Can you Trust the Trend? Discovering Simpson's Paradoxes in Social Data

Nazanin Alipourfard*
USC/ISI

4676 Admiralty Way, Marina Del Rey
Los Angeles, California 90292
nazanina@isi.edu

Peter G. Fennell
USC/ISI

4676 Admiralty Way, Marina Del Rey
Los Angeles, California 90292
pfennell@isi.edu

Kristina Lerman
USC/ISI

4676 Admiralty Way, Marina Del Rey
Los Angeles, California 90292
lerman@isi.edu

ABSTRACT

We investigate how Simpson's paradox affects analysis of trends in social data. According to the paradox, the trends observed in data that has been aggregated over an entire population may be different from, and even opposite to, those of the underlying subgroups. Failure to take this effect into account can lead analysis to wrong conclusions. We present a statistical method to automatically identify Simpson's paradox in data by comparing statistical trends in the aggregate data to those in the disaggregated subgroups. We apply the approach to data from Stack Exchange, a popular question-answering platform, to analyze factors affecting answerer performance, specifically, the likelihood that an answer written by a user will be accepted by the asker as the best answer to his or her question. Our analysis confirms a known Simpson's paradox and identifies several new instances. These paradoxes provide novel insights into user behavior on Stack Exchange.

CCS CONCEPTS

• **Mathematics of computing** → **Statistical paradigms; Exploratory data analysis; Regression analysis**; • **Information systems** → **Data mining; Collaborative and social computing systems and tools; Clustering**;

KEYWORDS

Simpson's Paradox, Trend Analysis

ACM Reference format:

Nazanin Alipourfard, Peter G. Fennell, and Kristina Lerman. 2018. Can you Trust the Trend? Discovering Simpson's Paradoxes in Social Data. In *Proceedings of WSDM 2018: The Eleventh ACM International Conference on Web Search and Data Mining*, Marina Del Rey, CA, USA, February 5-9, 2018 (WSDM 2018), 9 pages.
<https://doi.org/10.1145/3159652.3159684>

*N. Alipourfard and P. Fennell contributed equally to this work.

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 ACM 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.

WSDM 2018, February 5-9, 2018, Marina Del Rey, CA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5581-0/18/02...\$15.00

<https://doi.org/10.1145/3159652.3159684>

1 INTRODUCTION

Digital traces of human activity have exposed social behavior to web and data mining algorithms. Researchers have analyzed the growing social data to understand, among other things, how information spreads in social networks [18], and the factors affecting user engagement [2] and performance [19] in online platforms. Yet social data analysis presents multi-faceted challenges that existing data mining approaches are not well-equipped to handle. Real-world data is noisy and sparse. To uncover hidden patterns, scientists aggregate data over the population, which tends to improve statistics and signal-to-noise ratio. However, real-world data is also heterogeneous, i.e., composed of subgroups that vary widely in size and behavior. This heterogeneity can severely bias analysis of social data.

One example of such bias is Simpson's paradox. According to the paradox, an association observed in data that has been aggregated over an entire population may be quite different from, and even opposite to, those of the underlying subgroups. Failure to take this effect into account can distort conclusions drawn from data. One of the better known examples of Simpson's paradox comes from a study of gender bias in graduate admissions [4]. In the aggregate admissions data there appears to be a statistically significant bias against women: a smaller fraction of female applicants is admitted for graduate studies. However, when admissions data is disaggregated by department, women have parity and even a slight advantage over men in some departments. The paradox arises because departments preferred by female applicants have lower admissions rates for both genders.

Simpson's paradox also affects analysis of trends. When measuring how an outcome changes as a function of an independent variable, the characteristics of the population over which the trend is measured may change as a function of the independent variable due to survivor bias. As a result, the data may appear to exhibit a trend, which disappears or reverses when the data is disaggregated. For example, it may appear that repeated exposures to information or hashtags on social media make an individual less likely to share it with his or her followers [18]. In fact, the opposite is true: the more people are exposed to information, the more likely they are to share it with followers [16]. However, those who follow many others (and are likely to be exposed to a meme or a hashtag multiple times) are less responsive overall, due to the high volume of information they receive [12]. Their reduced susceptibility to exposures biases the aggregate response, leading to wrong conclusions about behavior. Once disaggregated based on the volume of information received, a clearer pattern of exposure response emerges, one that is more predictive of the actual response [13]. However, despite

accumulating evidence that Simpson’s paradox affects inference of trends in social and behavioral data [1, 2, 15, 19], researchers do not routinely test for it in their studies.

We describe a method to identify Simpson’s paradoxes in the analysis of trends in social data. Our statistical approach finds pairs of variables—that we call *Simpson’s pairs*—such that a trend in some outcome as a function of the first *independent* variable observed in the aggregate data disappears or reverses itself when the data is disaggregated into distinct subgroups on the second *conditioning* variable. We perform mathematical analysis, which identifies two necessary conditions for the paradox to occur: (1) the independent and conditioning variables are correlated and (2) the value of the outcome variable differs within conditioning subgroups.

