User Sentiment as a Success Metric: Persistent Biases Under Full Randomization

Ercan Yildiz*
meyildiz@gmail.com
Google, Inc
Mountain View, CA

Joshua Safyan
jsafyan@gmail.com
Google, Inc
Mountain View, CA

Marc Harper
marc.harper@gmail.com
Google, Inc
Mountain View, CA

ABSTRACT

We study user sentiment (reported via optional surveys) as a metric for fully randomized A/B tests. Both user-level covariates and treatment assignment can impact response propensity. We show that a simple mean comparison produces biased population level estimates and propose a set of consistent estimators for the average and local treatment effects on treated and respondent users. We show that our problem can be mapped onto the intersection of the missing data problem and observational causal inference, and we identify conditions under which consistent estimators exist. Finally, we evaluate the performance of estimators and find that more complicated models do not necessarily provide superior performance as long as models satisfy consistency criteria.

KEYWORDS

A/B test; causal inference; metrics; sentiment metric; experiment; online product

ACM Reference Format:

Ercan Yildiz, Joshua Safyan, and Marc Harper. 2020. User Sentiment as a Success Metric: Persistent Biases Under Full Randomization. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20), August 23–27, 2020, Virtual Event, CA, USA*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394486.3403340>

1 INTRODUCTION

Experimentation via A/B testing is widespread practice especially among technology companies [14, 25, 28]. A typical A/B test splits a population in two or more distinct groups at random, exposing each to different experiences. The experiment arm is then evaluated by a set of predefined metrics [25]. Engagement-based metrics are frequently used, and the entire process can be automated either by in-house infrastructure [25, 28] or by third party services [19].

Organizations have also measured user sentiment and attitudes at scale since the late 1980s [2, 7, 8], mainly driven by the relationship between consumer attitudes, long term loyalty, and revenue. Sentiment changes in an undesirable direction have also been shown to result in decreased engagement and lower long-term retention [8].

*All authors contributed equally to this research.

This work is licensed under a Creative Commons Attribution International 4.0 License.

KDD '20, August 23–27, 2020, Virtual Event, CA, USA
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7998-4/20/08.
<https://doi.org/10.1145/3394486.3403340>

Improvements in survey technology have enabled faster and more targeted feedback. Historically, phone and mail surveys as well as in-person studies were the primary source of consumer data. Email-based surveys brought increased reach and decreased the cost of collecting user feedback. More recently, companies have brought surveys into UIs and started capturing user sentiment during the experience [17]. An example of a UI-based survey is given in Figure 1.

Figure 1: Sample UI based survey

In recent years, companies have begun merging experimentation and user sentiment by utilizing sentiment as a success metric for product launches. This has become possible due to three major advantages of UI-based surveys over traditional survey mediums 1) mapping survey responses to treatment exposure is straightforward as surveys are attached to experiences, 2) capturing user sentiment in close temporal-proximity to the actual user experience increases validity of responses, 3) UI surveys have the potential to reach many users and have higher response rates compared to traditional channels [17].

Employing sentiment-based metrics in an A/B setting has its own challenges. First, unlike a typical engagement metric, sentiment data is incomplete in most cases. In particular, we have the ability to observe levels of engagement across all users in an A/B test¹ while sentiment is measured through optional surveys to which

¹Clearly, some users may not engage, though this is an observation of no engagement and differs from survey non-response.

not all users respond. Second, response behavior may be correlated with sentiment as well as treatment exposure. As a result, naive comparisons of sentiment-based metrics risk biased conclusions. This is troubling since the experiment platforms utilized by the majority of companies use naive mean comparisons to estimate treatment impact.

Accordingly, we focus on a scenario where user sentiment reported via optional surveys is used as a success metric for a fully randomized A/B test. There exists a set of user-level confounders which might impact the sentiment as well as the propensity to respond to the survey. Moreover, we allow treatment/control assignment to impact the probability of response. Our goal is to estimate both the population-level average treatment effect (ATE), and the local treatment effect on users that were treated and responded to the survey (ATETR).

2 RELATED WORK AND OUR CONTRIBUTION

Our setting relates to the well-known missing data and survey non-response problems [4, 11, 12]. Missing data literature focuses on the problem of estimating a population statistic from an incomplete dataset [3, 13, 15, 22]. In this scenario, there is no experimentation; only a single population with missing observations [13]. Typically, the data is assumed to be missing at random (MAR), i.e., missing data can be fully explained by a set of observed variables [22]. Under the MAR assumption, a rich set of consistent estimators have been proposed including outcome regression, inverse propensity weighing, and various doubly-robust methods [3, 13]. Though related, our problem departs from the traditional missing data problem because we have two distinct groups that are exposed to different experiences which impact both outcomes and the likelihood of missing data.

