# FairFoody: Bringing in Fairness in Food Delivery 

Anjali Gupta, Rahul Yadav, Ashish Nair, Abhijnan Chakraborty, Sayan Ranu, Amitabha Bagchi<br>Department of Computer Science and Engineering<br>Indian Institute of Technology Delhi<br>\{anjali,Rahul.Yadav.cs517,Ashish.R.Nair.cs517,abhijnan,sayanranu,bagchi\}@cse.iitd.ac.in


#### Abstract

Along with the rapid growth and rise to prominence of food delivery platforms, concerns have also risen about the terms of employment of the "gig workers" underpinning this growth. Our analysis on data derived from a real-world food delivery platform across three large cities from India show that there is significant inequality in the money delivery agents earn. In this paper, we formulate the problem of fair income distribution among agents while also ensuring timely food delivery. We establish that the problem is not only $N P$ hard but also inapproximable in polynomial time. We overcome this computational bottleneck through a novel matching algorithm called FAIRFOODY. Extensive experiments over real-world food delivery datasets show FAIRFOODY imparts up to 10 times improvement in equitable income distribution when compared to baseline strategies, while also ensuring minimal impact on customer experience.


## Introduction

Food delivery platforms like DoorDash, Zomato, GrubHub, Swiggy and Lieferando have become popular means for people to order food online and get delivery at their doorsteps, increasingly so due to the covid-19 pandemic related restrictions (Pengonda 2021). Typically, when a customer orders food from a particular restaurant a delivery agent is assigned to pick up the food from the restaurant once ready and deliver it to the customer. Thus, apart from providing business opportunities to the restaurants, food delivery platforms also provide livelihood to thousands of delivery agents. However, recent media reports have highlighted a range of issues faced by these delivery agents: poor working conditions, long working hours, non-transparent job allocations, and meagre pay (D'Souza 2021; Anab 2021; Murthy 2021). Due to the 'gig' nature of delivery jobs, a delivery agent typically gets a small fixed commission per order (except occasional tips and other incentives), and few employment benefits. A recent survey by the non-profit Fairwork (Fairwork 2021) found that despite the gig labels associated with this work (denoting part time engagements in addition to more stable jobs), most delivery agents in developing countries actually work full time on these platforms, depending on them entirely for their livelihood (Fairwork 2021). Fairwork also

[^0]found that none of the food delivery platforms in India guarantee local minimum wage to the delivery agents even if they work for more than 10 hours per day (Fairwork 2021).

Increasing the pay of delivery agents is a complicated proposition. If we decrease the number of agents so that the per-agent pay increases, there is a danger of increased customer wait time, which is anathema in the food delivery sector. Charging higher commission from restaurants or offering higher pay-per-delivery may disproportionately affect smaller restaurants and decrease their customer base. Recently such concerns led the city of Chicago to cap the delivery charge at $10 \%$ of the order value (City 2020). In some cities, restaurants are offering their own delivery services to cut back on the platform charges (Anand, Borah, and Majumdar 2021). In summary, it is not easy to increase the pool of money available to remunerate the delivery agents. However, apart from the issue of a small pool of money, the Fairwork report suggests that there is also high variability in pay - earnings on a platform can vary widely across agents (Fairwork 2021). To investigate this issue further, we perform an in-depth analysis of data obtained from a major food delivery platform for three Indian cities. Our analysis shows that there is significant inequality in the amount of money different agents earn from the platform. Interestingly, we see that the number of working hours of an agent, or her hours of operation cannot explain this inequality - it is rather the catchment area which makes a major difference.

In this paper, we present FAIRFOODY, the first algorithm to address the food delivery order allocation problem with fairness of pay distribution as a goal. We address the issue of catchment-based inequality by doing away with the restriction that a delivery agent must work within a single zone. Our algorithm, FAIRFOODY, does not, however, make any special effort to move agents across zones. In fact, in order to ensure timely delivery it limits the range from which a delivery agent can be picked to deliver a particular order. But, nonetheless, to ensure a more equitable pay distribution, FairFoody ends up creating a more uniform geographical distribution of agent activity.
The challenge we faced in designing FairFoody is that at any given time the number of orders may not be large enough to ensure that every idle agent is kept busy, and, hence, fairly remunerated. To deal with this issue we drew on Fairwork's finding that a number of delivery agents treat
food delivery as a full-time job to relax this temporal constraint: since the orders arrive throughout the day, we amortize fairness over a longer period of time than trying to be fair at each assignment. In our scheme if an agent does not get her fair share of assignments in a given time period, she may still make up for it on subsequent periods, and get a fair income over a longer term. With this relaxation FAIRFOODY is able to fairly remunerate those delivery agents who spend a significant period of time on the platform, and, consequently, rely on it to provide them with a living wage.

In summary, our key contributions are as follows:

- In-depth investigation of food-delivery data: We perform the first in-depth study of real food delivery data from large metropolitan cities and establish that the income distribution shows high levels of inequality.
- Problem formulation and algorithm design: We formulate the problem of fair income distribution in food delivery assignment, without compromising on customer experience. We show that the problem is not only $N P$-hard, but also inapproximable in polynomial time. To mitigate this computational bottleneck, we develop an algorithm called FAIRFOODY that uses bipartite matching on a data stream to perform real-time fair assignment of orders. ours is the first proposal to ensure fairness in food-delivery.
- Empirical evaluation: We evaluate FairFoody extensively on real food delivery data across a range of metrics, and establish that it is successful in its dual objective of fair income distribution and positive customer experience. A detailed version of our work containing proofs of important results, illustrative examples and additional experiments is available at (Gupta et al. 2022).


