

# Formation of Social Ties Influences Food Choice: A Campus-Wide Longitudinal Study

KRISTINA GLIGORIĆ, EPFL, Lausanne, Switzerland

RYEN W. WHITE, Microsoft Research, Redmond, Washington, United States

EMRE KICIMAN, Microsoft Research, Redmond, Washington, United States

ERIC HORVITZ, Microsoft Research, Redmond, Washington, United States

ARNAUD CHIOLERO\*, Population Health Laboratory, University of Fribourg, Fribourg, Switzerland

ROBERT WEST, EPFL, Lausanne, Switzerland

Nutrition is a key determinant of long-term health, and social influence has long been theorized to be a key determinant of nutrition. It has been difficult to quantify the postulated role of social influence on nutrition using traditional methods such as surveys, due to the typically small scale and short duration of studies. To overcome these limitations, we leverage a novel source of data: logs of 38 million food purchases made over an 8-year period on the École Polytechnique Fédérale de Lausanne (EPFL) university campus, linked to anonymized individuals via the smartcards used to make on-campus purchases. In a longitudinal observational study, we ask: How is a person's food choice affected by eating with someone else whose own food choice is healthy vs. unhealthy? To estimate causal effects from the passively observed log data, we control confounds in a matched quasi-experimental design: we identify focal users who at first do not have any regular eating partners but then start eating with a fixed partner regularly, and we match focal users into comparison pairs such that paired users are nearly identical with respect to covariates measured before acquiring the partner, where the two focal users' new eating partners diverge in the healthiness of their respective food choice. A difference-in-differences analysis of the paired data yields clear evidence of social influence: focal users acquiring a healthy-eating partner change their habits significantly more toward healthy foods than focal users acquiring an unhealthy-eating partner. We further identify foods whose purchase frequency is impacted significantly by the eating partner's healthiness of food choice. Beyond the main results, the work demonstrates the utility of passively sensed food purchase logs for deriving insights, with the potential of informing the design of public health interventions and food offerings, especially on university campuses.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**; • **Applied computing** → **Health informatics**; • **Information systems** → **Data mining**; **Social networks**; • **Mathematics of computing** → **Causal networks**.

Additional Key Words and Phrases: nutrition; behavioral analysis; causal analysis; university; public health

\*Also affiliated with Institute of Primary Health Care (BIHAM), University of Bern, Bern, Switzerland, and School of Population and Global Health, McGill University, Montreal, Canada.

Authors' addresses: Kristina Gligorić, EPFL, Lausanne, Switzerland, [kristina.gligoric@epfl.ch](mailto:kristina.gligoric@epfl.ch); Ryen W. White, Microsoft Research, Redmond, Washington, United States, [ryenw@microsoft.com](mailto:ryenw@microsoft.com); Emre Kiciman, Microsoft Research, Redmond, Washington, United States, [emrek@microsoft.com](mailto:emrek@microsoft.com); Eric Horvitz, Microsoft Research, Redmond, Washington, United States, [horvitz@microsoft.com](mailto:horvitz@microsoft.com); Arnaud Chiolero, Population Health Laboratory, University of Fribourg, Fribourg, Switzerland, [achiolero@gmail.com](mailto:achiolero@gmail.com); Robert West, EPFL, Lausanne, Switzerland, [robert.west@epfl.ch](mailto:robert.west@epfl.ch).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2021/4-ART184 \$15.00

<https://doi.org/10.1145/3449297>

**ACM Reference Format:**

Kristina Gligorić, Ryen W. White, Emre Kiciman, Eric Horvitz, Arnaud Chiolero, and Robert West. 2021. Formation of Social Ties Influences Food Choice: A Campus-Wide Longitudinal Study. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 184 (April 2021), 25 pages. <https://doi.org/10.1145/3449297>

**1 INTRODUCTION**

Nutrition plays a key role in people's daily lives and is a major determinant of long-term health [27, 33]. Interventions and policies that promote healthier diets are therefore important public health priorities [87]. In designing such interventions, situational food norms, including social influences, play a prominent role, as they are theorized to have a powerful effect on food intake [18, 38, 39, 52, 76].

Despite the postulated importance of social factors, measuring how dietary behaviors are affected by the behaviors of others remains challenging. On the one hand, experimental studies to date have been limited to observing people in small-scale scenarios with a short duration [62, 63]. On the other hand, observational studies have relied on survey-based methods [16], employing questionnaires [88] and personal food journals [19, 20], which are costly to organize and prone to biases [11].

Furthermore, making causal inferences regarding the influence of social ties on food intake faces the challenge of numerous confounding factors. Although similarities in diet and eating behaviors among persons connected via social links (e.g., friends, family and peers) have been observed in a number of experimental and survey-based studies [30, 36, 49, 58, 71, 77], it is not clear whether the similarity in the consumption patterns arises from social influence, or if confounding factors, such as self-selection in tie formation (homophily) and environmental influences, can explain the similarity [4, 45, 73, 74]. In real-world settings, it remains challenging to measure and disentangle properties that are relevant in the context of food consumption, such as attributes of the individuals and of the environment (e.g., food options available in different locations and settings). Researchers have only recently been addressing this gap by studying social media and other digital traces of human behavior in the context of food consumption [1, 14, 57].

In order to shed new light on the influence of social factors on food choice, we harness a novel data source: logs of 38 million food purchases made over an eight-year period on a major university campus, linked to anonymized individuals (students and staff) via the smartcards they use to make on-campus purchases. The large scale and long duration of the data enables studies with greater statistical power, compared to prior setups, and allows for reducing the influence of confounding factors, and thus for identifying the causal effect of social influence, by carefully selecting a suitable subset of individuals whose food choice behaviors are monitored throughout the individuals' long-term transaction histories.

Based on this dataset, we design a longitudinal observational study to address the question of how a person's food choice is affected by eating with someone else whose own food choice is healthy vs. unhealthy. To estimate causal effects from the passively observed log data, we control confounds in a matched quasi-experimental design. where we identify focal users who at first do not have any regular eating partners but then start eating with a fixed partner regularly, and we match focal users into comparison pairs such that paired focal users are nearly identical with respect to covariates measured before acquiring the eating partner, but the two focal users' new eating partners diverge in the healthiness of their respective food choice.

**Research questions.** Specifically, we seek to answer the following questions:

- (1) How is the *overall healthiness* of a focal person's food choice affected by the healthiness of an eating partner's food choice? Does the focal person's food choice change, and if so, in what direction?

- (2) How is a focal person's choice of *specific food categories* affected by the healthiness of an eating partner's food choice? Does the distribution over food categories change, and if so, what items are purchased more, and what items less?

**Summary of main findings.** Regarding question 1, we observe that, when a focal person acquires a new eating partner, the healthiness of the focal user's food choice shifts significantly in the direction of their new eating partner's dietary patterns. In a difference-in-differences analysis of 415 comparison pairs of focal persons (identified among a total of around 39,000 persons in eight years' worth of log data), which carefully controls for a number of confounding covariates, we find clear evidence of social influence: focal persons acquiring a healthy-eating partner change their habits significantly more in the direction of healthy foods than focal persons acquiring an unhealthy-eating partner. We quantify the robustness of this finding in a sensitivity analysis, and we provide further evidence by observing a dose-response relationship between the difference in exposures and the difference in effects.

Regarding question 2, we observe that focal persons who start eating with healthy-eating partners show an increase in the purchase of coffee and lunch meals, items generally purchased in large numbers, with the strongest effect. On the other hand, items purchased at higher rates by the matched counterparts, who start eating with unhealthy-eating partners, loosely form a cluster of potentially unhealthy items that should not be eaten in large quantities (soft drinks, drinks from vending machines, condiments, pizza, kebabs, and crêpes).

**Implications.** Students and staff consume large amounts of food on campuses, daily and globally. The present work shows the value of employing novel methods relying on population-scale digital traces to measure social influence on food choice behaviors in this context. The derived insights have the potential to support interventions aimed at encouraging more healthy and sustainable dietary habits in university environments and beyond.

## 2 RELATED WORK

### 2.1 Social influence and diet

Social influence on dietary habits is an active area of research [39, 76]. Food consumption has been found to be influenced by eating with others [38], and the food choices of others, including people one does not know, have been observed to influence food choices, even when not consciously recognized [17, 62]. Particular attention has been given to understanding the governing psychological mechanisms, including the seeking of dish uniformity driven by the goal of regret minimization, or the seeking of dish variety driven by self-presentation [5, 26, 53].

Although the underlying mechanisms are not fully understood, uniformity seeking is observed across a range of studies. For example, it is observed that the quantity dimension is used to communicate gender identity, and the food-type dimension to ingratiate the co-eater's preferences by matching the other's presumed choice, following gender-based stereotypes about food [13]. Such social norms, including the influence of peers, have tremendous potential for understanding dietary patterns and designing public health interventions [18, 52, 61, 63].

In our work, we monitor behaviors outside of experimental setups. While previous efforts in this area have focused on specific behaviors (e.g., buying a dessert or not), having access to a multi-year history of all transactions made on a large campus allows us to observe behavioral changes for longer time periods and in a more fine-grained way, by measuring a wide set of purchasing behaviors that occur in the real world.

## 2.2 The special case of children and adolescents

A large fraction of the transactions recorded in our logs were made by students, i.e., adolescents and young adults. Focusing on similar age groups, social influence in dietary habits has been examined in the context of school children [9, 30, 58, 71] and adolescents [24, 25, 77], who are theorized to be most susceptible to social pressures. In particular, effects of peer influence have been observed in children and adolescents' diets as well as activity patterns [6, 70].

Systematic reviews of social network analyses of young people's eating behaviors and body weight reveal consistent evidence that school friends are significantly similar in terms of their body mass index. Friends with the highest body mass index appear to be most similar [32]. Prior work further reveals that the family context is essential when implementing healthy eating interventions, as parents, not friends, are the most prominent influencers of adolescents' healthy eating [43, 59].

## 2.3 Contagiousness of unhealthy behavior

Previous work has particularly been focused on unhealthy behaviors and their contagious effects, observing that obesity [16], overeating [50], fast food [81], high-fat [29, 37], and alcohol and snack consumption [56, 88] are contagious. In fact, the strongest evidence of social influence in food choices has been found for unhealthy behaviors (e.g., snack foods) [10, 22]. Beyond food consumption, peer influence and social norms have been observed to play a role in unhealthy weight-control behaviors among adolescent girls: self-induced vomiting, laxatives, diet pills, and fasting were all shown to be contagious among adolescent girls [28]. A rich literature exists on tackling the problem of unhealthy behaviors through interventions with the goal of promoting healthy dietary habits and physical activity [31], losing weight [41], reducing the risk of chronic illnesses [34], and reducing food waste [60].

