Market Reactions to Railroad Accidents

Jerimiah Koll

The University of Chicago

ABSTRACT

Railroad accidents have become more common and seen in recent years. These accidents impose many costs on the railroad, from legal costs, cleanup, damages, and shifts in usership. While there has been a great deal of research into the effects of aviation accidents on airline stock prices, I only found two papers investigating the effects railroad accidents on stock prices of the railroad. However, there are methodological concerns for these previous studies. As such I implement an event study to investigate the effect of railroad accidents on the stock price of the railroad. I find little impact on the stock prices of the railroads, contradicting previous research. I discuss potential causes of these differences and include areas for further research involving availability of information.

INTRODUCTION

Railroad accidents have been the subject of much attention in recent months. Ever since the Norfolk and Southern derailment in East Palestine railroad accidents are reported on more often. While the causes of accidents are well understood, both at the individual and industry level, there's little in the way of understanding how an accident affects the railroads' balance sheet. Accidents are costly for the railroad as they pay for any cleanup, damages, legal fees, and may lose service and customers.

Given that railroad accidents and derailments cause tens of millions of dollars to railroads annually, one would expect existing literature on this and similar topics. Anecdotal evidence seems to indicate that these costs are reflected in the stock price of the railroad, and that railroad markets are efficient. Norfolk and Southern stock prices fell 6% over one week after a February 3rd derailment in East Palestine. Another instructive case is the Montreal, Maine, and Atlantic, which went bankrupt after the Lac-Megantic rail disaster which killed 47 people, released upwards of 100,000L of crude oil into a nearby river, and destroyed the entirety of downtown Lac-Megantic.

Unfortunately, there has been little investigation of this question in economics journals. Walker et. al. (2006), inspired by methods used to estimate the effects of aviation disasters on the stock market uses a synthetic control design to estimate the effect of an accident on the railroad that had the accident. They find that railroads suffer a 2-3% decrease in stock prices after roughly three trading days after the accident. Walker et. al. (2007) comes to similar conclusions and notes that much of this variation in stock prices comes from potential legal liability to the railroad.

I contrast this synthetic control as previously implemented with an event study, deciding that an event study is a better choice for model identification. I implement an event study with heterogeneous treatment effects from damages, injuries, deaths, HAZMAT releases, and accident cause. I exploit the fact that accidents occur more or less randomly across the US rail network and therefore serve as potential sources of exogeneous variation in the stock price. If the effect length is short or non-existent, then these markets are responding to new information rapidly, and are efficient on some level. With this event study I find little evidence of railroad specific or industry-wide effects for a variety of effect lengths. Across a variety of event horizons and model specifications there was little explanatory power for the day-to-day changes in stock prices, which indicates that railroad stock prices are unaffected accidents. These results can also indicate a level of market efficiency, as knowledge of the accident doesn't affect the underlying profitability of the firm very much.

This event study approach is surprisingly new in the literature, as previous research primarily uses synthetic control methods. Synthetic control methods can face difficult decisions in generating counterfactual data is a difficult process, and this problem is exasperated by potentially problematic decisions with regard to a chosen counterfactual data generating process. By implementing an event study, estimates of treatment effects are point identified, which allows causal interpretation of results where previous methods may have been lacking.

THEORY AND BACKGROUND

The argument for these accidents changing the stock price is simple, these exogenous shocks impose costs on the company. The railroad has to replace any damaged or destroyed locomotives or cars, pay for any cleanup, which can be particularly costly in cases of HAZMAT spills, legal fees depending on the cause of the accident, as well as potential damage to nonrailroad property. Accidents can also have additional costs to the railroad, by damaging locomotives and rail, or creating a perception of poor safety culture around a company, both of which could lead to a loss of customers. Accidents could also impose costs on other railroads, as an accident could lead to a reevaluation of the costs and benefits of shipping goods via rail relative to other methods. This would lead to a shift away from rail shipping and reduce stock prices across the industry.

