Supplemental Information: Whose Time Do Traffic Signals Save? Pedestrian–Vehicle Delay Tradeoffs Across Sydney Intersections

David M. Levinson
School of Civil Engineering, The University of Sydney
david.levinson@sydney.edu.au

1 Linked Data and Calibration

Study date and weather. The primary analysis uses Tuesday, 18 September 2018, selected arbitrarily. Sydney (Observatory Hill) recorded a minimum temperature of 10.2 °C and a maximum of 21.7 °C for that date. Rainfall was zero in both 24-hour reporting periods ending at 9 am on 18 and 19 September [Bureau of Meteorology, 2018]. These observations indicate mild, dry conditions at this station; they do not measure exposure or walking conditions at every intersection. The analysis does not estimate weather effects. The separate 14–20 September panel assesses variation across seven dates at 35 sites with complete inputs; it does not establish seasonal or network representativeness.

Analysis samples. The current-condition comparison uses 2,226 matched site-hours at 168 sites. For the 30-second experiment, 161 sites satisfy the other eligibility criteria, and the counterfactual flow-to-capacity screen excludes 66, leaving 95. Within the retained 95 sites, 66 save pedestrian time while adding vehicle delay and 29 save time for both modes. The two counts of 66 describe different groups. Main-text Table 1's aggregate allocation ratio uses the 66 vehicle-cost sites; its all-site change uses all 95 eligible sites. The 35-site seven-day panel is selected separately by input availability and does not apply the congestion screen.

Table 1 collects the operational terms used throughout the article and this Supplemental Information (SI).

The current-condition comparison joins:

- calibrated pedestrian-crossing and delay estimates by site-hour;
- decoded SCATS vehicle-detector counts by approach-hour;
- vehicle signal-group and pedestrian Walk windows extracted from SCATS histories;
- a consolidated phase timeline from the same histories;
- spatial detector, lane, path, phase, and signal-group features; and
- site and street metadata.

The first-actuation calibration is published in *Findings* [Levinson, 2026b]; the network application is under review [Levinson, 2026a]. The 14 manual-observation cases and 24 video-observation

Table 1: Operational terms

Term	Meaning in this study
Site or intersection	One signalised location operated by one SCATS controller.
Approach	One incoming intersection leg and travel direction for vehicles.
Movement	One pedestrian crossing direction or vehicle path through an intersection.
Phase	A controller interval serving a compatible set of movements. A phase boundary is the time when the controller changes phases.
Signal group	The controller output governing a particular set of signal displays.
Call or actuation	A pedestrian-button request accepted by the controller. A locked call remains active until served.
Walk and restrictive interval	Walk permits pedestrians to start crossing. The restrictive interval used here begins when flashing-don't-walk starts and ends at the next Walk.
Arrival event and crossing	An arrival event is one independently arriving pedestrian or group; a crossing counts one pedestrian.
Site-, approach-, or movement-hour	The named spatial unit observed during one clock hour.
Detected and estimated count	Vehicle counts cover represented detector lanes. Pedestrian crossings are estimated from first-actuation timing and calibration factors.
Passage and trip	A crossing or detected vehicle is one intersection passage, not a complete trip. One walking trip may include multiple crossings.
Current and counterfactual	Current uses the reconstructed observed signal timing. Counterfactual uses the hypothetical timing after eligible phase boundaries are moved.

cases form one 38-case transfer distribution. The mean factor estimates and empirical central 95% intervals are:

- pedestrian-crossing adjustment, $\overline{\widehat{\alpha}}_N = 2.203$ (1.417–3.138);
- delay adjustment assuming 100% compliance, $\overline{\widehat{\alpha}}_{D,100} = 2.212$ (1.217–3.671); and
- observed-behaviour delay, $\overline{\widehat{\alpha}}_{D,\mathrm{beh}} = 1.655$ (0.886–2.738).

Each transfer calculation applies the count, full-compliance delay, and observed-behaviour delay factors from one calibration case. Keeping the factors together preserves their relationship when calculating delay per crossing. The central 95% range is the percentile range across the 38 transferred results; it is not a confidence interval and does not include SCATS event-parsing, detector-mapping, or model-specification uncertainty. The calibration sites and dates differ from this application, so operational use would require contemporaneous local calibration or evidence that these factors transfer.

2 Vehicle Detector Mapping

The controller source numbers vehicle detectors, but the available map file gives each detector location only an ordered map-feature identifier. The workflow therefore treats feature order within an intersection as the detector's numbered *slot*. It links each detector point to the nearest mapped lane, determines whether the lane runs toward or away from the intersection centre, and uses mapped vehicle paths to identify the phases and signal groups that serve the lane.

Only one active detector is retained for each inferred inbound lane: the detector closest to the intersection. This reduces double counting where the same lane has both an upstream detector and one near the stop line. The approach count N^v_{ash} is the sum over those selected represented lanes. It is not an independently verified count of every vehicle entering the intersection.

