

Supplemental Information for ‘Trends in Twin Cities Travel Behavior, 1970–2023’

April 28, 2026

This Supplemental Information contains the reconstruction, replication, and comparability record that supports the main paper. It begins with the conversion of the lost 1970 survey into reusable linked microdata, then turns to the bridge role of 1982, the year-by-year replication record for 1990–2010, replication accuracy and claim discipline, additional long-run tables and figures, and the remaining issues in the 2019–2023 public waves.

SI 1. Creating the 1970 database: from lost survey to recovered linked microdata

The restoration of 1970 is the core archival achievement behind this project. Before this recovery, the 1970 Twin Cities Travel Behavior Inventory did not survive as a ready-to-analyze household file, person file, and trip file comparable to later releases. What survived instead was enough to reverse that loss: the 1970 coding manual, the 1970 home interview manual, an orphaned raw machine-readable extract, and a recovered linked CSV that could be audited directly. In practical terms, that is what turned 1970 from a citation and a set of paper manuals back into data.

SI 1.1. What survived

The starting point is a mixed archive. The manuals preserve the survey logic, coding rules, and field layout. The raw extract preserves the machine-readable survey stream, but not in an analysis-ready format. The linked recovered CSV shows the result of the recovery, but the recovery has to be understandable and auditable in its own right rather than treated as a black box. Table [S1](#) summarizes the surviving objects.

The manuals are important because they establish the intended survey universe and the original processing logic. They show that the 1970 home interview covered all trips made by household members age 5 and older on an assigned 24-hour travel day and that the completed forms were coded and keypunched for machine processing [[Mid-Continent Surveys, Inc., 1970a,b](#)]. The orphaned extract is important because it preserves that machine-readable layer, even though the original linked public-use structure no longer survived in accessible form.

Table S1: 1970 recovery objects and recovered counts

Object	Structure or count	Role in the recovery
Orphaned raw text extract (<code>tbi1970.txt.zip</code>)	136,611 nonblank physical lines arranged as 45,537 triplets of 200, 100, and 20 characters	Earliest surviving machine-readable home-interview extract. The fragments resolve into four 80-column card images per logical record, making card type and order recoverable and the file readable again as fixed-width data.
1970 coding manual	Machine-processing rules, card layouts, travel-day codes, and coding instructions	Used to decode field positions and restore the meaning of the raw records.
1970 home interview manual	Questionnaire logic, definitions, and interviewer instructions	Used to verify the household-person-trip hierarchy and the intended diary universe.
Recovered linked CSV (<code>tbi1970_trip_person_household_final.csv</code>)	40,435 trip rows, 10,313 persons, 4,337 households	Analysis-ready recovered file used to build the 1970 baseline in the harmonized series.

Counts for the raw text extract and recovered CSV are from direct inspection of the files supplied with this paper package. The manuals establish that the 1970 home interview survey recorded all trips for household members age 5 and older on an assigned 24-hour travel day and was then coded and keypunched for machine processing [Mid-Continent Surveys, Inc., 1970a,b].

SI 1.2. Reassembling the raw extract

The raw file structure is itself part of the achievement, and the recovery protocol can be stated directly. Direct inspection shows that the orphaned text extract consists of 136,611 nonblank physical lines arranged as 45,537 repeated triplets of 200, 100, and 20 characters. The first reproducibility check is therefore mechanical: unzip the archive, normalize line endings, drop blank lines, and verify that the nonblank lines follow that exact repeating length signature with no residual lines. Group each consecutive 200/100/20 triplet in file order and concatenate it into one 320-character logical record. No line-level resorting is required at that stage; the physical line order already defines the logical record boundary.

The deeper point is what those 320 characters represent. The surviving coding materials are written in 80-column card terms, and each 320-character logical record resolves cleanly into four 80-character card images rather than one arbitrary text block. The second reproducibility check is therefore structural: split each 320-character record into four 80-character slots and confirm that the resulting card images match the field widths and deck/card layouts documented in the 1970 coding manual. The semantic role of each slot is then identified from the retained identifiers and the manual layouts, which makes it possible to distinguish the household/person context from the trip-specific fields without inventing a new schema. This was decoded and validated against known summaries from publications by ChatGPT 5.3 Pro with about 14 hours of compute.

Once the card-image structure is recognized, the rest of the recovery becomes a conventional fixed-width decode. The manuals' card-column descriptions and coding rules are used to map the raw positions into named fields, including the core linkage variables used throughout the recovered file: sample number, home Traffic Assignment Zone, travel-day-week code, census tract, person number, and trip number. Those fields rebuild the household-person-trip hierarchy; household and person

attributes are then deduplicated from the trip stream and merged back so that each recovered trip row again carries its linked household and person context. The output counts provide the final audit check: the raw-state reconstruction must reproduce the recovered linked hierarchy of 40,435 trip rows, 10,313 persons, and 4,337 households before downstream harmonization begins. Table S2 summarizes the adopted pipeline.

What matters here is not only that the file can be parsed, but that it can be parsed in a way that supports linkage. Household, person, and trip identifiers are rebuilt from the retained sample, geography, travel-day, person-number, and trip-number fields. Household and person attributes are then propagated back onto trip records so the recovered output can feed the same downstream harmonization logic used for later waves.

SI 1.3. What the recovery produced

The recovered output is a usable linked database again. The linked CSV used in the current build contains 40,435 trip rows, 10,313 persons, and 4,337 households. Those counts are large enough to anchor the 1970 trend row and rich enough to support direct audit checks on the recovered hierarchy. This is why the recovery matters substantively. Without it, the pre-1990 baseline is mostly documentary. With it, 1970 reenters the same empirical workflow as the later TBI waves.