However, there's also reason to believe that accidents have little impact on the stock price of the railroad. First, if potential investors have good knowledge of the likelihood of an accident, any potential accident should be priced into the current valuation of the railroad. Of course, this isn't to say that there would not be a shift, as the valuation moved from the expectation of an accident to the actual realization of the accident, but this is reason to believe that such effects would be smaller than initially thought. In the exact opposite direction, if investors don't hear about an accident, there cannot be an effect on the stock price of the railroad. This may be the case for small and lower cost accidents, but unlikely for large accidents that cost the railroad millions of dollars. Third, all major railroads in the US and Canada are insured, so many of the costs of an accident are covered and not realized by the railroad. While insurance premiums may increase due to accidents, such effects are likely to be small, again, reducing the potential effect of the accident on the price. Fourth, major railroads hold large cash reserves in addition to their insurance. While accidents may be costly to the railroad, they can simply dip into this fund, and profitability is unchanged, so stock prices are unaffected. Finally, the vast majority of track in the

US and Canada is operated by six companies.[1] Many areas are served by one or two railroads so any substitution between one railroad to another becomes impossible.

Another potential question is market efficiency. If stock prices move rapidly after an accident, knowledge of the accident has moved through the market and has been incorporated into the valuation of the company. This is also the case if there is no noticeable movement in the stock price of the railroads after the accident. Slow and sustained movement of the stock price after an accident indicates that not all investors have good knowledge of the railroad, as the effect of an accident takes a while to move throughout the system. Such a case would indicate inefficient railroad asset markets.

LITERATURE REVIEW

Synthetic control methodology has been applied often in finance and has been used often in estimating the effects of aviation accidents on airline stock price. Ho et. al (2013) finds airline stock prices fall by roughly 2%. They find significant variation in spillover effects. Spillover effects happen when the reduction in stock prices in one airline results in a shift in other airlines' stocks. They find that small accidents with few deaths actually benefited competitors, likely because people are shifting towards these other airlines that have a better perceived safety. Large accidents with many deaths see the industry take a hit as a whole, likely because it results in a reevaluation of the risks of air travel as a whole.

---

[1] These six railroads are "Class 1" railroads. Class 1 railroads are railroads that annual revenue greater than $250,000,000 1992 USD. As of mid-2023 the class 1 railroads operating in the US are Amtrak, Burlington Northern and Santa Fe, Canadian National, Canadian Pacific, CSX, Norfolk and Southern, and Union Pacific. Besides Amtrak, all the Class 1 railroads are exclusively freight railroads.

There is reason to believe that this relationship would not carry over to the railroad industry, particularly that of the US. First, there is far less competition in the railroad industry, the vast majority of railroad miles are operated by the Class I railroads, so people shipping by rail have little choice between rail carriers, this means future profitability is only affected in as much as the accident imposes direct costs on the company. Second, railroad customers are probably more rational actors than airline customers. Individuals are less likely to rationally evaluate the risks of an accident when flying than a company would be when shipping goods by rail. Finally, railroad accidents are likely less publicized than aviation accidents, in part because aviation accidents just look worse than a comparatively bad railroad accident. This means that investors don't have the knowledge of the accident so stock prices will move slowly, if at all.

Up until now there have been two papers investigating the effect of a railroad accident on the stock price of the railroad. Walker et. al. (2006) and Walker et. al. (2007) answers this question using what is essentially synthetic control. First, they regress stock price pre-accident on S&P 500 to get a projection of the price if the accident did not happen. This should theoretically sweep out any potential effect of the accident on the stock price. This predicted stock price is then used as an untreated group to estimate any effect on the stock price by the accident. They do this by regressing the difference between the actual and predicted prices on accident attributes. They perform this analysis on 26 major rail accidents from January 1993 to December 2003. Crucially, the dependent variable is the sum of the difference between the actual and predicted prices after the accident. This method finds that accidents result in a roughly 2% drop in the stock price of the railroad in the first 2-3 days after the accident, with the stock being indistinguishable from the previous levels beyond a week at $p = 0.05$. They also find that the market is able to discriminate based on the severity of the accident. More damages, fatalities, injuries, HAZMAT spills, and accidents caused by negligence all result in more severe hits to the stock price, while accidents that can partially be attributed to poor weather conditions and mechanical issues with the train see a milder reaction. The second of these papers uses similar methodology to the previous one and finds similar results, focusing on potential legal costs of the accident as explaining the variation in stock prices after the accident.