Our problem also intersects with causal inference, where one is interested in estimating the effect of a treatment using observational and/or quasi-experiment data [3, 5, 9, 10, 20, 21]. In the absence of a randomized controlled trial (RCT), the relationship between the treatment assignment and the outcome of interest may be confounded. The causal inference literature includes a wide range of methods to adjust for confounding in treated and non-treated populations based on a set of observed variables [3, 5, 20, 21]. While we rely heavily on causal inference, our setting also deviates from a typical causal problem. In particular, treatment assignment in our case is fully randomized, so there is no treatment selection bias. However, since responses are optional, there is still non-response bias.

We propose a framework at the intersection of missing data and causal inference to adjust for differential response bias in surveys in the A/B setting. We map ATE estimation into a missing data problem across two (or more) populations, and ATETR estimation into a causal inference problem by leveraging methods for estimating the average treatment effect on the treated (ATET). Additionally, we provide a simulation study using existing models to demonstrate a stylized application of our method. Finally, in the case of ATETR, we argue that a simple regression based covariate control may provide comparable performance to more complicated and resource heavy techniques.

3 PROBLEM SETUP

Consider a scenario where user sentiment is used to evaluate a treatment. For the sake of simplicity, we assume that sentiment is collected via a single question². For instance, if we are interested in user satisfaction we may ask *How satisfied are you with the overall experience?* A user might pick one of the following { *Extremely Satisfied*, *Satisfied*, *Neither Satisfied Nor Dissatisfied*, *Dissatisfied*, *Extremely Dissatisfied* } [1, 26]. It is common practice to code favorable responses higher, e.g, "Extremely Satisfied" is commonly coded as 5 and "Extremely Dissatisfied" is coded as 1. Additionally, this scale can be simplified using a Top 2 Box score, i.e, {4, 5} (*Satisfied*, *Extremely Satisfied*) are mapped to 1 and {1, 2, 3} are mapped to 0 [16, 18]. We focus on Top 2 Box binary sentiment measures, though our proposed approach generalizes across different sentiment measures and coding practices. Furthermore, we restrict to the case where survey response is optional, which differentiates our problem from a typical RCT³.

Assume that there are N individuals in a population, and the population is randomly split into treatment and control groups. We will assume that all of the individuals in the population are exposed to either a treatment or a control condition. T_i is a dummy variable indicating whether the individual i is exposed to treatment or control, i.e, $T_i = 1$, if individual i is exposed to treatment, and $T_i = 0$ if individual i is exposed to control. ($Y_i(1)$, $Y_i(0)$) captures the (scalar) sentiment measure for individual i under treatment and control conditions, respectively. Finally, ($\Delta_i(1)$, $\Delta_i(0)$) are indicators for survey response behavior under treatment and control conditions. In other words, for a user assigned to condition T_i , $\Delta_i(T_i)$ will be equal to 1 if the user responds to the survey.

We are interested in estimating two distinct quantities: the average treatment effect on the overall population (ATE), i.e:

$$E[Y(1) - Y(0)], \quad (1)$$

and the average treatment effect on individuals who are exposed to the treatment and respond to the survey (Average Treatment Effect on Treated and Responded - ATETR), i.e:

$$E[Y(1) - Y(0) | T = 1, \Delta(T) = 1]. \quad (2)$$

The former resembles an engagement metric, which is observed across all individuals, by aiming to provide population level impact. Given the wide adoption of engagement-based metrics for product evaluations, comparability and consistency with such metrics is an important factor. However, stricter assumptions are required to extrapolate from the observed responses to the overall population. Such assumptions may be problematic given that average online surveys have a response rate lower than 10% [6]. While the latter measure does not require estimating the sentiment of survey non-respondents, it is more difficult to compare to engagement type metrics since the ATETR estimates a local effect rather than a population effect.

Going forward, we utilize two assumptions depending on the measure of interest.

²It is straightforward to expand our results to the case where sentiment is measured by a series of questions.

³If we could mandate survey response we would simply compare the averages of the treatment and control groups to calculate the average treatment effect.

ASSUMPTION 1 (MISSING AT RANDOM). *There exists a set of observed variables $X_i(T_i)$, $T_i \in \{0, 1\}$ such that:*

$$\Delta_i(T_i) \perp Y_i(T_i) \mid X_i(T_i), \quad (3)$$

$$P(\Delta_i(T_i) = 1 \mid X_i(T_i)) > 0 \quad (4)$$

$$\forall i \in \mathcal{N} = \{1, \dots, N\}.$$

The first equation of Assumption 1 is known as Missing At Random (MAR), and says that systematic differences between the missing and observed values can be entirely explained by a set of observed variables [13, 15, 22]. The second equation of the assumption requires that the probability of observing complete data is always positive for a given individual i . Since $X_i(T_i)$ is a function of the assignment, i.e., T_i , it is possible that different sets of variables can capture confoundedness in treatment and control groups. We present the causal graph representing Assumption 1 in Figure 2. The equivalent of the causal graph for $T = 0$ can be constructed simply by replacing (1)'s with (0)'s.