There is a heated debate about whether unhealthy behaviors are indeed contagious, or whether the observed similarities should instead be attributed to homophily, i.e., people's tendency to form ties with others who are similar to oneself to begin with. Disentangling social influence from homophily poses a fundamental challenge. Without strong assumptions about the structure of ties or the ability to measure confounding factors, homophily and contagion are generically confounded (i.e., the effect of social influence cannot be identified) [4, 73, 74].

Our work attempts to minimize the effect of confounding variables in previously infeasible ways. Based on the rich transaction data, we measure a set of relevant confounding variables and carefully control for them in our quasi-experimental setup.

## 2.4 Nutrition monitoring and modeling based on digital traces

Social media has emerged as a promising source of data for studies on monitoring food consumption. For instance, it has been shown that Twitter has tremendous potential to provide insights into food choices at a population scale [1]. Researchers have also studied specific dietary issues and behaviors: reports of eating disorders [14, 57], dietary choices, and nutritional challenges in food deserts, i.e., places with poor access to healthy and affordable food [23]. Another active area of research has been focused on improving methods for monitoring food consumption, relying on mobile phones [19, 20] and wearable devices, to recognize when the eating activities take place [80].

Recent related research has also demonstrated the value of monitoring and modeling of nutrition using other kinds of large-scale digital traces [35], such as grocery store purchase logs [3, 12], online recipes [64], logging-based smartphone applications and wearables [2, 8], reviewing platforms [15], search engine logs [83, 85], social media such as Twitter [1, 51, 86] or Instagram [55, 75], crowdsourcing platforms [40], and geo-location signals [69].

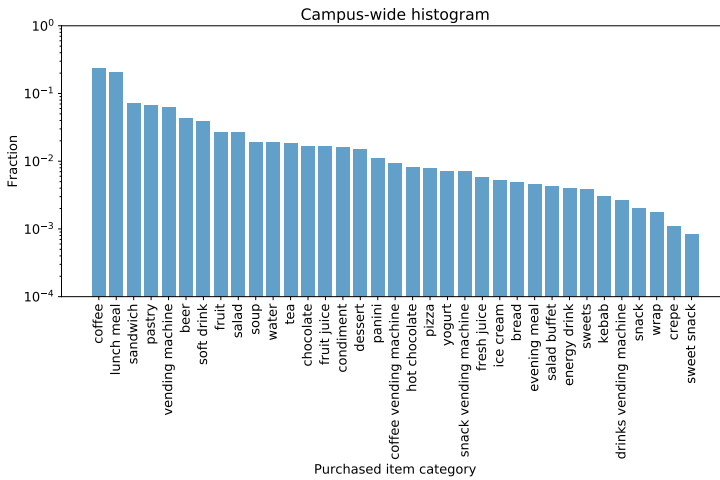


Fig. 1. Distribution of purchases across item categories, across a campus. Most purchased items are coffee and lunch meals.

Finally, large-scale passively sensed signals have been harnessed in university campus environments to measure factors of well-being outside of nutrition [7, 49, 54, 72, 78]. Recent preliminary insights point towards the feasibility and the potential of automatically inferring social interactions from behavioral traces for campus-centric applications [78].

To summarize, while large-scale digital traces are promising for monitoring and modeling nutrition, little is known about how these passively sensed behavioral signals can be used for understanding the factors that govern food consumption in campus settings. Our longitudinal study aims to bridge this gap by analyzing large-scale, long-term purchase data.

### 3 MATERIALS AND METHODS

#### 3.1 Transaction log data

This work leverages an anonymized dataset of food purchases made on the École Polytechnique Fédérale de Lausanne (EPFL) university campus. The data spans 8 years, from 2010 to 2018, and contains about 38 million transactions, of which about 18 million were made with a badge that allows linking to an anonymized person’s ID. The data includes 38.7k users, who, on median, are observed for a time period spanning 578 days and make 188 transactions. Each transaction is labeled with the time it took place, information about the sales location (shop, restaurant, vending machine, or café), the cash register where the transaction took place, and the purchase items. Items are associated with unstructured textual descriptions (e.g., “coffee”, “croissant”, “Coca-Cola can”). The unstructured textual descriptions were additionally manually mapped to categorical labels (such as “meal”, “drink”, or “dessert”) by a research assistant, who labeled the 500 most frequently purchased items, which account for 95.4% of the total volume of item purchases observed in the dataset. The distribution of purchases across categories is shown in Figure 1.

Purchases are not evenly spread over the course of the year, but, as expected, are higher during semesters, and lower during the breaks between semesters (Figure 2, left).

This work also leverages a smaller-size enriched transactional dataset gathered during a campus-wide sustainability challenge, for which 1,031 consenting participants formed 278 teams in order to compete in taking sustainable actions (e.g., taking the stairs instead of the elevator, or consuming a

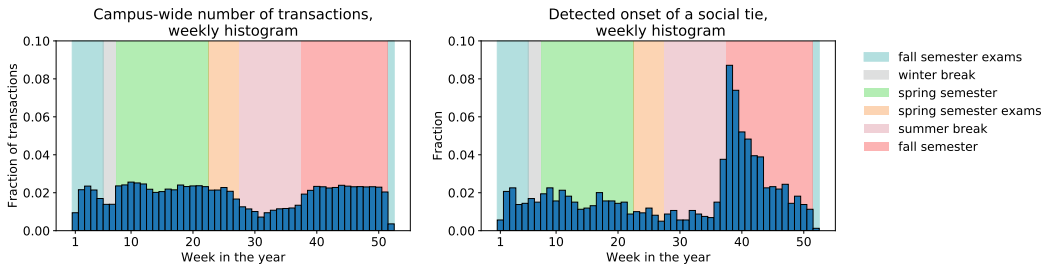


Fig. 2. **Left:** Annual distribution of food purchases. The trends mirror the university schedule: the number of transactions drops at the end of the spring semester (around week 25), and increases again at the start of the fall semester (around week 40). A similar pattern is observed before the beginning of the spring semester (around week 10). **Right:** Annual distribution of detected onsets of social ties. The ties emerge disproportionately often when classes start at the beginning of the fall semester (by a factor of 3.5 times, compared to a baseline sampled at random from the distribution of purchases).

vegetarian meal). This data was not used for our analyses, but only for assessing the accuracy of our heuristic method for inferring frequent eating peers (described next).

### 3.2 Inference of co-eating onset from proximity in transaction logs

To measure the effect of the emergence of new social ties, we first infer frequently co-eating persons based on the proximity in the transaction logs. Frequently co-eating persons are likely to share a social tie, i.e., they are persons likely to be friends, colleagues, or classmates who often eat together. Previous work has shown that such social ties can be reliably inferred from geospatial proximity [21].

To infer frequent eating peers, we monitor a sequence of transactions made on the same day with the badge in the queue of a fixed cash registry, in a given shop. We identify situations when two individuals are adjacent in the queue and make a transaction within one minute between each other, with no one in between them. We use a lower threshold of 10 such high-confidence proximity indicators to infer a likely social tie. The first appearance of proximity in the logs is then considered to be the onset of co-eating. We observe a spike in tie formation coinciding with the start of classes in the fall (Figure 2, right).

Furthermore, we evaluate the precision of our heuristic by comparing the inferred co-eaters with ground-truth team membership information from the sustainability challenge. We observe that team membership in the sustainability challenge, a ground-truth indicator of a social tie, is correlated with sharing an inferred tie based on the transaction logs: out of all the pairs of individuals from the sub-population taking part in the sustainability challenge who are detected as frequent eating partners, 72% are also members of the same team.

### 3.3 Inference of nutritional properties from raw transaction logs

We infer a set of summary nutritional properties from raw transaction logs by relying on a set of pre-established criteria. We derive healthiness labels based on food-pyramid recommendations [84]. Products that should be consumed in the least amounts possible, i.e., items at the top of the Swiss food pyramid (with high amounts of saturated fats, salt, added sugars, refined grains, and highly processed foods) were considered as “unhealthy” (e.g., sodas, chips, candies, and chocolate bars). Other products that are not at the top of the Swiss food pyramid are considered to be “healthy”

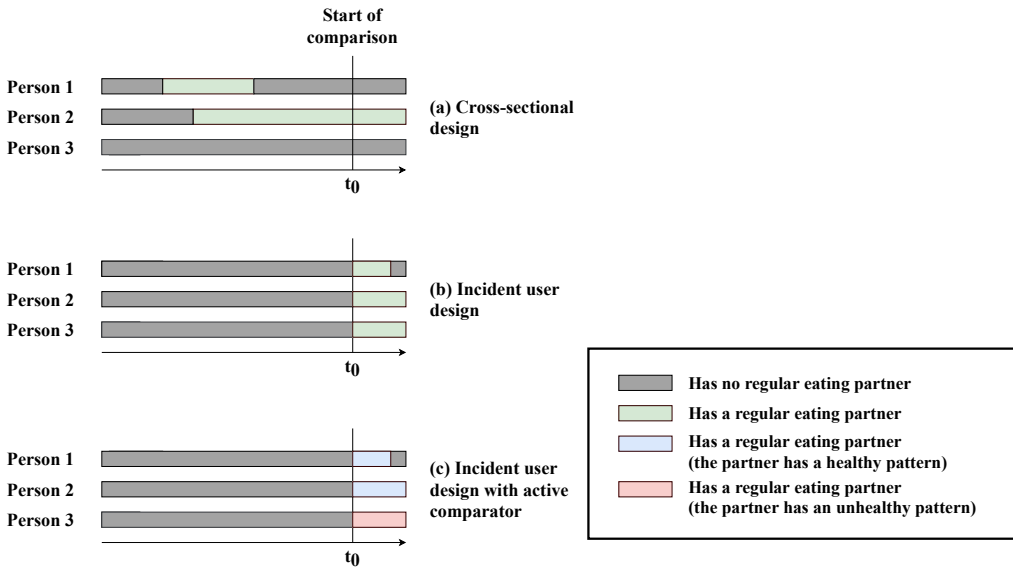


Fig. 3. Study design diagrams. We illustrate three potential observational study designs to estimate the effect of eating with other persons on food choices, (a) *cross-sectional design*, (b) *incident user design*, and (c) *incident user design with active comparators*. At different points in time, a person either does not have a regular eating partner (marked in gray), or she does (marked in green, red, or blue). A cross-sectional design observes food consumption at the start of the monitored period, a fixed time  $t_0$ , which is the same across all participants. Incident user design isolates the effect of the onset of co-eating with another person on subsequent food consumption. In incident user design, time is tracked relative to the moment of onset  $t_0$ , which may be different across participants. The active comparator design additionally allows for comparisons of the effect of onset among persons who all start to eat with someone, but their partners have different characteristics (marked in red and blue). The present paper is based on an incident user design with active comparators (presented in more detail in Figure 4).