METHODOLOGY

Given the previous literature's reliance on synthetic control methods, I hope to explain my deviation from the previous work when implementing an event study. Synthetic control methods have certain assumptions necessary for proper identification and can fall to spurious correlation quite easily if these assumptions are not met. Leaving the discussion there would be disingenuous, so I include a deeper discussion of potential shortfalls of synthetic control methods as they are currently implemented.

Synthetic control methods are only as good as the choice of control units. Having the only control unit as the S&P 500 (or TSX 300 for the Canadian railroads) may cause problems when trying to generate post-treatment data, even if there is good performance on the pre-treatment data. Synthetic control takes a list of data that is similar to the treated group and matches it based on observed characteristics. By reducing this list of data to one unit, this becomes DID between the treated group and a scalar multiple of the control, except there is little reason to believe that the observables of these two groups will match, making identification impossible. The smaller the list of potential controls, the worse fit there is going to be on observables of the railroad.

In the same vein, selecting on more observables is generally preferred to matching on fewer. But in this case, the only observable is the price. We can match prices as much as we want, but without similar characteristics, there is no reason to believe that the synthetic data is going to behave in the same way the true counterfactual would. Whether or not variation in stock prices is exogeneous after these controls is the major identifying assumption of synthetic control methods and given that selection is done using one control variable, this seems unlikely.

Third, there are endogeneity concerns when using the full S&P 500 or TSX 300 as the control group. All large railroads are listed in their respective index, so the index is going to be in some way dependent on the railroads' stock price. I grant that such endogeneity concerns are small, as it's one company on the S&P 500, but when alternative methods exist that don't suffer from these problems, there needs to be a good reason to use them as the control group.

Given identification concerns with synthetic control methods as currently used in the literature, I estimate the effect of the railroad accident by implementing an event study methodology. I take stock prices of the class one railroads three days before and after the accident. The treated group is then the railroad that had the accident and the untreated group is the rest of the railroads. I only take data around the accident, to prevent a few problems. First, stock prices months before the accident have very little to do with the stock prices a day before given stock prices three days before. This is because stock prices behave in like Markov chains. Assuming that there are no exogenous shocks, our expectations of future stock prices are determined by the current price, and knowledge of prices before isn't useful.

Second, I do not want potential effects longer than three days tainting my untreated group. If the true effect length is five days long, then the untreated group has the last two days of a treatment effect included, and the model is misidentified. This truncation prevents estimating the time baseline equation as a combination of treated and untreated observations.

$$Y_{it} = \gamma_0 + \gamma_1 X_i + \gamma_2 X_t + \beta_1 D_{it} + \beta_2 Z_{it} + \beta_3 D_{it} Z_{it} + \epsilon_{it}$$

Here, Y_{it} is a measure of the stock price of the railroad, X_i is a fixed effect that controls for specific railroads. Likewise, X_t is a time fixed effect, estimated each year. D_{it} is an indicator variable that is 1 when the i^{th} railroad has had an accident in the past three days.[2] If the accident took place after 12pm EST, I encode this as happening the following day, as the market likely doesn't have the time to respond to the event. I chose a 3 day event horizon as previous literature finds the largest cumulative effect on the third day. If any effect exists, we should find it by the third day. Finally, z_{it} is a vector of accident characteristics, these are number of injuries, number of deaths, logged damages, an indicator variable for whether or not there was a HAZMAT leak, relative time indicators for up to three days after the accident, and indicators for accident cause. Accident causes are split into six categories, mechanical, infrastructure, negligence, user error, outside factors, and unknown.

By having accident characteristics as their own variables in the model, independent of the treated group, lets me estimate the effect of the accident on the untreated group. This separates treatment effects on the treated group from any potential treatment effect on the untreated group. This is because there may be spillover effects, where an accident on one railroad affects the stock price of another. Again, these effects may exist if an accident shifts traffic away from the railroad

[2] Additional robustness checks using 1-, 5-, and 10-day event horizons are presented in the appendix with similar results.

that had the accident or railroading entirely. If there are no spillover effects, these coefficients should be zero as it implies that the untreated stock prices move normally after an accident.

