A Report on "Subways, Strikes, and Slowdowns: The Impacts of Public Transit on Traffic Congestion" by Anderson (2014)

Reviewer 2

February 25, 2026

v1

Disclaimer

This report was generated by large language models, overseen by a human editor. It represents the honest opinion of The Catalogue of Errors Ltd, but its accuracy should be verified by a qualified expert. Comments can be made here. Any errors in the report will be corrected in future revisions.

> I am wiser than this person; for it is likely that neither of us knows anything fine and good, but he thinks he knows something when he does not know it, whereas I, just as I do not know, do not think I know, either. I seem, then, to be wiser than him in this small way, at least: that what I do not know, I do not think I know, either.
>
> Plato, *The Apology of Socrates*, 21d

> To err is human. All human knowledge is fallible and therefore uncertain. It follows that we must distinguish sharply between truth and certainty. That to err is human means not only that we must constantly struggle against error, but also that, even when we have taken the greatest care, we cannot be completely certain that we have not made a mistake.
>
> Karl Popper, 'Knowledge and the Shaping of Reality'

Overview

Abstract Summary: This paper investigates the impact of public transit on traffic congestion, predicting that transit riders are likely to be individuals commuting along routes with severe roadway delays, thus having high marginal impacts on congestion. The study tests this prediction using data from a 2003 Los Angeles transit workers' strike, finding that average highway delay increases significantly when transit service ceases.

Key Methodology: The study uses a simple choice model calibrated with Los Angeles metro area data and tests predictions using freeway speed data from a 2003 Los Angeles transit strike, employing a regression discontinuity (RD) design.

Research Question: What are the impacts of public transit on traffic congestion, and do these impacts vary based on the heterogeneity of congestion levels faced by commuters?

Summary

Is It Credible?

This article argues that public transit provides much larger congestion relief benefits than previously believed. Using a regression discontinuity design based on a 35-day transit strike in Los Angeles in 2003, the author estimates that average highway delay increases by 47 percent when transit service ceases (p. 2779). To explain this unexpectedly large effect, the article introduces a choice model demonstrating that transit riders self-select from the most heavily congested routes, meaning their removal from the road network has a disproportionately high marginal impact on traffic speeds. Building on these short-run empirical findings and theoretical insights, the article ultimately claims that the long-term congestion relief benefits of rail transit are large enough to justify the massive capital costs of building such systems (p. 2791).

The core empirical finding—that a temporary halt in transit service causes severe short-run congestion—is highly credible. The regression discontinuity design is well-executed, and the use of a sudden, exogenous strike provides a clean natural experiment. The author also conducts convincing falsification tests, showing no similar traffic spikes on neighboring county freeways or during the same calendar weeks in the following year (pp. 2784–2785). However, the magnitude of the aggregate monetary benefits derived from this finding is less certain. The empirical data comes exclusively from freeway loop detectors, which cover only about half of the vehicle miles traveled in the region. To calculate the upper-bound congestion relief benefit of $4.1 billion, the article assumes that ceasing transit service increases delays on arterial roads by the exact same amount that it increases delays on freeways (p. 2787). Because there is no direct data on arterial traffic, this extrapolation introduces a substantial degree of uncertainty into the largest benefit estimates,

though the author justifies this by noting that arterial roads lack the ramp metering that protects freeways from saturation.

The theoretical model provides an elegant and logically sound explanation for the empirical results. By incorporating heterogeneity in driving delays, the model successfully bridges the gap between low overall transit ridership and high congestion relief. However, the model's explanatory power relies on a calibrated counterfactual. It predicts that if rail passengers chose to drive, they would face delays of 3.19 minutes per mile, which is more than six times the average delay faced by existing drivers (p. 2771). This crucial figure is not directly observed but is instead generated by calibrating the model to match observed transit market shares. Furthermore, while the model's quantitative predictions align with prior literature, the framework does not fully integrate the complex network spillovers and driver rerouting that clearly occur in the empirical data.