Table 2: Core notation

Symbol	Definition	Unit
s, h, a, e, k	Site, clock hour, vehicle approach, eligible first actuation, and editable phase-boundary indices	–
p, v	Pedestrian and vehicle labels	–
$q \in \{100, \text{beh}\}$	Full Walk-compliance and observed-behaviour pedestrian-delay targets	–
$\widehat{N}^{p,e}_{sh}, \widehat{N}^{p}_{sh}$	Estimated arrival events and pedestrian crossings	events; persons
N^{v}_{ash}	Detected vehicle count	vehicles
$\widehat{D}^{p,e}_{sh}, \widehat{D}^{p}_{shq}$	Unexpanded event-delay and calibrated pedestrian delay	event-s; person-s
$\widehat{D}^{v}_{sh}$	Estimated vehicle delay	vehicle-s
$\widehat{\bar{d}}^{p}_{shq}, \widehat{\bar{d}}^{v}_{sh}$	Estimated average pedestrian and vehicle delay	s/crossing; s/vehicle
$\Delta\widehat{D}^{p}, \Delta\widehat{D}^{v}$	Pedestrian person-hours saved and vehicle-hours added	person-h; vehicle-h
ρ	Detected flow divided by represented green-time capacity	–
$\widehat{R}_{PV}$	Pedestrian person-hours saved per vehicle-hour added	persons/vehicle
ω	Mean vehicle occupancy	persons/vehicle
θ_p, θ^*_p	Pedestrian-wait policy weight and its break-even threshold	–

Table 3: Detector-slot mapping coverage in the analysis samples. "Complete sites" have a spatial mapping for every active detector slot; this does not independently verify the inferred slot identities or prove that every physical inbound lane has a detector.

Sample	Sites	Active slots	Mapped	Mapped (%)	Complete sites
Current condition	168	1,990	1,872	94.1	100
30-second primary	95	1,074	1,008	93.9	55

The archive-wide mapping rate is not the relevant analysis denominator. The 9,811 active detector-day records cover 1,039 sites, whereas spatial detector geometry maps 2,334 active records at 221 sites. The matched analysis is drawn from those mapped sites. Table 3 reports coverage within the actual analysis samples.

Every mapped site has consecutive ordered detector shape identifiers, but detector-slot assignment remains inferred until checked against controller personality or configuration data. Unmapped active slots may represent inbound, outbound, or duplicate upstream detectors; the available geometry does not distinguish those cases. If an unmapped slot is an omitted inbound lane and its delay change resembles that of the mapped lanes, vehicle delay added is understated, $\widehat{R}_{PV}$ is overstated, and θ^*_p is understated. That direction is conditional: the data cannot identify what the unmapped slots represent. The complete-slot subset in 7 is therefore a composition sensitivity, not a lane-coverage correction.

3 Vehicle Signal Delay

For approach a at site s during clock hour h, green intervals are merged across serving signal groups. At each second for which the next serving green is known, a hypothetical arrival receives zero signal delay during green and the remaining time to the next green during red. Let $\mathcal{T}_{ash}$ contain these usable seconds, let $T_{ash} = |\mathcal{T}_{ash}|$ be their duration, and let N^{v}_{ash} be the detected vehicle count. The

Table 4: Matched current condition on 18 September 2018. Pedestrian counts are estimated crossings; vehicle counts are detected vehicles.

Quantity	Pedestrian, 100%	Pedestrian, behaviour	Vehicle
Estimated crossing or detected-vehicle count	328,421	328,421	4,581,992
Estimated total delay (person-h; vehicle-h)	3,632.4	2,717.1	13,279.3
Estimated average delay (s/crossing; s/vehicle)	39.8	29.8	10.4
Crossing-to-vehicle average-delay ratio	3.82	2.85	1.00
Aggregate current pedestrian-person-hour to vehicle-hour ratio	0.274	0.205	1.000

function $r^v_{ash}(t)$ is the delay assigned to an arrival at time t. Under uniform arrivals and no residual queue,

$$\widehat{\bar{d}}^v_{ash} = \frac{1}{T_{ash}} \int_{\mathcal{T}_{ash}} r^v_{ash}(t)\, dt, \qquad \widehat{D}^v_{ash} = N^v_{ash} \widehat{\bar{d}}^v_{ash}. \tag{1}$$

The interval before the first observed green is excluded because its start may fall before the data extract. Apparent red periods longer than three times the site's usual signal cycle are also excluded as likely missing or truncated history; the network median cycle length is used when a site's value is unavailable. The final observed green interval is included. Time after its end is excluded because the next green onset is unknown.