Figure 2: Causal graph representing Assumption 1. $X(T)$ influences both user sentiment $Y(T)$ and propensity to respond $\Delta(T)$. There is no other path between user sentiment and propensity to respond. Therefore, $Y(T)$ and $\Delta(T)$ are independent conditioned on $X(T)$.

ASSUMPTION 2. *There exists a set of observed variables X_i such that:*

$$(\Delta_i(0), \Delta_i(1)) \perp Y_i(0) \mid X_i \quad (5)$$

$$P(\Delta_i(0) = 1 \mid X_i) > 0, P(\Delta_i(1) = 1 \mid X_i) > 0 \quad (6)$$

The first equation of Assumption 2 requires the joint survey response behaviors to be independent of the sentiment under the control condition given the confounders. We require this assumption to estimate the ATETR, $E[Y(1) - Y(0) \mid T = 1, \Delta(T) = 1]$. Specifically, we want to estimate how those who were assigned to the treatment group and responded to the survey ($Y_i(0), \Delta_i(0) = 1$) would have responded under control. Note that Assumption 2 does not impose any constraints on the conditional independence of $Y_i(1)$ since we observe the responses from the users who were treated and responded to the survey, thus we only need to estimate the counterfactual sentiment for treated and respondent users under control.

In practice, one may be unable to adequately model the selection bias to adjust for factors influencing both sentiment and response behavior. A sensitivity analysis [24] should be performed to assess the susceptibility of one's estimates to violations of these assumptions.

We show the causal graph representing Assumption 2 in Figure 3.

Figure 3: Causal graph representing Assumption 2. X influences both user sentiment under control $Y(T = 0)$, and propensity to respond under both treatment and control conditions $\Delta(T = 0), \Delta(T = 1)$. There is no other path between user sentiment and propensity to respond. Therefore, $Y(T = 0)$ and $\Delta(T = 0), \Delta(T = 1)$, are independent conditioned on X .

4 ESTIMATION

We propose mechanisms to estimate the measures of interest: ATE and ATETR.

4.1 Sentiment Average Treatment Effect (ATE)

We are interested in estimating the Average Treatment Effect on user sentiment over the entire population. Our first results enables parallelization of estimation of ATE by mapping it onto two missing data problems:

LEMMA 4.1. *Let θ_1 denote a consistent estimator of $E[Y(1) \mid T = 1]$, and θ_0 denote a consistent estimator of $E[Y(0) \mid T = 0]$. Then, $\theta_1 - \theta_0$ is a consistent estimator of ATE.*

The proof of the lemma is straightforward, and follows from the linearity of expectation, linearity of convergence in probability, and randomized treatment assignment i.e., $E[Y(1) \mid T = 1] = E[Y(1) \mid T = 0] = E[Y(1)]$. Lemma 4.1 suggests that one can focus on the treatment and control groups separately and treat each case as an independent missing data problem. Independence results from the fully randomized treatment/control assignment even though survey measures suffer from response bias. Lemma 4.1 is powerful since it enables parallelization of the estimation.

While we have mapped our ATE estimation into two independent missing data problems, we need further assumptions in order to provide consistent estimators, i.e., θ_1, θ_0 . In particular, we rely on Assumption 1 (also known as MAR [3]) which guarantees consistency for a wide range of estimators including outcome regression

(OR), inverse propensity weighting (IPW), and adversarial balancing (AB) estimators⁴.

As discussed in detail in [3], Assumption 1 enables us to infer sentiment measures for users who have not responded to the survey in the treatment (control) group by simply utilizing the responses of the users who have responded to the survey in the treatment (control) group.

As an example, we focus on the OR estimator and demonstrate how θ_1, θ_0 can be calculated. We start with a set of equalities:

$$\begin{aligned} E[Y(1)] &= E_X[E_Y[Y(1)|X]] \\ &= E_X[E_Y[Y(1)|X, T = 1, \Delta(1) = 1]] \end{aligned} \quad (7)$$

where the first line follows from law of iterated expectation, and the second from Assumption 1 and the fact that treatment assignment is random. Equation (7) shows that building an estimator $E[Y_i(1)|X_i] = f(X_i)$ trained on observed responses only ($\{i : T_i = 1, \Delta_i(1) = 1\}$) provides a consistent estimator for $E[Y(1)]$ provided that $f(\cdot)$ is correctly specified.

$$\theta_1 = |\{i : T(i) = 1\}|^{-1} \sum_{i:T(i)=1} f(X_i)$$

is the estimator of interest, where $|\cdot|$ is the cardinality of the set. θ_0 can be calculated in a similar fashion. Interested readers can refer to [3] for detailed discussions.

We implemented several estimators (including outcome regression) for comparison purposes, and discuss performance in Section 5.

4.2 Sentiment Average Treatment Effect on Treated and Responded (ATETR)

Now we estimate the Average Treatment Effect on Treated and Responded given in Equation (2).