The article's most ambitious claim is that rail transit generates positive net benefits over the long term, challenging a prevailing consensus in transportation economics. This conclusion requires extrapolating the effects of a 35-day strike to evaluate a permanent infrastructure system. As the article acknowledges, a long-term shutdown of public transit would likely induce behavioral adaptations that are not feasible in the short run, such as workers moving closer to their jobs, firms relocating, or commuters purchasing vehicles (p. 2788). While the author uses long-run elasticities from studies of gasoline prices and congestion charges to simulate these adaptations and bound the uncertainty, estimating the true long-run equilibrium from a short-run shock remains inherently speculative.

Despite these limitations in extrapolating to long-run capital benefits, the article makes a highly valuable contribution to public policy analysis. It convincingly demonstrates that when policymakers implement second-best policies to mitigate externalities, using simple population averages will severely underestimate a program's effectiveness if the mitigating activity is disproportionately adopted by

those who generate the greatest externalities.

The Bottom Line

Anderson provides highly credible empirical evidence that public transit significantly reduces short-run traffic congestion, leveraging a clean natural experiment from a 2003 Los Angeles transit strike. The article's theoretical explanation—that transit disproportionately attracts commuters from the most congested routes—is logically sound and elegantly explains why a system with low overall ridership can provide massive congestion relief. However, the article's broader claims about the long-term, positive net benefits of rail infrastructure are more speculative, as they rely on extrapolating the effects of a temporary strike to permanent behavioral and structural adaptations, and on assuming that unmeasured arterial roads experience severe delays similar to monitored freeways.

Potential Issues

Extrapolation from a temporary event to permanent system value: The article's central conclusions about the long-term, multi-decade value of the Los Angeles rail system are derived from the empirical effects of a temporary, 35-day transit strike. This represents a potential mismatch between the nature of the evidence and the scope of the conclusion. The article acknowledges this limitation, stating that "The impact on congestion of a long-term shutdown of public transit is likely different than the short-run effect of temporarily shutting down public transit" (p. 2788). To bridge this gap, the analysis relies on a simulation model that uses long-run vehicle-miles-traveled (VMT) elasticities from studies of gasoline prices and the London congestion charge (pp. 2788–2789). However, it is uncertain whether these elasticities accurately represent the complex, long-run adaptations to a permanent transit shutdown, which could include residential and firm relocation, land-use changes, and political pressure for road network expansion. While the author's simulation is a rigorous attempt to bound the uncertainty, the conclusion that "the long-run congestion relief benefits of transit service are at least half the size of the short-run benefits" (p. 2790) is based on this speculative modeling rather than direct empirical evidence from the strike itself.

Generalization of congestion effects from freeways to all roads: The article's empirical analysis is based exclusively on data from freeway loop detectors, which account for 53 percent of vehicle miles traveled in the study area (p. 2774). The article is transparent about this limitation, noting that "we only have data on freeway delays and do not observe delays on arterial roads" (p. 2786). However, the article's upper-bound benefit estimate of $4.1 billion relies on the critical assumption that "ceasing transit service increases delays on arterial roads by the same amount that it increases delays on freeways" (p. 2787). The author defends this assumption by noting that arterial roads lack ramp meters, which restrict flow onto freeways, potentially making

arterial congestion even worse than freeway congestion. While this logic is sound, the absence of any data on the 47 percent of traffic on arterial roads means the article's largest benefit estimates rely heavily on this theoretical justification rather than empirical observation.

Dependence of the theoretical model on a calibrated counterfactual: The core explanatory mechanism of the theoretical model relies on a calibrated, but unobserved, counterfactual. The model predicts that the average delay for a rail passenger, *if they chose to drive*, would be 3.19 minutes per mile—more than six times the delay faced by existing drivers (p. 2771). This crucial figure is not an assumption, but rather an output of a calibration exercise where the distribution of transit access times is varied until the model's predicted transit share matches the observed share (p. 2767). While this is a standard method for arguing for a model's structural validity, the argument explains an observed fact (transit ridership) with an unobserved "fact" (the high counterfactual delay) generated by its own calibrated structure. This makes the explanation dependent on the model's ability to correctly infer unobserved preferences and constraints from aggregate data.

Incomplete integration of network spillover effects: The article's theoretical model focuses on a simple selection mechanism where former transit riders cause congestion on the specific freeways that parallel their old transit routes. However, the empirical results show a substantial and statistically significant 29 percent increase in delay on freeways that do not parallel a major transit line (p. 2779). This finding indicates the presence of wider network spillovers, likely from drivers rerouting across the entire system to avoid the most congested corridors. The article transparently reports this finding, and its empirical regression discontinuity design correctly captures these system-wide effects. However, the simple choice model presented does not account for such rerouting, making it a useful illustration of the selection mechanism but an incomplete explanation of the full, complex dynamics observed in the empirical data (p. 2766).