The detected count covers the full hour, so Equation 1 assumes that average delay during the usable seconds also represents the excluded seconds. The estimate does not represent vehicles still queued when green ends, groups of vehicles progressing together from an upstream signal, queues extending into an upstream intersection, or deliberate coordination between signals. A simple capacity check assumes 1,800 vehicles can pass per hour of green in each represented lane. Thus, the flow-to-capacity ratio is detected hourly flow divided by 1,800 times the number of represented lanes and their green-time share. The matched analysis retains an approach-hour only when this ratio is below 0.75 and at least 3,500 of the hour's 3,600 seconds can be evaluated. The 0.75 threshold is a conservative modelling screen, not an estimated breakdown point: as flow approaches estimated capacity, residual queues become more plausible and violate the model's no-residual-queue assumption. Among pedestrian-positive site-hours, it removes 1,654 otherwise qualifying vehicle approach-hours. It removes no entire current-condition site or matched site-hour because each retains at least one qualifying approach. The reduction from 174 pedestrian sites to 168 matched sites is instead due to other vehicle estimability or coverage requirements. These assumptions are more consequential in the counterfactual because shortening greens can create residual queues that are absent from the fixed-count model.

4 Matched Sample and Calibration Transfer

Pedestrian movements and qualifying vehicle approaches are aggregated to common site-hours. Vehicle aggregation is count-weighted:

$$N^v_{sh} = \sum_a N^v_{ash}, \qquad \widehat{D}^v_{sh} = \sum_a \widehat{D}^v_{ash}, \qquad \widehat{\bar{d}}^v_{sh} = \frac{\widehat{D}^v_{sh}}{N^v_{sh}}. \tag{2}$$

Pedestrian movement-hours are aggregated analogously by summing estimated crossings and person-delay before division. The matched file contains 2,226 site-hours at 168 sites.

Average pedestrian delay varies with $\widehat{\alpha}_D/\widehat{\alpha}_N$, while total pedestrian person-hours vary with

Table 5: Symmetric cross-modal flow–delay regressions across 2,226 matched site-hours at 168 sites. Outcomes and flows enter as natural logarithms. Both models include site and clock-hour fixed effects. Standard errors are clustered by site.

	Average pedestrian delay			Average vehicle delay		
Predictor	Estimate	SE	p	Estimate	SE	p
ln(detected vehicle flow)	0.0673	0.0313	0.032	0.2453	0.0963	0.011
ln(estimated pedestrian flow)	−0.1615	0.0219	< 0.001	0.0475	0.0231	0.040
Site fixed effects		Yes			Yes	
Clock-hour fixed effects		Yes			Yes	
Observations		2,226			2,226	
Sites (clusters)		168			168	
Within R^2		0.105			0.058	

$\widehat{\alpha}_D$. Vehicle quantities do not vary with pedestrian calibration. Applying all 38 case-specific factor sets gives an average-delay-ratio interval of 2.685–4.788 and a pedestrian-person-hour to vehicle-hour interval of 0.151–0.454 assuming 100% compliance. The corresponding observed-behaviour intervals are 1.732–3.916 and 0.110–0.339.

The probability densities are calculated across site aggregates. Sites contribute between 1 and 22 matched hours (median 15). A site's total delay therefore reflects both the delay accumulated during the observed hours and the number of hours available for that site. By contrast, average-delay distributions divide total delay by estimated crossings or detected vehicles, so they compare delay per passage rather than accumulated delay.

5 Cross-modal flow–delay associations

Across the 2,226 matched site-hours, we also test whether flow in one mode is associated with average delay in the other. Variables enter as natural logarithms. Regressions include site and clock-hour fixed effects, and standard errors are clustered by site. Calibration multipliers cancel from these log associations because they scale every observation by the same constant.

Controlling for pedestrian flow, a 10% increase in detected vehicle flow is associated with approximately 0.67% higher average pedestrian delay (elasticity 0.067, SE 0.031, $p = 0.032$). Symmetrically, controlling for vehicle flow, a 10% increase in estimated pedestrian flow is associated with approximately 0.47% higher average vehicle delay (elasticity 0.047, SE 0.023, $p = 0.040$). These are descriptive associations under observed timing, not causal effects or tests of the counterfactual delay model. Because each average-delay outcome contains its own mode's flow in the denominator, the own-mode flow coefficients serve primarily as adjustment terms.

Notably, higher vehicle flow is associated with higher average delay for both pedestrians and vehicles, whereas higher pedestrian flow is associated with lower average pedestrian delay and only a small increase in average vehicle delay. These patterns are consistent with causal interactions under adaptive signal control, but the observational data do not identify their causal magnitudes.

6 Phase Timelines and Pedestrian-Event Alignment

SCATS history-export software is run in a reproducible wrapper that combines its output into one phase timeline for every successfully processed intersection. The resulting file contains 409,564

phase intervals at 229 sites. All 168 sites in the descriptive pedestrian–vehicle sample have a phase timeline.

For each included first actuation, the method finds the next Walk start on the same controller clock as the phase timeline. The phase active at Walk is treated as the phase serving that pedestrian movement, and the immediately preceding phase is the one that could be shortened. If Walk appears up to three seconds before an independently mapped serving phase begins, it is assigned to that phase to allow for small reporting offsets between event streams.