Similar to its counterpart in the causal inference literature, *i.e.*, Average Treatment Effect on Treated (ATET), ATETR requires reconstructing counterfactuals for a subset of users. In our case, the subset is those who were assigned to the treatment group and chose to respond to the survey.

There is a consistent estimator for ATETR, defined as follows:

LEMMA 4.2. *Given the definition of ATETR in Equation (2) and Assumption 2,*

$$\begin{aligned} ATETR &= E[Y(1) - Y(0)|T = 1, \Delta(T) = 1] \\ &= E[Y(1)|T = 1, \Delta(T) = 1] \end{aligned} \quad (8)$$

$$- E[Y(0)|T = 1, \Delta(T) = 1] \quad (9)$$

Moreover, $|\{i : T_i = 1, \Delta_i(1) = 1\}|^{-1} \sum_{i:T(i)=1, \Delta(i)=1} Y_i(1)$ is a consistent estimator of Equation (8), and Equation (9) is equal to:

$$E_X[E_{X, \Delta(1)=1}[E_Y[Y(0)|\Delta(0) = 1, X]|\Delta(1) = 1, X]]. \quad (10)$$

Proof of Lemma 4.2 is given in Appendix A.

As an example, we will once again focus on the OR estimator and demonstrate how ATETR can be calculated.

Equation (8) can simply be estimated by taking the average sentiment among users that were assigned to the treatment group and responded to the survey ($\{i : T_i = 1, \Delta_i(1) = 1\}$). Equation (9) (or equivalently Equation (10)) can be estimated by a 3 step process:

- (1) Focus on survey respondents under the control assignment, *i.e.*, $\{i : T_i = 0, \Delta_i(0) = 1\}$, and train a model for $E[Y_i(0)|X_i] = g(X_i)$.
- (2) Construct counterfactual responses for users that were assigned to the treatment and responded to the survey, $E[Y_i(0)|T_i = 1, \Delta_i(1)]$, $\forall i \in \{i : T_i = 1, \Delta_i(1) = 1\}$.
- (3) Estimate $E[Y(1)|T = 1, \Delta(1) = 1]$ using $|\{i : T(i) = 1, \Delta_i(1) = 1\}|^{-1} \sum_{i:T_i=1} g(X_i)$.

In other words, training a model on the control group and using the trained model to predict counterfactuals over the treatment group provides a consistent estimator. Finally, the difference between the two provides a consistent estimator for ATETR⁵.

Similar to ATE, we implemented several estimators (including outcome regression) for comparison purposes, and discuss performance in Section 5.

5 SIMULATIONS

We applied our methodology to a simulated dataset. A causal graph representing the generative process is given in Figure 4. Latent variables ($X = [X_1, X_2]$) generate both individual sentiment Y and response behavior under the treatment and control conditions $\Delta(1), \Delta(0)$. User sentiment Y is the same across treatment/control groups. However, response behaviors $\Delta(0), \Delta(1)$ differ under treatment vs. control and correlate with sentiment, so observed sentiment differs.

Figure 4: Causal graph representing our data generation process. Latent variables ($X = [X_1, X_2]$) generate both individual sentiment Y and response behavior under the treatment and control conditions $\Delta(1), \Delta(0)$.

Details of the data generation process (DGP) are given in Appendix B. There are 4 important points regarding the DGP:

- (1) Both Assumptions 1 and 2 hold.
- (2) There is no difference in the population level sentiment, $E[Y(1)] = E[Y(0)]$.
- (3) There are differences in the observed sentiment: $E[Y(1)|\Delta(1) = 1] - E[Y(0)|\Delta(0) = 1] \approx 0.07$.
- (4) Response rates under the treatment and control are also different: $E[\Delta(1)] \approx 0.15$, and $E[\Delta(0)] \approx 0.09$. These numbers

⁴Additional assumptions for model specifications are required for parametric models.

⁵Additional assumptions for model specifications are required.

Metric	Value
$E[\Delta(1)]$	0.15
$E[\Delta(0)]$	0.09
$E[Y(1) - Y(0)]$	0
$E[Y(1) - Y(0) T = 1, \Delta(1) = 1]$	0
$E[Y(1) \Delta(1) = 1] - E[Y(0) \Delta(0) = 1]$	0.07

Table 1: Population level statistics based on DGP in Appendix B.

correspond to 15% and 9% survey response rates, which are consistent with lower response rates typically observed for online surveys.

We run 5K Monte Carlo simulations with $N = 10K$ ($\sim 5K$ each in treatment and control). 10K is used as the per simulation population size to reflect the small sample sizes typically observed during sentiment studies. Summary statistics are available in Table 1.

5.1 Sentiment Average Treatment Effect (ATE)

We implemented four different estimators and compared their performance under two different scenarios.

In scenario 1, we assume that we can observe the actual values of the confounding factors, *i.e.*, X_1, X_2 . This allows us to correctly specify a model to adjust for all of the confounding variables, and Assumption 1 holds.

