighborhood. The correlation between the HHI and the number of unique trade groups is -0.7.

Table A1 in Appendix reports the results. The interaction between specialization and the HHI-based diversity measure is negative and statistically significant for employment, particularly at longer horizons, and negative though imprecisely estimated for sales. Since a higher HHI corresponds to lower diversity, the negative sign on the interaction term is consistent with the baseline finding that neighborhood diversity benefits specialized firms. The results are therefore robust to the choice of the neighborhood diversity measure.

5.3.2 Full-sample results using directory-based outcomes

Our baseline regressions use sales and employment from the TDB data, which limits the sample that could be matched to TDB records. Because the full population of intermediate wholesalers recorded in the Tsukiji directories numbers over 1,000 firms, a natural concern is whether our findings extend beyond the matched subsample.

To address this concern, we use two outcomes that are available for every firm in the directory without requiring a TDB match: a dummy for exit between 1990 and 1995, and the change in the number of shops over the same period. The number of shops can be considered a measure of a firm's performance, as firms extend their shops by purchasing the right of stall-holding from other firms when their business grows. We confirmed that the number of shops is strongly positively correlated with sales and employment using the subset of firms that can be linked to the TDB dataset. If specialized firms assigned to diverse neighborhoods experience genuinely better performance, we would expect them to be less likely to exit and more likely to expand their shop holdings.

Table A2 reports the results. Columns (1)–(3) use the full sample of 1,064 intermediate wholesalers.

The interaction between specialization and neighborhood diversity is negative and significant for exit and positive and significant for the change in the number of shops, indicating that specialized firms in diverse neighborhoods are less likely to exit and more likely to expand. These results confirm that the neighborhood diversity effect documented in the baseline extends to the full population of wholesalers and is not an artifact of the TDB matching process. Columns (4)–(6) repeat the exercise on the TDB-matched subsample. The point estimates are similar in sign and magnitude, but the coefficients are no longer statistically significant, consistent with the substantially smaller sample size.

5.3.3 Front versus back neighbors

A distinctive feature of the Tsukiji market building is that shops are arranged in rows separated by horizontal corridors through which buyers walk. Shops facing each other across a corridor share the flow of buyers, but shops that are back-to-back do not: the backyard area between them has no corridor and buyers cannot pass through it (Figure 13). This physical layout allows us to distinguish the shopping-externality channel from other forms of neighborhood effects. If the diversity effect operates through shared buyer flow, it should not arise from back neighbors (those behind the shop). If, instead, the effect were driven by knowledge spillovers or other interactions between shop owners and workers—which are equally possible across the backyard—we would expect to observe an effect from back neighbors as well.

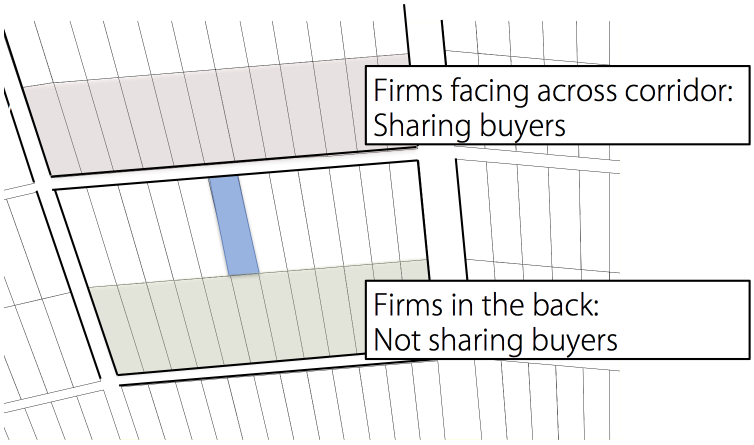


Figure 13: Front and back neighbors

Note: The figure illustrates the arrangement of shops in the Tsukiji market building. Shops facing each other across a horizontal corridor (front neighbors) share the flow of buyers. Shops that are back-to-back (back neighbors) are separated by a backyard area with no corridor for buyers.

