Excess success in
“Don’t count calorie labeling out: Calorie counts on the left side of menu items lead to lower
calorie food choices”

Gregory Francis¹, Evelina Thunell^{1,2}

¹Department of Psychological Sciences, Purdue University, West Lafayette, USA

²Department of Clinical Neuroscience, Karolinska Institutet, Stockholm, Sweden

Correspondence should be addressed to GF: gfrancis@purdue.edu, +1 765 494 6934

Abstract

Based on findings from six experiments, Dallas, Liu & Ubel (2019) concluded that placing calorie labels to the left of menu items influences consumers to choose lower calorie food options. Contrary to previously reported findings, they suggested that calorie labels do influence food choices, but only when placed to the left because they are in this case read first. If true, these findings have important implications for the design of menus and may help address the obesity pandemic. However, an analysis of the reported results indicates that they seem too good to be true. We show that if the effect sizes in Dallas et al. (2019) are representative of the populations, a replication of the six studies (with the same sample sizes) has a probability of only 0.014 of producing uniformly significant outcomes. Such a low success rate suggests that the original findings might be the result of questionable research practices or publication bias. We therefore caution readers and policy makers to be skeptical about the results and conclusions reported by Dallas et al. (2019).

1 Introduction

Many scientists take significant results that are replicated across multiple studies as strong support for their conclusions. However, this interpretation requires that the studies have high power. For example, when conducting six independent studies, each with a power of 0.5, one should expect only about half of the studies to produce significant results. It would be very rare for all six studies to produce significant results, namely $0.5^6 \approx 0.016$. When such excess success is observed in a publication, readers should suspect that the experiments were carried out using questionable research practices (John, Loewenstein, & Prelec, 2012; Simmons, Nelson, & Simonsohn, 2011) or that some experiments with non-significant results were run but not reported (publication bias; Francis, 2012b). A set of studies with too much success likely misrepresents reality, and conclusions from such studies should be discounted until non-biased investigations can be performed.

Here, we use a Test for Excess Success (TES) analysis (Francis, 2013; Ioannidis & Trikalinos, 2007; Schimmack, 2012) to show that the results of a recent article by Dallas et al. (2019) seem overly successful. Based on six successful studies, Dallas et al. (2019) conclude that consumers opt for lower calorie food items when calorie information is displayed on the menu – but only when it is placed so that it is read before the food names. According to the authors, previous failures to show an effect of calorie labeling on menus can be attributed to the fact that the calorie information was placed to the right of, and were thus read after, the food names. Dallas et al. (2019) correctly argue that their conclusions could have important implications for policy making to address the obesity crisis in America and elsewhere (Sunstein, 2019). However, the implications are only valid if the conclusions are valid, and the excess success of their findings undermines the credibility of the conclusion.

2 Material and methods

The mean, standard deviation, and sample size for each condition in each experiment reported by Dallas et al. (2019) is reproduced in Table 1, together with the key hypotheses used to support their theoretical conclusions. Each of the six reported experiments fully satisfied the hypotheses, thereby providing uniform support for the conclusions. Namely, in five of the studies (studies 1, 2, S1, S2, and S3), participants ordered fewer calories when calorie labels were placed to the left of food names, as compared to when they were placed on the right side. The remaining study (study 3) reported a corresponding effect for Hebrew readers (who read from right to left rather than left to right), so that calorie labels on the right (vs. left) led to lower calorie choices. Studies 1, S2, and S3 included a third condition with no calorie labels, and the number of calories ordered in the left condition was also lower than this no label condition.

We evaluated the plausibility of all six studies producing uniform success, by computing estimates of experimental power for replications of each of the studies. These power estimates are based on the statistics reported by Dallas et al. (2019), so our analysis starts by supposing that the reported findings are valid and accurate. Since the studies are statistically independent, we can then compute the probability of the full set of studies being uniformly successful by multiplying the estimates of the individual studies.

We first describe how to calculate the experimental power of the studies that used a single hypothesis test (study 2 and study S1). In this case the reported t value and sample sizes can be converted to a Hedges' g standardized effect size (Hedges' g is similar to Cohen's d , but with a correction for small sample sizes). For example, the conclusion of study 2 in (Dallas et al., 2019) is based on a significant two-sample t -test between the left and right calorie conditions, with $g = 0.25$. Based on this value, the power of a future experiment for any given sample size is easy to calculate. We used the *pwr* library (Champely et al., 2018) in

R to compute power for a replication experiment that uses the same sample sizes as the original study. Alternatively, the power could be computed from the means, standard deviations, and sample sizes by using the on-line calculator in Francis (2018) or similar tools. The same procedure applies to study S1.

