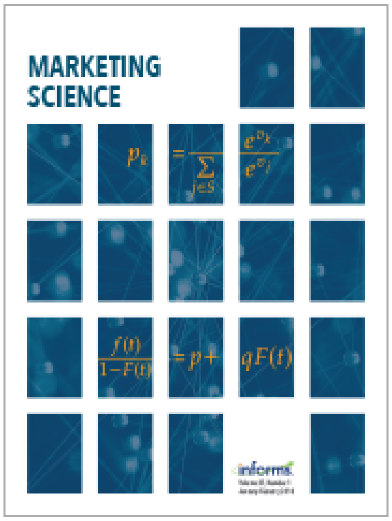

Marketing Science

Rejoinder: Heterogeneous Impact of Brands' Support for Black Lives Matter on Consumer Responses

Yang Wang, Marco Shaojun Qin, Xueming Luo, Yu (Eric) Kou

MARKETING SCIENCE
Vol. 41, No. 6, November–December 2022, pp. 1053–1056
ISSN 0732-2399 (print), ISSN 1526-548X (online)
https://pubsonline.informs.org/journal/mksc

Rejoinder: Heterogeneous Impact of Brands' Support for Black Lives Matter on Consumer Responses

Yang Wang,[a,b] **Marco Shaojun Qin,**[a,b] **Xueming Luo,**[a,b,*] **Yu (Eric) Kou**[a,b]

[a] Department of Marketing, Fox School of Business, Temple University, Philadelphia, Pennsylvania 19122; [b] Global Center for Artificial Intelligence & Business Analytics, Fox School of Business, Temple University, Philadelphia, Pennsylvania 19122
*Corresponding author
Contact: yangwang@temple.edu (YW); shaojun.qin@temple.edu, https://orcid.org/0000-0002-2787-5929 (MSQ); xueming.luo@temple.edu, https://orcid.org/0000-0002-5009-7854 (XL); yu.kou@temple.edu (Y(E)K)

Received: August 31, 2022
Accepted: August 31, 2022
Published Online in Articles in Advance: October 6, 2022

https://doi.org/10.1287/mksc.2022.1407

Abstract. We discuss three main lines of comments on our paper: internal validity of the research given consumers' potential misunderstanding of Black Lives Matter (BLM), challenges to the causal inference strategy, and interpretations of our results. We address them with additional robustness checks and reaffirm our view of the original findings as offering a panel of evidence with a consistent storyline: brands' support for BLM on social media induces, on average, a sizable negative effect on consumer responses. Whereas the commentary focuses on the average effect of BLM treatment, we highlight that the negative effect is highly heterogenous and can be mitigated when brands' BLM support is directed at Democratic consumers and delivered in a context where their BLM support can be interpreted as more authentic (e.g., no bandwagon, self-promotion, or cheap talk; consistent history of prosociality; and aligned brand prosocial mission). Furthermore, we want to emphasize the nascent nature of research on BLM and social justice movements for marketing, and our work represents a first empirical deep dive into the nature of consumer responses to the branded support of the BLM movement on social media. It is our hope that this research and the accompanying commentaries can inspire a wealth of new marketing perspectives on this critically important topic.

History: K. Sudhir served as the senior editor for this article.

Keywords: Black Lives Matter (BLM) • social media • brand management • causal inference • machine learning

1. Introduction

We are grateful for the thoughtful historical perspective and critical commentary by Thomas and Chintagunta (2022). The commentary provides three main comments: (1) questions of internal validity stemming from consumers' misunderstanding of the Black Lives Matter (BLM) movement and its implications for any empirical investigation in this domain, (2) challenges to difference-in-differences (DID) assumptions: differences and similarities between Instagram ("treated") and Twitter ("control") that may cause violations to the DID identification assumptions and the stable unit treatment value assumption (SUTVA), and (3) interpretations of the results. Next, we discuss each in order.