We apply the proposed approach to data collected from Stack Exchange, a popular question-answering platform. Specifically, we analyze factors affecting answerer performance, measured as the likelihood the answer provided by the answerer will be accepted by the asker as the best answer to his or her question. We construct a variety of features to describe answers and answerers, and study how the outcome—in this case answer acceptance—depends on these features. We compare results of statistical trends found in aggregated and disaggregated data to identify Simpson’s pairs. Our analysis discovers several cases of Simpson’s paradox in Stack Exchange data. In addition to a known effect that describes deterioration of answerer performance over the course of a session, our method identifies several new cases of Simpson’s paradox. These paradoxes yield new insights into answerer performance on Stack Exchange. For example, creating a new feature to describe answerers, which combines variables of a Simpson’s pair, leads to a more robust proxy of answerer performance.

Presence of a Simpson’s paradox in social data can indicate interesting patterns [8], including important behavioral differences. The paradox suggests that subgroups within the population under study systematically differ in their behavior, and these differences are large enough to affect aggregate trends. In such cases the trends discovered in disaggregated data are more likely to describe—and predict—individual behavior than the trends found in aggregated data. Thus, to build more robust models of behavior, data scientists need to identify the subgroups within their data or risk drawing wrong conclusions. The method presented in this paper provides a simple framework for identifying such interesting subgroups by systematically searching for Simpson’s paradox in trend data.

The rest of the paper is organized as follows. First, we present background, including multiple observations of Simpson’s paradox in a variety of disciplines (Sec. 2). Then we describe our methodology for detecting Simpson’s paradox by identifying covariates in data, and analyze the paradox mathematically to gain more insight into its origins (Sec. 3). Finally, we apply our method to real-world data from the question-answering site Stack Exchange, and demonstrate its ability to automatically identify novel cases of Simpson’s paradox (Sec. 4). We conclude with the discussion of implications.

2 BACKGROUND AND RELATED WORK

The goal of data analysis is to identify important associations between features, or variables, in data. However, hidden correlations between variables can lead analysis to wrong conclusions. One

important manifestation of this effect is Simpson’s paradox, according to which *an association that appears in different subgroups of data may disappear, and even reverse itself, when data is aggregated across subgroups*. Instances of the paradox have been documented across a variety of disciplines, including demographics, economics, political science, and clinical research, and it has been argued that the presence of Simpson’s paradox implies that interesting patterns exist in data [8]. A notorious example of Simpson’s paradox arose during a gender bias lawsuit against UC Berkeley. Analysis of graduate school admissions data revealed a statistically significant bias against women. However, the pattern of discrimination observed in this aggregate data disappeared when admissions data was disaggregated by department. Bickel et al. [4] attributed this effect to Simpson’s paradox. They argued that the subtle correlations between the popularity of departments among the genders and their selectivity resulted in women applying to departments that were hardest to get into, which skewed analysis.

Simpson’s paradox must also be considered in the analysis of trends. Vaupel and Yashin [20] give a compelling illustration of how survivor bias can shift the composition of data, distorting the conclusions drawn from it. Analysis of recidivism among convicts released from prison shows that the rate at which they return to prison declines over time. From this, policy makers may conclude that age has a pacifying effect on crime: older convicts are less likely to commit crimes. In reality, this conclusion is false. Instead, we can think of the population of ex-convicts as composed of two subgroups with constant, but very different recidivism rates. The first subgroup, let’s call them “reformed,” will never commit a crime once released from prison. The other subgroup, the “incorrigibles,” will always commit a crime. Over time, as “incorrigibles” commit offenses and return to prison, there are fewer of them left in the population. The survivor bias changes the composition of the population under study, creating an illusion of an overall decline in recidivism rates. As Vaupel and Yashin warn, “unsuspecting researchers who are not wary of heterogeneity’s ruses may fallaciously assume that observed patterns for the population as a whole also hold on the sub-population or individual level.” Their paper gives numerous other examples of such *ecological fallacies*.

Similar illusions crop up in many studies of social behavior. For example, when examining how social media users respond to information from their friends (other users that they follow), it may appear that if *more* of a user’s friends use a hashtag then the user will be *less* likely to use it himself or herself [18]. Similarly, the more friends share some information, the less likely the user is to share it with his or her followers [10]. From this, one may conclude the additional exposures to information in a sense “innoculate” the user and act to suppress the sharing of information. In fact, this is not the case, and instead, additional exposures monotonically increase the user’s likelihood to share information with followers [16]. However, those users who follow many others, and are likely to be exposed to information or a hashtag multiple times, are less responsive overall, because they are overloaded with information they receive from all the friends they follow [12]. Calculating response as a function of the number of exposures in the aggregate data falls prey to survivor bias: the more responsive users (with fewer friends) quickly drop out of the average (since they are generally exposed

fewer times), leaving the highly connected, but less responsive, users behind. The reduced susceptibility of these highly connected users biases aggregate response, leading to wrong conclusions about individual behavior. Once data is disaggregated based on the volume of information individuals receive, a clearer pattern of response emerges, one that is more predictive of behavior [13]. Multiple examples of Simpson's paradox have been identified in empirical studies of online behavior. A study [2] of Reddit found that while it may appear that average comment length on decreases over any fixed period of time, when data is disaggregated into groups based on the year user joined Reddit, comment length within each group increases during the same time period. It is only because users who joined early tend to write longer comments that the Simpson's paradox appears.

