

The Delay Is Seated: Cabin Egress Is a Makespan, Not a Queue

Ellis R. Marchbank

Institute of Transport Operations, University of Cleeve, e.marchbank@cleeve.ac.uk

Received 3 June 2026; revised 21 August 2026; accepted 2 September 2026

Abstract

The aeroplane stops and everyone stands up, and then nobody moves for eleven minutes. We show that those minutes are not caused by the people who stood up; they are caused, to the second, by the people who did not.

The standard account treats a cabin as a queue served at a fixed rate by its door, so that promptness cannot matter and haste is merely undignified. That account requires the aisle to be congested. Using ceiling-mounted depth sensing across 163,412 reconstructed egress trajectories, we find that it is not: aisle occupancy peaks at 0.41 of jam density and approaches saturation on no sector we observed. Transit is therefore free-flow, and total egress time is not a service time but a *makespan* — the last departure — which is set by whichever passenger stands up last, and by nobody else. Ninety-six per cent of a passenger's post-landing wait elapses at their own seat, before they enter the aisle at all.

The prediction is unusually sharp, and we tested it. In a pre-registered three-arm trial of 612 sectors, a confederate rising promptly from the rearmost occupied row shortened cabin egress by 187 s, while an identical confederate in a middle row changed nothing at all (95% CI −11 to +19 s). Nobody is waiting for the aisle. The aisle is waiting for them.

Keywords: cabin egress; makespan scheduling; coordination equilibria; free-flow pedestrian dynamics; heuristic inquiry

1 Introduction

An aircraft arrives at its stand and the seatbelt sign is extinguished. Within thirty seconds, by our measurements, a median of 71% of passengers are out of their seats. A median of 9% have entered the aisle. The rest stand where they are, bent beneath the overhead lockers, while the cabin takes a further nine to twelve minutes to empty.

The behaviour is universally observed and almost as universally deprecated. A substantial popular literature treats the immediate riser as impatient, and the passenger who remains seated with a book until the aisle has cleared as having correctly understood the situation. The reasoning behind that

verdict is straightforward and, we will argue, wrong at its first step. It holds that a cabin empties at a rate fixed by its door; that the number of passengers ahead of you is therefore the only quantity that matters; and that standing merely relocates a wait without shortening it.

That argument is valid. Its premise is not. A fixed egress rate is a property of a *congested* channel, and we find no evidence that the aisle of a single-aisle aircraft is ever congested. Across 1,144 instrumented sectors, peak linear occupancy of the aisle averages 0.41 of jam density and reaches 0.55 on no sector, at no moment, on no airframe. There is no queue in the technical sense. There is a corridor with people standing still in it, which is a different object and obeys different arithmetic.

If transit is free-flowing, then a passenger's departure time is the moment they enter the aisle, plus a short and nearly constant walk. Total egress time is then not a sum but a maximum: a *makespan*. And a makespan has a property that a service rate does not. It is set by exactly one element of the set, and is perfectly insensitive to every other. Entering the aisle early cannot help anybody — except the one passenger for whom it is decisive, and that passenger is whoever enters it last.

The moral consequence is immediate, and we state it without enthusiasm. The passenger who enters the aisle at wheels-stop has, at that instant, discharged their entire causal contribution to the delay. The passenger who waits for the aisle to clear has not been spared the delay. They have produced the whole of their own, and, if they are the last to rise, the cabin's as well.

Contributions

1. We report the first continuous occupancy measurements of the ordinary (non-emergency) cabin aisle, and find no saturation (§5.1).
2. We derive the *Makespan Identity* for unsaturated egress and its corollary that $\partial M / \partial \tau_i = 0$ for all but one passenger (§3).
3. We decompose passenger wait and find that 96% of it elapses at the passenger's own seat (§5.3).

4. We test the identity’s distinctive prediction — that intervening on a non-terminal passenger does exactly nothing — in a pre-registered three-arm trial (§5.4).
5. We give a coordination account of why the slow equilibrium is stable, turning on a signal that is systematically read backwards (§6).

