KORDIS RESEARCH · WORKING PAPER SERIES · KRS-2026-03

Notification Lead Time and the Recoverability of Unfilled Shifts

A race-condition model of same-day coverage, and the structural ceiling on fill rate

Patrick Underwood

Kordis Research, Toronto, Ontario

15 July 2026

This paper is issued in the Kordis Research working paper series. It has not been peer reviewed and should be cited as a working paper. Kordis develops operations software for temporary staffing agencies and therefore has a commercial interest in several of the questions examined here; the disclosure at the end of the paper states that interest in full, together with the source verification protocol. No external funding was received and no client data was used. Correspondence: research at kordisai.com.

Abstract

Fill rate is the headline operating metric of shift-based labour supply and it is reported unconditionally, which makes it close to incomparable across operations. We ask what the joint distribution of absence notification lead time and worker response latency implies about the attainable ceiling on same-day coverage, and under what conditions pool depth fails to substitute for lead time. The review returns a clean negative result: we could find no published distribution of absence notification lead time in any source class, and four retrieved sources return explicit negative findings on that variable. The reason is structural, in that the Canadian and United States absence measures are weekly-reference-period stocks carrying no notification timestamp. We then formulate shift recovery as a race between two distributions, using standard results on point-process superposition and parallel-system reliability that the operations literature has not assembled for this problem. Lead time enters the fill probability twice, through the response distribution and through the number of workers reachable inside the window, while pool depth enters once and is bounded above by a correlated component. Under a central parameterisation the probability of filling a shift at thirty minutes' notice is under three per cent whether the eligible pool holds five workers or a hundred and twenty.

Keywords: fill rate; absenteeism; notification lead time; response latency; shift scheduling; reserve staffing; queueing with deadlines; predictable scheduling; temporary help services; Canada

JEL classification: C41, J22, L84, M11, M54

1 Introduction

Two staffing operations report the same fill rate. The first is told about its vacancies the evening before; the second is told forty minutes before the shift starts. Nothing in the published metric distinguishes them, and this paper argues that almost everything that matters does.

The claim is not that lead time is one factor among many. It is that lead time is the variable through which the others operate, that its effect is strongly non-linear, and that in the short-notice regime where staffing agencies actually earn their margin, the size of the available workforce is close to irrelevant. That is a strong claim, so the paper states it as a model with every parameter exposed, and reports honestly that the data needed to test it does not exist.

Research question. In shift-based labour supply, what does the joint distribution of absence notification lead time and worker response latency imply about the structural ceiling on same-day fill rate? Under what conditions does adding pool depth fail to substitute for lead time?

2 Evidence base and review protocol

This paper is a structured evidence review. It is not a systematic review in the PRISMA sense, and it is not a statistical meta-analysis: we do not pool effect sizes into a summary estimate, because the studies bearing on the question measure different constructs on different populations and a pooled estimate would be uninterpretable. Where several studies report an estimate of the same parameter we tabulate them side by side with the number of independent samples and the total participants, and we let the reader see the dispersion rather than average it away. The protocol is set out here so that the review can be audited and repeated.

2.1 Scope and inclusion

The scope is the occupational health literature on absence, the operations research literature on staffing under absence uncertainty, the survey methodology literature on contact and response, the sociological literature on schedule instability, applied probability texts covering point processes and system reliability, official Canadian, United States and United Kingdom absence statistics, and published fill-rate reporting from public-sector and commercial sources. A source was included if it bore directly on one of the constructs in the research question, or if it established the absence of evidence on one of them. Recency was not an inclusion criterion: the foundational results in several of these literatures are decades old and remain the best available, and where a newer estimate supersedes an older one we report both and say which is current.

2.2 *Search*

The search was iterative rather than a single indexed query. It proceeded from the constructs in the research question outwards, through citation chains, through the reference lists of the meta-analyses and reviews retrieved, and through the publication lists of the statistical agencies and regulators with jurisdiction over the setting. Because the search was iterative, we do not report a records-screened count; a screening count implies a single reproducible query, and reporting one here would overstate the formality of the method. This is the principal respect in which the review falls short of a systematic review, and it is stated rather than concealed. Searching concluded in July 2026.

2.3 *Verification*

Every source cited was retrieved and read at the address given in the reference list. No figure, effect size, quotation, statutory provision or interval appears anywhere in this paper unless it was read in a retrieved document. Where a publisher platform refused access, we substituted an institutional repository, an author-hosted manuscript, or a bibliographic index record, and where only an abstract could be reached we report the design without reporting quantities that the abstract does not carry. Sources we expected to exist and could not verify were omitted rather than cited from memory, and the substantive omissions are named in the limitations section. Several sources are cited specifically for what they do *not* contain; in each such case the document was retrieved and the absence confirmed rather than assumed.