## Inequality in Delivery Agents' Income

## Dataset Used

We use six days of food delivery data from three large Indian cities, provided to us by a major food delivery service provider in India. Table 1 summarizes the dataset characteristics. The dataset consists of three components: trajectories of the delivery vehicles, the road network of each city (obtained from OpenStreetMap. org), and metadata describing various factors such as the vehicle IDs, information on each received order, locations of restaurants and customers, mean food preparation time in each restaurant, average speed in each road segment at different hours, etc. We match the vehicle GPS pings to the road network to obtain network-aligned trajectories (Newson and Krumm 2009).

## Inequality in Payment Distribution

We perform allocation using FoodMatch (Joshi et al. 2022) on the described dataset and compute the payment earned by different delivery agents over the course of these 6 days. Fig. 1(a) shows the Lorenz curve of their income, where the $y$-axis represents the cumulative percentage of total income and the $x$-axis represents the cumulative percentage of all agents. The diagonal line in Fig. 1(a) is the equality line; the further the income distribution from this line, the higher is the inequality. We see in Fig. 1(a) that there is high inequality in the income earned by the agents; the top $10 \%$

|  | City A | City B | City C |
| :--- | :---: | :---: | :---: |
| \# Restaurants | 2085 | 6777 | 8116 |
| \# Vehicles | 2454 | 159160 | 10608 |
| \# Orders | 23442 | 112745 | 112745 |
| Food prep. time (avg.in min) | 8.45 | 9.34 | 10.22 |
| \# Nodes | 39 k | 116 k | 183 k |
| \# Edges | 97 k | 299 k | 460 K |

Table 1: Summary of the dataset.
earners get $50 \%$ of the total payment while the bottom $60 \%$ only get $10 \%$ of it. Such a high inequality can starve many agents from getting decent income and would force them to quit the platform.

## What Drives Inequality?

Next, we try to uncover the sources of such high inequality. I. Number of working hours: It is reasonable to assume that pay variability springs from variability in the number of hours worked. However this is not the case in our data. We normalized payment by the number of active hours worked: Fig. 1(b) shows the Lorenz curve of the hourly income. We can see in Fig. 1(b) that the high inequality persists even after accounting for the activity levels.
II. Hours of operation: We may conjecture that top earners are active during lunch (11AM-2PM) or dinner (7PM11 PM ) times when order volumes are high, and the bottom earners are active during other periods of the day. Fig. 1(c,d,e) show the distribution of the operational periods (fraction of active times during lunch, dinner and all other time periods) for both the top $25 \%$ and bottom $25 \%$ earners. We can see that except City A, there is no noticeable difference in their activity patterns. It is not that the top earners were overwhelmingly more active during lunch and dinner times compared to the bottom earners. In fact, in City A, we can see that the bottom earners were more active during lunch times. Thus, difference in activity period is not a reasonable explanation for the inequality.
III. Geographical distribution: Next, we focus on the geographical spread of the order locations. Fig. 2 show heatmaps of restaurants' and customers' locations for the orders assigned to top earners and bottom earners, along with the location of the delivery agents when the orders were assigned to them. We can see a clear difference in the areas where top and bottom earners were active. For example, Fig. 2 shows that top earners deliver food mostly in the western part of City C, whereas the bottom earners are active more in the eastern part. It turns out that the order volume is higher in the western part, creating the inequality. We see similar trend in City A and City B as well.

Our finding here corroborates the delivery agents' experience as reported by Fairwork: agents often complained about not receiving orders in areas other than their chosen pickup zone (Fairwork 2021). Even when an agent delivers food outside their zone, they are not allocated any orders on their way back, thus incurring fuel costs on their return journey, without getting any payment (Fairwork 2021). We suggest that the delivery platforms can distribute opportunities more fairly among agents by allowing them to deliver orders in


Figure 1: (a-b) Lorenz curves for the distribution of agent incomes. (c-e) Distribution of operational periods for top earners and bottom earners.


Figure 2: Heat map of order locations: the customers' locations, restaurants' locations and locations of the assigned delivery agents. In (a-c) we show this location data with respect to the orders serviced by the top $25 \%$ agents and in (d-f) those serviced by bottom $25 \%$ agents based on income.
different parts of the city, as long as it does not negatively impact the waiting time for customers.

## Ensuring Fairness in Food Delivery

In this section, we define the concepts central to our work and formulate the problem of fair food delivery.

## Fairness Notions

Since food delivery platforms essentially distribute income opportunities among the delivery agents, the key question is: what would constitute a fair/just distribution? Fairness of distribution have been studied for a long time in Moral Philosophy, particularly in Distributive Justice (Lamont 2017). Next, we discuss few key principles from distributive justice and interpret them in the context of fair food delivery.
Strict Egalitarianism: The underlying idea behind this fairness principle is that people are morally equal, and hence
everyone should be treated equally (Arneson 2013). In food delivery context, this would mean that every delivery agent should earn the same income from the platform. To implement this in practice, the platform should pull together all delivery fees collected and then distribute them equally among the agents. However, such schemes are practically untenable; more so due to the gig nature of delivery jobs.
Difference Principle: In his seminal work on the theory of justice (Rawls 1971), John Rawls defined a system to be just if those affected by the system agree to be subjected to it. Rawls permit a departure from equality only if it provides 'greatest benefit to the least advantaged members of society'. In our context, this would translate into allocating the agent with the lowest income to a new order. However, this scheme would not consider the number of hours different agents work for the platform.
Proportional Equality: Ronald Dworkin opposed the idea of complete equality and argued for eliminating inequality that happens by sheer luck, but allowing the impact of people's choice or hard work (known as 'Luck Egalitarianism') (Arneson 2018). In the food delivery context, agents' incomes should be proportional to their effort, i.e., the number of hours they are working. In this work, we consider this notion of Proportional Equality, and propose to ensure that an agent's income is proportional to the number of hours they work for the platform.