(including non-sweetened beverages, fruits, vegetables, whole grains, meat, fish, and nuts). When insufficient information was available from the name of the product, “unclassifiable” was selected.

Two professional epidemiologists specialized in nutrition independently assessed each food item and categorized them into healthy vs. unhealthy vs. unclassifiable. The reviewers had access to the unstructured textual description of the item (e.g., “coffee”, “croissant”, “Coca-Cola can”). The reviewers did not have access to any other meta-information about the items. Disagreements were resolved by a third reviewer. Labels are used to create a healthiness score of a set of purchases by averaging individual product scores, coded numerically as 1 for healthy (25% of items), -1 for unhealthy (46% of items), and 0 for unclassifiable (29% of items).

### 3.4 Matched incident user design with active comparators

Recall that we are interested in determining whether and how eating with others impacts the nature of food consumption. As depicted in Figure 3(a), a naïve approach to answering those questions would be a *cross-sectional design*: at any given absolute point in time, some people are regularly eating with their peers (indicated with green) while others do not (indicated with gray). Starting from a certain absolute point in time  $t_0$ , by identifying persons with different habits, one could compare what is consumed by the persons who do not have a regular eating partner with what is

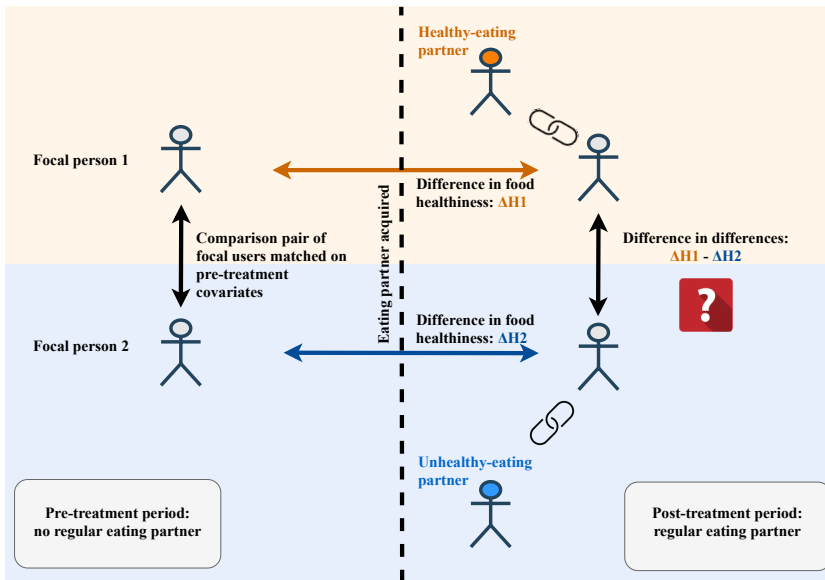


Fig. 4. The matched incident user design with active comparators on which the present study is based. We identify comparison pairs of focal persons 1 and 2, who are indistinguishable in the pre-treatment period and have no regular eating partners, until the moment of co-eating onset, when they each acquire a regular eating partner. Focal person 1 starts regularly eating with a healthy-eating partner, while focal person 2 starts regularly eating with an unhealthy-eating partner. The comparison pair of focal users is then observed in the pre-treatment period (no regular eating partner) and post-treatment period (regular eating partner). The effect of the co-eating onset is estimated using a difference-in-differences analysis.

consumed by the persons eating with a regular eating partner. One could also compare the food consumed by persons who are eating regularly with partners who have different habits.

The problem with this setup is that those persons who do not eat with others might have done so in the past (e.g., Person 1 in Figure 3(a)). Those who do eat with others might have been doing it for a long time or might have just initiated. Also, some people stop eating with others, whereas other people continue. It could be that those who stop do so because they prefer the diet they seek when eating alone (i.e., selection bias). Additionally, people who eat with others might differ in fundamental ways from those who do not.

For these two reasons, looking at everyone at the same moment in a cross-sectional way can be problematic. To overcome these challenges, we can turn to an *incident user design* (Figure 3(b)), which restricts the population to those people who newly initiate the treatment—starting to eat together with another person. We are interested in the causal effect on food consumption of initiating eating with a peer. Among people who had no regular eating partners in the past, what is the causal effect of starting to eat with a peer? In this way we isolate the causal effect of initiation. We restrict the observed population so that none of the persons have a history of eating with someone. Note how Person 1 in Figure 3(b) starts eating together with a regular partner, but then after a while no longer has a regular eating partner. This is not an issue because we are interested in the effect of the onset.

As opposed to the cross-sectional design, where time is absolute, the incident user design offers the flexibility of tracking time relative to an onset  $t_0$  that may be different for different participants. Although this design allows us to compare different treatments, the problem with this setup, which






Fig. 5. Daily fraction of purchases annotated as potentially unhealthy, tracked over five years. A seasonal pattern emerges. Drops in the daily fraction of unhealthy purchases coincide with between-semester breaks.

persists from the above-described cross-sectional design, is that, if the comparison group is “no treatment” (i.e., no initiation of co-eating), it is not apparent when the follow-up should start for the “no treatment” group. Additionally, selection bias remains and is not accounted for, as people who do not initiate might in other fundamental ways differ from those who do initiate.

Our study design addresses these challenges by implementing a variant of incident user design, *incident user design with active comparators* (Figure 3(c)). Here, before initiation, no user included in the study had a regular eating partner (i.e., was treatment-free). We compare the effect of initiating to eat with partners who have different habits among persons who all initiate to eat with someone (illustrated with blue and red in Figure 3(c)). Active comparator designs tend to involve significantly less confounding [42, 48, 90], as people who eat with different kinds of others are more alike among themselves than when compared to people who do not have regular eating partners.

Our study design is illustrated in more detail in Figure 4. We identify persons (referred to as *focal persons*) who had no regular eating partners and, at a moment  $t_0$  specific to that focal person, initiate eating with someone (referred to as *eating partner*). Here, as defined in Section 3.2, a person qualifies as a focal person’s potential eating partner if the two were observed making subsequent purchases in the same queue within one minute of one another on at least 10 occasions in the entire dataset, and the onset of co-eating is defined as the first one of these occasions. We then isolate pre-treatment and post-treatment periods of the focal person’s food purchases comprising all transactions made six months before the first purchase together (moment  $t_0$ ) and six months after, respectively. We ensure that the focal person does not initiate eating with anyone else in the pre- and post-treatment six months. The length of the pre-treatment period is chosen so that it is feasible to expect that an individual will be present on campus given the typical stay in the logs (the total observed 12 months of pre- and post-treatment correspond to one school year).

Some persons initiate co-eating with a person who has a positive healthiness score in the aligned pre-treatment period. In contrast, some initiate co-eating with a partner who has a negative score. These are the two groups that we seek to compare (we refer to the two types of partners as *healthy-eating partner* and *unhealthy-eating partner*).

For a focal person who starts to eat together with a partner who has a healthy dietary pattern, an *active comparator* (or *counterpart*), will be another focal person who starts to eat together with a partner who has an unhealthy dietary pattern. The potential counterparts start to eat with their partner in the same month as the other counterpart. This is done in order to control for temporal confounds that might arise from a seasonal variation of food popularity: as seen in Figure 5,

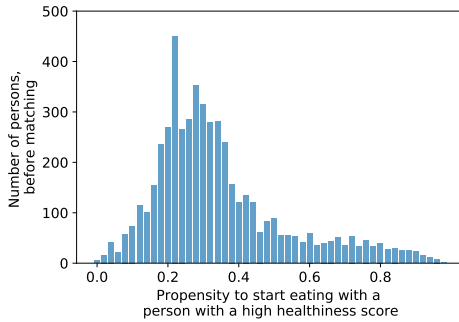


Fig. 6. Distribution (before matching) of propensity to start eating with a healthy-eating partner.




Fig. 7. Importance of most indicative features for predicting treatment assignment, i.e., initiation of eating with a healthy- vs. unhealthy-eating partner (shop names anonymized). Most important feature: pre-treatment healthiness score of focal person's purchases, which indicates homophily.

unhealthy foods are especially popular at certain times of the year. The healthiness of the partner's dietary pattern is determined according to its numeric value (greater or less than zero), and not relative to the focal person.

Comparing incident users with active comparators is an important step towards reducing the impact of biases. However, in the assignment of the type of treatment, there can still be confounding. For example, it might be the case that only people who already have healthy habits start eating together with a partner who has healthy habits, due to a preference for similar others. The influence of the partners would then be indistinguishable from the impact of selection biases caused by homophily.

Hence, we turn to a *matched* incident user design with active comparators. We introduce an improvement over the previously discussed setup, where the incident users are matched to the potential active comparators while additionally controlling for pre-treatment covariates. Our goal here is to balance potential confounding variables within pairs, to be able to observe how the onset of co-eating with partners with different dieting patterns is associated with subsequent changes in the focal person's dieting pattern. We achieve this by performing a propensity-score-based causal analysis. We approximate randomized treatment assignment by modeling the propensity to experience the assigned intervention, relying on a number of pre-treatment covariates describing the focal persons' eating profiles. Due to the balancing property of propensity scores [66], matching on propensities results in similar covariate distributions between groups that differ in their assigned interventions.

The covariates capture important dimensions of the pre-treatment dietary pattern of the focal person: *where* the food is purchased (what is the shop where the person most frequently buys food), *when* the food is purchased (what is the fraction of items occurring during lunchtime), *what* types of items are purchased (what fraction of purchased items are meals, and what is their estimated healthiness), and *how often* the person purchases food on campus (number of transactions). We measure these confounding covariates up to time  $t_0$ .