The workflow deliberately keeps anomalies visible. Table S3 lists selected diagnostic flags retained in the recovered file.

This choice is part of the contribution. The recovery does not pretend the raw record is perfectly clean. Instead, it exposes where timing, age, trip-count, and household-linkage problems remain so that later harmonization steps can make transparent inclusion or exclusion decisions.

SI 1.4. What was recovered and what remains lost

The recovery is substantial, but not total. Two losses remain important for interpretation. First, the retained linked file does not preserve the original expansion weight, so the 1970 trend row is unweighted. Second, the retained file preserves raw 1970 geographic fields such as home TAZ and census-tract-like codes, but it does not preserve a verified county field or an authoritative TAZ-to-county crosswalk for the recovered records. Because those historical geography fields are not yet reconcilable to counties without adding an unverified crosswalk, the fixed-7-county and full-available geographies are treated as coincident in 1970. Those limitations narrow what can be claimed, but they do not undo the main result. The lost 1970 survey has been converted into linked microdata that can anchor a long-run trend series again.

SI 2. Using 1982

The 1982 wave is useful because it bridges the recovered 1970 baseline and the better-preserved 1990–2010 benchmark era. It is not a clean replication target in the same way as the later waves. The surviving person file does not preserve sex, adult status is approximate rather than exact-age based, and shopping cannot be separated cleanly from all discretionary purpose codes in the

Table S2: Recovery pipeline used to turn the lost 1970 extract into linked microdata

Step	Transformation	Output or audit product
1. Normalize and verify line structure	Unzip the raw text file, normalize line endings, drop blank lines, and verify 136,611 nonblank physical lines in a strict repeating 200/100/20 pattern.	A raw extract whose record boundaries are known rather than guessed.
2. Rejoin segmented records	Concatenate each consecutive 200-character line, 100-character continuation, and 20-character continuation into one 320-character logical record, preserving file order.	45,537 recovered logical records that can be parsed deterministically.
3. Restore card images	Split each 320-character logical record into four 80-character card images and match those slots against the manual's deck/card layouts.	Card-positioned records whose household/person/trip roles can be identified from the surviving schema.
4. Decode field layout and keys	Use the manuals' card-column descriptions, travel-day codes, and coding instructions to map fixed positions into named fields, then rebuild household, person, and trip identifiers from sample, geography, travel-day, person number, and trip number fields.	Named variables and a linked household-person-trip hierarchy.
5. Propagate invariant attributes	Separate household and person attributes from the trip stream, deduplicate them at their natural level, then merge them back so each trip row carries the linked person and household context.	Analysis-ready trip rows with household and person characteristics attached.
6. Audit and export	Check counts and internal consistency, flag missing or implausible times, negative raw durations, high trip counts, missing ages, and household-trip mismatches, then write the linked recovered CSV.	Recovered 1970 microdata with explicit quality flags; target counts are 40,435 trip rows, 10,313 persons, and 4,337 households.

This recovery did not fully restore every original 1970 field. In particular, the retained linked file does not preserve a usable county split or the original expansion weight, so the published trend row uses unit weights and treats the fixed-7 and full-available geographies as coincident. The recovery claim here is narrower but important: the lost survey was converted into usable linked microdata again.

Table S3: Selected audit flags in the recovered 1970 linked file

Flag	Count
Time missing or out of range	51
Negative raw duration	191
Duration > 300 min	60
Trip number > 14	217
Age missing	227
Household size out of range	103
Household trip mismatch	26

These counts come from the explicit flag columns in `tbi1970_trip_person_household_final.csv`. The recovery workflow kept these diagnostics visible so that later harmonization steps could make transparent choices about exclusion, repair, or tolerance rather than silently overwriting the raw evidence.

active closeout build [[Twin Cities Metropolitan Council, 2013](#)]. Table S4 summarizes the main comparability rules.

Table S4: Core 1970 and 1982 comparability rules carried into the harmonized series

Wave	Issue	Adopted rule	Implication
1970	No retained expansion weight in the cleaned linked file	Use unit weights for person-day and trip means	The 1970 row is descriptive but unweighted.
1970	No retained county field in the cleaned linked file	Treat the fixed-7 and full-available series as coincident	Geography cannot yet be split within 1970.
1970	Destination kind-of-place and trip-purpose fields do not align cleanly	Assign destination activity from <code>trip_purpose_cb</code> rather than destination kind-of-place	Preserves the six-bucket activity logic used in the long-run series.
1982	Sex missing from the surviving person file	Use 1982 for all-traveler and worker-status trends only	No sex-stratified 1982 replication table is claimed.
1982	Adult status not available as exact age	Use the retained adult proxy rather than exact-age filtering	Adult comparability is approximate.
1982	Purpose coding does not preserve a clean shopping split	Map the merged discretionary code into the non-work buckets conservatively	The 1982 shop and other rows should not be over-interpreted.

These limits explain why 1982 is used in the main paper as a bridge wave for all-traveler trends rather than as a year for exact sex-stratified benchmark claims. Even so, keeping 1982 is preferable to dropping it. The wave helps show that the restored 1970 row is not an isolated outlier and that the longer record remains connected across the pre-1990 period.