2 Related Work

The literature on aircraft cabin egress is large, rigorous, and almost entirely concerned with the emergency case. Certification requires full evacuation within 90 seconds (Pellham, 2011), and the resulting body of work models a cabin under precisely the conditions in which the aisle is saturated: every passenger attempting to move at once, jam density reached within seconds, and throughput correctly described by a flow–density relation (Orsini et al., 2014; Dahlbeck, 2019). Under those conditions a service-rate model is appropriate, and the field’s confidence in it is well earned.

The ordinary case has inherited that model without inheriting its conditions. Boarding has been studied intensively — back-to-front, outside-in, and random boarding strategies are compared in Van der Ploeg and Ferrante (2016) and Ashworth and Prieto (2021) — but boarding is genuinely congested, because passengers stow luggage while standing in the aisle and block it. Disembarkation reverses the geometry: luggage is retrieved from a position the passenger already occupies, and the retrieving passenger does not obstruct anyone behind them who has not yet stood up. To our knowledge no published study has measured aisle density during ordinary disembarkation. The assumption of congestion appears to have been carried across from the boarding and evacuation literatures by analogy, and never checked.

Makespan objectives are standard in scheduling (Lindqvist, 2009) and the distinction between a sum objective and a maximum objective is elementary there. Its application to pedestrian egress appears to be new, and we suspect the reason is simply that nobody expected the non-saturation condition to hold.

Finally, our instrument requires comment. The hypothesis of this paper did not arise from the cohort; it arose from one observer’s sustained attention to their own experience, formalised afterwards. That is a recognised methodology (Moustakas, 1990), with recognised limits, and we treat it in §4.1 as hypothesis-generating only.

3 A Model of Cabin Egress

3.1 Notation

Let wheels-stop define $t = 0$. Let the aircraft carry n seated passengers, indexed $i = 1, \dots, n$, with passenger i seated in row $r_i \in \{1, \dots, R\}$ numbered from the door. The cabin door opens at $t = D$.




Figure 1: Geometry and notation for single-aisle egress. The passenger’s post-landing wait is dominated by τ_i , an interval spent in the seat and not previously counted as part of egress at all.

Definition 1 (Conversion time). The *conversion time* τ_i is the instant at which passenger i leaves their seat and occupies the aisle.

Definition 2 (Transit time). The *transit time* w_i is the interval from τ_i until passenger i crosses the door plane, comprising overhead-locker retrieval and the walk forward.

Standing is not conversion. A passenger who rises at their seat but does not step into the aisle has not converted, and §5.1 shows that for most passengers the two events are separated by several minutes.

Passenger i reaches the door plane at $\tau_i + w_i$ but cannot cross before the door opens, so departs at $d_i = \max(D, \tau_i + w_i)$, and the cabin is empty at

$$M = \max_{1 \leq i \leq n} d_i = \max_{1 \leq i \leq n} (\tau_i + w_i), \quad (1)$$

which we call the *egress makespan*. We further define the *seated residual* $R_i = \tau_i$: the portion of passenger i ’s post-landing wait spent at their own seat, whether sitting or standing over it.

3.2 The saturated case, for contrast

If the aisle is at jam density then transit is throughput-limited. Writing μ for the door-plane service rate, the classical account gives $M \approx D + n/\mu$: egress time scales with passenger count, and $\partial M / \partial \tau_i = 0$ for every i , because the queue is never starved. This is the model under which standing early is futile, and it is the correct model for an evacuation.

3.3 The unsaturated case

Let $\phi(t) \in [0, 1]$ denote instantaneous aisle occupancy as a fraction of jam density. We say egress is *unsaturated* on a sector if $\sup_t \phi(t) < \phi^*$, the critical density at which the flow–density relation departs from free flow. Under free flow, transit time is a function of distance alone:

$$w_i = w_0 + \beta r_i + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2), \quad (2)$$

with w_0 the locker-retrieval and door-crossing overhead — which varies between sectors with cabin-bag load, but not between passengers within a sector — and β the per-row walking cost. Note what (2) does not contain: any term in n , in ϕ , or in the conversion times of other passengers. This is the whole of the paper’s mathematical content, and everything below follows from it.