Before continuing I want to explain my choice for y_{it}. Usually, this would be logged to interpret coefficients as a percent change in prices. However, if stock price level is associated with probability of an accident, it's going to give us biased results without including a railroad size interaction term. This quickly becomes intractable as the number of regressors becomes large relative to the number of observations, removing much of the power of the model.

Thankfully, a solution exists. If we think of stock prices as a martingale, we can take the change from the day before as our y_{it} to get rid this problem. Any significant movement from zero can then be taken as caused by some exogeneous shock, and it also prevents the bias caused by railroad size. As such, y_{it} is the percent change in the stock price from the previous day. This means I can still interpret the coefficients one the treatment variables as causing a percent change to the stock price.[3]

DATA

I take the list of publicly available NTSB reports of railroad incidents on the six class one railroads in the US from January 1st, 2010 to December 31st, 2021, as the sample of accidents.[4] This is by no means an exhaustive list of railroad incidents, but the NTSB investigates railroad

[3] I run a regression to test if stock prices are martingales. It turns out that testing if the coefficient on the lagged percent change variable is equal to 0 is equivalent to testing if the coefficient of just the simple lagged stock price is equal to 1. I find a p-value on this coefficient of .427, so we fail to reject the null hypothesis that this coefficient is 0, which is equivalent to saying stock prices behave like martingales.

[4] Technically, this also includes Conrail Shared Assets Operations. However, Conrail is jointly operated by Norfolk and Southern and CSX. As such when accidents happen in the Conrail system my data indicates that both CSX and Norfolk and Southern had an accident.

incidents that result in HAZMAT spills, one or more deaths, two or more injuries, or a certain threshold of damages. These reports include data on damages, injuries, deaths, whether or not there was a HAZMAT leak, and probable cause of the accident.[5] I categorize these probable causes into infrastructure, mechanical, negligence, user error, outside, and unknown causes. While this is ultimately a judgement call on my part, the NTSB reports are fairly clear as to the probable cause of the accident.

I merge this accident level data with stock prices of the railroads over the same time period. I have stock price data from six railroads, these were Burlington Northern and Santa Fe, Canadian National, Canadian Pacific, CSX, Norfolk and Southern, and Union Pacific.[6] These are also the six largest railroads operating in the United States.

There is also a potential distinction between two types of accidents I would like to make. Many of these accidents take place in switchyards, in which during operations moving cars around a yard, an employee is struck and killed by a train. Of the 54 accidents in my sample, 19 of them are accidents in which only one person died, there was no HAZMAT leaks and no damages. If we believe there to be effects on the stock price of the railroad due to accidents imposing significant costs on the railroad due to legal costs, replacing capital, cleanup, and indemnity, its unlikely these accidents will register all that high on the radar of the investors. As

[5] The NTSB takes care to note that they do not determine the cause of the accident, only probable cause. As far as I am aware this is because there may be legal proceedings that happen because of an accident, and any cause of the accident is determined during these proceedings instead of in the NTSB investigation.

[6] Stock prices of the Burlington Northern and Santa Fe (BNSF) are not directly available, as it is owned by Berkshire Hathaway. It is estimated that BNSF makes up 40% of Berkshire Hathaway, so I am comfortable including it in the model. I present a model in the appendix that excludes accidents that involve BNSF and get similar results.

such I run two separate regressions, one with and one without these accidents included. Results for the models with these switchyard accidents removed are reported in the appendix.

Below I present summary statistics on the accidents. I present the mean number of injuries, deaths, and damages, as well as the number of accidents that had a HAZMAT spill. I also include a breakdown of accident causes. They fall into six different categories. Infrastructure, which are accidents caused by track defects. Mechanical, which is due to problems with either the locomotive or cars. Negligence and user error, which come down to an individual making a mistake that was the proximate cause of the accident. I assign an accident to be caused due to negligence based on a variety of factors, if there were drugs in the employees system, if the employee deviated significantly from company protocol, or if there was some level of gross disregard for safety. Outside factors are accidents caused by people not working for the railroad. Finally, there are two accidents in the sample that the NTSB has yet to release a final report on that includes causes and they are categorized as unknown.