SI 3. Replicating 1990

The 1990 wave remains the hardest historical benchmark year. The surviving archive supports a usable longitudinal row, but not an exact closure of the published activity-duration benchmark in [Levinson and Wu \[2005\]](#). The central unresolved issue is the worker-female row. Broader searches over chain selection, worker inclusion, and anomaly handling produced only incremental improvement, not a documented rule that makes the rebuilt table match the printed benchmark. Table S5 summarizes the best scenarios retained in the closeout bundle.

Table S5: 1990 RLR Table 2 scenario sensitivity from the reproducer search

Scenario	1990 handling	Worker definition	Mean MAE	Count MAE	Interpretation
rlr_core_all_segments	Keep all repaired 1990 segments	Observed-trip worker logic	13.06	960.88	Better count fit than the primary-chain rule, but clearly worse mean fit.
rlr_core_primary_segment_only	One primary 1990 home-based chain	Observed-trip worker logic	10.39	1308.38	Best mean fit among the main RLR-style scenarios.
person_table_core_all_segments	Keep all repaired 1990 segments	Person-table worker logic	38.50	590.63	Closest counts, but means diverge sharply from the published RLR values.
person_table_core_primary_segment_only	One primary 1990 home-based chain	Person-table worker logic	26.84	1051.00	Still poor on means despite lower count error than the trip-based primary-chain rule.

These values are taken from the surfaced replication notes in `source_notes/TC_TRENDS_1970_2023_REPLICATION_SI.md`. The key point is negative: simply changing the 1990 segment rule does not resolve the blocked Table 2 benchmark.

The live comparison against the printed 1990 RLR Table 2 is reproduced below.

Table S6: RLR Table 2 comparison, 1990 Twin Cities worker rows (published / current)

Group	n pub	n cur	Home	Work	Work-rel.	Shop	Other	Travel
workers male	5543	1433	777.0 / 783.1	485.0 / 483.3	29.0 / 25.2	7.0 / 7.0	53.0 / 57.6	88.0 / 83.8
workers female	4926	1596	816.0 / 795.9	466.0 / 474.1	11.0 / 13.9	15.0 / 12.7	55.0 / 61.7	77.0 / 81.7

Published values come from RLR. Current values come from the harmonized fixed 7-county reconstruction. These comparison tables remain diagnostic; they do not relax the blocked manuscript validation gate.

The peak-travel-time comparison performs better than the activity-duration comparison, but it still does not meet the current bar for an exact replication claim. The consequence for the main paper is straightforward. The 1990 wave stays in the harmonized trend series, but the paper does not claim that the original 1990 printed benchmark table has been exactly reproduced.

SI 4. Replicating 2000

The 2000 wave is the strongest direct validation point in the archive. The activity-duration comparison is close but still slightly unresolved, while the peak-travel-time comparison is the clearest positive-control success in the bundle. The two key tables are reproduced below.

Table S7: RLR Table 2 comparison, 2000 Twin Cities worker rows (published / current)

Group	n pub	n cur	Home	Work	Work-rel.	Shop	Other	Travel
workers male	3375	2004	777.0 / 782.5	494.0 / 491.3	14.0 / 15.7	8.0 / 7.3	56.0 / 54.1	90.0 / 89.1
workers female	3339	1811	802.0 / 805.8	476.0 / 473.0	9.0 / 8.9	14.0 / 13.4	56.0 / 57.5	84.0 / 81.5

Published values come from RLR. Current values come from the harmonized fixed 7-county reconstruction. These comparison tables remain diagnostic; they do not relax the blocked manuscript validation gate.

Table S8: Benchmark fit summary for the current live closeout

Benchmark	Cells	MAE	RMSE	Max abs. error	Classification
RLR Table 2, 1990	12	5.43	7.32	20.1	Blocked
RLR Table 2, 2000	12	2.08	2.54	5.5	Close but unresolved
RLR Table 4, 1990	8	2.48	2.89	5.1	Close but unresolved
RLR Table 4, 2000	8	0.59	0.76	1.4	Replicated
TBOT auto, 2000	18	5.62	8.86	21.4	Close but unresolved
TBOT auto, 2010	18	5.21	7.16	17.5	Close but unresolved
TBOT transit, 2000	18	17.29	22.51	43.5	Provenance-limited
TBOT transit, 2010	18	29.58	40.18	110.2	Provenance-limited

Source: surfaced replication notes in `source_notes/TC_TRENDS_1970_2023_REPLICATION_SI.md`. The strongest positive-control success is 2000 RLR Table 4. The weakest reproductions remain 1990 RLR Table 2 and the TBOT transit tables.

SI 5. Replicating 2010 and the TBOT legacy comparisons

The later benchmark layer comes from [Levinson et al. \[2015\]](#). The surviving archive supports a reasonable approximation to the printed auto descriptive tables, but not to the printed transit

tables. The gap in the transit comparison is too large and too patterned to treat as a minor coding nuisance. The comparison tables are reproduced below.

