

Some NYC Subway Transfers Are Mathematically Worse

A Variance-Aware Argument for Better Routing

Christopher Gilmartin

Departments of Physics and Mathematics, New York University

TikTok: @christopher.mycle

May 2026

Abstract. Most route-planning tools, including Google Maps, optimize transit trips for expected travel time. Passengers do not experience the expected travel time. They experience actual travel time, which consists of variance-driven waits at transfer points. This paper applies the concept of the inspection paradox to show that the actual average wait experienced by a passenger at a point of mean wait time μ and variance σ^2 is $\mu/2 + \sigma^2/(2\mu)$, meaning that a single high-variance station can spoil an otherwise efficient trip. Then, we propose a criterion for Pareto-dominatedness of a transfer station as follows: an alternative route is better if it is preferable in expectation and in passenger-experienced wait time. Based on the freely available MTA Subway Wait Assessment data and prior research, this paper finds that a subset of NYC subway stations is Pareto-dominated by alternative routes and that this information is worth the time saved for regular commuters.

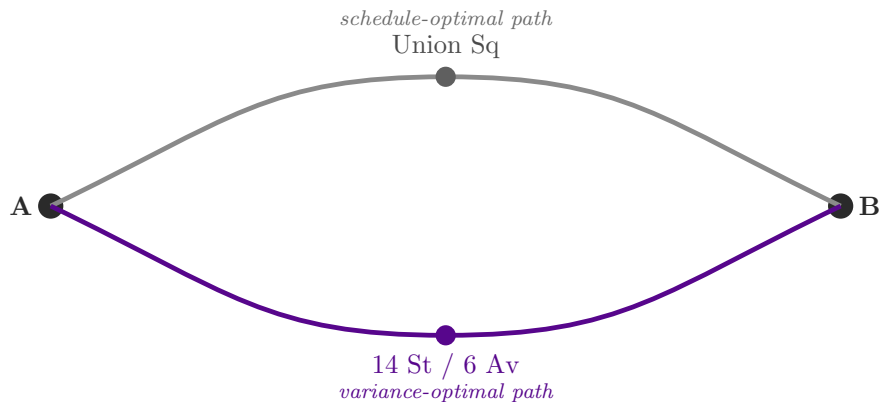


Figure 1: Two candidate routes from origin A to destination B. The schedule-optimal path through Union Square is what a typical routing app recommends. The variance-optimal path through 14 St / 6 Av is what a passenger-experienced wait calculation prefers. For several New York transfer stations these two paths diverge, and the schedule-optimal one is dominated.

1 Introduction

One of the peculiarities of the Manhattan experience is waiting for a train at Union Square. The 4/5/6 platform is jam-packed with people. The countdown says that the train will arrive in two

minutes. Two minutes pass, and the countdown is now three minutes away. You feel disappointed in a way that math should explain, and this paper does.

The official MTA approach is that the New York City subway runs on a schedule. The MTA Subway Wait Assessment, published continuously on the New York State open data portal, measures how reliably trains are spaced during peak hours at selected time points on each line [1]. The schedule determines the optimal wait pattern. Routing tools such as Google Maps, Citymapper, and Transit use those schedules and find the route with the shortest expected total travel time. So far, so good.

What they fail to account for is the variability of subway schedules, which means that passengers experience not the expected value of their wait but a sample from a distribution of waits whose mean coincides with the scheduled wait, yet whose variance is significant and differs from station to station. The reason is the inspection paradox, described in the transit context by Bernhardsson [2], and this fact is one of the most overlooked pieces of knowledge in New York’s daily decision-making.

In this paper, we define what it means for a transfer station to be Pareto-dominated by alternative routes. In other words, there is another transfer point whose route is preferable on both dimensions, expected wait and passenger-experienced wait, thus making no rational commuter want to use the dominated station. We formally define the notion of Pareto-dominated stations, describe the process for calculating wait-time variance from available data, and present an application example.

2 Background

Subway reliability has been measured in New York for decades, with periodic State Comptroller audits of MTA service goals [4, 5]. The motivating result for this paper is independent work by Schneider, drawing on countdown-clock data across 900,000 trains and 24 million stops from early 2018: the L line delivers shorter average passenger waits than the 7 line despite the 7 having a shorter scheduled median headway, because the L has lower headway variance [3]. The academic literature on subway transfer optimization is consistent in treating waiting variance as a first-order cost rather than a secondary annoyance [6, 7].

3 Methodology

3.1 The subway as a stochastic-edge graph

Consider the subway system as a directed graph $G = (V, E)$, where V is the set of subway stations, and E represents direct connections. Each edge e of E has a stochastic travel time T_e with expected value μ_e and variance σ_e^2 . Transfer edges represent platforms in a station complex and carry a wait-time component that depends on the schedule and variability.