Our first estimator is a naive mean comparison among respondents. By definition this is an inconsistent estimator and does not take into account observed confounders. The remaining estimators are based on Lemma 4.1, and use the observed covariates to correct for the underlying bias. We implemented outcome regression (OR) [3], inverse propensity weighing (IPW) [13], a doubly robust (DR) estimator based on OR and IPW [3], and an adversarial balancing based estimator (AB) using Logistic Regression as the classifier [20].

We measured the performance of each estimator via bias, median absolute error (MAE), and mean squared error (MSE)⁶ as discussed in [13]. Results are given in Table 2. See Figure 5 for the distribution of estimates.

Naive comparison incorrectly reports a 0.07 (7 percentage points) difference between treatment and control. Variation in the naive comparison is due to the stochasticity of the simulation. The OR estimator has the smallest bias, despite the fact that it suffers from higher variance. IPW has a large bias comparable to the naive estimator. The DR estimator's performance falls in between OR and IPW.

In scenario 2, we relax our assumption on the ability to observe the true confounders. Instead, we assume that we can only observe transformations of the true confounders as given in Eq. (15)-(16). This corresponds to a case where the researcher cannot directly observe X , but can still access a set of variables that are related to the confounders themselves (and thus will likely misspecify the

⁶Bias = $E[ATE - \hat{ATE}]$, MAE = $Median(|ATE - \hat{ATE}|)$, and MSE = $E[(ATE - \hat{ATE})^2]$

Figure 5: Performance of different ATE estimators when true confounders are fully observed. DR and AB has the highest variance.

model). Results and boxplots are given in Table 3 and Figure 6, respectively.

Figure 6: Performance of different ATE estimators when noisy confounders are observed. Increase in variances of OR and DR, AB retains performance characteristics.

Our first observation is the dramatic drop in the performance of OR. The estimator performs the worst in both bias and MSE.

Estimator	Bias	MAE	MSE
AB	2.8×10^{-2}	2.9×10^{-2}	1.5×10^{-3}
DR	-2.3×10^{-3}	6.1×10^{-2}	8.0×10^{-3}
IPW	5.7×10^{-2}	5.8×10^{-2}	6.2×10^{-3}
Naive comparison	7.1×10^{-2}	7.1×10^{-2}	5.8×10^{-3}
OR	8.8×10^{-4}	6.1×10^{-2}	8.1×10^{-3}

Table 2: Performance of different ATE estimators when true confounders are fully observed. Estimators are Adversarial Balancing (AB), Doubly Robust (DR), Inverse Propensity Weighing (IPW), Naive mean comparison and Outcome Regression (OR). The naive estimator has the largest bias while AB estimator has the best MSE performance.

Estimator	Bias	MAE	MSE
AB	3.4×10^{-2}	3.4×10^{-2}	1.8×10^{-3}
DR	6.2×10^{-3}	9.6×10^{-2}	2.0×10^{-2}
IPW	5.9×10^{-2}	6.0×10^{-2}	6.6×10^{-3}
Naive comparison	7.1×10^{-2}	7.1×10^{-2}	5.8×10^{-3}
OR	3.0×10^{-2}	1.0×10^{-1}	2.4×10^{-2}

Table 3: Performance of different ATE estimators when noisy confounders are observed. Estimators are Adversarial Balancing (AB), Doubly Robust (DR), Inverse Propensity Weighing (IPW), Naive mean comparison and Outcome Regression (OR). In contrast to Table 2, the performance of OR drops significantly while AB continues to provide a balance between variance and bias.

DR achieves the best performance in this case with respect to bias, however it once again suffers from large variance. The AB estimator seems to provide a balance between bias and variance.

5.2 Sentiment Average Treatment Effect on Treated and Responded (ATETR)

Similar to the ATE case, we have implemented five different estimators under two different scenarios.

In scenario 1, we assume that we can observe the actual values of the confounding factors, *i.e.*, X_1, X_2 . Our first estimator is a naive mean comparison among respondents. The naive estimator will return the same estimate for ATE and ATETR. We implemented outcome regression (OR) [3], inverse propensity weighing (IPW) [13], covariate control (CC), an adversarial balancing based estimator⁷ (AB) [20], and an entropy balancing[27]⁸ based method (EB) [9]. The main difference between ATE and ATETR estimation is the latter utilizes data among respondents only.

Results and boxplots are given in Table 4 and Figure 7, respectively. In our tests, the OR estimator has the largest bias. The estimator is trained using control-assigned survey respondents, and counterfactuals for the treatment assigned respondents are generated using the trained model. In our setting, the OR estimator consistently underestimates counterfactuals for the treatment-assigned respondents, resulting in a bias that is comparable with the naive

estimator. CC, EB and IPW have small biases; EB and CC estimators have the best MSE performance across all estimators.

Figure 7: Performance of different ATETR estimators when true confounders are fully observed. CC and EB outperforms AB in contrast to ATE.

In scenario 2, we relax our assumption on our ability to observe the true confounders, and we assume that we can only observe transformations of the true confounders as given in Eq. (15)-(16). Results and boxplots are given in Table 5 and Figure 8, respectively.

