



HADRON
GROUP

Local Roads Speed Change Impact Assessment


May 2026



Acknowledgement of Country

We respectfully acknowledge the traditional custodians of the diverse Aboriginal lands on which we operate and where our projects unfold. We pay our respects to their elders past and present, and recognize the enduring connection between Aboriginal and Torres Strait Islander peoples and their ancestral territories.

As we work on these lands, we strive to uphold the principles of respect, cooperation, and environmental stewardship, guided by a shared commitment to a harmonious future.



Executive Summary

Overview

Evidence widely supports that lowering speed limits on local roads significantly reduces the risk of serious injury or fatalities from crashes that occur. Transport for NSW and other transport agencies in Australia and abroad cite that lowering vehicle impact speeds from 50 km/h to 40 km/h roughly halves a pedestrian's risk of being killed in a crash¹. Despite this, public discussion of speed-limit reductions is consistently shaped by a single concern: that lower limits will materially lengthen everyday journeys.

This report tests that concern using simulation. A Local Government Area (Randwick) was chosen as a test case. Hadron's Safer Speeds Simulator was applied to quantify the travel-time impact of three 40 km/h intervention scenarios, defined by NSW residential zoning (R1–R2 only, R3–R4 only, and R1–R4 combined), against a base case of today's posted speed limits. Approximately 26,500 trips per scenario were simulated across all suburbs and four destination groups.

Summary findings

The headline finding is that **the overall travel-time impact is small**. Even under Scenario 3, the broadest intervention covering 71% of road kilometres in the LGA, the average trip gains **12 seconds** – approximately 2% of base travel time, and less than a single phase of a typical traffic signal. Half of all trips gain less than 2 seconds.

Three further observations are worth highlighting:

- The **high-density-only** scenario (R3–R4) adds an average of just 3 seconds per trip. R3–R4 zoned streets cover only 20% of road kilometres in the LGA, so trips traverse a short section of them; targeting this subset alone produces a barely-measurable change.
- **Longer trips absorb the change**. Trips under 2 km gain 8 seconds (4.7% of base time), while trips over 9 km gain 11 seconds (just 1% of base time). Long journeys spend most of their minutes on arterial roads – which keep their existing limits – so the residential reduction affects only the first and last few hundred metres.
- **Effects vary modestly by destination type**. Beach and leisure trips see the largest change (14 seconds, 1.9%), reflecting their reliance on residential approach roads, while park trips see the smallest (9 seconds, 1.2%). All four destination groups sit within a narrow 9-to-14-second band.

The community-safety case for lower speed limits is well-established and is not reiterated here. The purpose of this report is to provide council with a quantitative, transparent answer to the question that most often dominates community discussion: *"how much extra time will these changes add to my journey?"* The answer, on the evidence of 26,500 simulated trips, is small enough that most drivers are unlikely to notice the change against the ordinary variation in their daily commute.

¹ Speeding | Transport for NSW

Table of Contents

1.	Introduction	4
2.	Methodology	5
2.1	Scenarios tested	5
2.2	Analysis approach.....	5
2.3	Sample origins and destinations	6
2.4	Analysis outputs	7
3.	Results	8
3.1	Summary results	8
3.2	Results by destination group.....	10
3.3	Results by suburb	11
4.	Key findings	14
5.	Applications and uses	15
	Technical appendix	16
	Data sources and trip sampling	16
	Suburb-level results	17

1. Introduction

Background and context

Local governments across Australia have been reducing speed limits on local roads in recent years to improve the safety of their communities. Transport for NSW's evaluation of permanent 40 km/h zones reported a 100% reduction in fatal pedestrian crashes on roads that moved to 40 km/h, compared with a 19% reduction on comparable 50 km/h roads².

Similar interventions have been delivered by the City of Sydney, Inner West Council, Northern Beaches Council and the City of Melbourne, alongside international precedents in London, Paris, Edinburgh and Helsinki.

The experience across these implementations has been consistent: strong evidence base for the safety benefit, broad community support once limits are in place, and concentrated opposition during consultation – with the perception that lower limits will materially lengthen everyday journeys. That belief is widely held but rarely tested at municipality scale. Individual studies typically estimate the effect on a specific corridor or single journey; council deliberations seldom have access to a transparent, area-wide answer.

Document structure

The remainder of this report is structured as follows.

Chapter 2 sets out the methodology: the four scenarios tested, the modelling approach, the sampling of origins and destinations, and the analysis outputs produced.

Chapter 3 presents the results, beginning with summary findings at LGA level, then breaking down impacts by destination group (Section 3.2) and by origin suburb (Section 3.3).

Chapter 4 summarises the key findings and observations.

Chapter 5 outlines potential applications and uses of the analysis.

The **technical appendix** documents data sources, the trip-sampling pipeline, and per-suburb results.

