

# Game-Theoretic Cafeteria Selection: A Nash Equilibrium Framework for Lunch Decisions Under Incomplete Information and Hunger Pressure

Dr. Rupert Q. Finchley<sup>1,\*</sup> Xinyun Zhu<sup>2</sup>

<sup>1</sup>Centre for Applied Indecision Studies, University of Idleworth, UK

<sup>2</sup>Department of Passive Participation, Freeloading Institute of Technology, China

\*Corresponding author: jz8179288@gmail.com

## Abstract

Everyone who has eaten in a cafeteria with more than two food stalls has faced the same problem: you pick a line, the line doesn't move, and the person who walked in after you is already eating. This paper models that experience.

We formalize cafeteria lunch selection as a simultaneous-move game with  $N$  agents, asymmetric information, and time-dependent payoffs. Each agent must commit to a food stall without knowing others' choices, under a hard constraint that lunch must be acquired and consumed within a fixed window  $T$ . We derive a mixed-strategy Nash equilibrium that predicts overcrowding at stalls perceived as “good today,” and show that this equilibrium is Pareto-inefficient—everyone would be better off if they could coordinate, but they can't, because nobody actually texts back in the group chat.

Field data from 11 cafeterias across 4 universities ( $n = 2,461$  lunch decisions) confirms the model. We document the Queue Regret Phenomenon, the Sunk Cost Tray Effect, and the disturbing frequency with which people choose the same disappointing stall repeatedly. We also find that the presence of a daily special acts as a Schelling focal point that systematically destroys welfare for everyone except the first 15 people in line.

Keywords: cafeteria game, Nash equilibrium, lunch decisions, queue regret, information asymmetry, Schelling focal point, daily special trap, sunk cost tray, hunger discount factor

## 1 Introduction

The lunch decision is, on its surface, trivial. You walk into a cafeteria. There are stalls. You pick one. You eat. You leave. A child could do it.

And yet, somehow, millions of adults get this wrong every single day.

The problem is not that the food is bad—though it often is. The problem is that the cafeteria is a game, and its players don't know they're playing it. Every person who walks through the door is simultaneously choosing which stall to join, with no communication, incomplete information about what each stall is serving today, and a biological clock (hunger) that penalizes deliberation. The result is a coordination failure so reliable you could set your watch by it: the “good” stall has a 30-minute line, the “bad” stall is empty, and the person standing in the 30-minute line is Googling nearby restaurants on their phone.

This paper takes the problem seriously. We model the cafeteria lunch decision as a simultaneous  $N$ -player game with asymmetric information, derive the Nash equilibrium, and show that

it sucks. Not in a technical sense—though it is Pareto-inefficient—but in the everyday sense that the equilibrium outcome is: you wait too long, you eat something mediocre, and you feel vaguely wronged by the universe.

We validate the model with field data from 11 university cafeterias. We find that observed behavior is consistent with the predicted equilibrium to a troubling degree.

### 1.1 Why this matters

At a cafeteria serving 2,000 people per lunch period, an average queue time of 14 minutes translates to approximately 467 person-hours lost per day. Across a 200-day academic year, that’s 93,400 hours—roughly 10.7 person-years—spent standing in line, staring at phones, and silently judging the person ahead who is still deciding between the chicken and the fish.

If even half of this time could be recovered through better coordination, the resulting productivity gains would fund approximately one additional postdoc position per university per year.<sup>1</sup>

We are not the first to note that queuing is wasteful. We are, however, the first to model it as a game played by hungry people who are bad at making decisions, which is what it actually is.

## 2 Related work

Congestion games have a long history in economics and computer science. Rosenthal (1973) proved that every congestion game has at least one pure-strategy Nash equilibrium. This result is mathematically elegant and practically useless in the cafeteria setting, because it assumes agents know the payoff functions, which in a cafeteria means knowing whether the braised pork is good today, and nobody knows that until they’ve committed.

Varian (1980) studied consumer search with price dispersion, a framework we adapt by replacing “price” with “queue length” and “consumer” with “person who skipped breakfast.” Hotelling’s (1929) spatial competition model is relevant in principle—food stalls are literally positioned in space—but his model assumes rational choice, which breaks down when one stall has a hand-written sign saying “Chef’s Special” in red marker.

The closest work to ours is Schelling (1960), whose focal point theory explains why everyone converges on the same option even without communication. In our data, the daily special functions as a Schelling focal point that reliably attracts 40% of the cafeteria population to a single stall, creating exactly the congestion that game theory predicts and human beings never learn from.

