

An Observational Study of the Effect of Nike Vaporfly Shoes on Marathon Performance

Joseph Guinness, Debasmita Bhattacharya, Jenny Chen, Max Chen, Angela Loh

Cornell University

Abstract

We collected marathon performance data from a systematic sample of elite and sub-elite athletes over the period 2015 to 2019, then searched the internet for publicly-available photographs of these performances, identifying whether the Nike Vaporfly shoes were worn or not in each performance. Controlling for athlete ability and race difficulty, we estimated the effect on marathon times of wearing the Vaporfly shoes. Assuming that the effect of Vaporfly shoes is additive, we estimate that the Vaporfly shoes improve men’s times between 2.1 and 4.1 minutes, while they improve women’s times between 1.2 and 4.0 minutes. Assuming that the effect of Vaporfly shoes is multiplicative, we estimate that they improve men’s times between 1.5 and 2.9 percent, women’s performances between 0.8 and 2.4 percent. The improvements are in comparison to the shoe the athlete was wearing before switching to Vaporfly shoes, and represents an expected improvement rather than a guaranteed improvement.

1 Introduction

There is a growing consensus that Nike Corporation’s new line of marathon racing shoes, which are commonly referred to as Vaporflys, provide a significant performance advantage to athletes who wear them. While several different versions of the shoes have appeared in races, including the Vaporfly 4%, the Vaporfly Next%, and several prototype shoes, each iteration of the shoes has in common a carbon fiber plate stacked inside of a highly responsive foam sole.

Several research studies have investigated the magnitude of the Vaporfly performance benefit. Hoogkamer et al. (2018) and Barnes and Kilding (2019) tested highly trained distance runners in laboratory studies, measuring various biomechanical and physiological variables while subjects wore Vaporflys and several other shoes in trial runs on a treadmill. Although the measured benefits varied somewhat from athlete to athlete, both studies found a roughly 4% average reduction in energy expenditures while wearing Vaporflys, in comparison to other popular racing shoes such as the Adidas Adizero Adios Boost line of road running shoes, and Nike Zoom Matumbo track spikes.

The Upshot, a division of The New York Times, collected data from actual marathon performances recorded on Strava, a popular running log and GPS tracking website. Their study included hundreds of thousands of marathon performances, and dozens of different shoes.

The Upshot found that the Vaporflys imparted a 4 to 5% advantage in finishing time over an average shoe and a 1.5 to 2.5% advantage over any other shoe (Kealy and Katz, 2018, 2019) A study published by Wired Magazine found that a sample runners in the 2017 New York City Marathon were more likely to run the second half of the race faster than the first if they were wearing Vaporflys (Thompson, 2017).

Our study is most similar to the Upshot study in that we analyze data from marathon performances and compare people’s performances with and without the Vaporflys. However, our study differs in a few ways. First, instead of relying on a convenient sample of athletes who upload their data to Strava, we take an exhaustive sample of athletes who met a minimum performance standard at one of 43 of the largest marathon races in 2015 and 2016 in the US and Canada. Second, instead of relying on self-reported shoe data, we searched the internet for photos of races and visually identified the shoes that runners wore. Third, we focus only on athletes who performed at an elite level before the Vaporflys were released to the public. Thus, we are only considering accomplished runners with marathon experience who, most likely, settled on a suitable shoe before the Vaporflys were released. These runners are also those most likely to be affected by shoe regulations because many of them compete in national Olympic qualifying races subject to regulations.

2 Study Design

We selected athletes who recorded a sufficiently fast marathon time—men under 2:24 and women under 2:45—at a collection of 21 distinct marathon venues in 2015 or 2016, including the 2016 U.S. Olympic Marathon Trials, which were contested in Los Angeles on the same weekend as the 2016 Los Angeles Marathon. This amounts to 43 distinct races and resulted in a sample of 270 distinct women and 308 distinct men after matching names and our best effort to correct alternate spellings of names. We recorded these athletes’ performances in the same 21 marathon venues over the period 2015 to 2019, and searched publicly available online photographs, manually identifying whether or not each athlete was wearing a Nike Vaporfly shoe by visual inspection. All marathon times were downloaded from the website www.marathonguide.com.