² Transport for NSW, *Speed fact sheet*, 2018

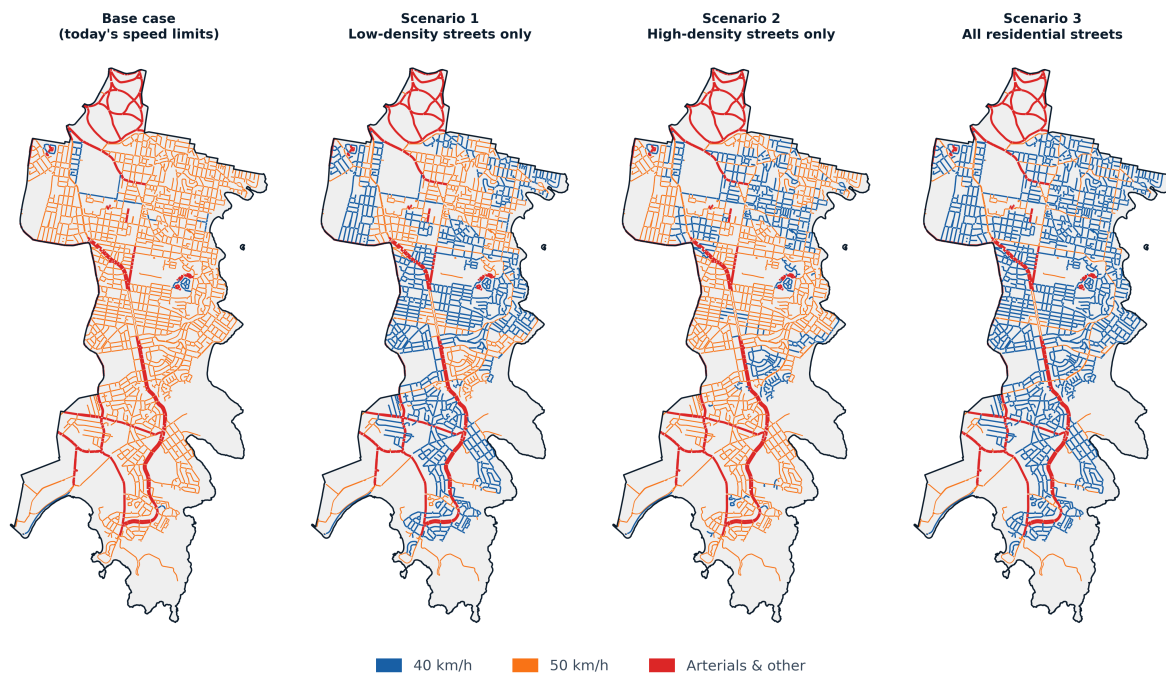
2. Methodology

2.1 Scenarios tested

The analysis includes a base case and three intervention scenarios. Each intervention scenario tests a target speed (40 km/h) across a selection of local roads. State roads are excluded from changes in all scenarios. Each scenario is explained in more detail below and illustrated in Figure 1.

- **Scenario 0 Base case:** existing posted speeds.
- **Scenario 1 Low-density streets only:** low-density residential streets (R1–R2) changed to 40 km/h (51% of LGA roads, 287 km). High-density (R3–R4), state & arterial roads unchanged.
- **Scenario 2 High-density streets only:** high pedestrian activity and high-density (R3-4) roads changed to 40kph (20% of LGA roads, 111 km). Low-density (R1-2), state & arterial roads unchanged.
- **Scenario 3 All residential streets:** high pedestrian activity and low- & high-density (R1–R4) roads changed to 40 km/h (71% of LGA roads, 398 km).

Figure 1: Scenarios included in the analysis



Source: SUMO scenario network clipped to the official Randwick LGA polygon (563 km of road).

To evaluate each scenario, example trips are modelled between origins within each Randwick suburb and a range of common destinations across the LGA and to key destinations outside the LGA. The approach is described in further detail below.

2.2 Analysis approach

The analysis uses **empty-network simulation**: each trip is simulated alone on the road network, with no other vehicles present, no congestion, no signal queuing, and no demand interaction between drivers. The simulation measures only the mechanical effect of a posted-speed change on a clear road – how much longer the same trip takes when residential streets sit at 40 km/h instead of today's posted limits.

This framing is deliberate. Real journey times are shaped by many factors – peak-hour traffic, school-zone slowdowns, parking, signal phasing, driver compliance – any of which can be larger than the effect of a 10 km/h speed-limit change. Modelling all of them together would obscure the specific question this report sets out to answer. By running every scenario on an empty network, the speed-limit change is the only variable moving between runs, so any difference in travel time is attributable to that change alone.

For each sampled origin, the simulator computes the fastest route to each destination point of interest – first under the base case, then under each intervention scenario. Routes are recomputed independently for each scenario, so where a slower residential street causes the model to choose a different path, that re-routing is captured. Travel times are recorded per trip; the difference between the intervention and base-case time is the metric reported throughout this study.

The simulation is built on **SUMO (Simulation of Urban Mobility)**, an open-source microscopic traffic simulator used internationally by transport agencies and researchers. Vehicle parameters are tuned for Australian urban driving, and the network is calibrated once against Google Maps free-flow times so that base-case durations match observed off-peak conditions; this calibration is held constant across all four scenarios. Full technical details are in the appendix.

Because the simulation excludes congestion, results in this report should be read as the **floor of the travel-time impact** – the unavoidable effect of a slower posted speed on a clear road. In peak-hour conditions, the absolute difference is generally smaller, not larger, because vehicles are already constrained by surrounding traffic and travelling well below the posted limit regardless.