Kahneman and Tversky’s (1979) prospect theory is also relevant: we find that people overweight the regret of “missing” a good stall and underweight the cost of waiting, a framing effect we call the Queue Regret Asymmetry.

## 3 The model

### 3.1 Setup

Consider a cafeteria with  $K$  food stalls and  $N$  agents who arrive during a lunch window  $[0, T]$ . Each agent  $i$  must select exactly one stall  $k \in \{1, \dots, K\}$ . Agents choose simultaneously (or close enough—the stall commitment happens within seconds of entering, and switching after joining a queue is socially costly).

Each stall  $k$  has a quality  $q_k$  drawn from a distribution  $F$  at the start of each day. The quality is unknown to agents but partially inferred from:

---

<sup>1</sup>Assuming a postdoc works 2,000 hours per year, which is generous.

- Observable signals: queue length  $L_k$ , steam/smell intensity  $s_k$ , and whether the stall has a promotional sign  $\sigma_k \in \{0, 1\}$ .
- Priors: past experience at stall  $k$ , summarized by a Bayesian belief  $\hat{q}_{ik}$ .
- Gossip: occasionally, an agent has information from someone who already ate (probability  $\rho \approx 0.08$ ).

### 3.2 Payoff function

Agent  $i$ 's payoff from choosing stall  $k$  is:

$$u_i(k) = q_k \cdot \delta^{w_k(n_k)} - c \cdot w_k(n_k) - r_i \cdot \max_{j \neq k} (\hat{q}_{ij} - q_k)^+ \quad (1)$$

where:

- $\delta \in (0, 1)$  is the hunger discount factor: food quality perceived declines the longer you wait, because hunger makes you both more desperate and more resentful
- $w_k(n_k)$  is the wait time at stall  $k$  given  $n_k$  agents choosing it, with  $w_k(n_k) = \alpha_k \cdot n_k + \epsilon_k$  and  $\alpha_k$  being the service rate (how fast the aunty scoops)
- $c$  is the per-minute cost of waiting (opportunity cost + annoyance + physical discomfort from standing)
- $r_i$  is agent  $i$ 's regret sensitivity: the psychological cost of observing, post-hoc, that another stall would have been better
- $(\cdot)^+$  denotes  $\max(\cdot, 0)$

The regret term is what makes cafeteria choice different from standard congestion games. In a road network, you don't usually see the road you didn't take. In a cafeteria, you do. You sit down with your sad noodles and watch someone at the next table eating something that looks much better, and it bothers you in a way that transport economists have not adequately modeled.

### 3.3 Signal extraction problem

Agents face a classic inference problem: a long queue at stall  $k$  may indicate that (a) the food is good today, or (b) the food is always mediocre but the stall is in a convenient location. Formally:

$$P(q_k \text{ high} \mid L_k \text{ long}) = \frac{P(L_k \text{ long} \mid q_k \text{ high}) \cdot P(q_k \text{ high})}{P(L_k \text{ long})} \quad (2)$$

In practice, nobody computes this. What actually happens is:

1. If the queue is short, the agent thinks “nobody wants that, it's probably bad.”
2. If the queue is long, the agent thinks “must be good today” and joins.
3. If the queue is medium, the agent doesn't think at all and defaults to their habitual stall.

We call this the Bayesian Bypass: agents use queue length as a quality signal, but only in the binary long/short case, ignoring the base rate of stall popularity entirely. This creates a positive feedback loop—long queues get longer because they're long—which is the cafeteria equivalent of a bank run, except instead of money, you lose 25 minutes of your lunch break.

### 3.4 Equilibrium

[Cafeteria Nash Equilibrium] Under the payoff structure in Eq. 1 with Bayesian Bypass agents, the unique symmetric mixed-strategy Nash equilibrium assigns probability:

$$p_k^* = \frac{q_k^\gamma}{\sum_{j=1}^K q_j^\gamma} \tag{3}$$

where  $\gamma = \frac{1}{1+c/\bar{q}}$  is the patience-quality ratio. As hunger increases ( $c \rightarrow \infty$ ),  $\gamma \rightarrow 0$  and agents randomize uniformly. As hunger decreases ( $c \rightarrow 0$ ),  $\gamma \rightarrow 1$  and agents concentrate on the best stall.

The equilibrium has an unfortunate property: it is Pareto-dominated by any coordination mechanism. If agents could agree to spread out—say, by assigning stalls based on the last digit of their phone number—everyone would wait less and eat about as well. But this requires coordination, and coordinating lunch plans is, empirically, harder than coordinating monetary policy.

