Enduring Terror: The Effects of the Boston Marathon Bombing on Competitive Performance*

Alexander Cardazzi† Zachary Rodriguez‡
Old Dominion University Union College

Preliminary Draft: please do not cite.

Abstract

The 2013 Boston Marathon bombing was the largest domestic terrorist attack in the United States since September 11. We study how direct exposure to the attack affected runners' subsequent performance and participation in competitive environments, exploiting the sudden division of the field into those who finished before the explosions and those who were diverted from the course. Diverted runners ran roughly nine minutes slower, a four percent decline, in future marathons, an effect that persists across other major events and over time. These performance losses impaired runners' ability to meet qualifying standards: participation fell in races gated by qualifying times (e.g., Boston) but was unchanged in races with open or lottery-based entry. Affected runners are thus not choosing to step away from competition; they are pushed out by their worse performance. Our findings represent a persistent, objectively measured cost of direct exposure to terrorism, help explain some of the negative economic consequences of terrorism in the literature, and underscore the need for sustained, long-term victim services.

Keywords: marathon, performance, terrorism, trauma

JEL Classification: J24, Z20, D91

*The authors would like to thank Alejandro Abarca, Adam Biener, Tomas Dvorak, James Flynn, Chengyuan Hua, Brad R. Humphreys, Kyutaro Matsuzawa, Bryan McCannon, Zachary Porecca, Younghwan Song, Jay K. Walker, Harris Wu, William Yurcik, Gary Wagner, Participants at the UMBC Sports Economics, VAE, & EEA Conferences, and a runner from the 2013 Boston Marathon, whom the authors will keep anonymous, for valuable feedback throughout. All errors are our own.

†Old Dominion University, Strome College of Business, Department of Economics, Norfolk, Virginia 23529, USA; Email: acardazz@odu.edu

‡Union College, Department of Economics, 807 Union St, Schenectady, NY 12308, USA; email: rodriguz@union.edu.

1 Introduction

While terrorist attacks are infrequent, the social and economic costs are devastating and well documented. On the aggregate, attacks deter investors (Abadie and Gardeazabal, 2008), depress labor markets (Benmelech et al., 2010; Brodeur, 2018), and ultimately, reduce GDP over time (Abadie and Gardeazabal, 2003). However, the micro-level mechanisms underpinning observed changes in economic output are less understood, as the current literature cannot distinguish between voluntary withdrawal from the labor market and a diminished capacity to perform. Indeed, research has shown that terrorist events reduce happiness (Clark et al., 2020), trust (Ahern, 2018), and well-being (Akay et al., 2020), and increase stress (Metcalfe et al., 2011) among those indirectly exposed to terrorism. Even larger psychological costs are borne by individuals with direct exposure. A rich literature on the effects of post-traumatic stress disorder, a common result of experiencing terrorism, finds that those who suffer from PTSD exhibit symptoms of depression and anxiety, leading to generally lower productivity (Searing et al., 2013; Montgomery-Marks et al., 2025). Exposure to violence also leads to preferences for certainty, which can persist for decades following the event (Bogan et al., 2013; Callen et al., 2014; Kim and Lee, 2014; Shai, 2022). Persistent risk aversion (or psychological trauma more broadly) may then lead to either underperformance in or selection out of competitive environments. Both of these outcomes would, on an aggregate level, appear as reduced output, earnings, and employment (Heitmueller, 2005). This is underscored by Brodeur and Yousaf (2025), who find reduced county-level earnings and employment following mass shootings, but, due to standard labor market data limitations, struggle to isolate whether these outcomes stems from worker choice or performance failure.

We address this gap by studying runners present at the 2013 Boston Marathon, the scene of the largest domestic terrorist attack in the United States since 9/11. Marathon running represents a uniquely suited laboratory setting for objectively measuring both performance and participation given its standardized and repeated task that is comparable across individuals, events, and years (Palacios-Huerta, 2025), which is not possible in traditional labor market datasets. In addition, marathon running closely parallels the selection and reward structure of the labor market. Qualifying for elite, World Marathon Majors like the Boston Marathon requires sustained high quality performance over months of training, with entry depending on beating age and gender specific time standards, not unlike the competitive screening job seekers face to secure a limited number of desirable positions. marathon runners also derive satisfaction from participating, achieving personal time goals, and improving upon prior performances (Markle et al., 2018), similar to the utility workers derive from compensation, job satisfaction, and career advancement.

We use official Boston Marathon data, collected from the Boston Athletic Association, and estimate the effects of experiencing the bombing on future performance and participation using difference-in-differences and event study estimators. Our identification strategy leverages the exogenous timing of the detonations, which divided the field based on race-day progress. Relative to finishers, diverted runners had a fundamentally different experience the day of the bombing, as being diverted contributed to the chaos around what had happened. These individuals were less likely to have met with their family, friends, and loved ones, have access to their cell phones, and were able to neither cool down properly nor stay warm as temperatures dropped, inducing high levels of stress. While finishers exited the area after a typical marathon experience, diverted runners were stranded as police locked down the vicinity. We provide empirical support for our identification strategy via a placebo analysis of the 2005 Boston Marathon, where we find no evidence that "would-be diverted" runners in a race with similar conditions exhibit the changes observed in our 2013 sample. Further, we are able to rule out injury as a result of proximity to the bombing as an explanation for our findings by dropping runners who were closest to the explosions from our analyses. This confirms that our findings are driven by the trauma of the attack itself rather than inherent differences between runners of varying speeds or physical ailment.

Our results indicate that being diverted caused finishing times to increase by an average of approximately nine minutes, a roughly four percent decline in performance, relative to comparable individuals who had already finished the race. We also find an equivalent reduction in performance for the diverted runners when participating in three other marathons (the New York City, Chicago and Marine Corps (Washington DC) marathons), suggesting the traumatic effect is not confined to the Boston Marathon. Diverted runners are also less likely to participate in the Boston and Chicago Marathons, where entry is often determined by competitive qualifying standards, but just as likely to participate in both the New York City and Washington DC Marathons. This pattern is consistent with a decreased ability to qualify rather than a deliberate choice not to race, providing evidence for the mechanism behind reduced economic outcomes in the aggregate. Lastly, to ease concerns about the reference group (finishers) potentially being contaminated by their own experience of the bombing, we also provide alternative estimates that compare non-finishers and finishers to a matched set of runners who did not attend the 2013 Marathon. We find that both non-finishers and finishers perform worse by about thirteen and four minutes, respectively, which further reinforces the difference between them of about nine minutes.

Notably, the magnitudes of our estimates are large and persistent. According to a back of the envelope calculation, we find that the effect size of about nine minutes is equivalent to aging 7.2 years for women and 2.3 years for men.[1] As an alternative reference point, consider Nike's Vaporfly shoes, released in 2017. Designed for marathon running, Nike's Vaporfly reduced finishing times by approximately 4%, the same magnitude we find, but were considered such an advantage that there were calls to ban the shoes from the Tokyo Olympics (Hoogkamer et al., 2018; Barnes and Kilding, 2019).

Our paper makes several important contributions. By observing the same runners performing an identical task across multiple years and events, we are able to cleanly estimate the effects of direct exposure to terrorism and better understand the mechanisms behind the findings established in the literature. First, we provide estimates of productivity losses among individuals directly exposed to a terrorist attack, measured in a setting where the task is standardized and performance is objectively and repeatedly recorded. Most of the existing literature on the economic consequences of terrorism relies on aggregate or subjective outcomes, survey-elicited preferences, or samples of indirectly exposed individuals (Abadie and Gardeazabal, 2003, 2008; Shany, 2025; Clark et al., 2020). Second, we document that the effects of exposure are persistent, lasting at least four years with no evidence of recovery. This stands in contrast to studies of indirect terrorism exposure, which find that well-being effects largely dissipate within weeks (Metcalfe et al., 2011; Clark et al., 2020). Our findings are more consistent with the literature on violence and risk aversion, which documents preference changes that can persist for decades (Bogan et al., 2013; Callen et al., 2014; Kim and Lee, 2014; Shai, 2022). Third, and perhaps most interestingly, we decompose the extensive margin response into voluntary withdrawal and involuntary exit. Our evidence indicates that diverted runners do not choose to stop competing, but rather their diminished performance renders them unable to meet qualifying standards, effectively pushing them out of future competition. The distinction between choosing to leave and being

[1]Using our sample of runners from 2009-2011, we regress time on age with and without runner specific fixed effects. The estimates presented here reflect 9 minutes divided by coefficients of 1.26 minutes/year for women and 4 minutes/year for men with fixed effects. Without the fixed effects, the coefficients shrink somewhat to 1.17 and 0.93. Note that our sample of men is already somewhat older than our sample of women, which is likely why the coefficient (from the estimation with fixed effects) is larger in magnitude for men than women.

forced out is critical for understanding how trauma operates in labor markets but is rarely identifiable in employment data. Further, these results can help inform policy decisions and explain the mechanisms of existing results in the literature with aggregate data (Brodeur and Yousaf, 2025). Our findings underscore the need for sustained, long-term support services for individuals who directly experience terrorist events, as the costs to their productive capacity extend well beyond the immediate aftermath of an attack.

The remainder of the paper is organized as follows. Section 2 describes the data sources and sample construction. Section 3 presents our identification strategy and empirical framework, while Section 4 discusses our main results and robustness checks. Section 5 concludes.

2 The 117$^{\text{th}}$ Boston Marathon

The Boston Marathon is an annual marathon held on Patriots' Day (the third Monday of April) in Boston, Massachusetts. Founded in 1897 and inspired by the marathon at the 1896 Summer Olympics, it is the world's oldest annual marathon and one of the most prestigious road racing events, serving as one of the seven World Marathon Majors (Boston Athletic Association, 2014). The event attracts approximately 500,000 spectators along the route, making it New England's most viewed sporting event. Starting with 15 participants in 1897, the race has grown to an average of about 30,000 registered participants each year.

The 117$^{\text{th}}$ Boston Marathon took place on April 15, 2013. The race had 26,800 registrants, of whom 23,300 started, a typical turnout given a qualification process widely considered one of the most demanding in road racing.[2] At 2:49 PM, with roughly 17,600 of the 23,000 starters already finished, two bombs detonated 12 seconds apart near the finish line, killing three people and wounding hundreds more, including 17 amputees (Arsenault and Krueger, 2023).[3] The race was halted within minutes, and the roughly 5,700 runners still on the course, most already past mile 20, were diverted before reaching the finish. Figure 1 shows the bomb locations along the route and the two diversion points: runners east of Massachusetts Avenue were redirected to Boston Common, and all others to Kenmore Square.

The Boston Marathon is organized into different "waves" of runners of similar ability (as proxied by qualifying time relative to age and gender). The fastest runners are placed at the front of the race while the slowest are placed toward the end. At the very back of the race are "charity" runners, who need not qualify but instead raise money over a certain threshold to participate. Organizing the Marathon into waves is done to prevent slower runners from getting in the way of (or "boxing in") faster runners. The 2013 Boston Marathon was organized into three waves. By the time of the bombing, approximately 75.7% of all runners had finished the marathon. However, 98% and 96% of Wave I and Wave II runners had completed the race by this point, whereas only 36% of the runners in Wave III had crossed the finish line. Within Wave III, 55.6% of traditional, non-charity runners, and 26.5% of charity runners, had finished the race. We focus the remainder of the paper on the approximately 2,500 non-charity[4] Wave III runners who competed in the 2013 Boston Marathon.[5]

[2] The 2010, 2011, and 2012 Marathons had similar registrant counts, with turnout rates of 85%, 89%, and 81%, placing 2013 squarely in the middle of this range.