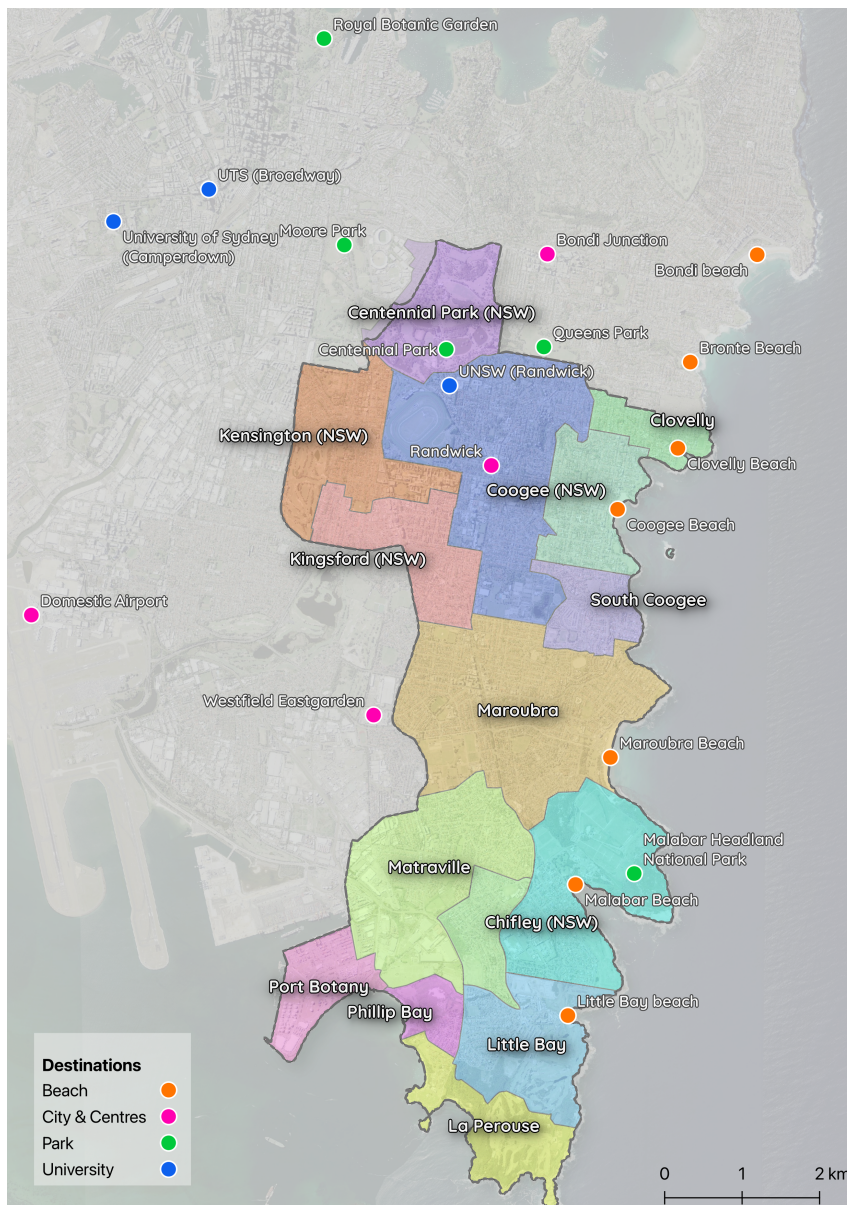
2.3 Sample origins and destinations

For each suburb within the Randwick LGA, we placed 100 sample trip-origin points inside the suburb boundary, distributed in proportion to where residents actually live, so the more populated parts of a suburb get more points. Trips are routed to the following key destinations:

- **Beaches:** Coogee, Clovelly, Maroubra, Malabar, Little Bay/La Perouse, Bronte, and Bondi.
- **City and major centres:** Randwick Centre, Sydney CBD, Bondi Junction, Westfield Eastgardens, and the Domestic Terminal.
- **Universities:** UNSW (Randwick), University of Sydney (Camperdown), and UTS (Broadway).
- **Parks:** Centennial Park, Queens Park, Moore Park, Malabar Headland National Park, and Royal Botanic Garden.

The suburb extents and points of interest are outlined below in Figure 2.

Figure 2: Suburb boundaries and key points of interest



2.4 Analysis outputs

The key outputs from the analysis include the following simulation outputs:

- **Average added travel time** per trip in absolute and relative terms, presented by destination groups and scenarios
- **A cumulative distribution** showing the spread of impact across all trips with key percentile thresholds
- **Breakdowns by trip length** group to identify where the impact concentrates
- **A suburb ranking and choropleth** to identify the origin suburbs most and least affected.
- **A worked example** from a representative suburb, illustrating sampled trips to the relevant destination under the base case and intervention scenario

3. Results

3.1 Summary results

Across all 26,543 simulated trips, changing residential streets to 40 km/h added an average of 12 seconds per trip in Scenario 3 (All residential streets), which changes the largest share of the road network. Restricting changes to high-density streets only added 3 seconds on average, which is less than the time spent at a single traffic light.

Figure 3: Summary analysis results

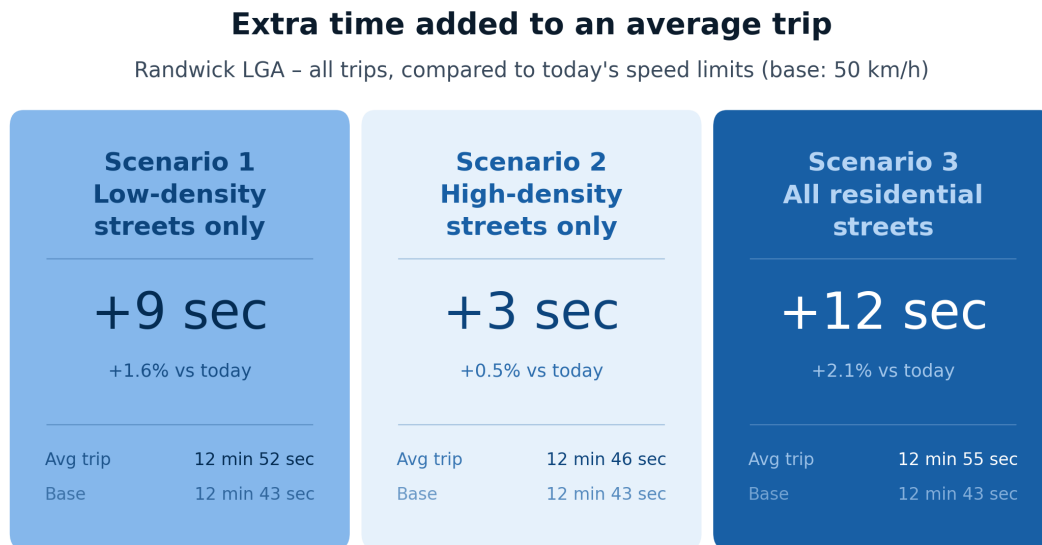
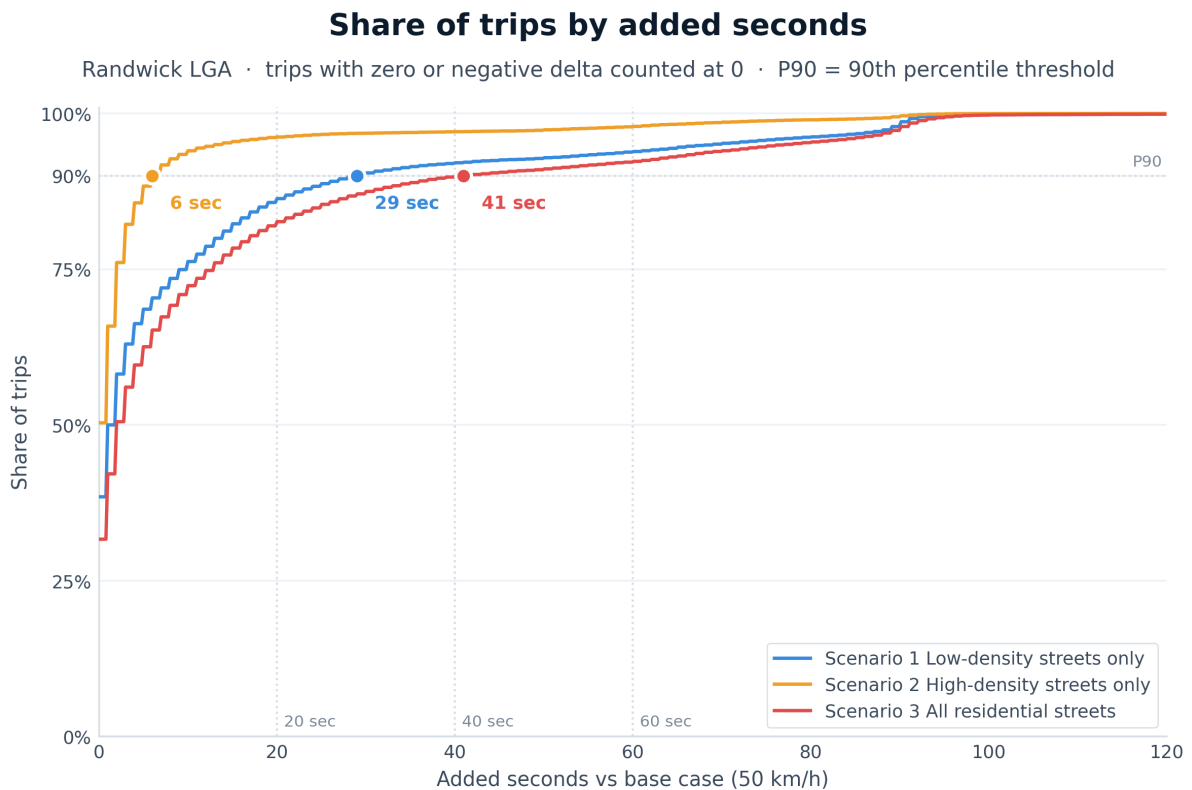