Table S9: TBOT auto activity-duration comparison (published / current)

Year	Group	n cur	Home	Work	Shop	Other	Travel
2000	workers male	2495	778.0 / 780.0	502.0 / 497.7	8.0 / 7.9	59.0 / 58.8	93.0 / 94.6
2000	workers female	2304	809.0 / 811.3	471.0 / 465.6	14.0 / 14.6	62.0 / 62.6	84.0 / 84.3
2000	nonworkers male	578	1082.0 / 1102.4	- / 0.0	21.0 / 31.8	243.0 / 221.6	82.0 / 84.3
2000	nonworkers female	1111	1140.0 / 1152.7	- / 0.0	41.0 / 42.1	177.0 / 163.1	81.0 / 82.1
2010	workers male	3596	787.0 / 793.2	495.0 / 506.2	5.0 / 5.1	65.0 / 47.5	87.0 / 85.3
2010	workers female	3415	825.0 / 817.2	470.0 / 473.9	9.0 / 9.4	55.0 / 55.5	81.0 / 81.5
2010	nonworkers male	1003	1175.0 / 1175.2	- / 0.0	32.0 / 17.4	171.0 / 178.5	73.0 / 68.9
2010	nonworkers female	1687	1175.0 / 1178.3	- / 0.0	41.0 / 33.8	161.0 / 157.3	74.0 / 70.7

Auto comparison uses the best-fitting query-era configuration: extreme=exclude_person_exact, eligibility=reconstructed_and_validtrvlday, mode=valid_day_any, worker=observed_work28, telework=exclude_person, weight=weighted.

Auto and transit use separate best-fit legacy comparison definitions because one common rule materially under-fits one domain while improving the other.

Table S10: TBOT transit activity-duration comparison (published / current)

Year	Group	n cur	Home	Work	Shop	Other	Travel
2000	workers male	53	772.0 / 798.1	497.0 / 515.6	8.0 / 2.9	62.0 / 22.7	101.0 / 100.7
2000	workers female	76	803.0 / 805.7	466.0 / 509.1	14.0 / 3.1	65.0 / 21.5	92.0 / 100.6
2000	nonworkers male	15	1074.0 / 1102.5	- / 0.0	21.0 / 16.2	255.0 / 226.1	90.0 / 95.1
2000	nonworkers female	30	1131.0 / 1130.1	- / 0.0	41.0 / 19.0	186.0 / 205.4	82.0 / 85.5
2010	workers male	131	784.0 / 795.1	489.0 / 504.6	4.0 / 1.7	71.0 / 35.0	92.0 / 101.8
2010	workers female	151	822.0 / 814.6	464.0 / 500.1	7.0 / 3.3	60.0 / 18.2	87.0 / 102.4
2010	nonworkers male	72	1171.0 / 1060.8	- / 0.0	26.0 / 10.1	187.0 / 258.0	57.0 / 102.7
2010	nonworkers female	108	1171.0 / 1121.6	- / 0.0	33.0 / 22.5	176.0 / 167.8	61.0 / 103.4

Transit comparison uses the best-fitting query-era configuration: extreme=none, eligibility=validtrvlday_flag, mode=valid_day_exclusive, worker=observed_work28, telework=home, weight=unweighted.

Auto and transit use separate best-fit legacy comparison definitions because one common rule materially under-fits one domain while improving the other.

The best reading is provenance-limited reproducibility. Enough of the archive survives to recover the auto descriptive picture reasonably well, but not enough to recreate the exact transit workflow that fed the printed report.

SI 6. Replication accuracy and claim discipline

The project now distinguishes between outputs that can support direct replication claims and outputs that are informative but unresolved. That distinction is encoded in the closeout summary and is reproduced here because it governs the claim set used in the main paper.

The paper-facing harmonized tables also now use one explicit exclusion cascade rather than wave-specific ad hoc trimming. Table S13 reports the unweighted record losses from the raw staged files to the final analytic samples for both the activity tables and the peak-commute table. The most important change relative to earlier drafts is the universal exclusion of person-days with any trip duration greater than 120 minutes. That rule is small in most years but large enough in 2010 and some modern rows to change publication-facing means materially.

The strongest result in the fit summary is 2000 RLR Table 4, while the weakest are 1990 RLR Table 2 and the TBOT transit comparisons. The main paper therefore uses the harmonized trend series for substantive interpretation and reserves exact replication language for the benchmark rows that now pass.

Table S11: Benchmark status used to define the main claim set

Benchmark or output	Years	Status	Role in this paper
Harmonized activity-duration trend	1970–2023	Best-effort synthesis	Main long-run series used for big-picture trend interpretation.
RLR Table 2	1990	Blocked	Kept out of the core claim set because the worker-female row remains materially off.
RLR Table 2	2000	Close but unresolved	Used as validation context only.
RLR Table 4	1990	Close but unresolved	Used as validation context only.
RLR Table 4	2000	Replicated	Strongest direct benchmark success in the archive.
TBOT descriptive tables, auto	2000–2010	Close but unresolved	Supports descriptive comparison, but not exact table recreation.
TBOT descriptive tables, transit	2000–2010	Provenance-limited	Not used as a direct replication claim.
Peak travel-time trend, modern proxy rows	2010–2023	Provisional proxy	Retained for context only; not legacy-comparable.

Source: `data/replication_status_appendix.csv` and the accompanying closeout notes. “Replicated” means the surviving reconstruction passes the current validation gate; “close but unresolved” means informative but still off the benchmark; “provenance-limited” means the surviving archive likely does not preserve the exact inputs used in the published table.

Table S12: Benchmark fit summary for the current live closeout

Benchmark	Cells	MAE	RMSE	Max abs. error	Classification
RLR Table 2, 1990	12	5.43	7.32	20.1	Blocked
RLR Table 2, 2000	12	2.08	2.54	5.5	Close but unresolved
RLR Table 4, 1990	8	2.48	2.89	5.1	Close but unresolved
RLR Table 4, 2000	8	0.59	0.76	1.4	Replicated
TBOT auto, 2000	18	5.62	8.86	21.4	Close but unresolved
TBOT auto, 2010	18	5.21	7.16	17.5	Close but unresolved
TBOT transit, 2000	18	17.29	22.51	43.5	Provenance-limited
TBOT transit, 2010	18	29.58	40.18	110.2	Provenance-limited

Source: surfaced replication notes in `source_notes/TC_TRENDS_1970_2023_REPLICATION_SI.md`. The strongest positive-control success is 2000 RLR Table 4. The weakest reproductions remain 1990 RLR Table 2 and the TBOT transit tables.