Study	Supporting hypotheses	Calorie Placement			Success rate (power)
		Left	Right	No label	
1	Main effect of calorie information $\mu_{\text{left}} < \mu_{\text{right}}$ $\mu_{\text{left}} < \mu_{\text{nocalories}}$	$\bar{x} = 654.53$ $s = 390.45$ $n = 45$	$\bar{x} = 865.41$ $s = 517.26$ $n = 55$	$\bar{x} = 914.34$ $s = 560.94$ $n = 50$	0.4594
2	$\mu_{\text{left}} < \mu_{\text{right}}$	$\bar{x} = 1249.83$ $s = 449.07$ $n = 143$	$\bar{x} = 1362.31$ $s = 447.35$ $n = 132$	--- --- ---	0.5426
3	Main effect of calorie information $\mu_{\text{left}} > \mu_{\text{right}}$ $\mu_{\text{nocalories}} > \mu_{\text{right}}$	$\bar{x} = 1428.24$ $s = 377.02$ $n = 85$	$\bar{x} = 1308.66$ $s = 420.14$ $n = 86$	$\bar{x} = 1441.45$ $s = 382.43$ $n = 83$	0.3765
S1	$\mu_{\text{left}} < \mu_{\text{right}}$	$\bar{x} = 185.94$ $s = 93.92$ $n = 99$	$\bar{x} = 215.73$ $s = 95.33$ $n = 77$	--- --- ---	0.5358
S2	Main effect of calorie information $\mu_{\text{left}} < \mu_{\text{right}}$ $\mu_{\text{left}} < \mu_{\text{nocalories}}$	$\bar{x} = 1182.15$ $s = 477.60$ $n = 139$	$\bar{x} = 1302.23$ $s = 434.41$ $n = 141$	$\bar{x} = 1373.74$ $s = 475.77$ $n = 151$	0.5667
S3	Main effect of calorie information $\mu_{\text{left}} < \mu_{\text{right}}$ $\mu_{\text{left}} < \mu_{\text{nocalories}}$	$\bar{x} = 1302.03$ $s = 480.02$ $n = 336$	$\bar{x} = 1373.15$ $s = 442.49$ $n = 337$	$\bar{x} = 1404.35$ $s = 422.03$ $n = 333$	0.4953

Table 1. Supporting hypotheses, statistical properties, and estimated success rates for the tests in the TES analysis of the six studies in (Dallas et al., 2019).

It is more complicated to estimate the success rate when a study's conclusions depend on multiple hypothesis tests. Studies 1, 3, S2 and S3 in Dallas et al. (2019) are based on at least three significant hypothesis tests. We describe the procedure for study 1, which is representative of our approach. In study 1, a significant ANOVA was required to indicate a difference across conditions (left, right, or no calorie labels). In addition, the conclusions required both a significant contrast between the left and right calorie label conditions and a significant contrast between the left and no calorie label conditions. Because multiple tests are required for the results to fully support the conclusions, there is no single standardized effect size that can be used to compute the success rate (power) of the study. Instead, we ran simulated experiments that drew random samples of the same size as the original study from normal distributions with population means and standard deviations matching the statistics reported by Dallas et al. We then performed the three tests that were used in the original study on the simulated data. The process was repeated 100,000 times to give a reliable measure of the proportion of simulated experiments that found significance for all three tests. This proportion was then used as an estimate of the overall power (success rate) of the study.

The same procedure was used for studies 3, S2, and S3, using the respective reported statistics. Some of these studies comprised additional mediation analyses, which we did not include in our TES analysis. Since all of the mediation effects were in agreement with the conclusions in Dallas et al. (2019), including them in our analysis could only further reduce the estimated power. Therefore, our power estimations are likely higher than the true success rate.

Simulation source code (R core team, 2017) for all of the analyses is available at the Open Science Framework.

3 Results

The rightmost column in Table 1 shows the estimated success rate for each of the studies in Dallas et al. (2019). Each study has a success rate of around 0.5, so replication studies with the same sample size as the original studies should produce significant results only about half of the time, assuming that the population effects are similar to the reported sample effects. Thus, even if the position of calorie labels does influence food selection, it is very unlikely that six studies like these would consistently show such an effect. Indeed, the probability of all six studies being successful is the product of the success rates in Table 1, which is only 0.014.