To implement this test, we redefine the neighborhood to include only back neighbors and re-estimate the baseline specification. Table A3 in Appendix reports the results. In sharp contrast to the baseline es-

timates, the interaction between specialization and back-neighbor diversity is small, statistically insignificant, and inconsistent in sign across all horizons and both outcomes. The absence of any effect from back neighbors provides direct evidence that the neighborhood effects we estimate operate through the sharing of buyer flow rather than through knowledge spillovers or other proximity-based interactions.

6 Conclusion

This paper empirically investigates neighborhood effects among intermediate wholesalers in the Tokyo Tsukiji Fish Market. We address the identification difficulty originating from the self-selection of firm location by exploiting a unique feature of the market: shop locations are periodically reassigned by relocation lotteries. Using this quasi-random variation, we estimate the causal effect of the diversity of neighboring firms on firm performance. We find that the diversity of the neighboring firms positively affects the performance of specialized firms. Then, exploiting information on customer composition, we show that the diversity effect is concentrated among wholesalers whose main clients are sushi restaurants—buyers known to visit multiple wholesalers in a single trip. Together, these results provide direct evidence that the neighborhood effects we estimate arise from shopping externalities rather than knowledge spillovers or other proximity-based interactions. Overall, our analysis shows that the complementarity of products between specialized diverse stores, a mechanism central to theoretical models of retail clustering, is an empirically important factor behind urban agglomeration.

A Robustness Tables

References

Arakawa, Kiyoshi, "A Model of Shopping Centers," *Journal of Regional Science*, 2006, *46* (5), 969–990.

Benmelech, Efraim, Nittai Bergman, Anna Milanez, and Vladimir Mukharlyamov, "The Agglomeration of Bankruptcy," *Review of Financial Studies*, 2018, *32* (7), 2541–2586.

Bester, Theodore C., *Tsukiji: The Fish Market at the Center of the World*, Berkeley: University of California Press, 2004.

Cosman, Jacob, "Industry Dynamics and the Value of Variety in Nightlife: Evidence from Chicago," 2014. Mimeo.

Table A1: Robustness: Specialization and neighborhood diversity (HHI)

	(1)	(2)	(3)
	(1989–1991)	(1989–1993)	(1989–1995)
Panel A: Δ log Sales			
Specialization × Neighborhood HHI	−1.004	−0.745	−1.663
	(1.095)	(1.031)	(1.290)
Other covariates	✓	✓	✓
N	242	240	238
Panel B: Δ log Employment			
Specialization × Neighborhood HHI	−2.203**	−2.370**	−3.876***
	(1.083)	(1.093)	(1.115)
Other covariates	✓	✓	✓
N	241	239	234

Note: Standard errors clustered at the neighborhood region level are in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively. The specification is the same as in Table 4, except that neighborhood diversity is measured by the Herfindahl–Hirschman Index (HHI) of the distribution of shops across trade groups within the neighborhood region. A higher HHI indicates lower diversity, so a negative coefficient on the interaction term corresponds to a positive effect of diversity for specialized firms.

Table A2: Robustness: Specialization and neighborhood diversity, exit and changes in number of shops

	Full sample			TDB-matched sample		
	(1)	(2)	(3)	(4)	(5)	(6)
	Exit	Δ No. of shops (all)	Δ No. of shops (survivors)	Exit	Δ No. of shops (all)	Δ No. of shops (survivors)
Specialization × Neighborhood Diversity	−0.0116***	0.0369***	0.0148*	−0.00822	0.0366	0.0186
	(0.00376)	(0.0115)	(0.00813)	(0.00602)	(0.0283)	(0.0284)
Other covariates	✓	✓	✓	✓	✓	✓
N	1064	1064	951	242	242	234

Note: Standard errors clustered at the neighborhood region level are in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively. Columns (1) and (4) report linear probability models for exit between 1990 and 1995. Columns (2) and (5) report the change in the number of shops between 1990 and 1995 for all firms (coding exiters as losing all shops). Columns (3) and (6) restrict to firms that survived until 1995. Columns (1)–(3) use the full sample of intermediate wholesalers; columns (4)–(6) restrict to the TDB-matched sample used in the baseline regressions.