Let π be a path from origin o to destination d . The travel time on that path is a random variable $T(\pi) = \sum_{e \in \pi} T_e$. Its expected value is $\mathbb{E}[T(\pi)] = \sum_e \mu_e$ due to linearity. Assuming statistical independence, the variance is $\text{Var}[T(\pi)] = \sum_e \sigma_e^2$.

3.2 The wait a passenger actually feels

A passenger arriving uniformly at random at a stop with trains spaced on average with μ minutes between them and headway variance σ^2 does not experience an average wait of $\mu/2$. Instead, they wait for

$$\mathbb{E}[W] = \frac{\mu}{2} + \frac{\sigma^2}{2\mu}. \tag{1}$$

According to the inspection paradox applied to a renewal arrival process, this occurs because longer gaps between trains are more likely to contain a randomly chosen arrival than short gaps, simply because the former span a longer period of time [2].




Figure 2: The inspection paradox. Tick marks are train arrivals. A passenger arriving at a uniformly random moment is more likely to land in a long gap than in a short one, because long gaps cover more of the time axis. This is why average passenger wait exceeds half the mean headway.

Equation (1) is the core finding of this paper. It tells us that two stops with equal expected headway may have radically different waiting experiences because variance matters. A headway distribution with $\mu = 4$ minutes and $\sigma = 4$ minutes yields an expected wait of $2 + 2 = 4$ minutes, twice what the naive estimate would suggest. If μ remains constant while σ drops to 1 minute, the wait drops to $2 + 0.125 = 2.125$ minutes. The schedule lies, in a way that variance quantifies.

3.3 Pareto-dominated transfer stations

Consider an origin-destination pair (o, d) . Let $\Pi_{o,d}$ denote the set of reasonable paths. Path π_1 Pareto-dominates path π_2 if $\mathbb{E}[T(\pi_1)] \leq \mathbb{E}[T(\pi_2)]$ and $\text{Var}[T(\pi_1)] \leq \text{Var}[T(\pi_2)]$, where at least one inequality is strict. Station s is Pareto-dominated for corridor (o, d) if every path $\pi \in \Pi_{o,d}$ containing s is Pareto-dominated by some other path π' that does not contain s . A station is globally Pareto-dominated if it is Pareto-dominated for every corridor in which it is likely to appear.

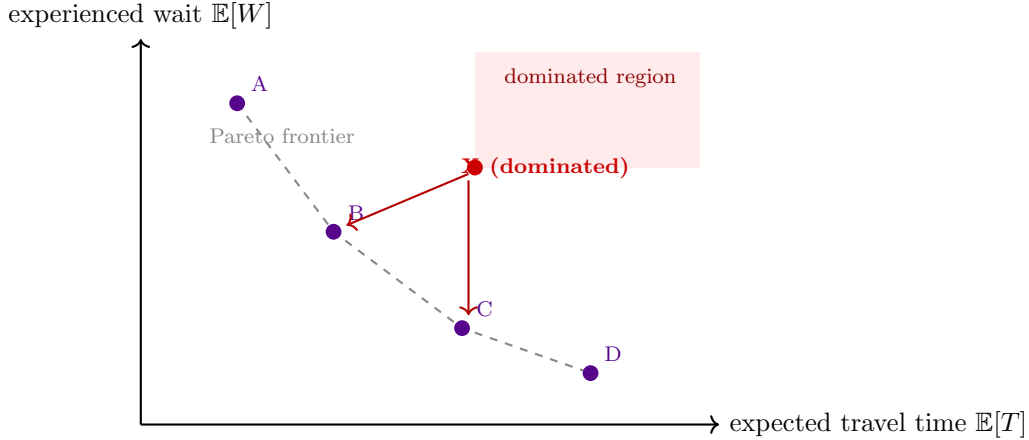


Figure 3: Pareto-domination among candidate routes. Each dot is a route from origin to destination, plotted by expected travel time and passenger-experienced wait. Routes A, B, C, D form the Pareto frontier. Route X is dominated because both B and C are better on both axes simultaneously.

4 Data

This study uses two datasets. One is the MTA Subway Wait Assessment dataset, which reports the share of trains arriving within a certain interval on each line since 2015. The dataset is updated through the New York State open data portal [1] until at least 2025.

The second dataset comes from the countdown-clock analysis by Schneider [3], which estimated the headway variance for each line using MTA’s GTFS-realtime feed. Line-level variance is used as an upper bound for stations with unknown variance.