Figure 4 below shows how the added time is distributed across all trips. Under Scenario 2 (high-density streets), 9 out of 10 trips add less than 6 seconds. Even under Scenario 3 All residential streets, 9 in 10 trips add less than 41 seconds. The steep rise at the left of each curve confirms that most trips are affected only minimally, with a small number of longer trips pulling the tail to the right.

Figure 4: Share of trips by travel time impacts



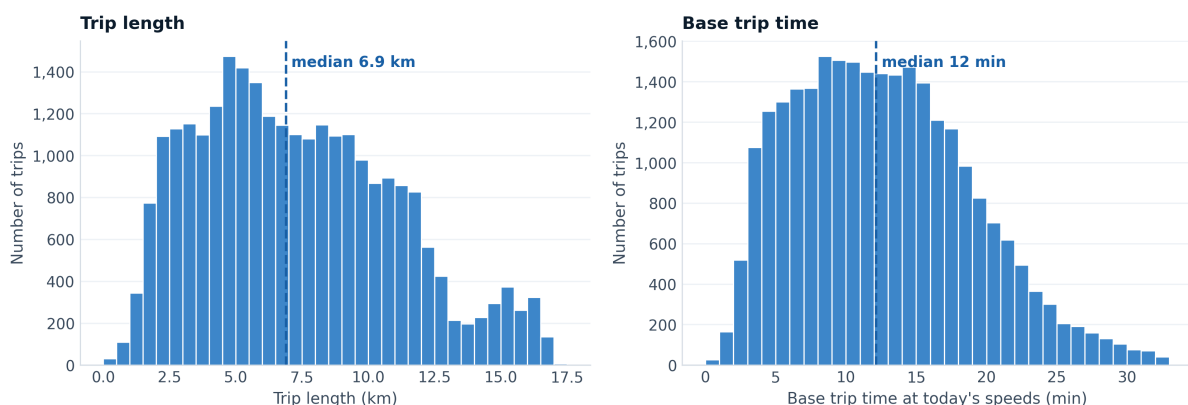
Source: Randwick SUMO model, n = 26,543 trips. Curves clipped at 120 seconds.

Figure 5 shows the distribution of the simulated origin–destination trips that underpin every scenario in this report. Trip lengths cluster tightly between roughly 2 and 12 kilometres, with a median of about 7 kilometres, and the long tail of trips above 12 kilometres is minimal. At current speed limits, these trips translate into base journey times of 5 to 17 minutes, with the typical trip taking around 12 minutes.

Figure 5: The distribution of trips length and time

What does an average trip look like?

Distribution across 26,543 simulated origin-destination trips. Most trips fall between 4 and 10 km, taking 8 to 17 minutes at today's speeds.




Source: Randwick SUMO model. Mean trip length 7.3 km · Mean base time 13 min. The top 0.5% of values are clipped for readability.

Counter-intuitively, longer trips are less affected. Trips under 2 km add 8 seconds on average – 4.7% of travel time. Trips over 9 km add 11 seconds – just 1% of travel time. Long journeys spend most of their minutes travelled on arterial roads, which keep their existing speed limits, so the lower residential limits affect only the first and last few hundred metres.

Figure 6: Average extra time per trip and a share of base trip time

Longer trips absorb the change

Scenario 3 All residential streets at 40 km/h - the same handful of seconds matters less the further you travel.



Source: Randwick SUMO model, n = 26,543 trips. Left bars: mean added seconds per band. Right bars: each trip's delay as a share of its own base time, averaged across the band.

3.2 Results by destination group

Travel time impacts vary by destination type but remain minimal across all scenarios. The metric used for each destination type reflects how residents actually use it: for beaches and parks we report the minimum travel time to the nearest option, while for universities and city centres we report the median travel time across all relevant destinations, since residents are typically heading to a specific campus or centre rather than the closest one.

On that basis, beach and leisure trips see the largest change under Scenario 3 All residential streets, with an average of 14 seconds added per trip. This reflects that the shortest path to the nearest beach typically passes through more residential streets than other journey types.

City centre and university trips add 11–12 seconds in Scenario 3. Park trips are the least affected, adding just 9 seconds on average.

Notably, the high-density streets scenario adds only 2–4 seconds across all destination types. This is because R3–R4 zoned streets cover roughly 20% of road kilometres in the LGA – about a third of what is changed under Scenario 3 – so most trips traverse only a short section of these streets. Targeting high-density residential streets alone therefore has a limited effect on travel times regardless of where people are going.