Our criteria for inclusion in the study were meant to satisfy certain objectives. First, we wanted to study elite and sub-elite athletes, since shoe regulations are motivated by performance advantages for athletes in this group. Second, we wanted to study athletes who had achieved success in the marathon before the Nike Vaporfly shoes had been released to the public. This ensures that inclusion in the study is unrelated to whether an athlete was wearing the shoes in the race where they qualified for inclusion in the study. This is important because, if any shoe effect exists, the magnitude of the effect may differ among different athletes. If we were to use performances potentially aided by the shoes to select the athletes, that might have biased our sample towards athletes who benefit most from the shoes.

To identify shoes worn by the runners, we used photos posted on public websites such as marathonfoto.com, marathon-photos.com, sportphoto.com, and flashframe.io. We also collected photographs from social media sites such as facebook.com and instagram.com. We

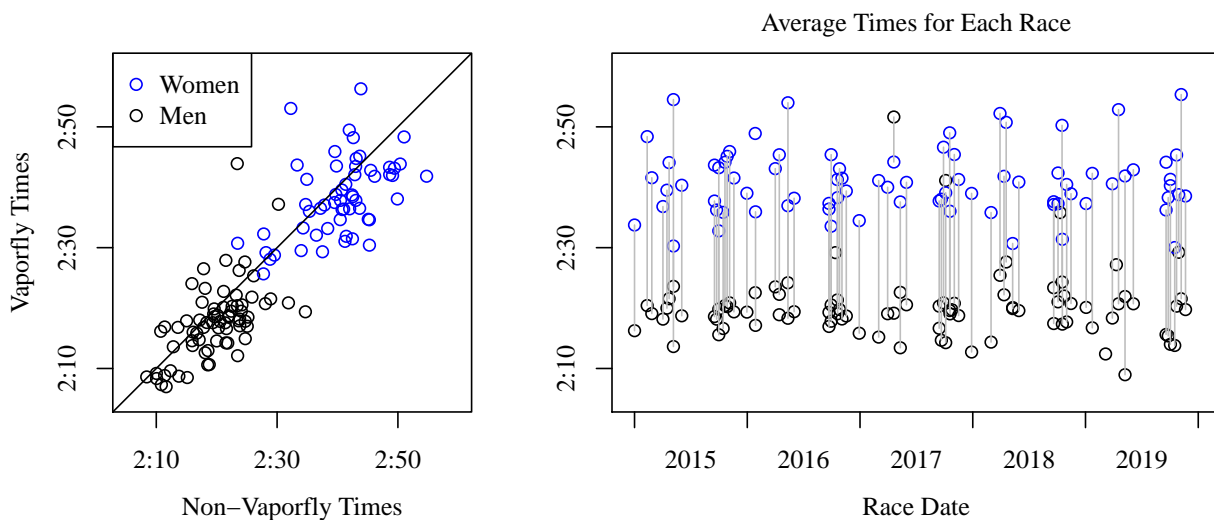


Figure 1: Averages of individual runner times, and averages of times from individual races for performances in our sampled dataset.

assumed that Vaporfly shoes were not worn in 2015 or 2016 by any runners except for a few that were reported to have worn prototypes in the 2016 US Olympic Trials Marathon. Identification of shoes via photos is a manual process that is subject to error. We have made all of our shoe identifications publicly available and will update this paper with new data if we are made aware of any errors in shoe identification. We identified the shoes worn in 840 of 880 (95.5%) men’s performances in our dataset and in 778 of 810 (96.0%) women’s performances.

3 Data Exploration

In Figure 1, we plot some summaries of the data, exploring how the times vary depending on runner, shoe, and marathon. In the left plot, we see that most runners’ average time in Vaporfly shoes is faster than their average time in non-Vaporfly shoes. Specifically, 53 of 71 men (74.5%) who switched to Vaporflys ran faster in them, and 40 of 56 women (71.4%) who switched to Vaporflys ran faster in them.