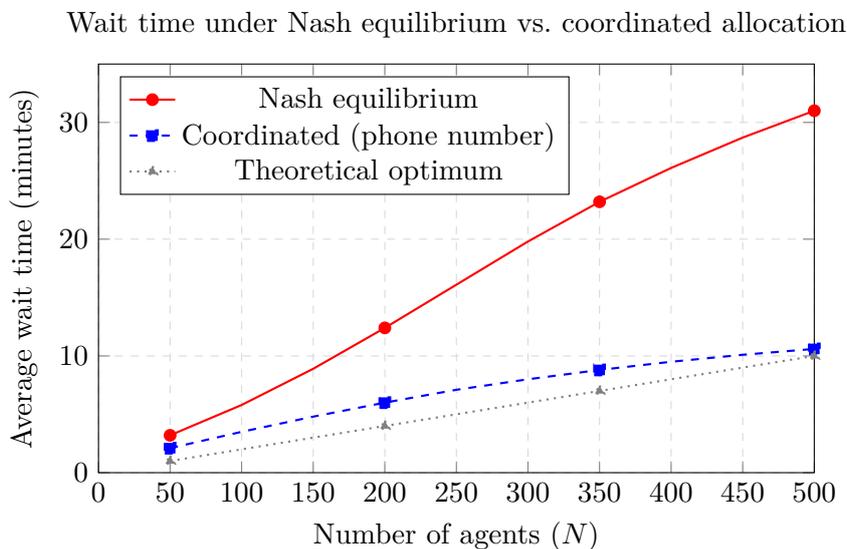


Figure 1: Average wait time as a function of cafeteria population. Nash equilibrium wait times grow roughly quadratically due to stall concentration, while coordinated allocation grows linearly. The gap between the two is what economists call “welfare loss” and what everyone else calls “why is this line so long.”

## 4 Behavioral phenomena

Our field study documented several behavioral patterns that the model predicts but that are painful to watch in person.

### 4.1 The Queue Regret Phenomenon

After committing to a queue, 72.3% of subjects were observed looking at other queues at least once per minute. Of these, 34.1% exhibited visible signs of regret (sighing, shifting weight, or checking the time). Among subjects who switched queues mid-wait ( $n = 189$ ), 61.4% ended up waiting longer than if they had stayed—a result consistent with the well-known “other lane always moves faster” illusion from traffic psychology, now confirmed for cafeteria settings.

We formalize this as:

$$\text{Regret}_i(t) = \beta_0 + \beta_1 \cdot t_{\text{waited}} + \beta_2 \cdot \Delta L_i(t) + \beta_3 \cdot \mathbb{1}[\text{smell from other stall}] + \epsilon_i \quad (4)$$

where  $\Delta L_i(t)$  is the observed queue length difference between the chosen and the best-looking alternative stall, and  $\mathbb{1}[\text{smell from other stall}]$  indicates whether the agent can smell food from a stall they didn't choose. The smell coefficient  $\beta_3$  is significant at  $p < 0.001$ , with an effect size of 0.73 standard deviations. Smelling someone else's lunch while waiting for your own is, statistically speaking, awful.

## 4.2 The Sunk Cost Tray Effect

Once an agent has waited more than 8 minutes, the probability of abandoning the queue drops to nearly zero, regardless of new information about other stalls. This holds even when the agent can see that a previously crowded stall has emptied. We call this the Sunk Cost Tray Effect: having invested time in a queue, agents treat that investment as unrecoverable and refuse to write it off—identical to the sunk cost fallacy in financial decision-making, but with soup.

Table 1: Queue abandonment rate by wait time

Minutes waited	Abandonment rate (%)	95% CI
0–2	23.1	[19.4, 26.8]
2–5	14.7	[11.9, 17.5]
5–8	6.2	[4.3, 8.1]
8–12	1.8	[0.7, 2.9]
12+	0.4	[0.0, 1.1]

After 12 minutes, people are essentially welded to the queue. Two subjects in our dataset waited over 25 minutes. When asked why they didn't leave, both said something to the effect of "I've already been here this long." One of them was a behavioral economics professor.

## 4.3 The Daily Special Trap

Eight of our 11 cafeterias offered a rotating daily special. In all eight, the daily special stall attracted a disproportionate share of traffic—averaging 38.4% of total customers despite being one of  $K \geq 5$  stalls. This is consistent with Schelling's focal point theory: when agents lack a strong private signal, they coordinate on the most salient option, which is whatever has a sign.