⁷We use logistic regression for the discriminator.

⁸We utilized first moment based balancing for EB estimator.

Estimator	Bias	MAE	MSE
AB	3.2×10^{-2}	3.3×10^{-2}	1.7×10^{-3}
CC	7.7×10^{-3}	1.9×10^{-2}	7.8×10^{-4}
EB	4.3×10^{-3}	1.8×10^{-2}	7.1×10^{-4}
IPW	7.7×10^{-3}	1.9×10^{-2}	1.0×10^{-3}
Naive comparison	7.1×10^{-2}	7.1×10^{-2}	5.8×10^{-3}
OR	7.8×10^{-2}	7.8×10^{-2}	6.8×10^{-3}

Table 4: Performance of different ATETR estimators when true confounders are fully observed. Estimators are Adversarial Balancing (AB), Covariate Control (CC), Entropy Balancing (EB), Inverse Propensity Weighing (IPW), Naive mean comparison and Outcome Regression (OR). EB has the best performance followed by simple CC estimator.

Estimator	Bias	MAE	MSE
AB	4.1×10^{-2}	4.1×10^{-2}	2.5×10^{-3}
CC	2.6×10^{-2}	2.8×10^{-2}	1.5×10^{-3}
EB	1.7×10^{-2}	2.3×10^{-2}	1.2×10^{-3}
IPW	-3.3×10^{-1}	3.1×10^{-2}	5.8×10^{-1}
Naive comparison	7.1×10^{-2}	7.1×10^{-2}	5.8×10^{-3}
OR	1.2×10^{-1}	1.2×10^{-1}	1.6×10^{-2}

Table 5: Performance of different ATETR estimators when noisy confounders are observed. Estimators are Adversarial Balancing (AB), Covariate Control (CC), Entropy Balancing (EB), Inverse Propensity Weighing (IPW), Naive mean comparison and Outcome Regression (OR). Similar to the results in Table 4, EB outperforms across all measures. Similar to Table 4, simple CC estimator provides a comparable performance to EB.

CC suffers from a dramatic drop in performance when compared to the fully observed covariates scenario. We also observe dramatic bias and higher levels of variance in IPW as well as a large number of estimates with extreme values. The performance of the AB estimator seems comparable between the two scenarios. While the EB estimator suffers from increased bias, its bias and MSE are nevertheless the lowest.

6 CONCLUSION

In this paper, we focused on a scenario where user reported sentiment is used as a success metric for a fully randomized A/B test. Sentiment is assumed to be collected via a survey, and users have the option to not respond to the survey.

For both the Average Treatment Effect (ATE) and Average Treatment Effect on Treated and Responded (ATETR) we showed that the naive estimator, i.e a simple mean comparison, results in biased estimates. This is troubling since the experiment platforms employed by many companies rely on simple mean comparisons for A/B analysis. Instead, we proposed consistent estimators for the ATE and ATETR.

We have already utilized our technique on user generated data. Similar to the simulations in Section 5, our framework has produced significantly different results when compared to naive comparisons, even changing the sign of the estimated impact in at least one case.

In the case of ATE, we showed that our problem can be broken down into two independent missing data problems. In other words, we can focus on treatment and control sets separately, and estimate population-level sentiment for both groups. Then we show that a simple difference between the two quantities provides a consistent estimator of the ATE. This also enables parallel execution for more efficient computation.

We simulate the problem with both fully observed and transformed confounders, observing significant performance decreases between these two settings, highlighting the potential sensitivity of causal inference techniques to model misspecification. The entropy and adversarial balancing estimators seem to provide a reasonable balance between bias and variance in the transformed confounders scenario.

In the case of ATETR, we showed that one can focus on respondent data only. Our problem can be mapped onto a standard causal analysis setting, i.e by restricting to the subset of respondent users in treatment and control and estimating the average treatment effect on the treated. This enables us to leverage existing techniques in the causal analysis literature.

Similarly to ATE, we observe significant performance decreases between the two ATETR settings. In particular, IPW suffers from dramatic performance degradation (potentially due to extreme values of predicted propensities). Interestingly, a simple regression-based covariate control (CC) was one of the best estimators for the

Figure 8: Performance of different ATETR estimators when noisy confounders are observed. Estimators are Adversarial Balancing (AB), Covariate Control (CC), Entropy Balancing (EB), Inverse Propensity Weighting (IPW), Naive mean comparison and Outcome Regression (OR). Similar to the results in Table 4, EB and CC outperform across all measures. IPW suffers from large variance, OR performs worse than naive estimator.

fully observed case, suggesting more complicated techniques do not necessarily provide superior performance.

Our novel framework combined with simple regression or balancing techniques has the potential to be integrated with existing experimentation platforms to provide robust, efficient, and automatic impact calculations for sentiment metrics.