Figure 7: Summary travel time impacts by destination type

Extra travel time by destination type
Seconds added vs base case (50 km/h) – % change shown below

	Base case (50 km/h)	Scenario 1 Low-density streets only	Scenario 2 High-density streets only	Scenario 3 All residential streets
	12 min 43 sec	12 min 52 sec	12 min 46 sec	12 min 55 sec
Beach & leisure	11 min 58 sec	+11 sec +1.5%	+3 sec +0.4%	+14 sec +1.9%
City & centres	11 min 32 sec	+9 sec +1.3%	+2 sec +0.3%	+11 sec +1.6%
Park & green space	13 min 10 sec	+8 sec +1.0%	+2 sec +0.3%	+9 sec +1.2%
University & education	15 min 17 sec	+9 sec +1.0%	+4 sec +0.4%	+12 sec +1.3%

Source: Randwick SUMO model. Cell shade scales with seconds added.

3.3 Results by suburb

As illustrated in Figure 8, Scenario 2 high-density streets (orange) consistently produces the smallest impact across every suburb, in most cases adding just 1 to 5 seconds. By contrast, Scenario 1 low-density streets and Scenario 3 all-residential streets produce very similar amounts of added travel time in almost every suburb, with the orange and red markers landing close together. This is because low-density streets make up the bulk of Randwick's residential network, so once their speed limits drop, the additional reduction on high-density streets contributes only a small additional delay.

Figure 8: Additional travel time by suburb across all scenarios

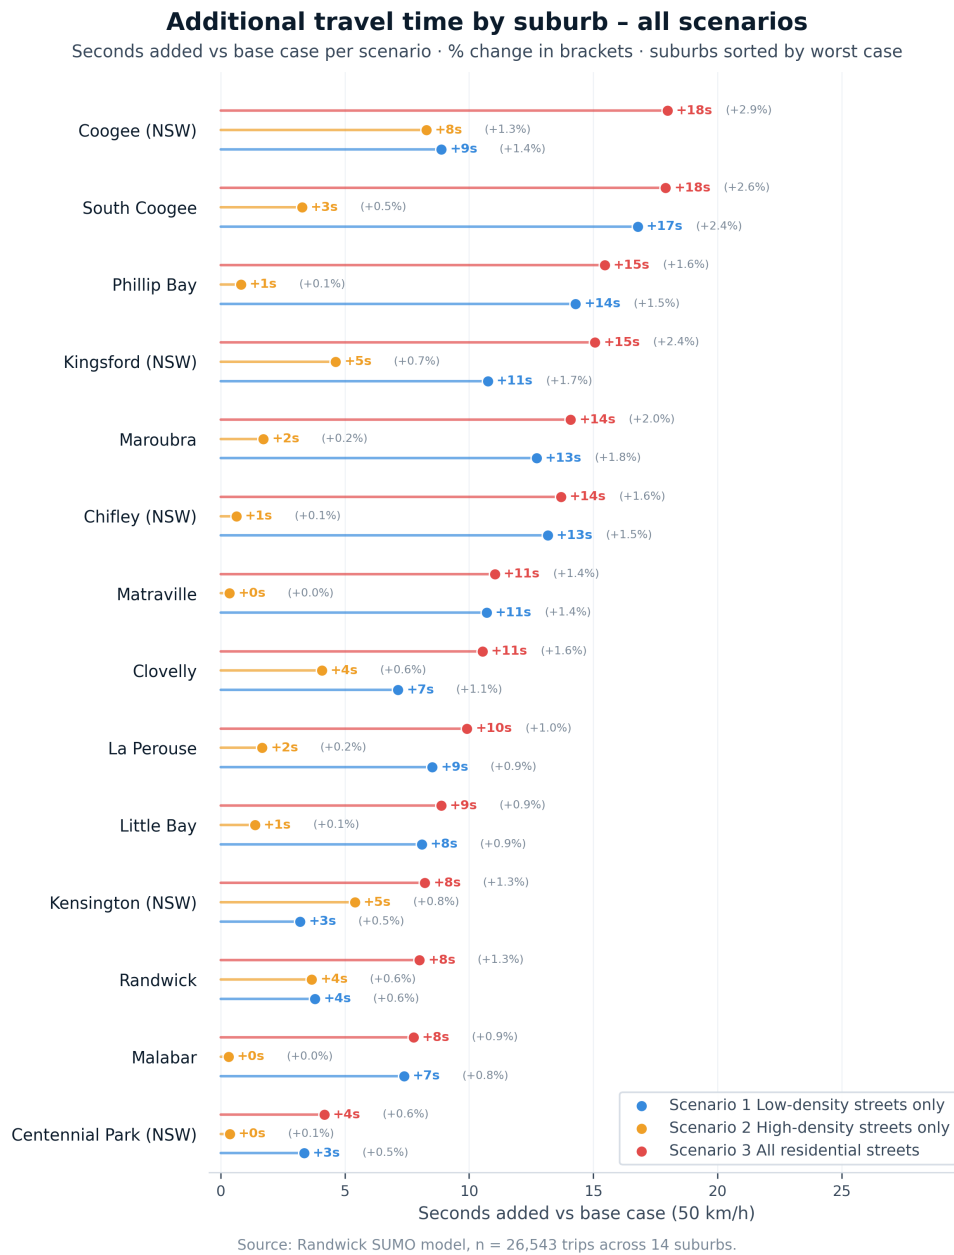
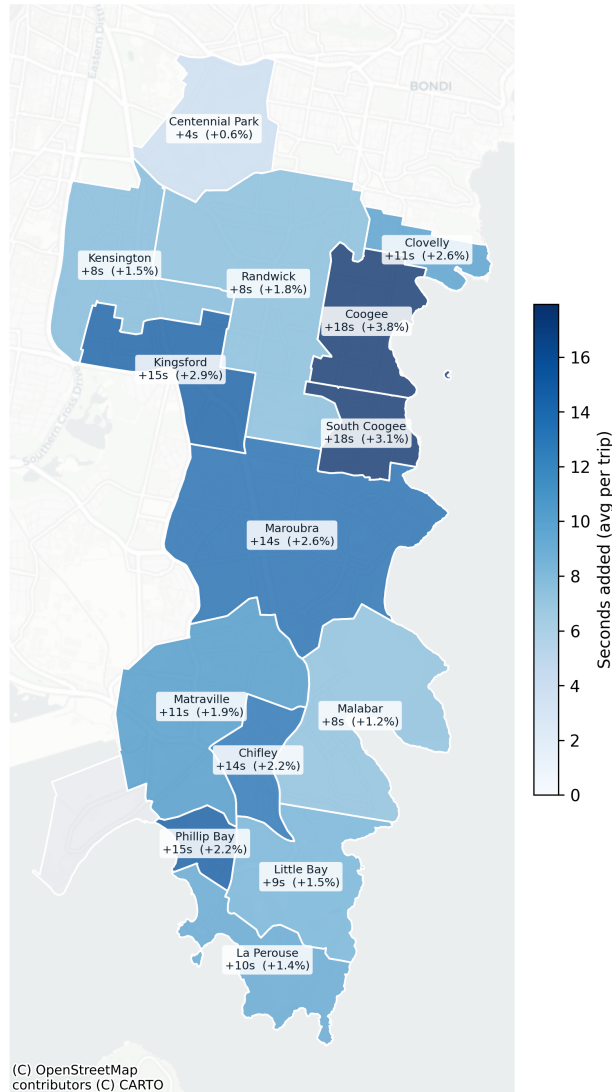