[3] The bombings were carried out by two brothers, Tamerlan and Dzhokhar Tsarnaev; some accounts suggest the timing was chosen to maximize casualties (Kirk, 2013).

[4] Our results are qualitatively similar when we include charity runners.

[5] In our cleaning of the data, we also remove a handful of runners from our sample who ran particularly poorly in 2012 due to the unique circumstances (heat) and obtained times over 5 hours and 15 minutes.

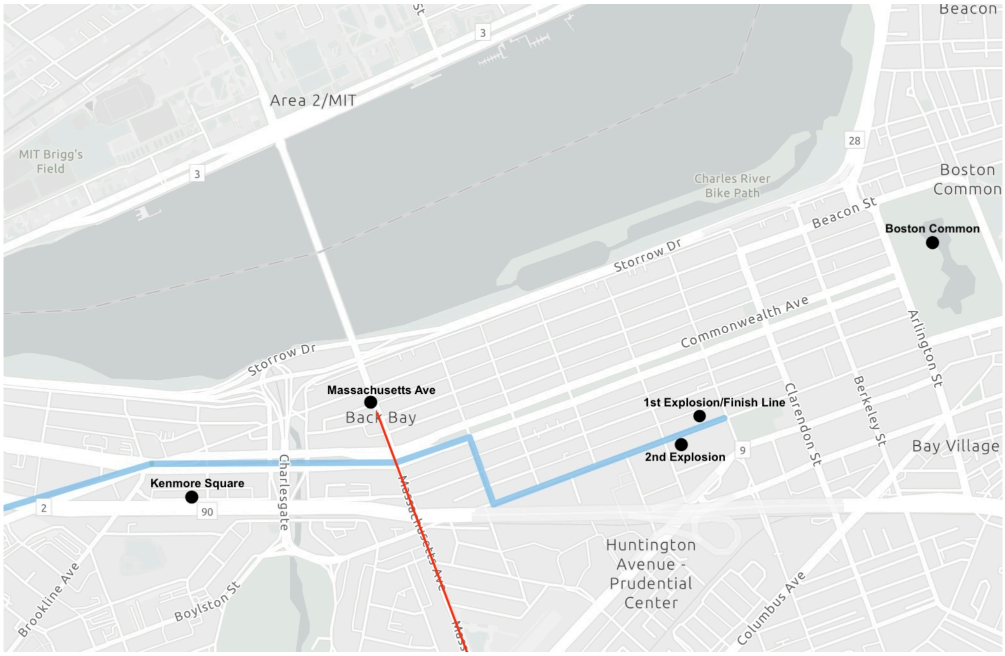

Figure 1: The 117th Boston Marathon Route and Diversion Locations Post-Bombing

Note: After both bombs exploded, some runners had finished and some were still running. Emergency workers diverted those still running based on their proximity to Massachusetts Avenue. If runners were east of Mass Ave, they were sent toward Boston Common, and if runners were west of Mass Ave, they were directed to Kenmore Square.

To study these runners, we collected data directly from various official Boston Athletic Association (BAA) sources. First, we gathered finish and "split" times, which are times at each 5 kilometer checkpoint along the race in addition to the halfway mark (at 21.1 kilometers), from https://www.baa.org/races/boston-marathon/results/. These data are available from 2001 to 2025, though we only collected information through 2017. For each runner who completed the Marathon, we observe their name, bib number, age, and performances (split times). We use the recorded split times to measure the time taken to run each approximate quarter of the race (0-10 km, 10-21.1 km, 21.1-30 km, and 30-42.2 km).

Second, we collected registration information for the 2010 (the earliest available) through 2017 Marathons from https://results.baa.org/. These data are important since the previous source only contains information on finishers. Notably, the results data source is missing information about runners who did not complete the race due to being diverted in 2013. The registration data contain names and bib numbers for all registrants regardless of whether they finished the marathon, dropped out midway (or were diverted), or did not show up altogether. In addition, we collect recorded (projected if diverted) split times in the 2013 Marathon for each runner who started the race from this source.

For our main analyses, we construct a panel of the 2013 non-charity Wave III runners spanning 2001–2017. The panel contains cleaned names and bib numbers from the registration data source so we can capture finishers and non-finishers alike. To minimize Type I errors when attributing information to runners across marathon years, we match individuals on both name and birth year. Registration data, however, do not contain age information, which we had to merge in from external sources. First, we obtain ages for finishers using race results data by matching on bib number, although this approach is not feasible for diverted runners. For these runners, we instead rely on an alternative source—since made unavailable by the BAA—that contains each registrant's bib number and age.[6] Once names and birth years are assigned, we merge in race results, split times, and registration information from other Boston Marathon years. This results in a final dataset with a total of 43,605 observations (2,565 runners over 17 years). It is important to note that we do not observe a completed marathon time for each runner in each year. In total, we observe 7,788 marathon times including 1,102 unfinished marathons with projected times from 2013.

Table 1 provides summary statistics on the finishers and non-finishers in our sample. Since the bombs occurred toward the end of the race, non-finishers were necessarily slower than finishers of the race. This is not particularly surprising, however, since these individuals also faced slower qualifying time requirements and had larger bib numbers on average. Non-finishers are also older on average and more likely to be male, but about equally likely to be from outside the United States.[7] Figure 2 illustrates the relationship between bib number and official/projected time of finish (time of day rather than time taken to run the race). The bombings occurred at 2:49 PM, which is represented by the dashed horizontal line. Finishers are represented by open circles beneath this line and non-finishers are filled in circles above the line. Non-finishers exist throughout Wave III, but it is clear that bib number is a strong predictor of whether a runner finished the marathon prior to the explosion.

[6]An archived snapshot of this source exists on the Wayback Machine: https://web.archive.org/web/20201112002928/http://registration.baa.org/2013/cf/Public/iframe_EntryLists.cfm. A 2025 version can be found here: https://registration.baa.org/2025/cf/Public/iframe_EntryLists.cfm. Downloads for 2025 are, are the time of writing, available at the following URL: https://registration.baa.org/2025/cf/Media/iframe_EntryLists.cfm.

[7]As a note, Table 1 suggests some overlap in times of finish between the two groups, which is of course impossible, and most likely due to projection error on the part of the BAA. Out of our sample of 2,565 runners, 74 have finish times that are inconsistent with their finish status according to the BAA. Our results are not sensitive to the classification of these 74 runners.

Table 1: Summary Statistics for 2013 Non-Charity Wave III Runners

Race Status		Mean	St. Dev.	Min	Median	Max
Finishers	Fin. Time	234.47	9.07	192.72	235.03	255.62
	Qual. Time	238.33	9.87	195	235	295
	Bib Num.	19192.77	769.07	18000	19113	21110
	Male	0.18	0.39	0	0	1
	Age	51.68	7.32	25	50	80
	Foreign	0.2	0.4	0	0	1
Non-Finishers	Fin. Time	271.44	22.65	237.35	264.25	484.98
	Qual. Time	246.47	15.8	185	240	310
	Bib Num.	20008.29	833.91	18002	20185.5	21117
	Male	0.31	0.46	0	0	1
	Age	57.04	8.67	23	57	82
	Foreign	0.2	0.4	0	0	1

Notes: N = 1,463 and 1,102 for finishers and non-finishers, respectively. Finish times for the non-finishing group are projections from the BAA rather than official times. Including the 2013 diversions, our sample includes a total of 7,788 observations of runners who completed a marathon. The average number of times we observe a particular runner is 3.04 times, but we observe some athletes a maximum number of 9 times and others a minimum of only 1 (in 2013) time.

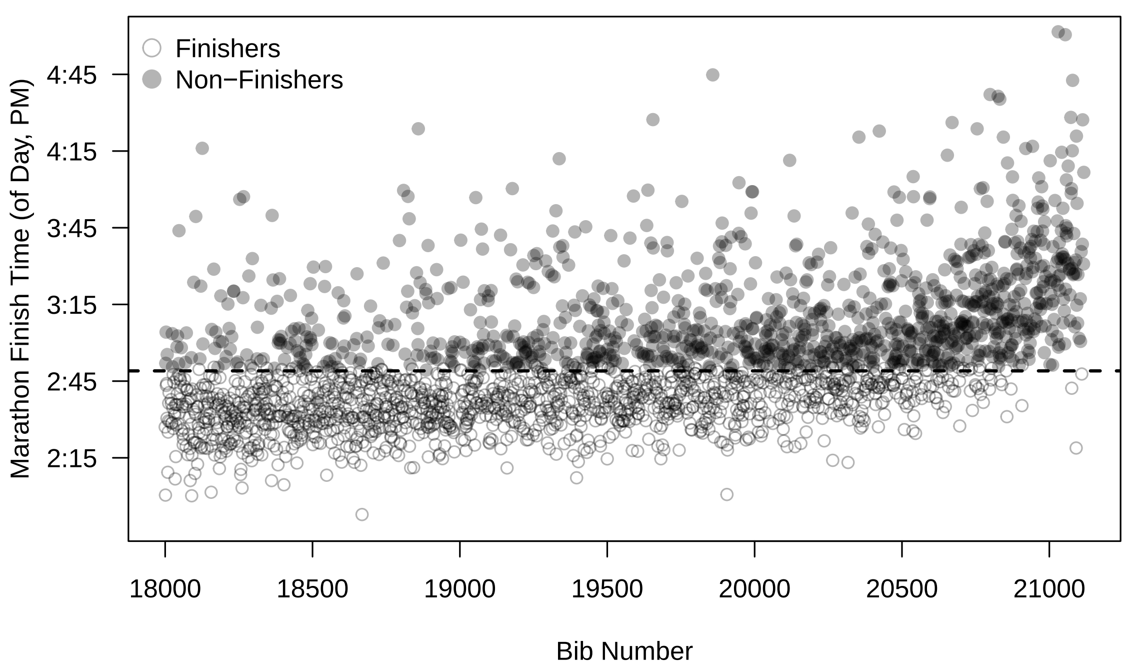

Figure 2: Time of Finish in 2013 by Bib Number and Treatment Status

Notes: This Figure illustrates the time of finish and bib number for each non-charity runner in Wave III of the 117$^{\text{th}}$ Boston Marathon. Times for the non-finishers are based on projections made by the BAA. Note that a few (74) observations fall on the "wrong" side of the cutoff due to projection inaccuracies. These cases are rare and do not affect the empirical results.

Figures 3a and 3b represent average finish time and registration by year and finish status in 2013. Prior to 2013, we find that the 2013 non-finishers are approximately 15 minutes slower than the 2013 finishers. This difference is consistent even in the 2012 Boston Marathon, which was an atypically hot day for the race as temperatures reached a high of 87°F (31°C). Temperatures do not often hit even 70°F, let alone mid to high 80s. The only Marathons to be as hot as 2012 in the last 50 years are 2004 (86 °F), 1987 (87°F), and 1976 (96°F).[8]. The difference in finish times in 2013 depicted in Figure 3a is obviously quite large, though this is by construction. Following 2013, we find an increase in the gap from about 15 minutes to about 30 minutes. This simple analysis of the raw data suggests that diverted runners performed worse after the bombing relative to their counterparts in Wave III.