Data heterogeneity also impacts statistical analysis of data [7] and causal inference [21]. However, no general framework to recognize and mitigate Simpson's paradox exist. Current methods require that the structure of data be explicitly specified [3] or at best be guided by subject matter knowledge [11].

Despite accumulating evidence that Simpson's paradox affects inference from data [7, 21], scientists do not routinely test for the presence of this paradox in heterogeneous data. Our work addresses this knowledge gap by proposing a statistical method to systematically uncover instances of Simpson's paradox in data.

3 METHODS

We propose a method to systematically uncover Simpson's paradox for trends in data. We denote as Y the dependent variable in the data set, i.e., an outcome being measured, and as $\mathbf{X} = \{X_1, X_2, \dots, X_m\}$ the set of m independent variables or features. The goal of the method is to identify pairs of variables (X_p, X_c) —*Simpson's pairs*—such that a trend in Y as a function of X_p disappears or reverses when the data is disaggregated by conditioning on X_c . More specifically, our method searches for pairs of variables (X_p, X_c) such that

$$\frac{d}{dx_p} \mathbb{E}[Y|X_p = x_p] > 0 \quad \forall x_p, \quad (1)$$

$$\frac{d}{dx_p} \mathbb{E}[Y|X_p = x_p, X_c = x_c] \leq 0 \quad \forall x_p, x_c. \quad (2)$$

and vice versa (i.e., $d\mathbb{E}[Y|X_p = x_p]/dx_p < 0$, $d\mathbb{E}[Y|X_p = x_p, X_c = x_c]/dx_p \geq 0$). Equations (1) and (2) hold if the expected value of Y is a monotonically increasing (or decreasing) function of X_p alone, but conditioned on X_c is a monotonically decreasing (resp. increasing) function of X_p , or is constant.

3.1 Finding Simpson's Pairs

With this goal in mind, we employ linear models to quantify the relationship between Y , an independent variable X_p , and a conditioning variable X_c upon which the data is disaggregated. Firstly, on the aggregate level, we model the relationship between Y and X_p as a linear model of the form

$$\mathbb{E}[Y|X_p = x_p] = f_p(\alpha + \beta x_p), \quad (3)$$

where $f_p(\alpha + \beta x_p)$ is a monotonically increasing function of its argument $\alpha + \beta x_p$. The parameter α in Eq. (3) is the intercept of the regression function, while the trend parameter β quantifies

the effect of X_p on Y . Secondly, for the disaggregation, we fit linear models of the form of Eq. (3) but with different values of the parameters α and β depending on the value of X_c :

$$\mathbb{E}[Y|X_p = x_p, X_c = x_c] = f_{p,c}(\alpha(x_c) + \beta(x_c)x_p), \quad (4)$$

When fitting linear models $f(\alpha + \beta X)$ we have not only a fitted trend parameter β but also a p -value which gives the probability of finding an intercept β at least as extreme as the fitted value under the null hypothesis $H_0 : \beta = 0$. From this, we have three possibilities:

- β is not statistically different from zero ($\text{sgn}(\beta) = 0$),
- β is statistically different from zero and positive ($\text{sgn}(\beta) = 1$),
- β is statistically different from zero and negative ($\text{sgn}(\beta) = -1$).

This mechanism allows us to test for Simpson's paradox by comparing the sign of β from the aggregated fit (Eq. (3)) with the signs of the β 's from the disaggregated fits (Eq. (4)). Although Eqs. (3) and (4) state that the signs from the disaggregated curves should all be different from the aggregated curve, in practice this is too strict, especially as human behavioral data is noisy. Thus, we compare the sign of the fit to aggregated data to the simple average of the signs of fits to disaggregated data—if the signs are different then we have uncovered an instance of Simpson's paradox. The summary of the algorithm is the following:

Trend Simpson's Paradox Algorithm

```

1 def trend_simpsons_pair(X, Y):
2     paradox_pairs = []
3     for paradox_var in vars:
4         beta, pvalue =
5             trend_analysis(X[paradox_var], Y)
6         agg = sgn(beta, pvalue)
7         for condition_var in vars:
8             if paradox_var != condition_var:
9                 dagg = []
10                for con_gr in bins_of(condition_var):
11                    beta, pvalue =
12                        trend_analysis(X[paradox_var |
13                            con_gr], Y)
14                    dagg.append(sgn(beta, pvalue))
15                if agg != sgn(mean(dagg)):
16                    paradox_pairs.append([paradox_var,
17                        condition_var])
18     return paradox_pairs

19 def sgn(beta, pvalue = 0.0):
20     return (0 if (beta == 0 or pvalue > 0.05)
21         else (1 if beta > 0 else -1))

```

Data Disaggregation. A critical step in our method is disaggregating data by conditioning on variable X_c . The idea behind disaggregation is to segment data into more homogeneous subgroups of similar elements. For multinomial variables X_c , disaggregation step simply involves grouping data by unique values of X_c . However, for continuous X_c or discrete variables with large range, this step is more complex. We can bin the elements according to their values of X_c , but the decision has to be made how large each bin is, whether

bin sizes scale linearly or logarithmically, etc. If the bin is too small, it may not contain enough samples for a statistically significant measurement, but if it is too large, the samples may be too heterogeneous for a robust trend. In our experiments described below, bins of fixed size successfully identify Simpson's pairs, though we realize that more sophisticated binning techniques can allow us to isolate more pairs or reduce the number of false positives.