Table 1: Accident Summary Statistics

Number of Accidents	Average Injuries	Average Deaths	Average Costs	Accidents with Spills
54	2.83	0.78	$3,455,055	18

Table 2: Accident Causes

Infrastructure	Mechanical	Negligence	User Error	Outside Factors	Unknown
14	3	14	20	1	2

Finally, I present the common trends assumption. I plot the means of the treated of the untreated groups day to day percent change three days prior to the accident. We expect the treated and untreated values to move in parallel prior to treatment at day zero.

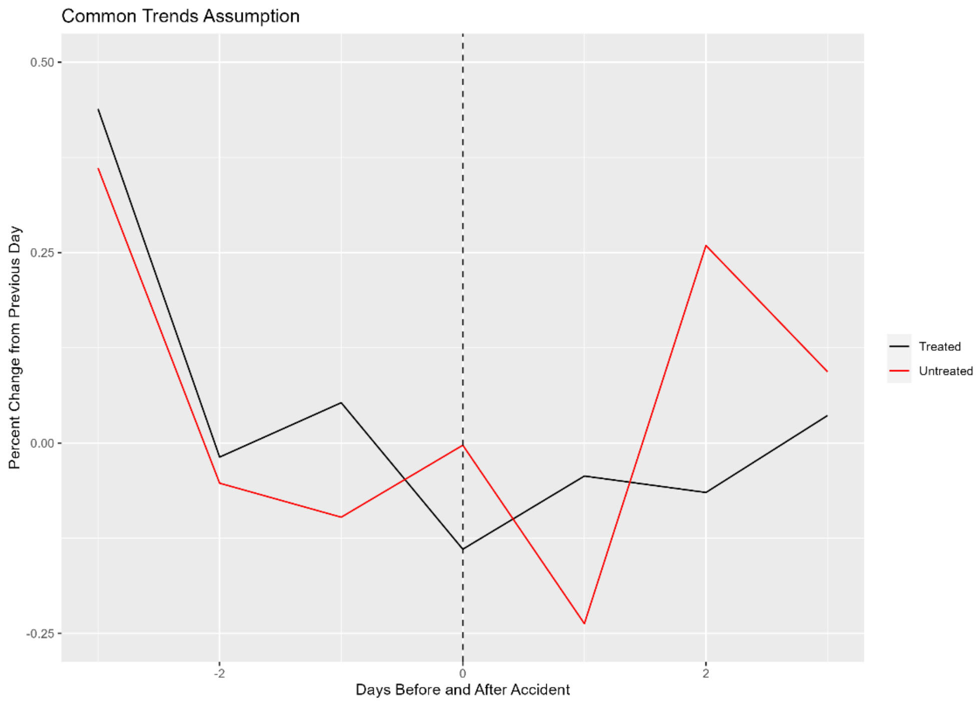

I get good performance on the common trends assumption as the stock prices don't vary by more than 0.5% on any day prior to the accident and it's difficult to tell these two apart. After all, accidents are random events that are impossible to predict down to the day, so there is no reason for railroads that have an accident to move differently prior to the accident.

RESULTS

After running the regression, I find that very few of the treatment variables are significant. Coefficients that theory predicts would be negative are not significant, and in many cases positive. For example, our estimates find that accidents that had a HAZMAT leak the railroad that had the accident saw increased stock prices relative to an accident that did not have a HAZMAT leak. This obviously makes no sense, if we think accidents have an effect on stock prices because it changes future profitability, then a HAZMAT leak necessarily affects future profitability in a negative way because of cleanup costs and potential investors seeing this as an increased risk which could change stock prices.