## Background: The Food Delivery Problem

Next, we formulate the problem of food delivery without any fairness consideration (Joshi et al. 2021). In general, the objective is to allocate orders to delivery agents such that the waiting time for customers is minimized.

Definition 1 (Road Network). A road network is a directed, edge-weighted graph $G=(V, E, \beta)$, where $V$ is the set of nodes representing regions, $E=\{(u, v): u, v \in V\}$ is the set of directed edges representing road segments connecting regions, and $\beta:(E, t) \mapsto \mathbb{R}^{+}$maps each edge to a weight at time $t$. The edge weight at time $t$ denotes the expected time required to traverse the corresponding road at time $t$.

We use the notation $S P\left(u_{i}, u_{i+1}, t\right)$ to denote the length of the shortest (quickest) path from $u_{i}$ to $u_{i+1}$ at time $t$.

Definition 2 (Food Order). $A$ food order $o=\left\langle o^{r}, o^{c}, o^{t}, o^{p}\right\rangle$ is characterized by four features. $o^{r} \in V$ denotes the restaurant location (pick-up node), $o^{c} \in V$ is the customer location (drop-off node), $o^{t}$ is the time of request and $o^{p}$ is the (expected) food preparation time.

Let $v$ be a food delivery vehicle. We use $O_{t}^{v}$ to denote the orders assigned to $v$. Furthermore, $\operatorname{loc}(v, t)$ denotes the node that is closest to $v$ at time $t$. We assume all vehicles have a maximum carrying capacity of MaxO. Given a set of orders in $O_{t}^{v}$, a route plan is a permutation of $\left\{o_{i}^{r}, o_{i}^{c}: 1 \leq i \leq m\right\} \subset V$ such that for each $i, o_{i}^{r}$ appears before $o_{i}^{c}$ in the permutation. The length of a route plan $R P=\left\{u_{1}, \cdots, u_{m}\right\}$ is $\sum_{i=1}^{m-1} S P\left(u_{i}, u_{i+1}, t\right)$. The quickest route plan is, therefore, the one with the smallest length. We assume that all vehicles always follow the quickest route plan, and hence any reference to a route plan is implicitly assumed to be the quickest one.
Definition 3 (Order assignment). Given a set of orders $O$ and vehicles $\mathcal{V}$, an order assignment function $A: O \rightarrow \mathcal{V}$ assigns each order $o \in O$ to a vehicle $v \in \mathcal{V}$. An order o may be assigned to $v$ at time $t$ only if $\left|O_{t}^{v}\right|<$ MaxO.

Once order assignments are done, the first-mile distance, firstMile $(o, v)$, of order $o$ is the distance from $v=A(o)$ 's current location $\operatorname{loc}\left(v, o^{t}\right)$ to the pick-up location $o^{r}$ in the route followed by $v$. Similarly, lastMile $(o, v)$ is the lastmile distance from $o^{r}$ to drop-off location $o^{c}$.

Definition 4 (Expected Delivery Time). The expected delivery time of order $o$ when assigned to vehicle $v=A(o)$ :

$$
\begin{align*}
E D T(o, v) & =\max \left\{\operatorname{time}(A(o))+\text { firstMile }(o, v), o^{p}\right\} \\
& +\operatorname{lastMile}(o, v) \tag{1}
\end{align*}
$$

Here, $\operatorname{time}(A(o, v))$ denotes the computation time taken by the assignment algorithm. To explain $\operatorname{EDT}(o, v)$, the time to prepare food, and the time to assign a vehicle and reach the restaurant can progress in parallel. Thus, we take the maximum of these two components.
Problem 1 (The food delivery problem (FDP)). Given a set of orders $O$ and vehicles $\mathcal{V}$, if $\mathcal{A}$ is the set of all possible assignments of $O$ to $\mathcal{V}$, find the assignment $A$ that minimizes the average expected delivery time.

$$
\begin{equation*}
\arg \min _{A \in \mathcal{A}}\left\{\frac{1}{|O|} \sum_{\forall o \in O} E D T(o, A(o))\right\} \tag{2}
\end{equation*}
$$

At this juncture, we highlight two practical constraints. (1) In the real world, we work with a data stream of orders and vehicles instead of sets. The typical strategy to circumvent this issue is to accumulate orders over a time window $\Delta$ and assign this order set to available vehicles (Joshi et al. 2021; Reyes et al. 2018). This process is then repeated over the data stream.
(2) A food delivery service provider guarantees a Servicelevel agreement (SLA) of delivering the order within a stipulated time. This SLA is needed since food goes stale within a short time duration. Thus, it is desirable that the expected delivery time of all orders is within the SLA threshold $\Omega$.