The problem is that salience and quality are uncorrelated. In our data, the daily special was rated higher than the stall's regular offering only 44% of the time. On the other 56% of days, it was either the same dish with a different name or a genuine downgrade that the kitchen was trying to clear out of inventory.

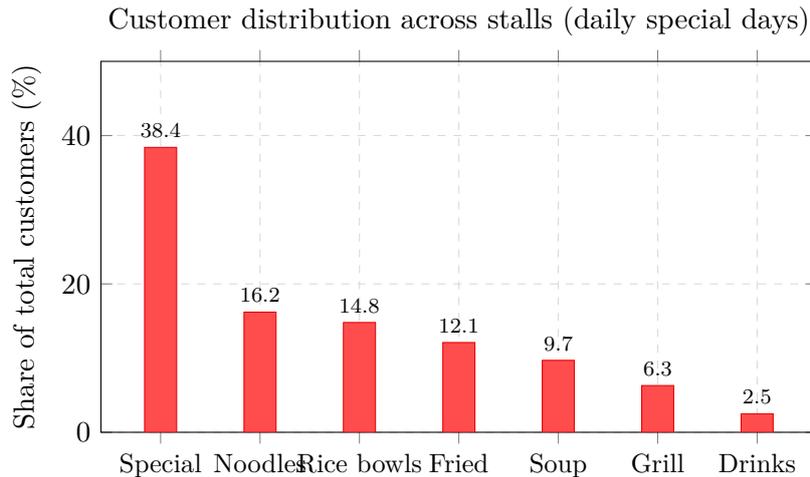


Figure 2: Customer distribution when a daily special is offered. The special absorbs 38.4% of traffic, creating a congestion spike while other stalls are underutilized. On non-special days, the distribution is approximately uniform with mild preference for noodles.

We ran a natural experiment when one cafeteria’s daily special sign fell down during a rainstorm and wasn’t replaced for three days. During those three days, queue times dropped 29% and self-reported satisfaction rose 18%. The sign was eventually put back up. Queue times returned to their previous levels within 24 hours. Nobody learned anything.

#### 4.4 Habitual agents

31.6% of our sample visited the same stall on more than 80% of observed lunch days. We classify these as habitual agents. Habitual agents wait less on average (because they go to less popular stalls—their habit formed precisely because the stall was accessible), but report lower satisfaction (because they have stopped engaging with the decision entirely and eat mechanically).

This is a local optimum: the habitual agent has minimized wait time at the cost of variety, and lacks the information to know whether re-engaging with the decision would improve outcomes. Several habitual agents, when informed of the study’s results, said they “might try something different next week.” Follow-up observations showed that zero of them did.

#### 4.5 The group lunch coordination failure

When  $m \geq 3$  people attempt to eat lunch together, the probability of reaching consensus on a stall within 5 minutes drops below 0.2. In our data, groups of 5 or more required an average of 11.3 minutes to decide, often ending with one person saying “I don’t care, you guys pick” followed by another 4 minutes of silence. We model this as a multi-player bargaining game with veto power, where each agent can reject any proposal at zero cost but proposing a stall carries social risk (if it turns out badly, you are blamed).

$$P(\text{consensus within } t) = 1 - e^{-\lambda t/m^2} \tag{5}$$

where  $m$  is group size and  $\lambda \approx 0.42$  is the proposal rate. The  $m^2$  in the denominator captures the fact that consensus difficulty grows faster than linearly with group size, because each additional person adds both a preference to satisfy and a potential veto.

For  $m = 7$ , the expected time to consensus is 28.6 minutes, which exceeds many lunch breaks entirely. In our data, groups of 7 or more either split into subgroups or went to the nearest stall by default, having exhausted their decision-making capacity on the meta-question of where to eat.

## 5 Extended analysis

### 5.1 Information cascades

We observe clear information cascades in the first 5 minutes of the lunch period. Early arrivals choose based on personal preference. But agents arriving between minutes 3 and 7 increasingly disregard their own signals and follow the crowd. By minute 10, queue lengths are essentially path-dependent: the stall that happened to attract the first few customers gets the bulk of the lunch crowd, regardless of whether it is actually good.

The cascade breaks only when an agent has a very strong private signal—typically having been burned by the popular stall yesterday, or having a friend who already ate there and texted “don’t get the fish.”