Theorem 3 (Makespan Identity). *If egress is unsaturated, then*

$$M = \max_i (\tau_i + w_0 + \beta r_i + \varepsilon_i) = \max(D, \tau_{i^*} + w_{i^*}), \quad (3)$$

where $i^* = \arg \max_i (\tau_i + w_i)$.

Proof. Immediate from (1) and (2). Under free flow no passenger’s transit time is a function of any other’s, so each d_i is determined by passenger i alone and no interaction term survives the maximisation. The door constraint enters only as a floor, since no passenger may cross before D . $\square$

Corollary 4 (Insensitivity). *For all $i \neq i^*$, $\partial M / \partial \tau_i = 0$, and $\partial M / \partial \tau_{i^*} = 1$.*

Corollary 4 is the paper’s operative result and deserves plain statement. Under an unsaturated aisle, the time at which the cabin empties is a function of one passenger’s decision and of no other passenger’s decision. Delaying your own departure by five minutes leaves every other passenger’s departure time exactly unchanged. It moves M only if you are the last to rise, and then by the full five minutes — a cost borne not by the other passengers, who have already gone, but by the aircraft, the crew and the turnaround.

Three qualifications. The derivatives are defined away from ties, and $\partial M / \partial \tau_{i^*} = 1$ only where $\tau_{i^*} + w_{i^*} > D$; below that the door floor binds and the derivative is zero for everyone. They are also partial, holding all other conversion times fixed, and within a seat row that is not innocuous: a passenger at the window cannot convert before the passenger in the aisle seat beside them has. We therefore treat the seat row, not the individual, as the unit of intervention in §5.4. Finally, Theorem 3 is conditional: it says nothing about real cabins unless the non-saturation antecedent holds empirically. Testing that antecedent is the subject of §5.1, and we take it to be the paper’s principal empirical burden.

4 Methods

Evidence is drawn from four sources of descending subjectivity and ascending sample size. We describe them in that order because that is the order in which they were obtained, and because the first does not test anything.

4.1 The instrument

The hypothesis originates in sustained first-person observation, conducted and analysed according to the heuristic-inquiry protocol of Moustakas (1990): initial engagement,

immersion, incubation, illumination, explication, and creative synthesis. The author recorded 1,187 sectors between 2006 and 2025, logging seat assignment, wheels-stop, personal conversion time, personal door-crossing time, and a hand count of standing passengers at 30-second marks.

Moustakas requires that the illumination be reported rather than reconstructed, and we do so. On 14 March 2019, on a remote stand at Gatwick, the author had been seated for nine minutes and was looking forward along the aisle. The aisle was not full. It was empty, and stationary. The eighteen passengers standing in it were spaced roughly a row and a half apart, and none of them was moving, because the passengers ahead of them had not yet stood up. The delay was not being caused by the people in the aisle; it was being caused by the people beside them.

We are explicit about what this tier is and is not. It is one unblinded observer with a hypothesis, and its internal validity is accordingly poor. It generated the question. It cannot answer it, and we make no inferential use of it below.

4.2 The instrumented cohort

Ceiling-mounted time-of-flight depth sensors (Radcliffe and Stoyanova, 2022) were installed on 214 airframes across four European carriers (three narrowbody types and one twin-aisle type), sampling the full aisle length at 4 Hz. Over 1,144 sectors between March 2023 and January 2026 we reconstructed 163,412 individual egress trajectories, yielding τ_i , w_i , r_i and the continuous occupancy trace $\phi(t)$. Sensors recorded continuously; trajectory reconstruction is computationally expensive, and the 1,144 analysed sectors were drawn at random from the recorded set, stratified by airframe type, carrier and time of day.

Trajectory reconstruction under partial occlusion is the principal methodological risk, and we treat it at length. Depth returns from the aisle are ambiguous when two passengers occupy adjacent rows; we resolved these by a constrained assignment step that forbids trajectory crossing, which is physically justified in a 51 cm aisle. A validation subset of 4,100 trajectories was independently hand-annotated from cabin video by two coders blind to the study hypothesis. Agreement with the automated pipeline was high ($\kappa = 0.91$); conversion times agreed to within 1.4 s (RMSE), and transit times to within 2.9 s. The residual disagreement is concentrated in the rearmost four rows, where sensor grazing angle is worst, and we flag below where this could matter.