The left plot does not tell the whole story because it might be the case that runners who switched to Vaporflys did so when they ran on faster marathon courses. Some courses, such as the Boston Marathon course, have hills or often have poor weather, while others are flat and fast. So it is important to use the data to attempt to understand the difficulty of each course. In the right plot, we display the average times at each marathon race for all of the athletes in our sampled dataset. We can see that there is quite a bit of variability, but this plot is also somewhat unsatisfying because we do not know if stronger athletes preferred to run at easier or more difficult courses, which would influence the averages.

To get a satisfactory estimate of the effect of Vaporfly shoes, we need to analyze all of the data holistically, controlling for the strength of each runner and the difficulty of each marathon course. In the next section, we describe a statistical model intended for that purpose.

4 Statistical Model

We seek to model the dependencies among performances from each runner and among performances from each race. This will allow us to appropriately leverage the information in our dataset and arrive at an accurate estimate of the effect of the Vaporfly shoes. To allow for the possibility that men and women have different performance characteristics, we analyze data from the two sexes separately. We assign each performance a label between 1 and n ($n = 840$ for men, $n = 778$ for women). Each runner is assigned an label between 1 and the number of runners R ($R = 308$ men, $R = 270$ women), and we assign each marathon a label between 1 and the number of marathons M , ($M = 106$). We summarize our notation for the data here:

$$\begin{aligned}
 y_i &= \text{marathon time in minutes for performance } i \\
 x_i &= \begin{cases} 1 & \text{if Vaporfly shoes worn in performance } i \\ 0 & \text{if Vaporfly shoes not worn in performance } i \end{cases} \\
 j(i) &= \text{athlete who completed performance } i \\
 k(i) &= \text{marathon associated with performance } i
 \end{aligned}$$

A statistical model is a family of probability distributions that encodes the assumptions we make about the processes generating the data. Models generally include unknown parameters relevant to the questions posed in the study. The goal of the analysis is to use the data to make inferences about these parameters. We consider the following two models for the performances y_1, \dots, y_n :

$$\begin{aligned}
 \text{Untransformed:} & \quad Y_i = b_0 + b_1 x_i + A_{j(i)} + C_{k(i)} + Z_i \\
 \text{Log Untransformed:} & \quad \log Y_i = b_0 + b_1 x_i + A_{j(i)} + C_{k(i)} + Z_i
 \end{aligned}$$

with each of the individual terms defined in the following table

Terms	Assumptions	Description
b_0, b_1	non-random parameters	$b_1 =$ Vaporfly effect
A_1, \dots, A_R	$\overset{ind}{\sim} N(0, \sigma_1^2)$	runner effects
C_1, \dots, C_M	$\overset{ind}{\sim} N(0, \sigma_2^2)$	marathon effects
Z_1, \dots, Z_n	$\overset{ind}{\sim} N(0, \sigma_3^2)$	residual effects

The primary parameter of interest is b_1 , which is the effect of the Vaporfly shoes. The model assumes that, all else held constant, switching to Vaporfly shoes changes the response by b_1 . We do not attempt to model Vaporfly effects that vary among individual runners.

The interpretations of the parameters are different depending on whether we take a log transformation of the times or not. When modeling untransformed times, the effect of Vaporfly shoes is additive, meaning that we expect the time to change by adding b_1 , and when modeling log-transformed times, the effect is multiplicative, meaning that we expect the time to change by multiplying by $\exp(b_1)$.

Aside from b_0 and b_1x_i , the rest of the terms are random effects. Each of the R runners has its own offset term A_j to account for the fact that runners have differing abilities, and each of the M marathon races has its own offset term C_k to account for the fact that there are conditions unique to each race (e.g. course difficulty or weather) that influence all runners in the race. The final term Z_i accounts for any other factors that affected the performance. We also considered including time-varying runner effects, to allow each runner’s fitness to improve or decline over time independently of the fitness of other runners, but the time-varying runner effects did not significantly improve the model or change the estimates of the Vaporfly effects, so we omit them here to streamline the definition of our model.