For the timing experiment, pedestrian events and vehicle approaches are restricted to exactly the same site-hours. Where the independent spatial pedestrian-phase lookup is available, 11,159 event assignments agree with it. Sixty-three checkable events disagree and are excluded, a 0.56% mismatch rate among 11,222 checkable events. A further 69,534 events have direct temporal alignment but no separate spatial pedestrian-phase lookup. In total, 80,693 events align, representing 99.96% of candidate pedestrian delay in those hours.

Vehicle movement paths map each phase to its signal groups. A transition is editable only when at least one target vehicle green begins and at least one preceding vehicle green ends within 30 seconds of the observed phase boundary. Both sets of green endpoints are moved by the same amount, preserving their observed intergreen.

7 Network Counterfactual

Let k index an editable phase boundary and $\mathcal{E}_k$ its eligible first actuations. For actuation e, R_e is the wait to Walk; L_{k-1} is the preceding phase duration; $c \in \{15, 20, 30, 45\}$ seconds is the tested wait cap; and $L_{\min} = 10$ seconds is the retained preceding-phase floor. The boundary-level shift is

$$\delta_k = \min \left\{ \max_{e \in \mathcal{E}_k} \max(R_e - c, 0),\ \max(L_{k-1} - L_{\min}, 0) \right\}. \tag{3}$$

Thus, when several movements call the same boundary, the largest required shift is applied once. The floor may prevent the boundary from attaining c. Pedestrian and vehicle delay are then recomputed on the altered signal-group timeline. The 10-second floor is an analytical safeguard, not a claim that every site-specific controller minimum is represented.

At least one boundary can be shifted under the 30-second cap at 165 of the 168 paired sites. This numerical match to the 165 current-condition sites where pedestrian average delay exceeds vehicle average delay is coincidental; the two counts describe different properties. The primary criteria retain 85, 89, 95, and 100 sites at caps of 15, 20, 30, and 45 seconds, respectively; 82 sites qualify under every cap. A site must satisfy all of the following:

- at least 80% of candidate pedestrian delay aligns to the phase history;
- at least 10 phase boundaries are shifted;
- current vehicle approach-hours have at least 3,500 seconds of signal exposure and a flow-to-capacity ratio below 0.75; and
- no retained approach-hour reaches 0.75 under the counterfactual.

Among sites satisfying the other primary criteria, the counterfactual flow screen excludes 76 of 161 sites at the 15-second cap, 72 of 161 at 20 seconds, 66 of 161 at 30 seconds, and 59 of 159 at 45 seconds. The 30-second primary sample therefore contains the remaining 95 sites.

Table 6: Network 30-second-cap counterfactual. Totals are calculated from unrounded components; displayed components may therefore not sum exactly to displayed totals.

Site class	Sites	Pedestrian person-h saved		Vehicle h changed
		100%	Observed behaviour	
Vehicle cost	66	335.96	251.31	+1,049.25
Vehicle saving	29	103.74	77.60	−163.32
All primary sites	95	439.70	328.90	+885.93

Table 7: Fixed-demand delay in the 95-site primary sample

Mode and pedestrian target	Current	Counterfactual	Change
	seconds per estimated crossing or detected vehicle		
Pedestrian, assuming 100% compliance	38.05	26.46	−11.59
Pedestrian, observed behaviour	28.46	19.79	−8.67
Vehicle	9.55	11.01	+1.46

Let subscript 0 denote current timing and subscript c the wait-cap counterfactual. At vehicle-cost sites,

$$\Delta\widehat{D}^p = \widehat{D}^p_0 - \widehat{D}^p_c, \qquad \Delta\widehat{D}^v = \widehat{D}^v_c - \widehat{D}^v_0, \qquad \widehat{R}_{PV} = \frac{\Delta\widehat{D}^p}{\Delta\widehat{D}^v}. \tag{4}$$

The first difference is pedestrian person-hours saved, the second is vehicle-hours added, and $\widehat{R}_{PV}$ reports pedestrian-hours saved per vehicle-hour added. To expose rather than conceal the value judgment, let θ_p weight pedestrian waiting relative to vehicle-occupant waiting. The weighted balance is $\theta_p\Delta\widehat{D}^p - \omega\Delta\widehat{D}^v$, which is positive when $\theta_p\widehat{R}_{PV} > \omega$. The special case $\theta_p = 1$ is reported only as an equal-clock-time sensitivity, not as a welfare measure. It omits pedestrians' exposure to weather and traffic, vulnerability as unprotected road users, and discomfort relative to people inside vehicles.

Among the 66 vehicle-cost sites, the median 100%-compliance delay-transfer ratio is 0.274 (interquartile range 0.140–0.648). This site median happens to equal the aggregate current-delay ratio in 4; the two quantities have different samples and meanings. At nine sites the delay-transfer ratio exceeds 1.2 under every one of the 38 case-specific 100%-compliance delay factors.