We use a random forest model that predicts the type of treatment based on pre-treatment covariates of the focal person (area under the ROC curve: 0.87). This implies that past purchases allow us to accurately predict whether the tie will be formed with a healthy- or an unhealthy-eating partner, and that confounding is a real problem that needs to be addressed. The distribution of the propensity to start eating with a partner who has a high healthiness score is presented in Figure 6. We also examine the feature importances in predicting the treatment assigned, i.e., the initiation of eating with a partner who has a healthy or unhealthy eating pattern (Figure 7). We observe that

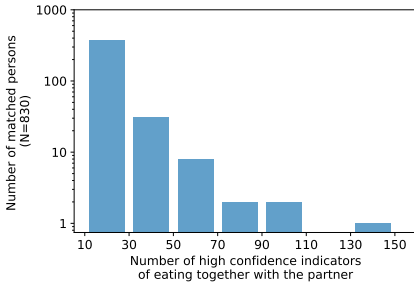


Fig. 8. Histogram of the number of high confidence indicators of eating together with their respective partners, for matched focal persons.




Fig. 9. Histogram of the partner's pre-treatment healthiness score, across matched persons. Orange bars correspond to healthy-eating partners, blue bars to unhealthy-eating partners. A margin of 0.1 is ensured to differentiate the treatments.

the focal person's pre-treatment healthiness score is in fact the most important predictor of the type of partner the focal person will start to eat with, pointing at homophily.

Focal persons in the two sets are then matched while ensuring that two potential matches have propensity scores (likelihoods of receiving the treatment) within a caliper of 0.1. The size of the caliper was chosen so that balance in covariates is achieved. Moreover, an exact match on the sign of the mean pre-treatment healthiness score and the most frequented shop is required to achieve tight control. We then create matched pairs based on possible candidates by performing maximum weight matching on the weighted bipartite graph, where nodes are focal persons, and the weights use similarity based on the Mahalanobis distance in covariates. We maximize the total similarity to find a maximal matching.

The result is a set of matched pairs of focal persons, indistinguishable up to the moment of initiation, who initiated co-eating with partners with different dietary patterns in the same month. This approach yielded 415 matched pairs of 830 focal persons who started to eat with different partners. We require at least 10 high confidence indicators of eating together with the partner (Figure 8). Partners' distribution of pre-treatment healthiness scores is shown in Figure 9.

Our matched analysis then moves on to comparing focal people who initiate co-eating with a person with a healthy dieting pattern, to their counterparts who have the same dieting patterns up to the moment of initiation, but initiate co-eating with a partner who has an unhealthy dieting pattern. The post-treatment patterns are then compared across treatments within the matched population.

Table 1. To ensure that matched persons are comparable, we evaluate the balance of their pre-treatment covariates, via the standardized mean difference (SMD) across covariates in the two matched groups.

Pre-treatment covariate	SMD before matching	SMD after matching
Preferred shop (i.e., where the largest fraction of pre-treatment transactions is made)	exact match required	
Pre-treatment percentage of lunchtime transactions	0.109	0.045
Pre-treatment percentage of meal transactions	0.207	0.075
Pre-treatment mean healthiness score	0.301	0.023
Pre-treatment mean weekly number of transactions	0.071	0.023




Fig. 10. The effect of the focal person's pre-treatment covariates, the partner's pre-treatment healthiness score, and the number of detected high-confidence indicators of eating together on the focal person's post-treatment healthiness of purchased items. The effects are estimated with linear regression ( $R^2 = 0.194$ ); 95% confidence intervals are approximated as two standard errors. Significant coefficients ( $p < 0.05$ ) are marked with an asterisk (\*). The focal person's own healthiness score and their eating partner's healthiness score are the only two statistically significant factors associated with the focal person's post-treatment healthiness score.

Before moving on to the analysis of the outcomes, we ensure that the matched persons are comparable by measuring the balance of their pre-treatment covariates (Table 1). We use the standardized mean difference (SMD) across covariates in the two groups to measure the balance. We observe that matching greatly reduces the SMD, as the largest SMD across covariates (the one of the pre-treatment healthiness score of the focal person) changes from 0.301 before matching, to 0.023 after matching. Groups are considered balanced if all covariates have SMD lower than 0.2 [44], a criterion that is satisfied here.

## 4 RESULTS

Recall our research question: we want to understand how a person's food choices are affected by the healthiness of a co-eating partner's food choices. Do people's choices change and, if so, in what direction (i.e., towards more or less healthy)?

### 4.1 Regression analysis of pooled data

First, we aim to determine if there are any significant differences between the outcomes of the matched focal persons. Is the pre-treatment healthiness of the partner predictive of post-treatment healthiness of the focal person? We start by performing a regression estimation of the effect of the partner's pre-treatment healthiness score on the focal person's post-treatment score.

We fit a model where the focal person's post-treatment healthiness score is the dependent variable. We include the focal person's pre-treatment healthiness score, the partner's pre-treatment healthiness score, the number of high-confidence indicators of eating together, and the focal person's pre-treatment covariates as the independent variables. The focal person's pre-treatment covariates (number of transactions, percentage of transactions that are meals, percentage of lunchtime transactions, and the pre-treatment healthiness score) are already controlled for by matching, but they are included in the model to account for possible residual confounding. The predictors and the outcome are standardized, so the coefficients are interpreted as increases in the healthiness score per standard deviation of the predictor.

In Figure 10, fitting the linear regression, we measure a significant positive effect of 0.13 (95% CI [0.07, 0.19]) of the partner's pre-treatment healthiness score. The focal person's own pre-treatment healthiness is the strongest predictor of post-treatment healthiness (coefficient 0.43, 95% CI [0.36, 0.49]). This is the first indicator that the pre-treatment score of the partner is associated with the focal person's patterns.

Table 2. Contingency table counting number of pairs of matched focal persons in each condition. Post-treatment healthiness score is compared to pre-treatment score to determine if there was an increase; in columns, for focal persons who start to eat with healthy-eating partners, and in rows, for their matched counterparts, i.e., focal persons who start to eat with unhealthy-eating partners.

		<i>Focal person with a healthy-eating partner</i>		<b>Total pairs</b>
		<b>Increase</b>	<b>No increase</b>	
<i>Focal person with an unhealthy-eating partner</i>	<b>Increase</b>	126	67	193
	<b>No increase</b>	103	119	222
<b>Total pairs</b>		229	186	<b>415</b>

## 4.2 Contingency-table analysis

Next, to obtain fine-grained insights about patterns taking place at the pair level, we analyze the outcomes with a contingency table. We binarize the outcome to look for either an increase or no increase post-treatment, compared to pre-treatment. Four possibilities exist for any matched pair: both increased, only one or the other increased, and none increased. The contingency table is presented as Table 2. The table counts the frequency of the four possible results. Using a chi-squared test, we reject the null hypothesis of no treatment effect ( $p = 0.00017$ ).

It is particularly informative to observe the discordant pairs (off-diagonal entries in the contingency table) among the matched pairs. Such pairs correspond to situations when the outcome (increase or no increase) differs in the matched pair. The intuition is the following: if there is no effect, the two types of discordant entries should be balanced. However, we observe that in 103 pairs, the focal person with a positive intervention increased, and the focal person with a negative intervention did not. The reversed situation, in comparison, occurs in 67 pairs. We test the null hypothesis of no effect in a paired randomized experiment using McNemar's test [46], which relies directly on the evidence that comes from the discordant pairs (their number and the ratio between them). Here, too, we reject the null hypothesis of no treatment effect ( $p = 0.007$ ).

## 4.3 Difference-in-differences analysis

We move on to further exploit the matched setup in order to estimate the difference-in-differences [47] effect for pairs of matched focal persons. The idea is to first calculate the difference between post-treatment and pre-treatment healthiness scores for each focal person separately. Then, we can calculate the difference in treatment effects between two matched focal persons in each pair. Averaging the differences in differences across all pairs yields the overall treatment effect.

**Regression model.** In practice, following the standard approach, we estimate the difference-in-differences effect via a regression model. Here, each focal user adds two data points (one pre-treatment, one post-treatment), each of which specifies, as predictors, the type of partner with whom the focal user started to eat as a treatment (healthy-eating or unhealthy-eating) and the time period (pre- or post-treatment); and, as the outcome, the healthiness score of the focal user's food choice during the respective period. Each matched pair thus contributes four data points, and the modeled dataset consists of  $4 \cdot 415 = 1,660$  data points. Formally, the model takes the following form:

$$y_{it} = \alpha + \beta \cdot \text{healthy\_treatment}_i + \gamma \cdot \text{treated}_t + \delta \cdot (\text{healthy\_treatment}_i \cdot \text{treated}_t) + \text{error}_{it}, \quad (1)$$

where the dependent variable  $y_{it}$  is the focal user  $i$ 's healthiness score in period  $t$ , and the independent variables indicate whether  $i$ 's partner has a positive or negative pre-treatment healthiness score ( $\text{healthy\_treatment}_i = 1$  or  $0$ , respectively) and whether the respective data point captures the pre-

or post-treatment period (treated<sub>*t*</sub> = 1 or 0, respectively). The coefficient  $\delta$  of the interaction term, then, is the difference-in-differences effect of starting to eat with a healthy- vs. unhealthy-eating partner.

**Results.** Calculating the average difference-in-differences effect with a linear regression across all matched focal persons, we observe a larger post-treatment increase in focal persons with healthy-eating partners compared to the post-treatment increase in matched counterparts,  $\delta = 0.051$  (95% CI [0.021, 0.076],  $R^2 = 0.07$ ). This means that, accounting for possible temporal drifts between post-treatment and pre-treatment that are not associated with the initiation, focal persons starting to eat with a healthy-eating partner significantly diverge from their matched counterparts starting to eat with an unhealthy-eating partner. Quantitatively, the effect size of  $\delta = 0.051$  means that, compared to matched counterparts who start eating with an unhealthy-eating partner, focal persons who start eating with a healthy-eating partner increase their healthiness score by an additional 5.1% of the full range spanning from a neutral healthiness score (i.e., 0) to the maximum healthiness score (i.e., 1).

Similarly, to estimate the effect of social tie formation on the absolute number of healthy and of unhealthy purchased items, we repeat the regression analysis described in Eq. 1, but now with different dependent variables that capture the *total number of healthy and the total number of unhealthy items purchased* by focal user *i* in period *t*. We observe that the focal persons who start to eat with a partner with a healthy pattern purchase an additional 5.71 (95% CI [3.21, 8.21],  $R^2 = 0.17$ ) healthy items, and an additional  $-1.13$  (95% CI  $[-3.04, 0.78]$ ,  $R^2 = 0.12$ ) unhealthy items in the six months following the tie formation, compared to their matched counterparts.

**Sensitivity analysis.** The above finding relies on the assumption that there are no unobserved variables that create the differences between the matched focal persons that could explain away the measured effect. Sensitivity analysis is a way of quantifying how the results of our calculations would change if the assumptions were violated to a limited extent. If the conclusions of the study would change little, then the study is insensitive to a violation of the assumptions, up to the specified limited extent. In contrast, if the conclusions would change substantially, then the study is sensitive to a violation of the assumption.