Table 3: Model Coefficents

	Dependent variable:	
	Untreated	Treated
Fault	0.022 (0.182)	−0.277 (0.339)
RT1	0.051 (0.155)	0.184 (0.314)
RT2	−0.177 (0.155)	0.479 (0.313)
RT3	0.316** (0.154)	-
Deaths	0.054 (0.083)	0.014 (0.191)
Injuries	−0.007 (0.005)	0.001 (0.011)
HAZMAT	0.054 (0.215)	0.281 (0.473)
log(Damages)	−0.012 (0.010)	−0.002 (0.023)
Infrastructure	0.073 (0.185)	−0.137 (0.433)
Mechanical	−1.257*** (0.337)	0.055 (0.779)
Negligence	−0.066 (0.166)	−0.202 (0.379)
Outside Factors	−0.431 (0.408)	−0.126 (0.969)
Unknown Cause	−0.101 (0.297)	−0.423 (0.699)

Although the model does have some level of joint significance across all included regressors, the lack of significance on the treatment variables indicates that accidents have little bearing on the stock prices of the railroad. Accidents that were caused by mechanical problems on the train are associated with lower stock prices across the industry, but this may just be spurious correlation. Given the total lack of significance on accident level coefficients, I argue that there is little reason to believe that accidents on the railroad affect their stock price. Below I illustrate the average response based on the main model.

Table 4: Predicted Response for Mean Accident

	0	1	2	3
Untreated	−0.135	−0.083	−0.312	0.182
Treated	−0.434	−0.198	−0.133	−0.118

There are many potential reasons for these null results. Again, railroads are insured, and for the accidents they aren't insured for, they also have large cash stores to deal with these accidents. As such, any costs to the railroad will be small and won't affect the bottom line of the railroad all that much. There could also be no alternative for these railroads, so they may face costs in the present, but there won't be any shift in the usage of the railroad. As this is just an upfront cost there won't be any shift in the stock price because there's no shift in utilization and therefore no shift in profitability and so stock prices are the same. Again, the railroad industry is incredibly concentrated, the class 1s operate the majority of freight miles in the US and even the smallest of them have thousands of locomotives, and tens of thousands of miles of track and employees, at those levels, the direct cost to a company for any given accident is likely to be small. However, given there is a fair amount of evidence that airlines see stock prices shift after an accident, what accounts for the difference here?

First, aviation accidents are worse, many of the rail accidents involve relatively low numbers of injuries or deaths, as well as the cost to replace damaged equipment and clean up. This leads to the second point, knowledge of aviation accidents is more common than knowledge of rail accidents. If both investors and customers do not learn about these accidents when they happen, prices will remain unchanged. Additionally, assuming that rail and aviation accidents were both widely known, the customers of railroads tend to be large companies moving freight, while the customers of airlines are an average person off the street. These large companies are going to be more knowledgeable about the risks involved in railroads relative to an average persons' knowledge of the risk of flying, so the company doesn't shift its usage while an average person might do so. These different customer responses then drive market reactions accordingly. Finally, airlines are a far less concentrated market than the railroads, so more alternatives exist, and accidents impose more costs on the airline than it would on a railroad.

Whether or not these markets are behaving efficiently unfortunately cannot be answered directly by the data. There are two cases here, one in which the knowledge of the accident is widespread among investors, in which case no movement indicates market efficiency. The accident is taken into account when investors make a decision, it just doesn't register relative to other reasons to buy or sell the stock. The other case happens when knowledge isn't widespread, and so no movement happens, indicating inefficiency as knowledge of certain events isn't priced into the stock.

CONCLUSION

Care should be taken before interpreting these results as saying railroad accidents never affect the stock price of the railroad. While it may be true that small accidents don't move the needle for investors, large accidents can definitely impose costs on a railroad beyond their ability to pay. The Maine, Montreal, and Atlantic railroad went bankrupt as a direct consequence of the damage the Lac-Megantic rail disaster caused. Norfolk and Southern has already spent hundreds of millions of dollars on cleanup due to the East Palestine derailment, to say nothing of on going legal cases and damage to their reputation. These particularly nasty accidents are extremely rare, which makes estimating their true effect on the railroad difficult. More broadly, the sample size of this analysis is lower than ideal. There are only fifty-four accidents that the NTSB investigated from January 2010 to December 2021. This means any conclusions are going to have low statistical power.