## Problem Formulation: Fair Food Delivery

Problem 1 optimizes only the customer experience and does not incorporate the driver experience. As discussed earlier, a system would be fair if it equally distributes the timenormalized income among all delivery agents. We therefore formalize the notion of fairness for delivery vehicle.
Definition 5 (Time-Normalized Vehicle Income). Given any two time points $t_{1}<t_{2}$ and a vehicle $v$, let aT $\left(v, t_{1}, t_{2}\right)$ be the time that $v$ was available in time interval $\left[t_{1}, t_{2}\right]$, and let $d T\left(v, t_{1}, t_{2}\right)$ and $w T\left(v, t_{1}, t_{2}\right)$ be the total time spent driving and waiting at restaurant respectively by vehicle $v$ in $\left[t_{1}, t_{2}\right]$. Then, if a $T\left(v, t_{1}, t_{2}\right)>0$ we say that $v$ 's timenormalized income in $\left[t_{1}, t_{2}\right]$ is defined as:

$$
\begin{equation*}
\operatorname{inc}\left(v, t_{1}, t_{2}\right)=\frac{w_{1} \cdot d T\left(v, t_{1}, t_{2}\right)+w_{2} \cdot w T\left(v, t_{1}, t_{2}\right)}{a T\left(v, t_{1}, t_{2}\right)} \tag{3}
\end{equation*}
$$

where $w_{1}, w_{2}$ are payment parameters decided by the food delivery company.

Typically, $w_{1}>w_{2}$. Note that $\mathrm{aT}\left(v, t_{1}, t_{2}\right)-$ $\left(\mathrm{dT}\left(v, t_{1}, t_{2}\right)+\mathrm{wT}\left(v, t_{1}, t_{2}\right)\right)$ is the time during this interval when $v$ was either available, but had no orders assigned to it, i.e., it was idle.
Problem 2 (Fair Income Distribution in Food Delivery). Given a set of orders $O$ beginning at time 0 and ending at time $T_{m}$ and a set of available vehicles $\mathcal{V}$, if $\mathcal{A}$ is the set of all possible assignments of $O$ to $\mathcal{V}$, give an algorithm to find an assignment $A$ to minimize the income gap.

$$
\begin{equation*}
\arg \min _{A \in \mathcal{A}}\left\{\max _{v \in \mathcal{V}}\left\{\operatorname{inc}\left(v, 0, T_{m}\right)\right\}-\min _{v \in \mathcal{V}}\left\{\operatorname{inc}\left(v, 0, T_{m}\right)\right\}\right\} \tag{4}
\end{equation*}
$$

Optimizing Problem 2 may lead to an increase in delivery time, and hence affect the customer experience. We therefore aim to minimize Problem 2 under a bounded increase in delivery time. Formally,
Problem 3 (The Fair Food Delivery Problem). Find an assignment A minimizing Problem 2 under the constraint:

$$
\begin{equation*}
\operatorname{Dist}(A(o), o) \leq \Gamma \times \text { nearDist }_{o}, \forall o \in O \tag{5}
\end{equation*}
$$

where $\operatorname{Dist}(o, A(o)$ is the road network distance (shortest path) of order o from its assigned vehicle $v=A(o)$, nearDist ${ }_{o}$ is the distance to the nearest vehicle from order $o$ and $\Gamma>1$ is a threshold.

Theorem 1. There is no PTIME algorithm that can approximate Prob. 2 within any constant factor $c$, where $0<c \leq 1$, unless $P=N P$.

Proof. Provided in (Gupta et al. 2022).
Corollary 1. There exists no PTIME algorithm that can approximate Prob. 3 within any constant factor $c$, where $0<c \leq 1$, unless $P=N P$.

Proof. Any instance of Prob. 2 reduces to an instance of Prob. 3 for $\Gamma>$ diameter $\times$ min_distance, where diameter indicates the longest shortest path in the road network and min_distance is the minimum distance between any two points in the network.

## Our Proposal: Fairfoody

We propose a heuristic algorithm FAIRFOODY to solve the fair food delivery problem. It builds a weighted bipartite graph with vehicles in one partition and clusters of orders in the other. The weights of the edges are computed such that finding a minimum weight matching (Kuhn 1955; Munkres 1957) in this bipartite graph optimizes the criterion of Prob. 2, while also ensuring a good solution with respect to Prob. 1. We now elaborate on the key steps.
Definition 6 (Shortest Delivery Time (Joshi et al. 2021)). The shortest delivery time for an order o is $S D T(o)=o^{p}+$ $S P\left(o^{r}, o^{c}, o^{t}\right)$.
$S D T(o)$ is a natural lower bound on EDT (Eq. 1).
Definition 7 (Augmented Order Delivery Time). Consider a vehicle $v$ at time $t$. Suppose at this time we add a cluster of orders $O$ to $v$ 's route plan and $v$ is able to deliver all orders in this augmented route plan by time $t^{\prime}>t$, on the assumption that no new orders are added to $v$ 's route plan in time interval $\left(t, t^{\prime}\right)$. Then we define the Augmented Order Delivery Time $A O D T(O, v, t)$ as $t^{\prime}-t$.
Definition 8 (Augmented Order Payment). Consider a vehicle $v$ at time $t$ which may be currently idle or currently assigned some order. Suppose at this time we add a cluster of orders $O$ to $v$ 's route plan (which may be empty if the vehicle is idle) and $v$ is able to deliver all orders in this augmented route plan at time $t^{\prime}>t$ such that the total driving time in the interval $\left[t, t^{\prime}\right]$ is $t_{1}$ and the total waiting time is $t_{2}$. Then, on the assumption that no new orders are added to $v$ 's route plan in time interval $\left[t, t^{\prime}\right]$, we define the Augmented Order Payment $A O P(O, v, t)$ as $w_{1} \cdot t_{1}+w_{2} \cdot t_{2}$.