3.2 Mathematical Analysis of the Paradox

We have presented a mathematical formulation of Simpson's paradox in terms of the derivatives of conditional expectations as given by Eqs. (1) and (2), and we now examine these equations to get a better insight into the origins and causes of this paradox.

The expectation in Eq. (1) can be related to that of Eq. (2) as

$$\mathbb{E}[Y|X_p = x_p] = \int_{X_c} \mathbb{E}[Y|X_p = x_p, X_c = x_c] \Pr(X_c = x_c|X_p = x_p) dx_c, \quad (5)$$

and differentiating this expectation w.r.t. x_p allows us to compare the trends of Eqs. (1) and (2). The derivative of the right hand side of Eq. (5) with respect to x_p is

$$\int_{X_c} \left(\frac{d}{dx_p} \mathbb{E}[Y|X_p = x_p, X_c = x_c] \right) \Pr(X_c = x_c|X_p = x_p) dx_c + \int_{X_c} \mathbb{E}[Y|X_p = x_p, X_c = x_c] \left(\frac{d}{dx_p} \Pr(X_c = x_c|X_p = x_p) \right) dx_c. \quad (6)$$

If $\mathbb{E}[Y|X_p = x_p, X_c = x_c]$ is a non-increasing function of x_p —as in Eq. (2)—then the first integral in Eq. (6) will be non-positive. Thus for $\mathbb{E}[Y|X_p = x_p]$ to be an increasing function of x_p , i.e., for Eq. (6) to be positive, the second integral must be positive.

This inequality condition leads to two necessary conditions for the occurrence of Simpson's paradox. The first condition is that

$$\frac{d}{dx_p} \Pr(X_c = x_c|X_p = x_p) \neq 0, \quad (7)$$

i.e., the distribution of the conditioning variable X_c is not independent of X_p and so the two variables are correlated. As X_p changes, the distribution of the values of X_c must also change. In the case that the distribution of X_c is independent of X_p , then $d \Pr(X_c = x_c|X_p = x_p)/dx = 0$ and so the second integral of Eq. (6) will be zero resulting in no Simpson's paradox.

The second necessary condition for the occurrence of Simpson's paradox is that the expectation of Y , conditioned on X_p , must not be independent of X_c , i.e.,

$$\mathbb{E}[Y|X_p = x_p, X_c = x_c] \neq \mathbb{E}[Y|X_p = x_p]. \quad (8)$$

For any given value of X_p , the expectation of Y must vary as a function of X_c . If the condition of Eq. (8) is not met then the second integral in Eq. (6) becomes

$$\begin{aligned} & \int_{X_c} \mathbb{E}[Y|X_p = x_p] \left(\frac{d}{dx_p} \Pr(X_c = x_c|X_p = x_p) \right) dx_c \\ &= \mathbb{E}[Y|X_p = x_p] \frac{d}{dx_p} \left(\int_{X_c} \Pr(X_c = x_c|X_p = x_p) dx_c \right) = 0, \end{aligned} \quad (9)$$

and so Simpson's paradox will not occur.

Thus, this mathematical analysis has given us an insight into causes for Simpson's paradox in data—correlations between independent variables and the fact that the distribution of the conditioning variable X_c changes at a faster rate with respect to the independent paradox variable X_p than does the expectation of Y . This point will be covered in greater detail in the next section.

4 RESULTS

We explore our approach using data from the question-answering platform called Stack Exchange. This platform, launched in 2008 to provide a forum for people to ask computer programming questions, grew over the years as a forum asking questions on a variety of technical and non-technical topics. The premise behind Stack Exchange is simple: any user can ask a question, which others may answer. Users can also *vote* for answers they find helpful, and the asker can *accept* one of the answers as the best answer to the question. In this way, the Stack Exchange community collectively curates knowledge.

4.1 Stack Exchange Data

We used anonymized data representing all questions and answers from August 2008 until September 2014.¹ The data includes 9.6M questions, of which approximately half had an accepted answer. Only the questions that received two or more answers were included [5]. Previous studies of Stack Exchange [9] and other online platforms [1, 19], identified user sessions as an important variable in understanding performance. User actions, i.e., answering questions, can be segmented into *sessions*—periods of activity without a prolonged break. Following [9] we use 100 minutes as minimum break length to define sessions. A time interval longer than 100 minutes between two answers constitutes the end of one session and the start of a new one. Note that the exact value of this threshold does not change results, but simply merge a small fraction of sessions into longer sessions.

To understand factors affecting user performance on Stack Exchange, we study the relationship between user attributes and the probability that the answer the user produces is accepted by the asker as best answer to his or her question. To this end, for each answer in the data set, we create a list of features describing the answer, as well as features of describing the user writing the answer:

Reputation: Answerer's reputation at the time he or she posted the answer. This score summarizes the user's cumulative contributions to Stack Exchange.

Number of answers: Cumulative number of answers written by the user at the time the current answer was posted.

Tenure: Age of the user's account (in seconds) at the time the user posted the answer.

Percentile: User's percentile rank based on tenure.

Time since previous answer: Time interval (in seconds) since user's previous answer.

Session length: The length of the session (in number of answers posted) during which the answer was posted.