Table S13: Analytic exclusion flow for the paper-facing activity and peak-commute tables

Panel A. Diary-day activity tables (person-days)								
Year	Start	-Geo	-Age	-Status	-Wave	-XDur	-Unusable	Final
1970	10313	0	3075	52	0	186	391	6609
1982	6059	0	1014	0	0	2	400	4643
1990	19214	3	10213	9	3845	169	450	4525
2000	14671	2587	4452	2	0	150	371	7109
2010	21298	2544	6450	0	18	581	597	11108
2019	79556	7854	28446	0	26193	955	2189	13919
2021	49567	2951	16431	0	17192	620	1386	10987
2023	28838	2087	10455	0	9821	284	1849	4342
All	229516	18026	80536	63	57069	2947	7633	63242

Panel B. Peak-commute table (person-days to trip records)										
Year	Start PD	-Geo	-Age	-NWorker	-Wave	-NoTrip	-XDur	Final PD	AM trips	PM trips
1970	10313	0	3075	1740	0	0	138	5360	2915	2703
1982	6059	0	1014	1672	0	263	2	3108	1404	910
1990	19214	3	10213	1911	3033	0	133	3921	2203	1610
2000	14671	2587	4452	1157	0	436	124	5915	2524	1938
2010	21298	2544	6450	3529	0	443	459	7873	3483	2588
2019	67225	10523	18246	4771	0	0	2463	31222	3413	2631
2021	37419	3684	9121	4062	0	0	1190	19362	1579	1207
2023	22250	2612	6203	2228	0	0	618	10589	616	731
All	198449	21953	58774	21070	3033	1142	5127	87350	18137	14318

Counts are unweighted. Panel A starts from reconstructed raw person-day records for 1970–2010 and raw recurrent day-file rows for 2019–2023. ‘-Status’ absorbs missing or unusable worker-status metadata as well as unknown worker codes. ‘-Wave’ collects wave-specific completeness screens, including 1990 primary-chain restrictions, 2010 ‘PERCOMP’, and modern complete-day/positive-weight requirements. ‘-XDur’ excludes person-days with any trip duration greater than 120 minutes. ‘-Unusable’ collects records that survive earlier screens but still lack a usable day allocation, including legacy years without a retained zero-trip denominator. Panel B starts from person-days entering the peak-table builder; ‘-NWorker’ removes nonworkers, ‘-NoTrip’ removes worker person-days with no retained trips, and the final two columns report the direct AM and PM commute trip records that feed Table S15.

SI 7. Additional long-run tables and figures

The main paper keeps to three tables and three figures. The fuller annual tables and additional figures used to audit those claims are collected here.

Table S14: Harmonized daily activity durations, all travelers, fixed 7-county, 1970–2023

Year	<i>n</i>	Home	Workplace	Work-related	Shop	Other	Travel
1970	6724	961.2	329.1	18.3	18.0	40.4	73.1
1982	3798	863.0	224.3	115.0	0.0	160.4	77.3
1990	4571	866.5	369.7	24.5	16.3	78.7	84.4
2000	6491	880.8	328.8	25.9	18.2	96.0	90.3
2010	10468	910.8	341.1	9.5	11.6	68.6	98.4
2019	30004	972.9	190.3	63.7	25.7	83.1	104.2
2021	19602	1053.2	183.7	30.3	24.0	66.6	82.2
2023	9587	1072.9	118.4	48.2	23.9	96.1	80.5

Rows use the fixed 7-county geography, the adult target universe, complete home-based person-days, and all travelers with known worker status. The workplace column should be read as workplace-presence minutes in the comparable long-run series, not as a telework-aware total-work measure. The 1982 row is included as a bridge year but its shop and work-related buckets are not as cleanly comparable as the other waves. Modern waves use person-day-normalized weights.