4.3 The confederate trial

Theorem 3 makes a prediction that no congestion model makes: the effect of a prompt riser depends entirely on which row they occupy. Under a strict service-rate model the effect is zero from every row, since $M \approx D + n/\mu$ and no single conversion time enters it (§3.2); under any partially saturated variant an early riser reduces start-up and tail starvation by some small amount, but does so equally from any

row, because the binding constraint is the door and not the seat. Every congestion account we are aware of therefore predicts that two interventions differing only in seat row will coincide, whatever their common value. The makespan account predicts that they differ by the whole of the effect, and specifies in advance which way. The trial is designed around that contrast rather than around either arm alone.

We pre-registered a three-arm parallel-group design across 612 sectors, randomised at the sector level, which is also the level at which the outcome is measured:

- **Arm A** ($n = 204$): a confederate in the *aisle seat* of the rearmost occupied row rises and enters the aisle within 5 s of wheels-stop.
- **Arm B** ($n = 203$): an identical confederate, identically instructed, in the aisle seat of a row drawn uniformly from the middle third of the cabin.
- **Arm C** ($n = 205$): no confederate.

The primary outcome was M , measured as the interval from wheels-stop to the final door-plane crossing. Confederates carried a single cabin bag, wore ordinary clothing, and were instructed not to speak to other passengers or to hurry. Cabin crew were not informed of arm assignment. Trial sectors were flown on instrumented airframes, so M was recorded by the pipeline of §4.2 and not by the confederate. The trial was powered at 0.90 to detect a 60 s difference between any two arms at $\alpha = 0.05$, which required 196 sectors per arm. Analysis followed the pre-registered plan; no outcome was added after unblinding.

4.4 Retrospective turnaround records

For external validation we obtained on-blocks and door-clear timestamps, decomposed following Sørhaug (2017), for 2.4 million sectors from the same four carriers (2016–2025). These support no causal claim and are used only to characterise the distribution of M and to test whether it varies with load factor.

4.5 Ethics and registration

The protocol was approved by the University of Cleeve Research Ethics Committee (TO/2024/119) and by the safety departments of all four participating carriers. The trial was pre-registered on 4 November 2024, before any confederate sector was flown; the registration fixed the primary outcome, the three arms, the allocation ratio, and the between-arm contrast reported in §5.4. Depth returns were reduced to anonymous point clouds on the aircraft and no imagery left it; the validation subset was hand-annotated from cabin video retained under the carriers' existing security-recording consent and destroyed after coding. Individual passenger consent was not sought, and we accept that this is a substantive limitation: the intervention was non-interactive, involved no deception

Table 1: Evidence tiers.

Tier	Unit	Sectors	Inferential role
Instrument	observer-sector	1,187	hypothesis generation
Cohort	trajectory	1,144	measures ϕ , τ , w
Trial	sector	612	causal test of Cor. 4
Retrospective	sector	2.4 M	distribution of M

Periods: 2006–2025, 2023–2026, 2025, 2016–2025 respectively. Only the trial supports a causal claim.

of any individual, and posed no more than the risk of an ordinary sector, but passengers were nonetheless studied without their knowledge.

5 Results

5.1 The aisle does not saturate

Across all 1,144 instrumented sectors, the mean of the per-sector peak aisle occupancies was $\phi_{\max} = 0.41$ (SD 0.07). The largest single-sector value observed was 0.55, on a full A321 with a rear-door restriction. The critical density for a corridor of this width is $\phi^* = 0.72$ (Dahlbeck, 2019); no sector at any moment came within 0.17 of it, and the flow-density relation showed no departure from free flow anywhere in the observed range (Fig. 2). The door opened at a mean $D = 95$ s (SD 24).

This is a strong result and we want to be careful about what it means. It does not mean the aisle is empty; the aisle typically contains 15 to 20 standing passengers, against a jam capacity of approximately 49 for a 26-row cabin. It means those passengers stand a row and a half or more apart and are therefore not impeding one another. Density is low because passengers enter the aisle slowly, not because few of them enter it: a median 71% of passengers are out of their seats within thirty seconds of wheels-stop, and the median passenger then stands over their own seat for a further 6.8 minutes before converting. The distinction is invisible from inside the cabin, and we return to it in §6.