How could Dallas et al. (2019) find positive results in all their studies when this outcome was so unlikely? One possible explanation is publication bias: Perhaps Dallas et al. (2019) ran more than six studies, but did not report the studies that failed to produce significant outcomes. Such selective reporting is problematic because even non-existent effects will sometimes produce significant outcomes (false positives). If only significant studies are reported, the reader is given the false impression of a very strong effect (Francis, 2012a). Another possible explanation is that the reported analyses were a subset of the full range of conducted analyses. A typical example of this approach is when researchers use several methods for outlier exclusion and then only report the one that resulted in the most favorable outcome. As another example, researchers may run their analysis on a set of data and then decide whether to gather more data based on the outcome of this intermediate analysis (e.g., stop if the results are significant and otherwise gather more data). This procedure results in multiple tests, and the Type I error rate is inflated (Simmons et al., 2011).

The selective reporting strategies described above are unfortunately rather common (Johns et al., 2011). We cannot know how Dallas et al. (2019) achieved their excessively successful results, but we can conclude that this set of studies does not make a plausible argument in support of the authors' conclusions. We recommend that scientists generally

ignore the findings and conclusions of Dallas et al. (2019). New studies will be required to determine whether calorie label position actually has the hypothesized effect on food choices.

4 Designing new studies

A scientist planning new studies on calorie label position might be tempted to base a power analysis on the results of Dallas et al. (2019). We show how this can be done, but we also caution readers that findings exhibiting excess success likely overestimate the reported effect size (Simmons et al., 2011; Francis, 2012a). Therefore, this approach likely overestimates experimental power and underestimates the necessary sample sizes. Thus, our calculation of sample sizes should be considered a lower limit.

For simplicity, we consider only the main comparison in the studies of Dallas et al. (2019): a difference in calories ordered depending on whether the calorie information was presented before or after the food names (in terms of reading direction). We used the reported means and (pooled) standard deviation to compute a standardized effect size (Hedge's g) for each study, and ran a meta-analysis to pool the effect sizes across experiments (source code is at the Open Science Framework). As an aside, a meta-analysis might not actually be appropriate for these studies because they differed in a number of potentially important methodological details. For example, in some studies the instruction was to order an entrée and a drink, while in others the menu only included entrées, or entrées and desserts. Still, the reported standardized effect sizes have mostly overlapping confidence intervals, and the meta-analysis will give a rough estimate of the effect size that might exist for a new study.

Researchers who feel that the full meta-analysis is not appropriate might pool the data in other ways. The standardized effect size for the individual experiments in Dallas et al. (2019) varies from 0.15 to 0.45, with smaller effect sizes for the studies with larger sample sizes. In a meta-

analysis, studies with larger sample sizes carry more weight. Taking everything in to account, the pooled effect size is $g^* = 0.2366$.

The second column of Table 2 shows the sample size per condition needed to achieve a specified power in a single study, based on the pooled effect size. To achieve 80% power, a new study should use 282 participants per condition; only one out of the six studies in Dallas et al. (2019) had at least this many participants (study S3). To achieve 90% power, sample sizes larger than any of the six studies (377 participants per condition) are required. Since these sample sizes are based only on the main comparison of left versus right position of the calorie labels, new studies that also include the left vs. no labels comparison or mediation analysis will require even larger sample sizes.

As noted above, the excess success analysis suggests that the pooled effect size is based on studies that most likely overestimate the effect size. A cautious scientist might therefore want to suppose that the population effect is smaller than the meta-analysis estimate: say, by one half. The third column of Table 2 shows the corresponding required sample sizes. In this case, to achieve 80% power for detecting a difference between the two calorie label placements, a sample size of 1123 participants per condition is needed. For an experiment to have 90% power, it would need to have 1502 participants in each condition.

The power analyses provide some guidance for repeating the experiments in Dallas et al. (2019). However, there may be two additional concerns that scientists seeking to understand the impact of calorie label placement should consider when designing new experiments. First, calorie count is often related to price, which makes it difficult to determine whether a condition is influencing selection based on calorie count or price. In the five out of six studies of Dallas et al. (2019) where the price was displayed on the menu, the calorie count was positively correlated with the price (see the supplemental material at the Open Science Framework). While such a positive relationship has real world validity, verifying the order-

based information processing hypothesis requires that an experiment uses menus where calories and prices are uncorrelated. Second, in the menus used by Dallas et al. (2019), placing the calorie labels after the food item name tended to place the label next to the item price. The two number items next to each other introduces visual clutter that can make it difficult for viewers to parse out relevant information (Shive & Francis, 2008). It is possible that viewers better process the calorie information when it is presented before the food name simply because it is then presented far from the price information. Bleich et al. (2017) describe additional challenges in studying the impact of calorie labels on food choices.