Potential for sample selection bias from detector dropout: The analysis relies on data from active freeway loop detectors, and the number of active detectors in the sample was lower during the strike period (an average of 640) than during the non-strike period (an average of 720) (p. 2776). The article includes detector fixed effects "to ensure that changes in the composition of detectors in service do not bias our estimates" (p. 2776). While this is the standard econometric approach, a potential for selection bias remains if the reason for detectors dropping out of the sample is correlated with the treatment itself. For example, if detectors in the most heavily congested areas are more prone to failure during periods of extreme traffic, and the strike caused such extreme traffic, the sample of active detectors might be systematically less congested than the full set. The author notes that observations are dropped when detectors are out of service (p. 2776), but if the failure mechanism is correlated with the strike intensity, the fixed effects might not fully correct for the composition change.

Treatment of welfare losses from displaced trips: The empirical results show that hourly traffic flow per lane decreased by an average of 2.2 percent during peak hours as a result of the strike, indicating that extreme congestion reduced road throughput (p. 2783). The article's primary short-run benefit calculation focuses on the cost of increased travel time for the drivers who remained on the road. While the author acknowledges that "time-shifted departures and queue spillovers represent an additional potential welfare cost from the strike" (p. 2783), these costs are not integrated into the main short-run benefit figure. However, the article does address this in the long-run analysis, estimating the welfare loss from individuals who stop traveling to be "second order," ranging from $13 million to $120 million per year (p. 2790). While this justifies their exclusion as economically insignificant relative to the billions in congestion costs, the primary short-run headline number technically presents an incomplete picture of the total immediate welfare impact.

Sensitivity of theoretical predictions to the congestion function: The theoretical

model's quantitative predictions of congestion effects rely on a specific functional form relating traffic volume to delay: delay = a • traffic volume^3.7, which is taken from Parry and Small (2009) (p. 2770). The exponent of 3.7 is a critical parameter that determines the non-linear nature of congestion. The article notes that its predicted effects are similar to those in Parry and Small (2009) because it uses similar parameters, and it incorporates short-run VMT elasticities from Small and van Dender (2007) to solve for equilibrium (p. 2788). While the author tests the sensitivity of the model to the delay multiplier in the appendix (Table A1), the absolute magnitude of the model's theoretical predictions still rests heavily on this specific functional form from the literature.

Presentation and transparency issues: The article contains several minor presentation issues. First, the back-of-the-envelope calculation of agglomeration benefits relies on a "12 percent speed increase" (p. 2787), but the derivation of this figure from the article's delay estimates is not explicitly shown. Second, the calculation of "Present value of net benefits" in Table 11 (p. 2791) contains minor discrepancies; for example, the reported net benefit is $0.1 billion higher than the difference between the reported gross benefits and costs. Finally, the supplementary appendix contains a formatting error from an uncorrected cross-reference, reading "Section Error! Reference source not found." (supplementary p. 1). These minor issues do not affect the article's substantive conclusions.

Future Research

Measurement of arterial congestion: Future work could utilize modern, ubiquitous GPS and smartphone location data to directly measure the impact of transit disruptions on arterial roads and local streets. This would eliminate the need to extrapolate freeway loop detector data and provide a complete, empirical picture of system-wide congestion relief, allowing for more precise aggregate benefit calculations.

Long-run behavioral adaptations: Researchers could exploit permanent changes to transit networks, such as the opening of new rail lines or permanent system closures, to estimate long-run elasticities of travel demand and residential sorting. Comparing these long-run structural effects to short-run disruptions like strikes would clarify exactly how much behavioral adaptation attenuates the initial congestion relief benefits of transit over time.

Modeling network spillovers: Future theoretical models could integrate spatial network dynamics to account for driver rerouting during transit disruptions. By combining commuter heterogeneity with a structural model of route choice, researchers could better isolate the direct effect of displaced transit riders from the indirect, system-wide effects of drivers shifting their commutes to avoid newly congested corridors.