2.4 *Classification*

Sources are classified by publication type on retrieval, and the class travels with the source: it is stated in the note to any figure or table that draws on it, and the full assignment is tabulated in the appendix. The distinction that matters most here is between evidence produced by parties with a commercial interest in the answer and evidence produced by parties without one. Of the 42 sources in this paper, 26 are peer-reviewed or peer-reviewed-adjacent, 7 are from statistical agencies, government departments or regulators, and 9 are from industry bodies, research centres, or vendor and practitioner publications.

Table 1. Composition of the evidence base, by source class.

Source class	Sources	Share (%)
Peer-reviewed	23	55
Monograph or courseware	3	7
National statistical agency	5	12
Government department or agency	2	5
Research centre	6	14
Industry body	2	5
Author's own research programme	1	2
Total	**42**	**100**

Notes: Classification is by publication type, assigned on retrieval and not revised afterwards. "Peer-reviewed" covers journal articles and peer-reviewed volume chapters; working papers, dissertations and open courseware are counted separately because they carry a different evidential weight. "Regulator or statute" covers a regulator's own published guidance and the text of a regulation. "Vendor or practitioner publication" covers material published by firms selling into the market under study, and by practitioner magazines; every figure taken from that class is identified as such at the point of use, in the figure or table note as well as in this appendix.

2.5 Negative findings

This review reports six findings of the form "this has not been measured." Such a claim is weaker than a claim about the world and we mark the difference throughout: what we assert is that we searched a stated body of literature by a stated method and did not find the result, not that the result does not exist. Where a retrieved document was examined specifically to confirm that it does not contain a quantity, that is recorded as a verified negative and the document is cited for the absence.

3 The variable nobody records

We looked for a published distribution of absence notification lead time from six directions: occupational health, human resources and absence management, nurse rostering, transit and airline reserve operations research, workforce-analytics vendors, and national statistical agencies. There is none. This is not a gap in the search. Four of the sources we verified return explicit negative findings on exactly this variable, and one of them explains why.

The Canadian work absence measure is constructed from the Labour Force Survey. Its incidence measure is "the percentage of full-time employees reporting some absence in the reference week," its inactivity rate is "hours lost as a proportion of the usual weekly hours of all full-time employees," and days lost per worker is "calculated by multiplying the inactivity rate by the estimated number of working days in the year" (Statistics Canada n.d.). The instrument is a weekly-reference-period stock. It has no notification timestamp, and the documentation states that it cannot distinguish avoidable from unavoidable absence. The US Current Population Survey measure has the same structure (BLS n.d.).

Table 2. What each available data source can and cannot tell us about notification timing. Four sources returned explicit negative findings on precisely the variable this paper is about.

Source	What it measures	Why it cannot answer the question
Official statistics		
Canadian Labour Force Survey work absence series	Hours lost as a proportion of usual weekly hours, converted to days lost per worker per year	The instrument is a weekly-reference-period stock with no notification timestamp, and it explicitly cannot distinguish avoidable from unavoidable absence (Statistics Canada n.d.).
US Current Population Survey absence table	Daily absence rate by industry and occupation	Same structural limitation: a rate, not an event stream (BLS n.d.).
Canadian night work and shift work releases	Incidence of regular night shifts by industry and demographic group	Contains no statistics on advance notice of work schedules (Statistics Canada n.d.).

Source	What it measures	Why it cannot answer the question
Employer and operational data		
Large multi-country employer absence survey, 1,280 respondents	Cost of absence, replacement-worker productivity by absence type	Contains no statistics on how much advance notice employers receive of absences (SHRM and Kronos).
Transit extraboard study, 14,305 garage-day-period observations over four years	Absence counts per garage-day-period and the reserve team size needed to cover them	Does not provide statistics on when absences are reported or notified relative to shift start times (Diab et al.).
Airline reserve crew planning	Reserve crew scheduling and the cost of uncovered trips	Does not state when absences become known relative to departure, nor describe the call-out process (Sohoni et al.).
Fill-rate reporting		
Public-sector nurse staffing reports	Fill rate as an hours ratio: actual staffed hours over planned hours, aggregated monthly	Not a per-request success rate and not conditioned on notice, so it cannot be inverted into a probability (NHS trust board paper).
Staffing industry fill rate	A vendor key performance indicator	Every result our searches returned was a software vendor or agency publication with no stated methodology and no defined denominator. We cite none of them.

Notes: Each cell in the third column that reports an absence is a verified negative finding, meaning we retrieved the document and looked. The pattern is consistent and it has a structural cause: absence is measured as a proportion of time lost, because that is what payroll and household surveys record. Nobody records the moment the vacancy became known, so nobody can condition on it.