Figure 9 below illustrates the travel time impacts by suburb for Scenario 3 All residential streets. Coogee and South Coogee stand out with the largest impacts under this scenario (+18 seconds each), while Centennial Park sees the smallest change (+4 seconds). Notably, the spread between scenarios is narrow for most suburbs, for 10 of 14 suburbs, Scenario 3 All-residential streets adds 15 seconds or less.

Figure 9: Travel time impact by suburb – Scenario 3 (All residential streets)

Travel time impact by suburb - Scenario 3
 Randwick LGA · Scenario 3 All residential streets at 40 km/h · seconds added vs base case



Source: Randwick SUMO model, n = 26,543 trips. Basemap: CartoDB Positron / OpenStreetMap contributors.

The variation between suburbs reflects the share of residential roads within each suburb's network where the speed limit actually changes – suburbs that are predominantly residential, like Coogee and South Coogee, see more of every trip run on 40 km/h streets and therefore pick up more added time, while suburbs threaded with arterials and through-roads are less affected.

4. Key findings

Across all simulated trips, the analysis finds that reducing residential street speeds to 40 km/h results in minor travel time impacts under all three scenarios tested. The findings below summarise the headline results.

Overall, the impact of reducing residential speeds to 40km/h on travel time is low.

Scenario 3, the scenario affecting the largest share of road kilometres, in which both low-density and high-density residential streets drop to 40 km/h, adds an average of 12 seconds per trip – far less than a single phase of a traffic light (typically 60 to 90 seconds for a full cycle). Scenario 2, which restricts the change to high-density streets only, adds an average of just 3 seconds. Scenario 1, which targets only low-density streets, sits closer to Scenario 3 because low-density streets make up the bulk of the residential network.

Impacts are consistent across destination types, with small differences driven by how much of each journey runs on low-density residential streets.

Beach and leisure destinations see the largest average change at 14 seconds under Scenario 3, because these journeys are predominantly routed through low-density residential streets, which account for the greatest proportion of the network change. City centre and university trips add 11 to 12 seconds, and park trips are the least affected at 9 seconds on average.

Suburb-level impacts are narrow.

The high-density streets scenario consistently produces the smallest impact across every suburb, typically adding 1 to 5 seconds. Under Scenario 3 (all residential streets), Coogee and South Coogee see the largest impacts at 18 seconds each, while 10 of the 14 suburbs add 15 seconds or less. The low-density and all-residential scenarios track closely across most suburbs, reflecting that low-density streets dominate the network change.

Longer trips absorb the change.

The same handful of seconds matters less the further you travel. Trips under 2 km gain only a few seconds in absolute terms but those seconds represent a larger share of the overall journey, while trips of 6 km or more pick up the largest absolute deltas yet would be less noticeable – the added time amounts to roughly 1.5% of the trip. Because Randwick's typical journey is 4 to 10 km long, most residents will experience the change as a marginal addition to a trip they already make in 8 to 17 minutes.

5. Applications and uses

The travel-time impact quantified in this report is one input into a broader decision-making framework. Combined with other datasets the analysis enables several practical applications.

Risk-prioritised intervention targeting. Overlaying the per-suburb and per-street travel-time impact with crash records and crash severity data identifies roads where the safety benefit of a 40 km/h limit is highest and the travel-time cost is lowest. Streets with frequent or severe crashes that also fall into low-impact suburbs offer the strongest "high-value" candidates for early implementation – measurable safety gain at minimal mobility cost. This is the application most likely to translate directly into a sequenced delivery program.

Equity and place-based prioritisation. Combining the suburb-level results with demographic indicators – ABS socio-economic indices, school catchment boundaries, density of vulnerable road users (children, older residents) – highlights areas where the same intervention delivers proportionally greater public benefit. This supports a place-based rather than a uniform rollout and aligns council's speed-limit policy with its broader social and equity objectives.

Sequenced rollout strategy. The asymmetry between scenarios – the high-density-only intervention adds approximately 3 seconds per trip while the all-residential intervention adds approximately 12 seconds – suggests a staged approach. Starting with the high-density streets first secures safety benefits in the densest, most pedestrian-active areas first, lowers political risk, and gives council a measured baseline of community response before considering broader implementation.

Community engagement. The LGA and suburb-level results translate an abstract policy question into specific, localised numbers that residents and stakeholders can engage with meaningfully. Presenting the simulated travel-time implications to the community concerned will help reduce the risk of debate being driven by anecdote or perception.

Extension to other precincts and LGAs. The methodology is portable. The same simulation pipeline can be re-applied at low marginal cost to other Sydney LGAs, neighbouring councils, or specific precincts within Randwick – for example, individual school catchments, town-centre clusters or arterial corridors – enabling consistent, comparable assessment across jurisdictions.

The report is intended to inform, not replace, council deliberation. The numbers it produces are most powerful when read alongside crash, demographic and engagement data, and integrated into an evidence base that supports a balanced safety-and-mobility policy framework.

Technical appendix

This appendix documents the data, sampling and modelling assumptions that underpin the Randwick speed-change impact analysis.