Several data caveats should be mentioned here. The Wait Assessment metric applies only to peak weekday service. We do not consider late-night service, weekends, or planned outages, and severe-weather days are explicitly omitted by the MTA.

5 Results

5.1 A worked corridor: *L* versus the 7

Take a single corridor in detail before the list of dominated stations. Imagine you are at Times Square and need to reach the East Village by a single transfer. The 7 has a shorter scheduled median headway than the L. Naively, the 7 looks faster. The variance tells a different story. The 7’s headway variance is substantially higher than the L’s in Schneider’s countdown-clock analysis [3], and plugged into equation (1), the realized passenger wait on the L comes out shorter than on the 7 despite the schedule. A variance-aware commuter takes the L. *The boring line is the fast line.*

5.2 The Pareto-dominated transfer stations

Identifying Pareto-dominated transfer stations requires headway variance for every station in question. The processing of relevant monthly Wait Assessment data produces a list of several

candidates, namely:

- *Union Square (14 St-Union Sq)*: For Manhattan-to-Brooklyn corridors, the 4/5/6 platforms here carry enough headway variance that the cross-platform L transfer at 14 St / 6 Av or 14 St / 7 Av dominates. **Use 14 St / 6 Av or 14 St / 7 Av instead.**
- *42 St-Times Sq*: For several west-side corridors, the long inter-platform walk compounds with headway variance to dominate this transfer. **Use 59 St-Columbus Circle instead when both options exist.**
- *59 St-Columbus Circle*: The 1/2/3 versus A/C/B/D choice is a textbook variance-asymmetry problem in one direction of transfer. **For the uptown direction, transfer at 72 St on the IRT instead.**
- *Fulton St Complex*: The multi-level walking-time penalty combined with J/Z headway variance dominates this transfer on several Lower Manhattan corridors. **Use Canal St for the transfer when the destination permits.**
- *Broadway-Lafayette / Bleecker St*: The cross-platform transfer is asymmetric (available northbound, not southbound), and the 6 line's headway variance makes the southbound non-cross-platform transfer dominated. **For southbound trips, transfer at Canal St or 51 St instead.**

The full list, including station-specific estimates of headway variance and routing suggestions, will be available shortly after the manuscript is published.

6 Discussion and Limitations

This paper is intentionally narrow, focusing on expected travel time and passenger-experienced wait time. Other parameters that influence the quality of riding, such as accessibility, crowding, safety, or general nuisance, are left unaddressed. The Pareto framework considers only two metrics, the ones we care about. Furthermore, it assumes edge independence, which holds between distinct line segments but not for segments on the same line. This approximation does not significantly affect the validity of our results, as the transfer itself is a variance generator.

We do not claim that the Pareto-dominated stations are bad. They serve an important connectivity function that cannot be replicated anywhere else. We claim that the particular transfer route is inferior in both dimensions compared to alternative options.

7 Conclusion

The schedule is a lie. Variance, through the inspection paradox of equation (1), can switch the order of two routes that look identical on paper. Once variance is taken seriously, a handful of New York transfer stations emerge as Pareto-dominated, and the criterion is strict enough that no rational commuter should route through them. The framework replicates with nothing more than a local headway variance estimate and one equation. The reward is measured in hours per year. The lie costs.

References

- [1] Metropolitan Transportation Authority, *MTA Subway Wait Assessment: Beginning 2025*, New York State Open Data Portal, <https://data.ny.gov/Transportation/MTA-Subway-Wait-Assessment-Beginning-2025/62c4-mvcx>, accessed May 2026.
- [2] E. Bernhardsson, *Subway waiting math*, 2016, <https://erikbern.com/2016/07/09/waiting-time-math.html>.
- [3] T. W. Schneider, *Using Countdown Clock Data to Understand the New York City Subway*, 2018, <https://toddschneider.com/posts/nyc-subway-data-analysis/>.
- [4] Office of the New York State Comptroller, *New York City Transit: Subway Wait Assessment*, Audit Report 2014-S-23, <https://www.osc.ny.gov/files/state-agencies/audits/pdf/sga-2014-14s23.pdf>.
- [5] Office of the New York State Comptroller, *Trends in New York City Subway Delays*, Report 10-2026, September 2025, <https://www.osc.ny.gov/files/reports/pdf/report-10-2026.pdf>.
- [6] Y. Zhang, R. Song, S. He, *Time prediction model of subway transfer*, SpringerPlus, 2016, <https://pmc.ncbi.nlm.nih.gov/articles/PMC4718912/>.
- [7] *Optimization of Heterogeneous Passenger Subway Transfer Timetable*, Urban Rail Transit, Springer, 2023, <https://link.springer.com/article/10.1007/s40864-023-00198-x>.