Figure 3b displays the fraction of finishers and non-finishers that appeared in races prior to and following 2013. From 2010 through 2012, we find an approximately equivalent fraction for finishers and non-finishers. In 2014, however, we find a substantially larger fraction for non-finishers than finishers, though both are higher than any of the 2010-2012 fractions. We attribute this to two possible explanations, first, the BAA invited back all of the diverted runners and waived the need for them to qualify or pay in 2014 and second, the citywide sentiment to overcome the effects of the bombing made the race more ceremonial than past races (PBS NewsHour, 2014; DeBree, 2014). However, following 2014, the temporal pattern for the group of finishers group seems to about mirror the pattern prior to the bomb, which is likely a natural trend. However, the probability of non-finishers returning is approximately 10 percentage points lower than that of the finishers, whereas it was about equal prior to 2013.

[8]Temperatures according to Boston Discovery Guide: https://www.boston-discovery-guide.com/boston-marathon-weather.html, accessed Jan 6, 2025.

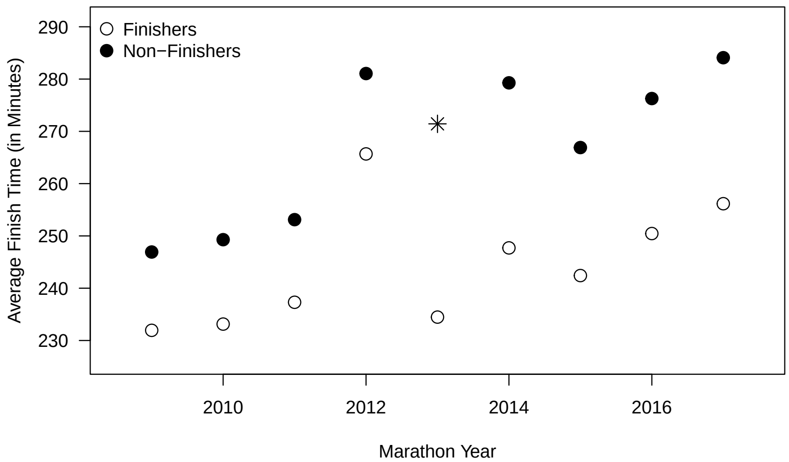

(a) Average Finish Time by 2013 Finishing Status and Year

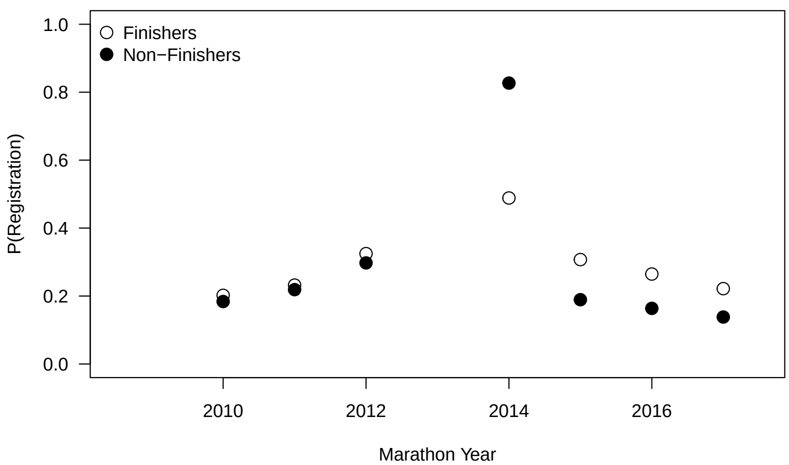

(b) Portion of 2013 Wave III Runners Registered in Other Boston Marathons by 2013 Finishing Status and Year

Figure 3: Main Study Outcomes by 2013 Finishing Status and Year

Notes: Panel (a) depicts the average finish times for both finishers and non-finishers of the 117$^{\text{th}}$ Boston Marathon across other Boston Marathons. Times for the non-finishers in 2013 are based on projections made by the BAA. Panel (b) depicts the fraction of runners from Wave III in 2013, by finishing status, who registered for other Boston Marathons in our sample. Values from 2013 are omitted from this graph since both groups necessarily have values of 1.

3 Empirical Strategy

Our research design identifies the effect of the Boston Marathon bombing on future performance and participation of the athletes involved by we comparing those runners who were diverted in 2013, and thus were more likely to experience the associated chaos, to those runners who completed the marathon without interruption, and thus experienced a highly regimented "cool down" process where they are guided away from the finish line by volunteers to prevent a pile-up and/or injuries, meet with their supporters, and leave the area.

Importantly, whether a runner was diverted was determined by their performance in 2013. Diverted runners were, mechanically, slower that day than runners who had finished the race prior to the explosions. Diverted runners were also older, more likely to be male, and registered with slower qualifying times (see Table 1). Diversion, in other words, was not randomly assigned. However, our identification does not require that diverted and non-diverted runners be comparable in levels, only that their outcomes would have evolved in parallel absent the attack. Formally, we assume that, conditional on runner fixed effects and age-gender-year cohorts, the outcomes of diverted runners would have moved in parallel with those of finishers, had the bombing not occurred. Non-random assignment into diversion is compatible with this assumption as long as whatever sorts runners into diversion does not also independently alter their subsequent trajectory.

An example of a violation of this assumption would be if a large portion of treated runners systematically sustained injuries between the registration deadline (mid-September) and race day (mid-April), sorting them into treatment by slowing them down on race day and separately altering their future outcomes via the same injury. In Sections 4.4 and 4.5, we test for, and rule out, such confounds via a matching exercise and by estimating the effect of a placebo bombing in the 2005 Boston Marathon. At the same time, diversion status was not purely a function of stable runner characteristics. Day-of performance shocks, quality of sleep the night before, what a runner ate, how hydrated they were, whether they got "boxed in" by the pack, mean that runners of similar underlying ability were pushed to either side of the diversion cutoff depending on how their race happened to go that day. This idiosyncratic variation creates meaningful overlap between the diverted and non-diverted groups, even though the two groups differ on average. Ultimately, we treat this overlap as support for the plausibility of parallel trends rather than as a substitute for it. If anything, mean reversion following transient performance shocks would cause diverted runners to improve relative to finishers after 2013, biasing our estimates toward zero.

3.1 Future Performance

As a first step, we model athlete i's performance (via finish time) in marathon t as:

$$\text{Time}_{it} = \delta\ (\text{DNF2013}_i \times \text{Post}_t) + \beta\ \text{Bib}_{it} + \mu_i + \tau_{a(i)g(i)t} + \epsilon_{it} \qquad (1)$$

where μ_i controls for runner-specific fixed effects (e.g., height and grit) while $\tau_{a(i)g(i)t}$ controls for the average performance of runners in the same age-gender cohort per race. In addition, we partial out the effects of runner i's assigned bib number for marathon t with β. Bib numbers are given out in increasing order of qualifying times and thus proxy for time-varying athlete ability that may be captured by neither μ_i nor $\tau_{a(i)g(i)t}$. ϵ_{it} represents a mean zero, possibly heteroskedastic, unobservable error term capturing all other factors affecting runner performance.[9]

δ represents our parameter of interest and is an estimate of the difference in the change in marathon times for those who were diverted in 2013 relative to the change in times for the group of runners who were not diverted. A positive estimate of δ would suggest that non-finishers in 2013 become relatively slower than finishers in subsequent races. δ, of course, is an intent to treat quantity since we cannot precisely observe which runners were exposed or quantify the extent to which they were "traumatized." Moreover, we expect this quantity to somewhat suffer from attenuation bias due to it being likely that some individuals in the

[9] We use heteroskedasticity robust standard errors throughout the study, which we find to be about 10% more conservative than clustering by runner.

control group were also nearby the explosions. A causal interpretation of δ requires the assumption that, in the absence of treatment, the times of diverted runners would have evolved in parallel to the times of runners who finished the 2013 marathon, conditional on age-gender cohort and pre-marathon expectations (i.e., bib number). Importantly, since we use performance times in minutes as our outcome, our assumption is explicitly that if the bombs had not detonated, the times of diverted runners would have increased the same amount, in minutes, as runners who completed the marathon. Instead, if we apply a log transformation, the assumption would subtly change to the following: if the bombs had not detonated, the times of diverted runners would have increased the same amount, as a rate, as runners who completed the marathon. We do not find differences in our results when we use levels compared to logs, though, which is due to initial average difference being only a few percentage points to start.

Including bib number as a control allows δ to be interpreted as the direct effect of the bombing on performance when running the Boston Marathon, rather than the total effect which could operate through qualification and race-day channels alike. Because bib numbers are assigned based on submitted qualifying times, they proxy for performance in marathons at least six months prior to the Boston Marathon, and thus could reflect any persistent effects of trauma that influence runner performance outside of Boston. Conditioning on bib number therefore holds constant the quality of a runner's qualifying performance and isolates the impact of trauma on performance specifically while running the Boston Marathon. Absent this control, δ would conflate the effect of trauma on race-day performance with its indirect effect operating through potentially altered qualifying times and bib assignments.

In another set of regressions, we vary the outcome to be either the time taken to complete different segments of the race ("splits") or the overall final ranking of the runner upon finishing. This allows us to explore whether different segments of the race experience different changes in speed. We use final rank as an outcome, even though it will be closely related to finish time, for two reasons. First, one can think of bib number as an expectation of final rank, making final rank the realized outcome of that expectation. Second, results from this regression give an alternative reference point for magnitudes.

3.2 Future Participation

Estimates from Equation (1) should be thought of as conditional on an athlete participating in the Boston Marathon in year t. Estimates of δ from this equation represent the effect of the bombing on the intensive margin, where athlete performance may be impacted. However, on the extensive margin, the rate at which athletes participate in the Boston Marathon may change, too. To explore this outcome, we model participation (via registration records) of runner i in marathon t as:

$$\text{Register}_{it} = \delta\ (\text{DNF2013}_i \times \text{Post}_t) + \mu_i + \tau_{a(i)g(i)t} + \epsilon_{it} \tag{2}$$

In this model, we use the same controls as Equation (1) except for β (since non-registered runners are not given bib numbers). We also use a modified dataset compared to the one used in Section 3.1. In this analysis, we do not remove observations of runners who did not complete the race, since this is part of what we are measuring. Instead we use the balanced panel of runners with an indicator for whether they had registered for the Boston Marathon in year t. In this Equation, δ represents the change in the probability of registering for the Boston Marathon for diverted runners relative to finishers.

Interpreting δ from Equation (2), however, should be done carefully. A negative estimate would indicate that non-finishers are less likely to register than finishers conditional on gender-age cohorts. However, broadly speaking, there are two main reasons why a runner may be less likely to participate in the Boston Marathon. First, each runner has to choose to participate in the Marathon. However, participation also depends on whether the runner can qualify for the Marathon. The data we observe, however, tells us whether runner i was given a bib for marathon t. From these data, δ cannot directly differentiate between these two components of registration. In the following subsection, we outline how we use data from non-Boston races to shed light on the mechanism behind changes in registration.

3.3 Performance & Participation Outside of Boston

Registering for the Boston Marathon requires a difficult-to-achieve qualification time. Therefore, estimates of δ in Equation (2) could be explained either by changes in the decisions of runners to register or by changes in (unobserved) qualification times. In addition, the interpretation of δ in Equation (1) as a manifestation of trauma could be threatened if runners treat the Boston Marathon differently than other marathons they participate in. As an example, it could be that runners who participate in the Boston Marathon consider the event more as a ceremony than as a competition.