Table 8: Calibration-transfer sensitivity at the 66 vehicle-cost sites

Pedestrian target	Person-h saved	Empirical interval	$\widehat{R}_{PV}$\$persons/vehicle)	Empirical interval
Assuming 100% compliance	335.96	184.87–557.47	0.320	0.176–0.531
Observed behaviour	251.31	134.59–415.88	0.240	0.128–0.396

Table 9: Fixed-count cap and composition sensitivity at $\omega = 1.2$ using the equal-clock-time special case $\theta_p = 1$. The first panel holds composition fixed at the 82 sites eligible under all four caps. The second restricts that sample to 52 sites where every active detector slot has a spatial mapping. Positive balances denote more pedestrian-hours saved than vehicle-occupant hours added. These balances omit exposure, vulnerability, and comfort differences. Totals are calculated from unrounded values.

Sample	Cap (s)	Pedestrian h saved 100%	Pedestrian h saved behaviour	Vehicle h changed	Equal-clock-time balance h 100%	Equal-clock-time balance h behaviour	Pareto sites
Common 82	15	453.1	338.9	+732.7	−426.2	−540.4	26
	20	423.2	316.6	+655.0	−362.8	−469.4	26
	30	359.7	269.1	+502.4	−243.2	−333.8	27
	45	263.2	196.8	+311.0	−110.1	−176.4	30
All active slots mapped (52)	15	297.3	222.4	+300.6	−63.4	−138.4	18
	20	275.9	206.4	+264.6	−41.6	−111.1	18
	30	229.8	171.9	+196.8	−6.4	−64.3	18
	45	159.7	119.5	+120.5	+15.1	−25.2	18

Within the common 82-site sample, the 52 complete-slot sites average 10.4 active detector slots and 21,668 detected vehicles per site; the 30 incomplete-slot sites average 12.3 slots and 23,296 vehicles. At the 30-second cap, calibrated pedestrian savings are similar at 4.42 versus 4.33 hours per site. Vehicle exposure differs only modestly, but the modelled delay increase is 0.63 seconds per vehicle at complete-slot sites versus 1.57 at incomplete-slot sites, yielding 3.79 versus 10.19 vehicle-hours added per site. The data identify this difference in modelled timing response but not whether it arises from geometry, signal timing, detector coverage, or another compositional feature. The subset remains a composition sensitivity, not a lane-coverage correction. The Pareto count of 18 in the 52-site subset is unchanged across all four caps.

No-flow-screen and pedestrian-weight sensitivities

Because SCATS adjusts green allocation adaptively, sustained oversaturation and spillback may be less common than a fixed-time capacity proxy suggests. Neither queues nor spillback are observed directly in these data. We therefore repeat the experiment while retaining the 3,500-second exposure and delay-estimability requirements but applying no flow-to-capacity exclusion to either current approach-hours or counterfactual sites. This restores the high-ratio approaches within the screened sites as well as the sites rejected only by the counterfactual flow screen.

For each cap and specification, Table 10 gives the pedestrian-wait weight $\theta_p^* = \omega\Delta\widehat{D}^v/\Delta\widehat{D}^p$ at mean vehicle occupancy $\omega = 1.2$. This is a deliberately low illustrative value between Transport for NSW estimates of 1.10 occupants for Sydney work trips and 1.46 for all weekday trips; the detector data do not measure occupancy [Bureau of Transport Statistics, 2014]. The low value favours the retiming case by lowering θ_p^*. The weighted balance is zero at θ_p^* and positive only when $\theta_p > \theta_p^*$. Thus, for example, the screened 30-second result requires pedestrian waiting to receive 2.42 times the weight of vehicle-occupant waiting under full compliance, or 3.23 times