2. Internal Validity

Thomas and Chintagunta (2022) argue that Wang et al. (2022) (henceforth WQLK) faces internal validity issues because, in order to identify the effect of brand support of BLM on consumers' social media engagement, we must first define what supporting BLM means to consumers and assess whether consumers' understanding of BLM is consistent with the intended goals of the movement. Whereas we appreciate the historical context provided in the commentary and agree with the conclusions about the intended goals of the BLM movement (and would have included a similar discussion on this in a longer format publication), we do not find consumers' potential misunderstanding of the intended purpose of BLM to be a threat to the "internal validity" of WQLK. In fact, as empirical researchers, our goal is precisely to assess how consumers' potential heterogeneous beliefs about BLM aggregate to measurable behavioral outcomes. We are not attempting to theorize a new construct of "BLM support" or measure the consumer views of this construct; rather, we are measuring consumers' behavioral responses to brands' voicing support for BLM, taking, prima facie, consumers' disagreement over BLM's purpose as the inherent ground truth. If all consumers are aligned with BLM's view of its own purpose, consumers would not react negatively to brands' support of BLM. It is precisely consumers' (current) disagreement over the value of BLM that makes supporting such movements a strategic dilemma for firms.

Making the analogy to studying the effect of a new brand positioning campaign, the enforcement of a normative understanding of BLM upon consumers would be akin to ensuring that all audiences believe the value proposition of the new campaign—but the whole difficulty lies in communicating the intended value proposition in a convincing manner. Thus, studying consumer responses to brands' BLM support with all the real-world complications inherent in the beliefs about BLM is in itself important and valid, just as studying the effectiveness of a new brand positioning strategy is important.

3. Challenges to DID Assumptions

Much of the methodological critique comes from the commentary authors' observation that Instagram and Twitter are simultaneously too dissimilar and similar. In fact, this is the fundamental tradeoff in DID analysis. Parallel trends could not possibly hold if Instagram and Twitter are not similar, yet if they are too similar, especially in audience composition, then there could be spillovers of the treatment. We think the most pointed comment made by Thomas and Chintagunta (2022) regarding the differences between Instagram and Twitter is the fact that one platform is on average gaining followers (Instagram) whereas the other is losing followers (Twitter). Thus, our parallel trends assumption may appear unnatural. Indeed, our parallel trends assumption is that the daily changes, rather than absolute level, of a brand's followers on Instagram and Twitter are parallel (in log scale). Practically, this means that forces that increase follower growth[1] on one platform should also do so on the other platform, even if one platform is on average losing followers whereas the other is on average gaining followers. To support this assumption, we want to point out our empirical evidence in the one-month pretreatment period (WQLK, figure 2) shows the status quo of the follower growth on both platforms remains stable and parallel in the absence of any interventions. This observation is consistent with the parallel trends identification assumption. Specifically, for the purposes of DID analysis, Twitter only needs to serve as a parallel counterfactual for Instagram's follower growth in the absence of any intervention, not that Twitter needs to remain parallel if also given the same intervention, as implied by Thomas and Chintagunta's (2022) Apple music example where black squares were posted to both Twitter and Instagram (the latter would be a stronger than necessary condition for causal inference, i.e., symmetric treatment effects).

As the commentary authors noted, we also do recognize that Twitter is not a perfect counterfactual for Instagram for many of the same reasons stated in their commentary. Therefore, we provided consistent evidence using the same brand's Instagram follower growth in the prior year as an alternative baseline in a year-over-year (YoY) DID analysis (WQLK, table 5, column 3). Thus, even without considering Twitter, we still find quantitatively similar results.

Another major concern to the causal inference strategy brought up by the commentary is the potential large overlap in Instagram and Twitter consumers (i.e., spillover due to violation of SUTVA). There are several points to be made here. First, if audiences recall a brand's Instagram support and carry over that reaction to Twitter, then it should make the engagement outcomes on the two platforms more similar, biasing our estimated treatment effects toward zero (i.e., our results for the BLM treatment effect would be more conservative). Second, our YoY DID results suggest that the spillover is not a significant concern given the similar quantitative treatment effects using alternate baselines that do not suffer from spillover and violations of SUTVA.

A third major comment is about the treatment of Blackout Tuesday itself. The commentary authors suggest that concurrent confounders (polarizing political tension, BLM protests, COVID-19, etc.) are simultaneously occurring and may confound the Blackout Tuesday treatment. For instance, the larger role of the ongoing political tension surrounding BLM could be driving the treatment effects. We agree that concurrent confounders may be problematic if (1) they did not apply to both Instagram and Twitter and (2) they did not occur in the pretreatment period. However, this is clearly not the case. BLM protests and COVID-19 occurred before Blackout Tuesday and should have been known to almost all consumers regardless of platform, yet parallel trends still appear to hold. Furthermore, our consistent triple differences results, which compares YoY Instagram DID to YoY Twitter DID, should difference out such contemporaneous confounder concerns. The YoY differences on both platforms would take care of any (possibly asymmetric) effects of platform-agnostic confounders such as real-world BLM protests whereas the DID between the two platforms uncovers the treatment effect net of those confounders.