5.2 The identity holds

With non-saturation established, we tested (3) directly. Regressing observed M on the predicted value $\max(D, \tau_{i^*} + w_{i^*})$ gives a slope of 0.994 (SE 0.006) and $R^2 = 0.987$; residual RMSE is 8.4 s against a mean makespan of 611 s. M was taken from the door-plane crossing record and the predictor from the reconstructed trajectories, so the two are not identically constructed; they are not independent either, and we return to this in §7.1. The terminal passenger i^* was the last to convert on 94% of sectors; on the remainder the last converter sat sufficiently far forward that a slightly earlier riser from the rear attained the maximum instead.

The last departure is not marginally last. Across the cohort the interval between the largest and the second-largest departure had a mean of 181 s (IQR 124–236), which is the




Figure 2: Aisle occupancy from wheels-stop to cabin-empty, pooled over 1,144 sectors (mean, with inter-quartile band). The trace does not approach the free-flow limit, still less jam density. The aisle contains 15–20 standing passengers throughout, against a jam capacity of about 49; they are too far apart to impede one another. The marked value is the mean of the per-sector peaks; the pooled trace peaks slightly below it, as the peaks do not coincide in time.




Figure 3: Egress on one representative sector (A320, 26 rows, 156 seated), one passenger drawn from each row. Twenty-five of the twenty-six bars could be moved anywhere in the window without changing the outcome. The cabin empties when the black bar ends.

quantity any intervention that displaces the terminal passenger can recover, and no more.

Passenger count contributed nothing detectable. In the retrospective record, load factor explains a partial R^2 of 0.03 of the variance in M after conditioning on aircraft type — a fuller aircraft is very slightly slower to empty. In the instrumented cohort, where M is decomposed against the identity rather than regressed on covariates, the corresponding share is 0.003; the two are not comparable, and the smaller is the better estimate because the larger absorbs whatever else covaries with load factor in the retrospective record. Either way the effect is two orders of magnitude smaller than that of a single passenger’s conversion time (Table 2).

5.3 The seated residual

Decomposing each passenger’s post-landing wait $T_i = \tau_i + w_i$ across the cohort, the pre-conversion component accounts for $\sum_i R_i / \sum_i T_i = 0.96$: mean transit time was 17.8 s (SD 5.2)

Table 2: Variance decomposition of the egress makespan.

Component	Share of $\text{Var}(M)$	95% CI
Terminal conversion τ_i^*	0.881	0.862–0.899
Door timing D	0.048	0.039–0.058
Row cost βr_i^*	0.031	0.024–0.039
Transit overhead w_0	0.018	0.013–0.024
Load factor	0.003	0.001–0.006
Unexplained residual	0.019	—

Narrowbody sectors only ($n = 1,056$); the 88 widebody sectors are excluded and discussed in §7.2. Shares are conditioned on aircraft type.

against a mean post-landing wait of 445 s. We report the aggregate rather than the mean of per-passenger ratios, which is heavily left-skewed by the minority who convert immediately; that mean is 0.91, its median 0.97, and its lower decile 0.72.

A passenger who experiences an eleven-minute disembarkation has therefore spent roughly ten and a half minutes of it at their own row and under half a minute walking, and — since the median passenger is out of their seat within thirty seconds and converts 6.8 minutes later — has spent most of those ten and a half minutes standing over the seat rather than sitting in it. This is not a marginal effect requiring statistical care; it is the overwhelming bulk of the phenomenon, and it has gone unmeasured because the interval before a passenger enters the aisle was not previously regarded as part of the egress process at all.

5.4 The confederate trial

Results conformed to prediction (Fig. 4). The control-arm mean makespan was 611 s. Relative to control, Arm A — the rearmost-row confederate — reduced the egress makespan by 187 s (95% CI –214 to –160; $p < 0.001$), a reduction of roughly 31%. Arm B — the middle-row confederate — produced an estimated effect of +4 s (95% CI –11 to +19; $p = 0.61$).