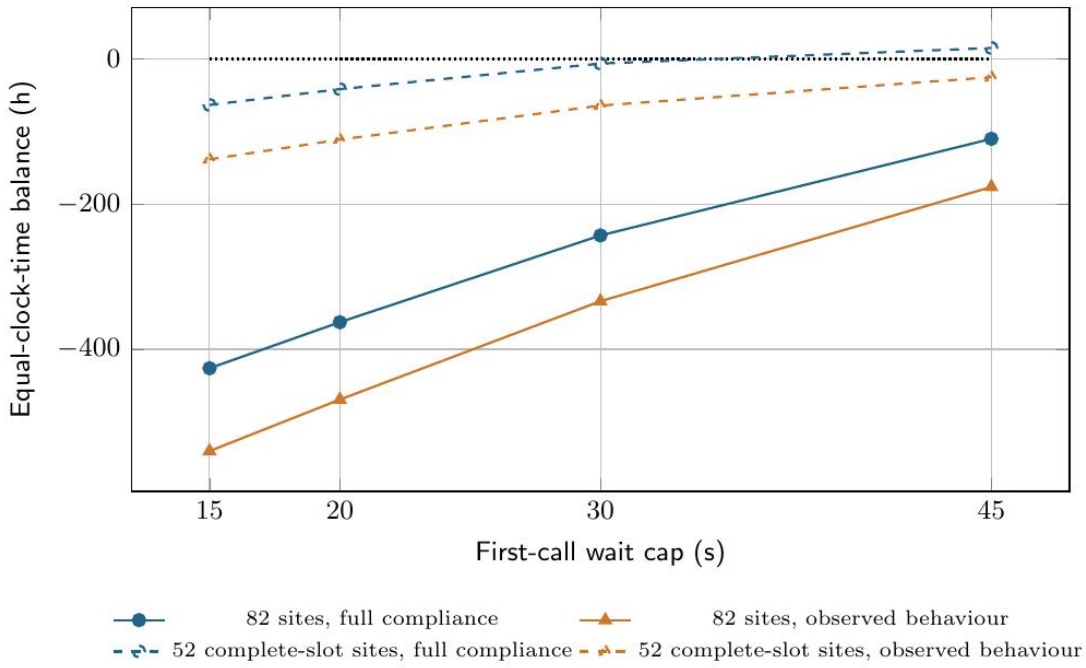


Figure 1: Cap and composition sensitivity for the equal-clock-time balance at occupancy 1.2. Solid lines hold the sample at 82 sites eligible under every cap; dashed lines restrict it to the 52-site complete-slot composition subset. Blue circles assume full compliance and orange triangles use observed behaviour. The dotted line is zero.

under observed behaviour. These are 142% and 223% above equal weighting. At occupancy 1.46, the corresponding thresholds would be 2.94 and 3.93.

On the original 95 sites at the 30-second cap, restoring 818 high-ratio approach-hours changes aggregate vehicle delay from +885.93 to +773.36 hours because those restored approaches save 112.57 vehicle-hours in aggregate. The full-compliance equal-clock-time balance changes from −623.42 to −488.33 hours; the pedestrian result is unchanged because the site sample is fixed. Restoring the approaches also raises the number of sites where both modes benefit from 29 to 36. Across all 161 otherwise-eligible sites, 1,650 of 10,694 current approach-hours have ratios at or above 0.75. Table 11 holds that current classification fixed when recomputing counterfactual delay.

Only 142 of 10,694 approach-hours (1.33%) have a current ratio above 3. At 30 seconds those rows save 53.53 vehicle-hours; all rows at or below 3 add 4,371.09 hours. Extreme ratios therefore slightly offset rather than drive the all-site vehicle-delay increase. At the site level, however, high-ratio rows can occur alongside lower-ratio approaches that add delay. In the untrimmed 95-site no-screen case, Site 750 contributes 183.60 of 773.36 vehicle-hours (23.7%), and the top 10 sites contribute 80.8%. Dropping all 19 sites with any current approach-hour above 3 leaves 76 sites, 348.89 full-compliance pedestrian-hours saved and 465.14 vehicle-hours added. The resulting θ_p^* is 1.60 under full compliance and 2.14 under observed behaviour. This is a whole-site composition sensitivity, not a correction. The screened primary result is less concentrated: its largest site contributes 9.6% and its top 10 contribute 58.4%.

The 161-site no-screen sample saves 1,330.79 pedestrian person-hours under full compliance and 995.46 under observed behaviour. Its equal-clock-time balances are $-3,850.27$ and $-4,185.61$

Table 10: Pedestrian-wait weights required for a positive aggregate balance. Screened samples apply the 0.75 rule. "Same sites" restores all otherwise estimable approaches at those screened sites. "All eligible" also admits sites meeting the other primary criteria. Totals are calculated from unrounded values.

Specification	Cap	Sites	Vehicle h changed	θ_p^*, 100%	θ_p^*, behaviour
Screened	15	85	+738.2	1.90	2.54
	20	89	+858.2	2.16	2.89
	30	95	+885.9	2.42	3.23
	45	100	+644.7	2.08	2.77
No screen, same sites	15	85	+677.1	1.74	2.33
	20	89	+772.0	1.94	2.60
	30	95	+773.4	2.11	2.82
	45	100	+610.4	1.96	2.63
No screen, all eligible	15	161	+5,584.9	4.30	5.75
	20	161	+5,189.0	4.18	5.59
	30	161	+4,317.6	3.89	5.20
	45	159	+3,056.2	3.47	4.64

Table 11: Thirty-second, 161-site no-screen vehicle decomposition. Approach classes use the current flow-to-capacity ratio. Totals are calculated from unrounded values.

Current approach class	Approach-h	Vehicles	Current s/vehicle	Changed s/vehicle	Vehicle h changed
Ratio below 0.75	9,044	4,512,731	10.55	14.30	+4,697.46
Ratio 0.75 to 3	1,508	1,319,723	35.33	34.44	−326.37
Ratio above 3	142	187,301	55.41	54.38	−53.53
All approaches	10,694	6,019,755	17.38	19.96	+4,317.55

hours. Maximum proxy ratios of 13.85 currently and 13.01 after retiming should not be interpreted as measured physical saturation; they instead expose detector, represented-lane, green-share, or model mismatch. Consequently, the no-screen result is a sensitivity to exclusion rules, not validation of the no-residual-queue assumption.

Assumption-driven local-count sensitivity

The counterfactual averages above retain current estimated crossings and detected vehicle counts as denominators to isolate the immediate effect of retiming. Counts could change in the longer run if people change route, departure time, crossing location, or travel mode. The present data do not estimate that response. The following calculation is retained only to show how an imposed response assumption would alter local counts.