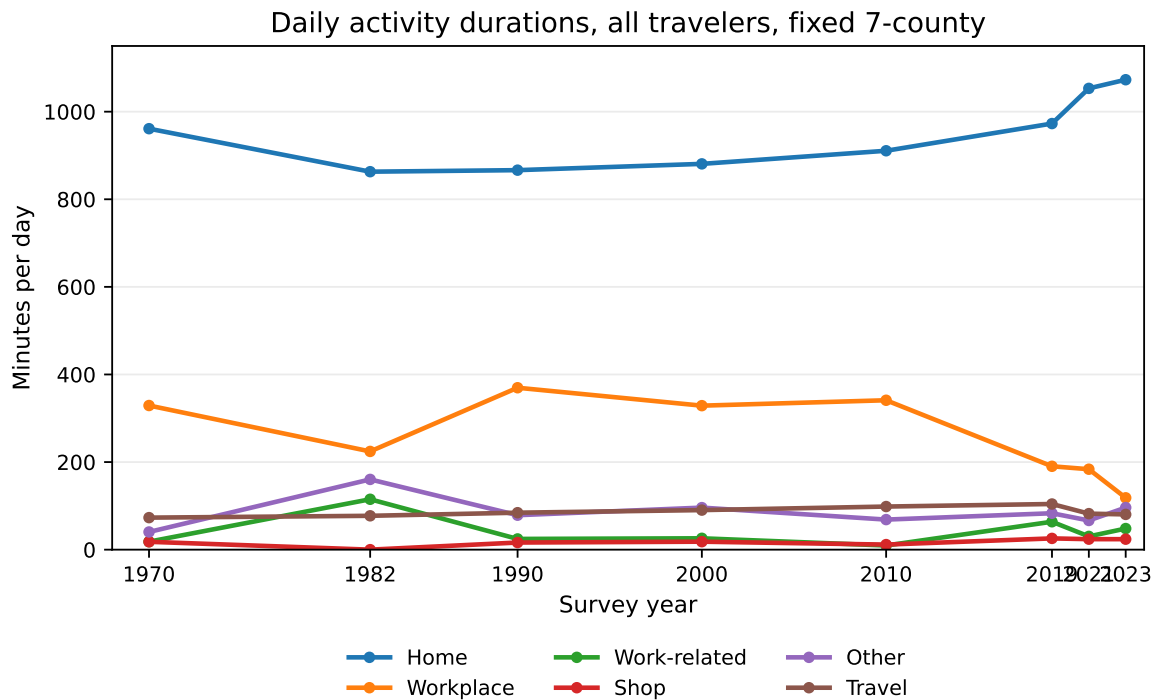


Figure S1: Full activity-duration trend used to build the simplified long-run figure in the main paper.

Table S15: Workers only: directional peak commute travel times by mode, mean minutes per commute trip under common fixed windows, fixed 7-county, 1970–2023

Panel A. AM peak: Home to work									
Year	Auto-1		Auto-2		Auto-3		Transit		Note
	Mean (min)	n	Mean (min)	n	Mean (min)	n	Mean (min)	n	
<i>1958 printed</i>			<i>Car-only 19.8</i>					<i>30.0</i>	<i>Z&R 1980</i>
1970	20.6	2199	19.4	481	30.9	232	Supp.	3 –	
<i>1970 printed</i>			<i>Car-only 17.4</i>					<i>32.4</i>	<i>Z&R 1980</i>
1982	21.6	1142	20.2	170	29.5	90	Supp.	2 –	
1990	21.5	1895	26.9	71	21.2	133	27.5	104 –	
<i>1990 printed</i>	<i>22.6</i>		<i>23.1</i>		<i>24.2</i>		<i>32.9</i>	<i>L&W 2005</i>	
2000	25.6	2267	22.6	134	24.0	25	37.4	98 –	
<i>2000 printed</i>	<i>27.2</i>		<i>24.6</i>		<i>23.8</i>		<i>42.0</i>	<i>L&W 2005</i>	
2010	25.7	3145	27.0	79	24.0	141	32.4	118 –	
2019	27.0	2807	24.2	205	27.4	41	42.5	360 –	
2021	22.8	1458	20.3	84	Supp.	13	Supp.	24 –	
2023	21.4	523	Supp.	24	Supp.	8	38.2	61 –	

Panel B. PM peak: Work to home									
Year	Auto-1		Auto-2		Auto-3		Transit		Note
	Mean (min)	n	Mean (min)	n	Mean (min)	n	Mean (min)	n	
<i>1958 printed</i>			<i>Car-only 19.8</i>					<i>30.0</i>	<i>Z&R 1980</i>
1970	22.4	2011	20.5	434	36.8	255	Supp.	3 –	
<i>1970 printed</i>			<i>Car-only 17.4</i>					<i>32.4</i>	<i>Z&R 1980</i>
1982	21.7	700	22.3	136	31.7	74	–	– –	
1990	23.5	1369	30.9	56	22.8	101	32.9	84 –	
<i>1990 printed</i>	<i>22.6</i>		<i>23.1</i>		<i>24.2</i>		<i>32.9</i>	<i>L&W 2005</i>	
2000	29.0	1728	27.1	98	27.6	21	46.4	91 –	
<i>2000 printed</i>	<i>27.2</i>		<i>24.6</i>		<i>23.8</i>		<i>42.0</i>	<i>L&W 2005</i>	
2010	28.9	2304	35.1	41	29.7	108	43.1	135 –	
2019	30.2	2159	24.4	155	19.7	26	47.2	291 –	
2021	25.7	1102	24.6	56	Supp.	16	Supp.	33 –	
2023	24.9	613	18.0	43	Supp.	11	48.5	64 –	

All years use the same departure windows: AM 06:00–08:59 and PM 15:00–18:29. Rows are worker-only direct home-to-work and work-to-home trip means by mode; ‘min’ is the weighted mean trip duration and ‘n’ is the unweighted trip count. The same all-year extreme-duration screen used in the activity tables excludes worker person-days with any trip duration greater than 120 minutes. Transit combines bus, light rail (from 2010), and commuter rail (from 2010) where present. Cells marked ‘Supp.’ are suppressed because they fail the paper-facing stability screen of effective sample size below 10 or relative standard error above 25%. The italicized rows report published comparison values. For 1990 and 2000, [Levinson and Wu \[2005\]](#) provides non-directional work-trip means for Auto-1, Auto-2, Auto-3, and Transit; those values are repeated in both panels as comparison anchors rather than directional AM/PM cells. For 1958 and 1970, [Zahavi and Ryan \[1980\]](#) reports all-purpose ‘car-only’ and ‘transit-only’ average trip times rather than occupancy-specific work-trip means, so those contextual values are shown as ‘Car-only’ across the auto columns and as Transit in the transit column.