The key assumption made in our analysis is that the treatment assignment is not biased, or in other words, that after balancing the pre-treatment covariates, the co-eating initiation with a healthy-eating vs. an unhealthy-eating partner is randomized (i.e., it is effectively decided by a coin flip). We measure by how much that assumption needs to be violated in order to alter our conclusion that there is a significant difference-in-differences effect on the healthiness of purchased items among the matched focal persons. Specifically, sensitivity analysis lets us answer the following question: if there is a violation of randomized treatment assignment (i.e., a deviation from fair 50/50 coin flip), how large would it need to be in order to alter the conclusion that the null hypothesis of no difference between focal persons can be rejected? This notion is quantified by the *sensitivity*  $\Gamma$ , which specifies the ratio by which the treatment odds of two matched persons would need to differ in order to result in a *p*-value above the significance threshold. We always have  $\Gamma \geq 1$ , with larger values of  $\Gamma$  corresponding to more robust conclusions.

For the chosen  $p = 0.05$ , we measure a sensitivity of  $\Gamma = 1.17$ , which implies that, within matched pairs, an individual's probability of being the treated one could take on any value between  $1/(1+\Gamma) = 0.46$  and  $\Gamma/(1+\Gamma) = 0.54$  without changing our decision of rejecting the null hypothesis of no effect. In other words, if the assignment of the treatment after matching were not approximating the ideal 0.5, but a third variable made some people more likely to initiate eating with a healthy-eating or an unhealthy-eating partner, the randomized treatment assignment would have to be violated by deviating from the fair 0.5 by at least four percentage points.

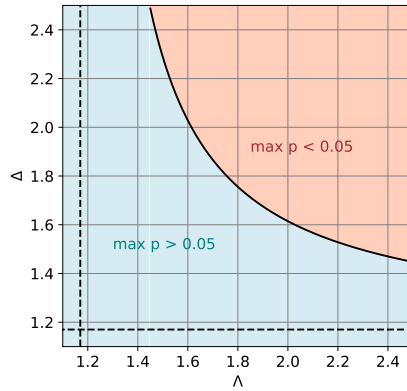


Fig. 11. Sensitivity analysis. For the sensitivity  $\Gamma = 1.17$ , the amplification  $(\Lambda, \Delta)$  is plotted (see text for explanation). Horizontal and vertical dashed lines indicate  $\Gamma$ , i.e., the asymptotic value of  $\Lambda$  for  $\Delta \rightarrow \infty$ , and vice versa.

Additionally, we conduct an amplification of the sensitivity analysis [67]. Amplification is particularly relevant when the concern is not about a violation of the randomized treatment assignment, but rather about the potential existence of a specific unobserved covariate. It then becomes useful to consider possible combinations of  $\Lambda$  and  $\Delta$ , two parameters describing the unobserved covariate, that would result in the measured  $\Gamma$ . The strength of the relationship between the unobserved covariate and the difference in outcomes within the matched pair is defined by  $\Delta$ , whereas  $\Lambda$  defines the strength of the relationship between the unobserved covariate and the difference in probability of being assigned a treatment. With these definitions, the sensitivity  $\Gamma$  can be expressed in terms of  $\Lambda$  and  $\Delta$ , as  $\Gamma = (\Lambda\Delta + 1)/(\Lambda + \Delta)$ .

The result of sensitivity analysis amplification is presented in Figure 11. For combinations of  $\Lambda$  and  $\Delta$  in the orange area, significant effects would be detected (leading to  $p < 0.05$ ), whereas for the combinations in the blue area, no significant effects would be detected (leading to  $p > 0.05$ ). An infinite number of  $(\Lambda, \Delta)$  combinations fall on the border; e.g.,  $(\Lambda, \Delta) = (2.0, 1.6)$  corresponds to an unobserved covariate that doubles the odds of treatment and multiplies the odds of a positive pair difference in the outcomes by 1.6.

Overall, we conclude that the study design is insensitive to small biases [65].

#### 4.4 Dose–response relationship

Next, we analyze the dose–response relationship in our matched setup. Similar focal persons initiate eating with differing partners. We observe systematic changes in the dieting patterns of the focal persons after the tie formation. But do more drastic difference-in-differences effects occur when the differences between partners are more drastic? In the case of a true causal effect, one would expect a dose–response effect where focal persons diverge more post-treatment if their partners diverged more pre-treatment.

Although large differences in the pre-treatment scores between matched focal persons’ partners are rare (Figure 12a), Figure 12b shows evidence of a dose–response relationship: the difference-in-differences effect is stronger when the partners are more different (i.e., the more extreme difference in partners leads to more extreme effect estimates). If there were other confounding factors that could explain the observed difference-in-differences effects, and those factors had nothing to do

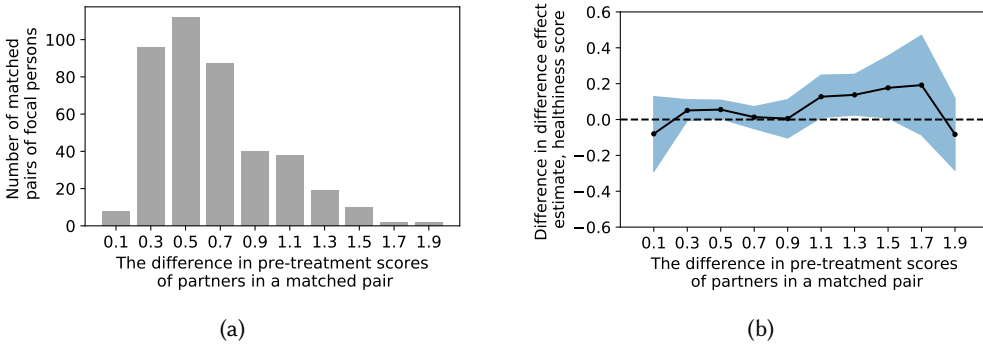


Fig. 12. Dose–response relationship. **(a)** Histogram of pre-treatment differences in healthiness scores between partners of paired focal users. **(b)** Difference-in-differences effect between focal persons in matched pairs, stratified by the pre-treatment differences in healthiness scores between the partners they were exposed to.

with the onset of eating together, we would not expect to find a dose–response relationship. The observed dose–response relationship thus further supports the conclusion of a causal effect.

#### 4.5 Stratification by pre-treatment healthiness

Additionally, we would like to understand for whom the treatment is effective. Are there changes across the board with respect to the initial healthiness, or only for specific sub-populations? For whom is the intervention most efficient? We again monitor the differences between post- and pre-treatment healthiness scores, but now stratified into quartiles by pre-treatment healthiness score of the focal person (Figure 13). Moreover, we repeat this analysis for post-treatment observation periods of varying length (3, 6, and 12 post-treatment months). In the aligned, post-intervention period, persons who start eating with partners with healthy dieting patterns are characterized with consistently higher healthiness scores compared to the matched counterparts, across strata of the focal person’s pre-treatment healthiness score. Note that the fact that the slopes are decreasing may be a simple regression to the mean. The key observation is that, within each stratum, when comparing the outcomes in orange and blue, people who initiate eating with a healthy-eating partner (orange) see a greater post- vs. pre-treatment difference compared to people who initiate eating with an unhealthy-eating partner (blue).




Fig. 13. Post-treatment increase in healthiness score, stratified by pre-treatment healthiness score of focal person (with 95% bootstrapped confidence intervals). The difference is shown separately for focal persons who start eating with a person with a positive (orange) vs. negative (blue) healthiness score. The difference is monitored in the first 3, 6, and 12 post-treatment months (left, center, and right panel, respectively).






Fig. 14. Estimated difference-in-differences effects of co-eating onset with healthy- vs. unhealthy-eating partner on frequency of purchased food categories (with 95% confidence intervals approximated as plus/minus two standard errors). Categories with a significant effect are marked in orange (positive) and blue (negative), whereas categories with no significant effect are marked in gray.

#### 4.6 Analysis of affected food-item categories

Finally, we set out to understand the influence of new co-eating partners on the rates at which categories of food items are subsequently purchased. Since we observed that the behaviors are modified, we now ask: what items are eaten more, and which less? What foods being purchased and eaten in group settings on campus have the largest influence on others?

To estimate category-specific difference-in-differences effects, we repeat the regression analysis described in Eq. 1, but now with a different dependent variable  $y_{cit}$ , which captures the number of items from food category  $c$  purchased by focal user  $i$  in period  $t$ . By fitting a separate regression for each food category  $c$ , we obtain category-specific effects  $\delta_c$ .

The estimated effects  $\delta_c$ , together with 95% confidence intervals, are presented in Figure 14. We observe that the focal persons initiating to eat with a healthy-eating partner purchase more coffee, lunch meals, coffee from vending machines, soup, fruit, dessert, tea, salad, and wraps, compared to their matched counterparts, who purchase more soft drinks, drinks from vending machines, water, condiments, pizza, kebabs, and crêpes. The values on the  $x$ -axis can be interpreted as the number of purchased items by which the matched focal persons diverge in the post-treatment period. For example, in the six months following tie formation, people who start eating with healthy-eating partners purchase, on average, around two additional meals and around four additional coffees, compared to the matched counterparts. The matched counterparts who start eating with unhealthy-eating partners, by contrast, on average purchase around one additional soft drink in the six months following tie formation.

Coffees and lunch meals are the items that see the largest increase after tie formation with a healthy-eating partner. These items are in general purchased in large numbers (Figure 2). Conversely,

items with the strongest effect among the matched counterparts, with the exception of water, loosely form a cluster of potentially unhealthy items that should not be eaten in large quantities. The remaining items with a significant positive effect, soups, fruits, desserts, tea, salad buffet, and wraps are overall less indicative of an unhealthy dietary pattern.

## 5 DISCUSSION AND CONCLUSIONS

We report on a longitudinal observational study of the effect of the formation of social ties on food choice, leveraging a novel source of data: logs of millions of food purchases made over an eight-year period on a major university campus. To estimate causal effects from the passively observed log data, we control confounds in a matched quasi-experimental design: we identify focal persons who start regularly eating with a fixed partner and match focal persons into pairs such that paired focal persons are nearly identical with respect to covariates measured before acquiring the partner, but the new eating partners diverge with respect to the healthiness of their respective food choice behaviors (before tie formation with the focal person).