5 Estimated Parameters

To fit the models, we used the `lmer` function, which is part of the `lme4` package (Bates et al., 2015) in the R programming language (R Core Team, 2018). The `lme4` package is a well-established piece of statistical software for fitting random effects models of the type we seek to estimate in this research study. Code and data for reproducing our results are available online at <https://github.com/joeguinness/vaporfly>.

We fit separate models for men and women, and additionally, separate models for the untransformed and log-transformed marathon times. The estimated parameters are summarized in the table below:

	men minutes	women minutes	men log minutes	women log minutes
	estimate (s.e.)	estimate (s.e.)	estimate (s.e.)	estimate (s.e.)
b_0	139.41 (0.43)	159.45 (0.63)	4.94 (0.003)	5.070 (0.0039)
b_1	-3.12 (0.62)	-2.59 (0.83)	-0.0221 (0.0042)	-0.0157 (0.0050)
σ_1	4.195	6.39	0.030	0.041
σ_2	2.328	3.17	0.016	0.019
σ_3	4.114	5.02	0.028	0.030

We see that for both men and women and for untransformed and log-transformed times, the Vaporfly effect is negative, indicating that the evidence supports the hypothesis that Vaporflys decrease, or improve, marathon times. Our best estimates of the additive effects are -3.12 minutes for men and -2.59 minutes for women. Using log-transformed data, our best estimates of the multiplicative effects are $\exp(-0.022) = 0.978$ for men, and $\exp(-0.0157) = 0.984$ for women, meaning that we expect men’s times to decrease by 2.18%, and women’s times to decrease by 1.56%, when wearing Vaporfly shoes, as compared to the shoes each athlete was wearing before switching to Vaporflys.

While our estimates suggest that the effect of Vaporfly shoes is greater for men, the estimates come with some uncertainty. In the following table, we include 90% confidence intervals for each of the Vaporfly effects, constructed using a normal approximation to the sampling distribution of the estimates.

men minutes	women minutes	men log minutes	women log minutes
(-4.144, -2.109)	(-3.959, -1.225)	(-0.029, -0.015)	(-0.024, -0.008)

Table 1: 90% confidence intervals for Vaporfly effects in each model.

None of the intervals contain zero, which indicates strong evidence for a non-zero Vaporfly effect. There is substantial overlap between the men’s and women’s confidence intervals, which leaves some uncertainty about which sex benefits most from Vaporfly shoes.

In random effects models such as the one we are using here, the estimates of the fixed effects, b_0 and b_1 , are calculated using the generalized least squares criterion. Generalized least squares attempts to triangulate all of the dependencies in the data, for example the fact that there are several performances for each runner and for each race, to arrive at a statistically optimal estimate of the effects. The estimates are linear combinations of the responses, for example

$$\hat{b}_1 = \sum_{i=1}^n c_i y_i,$$

where c_1, \dots, c_n are coefficients calculated using the covariance matrix of the random effects model. These coefficients can sometimes have seemingly counter-intuitive values. In the spirit of attempting to make sense of the magic of generalized least squares, and to promote its utility for this type of problem, we plot the coefficients for the estimates of the Vaporfly effects in Figure 2.

To help make sense of why generalized least squares picks these coefficients, consider four performances

- y_1 = Time for Runner 1 at Boston Marathon 2016
- y_2 = Time for Runner 2 at Boston Marathon 2016
- y_3 = Time for Runner 1 at Chicago Marathon 2017
- y_4 = Time for Runner 2 at Chicago Marathon 2017

The first runner (y_1 and y_3) did not wear Vaporflys, but the second runner (y_2 and y_4) switched to Vaporflys at Chicago in 2017. A reasonable estimate for the Vaporfly effect from these data might be the average of the Vaporfly performances minus the average of the non-Vaporfly performances,

$$y_4 - \frac{1}{3}(y_1 + y_2 + y_3),$$

which places a positive coefficient (1.0) on the Vaporfly performance and negative coefficients (-0.33) on the non-Vaporfly performances. However, a better estimate would consider how