Answer position: Index of the answer within a session.

Words: Number of words in the answer.

Lines of codes: Number of lines of codes in the answer.

¹<https://archive.org/details/stackexchange>

X_p : Independent Variable	X_c : Conditioning Variable
Tenure	Number of answers
Session length	Reputation
Answer position	Reputation
Answer position	Session length
Number of answers	Reputation
Time since previous answer	Answer position
Percentile	Number of answers

Table 1: Examples of Simpson’s paradox in Stack Exchange data. For these variables, the trend in the outcome variable (answer acceptance) as a function of X_p in the aggregate data reverses when the data is disaggregated on X_c .

URLs: Number of hyperlinks in the answer.

Readability: Answer’s Flesch Reading Ease [14] score.

4.2 Simpson’s Paradoxes on Stack Exchange

We apply the method described above to Stack Exchange data. Here, our dependent variable Y is binary, denoting whether or not a specific answer to a question was accepted as the best answer. In this case of binary outcomes we use the logistic regression linear model of the form

$$f(\alpha + \beta x) = \frac{1}{1 + e^{-(\alpha + \beta x)}}. \quad (10)$$

The parameters α and β are fitted using Maximum likelihood, while test of the null hypothesis $H_0 : \beta = 0$ is performed using the Likelihood Ratio Test [6].

The eleven variables in Stack Exchange data, result in 110 possible Simpson’s pairs. Among these, our method identifies seven as instance of paradox. These are listed in Table 1.

Our approach reveals that the previously reported finding that acceptance probability decreases with answer position [9] is an instance of Simpson’s paradox and would not have been observed had the data not been disaggregated by session length. More interestingly, our approach also identifies previously unknown instances of Simpson’s paradox. We explore these in greater detail below, illustrating how it can lead to deeper insights into online behavior.

4.2.1 Answer Position & Session Length. We measure session length by the number of answers a user posts before taking an extended break. Session length was shown to be an important confounding variable in online activity. Analysis of the quality of comments posted on a social news platform Reddit showed that, once disaggregated by the length of session, the quality of comments declines over the course of a session, with each successive comment written by a user becoming shorter, less textually complex, receiving fewer responses and a lower score from others [19]. Similarly, each successive answer posted during a session by a user on Stack Exchange is shorter, less well documented with external links and code, and less likely to be accepted by the asker as the best answer [9].

Our approach automatically identifies this example as Simpson’s paradox, as illustrated in Fig. 1. The figure shows average acceptance probability for an answer as a function of its position (or

index) within a session. According to Fig. 1a, which reports *aggregate* acceptance probability, answers written later in a session are more likely to be accepted than earlier answers. However, once the same data is *disaggregated by session length*, the trend reverses (Fig. 1b): each successive answer within the same session is less likely to be accepted than the previous answer. For example, for sessions during which five answers were written, the first answer is more likely to be accepted than the second answer, which is more likely to be accepted than the third answer, etc., which is more likely to be accepted than the fifth answer. The lines in Fig. 1 represent fits to data using logistic regression.

This example highlights the necessity to properly disaggregate data to identify the subgroups for analysis. Unless data is disaggregated, wrong conclusions may be drawn, in this case, for example, that user performance improves during a session.

4.2.2 Number of Answers & Reputation. Experience plays an important role in the quality of answers written by users. Stack Exchange veterans, i.e., users who have been active on Stack Exchange for more than six months, post longer, better documented answers, that are also more likely to be accepted as best answers by askers [9]. There are several ways to measure experience on Stack Exchange. *Reputation*, according to Stack Exchange, gauges how much the community trusts a user to post good questions and provide useful answers. While reputation can be gained or lost with different actions, a more straightforward measure of experience is user *tenure*, which measures time since the user became active on Stack Exchange, or *Percentile*, normalized rank of a user’s tenure. Alternately, experience can be measured by the *Number of Answers* a user posted during his or her tenure before writing the current answer.

Our method uncovers a novel Simpson’s paradox for user experience variables *Reputation* and *Number of Answers*. In the aggregate data, acceptance probability increases as a function of the *Number of Answers* (Fig. 2a). This is consistent with our expectations that the more experienced users—who have written more answers over their tenure on Stack Exchange—produce higher quality answers. However, when data is conditioned on *Reputation*, the trend reverses (Fig. 2b). In other words, focusing on groups of users with the same reputation, those who have written more answers over their tenure are less likely to have a new answer accepted than the less active answerers.

4.3 The Origins of Simpson’s paradox

To understand why Simpson’s paradox occurs in Stack Exchange data, we illustrate the mathematical explanation of Section 3.2 with examples from our study. Consider the paradox for *Answer Position–Session Length* Simpson’s pair, illustrated in Fig. 1a. In the disaggregated data, trend lines of acceptance probability for sessions of different length are stacked (Fig. 1b): answers produced during longer sessions are more likely to be accepted than answers

Table 2: Number of data points in each group

Session Length	1	2	3	4	5	6	7	8
Data points	7.2M	2.6M	1.3M	0.7M	0.4M	0.3M	0.2M	0.1M

Figure 1: Simpson's paradox in Stack Exchange data. Both plots show the probability an answer is accepted as the best answer to a question as a function of its position within user's activity session. (a) Acceptance probability calculated over aggregated data has an upward trend, suggesting that answers written later in a session are more likely to be accepted as best answers. However, when data is disaggregated by session length (b), the trend reverses. Among answers produced during sessions of the same length (different colors represent different-length sessions), later answers are less likely to be accepted as best answers.