We observe that the people who acquire a healthy-eating partner change their habits significantly more toward healthy foods than those acquiring an unhealthy-eating partner. We further identify foods whose purchase frequency is impacted significantly by the partner's healthiness of food choice: coffees and lunch meals are the items that see the largest increase on behalf of those who initiate eating with a healthy-eating partner, whereas the matched counterparts, with the exception of water, increase purchases of items that loosely form a cluster of potentially unhealthy items that should not be consumed in large quantities: soft drinks, drinks from vending machines, condiments, pizza, kebabs, and crêpes.

Our findings show that digital traces can be used as a valuable tool for monitoring dietary habits, and can provide valuable insights into the effects of social ties on dietary choices. This work establishes the feasibility of relying on transactional logs in order to monitor food consumption and derive meaningful insights about behavioral patterns taking place at a population scale. Such digital traces can complement small-scale field experiments, making it possible to observe large populations over long time periods.

By relying on a novel transactional dataset and using carefully designed quasi-experimental methodology, we confirm theories of social influence on food choice postulating that social influence plays a prominent role and has a powerful effect on food intake [18, 38, 39, 52, 76]. Conforming to the behavior of others is adaptive, and individuals find it rewarding [39]. Hence, dietary choices are expected to converge with those of our close social connections.

Additionally, eating norms are known to reduce the intake of unhealthy foods [61] in specific contexts. For example, exposure to social eating norms is known to result in a reduction in the weight of consumed high-energy snack foods [61]. The fact that persons exposed to a partner eating healthy foods eat fewer items that should not be eaten in large quantities (soft drinks, drinks from vending machines, condiments, pizza, kebabs, and crêpes) compared to the counterparts indicates the presence of such positive influence of others.

Finally, we contribute to the rich literature about social influences in eating by demonstrating that, beyond studying social modeling in specific situations [82], the naturally or experimentally occurring event of tie formation is an important dimension to consider in order to understand the mechanisms of social influence and their potential for influence and interventions.

**Implications.** The most imminent utility of measuring the impact of social tie formation on food choice relying on large-scale passively sensed transaction signals lies in its potential to inform public health interventions on campuses. Designing large-scale nutritional interventions is challenging and logistically complicated. Additionally, it is difficult to predict their impact through experimentation

due to ethical concerns stemming from the danger of eliciting undesirable effects. This work shows how observational insights based on passively sensed data can be used to evaluate the impact of potential interventions by estimating the impact of similar interventions that occur naturally, without external experimentation.

For instance, in order to incentivize healthy eating habits, university or corporate stakeholders might consider launching programs with disclosure and consent that help students or staff connect onsite with “lunch buddies” and incentivize consumption of meals in a company. Such programs leveraging peer-led nutritional interventions [89] would need to take into consideration self-selective disclosure and consent as confounding factors. First, by relying on passively sensed observational data and evaluating the impact of similar interventions, it would be possible to anticipate the impact on the involved individuals ahead of implementing any interventions in the real world. Second, it would be possible to optimize the pairing of people, by estimating on whom the effect of tie formation with others could be strongest. Finally, being informed about what products are most likely to be purchased in modified quantities after tie formation, the stakeholders could optimize the offering of products.

The question of whether the estimated impact of such naturally occurring interventions is identical to the true causal effect, and whether it is expected to mirror the impact of intentional externally induced interventions, remains. In what follows, we discuss the assumptions that would need to hold, and the limitations that should be considered.

**Causal assumptions.** We discuss what assumptions are necessary for our study design to let us isolate the causal effect of social-tie formation on food consumption. In particular, we consider the assumptions of the potential-outcomes framework [68] and the extent to which they can be assumed to hold in the present work.

*Stable unit treatment value.* Units are assumed not to interfere with each other. In other words, the treatment assignment of one unit does not affect the outcome of another unit, i.e., there is no “spillover” or “contagion.” Recall that we require initiation with no more than one peer during the monitored pre- and post-treatment periods. While this restriction ensures that there are no spillovers, our study is unable to capture complex network interactions occurring in on-campus settings. Future work should therefore generalize beyond the studied setup by identifying time-varying treatments and dynamic treatment regimes using g-formula methods [79, 91].

*Consistency.* The potential outcome under a treatment is assumed to be equal to the observed outcome when the actual treatment is received. In other words, the counterfactual outcome for treated units is the observed outcome for controls. While our study design attempts to compare people who are similar up to the moment of receiving the treatment by modeling the propensity to be treated, the assumption that the outcome of the matched person would be exactly the outcome of the treated person is untestable.

*Positivity.* The probability of every treatment for every set of covariates is assumed to be non-zero. In our study, it is reasonable to assume there are no people who could never possibly receive a given treatment.

*Ignorability.* Finally, the ignorability assumption refers to the absence of unmeasured confounding. By performing a sensitivity analysis, we have attempted to assess the possible impact of unobserved biases if the ignorable treatment assignment assumption is violated. This analysis leads us to conclude that our findings are insensitive to small biases.

**Limitations.** This study is subject to certain limitations, some of which suggest promising directions for future work. Inherent to observational studies, we recognize the inability to infer true causality. Controlling for all the possible confounds is fundamentally infeasible. Still, we make a step towards understanding the effects of a certain “treatment” on food consumption by developing

a quasi-experimental design based on propensity-score matching and difference-in-differences methods, whereby we seek to minimize biases due to *observed* confounding variables, enriched with a quantification of the danger of *unobserved* confounding variables via sensitivity analysis. Our work, therefore, provides insights based on passively sensed behavioral signals that go far beyond simpler correlational analyses.

We note that the inference of social ties might be imperfect. However, the fact that a large fraction of ties forms precisely at the beginning of the academic year (with the fall semester), when students are exposed to new fellow students (Figure 2), and the fact that there is a correlation with ground-truth team membership points towards reliability. In addition, we note that we might potentially be detecting the onset of co-eating with a delay (i.e., it actually occurs earlier than detected) if peers eat together as part of a larger group and are not directly adjacent in the queue, or if they use cash for the transaction. That said, we note that any potential delay in estimating the onset would lead to more conservative estimates of the effect of partnering up, as potential changes in the patterns would be counted as purchases before the tie formed.

*Construct validity.* We also consider the issue of the construct validity of our study design, i.e., the degree to which the obtained indirect measurements (transaction logs of food consumption) are reflecting the true phenomenon that is intended to be measured (actual food consumption). It is reasonable to assume that students and staff indeed consume the food that they purchase. However, one cannot eliminate the possibility of persons borrowing the card, or paying for items consumed by other people. Conversely, people on campus may consume food not recorded in the purchase logs (e.g., food purchased using cash, or food prepared at home or in off-campus restaurants). Future work should determine the extent to which the assumption that purchasing implies consumption, and vice versa, holds.

*External validity.* Future work should also determine external validity, that is, to what extent behaviors measured on campus reflect other settings off campus, e.g., food consumption—and the complex interplay with food consumption of their social ties at home—in families or in restaurants. People who primarily eat home-cooked food might present a skewed representation of their diet in the studied cafeteria purchase logs, and the fact that individuals might consume food prepared at home introduces unobserved variables into our analyses. For example, people who do not have regular eating partners in extended periods might eat alone because they prefer to eat a peculiar type of food that they cook at home and that is not served on campus, or measured via purchase logs.

**Ethical considerations.** Nutrition is a potentially sensitive personal behavior. To protect user privacy, the log data used here was accessed exclusively by EPFL personnel involved in this project, and stored and processed exclusively on EPFL servers. The data was obtained with approval from EPFL's Data Protection Officer and was anonymized before it was made available to the researchers for analysis. Finally, we note that our work was conducted retroactively on data that had been collected passively in order to support campus operations. Thus, our analysis did not influence users in any way.

**Future work.** This study opens the door for future research directions and potential follow-up studies of the mechanisms of social influence beyond those described, such as the role of the campus-wide sustainability challenge (cf. Section 3.1), and if eating with a healthy-eating partner is correlated with other behavioral changes (e.g., switching cafeterias). Future work should further understand the measured behavioral changes by understanding the underlying mechanisms and the social contexts in which the foods are consumed in modified quantities (e.g., socializing by meeting for a cup of coffee versus consuming a meal together). Future work should demonstrate whether modified habits are retained in those individuals who return later to eating alone for long

stretches, or, if habits are not retained, how the individual regresses to the state they were in before the tie formation, and if so, what is the rate of the decay back.

**Beyond food.** To the best of our knowledge, this is the first study to leverage large-scale transactional data to retrospectively evaluate the impact of implicit dietary behavioral interventions. We show that purchase logs are useful as an effective sensor to detect behavioral changes in individuals who interact with others. Beyond food purchase logs on campus, this work highlights how careful quasi-experimental comparisons leveraging passively sensed data can be used to measure complex interactions between individuals. The methodology of isolating influence through matched pairwise comparisons has the potential to allow for detecting behavioral changes also in areas relevant to other social computing scenarios. Ambitious questions of how individuals impact behaviors or skills of one another in the context of purchasing decisions, exercise habits, coaching or counselling, travelling, sleep patterns, online interactions, collaborative writing, or coding practice could all be tackled by identifying the onset of such interactions, and comparing outcomes in similar matched persons while controlling for confounds.

**Code and data.** Although the transaction logs cannot be shared publicly, for transparency we make all analysis code, as well as the food-item and healthiness-score annotations, publicly available: <https://github.com/epfl-dlab/tie-formation-food>

## ACKNOWLEDGMENTS

We would like to thank Nils Rinaldi, Aurore Nembrini, and Philippe Vollichard for their help in obtaining and anonymizing the data. We are also grateful to Jonas Racine and Kiran Garimella for early help with data engineering, and to our reviewers for their helpful suggestions. This project was funded by the Microsoft Swiss Joint Research Center.