ACKNOWLEDGMENTS

We are grateful to Micha Segeritz and David Sinclair for their feedback on this work. We also would like to thank Xin Fu for highlighting our work at KDD 2019 during the *Challenges, Best Practices and Pitfalls in Evaluating Results of Online Controlled Experiments* tutorial [23].

A PROOF OF LEMMA 4.2

First recall that ATETR is defined as $E[Y(1) - Y(0)|T = 1, \Delta(T) = 1]$ in Equation (2). By the linearity of the expectation operator:

$$ATE_{TR} = E[Y(1)|T = 1, \Delta(T) = 1] - E[Y(0)|T = 1, \Delta(1) = 1].$$

A simple average over the set of users who are assigned to the treatment ($T = 1$) and responded to the survey $\Delta(1) = 1$ is a consistent estimator of the first quantity, i.e., $E[Y(1)|T = 1, \Delta(T) = 1]$. This follows from the fact that assignment is random.

If we focus on the second term $\theta = E[Y(0)|T = 1, \Delta(T) = 1]$, the following set of equalities hold:

$$\theta = E[Y(0)|\Delta(1) = 1] \quad (11)$$

$$= E[E[Y(0)|\Delta(1) = 1, \Delta(0) = 1]|\Delta(1) = 1] \quad (12)$$

$$= E[E[E[Y(0)|\Delta(1) = 1, \Delta(0) = 1, X]|\Delta(1) = 1, X]] \quad (13)$$

$$= E[E[E[Y(0)|\Delta(0) = 1, X]|\Delta(1) = 1, X]]. \quad (14)$$

Equation (11) follows from the fact that treatment/control groups are randomly assigned, Equation (12) and (13) hold due to the law of iterated expectations, Equation (14) follows from Assumption 2.

B DATA GENERATING PROCESS FOR SIMULATIONS IN SECTION 5

Given our causal graph in Figure 4, we draw Monte Carlo simulations using the following data generating process (DGP):

$$X_1 \sim \text{Exp}(\lambda = 2), X_2 \sim \text{Exp}(\lambda = 3)$$

$$S = 2X_1 - 1.5X_2$$

$$Y \sim \text{Bin}\left(\frac{1}{1 - e^{-S}}\right)$$

$$R = X_1T - X_2T - X_2(1 - T) - 2$$

$$\Delta(T) \sim \text{Bin}\left(\frac{1}{1 - e^{-R}}\right)$$

where $\text{Exp}(\lambda)$ is an exponential distribution with rate λ , $\text{Bin}(\theta)$ is a binomial distribution with probability θ , T is a treatment indicator ($T = 1$ for treatment and $T = 0$ for control), Y is the sentiment, and $\Delta(T)$ is an indicator of whether a response was received under treatment T . Our DGP satisfies both Assumptions 1 and 2 since $X = [X_1, X_2]$ confound both response rates and sentiment. True sentiment Y is the same across treatment and control since parameter S is independent of the treatment indicator T . Response rates, however, are a function of both the confounders X and the treatment indicator T . Therefore, the observed sentiment for the two groups is different.

In some of our simulations, we relax the observability on the true confounders X_1, X_2 to study the effects of model misspecification. We instead assume that we can observe functions of the confounders, Z_1, Z_2 where:

$$Z_1 = e^{X_1/2} \quad (15)$$

$$Z_2 = X_2^2 - X_2 + X_1 * X_2 \quad (16)$$

This corresponds to a scenario where a researcher cannot directly observe the confounders but has access to a set of proxy variables for the confounders.

REFERENCES

- [1] Eugene W. Anderson. 1998. Customer Satisfaction and Word of Mouth. *Journal of Service Research* 1, 1 (1998), 5–17. <https://doi.org/10.1177/109467059800100102>
- [2] Eugene W. Anderson and Mary W. Sullivan. 1993. The Antecedents and Consequences of Customer Satisfaction for Firms. *Marketing Science* 12, 2 (1993), 125–143. <http://www.jstor.org/stable/184036>
- [3] Heejung Bang and James M. Robins. 2005. Doubly Robust Estimation in Missing Data and Causal Inference Models. *Biometrics* 61, 4 (2005), 962–973. <https://doi.org/10.1111/j.1541-0420.2005.00377.x> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1541-0420.2005.00377.x>
- [4] J Michael Brick and Graham Kalton. 1996. Handling missing data in survey research. *Statistical methods in medical research* 5, 3 (1996), 215–238.
- [5] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21, 1 (2018), C1–C68. <https://doi.org/10.1111/ectj.12097> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/ectj.12097>
- [6] Edith D. de Leeuw. 2012. Counting and Measuring Online: The Quality of Internet Surveys. *BMS: Bulletin of Sociological Methodology / Bulletin de Methodologie Sociologique* 114 (2012), 68–78. <http://www.jstor.org/stable/24311411>
- [7] Frederick F Reichheld. 2004. The One Number you Need to Grow. *Harvard business review* 81 (06 2004), 46–54, 124.