To shed light onto the mechanisms behind both performance at and participation in the Boston Marathon, we collect data from three other significant marathons in the United States: the New York City Marathon, the Chicago Marathon, and the Marine Corps Marathon (in Washington DC and Arlington County, Virginia). These marathons appear to be the most commonly attended marathons by runners who also attended the 2013 Boston Marathon[10], with all three being geographically close, large, and two (NYC and Chicago) being Abbott World Marathon Majors. However, the Boston Marathon remains uniquely prestigious, with the median time in Boston (227 minutes) being approximately 45 minutes, or 15%, faster than the median times at the other races (272 in CHI, 286 in DC, and 266 in NYC). Even at the 25th and 75th percentiles, times in Boston are 30 and 60 minutes, or 13% and 19%, faster, respectively.[11]

Using these data, we identify instances of runners from our main sample competing in these races. We find approximately 50% of the 2013 non-charity Wave III runners in at least one of these other events. The total number of runner-marathon observations is approximately 2,500. Then, we re-estimate the parameters in Equation (1) to explore the effects of the attack on diverted runners when competing outside of Boston. In addition, we estimate a triple difference model that formally tests whether the effect of the bombing is different in Boston relative to marathons in other cities. If performances in marathons outside of Boston do not change for diverted runners following the attack, then we should not expect a difference in the likelihood of being able to qualify for the Boston Marathon. Any change in registration patterns at Boston can then be attributable entirely to choices made by the diverted runners. However, if the change in performances outside of Boston is similar to the change in performance at Boston, then changes in registration patterns are likely impacted by a change in the likelihood of qualifying for the Boston Marathon. We further test for indirect evidence of changes in the decision to participate in the Boston Marathon by looking at the likelihood of participating in one of these other three major marathons over time. If runners appear more in other marathons but less in Boston, this could be evidence of runners substituting participating in Boston

[10]We searched for the name of each runner in our sample on https://www.athlinks.com/. Tabulating the results, we found these three marathons to be, by far, the most well attended for the names of the runners in our sample.

[11]Results from the 2012 NYC Marathon are missing as it was canceled due to the landfall of Hurricane Sandy. In addition, the website we use to obtain this data, Marathon Guide, is missing age information for runners from the 2015 Chicago Marathon.

with other marathons. However, if runners attend both Boston and other marathons less frequently, the associated trauma could be deterring runners from large scale races in general.

3.4 Future Performance of Finishers

Thus far, we have discussed estimating the effects of being diverted relative to having finished the race prior to the attack. However, since these estimates represent the difference between two groups who both had some experience with the event, any potential psychological damage sustained by those who finished is obfuscated. Some finishers may have remained nearby the finish line, though they are generally hurried away by volunteers as mentioned earlier, and were thus exposed to chaos the diverted runners may not have experienced. Or, it could be that finishers face lingering internal "what if" questions that could impact their future performances. Put differently, participating in the race could cause trauma A while being diverted could cause trauma B. Therefore, finishers suffer A while diverted runners suffer $A + B$ (so long as these traumas are additively separable), and δ from Equation (1) represents B since A is common, and thus differenced out, among the two groups. To estimate A, or changes in performance within the group of runners who completed the race, we propose a slightly modified empirical design.

To start, we consider each non-charity Wave III runner in 2013, regardless of finish status. Each time we observe one of these runners in another marathon between 2009 and 2017, we match them to the runner in that marathon with the closest bib number that is also of the same gender and age. The one caveat is that we ensure that the matched runner did not participate in the 2013 Boston Marathon. Overall, we have 6,568 observations of a 2013 non-charity Wave III runner participating in another Boston Marathon between 2009-2017 (but not 2013). However, there are some instances where we cannot find a match, in which case we remove that runner-marathon observation from the sample, leaving us with 5,955 runner-marathon observations. Since each of these observations is matched to another, the entire sample contains 11,910 runner-marathon observations.

We estimate the average effect of participating in the 2013 Marathon on performance using an equation very similar to Equation 1. We model athlete i's performance in marathon t as:

$$\text{Time}_{it} = \delta\ (\text{In2013}_i \times \text{Post}_t) + \beta\ \text{Bib}_{it} + \mu_{m(i)} + \tau_{a(i)g(i)t} + \epsilon_{it} \tag{3}$$

where In2013_i is an indicator equal to one if the runner had participated in the 2013 Boston Marathon. Since we do not match runners to the same other runner in each race, we use $\mu_{m(i)}$, instead of μ_i, to represent runner-match specific fixed effect. Each observation of a runner who participated in the 2013 Boston Marathon shares a single runner-specific fixed effect, much like in Equation (1), however, each of their matched observations also share a single, but distinct, fixed effect. Concretely, if runner i competed in the 2013 Boston Marathon, each observation of this runner shares the same fixed effect. Then, each of the runner-marathon observations matched to runner i's observations are tied together by sharing the same fixed effect (as if they were the same individual).[12]

Estimates of δ from this equation represent the average effect of participating in 2013, regardless of whether the runner had been diverted. To separately estimate the effect on finishers and non-finishers, we decompose the indicator In2013_i into two mutually exclusive indicators: DNF2013_i and Fin2013_i (which is $\text{In2013}_i - \text{DNF2013}_i$). Then, returning to A and B from earlier, δ_{DNF2013} and δ_{Fin2013} would represent

[12] It is possible to use runner fixed effects (μ_i), but there would be many singletons and the statistical power of the analysis would be substantially reduced.

$A + B$ and A, respectively. Therefore, $\delta_{\text{DNF2013}} - \delta_{\text{Fin2013}}$ should be approximately equal the estimate of δ in Equation (1).

4 Results

4.1 Future Performance

As a first step, we estimate the difference in running times for 2013 non-finishers relative to 2013 finishers, and display the coefficients in Figure 4. We use 2011 as our baseline comparison year due to the exceptional nature of the 2012 marathon's heat, though our results are not sensitive to this decision as the zero-point would simply shift to 2012. Prior to 2013, we do not find any significant change in the differences between the two groups, which, according to Figure 3a, is approximately 15 minutes. Said differently, relative to the average time taken to run the marathon for 2013 finishers (approximately 240 minutes), 2013 non-finishers consistently take about 6% longer in each year from 2009 to 2012. This difference being steady is important, because it assuages concerns about there being potentially different aging/performance trajectories for the treatment and comparison groups. As alternative evidence, applying a log transformation to the outcome variable, which subtly changes the parallel trends assumption as discussed earlier, does not change our results. In 2013, however, this difference increased by 21.76 minutes (which is based on projected finish times for non-finishers) although this is only an estimate of the aforementioned performance shocks we use to determine treatment. These shocks can work in two directions, however, where some runners received a negative shock that sorted them into the treated group and other runners received a positive shock that allowed them to complete the race prior to the bombing. This also explains why, in Figure 3a, there is an increase in the average time for the non-finishers and a decrease in the average time for the finishers.

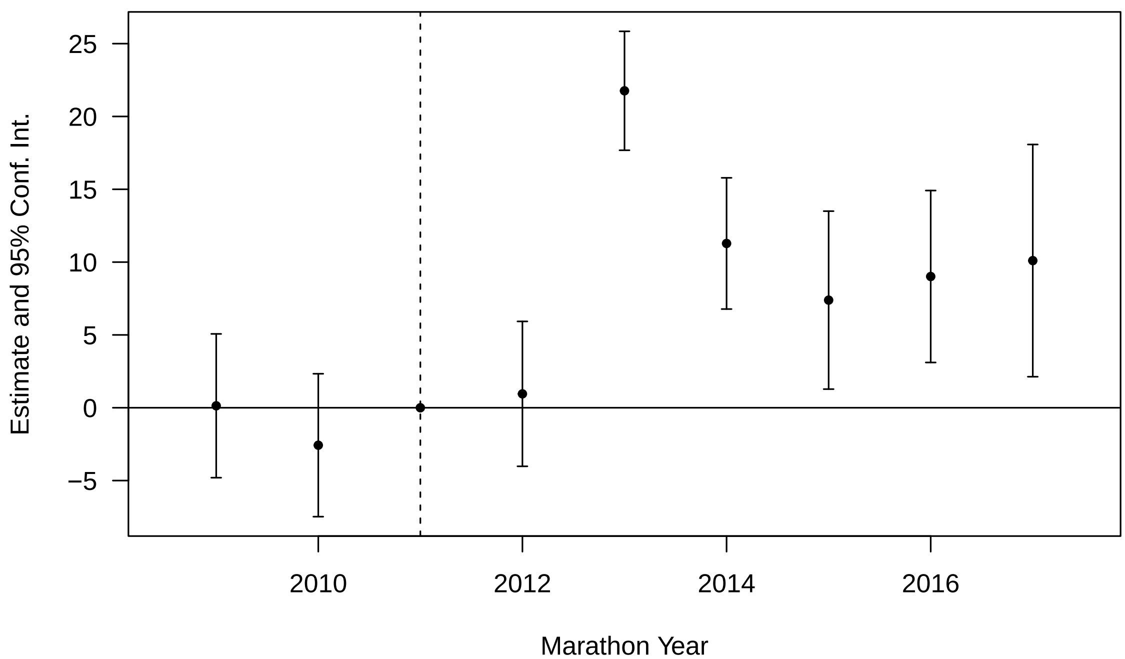

Figure 4: Event Study Coefficients From an Analysis of Finish Times

Notes: This Figure presents event study estimates and their corresponding 95% confidence intervals (heteroskedasticity robust). The outcome variable in this analysis is finish time, so these estimates represent the difference in finish times between finishers and non-finishers of 2013 relative to their difference in 2011. A positive coefficient in year t suggests that 2013 non-finishers are relatively slower in t than in 2011, relative to the 2013 finishers. The estimate for 2013 is partially based on projections for non-finishers made by the BAA.

After 2013, we assume that performance shocks are once again idiosyncratic, and what we estimated is the effect of being diverted in 2013 in subsequent races. We find a persistent, statistically significant increase in performance times in 2014, 2015, 2016, and 2017. In 2014, the BAA invited all diverted runners back to the Marathon in order to finish the race, and required neither a qualifying time nor a monetary payment.[13] It could be the case that a large portion of 2013's non-finishers only returned in 2014 due to the discount offered by the BAA and shirked in their training for the race, thus resulting in worse times. However, this point does not explain the pattern we find in 2015-2017 as these athletes once again need to qualify for the Marathon. Further, we control for each athlete's bib number, allowing us to somewhat partial out the effect of their qualifying time on performance in Boston.

The point estimates from 2014 through 2017 average to approximately 10 minutes, suggesting the initial difference of about 15 minutes grows to 25 minutes following the bombing. The standard deviation of times for 2013 finishers in non-2013 races is 24 minutes, making the effect size approximately 40% of a standard deviation. We also explore whether this effect is uniform or concentrated in a particular section of the race. In Figure 5, we break down the effect by approximate quarter of the race. Again, prior to the event, we find no change in the differences relative to 2011. However, in subsequent years, we find an interesting pattern where diverted runners exhibit progressively slower quarter-split times relative to runners who finished the

[13]Committing to running the Boston Marathon is not trivial. The Marathon itself costs approximately $250, plus airfare, hotel, time off of work (the Marathon is always on a Monday), etc.

2013 marathon. This pattern is mostly consistent in the four years following the bombing.