The mechanism deserves statement, since a single prompt riser cannot shorten a makespan merely by leaving early. A confederate in an aisle seat gates the two passengers inboard of them, who cannot convert until the seat is vacated; in the rearmost occupied row those passengers are, in control sectors, among the last in the cabin. Arm A therefore removes the largest departure from the maximum, and the makespan falls to the second largest. The estimate of 187 s is an estimate of that gap, and agrees with the 181 s measured directly in the cohort. Arm B performs the identical operation on a row whose departures were never attaining the maximum.

We draw attention to Arm B, because a null result is ordinarily weak evidence and here it is not. The interventions in Arms A and B were identical in every respect except seat row, and the pre-specified contrast between them was –191 s (95% CI –222 to –160; $p < 0.001$). No congestion account distinguishes the two, because none of them makes the effect




Figure 4: The confederate trial. The two interventions were identical in every respect except seat row. A congestion account requires the two to coincide; the makespan account requires them to differ by the whole of the effect, and says in advance which way.

of a prompt riser depend on where that riser was sitting; the makespan account holds that the seat row is the only thing that determines it. The confidence interval on Arm B excludes any effect larger than 19 s, against a predicted effect of 0.

6 The Courtesy Equilibrium

Corollary 4 raises an obvious question. If rising promptly is costless to others and occasionally decisive, and if the aisle is demonstrably free-flowing, why do most passengers remain in place, and why has the folk model survived contact with several billion sector-passengers of daily experience?

We model the decision as a coordination game. Passenger i chooses τ_i against a cost c of standing (discomfort, restricted headroom) and a belief $\hat{\phi}_i$ about aisle congestion. If $\hat{\phi}_i$ is high, waiting is a best response: entering a congested aisle relocates the wait without shortening it. If $\hat{\phi}_i$ is low, converting strictly dominates waiting for any passenger who would otherwise convert after $D - w_i$. Below that threshold the door floor binds, and converting is costly without being useful; we make no claim for it. Throughout we take a passenger to have converted *promptly* if $\tau_i \leq 30$ s. Two equilibria follow, each strict in the belief that sustains it:

- the *prompt equilibrium*, $\tau_i \approx 0$ for all i : the aisle is fed rapidly, moves continuously, and M is small;
- the *courtesy equilibrium*, τ widely staggered: the aisle is fed slowly, stands still, and M is large.

Equilibrium selection turns on what a seated passenger can actually observe. They cannot observe ϕ . They observe *motion*, and what they see is a line of standing people who are not moving. The inference from stillness to congestion is natural, and it is exactly backwards. A stationary aisle is not a full aisle. It is an under-fed one.

The signal is therefore self-confirming. Believing the aisle congested, the passenger remains seated; remaining seated, they withhold the conversion that would have fed the aisle; the aisle remains stationary; and the stillness they observe next is in part their own doing. The courtesy equilibrium is stable not despite being false but because the evidence against it is produced by belief in it (Whitlock, 2020) (Fig. 5).




Figure 5: Best-response dynamics for the conversion decision. Both fixed points are strict equilibria; observed cabins sit in the lower one. Its stability does not require passengers to be mistaken about their own interests, only about what a stationary aisle indicates.

7 Discussion

7.1 What we regard as established, and how firmly

The non-saturation finding (§5.1) is secure. It is a direct measurement with a hand-validated pipeline, it is consistent across all four airframe types and four carriers, and it does not depend on the model. The Makespan Identity is a mathematical consequence of that finding and holds conditionally on it; the empirical fit ($R^2 = 0.987$) is strong but is not independent evidence, since both sides of the regression are built from the same trajectories.

The trial (§5.4) is the load-bearing causal evidence; its Arm A effect is well established and its Arm B null informative rather than merely absent. We are more cautious about the equilibrium account of §6. It is consistent with our data and it explains the persistence of the folk model, but we did not measure $\hat{\phi}$ and we did not attempt to. It should be read as a proposed mechanism, not a demonstrated one.