- [8] Anders Gustafsson, Michael Johnson, and Inger Roos. 2005. The Effects of Customer Satisfaction, Relationship Commitment Dimensions, and Triggers on Customer Retention. *Journal of Marketing - J MARKETING* 69 (10 2005), 210–218. <https://doi.org/10.1509/jmkg.2005.69.4.210>
- [9] Jens Hainmueller. 2012. Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies. *Political Analysis* 20, 1 (2012), 25–46. <https://doi.org/10.1093/pan/mpr025>
- [10] Guido W. Imbens and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139025751>
- [11] G. Kalton. 1983. *Compensating for missing survey data*. Survey Research Center, Institute for Social Research, the University of Michigan. <https://books.google.com/books?id=vTpHAAAAAAAJ>
- [12] Graham Kalton and Ismael Flores-Cervantes. 2003. Weighting methods. *Journal of official statistics* 19, 2 (2003), 81.
- [13] Joseph D. Y. Kang and Joseph L. Schafer. 2007. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statist. Sci.* 22, 4 (11 2007), 523–539. <https://doi.org/10.1214/07-STS227>
- [14] Ron Kohavi, Alex Deng, Brian Frasca, Toby Walker, Ya Xu, and Nils Pohlmann. 2013. Online Controlled Experiments at Large Scale. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Chicago, Illinois, USA) (KDD '13). ACM, New York, NY, USA, 1168–1176. <https://doi.org/10.1145/2487575.2488217>
- [15] Roderick J A Little and Donald B Rubin. 1986. *Statistical Analysis with Missing Data*. John Wiley & Sons, Inc., New York, NY, USA.
- [16] Neil A. Morgan, Eugene W. Anderson, and Vikas Mittal. 2005. Understanding Firms' Customer Satisfaction Information Usage. *Journal of Marketing* 69, 3 (2005), 131–151. <http://www.jstor.org/stable/30162061>
- [17] Hendrik Muller and Aaron Sedley. 2014. HaTS: Large-scale In-product Measurement of User Attitudes: Experiences with Happiness Tracking Surveys. In *Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures: The Future of Design* (Sydney, New South Wales, Australia) (OzCHI '14). ACM, New York, NY, USA, 308–315. <https://doi.org/10.1145/2686612.2686656>
- [18] J.H. Myers. 1999. *Measuring Customer Satisfaction: Hot Buttons and Other Measurement Issues*. American Marketing Association. <https://books.google.com/books?id=PdlmQgAACAAJ>
- [19] Optimizely. 2019. *Optimizely*. <https://www.optimizely.com/>
- [20] Michal Ozery-Flato, Pierre Thodoroff, and Tal El-Hay. 2018. Adversarial Balancing for Causal Inference. *arXiv e-prints*, Article arXiv:1810.07406 (Oct 2018), arXiv:1810.07406 pages. arXiv:1810.07406 [cs.LG]
- [21] Judea Pearl. 2009. *Causality: Models, Reasoning and Inference* (2nd ed.). Cambridge University Press, New York, NY, USA.
- [22] Donald B. Rubin. 1976. Inference and Missing Data. *Biometrika* 63, 3 (1976), 581–592. <http://www.jstor.org/stable/2335739>
- [23] Xiaolin Shi, Somit Gupta, Pavel Dmitriev, and Xin Fu. 2019. Tutorial: Challenges, Best Practices and Pitfalls in Evaluating Results of Online Controlled Experiments. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery Data Mining* (Anchorage, AK, USA) (KDD '19). ACM, New York, NY, USA. <https://sites.google.com/view/kdd2019-exp-evaluation/>
- [24] Louisa H Smith and Tyler J VanderWeele. 2019. Bounding bias due to selection. *Epidemiology (Cambridge, Mass.)* 30, 4 (2019), 509.
- [25] Diane Tang, Ashish Agarwal, Deirdre O'Brien, and Mike Meyer. 2010. Overlapping Experiment Infrastructure: More, Better, Faster Experimentation. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Washington, DC, USA) (KDD '10). ACM, New York, NY, USA, 17–26. <https://doi.org/10.1145/1835804.1835810>
- [26] Eugene W. Anderson, Claes Fornell, and Roland T. Rust. 1997. Customer Satisfaction, Productivity, and Profitability: Differences Between Goods and Services. *Marketing Science* 16 (05 1997), 129–145. <https://doi.org/10.1287/mksc.16.2.129>
- [27] Xiaojing Wang, Jingang Miao, and Yunting Sun. 2019. A Python Library For Empirical Calibration. *arXiv preprint arXiv:1906.11920* (2019).
- [28] Ya Xu, Nanyu Chen, Addrian Fernandez, Omar Sinno, and Anmol Bhasin. 2015. From Infrastructure to Culture: A/B Testing Challenges in Large Scale Social Networks. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Sydney, NSW, Australia) (KDD '15). ACM, New York, NY, USA, 2227–2236. <https://doi.org/10.1145/2783258.2788602>