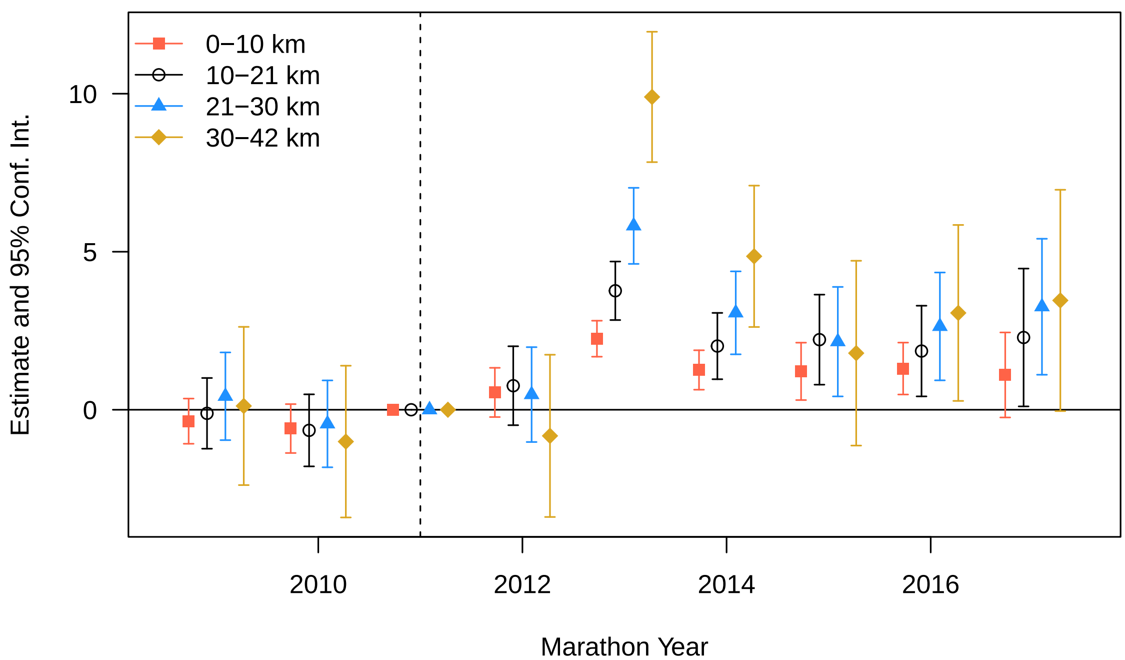

Figure 5: Event Study Coefficients From an Analysis of Split Times

Notes: This Figure presents event study estimates and their corresponding 95% confidence intervals (heteroskedasticity robust). The outcome variable in this analysis is a split time for quarter q of the marathon, so these estimates represent the difference in split times between finishers and non-finishers of 2013 relative to their difference in 2011. A positive coefficient in year t suggests that 2013 non-finishers are relatively slower in t than in 2011 over marathon section q, relative to the 2013 finishers. These results differ from those in Figure 4 in that they demonstrate whether differences in groups are especially pronounced in specific sections of the marathon. Further, while finish times in 2013 for non-finishers are based on projections, split times (aside from those in 2013) are not.

Table 2 tabulates δ from Equation (1) when varying the outcome variable and whether bib number is used as a control variable. It is important to note that we do not include data from 2013 in these regressions since those data are neither pre- nor post-period. The first two columns summarize the event study coefficients from Figure 4. Surprisingly, we find that controlling for bib number marginally increases the estimate of δ. This provides suggestive evidence that, although non-finishers are slower on race day, they are also given relatively lower bib numbers, suggesting that they register using relatively faster qualifying times. While this is initially counterintuitive, reversion to the mean offers a simple explanation. Runners who are traumatized, and consequentially slower, may be less likely to qualify for future Boston Marathons. Therefore, if a runner is indeed able to qualify, they may have done so via a "lucky" race, one where they received a positive performance shock. Then, once at the Marathon, they revert back to their mean, expected performance level, which is once again subject to the trauma associated with the 2013 Marathon. We provide more discussion about the idea of runners struggling to qualify for future Boston Marathons in Sections 4.2 and 4.3. We also find that the increase in finish times of approximately nine minutes for 2013 non-finishers corresponds to roughly 1,100 places within the race, highlighting the substantial impact on relative ranking.

Table 2: Effect of Being Diverted on Performance

	Finish Time		Finish Rank	
DNF2013 x Post	9.1748***	9.4617***	1097.4633***	1176.1726***
	(1.7651)	(1.7418)	(259.4237)	(251.5821)
Bib Number		0.0007***		0.2044***
		(0.0002)		(0.0355)
R2 Adj.	0.637	0.639	0.738	0.744
Num.Obs.	5223	5223	5223	5223

Notes: This Table contains difference-in-differences estimates of the effect of exposure to the Boston Marathon Bombing on future performance. While the first two columns explore the effects on time taken to complete the marathon, the third and fourth columns look at the overall ranking within the race. Both of these variables measure athlete success within the Boston Marathon in particular. We omit data from 2013 in each of these estimations since 2013 is neither pre- nor post- treatment. Without 2013, we have 5,223 observations, but our observations would increase to 7,788 if we included 2013's data. Heteroskedasticity robust standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.2 Future Participation

We report our estimates of the effect of the bombing on future Boston Marathon performance in Section 4.1. These results are conditional on runners registering for and competing in the Marathon, and can be thought of as an intensive margin. On an extensive margin, however, exposed runners may be more or less likely to register for future Marathons. In Figure 3b, we find suggestive evidence that the treated group is less likely to register in 2015-2017, but more likely to register in 2014. To test this more formally, we estimate the parameters in Equation 2 and present the estimates in Figure 6.

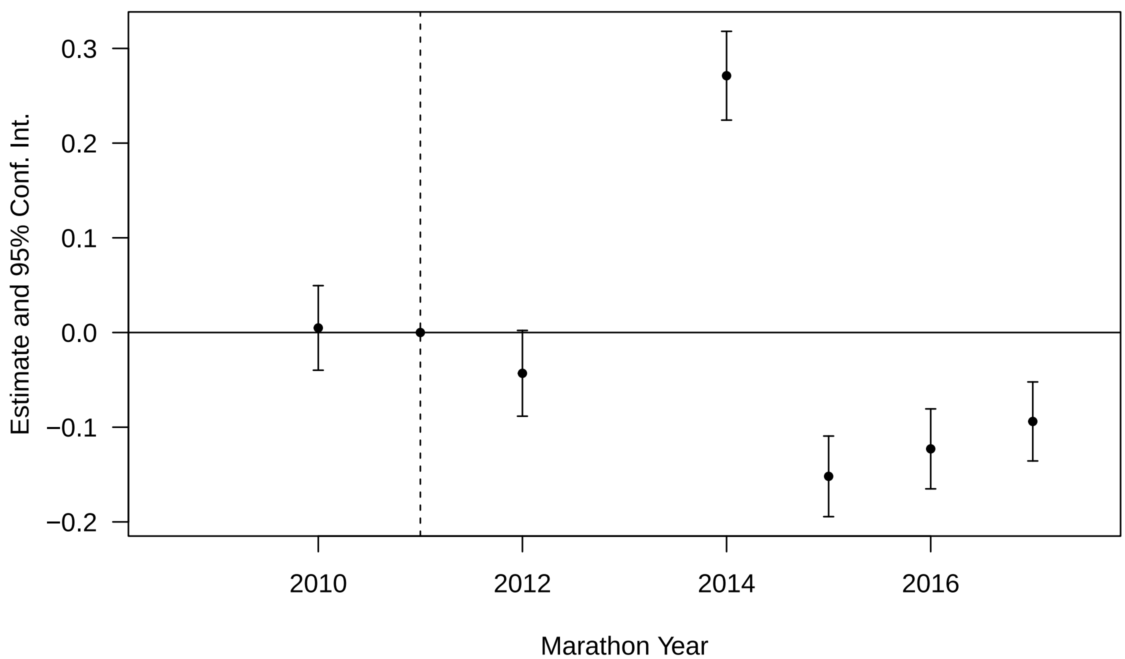

Figure 6: Event Study Coefficients From an Analysis of Registrations

Notes: This Figure presents event study estimates and their corresponding 95% confidence intervals (heteroskedasticity robust). The outcome variable in this analysis is whether a runner registered for marathon t, so these estimates represent the difference in the probability of registration between finishers and non-finishers of 2013 relative to their difference in 2011. A positive coefficient in year t suggests that 2013 non-finishers are relatively more likely to register in t than in 2011, relative to the 2013 finishers.

In Figure 6, we find a similar trend to what we see in Figure 3b: a substantial increase in the probability of returning in 2014 followed by a 10 percentage point reduction in the probability of returning in 2015, 2016, and 2017. Again, the sharp increase in 2014 is explained by the BAA having invited all diverted runners back to the Marathon, and neither requiring a qualifying time nor a monetary payment. Interestingly, the magnitude of the negative effect attenuates in the later post-bombing years. A potential explanation for the attenuation at the end of the sample is decay in the probability of registration in the comparison group rather than a recovery among exposed runners. Exposure to the bombing may have accelerated permanent exit among runners in the treated group, while similar exits occur more gradually in the control group. As a result, the control group increasingly resembles the treated group over time, mechanically narrowing the estimates.

In Section 3.2, we discuss how interpreting these results should be done carefully given the multiple reasons registration rates can change. First, and most intuitively, it could be that the runners' trauma simply caused a reduced demand for competing in large-scale marathons such as the Boston Marathon, but not for their typical training. If runners had developed a strong distaste for races like the Boston Marathon, they probably would not have nearly universally returned in 2014 despite the reduced costs. If, however, the reduced costs were enough to overcome their newly developed distaste, it would not necessarily explain their reduced performance conditional on participating. Of course, these could be two separate and independent trauma responses, but more evidence is required to say more. A second possibility is that simply having one's race disrupted is enough for them to develop a distaste for training (or running more broadly), thus reducing the probability of registering for Boston and lowering their performance. This possibility is subtle but important, as having a race prematurely ended due to factors outside of one's control after training for months on end could be a more meaningful contributor to worsening performance than bomb-related trauma. Again, with our results up to this point, we are unable to rule this out as a possibility. Another reason is that runners used free-entry into 2014 as an excuse to shirk during training, which would explain slower times in 2014, and as a substitute for future Boston Marathons. However, this explanation is only partially consistent with our findings so far, as it fails to explain the slower times in 2015-2017 (that are conditional on participating, which is difficult without proper training). This inconsistency allows us to rule out this explanation as a possibility. Each of these stories has to do with choice, but another mechanism remains: worse performances make it more difficult to qualify for future marathons.

Definitively determining which mechanism (choice or ability) is driving this registration result would require (inherently unobservable) information about the fastest qualifying time each runner could produce to determine whether they would qualify, and further information on whether they decided to submit their time to the BAA for qualification consideration. Unfortunately, neither of these items are available to the researcher. Instead, in Section 4.3, we provide additional results using data from other marathons that reveal an equivalent reduction in performance but only reductions in participation in cases where surpassing qualification standards is the most common way by which runners participate. This provides highly suggestive evidence that runners are not avoiding other high-profile marathons, but are instead struggling to qualify.