## REFERENCES

- [1] Sofiane Abbar, Yelena Mejova, and Ingmar Weber. 2015. You tweet what you eat: Studying food consumption through twitter. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*. 3197–3206.
- [2] Palakorn Achananuparp and Ingmar Weber. 2016. Extracting food substitutes from food diary via distributional similarity. (2016).
- [3] Luca Maria Aiello, Rossano Schifanella, Daniele Quercia, and Lucia Del Prete. 2019. Large-scale and high-resolution analysis of food purchases and health outcomes. *EPJ Data Science* 8, 1 (2019), 14.
- [4] Sinan Aral, Lev Muchnik, and Arun Sundararajan. 2009. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences* 106, 51 (2009), 21544–21549.
- [5] Dan Ariely and Jonathan Levav. 2000. Sequential choice in group settings: Taking the road less traveled and less enjoyed. *Journal of Consumer Research* 27, 3 (2000), 279–290.
- [6] Kylie Ball, Robert W Jeffery, Gavin Abbott, Sarah A McNaughton, and David Crawford. 2010. Is healthy behavior contagious: associations of social norms with physical activity and healthy eating. *International Journal of Behavioral Nutrition and Physical Activity* 7, 1 (2010), 86.
- [7] Kieron J Barclay, Christofer Edling, and Jens Rydgren. 2013. Peer clustering of exercise and eating behaviours among young adults in Sweden: a cross-sectional study of egocentric network data. *BMC Public Health* 13, 1 (2013), 784.
- [8] Mehrab Bin Morshed, Samruddhi Shreeram Kulkarni, Richard Li, Koustuv Saha, Leah Galante Roper, Lama Nachman, Hong Lu, Lucia Mirabella, Sanjeev Srivastava, Munmun De Choudhury, Kaya de Barbaro, Thomas Ploetz, and Gregory D Abowd. 2020. A Real-Time eating detection system for capturing eating moments and triggering ecological momentary assessments to obtain further context: system development and validation study. *JMIR mHealth and uHealth* 8, 12 (2020).
- [9] Leann Lipps Birch. 1980. Effects of peer models' food choices and eating behaviors on preschoolers' food preferences. *Child Development* (1980), 489–496.
- [10] DJ Blok, Pepijn van Empelen, FJ Van Lenthe, Jan Hendrik Richardus, and SJ De Vlas. 2013. Unhealthy behaviour is contagious: an invitation to exploit models for infectious diseases. *Epidemiology & Infection* 141, 3 (2013), 667–669.
- [11] Ann Bowling. 2005. Mode of questionnaire administration can have serious effects on data quality. *Journal of Public Health* 27, 3 (2005), 281–291.

- [12] David L Buckeridge, Katia Charland, Alice Labban, and Yu Ma. 2014. A method for neighborhood-level surveillance of food purchasing. *Annals of the New York Academy of Sciences* 1331, 1 (2014), 270–277.
- [13] Nicoletta Cavazza, Margherita Guidetti, and Fabrizio Butera. 2017. Portion size tells who I am, food type tells who you are: Specific functions of amount and type of food in same-and opposite-sex dyadic eating contexts. *Appetite* 112 (2017), 96–101.
- [14] Stevie Chancellor, Zhiyuan Lin, Erica L. Goodman, Stephanie Zerwas, and Munmun De Choudhury. 2016. Quantifying and predicting mental illness severity in online pro-eating disorder communities. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)*. 1171–1184.
- [15] Martin J Chorley, Luca Rossi, Gareth Tyson, and Matthew J Williams. 2016. Pub crawling at scale: tapping untappd to explore social drinking. In *Proceedings of the 10th International AAAI Conference on Web and Social Media (ICWSM)*.
- [16] Nicholas A Christakis and James H Fowler. 2007. The spread of obesity in a large social network over 32 years. *New England Journal of Medicine* 357, 4 (2007), 370–379.
- [17] Chelsea D Christie and Frances S Chen. 2018. Vegetarian or meat? Food choice modeling of main dishes occurs outside of awareness. *Appetite* 121 (2018), 50–54.
- [18] Emily IM Collins, Jason M Thomas, Eric Robinson, Paul Aveyard, Susan A Jebb, C Peter Herman, and Suzanne Higgs. 2019. Two observational studies examining the effect of a social norm and a health message on the purchase of vegetables in student canteen settings. *Appetite* 132 (2019), 122–130.
- [19] Felicia Cordeiro, Elizabeth Bales, Erin Cherry, and James Fogarty. 2015. Rethinking the mobile food journal: Exploring opportunities for lightweight photo-based capture. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*. 3207–3216.
- [20] Felicia Cordeiro, Daniel A. Epstein, Edison Thomaz, Elizabeth Bales, Arvind K. Jagannathan, Gregory D. Abowd, and James Fogarty. 2015. Barriers and negative nudges: Exploring challenges in food journaling. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*. 1159–1162.
- [21] David J Crandall, Lars Backstrom, Dan Cosley, Siddharth Suri, Daniel Huttenlocher, and Jon Kleinberg. 2010. Inferring social ties from geographic coincidences. *Proceedings of the National Academy of Sciences* 107, 52 (2010), 22436–22441.
- [22] Tegan Cruwys, Kirsten E. Bevelander, and Roel C.J. Hermans. 2015. Social modeling of eating: A review of when and why social influence affects food intake and choice. *Appetite* 86 (2015), 3–18.
- [23] Munmun De Choudhury, Sanket Sharma, and Emre Kiciman. 2016. Characterizing dietary choices, nutrition, and language in food deserts via social media. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing (CSCW)*. 1157–1170.
- [24] Kayla de la Haye, Garry Robins, Philip Mohr, and Carlene Wilson. 2010. Obesity-related behaviors in adolescent friendship networks. *Social Networks* 32, 3 (2010), 161–167.
- [25] Kayla de la Haye, Garry Robins, Philip Mohr, and Carlene Wilson. 2011. How physical activity shapes, and is shaped by, adolescent friendships. *Social Science and Medicine* 73, 5 (2011), 719–728.
- [26] Kayla De La Haye, Garry Robins, Philip Mohr, and Carlene Wilson. 2013. Adolescents' intake of junk food: Processes and mechanisms driving consumption similarities among friends. *Journal of Research on Adolescence* 23, 3 (2013), 524–536.
- [27] Mary Delaney and Mary McCarthy. 2011. Food choice and health across the life course: A qualitative study examining food choice in older Irish adults. *Journal of Food Products Marketing* 17, 2-3 (2011), 114–140.
- [28] Marla E. Eisenberg, Dianne Neumark-Sztainer, Mary Story, and Cheryl Perry. 2005. The role of social norms and friends' influences on unhealthy weight-control behaviors among adolescent girls. *Social Science & Medicine* 60, 6 (2005), 1165–1173.
- [29] Gerda I.J. Feunekes, Cees de Graaf, Saskia Meyboom, and Wija A. van Staveren. 1998. Food choice and fat intake of adolescents and adults: Associations of intakes within social networks. *Preventive Medicine* 27, 5 (1998), 645–656.
- [30] Tara Finnerty, Sue Reeves, Jaqueline Dabinett, Yvonne M Jeanes, and Claus Vögele. 2010. Effects of peer influence on dietary intake and physical activity in schoolchildren. *Public Health Nutrition* 13, 3 (2010), 376–383.
- [31] Brianna Fjeldsoe, Maïke Neuhaus, Elisabeth Winkler, and Elizabeth Eakin. 2011. Systematic review of maintenance of behavior change following physical activity and dietary interventions. *Health Psychology* 30, 1 (2011), 99.
- [32] Adam Fletcher, Chris Bonell, and Annik Sorhaindo. 2011. You are what your friends eat: systematic review of social network analyses of young people's eating behaviours and bodyweight. *Journal of Epidemiology & Community Health* 65, 6 (2011), 548–555.
- [33] Emmanuela Gakidou, Ashkan Afshin, Amanuel Alemu Abajobir, Kalkidan Hassen Abate, Cristiana Abbafati, Kaja M Abbas, Foad Abd-Allah, Abdishakur M Abdulle, Semaw Ferede Abera, Victor Aboyans, et al. 2017. Global, regional, and national comparative risk assessment of 84 behavioural, environmental and occupational, and metabolic risks or clusters of risks, 1990–2016: a systematic analysis for the Global Burden of Disease Study 2016. *The Lancet* 390, 10100 (2017), 1345–1422.

- [34] Joel Gittelsohn, Megan Rowan, and Preety Gadhoke. 2012. Interventions in small food stores to change the food environment, improve diet, and reduce risk of chronic disease. *Preventing Chronic Disease* 9 (2012).
- [35] Samuel L Groseclose and David L Buckridge. 2017. Public health surveillance systems: recent advances in their use and evaluation. *Annual Review of Public Health* 38 (2017), 57–79.
- [36] Brook E Harmon, Melinda Forthofer, Erin O Bantum, and Claudio R Nigg. 2016. Perceived influence and college students' diet and physical activity behaviors: an examination of ego-centric social networks. *BMC Public Health* 16, 1 (2016), 1–10.
- [37] Roel CJ Hermans, Junilla K Larsen, C Peter Herman, and Rutger CME Engels. 2009. Effects of social modeling on young women's nutrient-dense food intake. *Appetite* 53, 1 (2009), 135–138.
- [38] Marion M Hetherington, Annie S Anderson, Geraldine NM Norton, and Lisa Newson. 2006. Situational effects on meal intake: A comparison of eating alone and eating with others. *Physiology & Behavior* 88, 4-5 (2006), 498–505.
- [39] Suzanne Higgs and Jason Thomas. 2016. Social influences on eating. *Current Opinion in Behavioral Sciences* 9 (2016), 1–6.
- [40] Patrick D. Howell, Layla D. Martin, Hesamoddin Salehian, Chul Lee, Kyler M. Eastman, and Joohyun Kim. 2016. Analyzing taste preferences from crowdsourced food entries. In *Proceedings of the 6th International Conference on Digital Health Conference (DH)*. 131–140.
- [41] Robert W Jeffery, Rena R Wing, Carolyn Thorson, Lisa R Burton, Cheryl Raether, Jean Harvey, and Monica Mullen. 1993. Strengthening behavioral interventions for weight loss: a randomized trial of food provision and monetary incentives. *Journal of Consulting and Clinical Psychology* 61, 6 (1993), 1038.
- [42] Eric S Johnson, Barbara A Bartman, Becky A Briesacher, Neil S Fleming, Tobias Gerhard, Cynthia J Kornegay, Parivash Nourjah, Brian Sauer, Glen T Schumock, Art Sedrakyan, et al. 2013. The incident user design in comparative effectiveness research. *Pharmacoepidemiology and Drug Safety* 22, 1 (2013), 1–6.
- [43] O Kaka, F Bert, C Botezatu, MR Gualano, and R Siliquini. 2018. How we make choices about food? Analysis of factors influencing food expenditure in Northern Italy. *European Journal of Public Health* 28 (11 2018).
- [44] Emre Kıcıman, Scott Counts, and Melissa Gasser. 2018. Using longitudinal social media analysis to understand the effects of early college alcohol use. *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)* (2018).
- [45] Gueorgi Kossinets and Duncan J Watts. 2006. Empirical analysis of an evolving social network. *Science* 311, 5757 (2006), 88–90.
- [46] Peter A Lachenbruch and Cornelius J Lynch. 1998. Assessing screening tests: extensions of McNemar's test. *Statistics in Medicine* 17, 19 (1998), 2207–2217.
- [47] Michael Lechner. 2011. The estimation of causal effects by difference-in-difference methods. *Foundations and Trends in Econometrics* 4, 3 (2011), 165–224.
- [48] Jennifer L Lund, David B Richardson, and Til Stürmer. 2015. The active comparator, new user study design in pharmacoepidemiology: historical foundations and contemporary application. *Current Epidemiology Reports* 2, 4 (2015), 221–228.
- [49] Anmol Madan, Sai T Moturu, David Lazer, and Alex Sandy Pentland. 2010. Social sensing: obesity, unhealthy eating and exercise in face-to-face networks. In *Wireless Health 2010*. 104–110.
- [50] Brent McFerran, Darren W. Dahl, Gavan J. Fitzsimons, and Andrea C. Morales. 2010. I'll have what she's having: Effects of social influence and body type on the food choices of others. *Journal of Consumer Research* 36, 6 (2010), 915–929.
- [51] Yelena Mejova, Sofiane Abbar, and Hamed Haddadi. 2016. Fetishizing food in digital age: #foodporn around the world. In *Tenth International AAAI Conference on Web and Social Media (ICWSM)*.
- [52] Saar Mollen, Rajiv N Rimal, Robert AC Ruiter, and Gerjo Kok. 2013. Healthy and unhealthy social norms and food selection. Findings from a field-experiment. *Appetite* 65 (2013), 83–89.
- [53] AE Munt, SR Partridge, and M Allman-Farinelli. 2017. The barriers and enablers of healthy eating among young adults: A missing piece of the obesity puzzle: A scoping review. *Obesity Reviews* 18, 1 (2017), 1–17.
- [54] Erik C Nook and Jamil Zaki. 2015. Social norms shift behavioral and neural responses to foods. *Journal of Cognitive Neuroscience* 27, 7 (2015), 1412–1426.
- [55] Ferda Ofli, Yusuf Aytar, Ingmar Weber, Raggi Al Hammouri, and Antonio Torralba. 2017. Is saki# delicious?: The food perception gap on Instagram and its relation to health. In *Proceedings of the 26th International Conference on World Wide Web (WWW)*. 509–518.
- [56] Mark A Pachucki, Paul F Jacques, and Nicholas A Christakis. 2011. Social network concordance in food choice among spouses, friends, and siblings. *American Journal of Public Health* 101, 11 (2011), 2170–2177.
- [57] Jessica A. Pater, Oliver L. Haimson, Nazanin Andalibi, and Elizabeth D. Mynatt. 2016. "Hunger hurts but starving works": Characterizing the presentation of eating disorders online. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)*. 1185–1200.