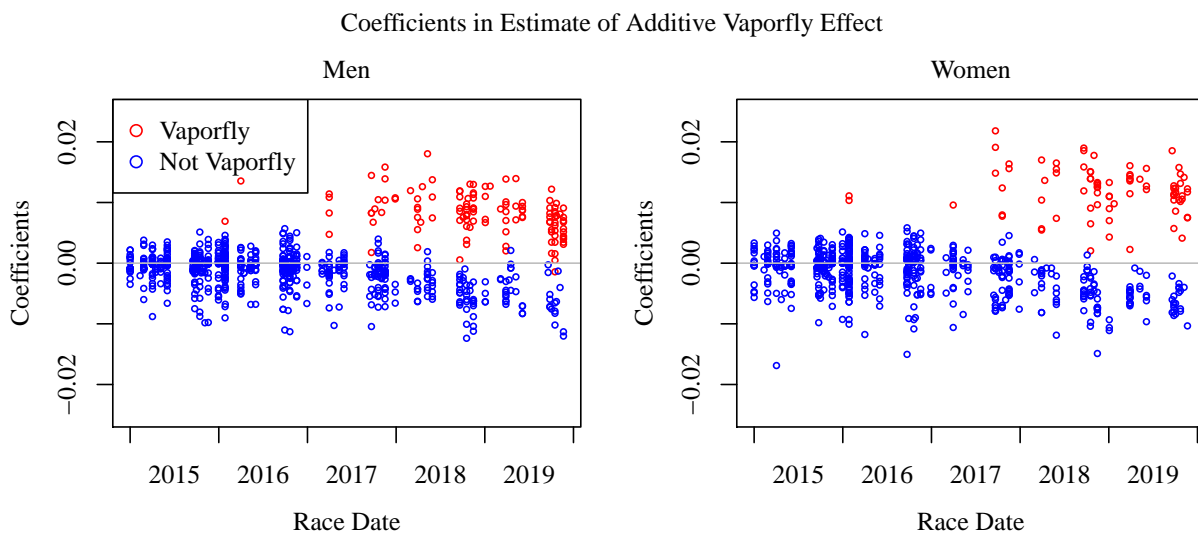


Figure 2: Coefficients for the generalized least squares estimate of the Vaporfly effect \hat{b}_1 for the untransformed data. Each point represents an individual performance y_i . The height of the point is the corresponding coefficient c_i .

much the second athlete’s advantage increased after switching to the Vaporflys,

$$(y_4 - y_3) - (y_2 - y_1).$$

The first difference $(y_4 - y_3)$ measures how much better the second athlete did in Chicago (wearing Vaporflys), and the second difference is how much better the second athlete did in Boston (not wearing Vaporflys). This estimate places a positive coefficient on the second runner’s Vaporfly performance in Chicago, a positive coefficient on the first runner’s Boston performance, a negative coefficient on the second runner’s Boston performance, and a negative coefficient on the first runner’s Chicago performance.

This pattern can be observed in Figure 2; early performances—before the Vaporfly appeared on the market—have either positive or negative coefficients, whereas later performances generally have positive coefficients when the Vaporfly is worn and negative coefficients when not worn. There is some variation in the magnitude of these coefficients, which we expect is due to the differing number of performances from each runner and from each race. Note that Figure 2 shows the coefficients in the generalized least squares estimate, not raw performances.

6 Discussion

By collecting data on marathon times and identifying shoes worn in a systematic sample of elite and sub-elite marathon runners, we studied how much a runner’s marathon time can be

expected to improve after switching to Vaporfly shoes. For men, the improvement is most likely somewhere between 2.1 and 4.1 minutes, or between 1.5% and 2.9%. For women it is likely between 1.2 and 4.0 minutes, or between 0.8% and 2.4%. To put these numbers into perspective, elite marathon runners cover more than half a mile in 3 minutes.