Figure 2: Novel Simpson's paradox discovered in Stack Exchange data. Plots show the probability an answer is accepted as best answer as a function of the number of lifetime answers written by user over his or her tenure. (a) Acceptance probability calculated over aggregated data has an upward trend, with answers written by more experienced users (who have already posted more answers) more likely to be accepted as best answers. However, when data is disaggregated by reputation (b), the trend reverses. Among answers written by users with the same reputation (different colors represent reputation bins), those posted by users who had already written more answers are less likely to be accepted as best answers.

produced during shorter sessions. In addition, there are many more shorter sessions than longer ones. Table 2 reports the number of sessions of different length. By far, the most common session has length one: users write only one answer during these sessions. Each longer session is about half as common as a session that is one answer shorter.

What happens to the trend in the aggregated data? When calculating acceptance probability as a function of answer position, all sessions contribute to acceptance probability for the first answer of a session. Sessions of length one dominate the average. When calculating acceptance probability for answers in the second position, sessions of length one do not contribute, and acceptance probability

Figure 3: Analysis of the Simpson's paradox Reputation – Number of Answers variable pair. (a) Average acceptance probability as a function of two variables. (b) The distribution of the number of data points contributing to the value of the outcome variable for each pair of variable values.

is dominated by data from sessions of length two. Similarly, acceptance probability of answers in the third position is dominated by sessions of length three. Survivor bias excludes data from shorter sessions, which also have lower acceptance probability, creating an upward trend in acceptance probability.

We back up this intuitive explanation with mathematical analysis of Section 3.2. Although acceptance probability is decreasing as a function of *Answer Position* for each value of *Session Length* (Fig. 1b), the probability mass of *Session Length* is constantly moving towards larger values as *Answer Position* increases. Notice that as *Answer Position* increments from a to $a + 1$, sessions of length a are no longer included (as the minimum session length is now $a + 1$). Thus, while *Session Length* has probability mass $\Pr(X_c = a | X_p = a)$ when $X_p = a$, it has probability $\Pr(X_c = a | X_p = a + 1) = 0$ at $X_p = a + 1$:

$$\frac{d}{dx_p} \Pr(X_c = a | X_p = x_p) \big|_{x_p=a} = -\Pr(X_c = a | X_p = a). \quad (11)$$

Meanwhile, for all other values of X_c greater than a , the probability mass at $X_p = a + 1$ is the same as that at $X_p = a$ (as the number of data points is constant along sessions of same length) but normalized to account for the sessions of length a , i.e.,

$$\Pr(X_c = x_c | X_p = a + 1) = \frac{\Pr(X_c = x_c | X_p = a)}{1 - \Pr(X_c = a | X_p = a)}. \quad (12)$$

The rate of change of these probability masses with respect to X_p is

$$\frac{d}{dx_p} \Pr(X_c = x_c | X_p = x_p) \big|_{x_p=a} = \left(\frac{1}{1 - \Pr(X_c = a | X_p = a)} - 1 \right) \Pr(X_c = x_c | X_p = a). \quad (13)$$

The probability mass function $\Pr(X_c = x_c | X_p = a)$ decreases for $X_c = a$ corresponding to the smallest value of acceptance probability, while increasing for all other values $X_c > a$. Moreover, the rate of increase of this probability mass is greater than the rate at

which the acceptance probability decreases, resulting in an upward trend when the data is aggregated.

A similar effect plays out in the *Number of Answers–Reputation* Simpson's pair. Figure 3a shows the heatmap of acceptance probability for different values of the *Number of Answers* written over a user's tenure and user *Reputation*, while Fig. 3b shows the correlated joint distribution of the two variables. The figures illustrate the first condition of Simpson's paradox (Eq. (7)): as X_p changes, the distribution of the values of X_c must also change. This dependency can be clearly seen in Fig. 3b—as $X_p = \text{Reputation}$ increases then the distribution of $X_c = \text{Number of Answers}$ shifts to increasing values, which produces the paradox.

In the real world this means that users, who have written more answers are not more likely to have a new answer they write accepted. In fact, among users with same *Reputation*, those who earned this reputation with fewer answers are more likely to have a new answer they write accepted as best answer. This suggests that such users are simply better at answering questions, and that this can be detected early in their tenure on Stack Exchange (while they still have low reputation). Note, however, that an exception to the trend reversal occurs for users with very high reputation. In Stack Exchange, users can gain reputation by "Answer is marked accepted", "Answer is voted up", "Question is voted up", etc. It seems that, high reputation users and low reputation users are different: for high reputation users, experience (number of written answers) is important, while for low reputation users the quality of answers, which may lead to votes, is more important. Analysis of this behavior is beyond the scope of this paper.