There are a few potential solutions to this problem. One, instead of sampling accidents from the NTSB, sample from the FRA. The FRA has a list of railroad accidents dating back many years and doesn't choose to investigate or not, so there are far more observations in the sample each year. This comes at the cost of including many more "small" accidents, in which the potential costs to the company are negligible. Another potential solution is incorporating data from other countries, in particular, including accidents that take place in Canada would be a natural extension given both Canadian National and Canadian Pacific operate in both countries. Finally, if data prior to 2010 became available, this would greatly expand the sample. Even if reports only go back to the foundation of Amtrak, that is still forty years of additional data. Such data may be available via FOIA requests, which would be an interesting way forward.

Further research into this field should take a combination of these approaches. Canadian records are relatively easy to come by, but many don't report exact levels of damages. FOIA requests would take some time, but I doubt these requests would be rejected. FRA records would be difficult to combine, especially their older records, but should yield more observations as well, theoretically going back as far as 1975.

Finally, including data about how widespread information about an accident is could lead to interesting results. Such a methodology would take say, the number of articles about a given railroad, and a spike around an accident would be indicative of greater knowledge of the accident. Because some accidents have more information more widely available than others, this could be a significant source of heterogeneity in the response to an accident, although how correlated this amount of information is with the severity of the accident is unknown, it is an interesting way forward. This would also help answer the question of market efficiency, as widespread information leans towards market efficiency while little information lends itself to inefficiency as information takes a while to move through the market. This is still limited however, as investors can see earning reports, so shocks to the stock price happen when stuff actually changes.

APPENDIX: ROBUSTNESS CHECKS

Burlington Northern Santa Fe (BNSF) stock prices are not directly available. In 2009 Berkshire Hathaway acquired the railroad for $44 billion. Since then, most estimates place BNSF as making up roughly 40% of Berkshire Hathaway holdings. I believe this is enough to still see any effect if it exists, but regardless, I estimate a model without BNSF accidents and BNSF stock prices. This model is identical to the model presented in the main text. This reduced model is estimated from 41 different accidents.

Table 5: Model With BNSF Accidents Removed

	Dependent variable:	
	Change	
FaultInd	0.018 (0.201)	−0.119 (0.402)
RT1	−0.118 (0.205)	−0.061 (0.363)
RT2	−0.389* (0.205)	0.307 (0.362)
RT3	−0.141 (0.204)	-
Deaths	0.234** (0.116)	−0.023 (0.244)
Injuries	−0.006 (0.006)	0.0005 (0.012)
Hazmat1	−0.048 (0.271)	0.396 (0.539)
log(Damages)	0.016 (0.014)	−0.008 (0.028)
Infrastructure	−0.094 (0.243)	0.044 (0.511)
Mechanical	−1.453*** (0.556)	−0.297 (1.106)
Negligence	−0.741*** (0.218)	0.123 (0.460)
Outside Factors	−0.400 (0.468)	−0.043 (0.992)
Unknown Cause	0.101 (0.341)	−0.441 (0.718)

Much like the main model, most accident characteristics are insignificant at the $\alpha = 0.05$ level. Number of deaths is significant, but positive, which could be due to investors shifting from the railroad that had the accident towards railroads that did not have an accident. However, deaths on the railroad with the accident is not significant. Accidents caused by mechanical problems on the locomotive result in roughly a 1.5% decrease in the stock price across the industry, and accidents caused by employee negligence result in a 0.74% decrease in prices across the industry. Again, the railroad at fault for the accident still sees no significant movement for these accidents beyond the level change.

Additional robustness checks are in the form of removing switching accidents. Switching accidents are accidents that take place during yard operations. Yard operations usually consist of moving freight cars between locomotives to either send further on the rails to another destination or to load and unload freight cars at their terminal destinations. Yard operations take place at low speeds, usually no faster than five miles per hour. Given that yard operations involve large amounts of coupling and uncoupling cars, incidents in which railroad employees are crushed or otherwise run over by a train aren't unheard of. Of the 54 accidents in the main sample, 21 of them are switching accidents of some kind. Given the nature of these accidents its unlikely that they will be reported to investors. Not to be glib, but compared to other railroad accidents, these are quite tame. They involve a singular death or injury, and nothing else. Events like these are not reported on in the same way that derailments on the main line are reported on, nor do they interfere with operations in the same way. As such, I believe the effects of such accidents to be quite small, if they exist at all. As such, I estimate a model without these switching accidents included using the same equation as the model presented in the main text.