Desired power	To detect $g = 0.2366$ $n_1 = n_2$	To detect $g = 0.1183$ $n_1 = n_2$
0.80	282	1123
0.85	322	1284
0.90	377	1502
0.95	465	1858
0.99	658	2626

Table 2. Sample sizes required for a new study investigating left/right placement of calorie labels to have a desired power.

5 Conclusions

Dallas et al. (2019) note that their findings may have important implications for policies regarding calorie labels and their possible impact on obesity. However, such implications are only valid if the reported data support their conclusions. Given the inherent variability in the data, some non-significant results are highly likely when conducting experiments like their six reported studies - even if the effect of calorie label position is real and similar in magnitude to what they report. The lack of non-significant results (excess success) indicates that something

has gone wrong during data collection, analysis, or interpretation, perhaps unbeknownst to the authors (Gelman & Loken, 2014). We therefore advise readers to be skeptical about the results and conclusions reported by Dallas et al. (2019).

Contributions

GF and ET contributed to conception and interpretation. GF developed the analyses and wrote the code. GF and ET wrote the article and approved the submitted version for publication.

Competing interests

The authors declare no competing interests.

Data accessibility

R scripts and supplemental material are available at the Open Science Framework

https://osf.io/xrdhj/?view_only=4c023e4678c14e1eb7e5aa9380fb5bb5

References

- Bleich, S. N., Economos, C. D., Spiker, M. L., Vercammen, K. A., VanEpps, E. M., Block, J. P., ... Roberto, C. A. (2017). A Systematic Review of Calorie Labeling and Modified Calorie Labeling Interventions: Impact on Consumer and Restaurant Behavior. *Obesity*, 25(12), 2018–2044. <https://doi.org/10.1002/oby.21940>
- Champely, S., Ekstrom, C., Dalgaard, P., Gill, J., Weibelzahl, S., Ford, C., & Volcic, R. (2018). *Package ‘pwr’: Basic Functions for Power Analysis Description* (pp. 1–22). pp. 1–22.
- Dallas, S. K., Liu, P. J., & Ubel, P. A. (2019). Don’t count calorie labeling out: Calorie counts on the left side of menu items lead to lower calorie food choices. *Journal of Consumer Psychology*, 29(1), 60–69. <https://doi.org/10.1002/jcpy.1053>
- Francis, G. (2012a). Publication bias and the failure of replication in experimental psychology. *Psychonomic Bulletin & Review*, 19(6), 975–991.
- Francis, G. (2012b). Too good to be true: Publication bias in two prominent studies from experimental psychology. *Psychonomic Bulletin & Review*, 19(2), 151–156.
- Francis, G. (2013). Replication, statistical consistency, and publication bias. *Journal of Mathematical Psychology*, 57(5), 153–169.
- Francis, G. (2018). Power for independent means. In *IntroStats Online* (2nd ed.). Retrieved from https://introstatsonline.com/chapters/calculators/mean_two_sample_power.shtml
- Gelman, A., & Loken, E. (2014). The statistical crisis in science. *American Scientist*, 102(6), 460–465. <https://doi.org/2014.111.460>
- Ioannidis, J. P. A., & Trikalinos, T. A. (2007). An exploratory test for an excess of significant findings. *Clinical Trials*, 4(3), 245–253. <https://doi.org/10.1177/1740774507079441>
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23(5), 524–

532. <https://doi.org/10.1177/0956797611430953>

R core team. (2017). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria.

Schimmack, U. (2012). The ironic effect of significant results on the credibility of multiple-study articles. *Psychological Methods*, 17(4), 551–566. <https://doi.org/10.1037/a0029487>

Shive, J., & Francis, G. (2008). Applying models of visual search to map display design. *International Journal of Human Computer Studies*, 66(2), 67–77.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11), 1359–1366.
<https://doi.org/10.1177/0956797611417632>

Sunstein, C. R. (2019, March). Putting the calorie count before the cheeseburger. *Bloomberg Opinion*. Retrieved from <https://www.bloomberg.com/opinion/articles/2019-03-10/calorie-counts-on-menus-might-work>