We made several assumptions that we believe to be reasonable, but nevertheless are open for debate and could be refined. We assumed that the expected Vaporfly effect (b_1) is the same for every runner of the same sex. Including prototypes, there are several different versions of Vaporflys, and it is logical to expect that newer versions improve upon older versions. Moreover, depending on biomechanical factors, some runners may benefit from Vaporflys more than others, or could have been wearing shoes that were not optimal for them before switching to Vaporflys. The models assume that, conditional on the runner and race, the marathon time follows a normal distribution. This may not be entirely appropriate because we believe that a runner is more likely to run 5 minutes slower than expected rather than 5 minutes faster; when things go wrong in a marathon, they can go really wrong.

The study did not include data on runners that did not finish their races. This data is more difficult to obtain and more difficult to model. Based on the Wired study, which found that people wearing Vaporflys generally ran better in the second half of the race, we believe that people wearing Vaporflys are less likely to drop out, so we are more likely to miss the very worst performances in non-Vaporfly shoes. Thus, we believe that including drop-outs would only strengthen our estimate of the Vaporfly effect.

The shoes were identified via a manual process of searching through photographs. We believe that this manual process for getting the shoe data is better (though labor intensive) than relying on self-reported shoes. Nonetheless, it is prone to error in misidentification of the person and the shoes. For example, Nike has an orange Vaporfly with a black Swoosh logo, and also a Zoom Fly with the same color scheme. Another example, some professional athletes ran in a neon yellow and pink prototype Vaporfly, which looks very similar to the neon yellow and pink Nike Zoom Streak 6. See Johnson (2020) for a more detailed analysis. Yet a further example, we identified one athlete who attempted to conceal the identity of his shoes by coloring in the white Swoosh on a blue pair of Vaporflys.

Some, perhaps many, of the athletes may have competed in marathons not included in the 21 marathons that we sampled. Missing these performances shouldn't bias our results, but our results could be strengthened if we are able to track down every performance from the athletes in the study.

It is possible that athletes are more likely to switch to Vaporfly shoes when they know they are ready to turn in a good marathon performance. Inversely, some athletes might not be willing to pay \$250 for shoes when they are out of shape. The Upshot study investigated this possibility by controlling for training volume; they did not see substantially different results. Further, our sample consists solely of highly accomplished runners. We believe that these athletes are generally wearing the best shoes available to them whenever they run a marathon.

We were able to identify shoes in nearly all, but not all marathon performances. Athletes have the ability to suppress photographs of themselves, for example by untagging themselves

in Facebook photos, or simply electing not to post pictures of themselves from their races. If athletes are more likely to suppress photos of poor performances in Vaporfly shoes, our estimated effect of Vaporfly shoes could be larger than it should be.

Acknowledgements

The authors thank Richard Heffron and Melissa Hardesty for discussions and providing comments on an early draft of the manuscript. We also express gratitude for those who take photos of major marathons and post them to social media, especially Karen Mitchell, Clay Shaw, and Malrie Sonier.

References

- Barnes, K. R. and Kilding, A. E. (2019). A randomized crossover study investigating the running economy of highly-trained male and female distance runners in marathon racing shoes versus track spikes. *Sports Medicine*, 49(2):331–342.
- Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.
- Hoogkamer, W., Kipp, S., Frank, J. H., Farina, E. M., Luo, G., and Kram, R. (2018). A comparison of the energetic cost of running in marathon racing shoes. *Sports Medicine*, 48(4):1009–1019.
- Johnson, R. (2020). Mechanical doping in 2016 – will anything be done about it? <https://tinyurl.com/umvmy7t>.
- Kealy, K. and Katz, J. (2018). Nike says its \$250 running shoes will make you run much faster. what if that’s actually true? <https://www.nytimes.com/interactive/2018/07/18/upshot/nike-vaporfly-shoe-strava.html>.
- Kealy, K. and Katz, J. (2019). Nike’s fastest shoes may give runners an even bigger advantage than we thought. <https://www.nytimes.com/interactive/2019/12/13/upshot/nike-vaporfly-next-percent-shoe-estimates.html>.
- R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Thompson, N. (2017). Do nike’s new marathon shoes actually make you run faster? <https://www.wired.com/story/do-nike-zoom-vaporfly-make-you-run-faster/>.