Table A3: Placebo: Back-neighbor diversity and specialization

	(1)	(2)	(3)
	(1989–1991)	(1989–1993)	(1989–1995)
Panel A: Δ *log Sales*			
Specialization × Back-Neighbor Diversity	−0.0212	−0.0044	−0.0282
	(0.0253)	(0.0300)	(0.0348)
Other covariates	✓	✓	✓
N	242	240	234
Panel B: Δ *log Employment*			
Specialization × Back-Neighbor Diversity	−0.00906	−0.0203	−0.0061
	(0.0176)	(0.0217)	(0.0226)
Other covariates	✓	✓	✓
N	241	239	229

Note: Standard errors clustered at the neighborhood region level are in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively. The specification is the same as in Table 4, except that neighborhood diversity is computed using only back neighbors—shops that are back-to-back with the focal shop and do not share a corridor for buyers.

Couture, Victor, "Valuing Consumption Benefits of Urban Density," 2015. Mimeo.

Duranton, Gilles and Henry G. Overman, "Testing for Localization Using Micro-geographic Data," *Review of Economic Studies*, 2005, *72* (4), 1077–1106.

Ellison, Glenn and Edward L. Glaeser, "Geographic Concentration in US Manufacturing Industries," *Journal of Political Economy*, 1997, *105* (5), 889–927.

Glaeser, Edward L., Jed Kolko, and Albert Saiz, "Consumer City," *Journal of Economic Geography*, 2001, *1* (1), 27–50.

Handbury, Jessie and David E. Weinstein, "Goods Prices and Availability in Cities," *Review of Economic Studies*, 2015, *82* (1), 258–296.

Henkel, Joachim, Konrad Stahl, and Uwe Walz, "Coalition Building in a Spatial Economy," *Journal of Urban Economics*, 2000, *47* (1), 136–163.

Holmes, Thomas J. and John J. Stevens, "An Alternative Theory of the Plant Size Distribution, with Geography and Intra and International Trade," *Journal of Political Economy*, 2014, *122*, 369–421.

Jardim, Eduardo, "All in the Mix: Spillovers and the Agglomeration of Neighborhood Retail," 2015. Mimeo.

Konishi, Hideo, "Concentration of Competing Retail Stores," *Journal of Urban Economics*, 2005, *58* (3), 488–512.

__ **and Michael T. Sandfort**, "Anchor Stores," *Journal of Urban Economics*, 2003, *53*, 413–435.

Koster, Hans R. A., Ilias Pasidis, and Jos van Ommeren, "Shopping Externalities and Retail Concentration: Evidence from Dutch Shopping Streets," *Journal of Urban Economics*, 2019, *114*, 103194.

Leonardi, Marco and Enrico Moretti, "The Agglomeration of Urban Amenities: Evidence from Milan Restaurants," *American Economic Review: Insights*, 2023, *5* (2), 141–157.

Matsuyama, Kiminori, "Making Monopolistic Competition More Useful," Working Paper E-92-18, Hoover Institution 1992.

__ , "Complementarities and Cumulative Processes in Models of Monopolistic Competition," *Journal of Economic Literature*, 1995, *33* (2), 701–729.

Miyauchi, Yuhei, Kentaro Nakajima, and Stephen J. Redding, "The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data," *Quarterly Journal of Economics*, November 2025, *140* (4), 2507–2570.

Oh, Ryungha and Jaeeun Seo, "What Causes Agglomeration of Services? Theory and Evidence from Seoul," 2023. Yale University, mimeo.

Rhodes, Andrew and Jidong Zhou, "Consumer Search and Retail Market Structure," *Management Science*, 2019, *65* (6), 2607–2623.

Schiff, Nathan, "Cities and Product Variety: Evidence from Restaurants," *Journal of Economic Geography*, 2015, *15* (5), 1085–1123.

Shoag, Daniel and Stan Veuger, "Shops and the City: Evidence on Local Externalities and Local Government Policy from Big-Box Bankruptcies," *Review of Economics and Statistics*, 2018, *100* (3), 440–453.

Silverman, Bernard W., *Density Estimation for Statistics and Data Analysis*, London: Chapman & Hall, 1986.

Tabuchi, Takatoshi, "Self-organizing Marketplaces," *Journal of Urban Economics*, 2009, *66* (3), 179–185.

Wolinsky, Asher, "Retail Trade Concentration due to Consumers' Imperfect Information," *Bell Journal of Economics*, 1983, *14* (2), 275–282.