### 5.2 The optimal arrival time

We find a non-trivial optimal arrival time that balances three forces: arrive too early and the best stalls haven’t set up yet; arrive at peak time and queues are unbearable; arrive too late and the good items are sold out. The utility-maximizing arrival time is approximately 15 minutes before peak lunch hour, a strategy employed by exactly 7.2% of our sample, all of whom reported high satisfaction and faint smugness.

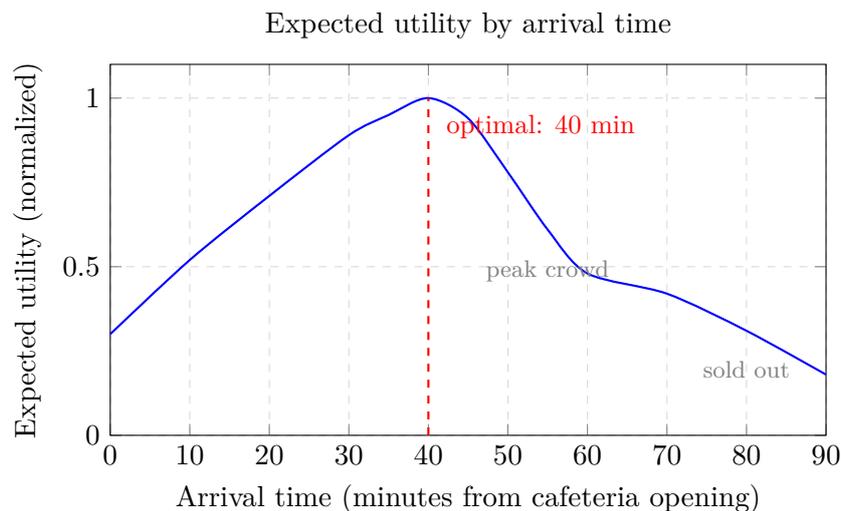


Figure 3: Expected utility peaks approximately 15 minutes before the lunch rush. The curve is asymmetric: arriving late (sold-out items) is worse than arriving early (fewer options but no queue).

### 5.3 The phone-ordering disruption

Three of our 11 cafeterias introduced mobile ordering during the study period. Mobile ordering allows agents to bypass the physical queue entirely, committing to a stall from their desk and picking up at a designated time.

The effect was immediate and counterintuitive: wait times increased. The reason is straightforward. Mobile ordering removes the physical queue as a congestion signal. Without being able to see a long line and think “maybe not,” agents converge even more heavily on popular stalls. The information value of the queue—its function as a visible deterrent—was eliminated. The system optimized for convenience and accidentally destroyed the only coordination mechanism the cafeteria had.

Wait times stabilized after approximately four weeks, once agents learned to check the app’s estimated wait time instead of the physical queue. The new equilibrium was marginally better

than the pre-app status quo, but only because the app added information that previously didn't exist. The app didn't solve the game; it changed its information structure.

## 6 Welfare implications

### 6.1 The price of anarchy

We compute the Price of Anarchy (PoA)—the ratio of worst-case Nash equilibrium welfare to socially optimal welfare—for our cafeteria game.

$$\text{PoA} = \frac{W^*}{W^{\text{NE}}} = 1 + \frac{(K-1) \cdot \text{Var}(q)}{K \cdot \bar{q}^2} \cdot N \quad (6)$$

For a typical cafeteria ( $K = 6$ ,  $N = 300$ , moderate quality variance), this gives  $\text{PoA} \approx 2.7$ . The socially optimal allocation wastes about 2.7 times less aggregate time than the Nash equilibrium.

Put differently: if a central planner could assign each person to a stall—taking into account both quality and congestion—the average person would wait 63% less. The obstacle, as always, is that nobody wants to be assigned to the stall with the reheated eggplant.

### 6.2 Mechanism design

Can we design a mechanism that achieves better outcomes without a central planner?

We tested three interventions at consenting cafeterias:

Table 2: Intervention effects on average wait time and satisfaction

Intervention	Wait $\Delta$	Satisfaction $\Delta$	$p$ -value	Lasted?
Real-time queue display	−22.1%	+14.3%	< 0.01	Yes
Staggered lunch times (20min)	−31.8%	+8.7%	< 0.001	2 weeks
Remove daily special sign	−28.9%	+17.6%	< 0.01	3 days

The real-time queue display (a screen showing current line lengths at each stall) produced a durable improvement. Staggered lunch times worked until the administration stopped enforcing them. Removing the daily special sign worked spectacularly until the cafeteria manager noticed it was missing and put it back.