4.3 Performance & Participation Outside of Boston

As a next step, we test for changes in performance of non-finishers in marathons outside of Boston. Specifically, we track the finishers and non-finishers in our sample in three other marathons in New York City, Chicago, and Washington D.C., and compare their performances in these races over time. Figure 7 contains estimates from an event study version of Equation (1), and shows that non-finishers become slower in these races compared to finishers. Interestingly, the estimate of δ from Equation (1) using these data is 9.76 minutes (with a standard error of 3.22)[14], which is very similar to the coefficient in the first column of Table 2, providing strong initial evidence that the reduction in performance may be the same across all marathons.

[14]This estimate is the result of a regression with runner fixed effects, year-age-gender fixed effects, and a bib number control. We interact bib number by marathon in Table 3, which yields a slightly higher estimate of approximately 12.

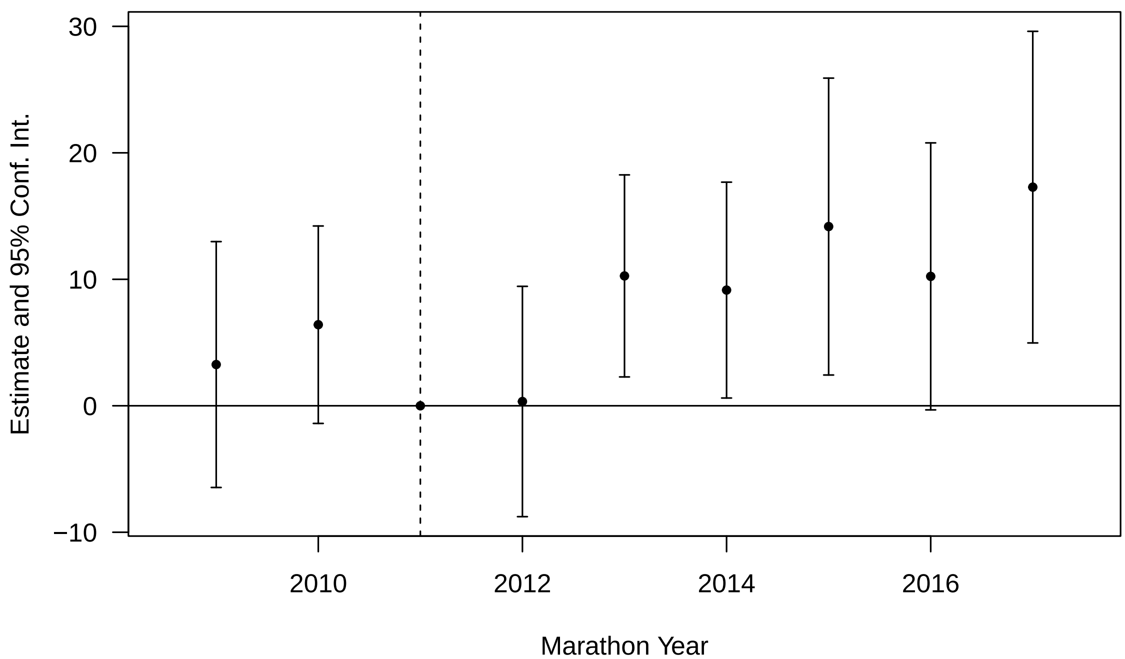

Figure 7: Event Study Coefficients From an Analysis of Finish Times From Other Marathons

Notes: This Figure presents event study estimates and their corresponding 95% confidence intervals (heteroskedasticity robust). The outcome variable in this analysis is finish time, so these estimates represent the difference in finish times between finishers and non-finishers of 2013 relative to their difference in 2011 when competing in one of the New York City, Chicago, or Marine Corps Marathons. A positive coefficient in year t suggests that 2013 non-finishers are relatively slower in t than in 2011, relative to the 2013 finishers.

To explicitly test whether the results in Figure 7 significantly differ from the results in Figure 4, we pool the data from Boston and the other marathons in a difference-in-difference-in-differences framework, and present the estimates in Table 3. Here, we find the triple difference interaction term to be statistically insignificant, suggesting diverted runners were not differentially affected in Boston relative to other races. Instead, the worsening performance of non-finishers occurs across all marathons, as indicated by the statistically significant interaction term between non-finishers and the post period.

Table 3: Effect of Being Diverted on Performance in Other Marathons

	At Boston	Outside of Boston	All Marathons
Bib Number (BOS)	0.0007***		0.0008***
	(0.0002)		(0.0002)
Bib Number (NYC)		0.0004**	0.0005***
		(0.0002)	(0.0001)
Bib Number (CHI)		0.0003	0.0005***
		(0.0002)	(0.0001)
Bib Number (DC)		0.0008**	0.0001
		(0.0004)	(0.0002)
DNF2013 x Post	9.4617***	11.9957**	7.8584***
	(1.7418)	(5.3876)	(2.5807)
DNF2013 x Boston			1.3214
			(2.2601)
DNF2013 x Post x Boston			1.9748
			(2.9658)
R2 Adj.	0.639	0.447	0.613
Num.Obs.	5223	2468	7691

Notes: This Table contains difference-in-differences estimates of the effect of exposure to the Boston Marathon Bombing on future performance in other marathons. Column 1 represents the baseline results for Boston while Column 2 pools together marathon results for the other races. Column 3 presents estimates from a triple differences specification. Standard errors are heteroskedasticity robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The result that diverted runners are slower in all marathons, not just the Boston Marathon, sheds light onto the mechanism behind the reduction in the probability of registration in Section 4.2. As mentioned, qualifying for the Boston Marathon is uniquely challenging due to its popularity and prestige. Moreover, beginning with the 2012 Boston Marathon, runners faced additional cutoffs below the published qualifying times, which were subtracted from each age-group's qualifying times. For example, the 2025 registration closed with a 6:51 cutoff, meaning that runners who beat the qualifying time by less than six minutes and 51 seconds were not accepted into the race despite technically qualifying. These cutoffs are not known prior to the deadline since they depend on both the quantity and quality of times submitted, as the BAA accepts the fastest runners per group until the predetermined number of spots are filled. The average cutoff time since 2012 is 2 minutes and 52 seconds, with the largest cutoff occurring in 2021 at 7 minutes and 47 seconds. Therefore, a reduction in speed due to psychological factors associated with the 2013 bombing of approximately 9 minutes (relative to the change for finishers) is certainly enough to remove a runner from contention barring a significant positive performance shock.

In addition to performing worse, and therefore being less likely to qualify for the Boston Marathon, runners may also choose to participate less overall or instead substitute towards other marathons aside from Boston (e.g., New York City, Chicago, or Washington DC).[15] We next look to see if non-finishers are less

[15] Anecdotally, the authors had a conversation with a non-finisher in 2013 (who is omitted from our sample due to being part of the "charity" group) who mentioned switching over to shorter races like half-marathons.

likely to participate in these other marathons after being diverted in 2013. We estimate δ from Equation (2) using data from each of the three other marathon races and present the estimates in Table 4. Interestingly, we only find significant reductions in the probability of returning to the Chicago Marathon but neither the New York City nor the Washington DC Marathons.

If runners' trauma caused them to avoid marathons in general, we would expect to see a reduction in the probability of attending each of the other marathons. If runners began to avoid only Boston, we would not expect to see reductions in any of the marathons. If runners chose to substitute towards other marathons, we would expect to see increases in the probability of attending other marathons. The pattern we find is consistent with neither avoidance of marathons overall, avoidance of Boston alone, nor substituting other marathons for Boston. In fact, we believe this pattern provides evidence of an impaired ability to qualify for high-profile marathons.

The qualification standards for the New York City Marathon are, in fact, even more difficult to achieve than the standards for the Boston Marathon. Aside from qualifying, runners can participate in the New York City Marathon either through the New York Road Runners organization[16] or via their lottery system. Therefore, in effect, participation in the New York City Marathon is not as dependent on a runner's ability to qualify as it is for the Boston Marathon. The Marine Corps Marathon, conversely, is open to all ability levels and does not require runners to qualify. Like the New York City Marathon, participation in Washington DC's major marathon does not hinge on runners being able to surpass an especially competitive qualification hurdle. The Chicago Marathon, however, uses relatively more attainable qualification standards than Boston, though only by about 5 minutes, which does not consider the additional demand-based cutoffs used by Boston described previously.[17] Therefore, much like for Boston, athletes who were previously able to qualify for Chicago (and avoid Chicago's lottery) may now struggle to do so given the effects of their trauma. Of course, given Chicago's slower standards, it is reasonable to expect a smaller coefficient than the one for Boston, which is precisely what we find in Table 4. Taken together, this provides strong evidence that runners were not avoiding large scale, high-profile events, but rather were less able to qualify for them (Boston, Chicago).

While these results are highly consistent with runners having struggled to qualify for future marathons, we cannot completely rule out the possibility that some runners may have avoided Boston in particular. We also cannot rule out the possibility that race disruption, not trauma, discouraged runners and caused them to develop a distaste for training, but it seems unlikely that runners would avoid training, but continue racing in high-profile events, especially in the long run. Moreover, in Section 4.4, we find that even race finishers, not just diverted runners, suffered slowdowns relative to individuals who were absent from Boston in 2013. These individuals were able to complete their races, but still suffered slowdowns, presumably due to their experience in 2013, providing additional evidence against the idea that race disruption is the key driving mechanism.

[16]Individuals must be members of the New York Road Runners, participate in nine "qualifying races," and volunteer at one "qualifying event" for guaranteed entry into the Marathon

[17]For example, the 2018 qualification standards for 50-54-year-olds were 10 minutes faster for Boston compared to Chicago while the standards for 55-59-year-olds were the same. However, in 2018, the average additional cutoff added to times for Boston was 3:23, making the totals 13:23 and 3:23.

Table 4: Effect of Being Diverted on Participation in Other Marathons

	Chicago	Washington DC	New York City	All
DNF2013 x Post	-0.0421***	-0.0056	0.0027	-0.0478***
	(0.0115)	(0.0081)	(0.0131)	(0.0163)
R2 Adj.	0.151	0.131	0.127	0.098
Num.Obs.	10576	11898	10576	11898

Notes: This Table contains difference-in-differences estimates of the effect of exposure to the Boston Marathon Bombing on future participation in other marathons. Columns 1-3 represent the probability of attending a specific marathon whereas Column 4 represents the probability of attending at least one. Heteroskedasticity robust standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.4 Future Performance of Finishers

In this subsection, we estimate the effects of participating in the third wave of the 2013 Boston Marathon by comparing those runners to runners who were absent. We estimate δ from Equation (3) and present the results in Table 5. Then, we present the corresponding event study estimates in Figure 8. In the first column of Table 5, we estimate δ, the effect of participating in the 2013 Marathon, to be an approximate 7.5 minute slowdown. This suggests that, on average, runners who participated in 2013 ran 7.5 minutes slower than someone of the same age, gender, and similar qualifying time. In the next column, we estimate the effects for finishers and non-finishers separately. We find that, despite them representing our comparison group in the main results, finishers are relatively slower by approximately 4 minutes while non-finishers are slower by almost 13. The difference we estimate in this Table, while using a different estimation strategy, is qualitatively similar to the estimates in both Tables 2 and 3. This is important, as it provides evidence that finishers also experienced reduced performance as a result of participation in the 2013 Boston Marathon.