Data sources and trip sampling

Road network and speed limits. The road network was built from OpenStreetMap, filtered to vehicle-relevant ways (motorways, primary through tertiary roads, residential, living streets, service and unclassified roads) and to the SUMO passenger vehicle class, so footways, cycleways, rail and tram links are excluded. Posted speed limits for the base case were sourced from OSM and cross-checked against the Transport for NSW speed-limit dataset; scenarios 1 to 3 then re-tag a defined subset of these roads to 40 km/h. Junctions sitting close together are merged, and a turn-speed cap of 30 km/h is applied above a 15° turn angle so vehicles don't accelerate unrealistically through bends. Left-hand-drive rules are enforced.

The classification of streets as low-density or high-density residential was derived from NSW state land use zone data, restricted to residential zone classifications.

Points of interest. POIs were drawn from OSM and grouped into four destination types relevant to Randwick residents: beach and leisure, city and centres, parks and green space, and university and education. The model bounding box extends roughly 500 m beyond the convex hull of POIs to avoid edge effects on routes that briefly exit and re-enter the LGA.

Origin sampling. Trip origins were sampled from ABS 2021 Mesh Block population counts, rasterised at 100 m resolution, with 100 origin points drawn per suburb weighted by resident population. Each origin is paired with a destination POI and snapped to the nearest road edge within a 300 m search radius; trips are dropped if no edge is found in range.

Simulation engine. Trips were simulated in SUMO (Simulation of Urban Mobility) version 1.26, the open-source microscopic traffic simulator maintained by the German Aerospace Center (DLR) and the Eclipse Foundation. Vehicles use the Krauss car-following model with parameters chosen for Australian urban driving: 3.0 m/s² acceleration, 3.09 m/s² comfortable deceleration, 9.0 m/s² emergency deceleration, 2.5 m minimum gap between vehicles, 1.0 s reaction time, and 10% speed variation between drivers. A global speed factor of 0.811 is applied to every vehicle; this was calibrated once against Google Maps free-flow times so the base-case durations match observed off-peak conditions and is then held constant across all four scenarios so the only thing changing between runs is the posted speed limit. Traffic signals are removed from the network, so the simulation represents a free-flow proxy – there is no demand-side modelling, no queueing and no signal coordination. Trips longer than 3,600 seconds are treated as failed routes and excluded.

Validation against Google Maps. The base-case network was validated against Google Maps Distance Matrix queries for the same origin–destination pairs, using the "best guess" duration in free-flow conditions (early-morning request, no traffic layer). Calibration was performed on the base case only; the three scenarios then run against this calibrated network with no further tuning, so any difference in trip duration is attributable to the speed-limit change rather than to model fit.

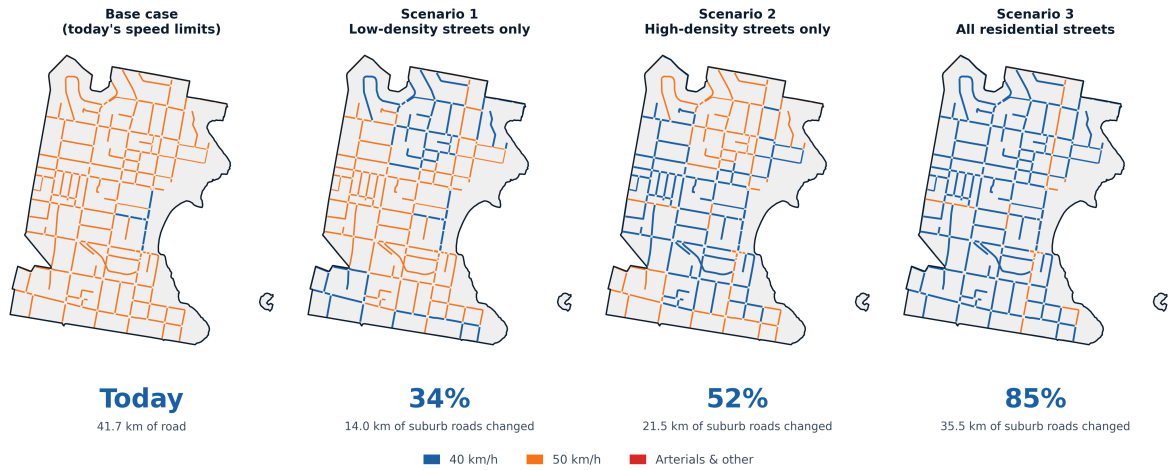
What this model does not capture. The simulation isolates one variable: the posted speed limit. It does not model rush-hour congestion, queueing at intersections, signal phasing, pedestrian crossings, parking manoeuvres, school

ones, or driver compliance. Real-world journeys are also shaped by these factors, which can dwarf the effect of a 10 km/h speed-limit change. The numbers in this report should be read as the direct mechanical effect of the speed change on a clear road – the floor of the impact, not the ceiling.

Suburb-level results

Coogee

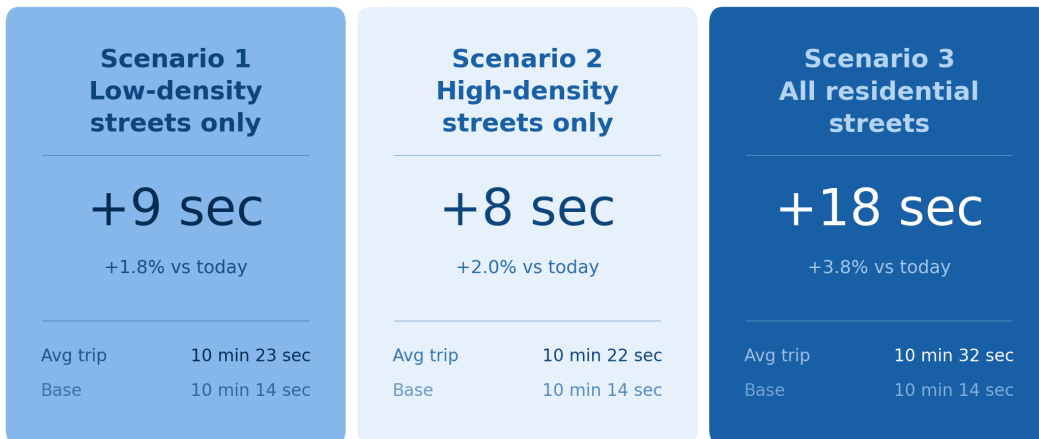
Share of Coogee road network changed by scenario



Source: SUMO scenario network clipped to the Coogee suburb polygon (41.7 km of road).