As discussed in our problem formulation (Prob 1), we partition the data stream of orders into windows of length $\Delta$, and allocate all orders that arrived within this window. This process is then repeated for each of the subsequent windows.
Definition 9 (Next-slot Normalized Income). At time $\ell \Delta$, i.e. at the beginning of the $\ell+1$ st window, consider a vehicle $v$ that is available. In addition, consider an unassigned order $o$. We define the next-slot normalized income of $v$ if it is assigned o as
$n s-i n c(v, o, \ell)=\frac{\operatorname{inc}(v, \ell \Delta) \cdot a T(0, \ell \Delta)+A O P(o, v, \ell \Delta)}{\ell \Delta+A O D T(o, v, \ell \Delta)}$.
Creating a weighted bipartite graph: At time $\ell \Delta$, let us assume that $\mathcal{V}_{\ell}$ is the set of available vehicles and $O_{\ell}$ is the set of unallocated orders. We create a weight bipartite graph ( $\left.U_{1}, U_{2}, E\right)$ where $U_{1}=\mathcal{V}_{\ell}$, i.e., one side of the partition is the set of available vehicles and $U_{2}$ is a cluster of orders.
Clustering orders: If $\left|O_{\ell}\right| \leq f \cdot\left|\mathcal{V}_{\ell}\right|$ where $f$ is a parameter chosen in $(0,1)$, we set $U_{2}=O_{\ell}$. Otherwise we perform a clustering on $O_{\ell}$ using Ward's method, i.e., we successively coalesce those two clusters whose being delivered by a single vehicle leads to the least increase in the extra delivery time. The cluster size is not allowed to cross MaxO. We stop either when the number of clusters falls to $f \cdot\left|\mathcal{V}_{\ell}\right|$ or when the increase in expected delivery time due to clustering crosses a threshold $\eta$. We denote the final clusters as $O_{\ell}^{\prime}$.

Edge weights in bipartite graph: Each edge of the bipartite graph is of the form $(v, O)$ where $O \subset O_{\ell}$ is either a single order (a singleton set) or a cluster of orders. If vehicle $v$ is at a distance not exceeding $\Gamma \times$ near $^{2}$ ist $_{O}$ (Recall Prob. 3), we set the weight as follows:

$$
\begin{equation*}
w(v, O)=n s-i n c(v, O, \ell)-\min _{v \in \mathcal{V}} \operatorname{inc}(v, \ell \Delta) \tag{7}
\end{equation*}
$$

To identify all vehicles whose distance not exceeding $\Gamma \times$ nearDisto in an efficient manner, we perform best-first search on the road network graph (See Alg. 1 in (Gupta et al. 2022)). Specifically, we start from the restaurant locations of each order in $O$ and visit all the nearby vehicles in best first search order till distance from source restaurant exceeds $\Gamma \times$ near $D i s t_{O}$. All vehicles beyond this boundary are assigned edge weight $\approx \infty$.

Finally, we run Kuhn-Munkres on the bipartite graph to obtain the allocation for the $\ell+1$ st window.
Theorem 2. The time complexity of the proposed algorithm is $\mathcal{O}(m \cdot n(q+\max (m, n)))$, where $m=\left|\mathcal{V}_{\ell}\right|, n=\left|O_{\ell}^{\prime}\right|$, $\mathcal{O}(q)$ is the time taken for shortest path computation.

Proof. Provided in (Gupta et al. 2022).

## Experimental Evaluation

In this section, we benchmark FAIRFOODY and establish:

- Fairness: FAIRFOODY imparts more than $10 X$ improvement in fairness over baselines.
- Cost of fairness: FAIRFOODY maintains a comparable delivery time as that of FoodMatch (Joshi et al. 2021). Code is available at https://github.com/idea-iitd/fairfoody.


## Evaluation Framework

Baselines: We consider (1) FoodMatch, and (2) Twosided Fairness (2SF) (Sühr et al. 2019). The code for both algorithms are obtained from the authors.
Metrics: The performance is quantified through:

- Gini coefficient: Gini coefficient is the ratio of the area that lies between the line of equality and the Lorenz curve over the total area under line of equality (Gastwirth 1972).
- DTPO: DTPO measures the average delivery time per order. DTPO allows us to quantify the cost of fairness.
- Percentage of SLA violations (SLA-V): We measure the percentage of orders not delivered within the promised time limit (Prob. 1). The time limit is set to 45 minutes.
- Spatial distribution distance $(\psi)$ : This metric quantifies the total variation distance (Levin, Peres, and Wilmer 2017) between the heatmaps of top- $25 \%$ and bottom- $25 \%$ drivers (Figs. 2 and 3). See (Gupta et al. 2022) for details.
Parameters: The default size of accumulation window $\Delta$ is 3 minutes. Clustering parameter $f$, and payment weights $w_{1}$ and $w_{2}$ (Def. 5) are set to $0.8,1.0$ and 0.8 respectively as their default values.