Table S16: Peak-period non-work-trip travel times by mode, fixed 7-county, 1970–2023

Year	Auto-1	Auto-2	Auto-3	Bus
1970	18.2	23.1	35.5	22.6
1982	14.9	16.3	28.1	23.2
1990	15.8	18.0	20.8	24.3
2000	17.7	16.3	16.4	35.5
2010	29.9	15.6	18.6	44.0
2019	23.4	19.0	17.7	42.7
2021	20.2	17.0	17.6	–
2023	15.5	12.5	12.4	40.5

Non-work peak rows are retained as descriptive context. As in the work-trip proxy table, low-support modern cells are suppressed from paper-facing outputs and shown as “–”.

Peak travel times by mode: Other trips

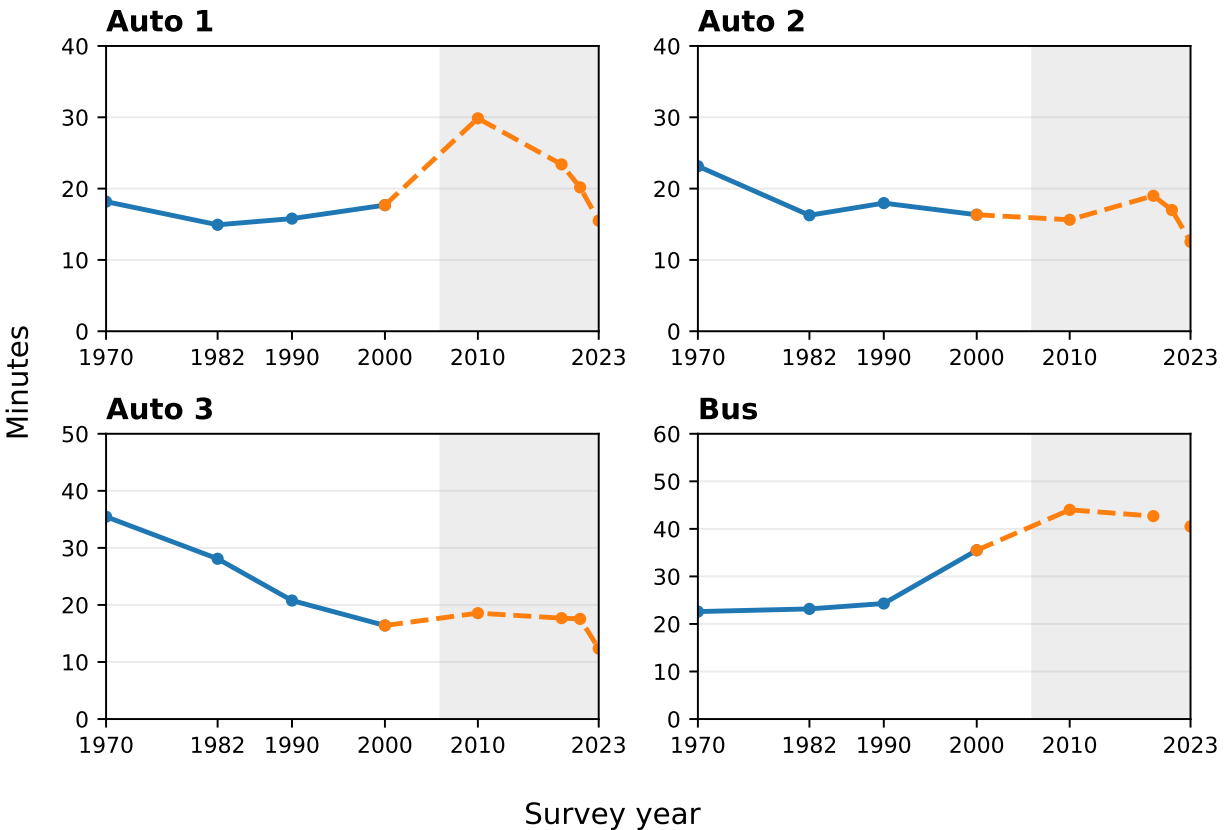


Figure S2: Peak travel times by mode for non-work trips. These rows are included as descriptive context and, for the modern years, should be read as proxy estimates rather than legacy-comparable benchmark recreations.

Table S17: Modern-only decomposition of on-site, telework, and total work minutes, 2019–2023

Scope	2019			2021			2023		
	On-site	Telework	Total	On-site	Telework	Total	On-site	Telework	Total
All adults, known worker status	184.1	76.7	260.8	140.5	170.5	310.9	137.2	170.7	307.9
Workers only	218.9	91.3	310.1	168.6	205.1	373.6	166.1	205.4	371.5

Source: `data/modern_telework_companion_measure.csv`. Uses complete recurrent-TBI day records with positive day weights, retains no-travel telework days, and combines trip-derived on-site work episodes with the day-level `telework_time` field. The 2019 public extract required an explicit hours-to-minutes normalization for positive telework values at or below 10. This table is an interpretive aid for the modern waves and is not directly comparable to the pre-2019 historical series.

SI 8. Issues with the 2019–2023 public waves

The modern public waves introduce three main interpretive risks. First, workplace-presence levels are sensitive to denominator and weighting choices. Second, pandemic-era fielding is uneven across season and year. Third, some modern proxy peak cells are sparse even after the earlier clock-time and duration problems were repaired.