For site s, mode u, and wait cap c, the illustration assumes a constant response to that site's average signal delay:

$$\frac{\widehat{n}_{su,c}}{n_{su,0}} = \left(\frac{\widehat{\bar{d}}_{su,c}}{\widehat{\bar{d}}_{su,0}}\right)^{-\eta_u}. \tag{5}$$

Here $n_{su,0}$ is the current local count, $\widehat{n}_{su,c}$ is the assumption-based counterfactual count, and η_u is an imposed positive elasticity magnitude. A value of zero reproduces fixed counts. The values $\{0, 0.25, 0.5, 1\}$ are scenarios, not estimates. All 16 pedestrian–vehicle combinations are retained in the output; Table 12 shows the four cases in which the two modes receive the same assumed response.

Table 12: Assumption-driven one-step local count sensitivity at the 95 sites eligible for the 30-second scenario. Factors are medians of site-specific count multipliers; counterfactual counts are sums across sites, so aggregate count ratios need not equal the median factors. Pedestrian baselines are estimated crossings; vehicle baselines are detected counts. Both counterfactual count columns are estimates.

$\eta_p = \eta_v$	Median pedestrian factor	Median vehicle factor	Counterfactual crossings	Counterfactual vehicles
0	1.00	1.00	136,592	2,191,910
0.25	1.08	0.99	149,854	2,080,745
0.50	1.18	0.97	164,785	1,988,119
1.00	1.38	0.94	200,710	1,846,382

These values are generated by Equation 5. The calculation treats average delay at this one signal as though it were the entire cost to which users respond, although signal delay at one intersection is only one part of a trip's time and cost. It then sums independent site responses without assigning displaced users elsewhere, conserving network flows, distinguishing among behavioural mechanisms, or feeding changed vehicle counts back into queues and delay. A behavioural forecast would instead require total trip costs, empirically supported response parameters, network reassignment, and updated signal and queue conditions. Table 12 is therefore a stress test of the fixed-count assumption, not an equilibrium, welfare, or demand estimate.

8 Seven-Day Input Availability and Validation

We ran the complete pedestrian, detector, signal-delay, and phase-boundary pipeline for 14–20 September 2018 at every current-condition site with the inputs needed on all seven dates. A site was input-ready when every active detector slot had a spatial mapping, at least one representative inbound approach detector was available, pedestrian movements had a phase lookup, a spatial phase file was present, and nonempty regional history and vehicle-count files were available for all seven dates. This screen tests whether the calculation can be run; it does not select sites using the 0.75 flow-to-capacity diagnostic. The 35 sites are availability-selected and their representativeness of Sydney intersections overall is unknown.

All 35 input-ready sites were analysed for each of the seven days, yielding 245 site-days and results for four wait caps at every site. To make the one-day and seven-day estimates comparable, we applied the same mean pedestrian-delay calibration factor derived from all 38 cases. This adjustment does not affect the vehicle results or whether delay increases or decreases at a site.

Table 13: Availability and completion of the seven-day analysis. The one-day classifications refer to 18 September 2018.

Criterion	Sites
Current-condition sample	168
Every active detector slot spatially mapped	100
All required seven-day inputs present	35
Also in the one-day 30-second primary sample	25
Also a one-day 30-second vehicle-cost site	17
Seven-day pipeline completed successfully	35

9 Seven-Day Counterfactual Results

Table 14 reports the seven-day results for all 35 sites with complete data. The analysis retains observations with high estimated flow relative to capacity so that every input-ready site remains in the comparison. At those observations, residual queues are more plausible and the modelled vehicle delays are less reliable; the flow-to-capacity ratios are therefore reported as diagnostics.

At the 30-second cap, estimated pedestrian delay falls at 32 of the 35 sites. Eighteen sites trade pedestrian savings against added vehicle delay, 14 reduce delay for both modes, and three have no modelled change. Twenty-two of 35 sites retain the same effect sign on all seven days. The maximum counterfactual hourly flow-to-capacity ratio remains below 0.75 throughout the week at eight sites. Fifteen sites have at least 12 matched hours on the 18 September reference day; six satisfy both conditions.

The one-day median among the 66 vehicle-cost sites is 0.274 pedestrian person-hours saved per vehicle-hour added. Among the three lower-flow vehicle-cost cases in Table 15, Site 822 is closest at 0.352, but its weekly pedestrian saving is only 0.494 hours and one day has no modelled change. Site 922 has a stable sign but a ratio of 0.718; Site 516 is also stable but has an unusually high ratio of 2.910. Site 645, the closest input-ready one-day candidate, completes the run but reaches a counterfactual flow-to-capacity ratio of 5.31, so it is not a suitable lower-flow example.

The all-site multiday results remain fixed-count timing experiments. They do not impose every controller personality minimum, model residual queues or platoon progression, optimise coordination, or estimate behavioural response. High flow-to-capacity values identify observations for which those limitations matter more; they are not measurements of spillback.