Table 5: Effect of Participation in 2013 on Performance

	Finish Time	
In2013 x Post	7.5123***	
	(1.1570)	
Fin2013 x Post		4.0874***
		(1.2633)
DNF2013 x Post		12.8906***
		(1.5578)
Bib Number	0.0030***	0.0030***
	(0.0002)	(0.0002)
R2 Adj.	0.517	0.519
Num.Obs.	11910	11910

Notes: This Table contains difference-in-differences estimates of the effect of participation in the 2013 Boston Marathon as a non-charity Wave III runner on future performance. The first column contains an estimate of the average effect of participating in the 2013 Boston Marathon, regardless of outcome. The second column decomposes this estimate into the effect for those who finished in 2013 and those who were diverted. We omit data from 2013 in each of these estimations since the control group is comprised of individuals who did not attend 2013's Marathon. Heteroskedasticity robust standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 8 displays the event study coefficients that correspond to the second column of Table 5. Again, we find a relatively stable pre-2013 difference for both groups. While we find evidence of increased finish times in 2014 for the non-finisher group, we find almost precisely no change for the finishers. However, we find that both groups see statistically significant increases in their times in 2015-2017 relative to comparable runners in those races. We discount slightly the coefficients from 2014 due to the potentially unique nature of the Marathon that year, but the evidence for the other years is quite strong: both groups see marked increases in times, though the non-finisher group's increase is greater than that of the finisher group.

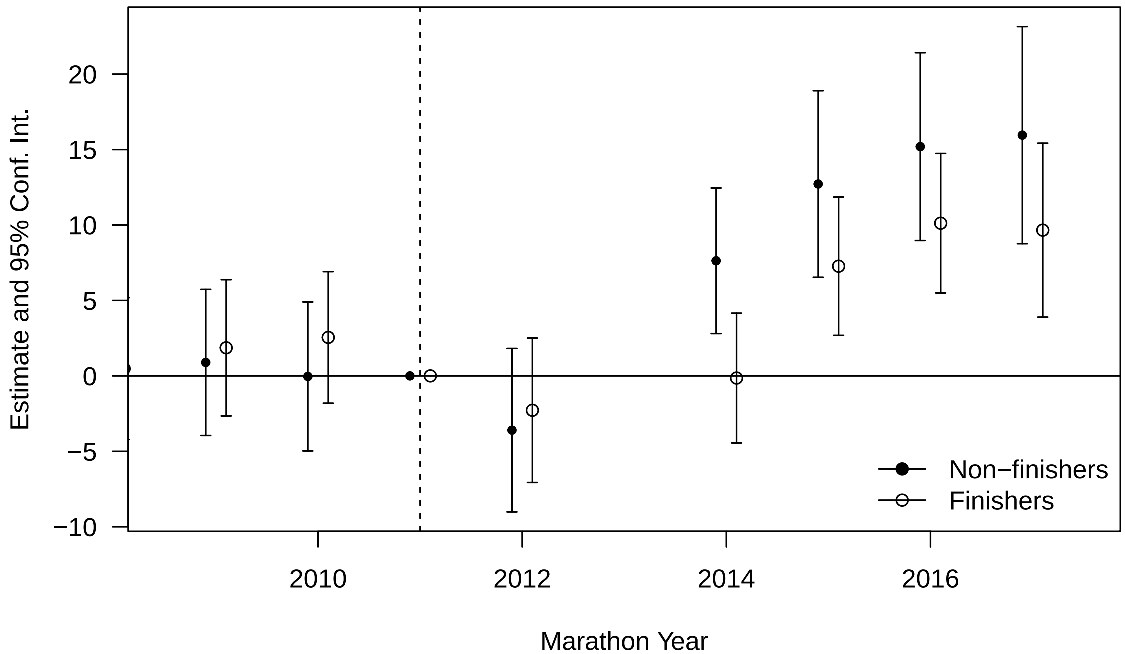

Figure 8: Event Study Coefficients From an Analysis of Finish Times for 2013 Boston Marathon Wave III Participants

Notes: This Figure presents event study estimates and their corresponding 95% confidence intervals (heteroskedasticity robust). The outcome variable in this analysis is finish time, so these estimates represent the difference in finish times between finishers/non-finishers and a matched runner of the same age, gender, and similar bib number relative to their difference in 2011. A positive coefficient in year t suggests that individuals are relatively slower in t than in 2011, relative to their matched runner.

4.5 Placebo Bombing

As a final step in our analysis, we apply our identification strategy to a completely different Boston Marathon and re-estimate our performance regressions. Of course, runners who happened to complete this marathon relatively slowly (in part due to what we propose are transient and idiosyncratic performance shocks) are not exposed to a bombing/diversion, allowing us to test the extent to which our identification strategy is the cause of our estimates (through some omitted confounding factor like injury).

We choose to study the 2005 Boston Marathon for our placebo analysis. The BAA publishes results back to 2001, which makes 2005 an attractive selection to maintain four years of pre-period data and minimize the overlap with the 2013 time frame of 2009 to 2017. We also note that the runners are mostly distinct as only 250 out of 2,565 runners in the 2013 analysis sample also ran in the 2005 Marathon. Serendipitously, 2004 was similar to 2012 with its exceptional temperature, making the conditions in 2005 even more similar to those in 2013.

To construct our placebo sample, we go through a very similar process as constructing our main analysis sample. First, we eliminate all 2005 charity runners. Next, we classify runners as "treated" if they run slower than a 4:10 marathon time (250 minutes), which was approximately the cutoff in 2013. Then, we select runners with bib numbers between 11,725 and 16,600, making the average finish time and percent of runners in the treated group equivalent to that in the 2013 sample (250 minutes and 43% treated).[18] Finally, we estimate the parameters in Equation (1) using this sample, and display the event study coefficients in Figure 9.

In this Figure, we find an expected large increase in the difference in times in 2005, which is akin to the large difference we find in 2013. However, unlike the other analysis, we do not find sustained differences following 2005. Tabulating these coefficients in the same way we do in Table 2 yields a coefficient of 1.58 with a standard error of 1.92. Of course, the difference between these two coefficients is approximately 8 minutes, which is obviously a practically significant difference. Performing a difference in means test between this coefficient and the coefficient in the second column of Table 2 yields a t-statistic of 3.036, suggesting a statistical, as well as a practical, difference between these two estimates.

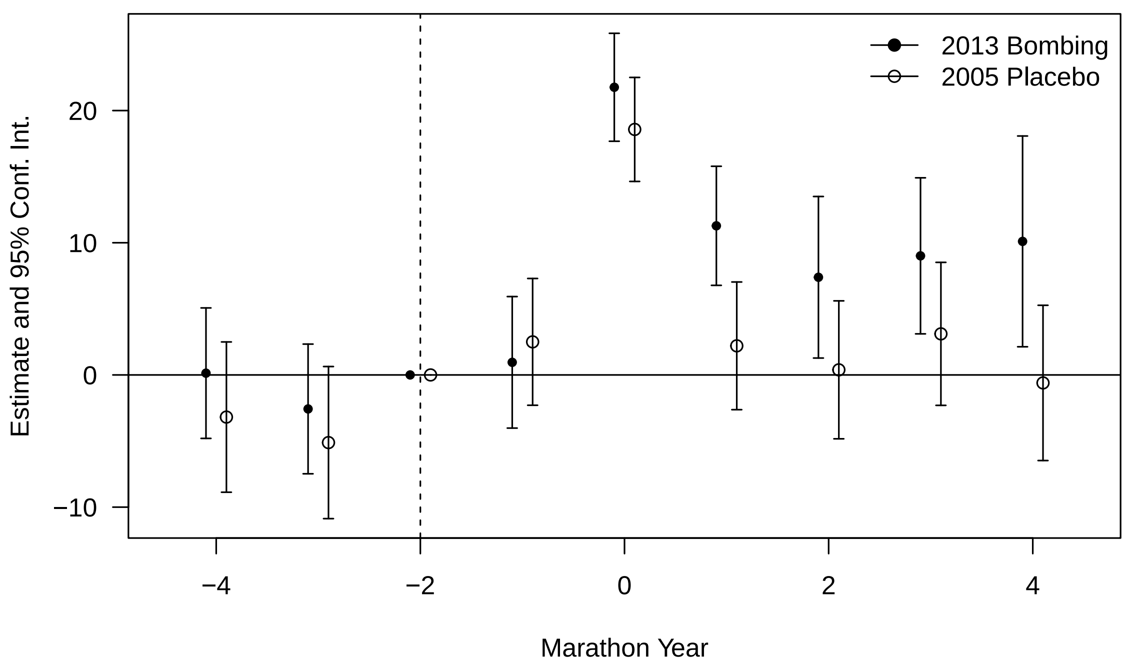

Figure 9: Event Study Coefficients From a Placebo Analysis of Finish Times

Notes: This Figure presents event study estimates and their corresponding 95% confidence intervals (heteroskedasticity robust). The outcome variable in this analysis is finish time, so these estimates represent the difference in finish times between "treated" and "control" individuals in 2005 relative to their difference in 2003. A positive coefficient in year t suggests that 2005 "treated" individuals are relatively slower in t than in 2003, relative to 2005 "control" individuals. We also include event study estimates and standard errors from Figure 4 as a comparison.

Lastly, we apply the matching strategy used in Section 4.4 to estimate the effects of the 2013 attack on both finishers and non-finishers on the sample of 2005 runners used above. We replicate the event study from Figure 8 and present the 2005 estimates in Figure 10. Here, we find consistent evidence of a null effect for both "finishers" and "non-finishers" in this context, providing further support that the 2013 findings are

[18]This sample contains a total of 4,875 runners, whereas the 2013 sample only contains 2,565 runners. We also construct a sample with an equivalent number of runners rather than a sample with equivalent summary measures and our results remain qualitatively similar.

indeed driven by psychological trauma suffered by the runners who participated in the Marathon.

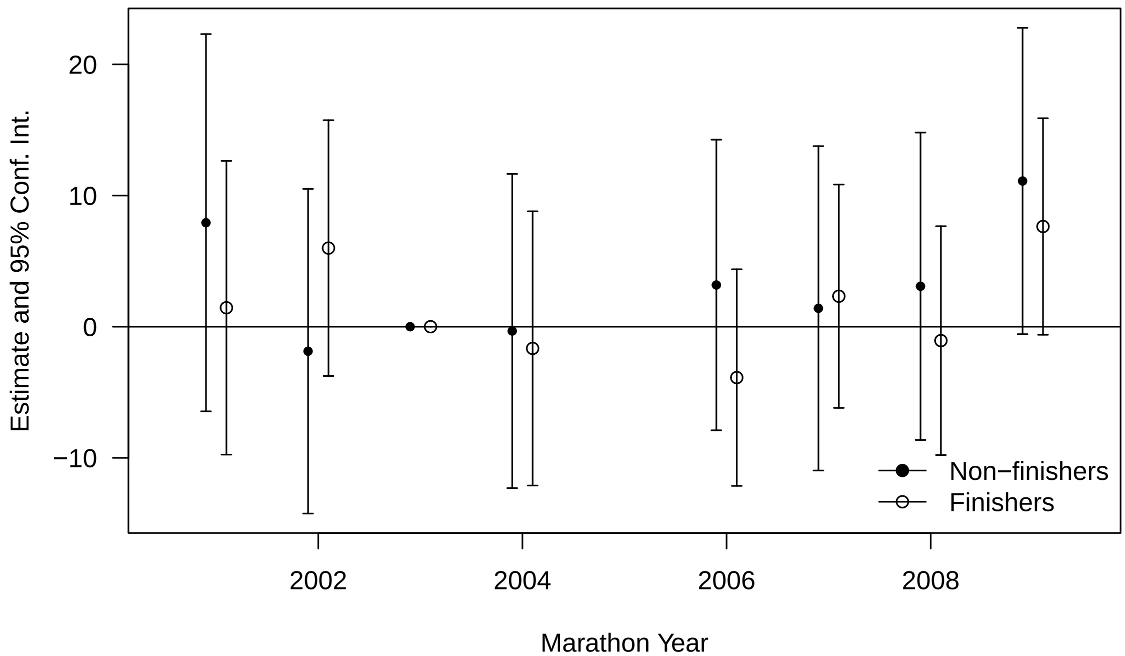

Figure 10: Event Study Coefficients From an Analysis of Finish Times for 2005 Boston Marathon Participants

Notes: This Figure presents event study estimates and their corresponding 95% confidence intervals (heteroskedasticity robust). The outcome variable in this analysis is finish time, so these estimates represent the difference in finish times between finishers/non-finishers and a matched runner of the same age, gender, and similar bib number relative to their difference in 2003. Of course, finishers/non-finishers are defined as whether the runner completed the 2005 Marathon faster/slower than 250 minutes. A positive coefficient in year t suggests that individuals are relatively slower in t than in 2003, relative to their matched runner.