Table 6: Model With Switching Accidents Removed

	Dependent variable:	
	Untreated	Treated
FaultInd	0.083 (0.250)	−0.884 (1.970)
RT1	−0.192 (0.908)	0.058 (0.395)
RT2	−0.694 (0.908)	0.424 (0.395)
RT3	−0.086 (0.908)	-
Deaths	0.075 (0.098)	0.087 (0.218)
Injuries	−0.006 (0.005)	−0.001 (0.012)
HAZMAT	0.093 (0.228)	0.457 (0.487)
log(Damages)	0.008 (0.059)	0.050 (0.128)
Infrastructure	0.009 (0.260)	−0.388 (0.577)
Mechanical	-1.212^{***} (0.374)	−0.212 (0.840)
Negligence	−0.007 (0.243)	−0.510 (0.541)
Unknown	−0.566 (0.454)	0.354 (1.058)

Again, the only significant coefficient is on accidents caused by mechanical problems. All other coefficients remain insignificant, and many have the wrong sign. If we expect an effect of the accident on the stock price to exist, it should be in accidents that have high costs to the railroad. Given this reduced sample of accidents that is made up of mostly collisions with other trains and derailments, these are the accidents we expect to see an effect in. Despite this, no such effect exists, so there is more evidence that railroad accidents do not affect the stock price of the railroad.

Finally, I implement robustness checks with a 1-, 5-, and 10-day event horizons, results are too long to include here and are presented in the attached R file. Regardless, the 1-, 5-, and 10-day event horizon models have similar results to the 3-day model. All have significant and negative estimates for the mechanical causes term, with much of the other terms being insignificant. The 5- and 10-day models have positive estimates for accident caused by unknown factors while the 1-day model this term is negative, and all three are significant. Regardless, I think most of this is due to random noise in the data, and not indicative of any real movement in the stock price of the railroad.

Across all of these models, an average accident doesn't move the stock price of the railroad by more than 1%. Given this small movement relative to the existing variation in the stock price, these variations are just not economically significant, even if there is some level of statistical significance attached. Of course, particularly bad accidents may have an impact, but there are far too few of these to draw any solid conclusions.

REFERENCES

Abadie, Alberto. "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects." *Journal of Economic Literature* 59, no. 2 (June 2021): 391–425. https://doi.org/10.1257/jel.20191450.

Akyildirim, Erdinc, Shaen Corbet, Marina Efthymiou, Cathal Guiomard, John F. O'Connell, and Ahmet Sensoy. "The Financial Market Effects of International Aviation Disasters." *International Review of Financial Analysis* 69 (May 1, 2020): 101468. https://doi.org/10.1016/j.irfa.2020.101468.

Carter, David A., and Betty J. Simkins. "The Market's Reaction to Unexpected, Catastrophic Events: The Case of Airline Stock Returns and the September 11th Attacks." *The Quarterly Review of Economics and Finance* 44, no. 4 (September 1, 2004): 539–58. https://doi.org/10.1016/j.qref.2003.10.001.

Ho, Jerry C., Mei Qiu, and Xiaojun Tang. "Do Airlines Always Suffer from Crashes?" *Economics Letters* 118, no. 1 (January 1, 2013): 113–17. https://doi.org/10.1016/j.econlet.2012.09.031.

"The Market's Reaction to Unexpected, Catastrophic Events: The Case of Airline Stock Returns and the September 11th Attacks.," n.d.

Walker, Thomas, Kuntara Pukthuanthong, and Sergey Barabanov. "On the Stock Market's Reaction to Major Railroad Accidents." *Journal of the Transportation Research Forum* 45, no. 1 (October 14, 2006). https://doi.org/https://doi.org/10.5399/osu/jtrf.45.1.871.

Walker, Thomas, Dolruedee Thiengtham, Onem Ozocak, and Sergey Barabonov. "The Role of Indemnification Agreements and Legal Liability in Railroad Disasters : A Financial Market Perspective." *International Journal of Managerial Finance* 3, no. 4 (October 2, 2007).