Table 14: Seven-day counterfactual totals across 35 input-ready sites, assuming full Walk compliance. Vehicle-delay changes are positive when delay is added. The balance uses $\theta_p = 1$ and $\omega = 1.2$; θ_p^* is the pedestrian-wait weight required for break-even. Totals are calculated from unrounded values.

Cap (s)	Boundaries shifted	Pedestrian h saved	Pedestrian delay reduction (%)	Vehicle h changed	Balance (h)	θ_p^*	Tradeoff sites	Both benefit sites
15	35,250	603.20	40.0	+1,154.71	−782.45	2.30	19	13
20	32,350	547.43	36.3	+1,022.65	−679.75	2.24	19	13
30	26,072	429.12	28.4	+794.32	−524.06	2.22	18	14
45	16,844	265.57	17.6	+523.59	−362.73	2.37	17	15

Table 15: Thirty-second results for sites with at least 12 matched reference-day hours and a maximum seven-day flow-to-capacity ratio below 0.75. Negative vehicle changes are hours saved. The delay-transfer ratio is reported only where vehicle delay is added.

Site and intersection	Matched h	Pedestrian h saved	Vehicle h changed	$\widehat{R}_{PV}$	Maximum f/c	Stable sign
138, Pacific Hwy / Longueville Rd	16	10.154	−8.195	–	0.351	Yes
516, Willoughby Rd / Albany St	17	51.073	+17.551	2.910	0.558	Yes
822, Concord Rd / Patterson St	15	0.494	+1.405	0.352	0.504	No
881, Parramatta Rd / Marlborough Rd	17	3.768	−14.497	–	0.712	Yes
922, Lyons Rd / Marlborough St	17	20.925	+29.144	0.718	0.408	Yes
1081, Lilyfield Rd / Mary St	14	2.039	−6.944	–	0.524	Yes

10 Generated Outputs

The primary network files include vehicle detector mappings, vehicle approach-hour counts and delay estimates, calibrated pedestrian site-hour estimates, paired site-hour comparisons, phase timelines, counterfactual event and site summaries, cap sensitivities, and local count-response sensitivities. The seven-day files include the input-availability screen, completion status, combined site–cap results, and site-level current, approach, and counterfactual outputs. The revision package also includes the cap-sensitivity summary, detector-coverage audit, 38-case recalibrations, conditional flow–delay estimates, the approach- and site-level high-ratio sensitivities, and hashes of source files.

The workflow uses the following components, including upstream programs and scripts supplied in the supplementary archive:

- phase extraction, `BroadPhaseExporter.java`;
- detector mapping, `build_vehicle_detector_spatial_mapping.py`;
- the network intervention, `network_pedestrian_priority_counterfactual.py`;
- the seven-day availability screen, `audit_seven_day_linkage_capability.py`;
- the generic site analysis and resumable all-site worker, `analyze_site_multiday.py` and `run_all_site_multiday.py`;
- pooled Paper 3 transfer calculations and figures, `build_paper3_figures.py`;

- revision cap and detector sensitivities, `build_revision_sensitivity.py`;
- the no-flow-screen rerun and extreme-ratio decomposition, `build_no_flow_threshold_sensitivity.py`;
- the no-flow and seven-day 38-case recalibrations, `recalibrate_no_flow_threshold_outputs.py` and `recalibrate_seven_day_outputs.py`;
- conditional flow–delay associations, `analyze_cross_modal_flow_delay.py`;
- the whole-site high-ratio trim, `build_site_trim_sensitivity.py`;
- the normalised Figure 1 densities, `build_normalized_delay_density.py`; and
- pedestrian-wait thresholds, `build_policy_weight_thresholds.py`.

The supplementary archive contains derived outputs and the analysis scripts listed above, but excludes four upstream programs: `BroadPhaseExporter.java`, `build_vehicle_detector_spatial_mapping.py`, `network_pedestrian_priority_counterfactual.py`, and `build_paper3_figures.py`. The density builder also requires an external figure-builder module supplied with `--builder`. Re-running the upstream extraction and counterfactual steps requires the raw controller histories and associated local SCATS extraction software.

References

Bureau of Meteorology. Greater Sydney in September 2018: Relatively cool conditions and average rainfall, October 2018. URL https://www.bom.gov.au/climate/current/month/nsw/archive/201809.sydney.shtml. Daily observations for Sydney (Observatory Hill); accessed 27 September 2026.

Bureau of Transport Statistics. Household travel survey report: Sydney 2012/13. Technical Report 2014/11, Transport for NSW, 2014. URL https://www.transport.nsw.gov.au/sites/default/files/media/documents/2017/r2014-11-hts-summary-report_12-13-1.pdf.

David M. Levinson. Estimating pedestrian flows and signal delay from Sydney traffic-signal event histories. Manuscript under review, 2026a.

David M. Levinson. Testing pedestrian flow and delay estimation from first-actuation latency. *Findings*, 2026b. doi: 10.32866/001c.166369. URL https://doi.org/10.32866/001c.166369.