7.2 An unresolved discrepancy

The identity fails on twin-aisle aircraft disembarking through two doors. Across 88 widebody sectors the observed makespan exceeded prediction by a mean of 210 s (SD 54; range 118–317), and the sign of the residual was stable across every airframe and carrier in that subset. This is the opposite of what the model implies: two doors halve the maximum walking distance and should make the identity tighter, not looser. We have examined the obvious explanations — door-timing asymmetry, crew procedures, sensor coverage of the second aisle — and none accounts for a residual of this size. We report it because it is large, and we do not have an explanation for it.

7.3 Limitations

Our cohort is European short-haul and may not generalise; boarding conventions and cabin-bag norms vary consider-

ably by region, and cabin-bag volume enters w_0 directly. The depth-sensing pipeline degrades in the rearmost rows, which is precisely where the terminal passenger is most often seated, and although the hand-validated subset showed no bias there, its coverage of those rows is thin. Confederates could not be blinded to their own arm assignment. The retrospective records measure door-clear, not final crossing, and the two differ by a small and probably non-constant offset. Finally, we did not randomise conversion times directly, which would be the cleanest possible test and which we could not obtain permission to attempt. Sectors are nested within airframes and carriers, and the intervals in Table 2 do not account for that clustering; they are therefore narrower than they should be. And the confederate was visible to the cabin as well as to their own row, so we cannot separate the gating effect of §5.4 from any effect of example, though the Arm B null bounds the latter at under 19 s.

7.4 Practical implication

We are reluctant to convert this into advice, since the finding concerns attribution rather than optimisation, and since the saving accrues to the turnaround and to the last passengers off, and not to the passenger who acts. We note only that the current distribution of disapproval is inverted with respect to the causal structure, and that the rearmost occupied row is the only part of the cabin from which individual promptness affects anything beyond the passenger's own departure.

8 Conclusion

A cabin does not empty at a rate. It empties at a moment, and the moment belongs to one person. Under the free-flow conditions we measure, the passenger who enters the aisle at wheels-stop has completed their contribution and can add nothing further to it, while the passenger who remains seated until the aisle clears is not avoiding a delay but producing one — their own in full, and the cabin's if they are last to rise.

Further work should measure $\hat{\phi}$ directly — seat-back eye-tracking during disembarkation would establish whether passengers are in fact reading stillness as density — and should extend the instrumented cohort to twin-aisle operations, where our identity currently fails and where we expect the explanation to bear on the conditions under which it holds at all.

Acknowledgements

The author thanks the cabin crews of four carriers for tolerating the sensor installation, and two anonymous coders for 4,100 hand-annotated trajectories.

References

- Ashworth, K. and Prieto, M. (2021). Outside-in boarding under realistic cabin-bag distributions. *Transportation Science*, 55(3):712–729.
- Dahlbeck, N. (2019). Flow–density relations for constrained pedestrian corridors below 60 cm width. *Journal of Pedestrian Dynamics*, 12(1):44–61.
- Lindqvist, P. (2009). *Scheduling with Release Dates: Sum and Maximum Objectives*. Corveth Academic Press, 2nd edition.
- Moustakas, C. (1990). *Heuristic Research: Design, Methodology, and Applications*. Sage Publications, Newbury Park, CA.
- Orsini, L., Bekker, H. and Nwachukwu, T. (2014). Ninety seconds: a review of certification evacuation modelling. *Aviation Safety Review*, 31(2):101–138.
- Pellham, R. (2011). Certification demonstration and its discontents. *Transport Regulation Quarterly*, 19(4):388–410.
- Radcliffe, J. and Stoyanova, I. (2022). Depth-sensing reconstruction of dense indoor pedestrian trajectories. *Sensors and Built Environments*, 8(2):210–233.
- Sørhaug, B. (2017). Turnaround time decomposition for European short-haul operations. *Journal of Air Transport Management*, 61:55–68.
- Van der Ploeg, S. and Ferrante, D. (2016). A comparison of boarding strategies under stochastic stowage times. *European Journal of Operational Research*, 249(1):233–245.
- Whitlock, A. (2020). Coordination failure and self-confirming beliefs in transient groups. *Games and Economic Behavior*, 122:88–114.