## Comparison with FOODMATCH and 2SF

Table 2 presents the performance of various algorithms across all three cities. In Fig. 1(b), we observed that FoodMATCH induces significant income disparity among drivers. This is reflected in the high Gini of Foodmatch across all


Figure 3: Heat map of order locations after applying FAIRFoody. In (a-c) we show this location data with respect to the orders serviced by the top $25 \%$ agents and in (d-f) those serviced by bottom $25 \%$ agents based on income.
cities in Table 2. In contrast, FAIRFOODY reduces Gini more than 10 times across three cities. This reduction in Gini, however, does not come at the cost delivery time or SLA violations (SLA-V). Specifically, there is minimal increase in DTPO and SLA-V. Similar to Gini, FaIrFoody is also significantly better in Income Gap. Overall, this shows that it is possible to ensure fairness without compromising on the customer experience.

We also compare with 2 SF at $\lambda=1$ and $\lambda=0$. We choose these two $\lambda$ values since they represent the two extremes; at $\lambda=1,2$ SF optimizes only the driver income gap, whereas $\lambda=0$ minimizes only the delivery time. Thus, $\lambda=1$ represents the best possible Gini by 2 SF . We observe that even in this scenario, FAIRFOODY is 9 times better on average. At $\lambda=0$, although 2 SF achieves low delivery times, it fails to

| City | Algorithm | Gini | Income <br> Gap | DTPO | SLA-V <br> $(\boldsymbol{\%})$ |
| :--- | :--- | :--- | :--- | :--- | :--- |
| $\mathbf{A}$ | FAIRFOODY | $\mathbf{0 . 0 3 5}$ | $\mathbf{2 0 . 5}$ | 15.4 | 0.33 |
|  | 2SF, $\lambda=1$ | 0.32 | 58 | 17.3 | 0.37 |
|  | 2SF, $\lambda=0$ | 0.526 | 55.9 | $\mathbf{1 5 . 2}$ | 0.33 |
|  | FoODMATCH | 0.518 | 59.2 | $\mathbf{1 5 . 2}$ | $\mathbf{0 . 3 2}$ |
| $\mathbf{B}$ | FAIRFOODY | $\mathbf{0 . 0 4 7}$ | $\mathbf{3 0 . 4}$ | 15.5 | $\mathbf{0 . 2 2}$ |
|  | 2SF $\lambda=1$ | 0.316 | 59.9 | 15.5 | 23.12 |
|  | 2SF, $\lambda=0$ | 0.471 | 59.4 | $\mathbf{1 4 . 5}$ | 19.62 |
|  | FOODMATCH | 0.512 | 59.3 | 15.4 | 0.23 |
| $\mathbf{C}$ | FAIRFOODY | $\mathbf{0 . 0 3 5}$ | $\mathbf{3 3 . 5}$ | 16.2 | 0.33 |
|  | 2SF, $\lambda=1$ | 0.323 | 58.3 | 16.9 | 12.05 |
|  | 2SF, $\lambda=0$ | 0.513 | 56.1 | $\mathbf{1 5 . 8}$ | 6.63 |
|  | FOODMATCH | 0.562 | 59.1 | 16.0 | $\mathbf{0 . 3 2}$ |

Table 2: Performance of FairFoody, FoodMatch, and 2 SF across various metrics. The unit of DTPO is in minutes. The best performance in each metric is highlighted in bold.

| City | Algorithm | 25 \% | 50 \% | 75 \% | 95\%0 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| A | FAIRFOODY | 38 | 51.5 | 70 | 88 |
|  | Foodmatch | 13 | 29 | 65 | 205.3 |
|  | $2 \mathrm{SF} \lambda=1$ | 20 | 44 | 80 | 142 |
| B | FAIRFOODY | 31 | 61 | 92 | 115 |
|  | FOODMATCH | 12 | 33 | 75 | 264 |
|  | $2 \mathrm{SF} \lambda=1$ | 23 | 53 | 90 | 156 |
| C | FAIRFOODY | 44 | 64 | 79 | 95 |
|  | Foodmatch | 13 | 31 | 74 | 250 |
|  | $2 \mathrm{SF} \lambda=1$ | 28 | 52 | 87 | 154 |

Table 3: Number of orders across all six days per driver at various percentiles $(\% 0)$.
satisfy SLA across a large portion of orders.
In Table 3, we further examine the number of orders delivered by drivers across various percentiles based on normalized income (Def. 5). It is clear from the data that FAIRFOODY achieves the most equitable distribution.

## Impact on Spatial Distribution

In Fig. 2, we showed that spatial distribution is a key driver of payment inequality. Now we will show that FAIRFOODY equalizes the spatial distributions across the pay range. Fig. 3 studies this question in City C. As visible, the heatmaps of the top $-25 \%$ and bottom $-25 \%$ are much more similar when compared to Fig. 2. This indicates that disparity in spatial distribution is correlated to income disparity.

To quantify this observation and examine whether the pattern holds across all cities, we compute $\psi_{P}$ across all combinations of cities and properties (See (Gupta et al. 2022) for the exact formula). Table 4 presents the results. Two key observations emerge from this experiments. First, both for FAIRFOODY and $2 \mathrm{SF}, \lambda=1$, the $\psi$ is lower across all cities and properties. This indicates that spatial distribution distance and Gini (as well as Income Gap) are indeed correlated. However, just minimizing spatial distribution distance is not enough in minimizing Gini. Specifically, although $2 \mathrm{SF}, \lambda=1$ has a lower spatial distance than FairFoody in City B , its Gini is significantly higher. (Table 2).