4. Interpreting DID Effect Size

There are several issues the commentary brings up about interpreting our findings. First, there is an issue of the importance of the effect size of the treatment effects from Blackout Tuesday participation. Because we have used log transformation for our dependent variable of *follower growth* (Y), to compute the effect size of BLM treatment on raw follower changes, we have to do the inverse operations:

Raw followers changed due to BLM treatment

$$= Sign(Y_t + coeff)EXP(Y_t + coeff) - Sign(Y_t)EXP(Y_t),$$

where *coeff* is about −1 in our DID estimator results.

For example, the effect size of BLM treatment on raw follower changes at the min (−8.36) of our dependent variable should be calculated as $[-\exp(|-8.36-1|) - (-\exp(|-8.36|)] = -7{,}342$ followers. Similarly, the effect size of BLM treatment on raw follower changes at the max (11.45) should be calculated as $\exp(|11.45-1|) - \exp(|11.45|) = -59{,}356$ followers. Following this computation process, Figure 1 presents the distribution of expected effect sizes measured as number of raw follower changes for brands on Instagram. The average effect size of BLM treatment is −407 followers. Therefore, we take this as evidence that our findings have substantially meaningful managerial significance.

A related interpretation issue regards the meaningfulness of studying follower change rather than follower change relative to total followers. Whereas a brand with a large number of followers might not care so much about not gaining an additional few hundred followers, follower growth is nevertheless an important key performance indicator (KPI) for brands.[2] Moreover, our data suggests that large brands may actually have a more difficult time growing their followers (see the negative correlation in Figure 2), perhaps due to saturation. Marketing practitioners also believe that influencing new customers during the honeymoon phase is a vital component to marketing strategy,[3] and that first impressions are important to the success of a brand.[4] Thus, even brands with large existing followings would remain interested in effects on follower growth.

A third interpretation issue is with regard to the permanence of the effect. We agree that the effects appear to be short lived. This is unsurprising given the constant stream of new content and the general

Figure 1. DID Effect Size in Terms of Raw Follower Changes

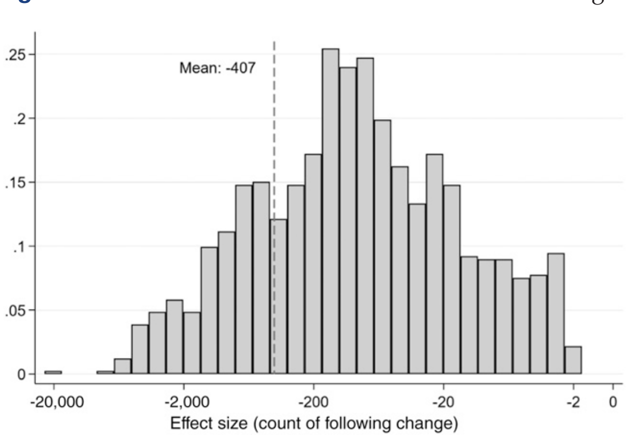

Figure 2. Scatter Plot of Brands' Growth Rate and Number of Total Followers

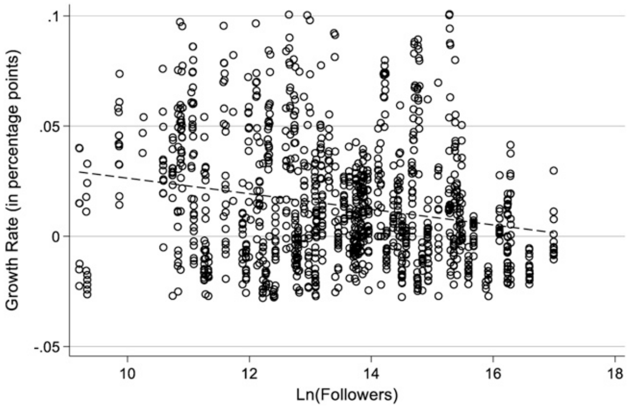

ephemeral nature of the reactions on social media. However, follower growth is a very coarse measure of consumer opinion, albeit one that we can fortunately observe. Nonetheless, we believe that it is reflective of the average negative receptivity from the audience even if most of that audience does not necessarily follow or unfollow brands. Thus, the perception effects are likely more widespread than can be inferred from the follower growth effects. Moreover, the fact that the unfollowing effect is relatively short lived does not make it unimportant if underlying perceptions of the brand are permanently changed for many consumers. We encourage future research to study opinion shifts possibly with YouGov data.