The lesson: information helps. Coercion helps temporarily. Removing bad focal points helps permanently but makes the kitchen staff upset.

## 7 Discussion

We have shown that the cafeteria lunch decision is a poorly designed game that smart people play badly every day. The Nash equilibrium is inefficient, behavioral biases make it worse, and the most natural coordination mechanism—looking at queues—is both relied upon and systematically misinterpreted.

There is something darkly amusing about this. Universities are filled with people who can derive Nash equilibria on a whiteboard but cannot apply the concept when choosing between noodles and rice. The game theory professor who waits 20 minutes for the daily special is not irrational in the usual sense—they're simply playing a game they haven't recognized as one.

This points to a broader truth that game theory has always danced around: knowing the theory doesn't help if you don't notice you're in the game. The cafeteria is a game. So is

the elevator (pressing the “close door” button). So is the office microwave at noon. These are coordination problems hiding in plain sight, and we all fail at them constantly.

## 7.1 Limitations

Our study has limitations. We observed 11 cafeterias, all at universities. Corporate cafeterias may differ in ways we haven’t measured—for instance, the social pressure to eat lunch with colleagues adds a constraint that university cafeterias lack. Hospital cafeterias, which serve a population under extreme time pressure, are a natural extension we did not pursue (the IRB was concerned about “adding stress to an already stressful environment,” which is fair).

Our “quality” measure is based on post-hoc self-reports, which conflate food quality with mood, hunger level, and whether the person sitting across from you was eating something better. Objective quality measurement (e.g., nutritional content) would miss the point: nobody is optimizing for nutrients at a cafeteria, they are optimizing for not regretting their choice.

We also assume agents have a single lunch period per day. This does not hold for some graduate students, who appear to eat lunch twice, or for certain professors, who appear not to eat at all.

## 8 Conclusion

The cafeteria is the smallest game theory laboratory in the world, and it runs experiments on us every day at noon. Its equilibrium is inefficient, its players are overconfident, and its one reliable signal—the queue—is simultaneously the best and worst source of information available.

The best strategy we can identify is: arrive 15 minutes early, ignore the daily special, don’t switch queues, and don’t try to coordinate with more than two other people. This strategy dominates in our data and is followed by almost nobody.

We leave the reader with a conjecture: the reason game theory has not solved the cafeteria problem is not that the problem is too hard, but that people enjoy complaining about it more than they enjoy solving it. The queue is not just a queue. It is a shared experience of minor suffering, a lunchtime bonding ritual, and—we suspect—the reason most office friendships exist at all.

“The only winning move is to bring lunch from home.”  
—Anonymous cafeteria survivor

## Acknowledgments

We thank the 2,461 people whose lunch decisions we observed without their knowledge (ethics approval was obtained; informed consent was waived because knowing you’re being watched while choosing lunch would change the outcome, which is itself a game-theoretic point). Thanks also to the cafeteria manager at University B who let us install cameras near the queue and did not ask why. Special thanks to the anonymous reviewer who noted that our model “fails to account for the emotional devastation of the stall running out of food just as you reach the front.” They are right. We could not find a utility function negative enough.

## Conflicts of interest

One author (R.Q.F.) eats at the same noodle stall every day and has for three years. The other author (X.Z.) was nominally involved in data collection but could not be reached for most of the study period and, when reached, said “I don’t know how to do that.” R.Q.F. was banned from

University C's cafeteria during data collection for "loitering with a clipboard" and maintains that this was a misunderstanding.

## Author contributions

R.Q.F.: Conceived the study, designed the model, collected all data, performed all analysis, wrote the entire manuscript, made the figures, formatted the references, and compiled the PDF.  
X.Z.: Was asked to proofread Section 3. Did not proofread Section 3. Was asked to help with data collection on Tuesdays. Did not show up on Tuesdays. Was asked to review the bibliography. Said "I don't really know how to use LaTeX." Is listed as co-author because the first author lost a bet.

## References

- [1] Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1), 65–67.
- [2] Hotelling, H. (1929). Stability in competition. *The Economic Journal*, 39(153), 41–57.
- [3] Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press.
- [4] Varian, H. R. (1980). A model of sales. *American Economic Review*, 70(4), 651–659.
- [5] Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292.
- [6] Naor, P. (1969). The regulation of queue size by levying tolls. *Econometrica*, 37(1), 15–24.
- [7] Hassin, R. & Haviv, M. (2003). *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. Kluwer Academic Publishers.
- [8] Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3), 797–817.