## Impact of Parameters

Clustering parameter $(f)$ : Lowering of $f$ leads to more clustering among orders. Fig. 4 analyzes variation in different performance metrics against $f$. Both Gini coeffi-

| City | Locations <br> of the | FOOD- <br> MATCH | 2SF <br> $\lambda=0$ | 2SF <br> $\lambda=1$ | FAIR- <br> FOODY |
| :--- | :--- | :--- | :--- | :--- | :--- |
| $\mathbf{A}$ | Vehicle | 0.543 | 0.500 | 0.293 | $\mathbf{0 . 2 3 1}$ |
|  | Customer | 0.490 | 0.437 | 0.268 | $\mathbf{0 . 2 3 5}$ |
|  | Restaurant | 0.469 | 0.442 | $\mathbf{0 . 1 6 9}$ | 0.278 |
| $\mathbf{B}$ | Vehicle | 0.394 | 0.394 | $\mathbf{0 . 1 6 9}$ | 0.191 |
|  | Customer | 0.348 | 0.312 | $\mathbf{0 . 1 5 2}$ | 0.182 |
|  | Restaurant | 0.354 | 0.342 | $\mathbf{0 . 1 3 5}$ | 0.174 |
| $\mathbf{C}$ | Vehicle | 0.386 | 0.403 | 0.243 | $\mathbf{0 . 1 9 3}$ |
|  | Customer | 0.318 | 0.310 | 0.217 | $\mathbf{0 . 1 8 5}$ |
|  | Restaurant | 0.318 | 0.382 | $\mathbf{0 . 2 0 6}$ | 0.239 |

Table 4: Comparison of spatial distribution $\psi_{P}$.


Figure 4: Impact of clustering parameter (f). Gini in Fig. 4(a) is computed on time-normalized income.


Figure 5: Impact of number of vehicles on various metrics. Gini in Fig. 5(a) is computed on time-normalized income.
cient (Fig. 4(a)) and Average delivery time per order (Avg. DTPO) (Fig. 4(c)) increase as we decrease $f$. Higher clustering prevents equitable distribution of orders (and, therefore, income). Hence, Gini increases. DTPO increases since larger clusters lead to less simultaneous delivery of orders through multiple drivers. SLA violations remain unaffected (Fig. 4(d)). There is no consistent trend across three cities in income gap (Fig. 4(b)).
Impact of number of vehicles: Instead of considering all vehicles that were available in our real dataset, we randomly sample a subset of $X$ vehicles. Next, we vary $X$ in the $x$ axis and observe the impact of various metrics in $y$-axis. Fig. 5 presents the results. We observe that Gini increases with increase in vehicles (Fig. 5(a)). This is expected since decreasing the number of vehicles generates higher demand and more scope to fairly distribute orders. It is also natural that higher vehicle available leads to lower delivery time (Fig. 5(c)) and less SLA-violations (Fig. 5(d)).
Further experiments: Please see (Gupta et al. 2022) for more experiments on this work.

## Related Work

Food order assignment: On the problem of food-delivery, FoodMatch (Joshi et al. 2021) is the only work to provide a realistic and scalable solution in food delivery domain. Other works on food delivery suffer from various unrealistic assumptions such as perfect information about arrival of orders (Yildiz and Savelsbergh 2019), ignoring the road network (Reyes et al. 2018), and ignoring food preparation time (Zeng, Tong, and Chen 2019).
Fairness in multi-sided platform algorithms: With the growing popularity of multi-sided platforms, a number of recent works have investigated the challenges of unfairness and bias in such platforms (Patro et al. 2020a; Dash et al.
2021). For example, (Edelman, Luca, and Svirsky 2017) looked into the likelihood of racial bias in Airbnb hosts' acceptance of guests, while (Lambrecht and Tucker 2016) looked at gender discrimination in job advertisements. Few works have also looked at how producers and customers treat each other as a group. (Chakraborty et al. 2017) and (Sühr et al. 2019) proposed strategies for two-sided fairness in matching situations, whereas (Burke 2017) categorised distinct types of multi-stakeholder platforms and their required group fairness qualities. Individual fairness for both producers and customers is addressed by (Patro et al. 2020b) in tailored suggestions in two-sided platforms. Despite these works on fairness in two-sided platforms, there has not been any studies on food delivery platforms. It is also worth noting that, as discussed in (Joshi et al. 2022), allocation algorithms for the cab service industry (Garg and Ranu 2018; Yuen et al. 2019; Ma, Zheng, and Wolfson 2013; Cheng, Xin, and Chen 2017) is not a natural fit food delivery.

## Conclusion

In this work, we focused on the unfairness issues faced by delivery agents in food delivery platforms. Using data from a large real-world food delivery platform, we showed that there exists high inequality in the income earned by the agents. To counter such inequality, we proposed an algorithm FAIRFOODY to assign delivery agents to orders ensuring that income opportunities are fairly distributed among the agents. Extensive experiments show that FairFoody outperforms state-of-the-art baselines in lowering inequality while ensuring minimal increase in delivery time. Given the increasing adoptions of food delivery platforms, ensuring fairness is the need of the hour and we hope that our work would lead to more followup works in this space.