- [58] Heather Patrick and Theresa A Nicklas. 2005. A review of family and social determinants of children's eating patterns and diet quality. *Journal of the American College of Nutrition* 24, 2 (2005), 83–92.
- [59] Susanne Pedersen, Alice Grønhoj, and John Thøgersen. 2015. Following family or friends. Social norms in adolescent healthy eating. *Appetite* 86 (2015), 54–60.
- [60] Christian Reynolds, Liam Goucher, Tom Quedsted, Sarah Bromley, Sam Gillick, Victoria K Wells, David Evans, Lenny Koh, Annika Carlsson Kanyama, Cecilia Katzeff, et al. 2019. Consumption-stage food waste reduction interventions—What works and how to design better interventions. *Food Policy* 83 (2019), 7–27.
- [61] Eric Robinson, Jackie Blissett, and Suzanne Higgs. 2013. Social influences on eating: implications for nutritional interventions. *Nutrition Research Reviews* 26, 2 (2013), 166–176.
- [62] Eric Robinson and Suzanne Higgs. 2013. Food choices in the presence of “healthy” and “unhealthy” eating partners. *British Journal of Nutrition* 109, 4 (2013), 765–771.
- [63] Eric Robinson, Jason Thomas, Paul Aveyard, and Suzanne Higgs. 2014. What everyone else is eating: A systematic review and meta-analysis of the effect of informational eating norms on eating behavior. *Journal of the Academy of Nutrition and Dietetics* 114, 3 (2014), 414–429.
- [64] Markus Rokicki, Christoph Trattner, and Elco Herder. 2018. The impact of recipe features, social cues and demographics on estimating the healthiness of online recipes. In *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM)*.
- [65] Paul R Rosenbaum. 2017. *Observation and experiment*. Harvard University Press.
- [66] Paul R. Rosenbaum and Donald B. Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, 1 (04 1983), 41–55.
- [67] Paul R Rosenbaum and Jeffrey H Silber. 2009. Amplification of sensitivity analysis in matched observational studies. *J. Amer. Statist. Assoc.* 104, 488 (2009), 1398–1405.
- [68] Donald B Rubin. 2005. Causal inference using potential outcomes: Design, modeling, decisions. *J. Amer. Statist. Assoc.* 100, 469 (2005), 322–331.
- [69] Adam Sadilek, Stephanie Caty, Lauren DiPrete, Raed Mansour, Tom Schenk, Mark Bergtholdt, Ashish Jha, Prem Ramaswami, and Evgeniy Gabrilovich. 2018. Machine-learned epidemiology: real-time detection of foodborne illness at scale. *npj Digital Medicine* 1, 1 (2018), 36.
- [70] Sarah-Jeanne Salvy, Kayla De La Haye, Julie C Bowker, and Roel CJ Hermans. 2012. Influence of peers and friends on children's and adolescents' eating and activity behaviors. *Physiology & Behavior* 106, 3 (2012), 369–378.
- [71] Sarah-Jeanne Salvy, Elizabeth Kieffer, and Leonard H Epstein. 2008. Effects of social context on overweight and normal-weight children's food selection. *Eating Behaviors* 9, 2 (2008), 190–196.
- [72] Yasaman S Sefidgar, Woosuk Seo, Kevin S Kuehn, Tim Althoff, Anne Browning, Eve Riskin, Paula S Nurius, Anind K Dey, and Jennifer Mankoff. 2019. Passively-sensed behavioral correlates of discrimination events in college students. *Proceedings of the ACM on Computer-Supported Cooperative Work & Social Computing (CSCW)* (2019), 1–29.
- [73] Cosma Rohilla Shalizi and Edward McFowland III. 2016. Estimating causal peer influence in homophilous social networks by inferring latent locations. (2016).
- [74] Cosma Rohilla Shalizi and Andrew C Thomas. 2011. Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods & Research* 40, 2 (2011), 211–239.
- [75] Sanket S Sharma and Munmun De Choudhury. 2015. Measuring and characterizing nutritional information of food and ingestion content in Instagram. In *Proceedings of the 24th International Conference on World Wide Web (WWW)*. 115–116.
- [76] Richard Shepherd. 1999. Social determinants of food choice. *Proceedings of the Nutrition Society* 58, 4 (1999), 807–812.
- [77] Clifford Stevenson, Glenda Doherty, Julie Barnett, Orla T Muldoon, and Karen Trew. 2007. Adolescents' views of food and eating: Identifying barriers to healthy eating. *Journal of Adolescence* 30, 3 (2007), 417–434.
- [78] V. Das Swain, H. Kwon, B. Saket, M. Bin Morshed, K. Tran, D. Patel, Y. Tian, J. Philipose, Y. Cui, T. Plötz, M. De Choudhury, and G. D. Abowd. 2020. Leveraging WiFi network logs to infer social interactions: A case study of academic performance and student behavior. arXiv:2005.11228
- [79] Sarah L Taubman, James M Robins, Murray A Mittleman, and Miguel A Hernán. 2009. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *International Journal of Epidemiology* 38, 6 (2009), 1599–1611.
- [80] Edison Thomaz, Irfan Essa, and Gregory D. Abowd. 2015. A practical approach for recognizing eating moments with wrist-mounted inertial sensing. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. 1029–1040.
- [81] Lukar E Thornton, Robert W Jeffery, and David A Crawford. 2013. Barriers to avoiding fast-food consumption in an environment supportive of unhealthy eating. *Public Health Nutrition* 16, 12 (2013), 2105–2113.
- [82] Lenny R Vartanian, Samantha Spanos, C Peter Herman, and Janet Polivy. 2015. Modeling of food intake: a meta-analytic review. *Social Influence* 10, 3 (2015), 119–136.



- [83] Simeon Vosen and Torsten Schmidt. 2011. Forecasting private consumption: survey-based indicators vs. Google trends. *Journal of Forecasting* 30, 6 (2011), 565–578.
- [84] Paul Walter, Esther Infanger, and Pascale Mühlemann. 2007. Food pyramid of the Swiss Society for Nutrition. *Annals of Nutrition and Metabolism* 51, 2 (2007), 15–20.
- [85] Robert West, Ryen W. White, and Eric Horvitz. 2013. From Cookies to Cooks: Insights on Dietary Patterns via Analysis of Web Usage Logs. In *Proceedings of the 22nd International Conference on World Wide Web (WWW)*. 1399–1410.
- [86] Michael J Widener and Wenwen Li. 2014. Using geolocated Twitter data to monitor the prevalence of healthy and unhealthy food references across the US. *Applied Geography* 54 (2014), 189–197.
- [87] Walter Willett, Johan Rockström, Brent Loken, Marco Springmann, Tim Lang, Sonja Vermeulen, Tara Garnett, David Tilman, Fabrice DeClerck, Amanda Wood, et al. 2019. Food in the Anthropocene: the EAT–Lancet Commission on healthy diets from sustainable food systems. *The Lancet* 393, 10170 (2019), 447–492.
- [88] Eveline J Wouters, Junilla K Larsen, Stef P Kremers, Pieter C Dagnelie, and Rinie Geenen. 2010. Peer influence on snacking behavior in adolescence. *Appetite* 55, 1 (2010), 11–17.
- [89] Calvin Yip, Michelle Gates, Allison Gates, and Rhona M Hanning. 2016. Peer-led nutrition education programs for school-aged youth: a systematic review of the literature. *Health Education Research* 31, 1 (2016), 82–97.
- [90] Kazuki Yoshida, Daniel H Solomon, and Seoyoung C Kim. 2015. Active-comparator design and new-user design in observational studies. *Nature Reviews Rheumatology* 11, 7 (2015), 437–441.
- [91] Jessica G Young, Lauren E Cain, James M Robins, Eilis J O'Reilly, and Miguel A Hernán. 2011. Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. *Statistics in Biosciences* 3, 1 (2011), 119.

Received October 2020; revised January 2021; accepted January 2021