We appreciate the commentary authors' speculations about alternative explanations—as opposed to audience backlash from the support. They offer two explanations. First, losing followers is exactly the outcome intended by Blackout Tuesday (which is designed to silence commercial content). Second, as documented by news articles, BLM supporters exhibited backlash to some Blackout Tuesday posts as they wrongly tagged the content with "#BLM" and made it difficult to find important BLM messages. Whereas we find these arguments plausible, the simpler explanation that consumers negatively reacted to Blackout Tuesday posts remains most consistent with our empirical data. In particular, our results in column 4 of table 1 in WQLK show that, even when we remove Blackout Tuesday, we still get persistent negative effects from other bandwagon BLM support, thus it seems unlikely that the effects are completely explained by the special purpose of Blackout Tuesday to silence other content. Moreover, as Democrats are more likely to support BLM, our finding that brands with Democratic-leaning consumer bases are less likely to lose followers due to Blackout Tuesday participation seems to rule out the explanation that the negative effect is simply caused by BLM supporters.

Whereas the commentary focuses on the average effect of BLM treatment, we highlight that the effect is highly heterogeneous with nuanced implications in six aspects: (i) lone-wolf BLM support leads to negligible effects, but large-scale BLM support such as Blackout Tuesday by many brands may lead to strong negative effects; (ii) brands' self-promotion exacerbates the negative effects; (iii) a constant historical record of prosocial posting on social media attenuates the negative effects; (iv) socially oriented brand missions reduce the negative effects; (v) the negative effects are stronger/weaker for brands with mostly Republican/Democratic customers; and (vi) cheap talk in support for BLM (but without financial donations) can lessen the negative impact for brands with mostly Republican consumers but exacerbate the negative impact for brands with mainly Democratic consumers. In sum, the negative effects can be mitigated and even become positive when brands' BLM support is directed at Democratic consumers and delivered in a context where their BLM support can be interpreted as more authentic (e.g., no bandwagon, self-promotion, or cheap talk; consistent history of prosociality; and aligned brand prosocial mission).

5. Discussion

The thoughtful commentary provided by Thomas and Chintagunta (2022) highlights the importance of conducting careful marketing research on BLM, racial justice, and the larger diversity, equity, and inclusion domain. Although WQLK is a solid first step to understanding the nuanced marketing implications for branded racial justice support, we acknowledge that ours, like many other pieces of empirical research, is limited by the context and availability of data. Under these constraints, our results on the main and heterogeneous effects of brands' support of BLM seem to suggest the consistent story that brand advocacy of racial justice movements can backfire unless it is approached from a demonstrably authentic perspective by avoiding bandwagon support and consistently advocating for prosocial causes both and on and off social media. We hope that future research can expand upon our work with even more nuanced examinations of the important racial justice topics.

Endnotes

[1] Follower growth as defined in WQLK, $Y_t = \mathrm{Sign}(F_t - F_{t-1})\ln(|F_t - F_{t-1}|)$, where F_t is followers in time t.

[2] See https://www.klipfolio.com/resources/kpi-examples/social-media/followers-target.

[3] See https://www.linkedin.com/pulse/honeymoon-period-new-customers-cem-ozguven/.

[4] See https://www.adroll.com/blog/how-can-brands-make-the-right-first-impression.

References

Thomas JS, Chintagunta PK (2022) Invited commentary—"How support for Black Lives Matter impacts consumer responses on social media." *Marketing Sci.* 41(6):1045–1052.

Wang Y, Shaojun Qin M, Luo X, Kou YE (2022) Frontiers: How support for Black Lives Matter impacts consumer responses on social media. *Marketing Sci.* 41(6):1029–1044.