## References

Anab, M. 2021. Delivery staffers battle odds, health risks to be on job. https://bit.ly/3jXa3Gi.
Anand, S. N.; Borah, P. M.; and Majumdar, M. 2021. The great Indian food delivery tussle. https://www.thehindu.com/life-and-style/food/indian-restaurants-go-for-direct-delivery/article34770229.ece.
Arneson, R. 2013. Egalitarianism. In Zalta, E. N., ed., Stanford Encyclopedia of Philosophy. Stanford University.
Arneson, R. J. 2018. Dworkin and Luck Egalitarianism. In The oxford handbook of distributive justice.
Burke, R. 2017. Multisided Fairness for Recommendation. ArXiv:1707.00093 [cs.CY].
Chakraborty, A.; Hannák, A.; Biega, A. J.; and Gummadi, K. 2017. Fair Sharing for Sharing Economy Platforms. In FATREC 2017.
Cheng, P.; Xin, H.; and Chen, L. 2017. Utility-Aware Ridesharing on Road Networks. In SIGMOD, 1197-1210.
City, C. 2020. Notice of Fee Caps on Third-party Food Delivery Services. https://bit.ly/38U25aM.
Dash, A.; Chakraborty, A.; Ghosh, S.; Mukherjee, A.; and Gummadi, K. P. 2021. When the umpire is also a player: Bias in private label product recommendations on e-commerce marketplaces. In ACM FAccT, 873-884.
D'Souza, P. M. 2021. Rising fuel prices, less earnings: Double whammy for food delivery agents amid lockdown. https://bit.ly/3zX0mNv.
Edelman, B.; Luca, M.; and Svirsky, D. 2017. Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment. American Economic Journal: Applied Economics, 9: 1-22.
Fairwork. 2021. Fairwork India Ratings 2020: Labour Standards in the Platform Economy. https://bit.ly/3tsDIua.
Garg, N.; and Ranu, S. 2018. Route recommendations for idle taxi drivers: Find me the shortest route to a customer! In $A C M K D D, 1425-1434$.
Gastwirth, J. L. 1972. The estimation of the Lorenz curve and Gini index. The review of economics and statistics, 306316.

Gupta, A.; Yadav, R.; Nair, A.; Chakraborty, A.; Ranu, S.; and Bagchi, A. 2022. FairFoody: Bringing in Fairness in Food Delivery. arXiv preprint arXiv:2203.08849.
Joshi, M.; Singh, A.; Ranu, S.; Bagchi, A.; Karia, P.; and Kala, P. 2021. Batching and Matching for Food Delivery in Dynamic Road Networks. In Proc. ICDE.
Joshi, M.; Singh, A.; Ranu, S.; Bagchi, A.; Karia, P.; and Kala, P. 2022. FoodMatch: Batching and Matching for Food Delivery in Dynamic Road Networks. ACM Transactions on Spatial Algorithms and Systems (TSAS), 8(1): 1-25.
Kuhn, H. W. 1955. The Hungarian method for the assignment problem. Naval research logistics quarterly, 2(1-2): 83-97.
Lambrecht, A.; and Tucker, C. 2016. Algorithmic Bias? An Empirical Study into Apparent Gender-Based Discrimination in the Display of STEM Career Ads. SSRN Electronic Journal.

Lamont, J. 2017. Distributive justice. Routledge.
Levin, D. A.; Peres, Y.; and Wilmer, E. L. 2017. Markov chains and mixing times. American Mathematical Soc.
Ma, S.; Zheng, Y.; and Wolfson, O. 2013. T-share: A largescale dynamic taxi ridesharing service. In ICDE, 410-421.
Munkres, J. 1957. Algorithms for the assignment and transportation problems. Journal of the society for industrial and applied mathematics, 5(1): 32-38.
Murthy, D. 2021. The 'customer is king' motto has invisibilised the food delivery rider. https://www.thenewsminute.com/article/customer-king-motto-has-invisibilised-food-delivery-rider-153591.
Newson, P.; and Krumm, J. 2009. Hidden Markov Map Matching Through Noise and Sparseness. In ACM SIGSPATIAL GIS, 336-343.
Patro, G. K.; Biswas, A.; Ganguly, N.; Gummadi, K. P.; and Chakraborty, A. 2020a. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In Proceedings of The Web Conference 2020, 1194-1204.
Patro, G. K.; Biswas, A.; Ganguly, N.; Gummadi, K. P.; and Chakraborty, A. 2020b. FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms. In $W W W$. ACM.
Pengonda, P. 2021. Why covid-19 has been a shot in the arm for online food delivery firms. https://www.livemint.com/market/mark-to-market/why-covid-19-has-been-a-shot-in-the-arm-for-online-food-delivery-firms-11611558506486.html.
Rawls, J. 1971. A theory of justice. Harvard university press. Reyes, D.; Erera, A. L.; Savelsbergh, M. W. P.; Sahasrabudhe, S.; and O'Neil, R. J. 2018. The Meal Delivery Routing Problem. Optimization Online.
Sühr, T.; Biega, A. J.; Zehlike, M.; Gummadi, K. P.; and Chakraborty, A. 2019. Two-sided fairness for repeated matchings in two-sided markets: A case study of a ridehailing platform. In $A C M K D D$.
Yildiz, B.; and Savelsbergh, M. 2019. Provably HighQuality Solutions for the Meal Delivery Routing Problem. Transportation Science, 53.
Yuen, C. F.; Singh, A. P.; Goyal, S.; Ranu, S.; and Bagchi, A. 2019. Beyond Shortest Paths: Route Recommendations for Ride-sharing. In Proc. WWW, 2258-2269.
Zeng, Y.; Tong, Y.; and Chen, L. 2019. Last-Mile Delivery Made Practical: An Efficient Route Planning Framework with Theoretical Guarantees. Proc. VLDB Endow., 13(3): 320-333.


[^0]:    Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