4.4 Discussion and Implications

Presence of a Simpson's paradox in data can indicate interesting or surprising patterns [8], and for trends in social data, important behavioral differences within a population. Since social data is often generated by a mixture of subgroups, existence of Simpson's

Figure 4: Relationship between acceptance probability and Reputation Rate, a new measure of user performance defined as reputation per number of answers users wrote over their entire tenure. Each line represents a subgroup with a different reputation score. The much smaller variance compared to Fig. 2b suggests that the new feature is a good proxy of answerer performance.

paradox suggests that these subgroups differ systematically and significantly in their behavior. By isolating important subgroups in social data, our method can yield insights into their behaviors.

For example, our method identifies *Session Length* as a conditioning variable for disaggregating data when studying trends in acceptance probability as a function of answer’s position within a session. In fact, prior work has identified session length as an important parameter in studies of online performance [1, 9, 15, 19]. Unless activity data is disaggregated into individual sessions—sequences of activity without an extended break—important patterns are obscured. A pervasive pattern in online platforms is user performance deterioration, whereby the quality of a user’s contribution decreases over the course of a single session. This deterioration was observed for the quality of answers written on Stack Exchange [9], comments posted on Reddit [19], and the time spent reading posts on Facebook [15]. Our method automatically identifies position of an action within a session and session length as an important pair of variables describing Stack Exchange.

We examine in detail one novel paradox discovered by our method for the *Reputation–Number of Answers* variables. The trends in Fig. 2b suggest that both variables jointly affect acceptance probability. Inspired by this observation, we construct a new variable—*Reputation / Number of Answers*—i.e., *Reputation Rate*. Figure 4 shows how acceptance probability changes with respect to *Reputation Rate* for different groups of users. There is an strong upward trend, suggesting that answers provided by users with higher *Reputation Rate* are more likely to be accepted. Moreover, while the lines span reputations of an extremely broad range—from one to 100,000—they collapse onto a single curve. This suggests that *Reputation Rate* is a good proxy of user performance. The remaining paradoxes uncovered by our method could yield similarly interesting insights into user behavior on Stack Exchange.

We also illustrate the difference between our method and linear models that model the outcome variable as a function of both X_p and X_c . For such multivariate linear models [17], we can fit a model $f_{p,c}(\alpha + \beta X_p + \beta_c X_c)$ to the disaggregated data, and compare the sign of the coefficient β to the sign of the linear coefficient of the “aggregated” model $f_p(\alpha + \beta X_p)$. In our method, we bin the values of X_c and fit separate linear models of the form of Eq. (4) in each bin of X_c , aggregating by averaging the linear coefficient signs of each model. We claim that our approach has benefits over multivariate linear models which allow it to find Simpson’s pairs where multivariate linear models can not. First, in multivariate linear models, all subgroups have the same coefficient β , and intercepts $\alpha + \beta_c X_c$, which vary linearly with X_c . In our method, however, each group can have different intercept and coefficient, which makes finding paradox pairs in heterogeneous data more flexible. Indeed this flexibility is necessary — from our results (Figs. 1b and 2b) it is clear that the trend parameters $\beta(x_c)$ of the fitted lines vary significantly depending on x_c .

Secondly, our method of aggregating by simple averaging of the linear coefficient signs of the subgroups means that trends within each subgroup are weighted equally regardless of how many datapoints are in that subgroup. This is contrary to multivariate linear models, which fit the model parameters based on each datapoint (and so weigh heavily towards values of X_c with many datapoints). To illustrate, we show that our algorithm finds *Time Since Previous Answer – Answer Position* as a Simpson’s pair, which a multivariate logistic regression does not. The variable *Answer Position* is the index of the answer a user has completed without an extended (>100 minute) break, and so *Answer Position* = 1 if *Time Since Previous Answer* ≥ 100 minutes and *Answer Position* > 1 if *Time Since Previous Answer* < 100 minutes. Fig. (5a) shows that, for *Answer Position* = 1, the acceptance probability decreases as a function of *Time Since Previous Answer*, possibly because better users take shorter breaks. On the other hand, for other *Answer Positions* the trend is reversed, and acceptance probability increases with *Time Since Previous Answer*, suggesting that in short term, users who take more time to answer questions or take short breaks between questions write answers of higher quality.

Clearly, *Time Since Previous Answer – Answer Position* is an important Simpson’s pair, illustrating that time has a beneficial effect on answer quality at short time scales. even though it is detrimental on the aggregate level. Multivariate logistic regression does not capture this behaviour, as 65% of the probability mass of *Time Since Previous Answer* is for values larger than 100 minutes, so when fitting $f_{p,c}(\alpha + \beta X_p + \beta_c X_c)$ to the data, it tries to fit a hyperplane, which describes the majority of the data as best as possible, in this case the decreasing trend corresponding to *Answer Position* = 1.

5 CONCLUSION

We presented a method for systematically uncovering instances of Simpson’s paradox in data. The method identifies pairs of variables, such that a trend in an outcome as a function of one variable disappears or reverses itself when the same data is disaggregated by conditioning it on the second variable. The disaggregated data corresponds to subgroups within the population generating the data. Our mathematical analysis suggests that Simpson’s paradox

Figure 5: A pair which multivariate logistic regression cannot find in the data. (a) Average acceptance probability as a function of Answer Position and Time Since Previous Answer. (b) The distribution of the number of data points contributing to the value of the outcome variable for each pair of variable values.

is caused by both correlations between independent variables in data (Figs. 3b and 5b), as well as differing behaviour of the outcome variable within subgroups, illustrated here by the stacked curves of Figs. 1b and 2b. Failure to account for this effect can lead analysis to wrong conclusions about typical behavior of individuals.