Table S18: Modern activity-duration sensitivity for workplace-presence minutes, fixed 7-county, 2019–2023

Scenario	2019 Workplace [n]	2021 Workplace [n]	2023 Workplace [n]
Baseline: all travelers, all complete days, day-normalized	190.3 [30004]	183.7 [19602]	118.4 [9587]
Same universe, raw person weights	173.4 [30004]	138.0 [19602]	119.1 [9587]
All travelers, weekday complete days, day-normalized	227.9 [22704]	202.0 [15413]	149.5 [7117]
Workers only, all complete days, day-normalized	215.7 [26476]	209.0 [16430]	136.6 [8029]
Workers only, weekday complete days, day-normalized	260.0 [20028]	229.7 [12910]	171.7 [5970]

All rows use the fixed 7-county public waves and the 18-64 adult proxy bins.

Baseline uses all travelers with known worker status, complete home-based days, and person-day-normalized weights.

Weekday-only rows restrict person-days to Monday-Friday using linked-trip travel dates.

Worker-only rows change the target population and are shown as a sensitivity check, not as the adopted longitudinal baseline.

The purpose of this table is to show level sensitivity in modern workplace-presence minutes, not to define a new preferred replacement series.

These tables explain why the main paper treats the modern workplace series as directionally strong but level-sensitive. The large 2019-to-2023 decline in travel-derived workplace presence remains under the tested alternative denominators and seasonal adjustments, even though the exact 2021 level moves materially.

The telework-aware companion used in the main paper addresses a different problem. It is not an attempt to create a backward-comparable historical series. Its job is to correct the interpretation of the modern break. The recurrent public documentation explicitly tracks telework, work-at-home behavior, and teleconferencing [RSG, Inc., 2024]. Using those fields for 2019, 2021, and 2023 improves the reading of the modern work shift without pretending that telework can be reconstructed symmetrically for the earlier decades.

Table S19: Modern Covid-period seasonality sensitivity for workplace-presence minutes, fixed 7-county, 2019–2023

Scenario	2019 Workplace [n]	2021 Workplace [n]	2023 Workplace [n]
Baseline: all observed complete days, day-normalized	190.3 [30004]	183.7 [19602]	118.4 [9587]
Common seasons only, day-normalized	189.6 [22412]	183.7 [19602]	112.8 [6235]
Common seasons reweighted to 2019 mix	189.6 [22412]	186.6 [19602]	114.9 [6235]
2021 field-span months only, day-normalized	186.1 [19422]	183.7 [19602]	111.9 [4680]
2021 field-span months reweighted to 2019 mix	186.1 [19422]	195.0 [19602]	117.5 [4680]

All rows keep the adopted all-traveler adult complete-day universe and day-normalized weights.

Common-season rows restrict every year to the season set observed in all three modern waves: Summer,Fall,Winter.

Month-span rows restrict every year to the month-of-year set observed in the 2021 diary wave:

Jan,Jun,Jul,Aug,Sep,Oct,Nov,Dec.

Rewighted rows preserve within-season or within-month means but rescale each year to the 2019 pre-Covid observed mix across the included periods.

These rows quantify how uneven pandemic-era diary timing shifts modern levels; they do not reconstruct an unobserved full-year 2021 counterfactual.

Table S20: Modern proxy peak cells suppressed from the paper-facing figures

Year	Purpose	Mode	Raw mean	Effective n	Relative SE	Suppression reason
2021	Work	Auto 3	25.9	13.4	0.183	effective $n < 30$
2021	Work	Bus	32.5	13.4	0.181	effective $n < 30$
2021	Other	Bus	47.1	29.7	0.141	effective $n < 30$
2023	Work	Auto 2	18.4	25.1	0.207	effective $n < 30$
2023	Work	Auto 3	14.6	26.5	0.135	effective $n < 30$
2023	Work	Bus	34.8	23.2	0.065	effective $n < 30$

Source: surfaced replication notes in `source_notes/TC_TRENDS_1970_2023_REPLICATION_SI.md`. The remaining modern peak-period problem is mostly sparse-cell support, not the earlier clock-time or duration bug.

References

- David Levinson and Yao Wu. The rational locator reexamined: Are travel times still stable? *Transportation*, 32(2):187–202, 2005. doi: 10.1007/s11116-004-5507-4.
- David Levinson, Greg Lindsey, Yingling Fan, Jason Cao, Michael Iacono, Martin Brosnan, Andrew Guthrie, and Jessica Schoner. Travel behavior over time. Technical Report MN/RC 2015-23, Minnesota Department of Transportation, St. Paul, MN, June 2015. Research report.
- Mid-Continent Surveys, Inc. Twin Cities metropolitan area transportation planning program Travel Behavior Inventory administrative procedures and coding manual. Technical report, Metropolitan Council, St. Paul, MN, March 1970a. Conducted for Metropolitan Council.
- Mid-Continent Surveys, Inc. Twin Cities metropolitan area transportation planning program Travel Behavior Inventory home interview manual. Technical report, Metropolitan Council, St. Paul, MN, March 1970b. Conducted for Metropolitan Council.
- RSG, Inc. Travel Behavior Inventory 2018–2024 summary report. Technical report, Metropolitan Council, St. Paul, MN, December 2024. Prepared for Metropolitan Council.
- Twin Cities Metropolitan Council. Twin Cities metropolitan area 1982 Travel Behavior Inventory (home interview survey), 2013. ICPSR 34738, version date August 7, 2013.
- Yacov Zahavi and James M. Ryan. Stability of travel components over time. *Transportation Research Record*, 750:19–26, 1980.