5 Conclusion

We examine how direct exposure to the 2013 Boston Marathon bombing affected the subsequent performance and participation of marathon runners who were present. Exploiting the sudden division of the field into those who finished before the explosions and those who were diverted from the finish line, we find that diverted runners ran approximately nine minutes slower, a roughly four percent decline, in subsequent marathons. The effect persists across other major marathons, indicating a generalized response rather than something specific to Boston. Diverted runners were also about ten percentage points less likely to register for future Boston Marathons, a decline we attribute to their deteriorating performance rather than an intention to avoid the race. Finishers, too, exhibit measurable performance declines relative to comparable runners who were not present in 2013, suggesting that our main estimates represent a lower bound on the full cost of exposure.

These findings extend the existing literature on the effects of terrorism on well-being and economic behavior (Gaibulloev and Sandler, 2023). First, we are able to measure concrete individual-level outcomes while most other studies rely on either subjective surveys or aggregated data. Marathon running provides an advantage over other settings since the task is well-defined, comparable, objectively measured, and repeated. In addition, individuals who participate in World Marathon Majors like the Boston Marathon take running seriously and invest substantial resources into competing, providing a clean parallel to the labor market. Second, we find persistent declines in both performance and participation, which represents an important distinction from other studies in the literature who find effects that are transitory (Metcalfe et al., 2011; Clark et al., 2020). Third, our results help to explain important findings in the literature of reduced earnings and employment following terrorist attacks (Brodeur and Yousaf, 2025). Our results suggest that these outcomes are driven by a reduced production capacity rather than an intentional choice to reduce labor supply.

Of course, several limitations qualify these conclusions. First, marathon runners are a selected population of healthier, higher-income, and more achievement-oriented individuals than the general public (Braschler et al., 2024; Partyka and Waśkiewicz, 2024). Further, terrorist attacks are heterogeneous, and the Boston Marathon Bombing is just a single example. Therefore, while the qualitative lesson that direct exposure produces persistent, involuntary declines in productive capacity, the specific magnitudes may not transfer to the broader workforce. Second, we observe finishing times only for the marathons in our data and only for runners who enter and complete them, so post-bombing performance is conditioned on continued participation. Runners who stop racing altogether drop out of our intensive-margin estimates, and if the most severely affected are also the most likely to exit, our performance estimates understate the true effect. Third, to participate in the Boston Marathon, a runner must both choose to participate *and* qualify for the race. Unfortunately, our participation measure only captures instances when both of these conditions are met. We infer that runners are pushed out since we see participation reductions in races where both conditions need to be met, but no change in races where only the choice to race needs to be met. However, to definitely decompose this, one would need to send yearly surveys to the runners to ask 1) if they would like to participate in the Boston Marathon and 2) what their fastest time of the year is. Lastly, given our data, we are unable to pin down the precise behavioral change responsible for the reduction in competitive performance. It could be that the bombing caused psychological distress that reduced the productive capacity of the athletes, or impacted their risk preferences causing them to adopt more conservative race strategies, change their training habits, etc. More research is needed to fully characterize the internal mechanisms behind the measured performance declines (Caliendo et al., 2015; Metzl and MacLeish, 2015).

Finally, our results speak directly to the design of policy responses to mass traumatic events. While prevention is the first-best response, perfect surveillance is prohibitively costly, and so support services must be in place for the unexpected victims of terrorism. The Marathon Bombing Victim Assistance program, established after the 2013 attack, provided both immediate and sustained support to victims, survivors, their families, and first responders, and our findings reinforce its central lesson: victim services must be flexible and sustained over long horizons to be effective (Naturale et al., 2017). The performance and participation gaps we document remain evident years after the attack and are present even among runners who finished before the detonations, suggesting that psychological trauma spills well beyond those most visibly injured and well beyond the acute phase of the crisis. Just as many physical injuries from terrorism (e.g., hearing loss, traumatic brain injury, chronic pain) emerge or worsen gradually (Beinecke, 2014), so too, our results suggest, do the productivity costs. Programs designed only for the immediate aftermath are likely to underestimate both the breadth and the duration of the damage they aim to address.

References

Abadie, A. and J. Gardeazabal (2003, March). The economic costs of conflict: A case study of the basque country. *American Economic Review 93*(1), 113–132.

Abadie, A. and J. Gardeazabal (2008). Terrorism and the world economy. *European Economic Review 52*(1), 1–27.

Ahern, K. R. (2018, February). The Importance of Psychology in Economic Activity: Evidence from Terrorist Attacks. Working Paper 24331, National Bureau of Economic Research.

Akay, A., O. Bargain, and A. Elsayed (2020). Global terror, well-being and political attitudes. *European Economic Review 123*, 103394.

Arsenault, M. and H. Krueger (2023). A decade after the marathon bombing, reflection and renewal. https://apps.bostonglobe.com/metro/2023/04/boston-marathon-bombing-10-years/. The Boston Globe, published April 14, 2023. Accessed: 2025-12-14.

Barnes, K. R. and A. E. Kilding (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.

Beinecke, R. H. (2014). Addressing the Mental Health Needs of Victims and Responders to the Boston Marathon Bombings. *International Journal of Mental Health 43*(2), 17–34.

Benmelech, E., C. Berrebi, and E. F. Klor (2010). The economic cost of harboring terrorism. *Journal of Conflict Resolution 54*(2), 331–353.

Bogan, V. L., D. R. Just, and B. Wansink (2013). Do psychological shocks affect financial risk taking behavior? a study of us veterans. *Contemporary Economic Policy 31*(3), 457–467.

Boston Athletic Association (2014). The First Boston Marathon. http://216.235.243.43/races/boston-marathon/boston-marathon-history.aspx. Archived April 24, 2014; accessed December 16, 2025.

Braschler, L., M. Thuany, C. A. B. de Lira, V. Scheer, P. T. Nikolaidis, K. Weiss, and B. Knechtle (2024). Personality of marathon runners: a narrative review of recent findings. *EXCLI journal 23*, 441.

Brodeur, A. (2018). The Effect of Terrorism on Employment and Consumer Sentiment: Evidence from Successful and Failed Terror Attacks. *American Economic Journal: Applied Economics 10*(4), 246–282.

Brodeur, A. and H. Yousaf (2025). On the economic consequences of mass shootings. *Review of Economics and Statistics 107*(1), 109–124.

Caliendo, M., D. A. Cobb-Clark, and A. Uhlendorff (2015, 03). Locus of control and job search strategies. *The Review of Economics and Statistics 97*(1), 88–103.

Callen, M., M. Isaqzadeh, J. D. Long, and C. Sprenger (2014, January). Violence and risk preference: Experimental evidence from afghanistan. *American Economic Review 104*(1), 123–48.

Clark, A. E., O. Doyle, and E. Stancanelli (2020). The Impact of Terrorism on Individual Well-Being: Evidence from the Boston Marathon Bombing. *The Economic Journal 130*(631), 2065–2104.

DeBree, C. S. (2014, March). Boston Marathon runners: ‘You have to go back’. https://www.phillyburbs.com/story/news/2014/03/29/boston-marathon-runners-you-have/17722298007/.

Gaibulloev, K. and T. Sandler (2023). Common myths of terrorism. *Journal of Economic Surveys 37*(2), 271–301.

Heitmueller, A. (2005). Unemployment benefits, risk aversion, and migration incentives. *Journal of Population Economics 18*(1), 93–112.

Hoogkamer, W., S. Kipp, J. H. Frank, E. M. Farina, G. Luo, and R. Kram (2018). A comparison of the energetic cost of running in marathon racing shoes. *Sports Medicine 48*(4), 1009–1019.

Kim, Y.-I. and J. Lee (2014). The long-run impact of a traumatic experience on risk aversion. *Journal of Economic Behavior Organization 108*, 174–186.

Kirk, C. (2013, April). Did the Boston Marathon Bomber Time the Explosions to Guarantee the Most Possible Casualties? https://www.slate.com/blogs/the_slatest/2013/04/16/boston_bombs_timing_were_the_marathon_attacks_timed_to_maximize_casualties.html.

Markle, A., G. Wu, R. White, and A. Sackett (2018). Goals as reference points in marathon running: A novel test of reference dependence. *Journal of Risk and Uncertainty 56*, 19–50.

Metcalfe, R., N. Powdthavee, and P. Dolan (2011). Destruction and Distress: Using a Quasi-Experiment to Show the Effects of the September 11 Attacks on Mental Well-Being in the United Kingdom. *The Economic Journal 121*(550), F81–F103.

Metzl, J. M. and K. T. MacLeish (2015). Mental illness, mass shootings, and the politics of american firearms. *American Journal of Public Health 105*(2), 240–249. PMID: 25496006.

Montgomery-Marks, P., S. Bandyopadhyay, C. B. Weisman, and T. Bose (2025). Economic burden of PTSD in the UK: a systematic review and economic analysis. *BMJ Open 15*(7), e084394.

Naturale, A., L. T. Lowney, and C. S. Brito (2017). Lessons Learned From the Boston Marathon Bombing Victim Services Program. *Clinical Social Work Journal 45*(2), 111–123.

Palacios-Huerta, I. (2025). The beautiful dataset. *Journal of Economic Literature 63*(4), 1363–1423.

Partyka, A. and Z. Waśkiewicz (2024). Motivation of marathon and ultra-marathon runners. a narrative review. *Psychology Research and Behavior Management*, 2519–2531.

PBS NewsHour (2014, April 21). In celebration and resilience, Boston Marathon runners race again. https://www.pbs.org/newshour/show/celebration-resilience-boston-marathon-runners-race. Video clip and transcript on PBS.org.

Searing, E. A., F. Rios-Avila, and J. D. Lecy (2013). The impact of psychological trauma on wages in post-conflict Bosnia and Herzegovina. *Social Science & Medicine 96*, 165–173.

Shai, O. (2022). Out of time? the effect of an infrequent traumatic event on individuals' time and risk preferences, beliefs, and insurance purchasing. *Journal of Health Economics 86*, 102678.

Shany, A. (2025). Too Scared for School?: Effects of Terrorism on Students' Achievement. *Journal of Human Resources 60*(3), 977–1018.