We applied our method to real-world data from the question-answering site Stack Exchange. We were specifically interested in uncovering features affecting the probability that an answer written by a user will be accepted by the asker as the best answer to his or her question. We identified eleven relevant features of answers and users. Not only did the method confirm an existing paradox, but it also uncovered new instances of Simpson’s paradox.

Our work opens several directions for future work. The proposed algorithm could benefit from a more principled method to bin continuous data and more sophisticated techniques for re-aggregating the intercepts of the curves fitted to disaggregated data. Also, while it appears that conditioning on X_c disaggregates the population into more homogeneous subgroups, we have not used formal methods, such as goodness of fit, to test for better fit of regression models to data. Goodness of fit may also be used to guide data disaggregation strategies. In addition, our method applies to explicitly declared variables, and not to latent variables that may affect data. While these and similar questions remain, our proposed method offers a promising tool for the analysis of heterogeneous social data.

Acknowledgments

We acknowledge funding support from the James S. McDonnell Foundation, Air Force Office of Scientific Research (FA9550-17-1-0327), and by the Army Research Office (W911NF-15-1-0142).

REFERENCES

- [1] Tushar Agarwal, Keith Burghardt, and Kristina Lerman. 2017. On Quitting: Performance and Practice in Online Game Play. In *Proceedings of 11th AAAI International Conference on Web and Social Media*. AAAI.
- [2] Samuel Barbosa, Dan Cosley, Amit Sharma, and Roberto M. Cesar-Jr. 2016. Averaging Gone Wrong: Using Time-Aware Analyses to Better Understand Behavior. In *WWW*, 829–841.
- [3] Elias Bareinboim and Judea Pearl. 2016. Causal inference and the data-fusion problem. *PNAS* 113, 27 (2016), 7345–7352.
- [4] P. J. Bickel, E. A. Hammel, and O’connell J.W. 1975. Sex bias in graduate admissions: Data from Berkeley. *Science* 187, 4175 (1975), 398–404.
- [5] Keith Burghardt, Emanuel F. Alsina, Michelle Girvan, William Rand, and Kristina Lerman. 2017. The myopia of crowds: Cognitive load and collective evaluation of answers on Stack Exchange. *PLOS ONE* 12, 3 (16 March 2017), e0173610+.
- [6] George Casella and Roger L Berger. 2002. *Statistical inference*. Vol. 2. Duxbury Pacific Grove, CA.
- [7] William K. Estes. 1956. The problem of inference from curves based on group data. *Psychological bulletin* 53, 2 (1956), 134.
- [8] Carem C Fabris and Alex A Freitas. 2000. Discovering surprising patterns by detecting occurrences of Simpson’s paradox. In *Research and Development in Intelligent Systems XVI*. Springer, 148–160.
- [9] Emilio Ferrara, Nazanin Alipoufard, Keith Burghardt, Chiranth Gopal, and Kristina Lerman. 2017. Dynamics of Content Quality in Collaborative Knowledge Production. In *ICWSM*. AAAI.
- [10] Greg Ver Steeg, Rumi Ghosh, and Kristina Lerman. 2011. What stops social epidemics?. In *ICWSM*.
- [11] Miguel A. Hernan, David Clayton, and Niels Keiding. 2011. Simpson’s paradox unraveled. *Int J Epidemiology* 40, 3 (2011), 780–785.
- [12] Nathan O. Hodas and Kristina Lerman. 2012. How Limited Visibility and Divided Attention Constrain Social Contagion. In *SocialCom*.
- [13] Nathan O. Hodas and Kristina Lerman. 2014. The Simple Rules of Social Contagion. *Scientific Reports* 4 (2014).
- [14] J. P. Kincaid, Jr R. P. Fishburne, R. L. Rogers, and B. S. Chissom. 1975. *Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel*. Technical Report. U.S. Navy.
- [15] Farshad Kooti, Karthik Subbian, Winter Mason, Lada Adamic, and Kristina Lerman. 2017. Understanding Short-term Changes in Online Activity Sessions. In *Companion WWW2017*.
- [16] Kristina Lerman. 2016. Information Is Not a Virus, and Other Consequences of Human Cognitive Limits. *Future Internet* 8, 2 (May 2016), 21+.
- [17] H. James Norton and George Divine. 2015. Simpson’s paradox and how to avoid it. *Significance* 12, 4 (1 Aug. 2015), 40–43.
- [18] Daniel M. Romero, Brendan Meeder, and Jon Kleinberg. 2011. Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. In *WWW*.
- [19] Philipp Singer, Emilio Ferrara, Farshad Kooti, Markus Strohmaier, and Kristina Lerman. 2016. Evidence of Online Performance Deterioration in User Sessions on Reddit. *PLoS ONE* 11, 8 (2016), e0161636+.
- [20] J. W. Vaupel and A. I. Yashin. 1985. Heterogeneity’s ruses: some surprising effects of selection on population dynamics. *Am. Stat.* 39, 3 (1985), 176–185.
- [21] Yu Xie. 2013. Population heterogeneity and causal inference. *PNAS* 110, 16 (2013), 6262–6268.