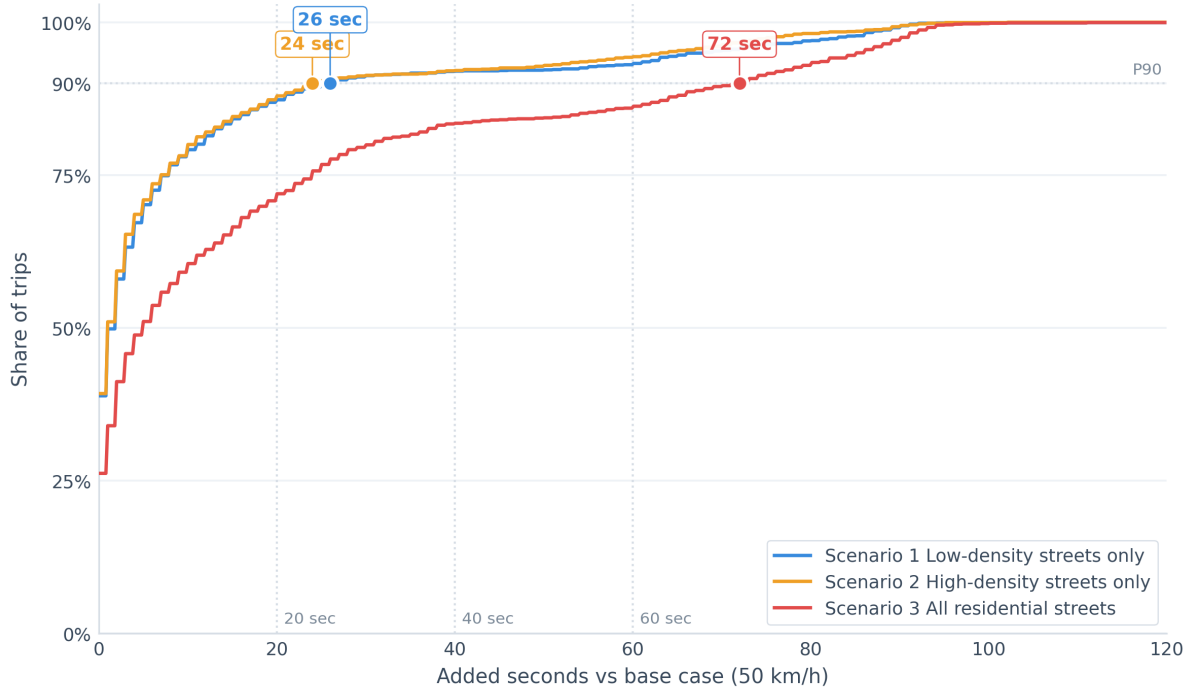
Extra time added to an average trip - Coogee

All trips originating in Coogee (n = 1,900) - compared to today's speed limits



Share of trips by added seconds - Coogee

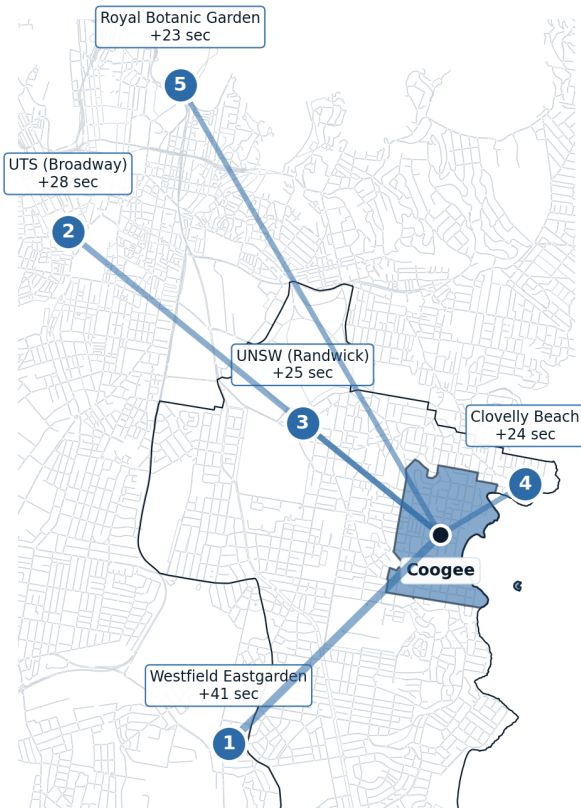
Trips originating in Coogee · trips with zero or negative delta counted at 0 · P90 = 90th percentile threshold



Source: Randwick SUMO model, n = 1,900 trips originating in Coogee. Curves clipped at 120 seconds.

Top 5 destinations affected - Coogee

Ranked by mean added seconds under Scenario 3 (all residential streets to 40 km/h)



1	Westfield Eastgarden	+41 sec
	City & centre	+6.1% vs today
2	UTS (Broadway)	+28 sec
	University	+3.5% vs today
3	UNSW (Randwick)	+25 sec
	University	+7.7% vs today
4	Clovelly Beach	+24 sec
	Beach	+7.4% vs today
5	Royal Botanic Garden	+23 sec
	Park	+2.3% vs today

Extra travel time by destination type - Coogee

Trips originating in Coogee (n = 1,900) · seconds added vs base case (50 km/h) · % change shown below

	Base case (50 km/h)	Scenario 1 Low-density streets only	Scenario 2 High-density streets only	Scenario 3 All residential streets
	10 min 14 sec	10 min 23 sec	10 min 22 sec	10 min 32 sec
Beach & leisure	9 min	+9 sec +1.6%	+4 sec +0.8%	+14 sec +2.6%
City & centres	9 min 58 sec	+14 sec +2.4%	+7 sec +1.2%	+22 sec +3.7%
Park & green space	11 min 1 sec	+7 sec +1.0%	+10 sec +1.5%	+17 sec +2.6%
University & education	12 min 9 sec	+7 sec +0.9%	+17 sec +2.3%	+23 sec +3.2%

Source: Randwick SUMO model. Cell shade scales with seconds added.

Example trip - Coogee to UNSW (Randwick)

Randwick City Council – map shows all residential streets at 40 km/h · base: 50 km/h