The same absence runs through the research literature. The field's canonical twenty-year review organises absence research explicitly by time frame, and in doing so identifies short-horizon absence as the least understood window, calling for "more explicit consideration of time frames, causal lags, and aggregation periods" (Harrison and Martocchio 1998). Nearly three decades later, the systematic reviews of absence predictors remain organised around psychosocial and health antecedents rather than around notification (Margheritti et al. 2025).

Fill rate has a mirror-image problem. It is published in two incompatible forms and neither is what the question requires. Public-sector nurse staffing reports give it as an hours ratio, actual staffed hours over planned staffed hours, aggregated monthly, with no per-request denominator and no conditioning on notice (Shrewsbury and Telford Hospital NHS Trust 2025). Staffing industry fill rate is a vendor key performance indicator; every result our searches returned was a software vendor or agency publication with no stated methodology and no defined denominator, and we cite none of them.

The gap, stated precisely

There is no peer-reviewed or government empirical fill rate conditioned on notification lead time, in any shift-based sector, including on-demand staffing platforms. There is no published distribution of notification lead time. And no study decomposes fill-rate variance into notification timing, pool depth, and effort. The consequence is that the industry's central operating metric is reported without the one covariate that plausibly dominates it.

4 What is known: absence rates and their composition

What we do know is how much absence there is, and roughly what it consists of. Canadian full-time employees lost 9.3 days per year on average, of which 7.7 were attributed to illness or disability and 1.6 to personal or family responsibilities, with transportation and warehousing among the highest-absence industries at 12.3 days (Statistics Canada 2012). The series continues to 2025 as table 14-10-0191-01 (Statistics Canada n.d.). On the daily-rate measure, the US absence rate was 3.2 per cent overall, 2.9 per cent in manufacturing and 2.9 per cent in transportation and utilities (BLS n.d.). A large United Kingdom employer survey covering 918 organisations and over 6.5 million employees reported 7.8 days per employee and 3.4 per cent of working time lost, described as the highest level in over a decade and about two days above the pre-pandemic figure of 5.8 days (CIPD 2023).

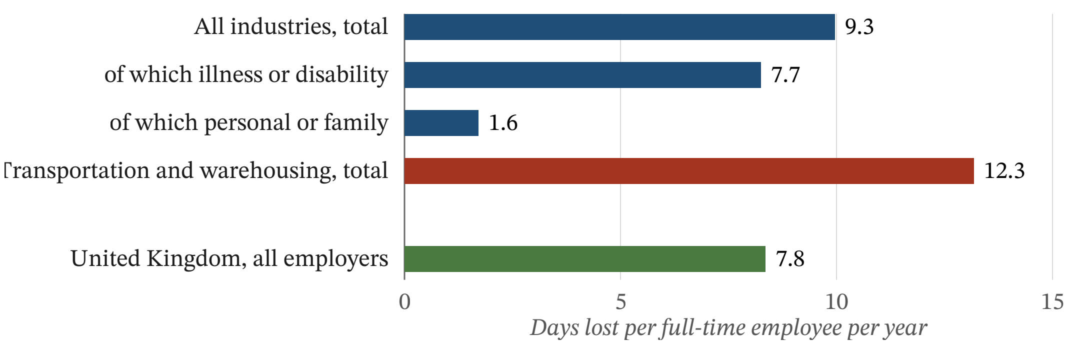

Figure 1. How much absence there is. Days lost per full-time employee per year, which is the form in which absence is officially published in Canada. Transportation and warehousing, the sector adjacent to light industrial staffing, is among the highest. The lower row is a comparable United Kingdom employer survey figure, included because it is the only source we verified that also reports what short-term absence consists of.

Source: Canadian figures: Statistics Canada, *Work Absence Rates 2011* (Catalogue 71-211-X). National statistical agency. United Kingdom figure: Chartered Institute of Personnel and Development with Simplyhealth (2023), *Health and Wellbeing at Work*, 918 organisations covering over 6.5 million employees. Industry body survey. The Canadian series continues as table 14-10-0191-01 through 2025, but its values do not render to an automated reader and we did not retrieve them; see section 11.

Two features of the composition matter for the model. The first is that short-term absence is dominated by conditions recognised hours rather than weeks in advance: minor illness, colds and influenza, stomach upsets and headaches, was in the top three causes of short-term absence for 94 per cent of respondents in the United Kingdom survey (CIPD 2023). This is a mechanism claim about lead time, and it is the closest thing to direct evidence on the notification distribution that we could verify: the physiology of the dominant cause implies that most notice arrives on a horizon of hours.

The second is that absence is partially predictable at the person level, which is the only way an operation can manufacture lead time that the notification process does not supply. A meta-analysis of seventeen studies in nursing found that prior sick leave predicted future leave with an odds ratio of 3.35 (95 per cent confidence interval 1.37 to 8.19), alongside musculoskeletal pain at 2.41 and back pain specifically at 3.05 (Gohar et al. 2021). At the shift level, an analysis of 601,282 shifts and 8,090 distinct sickness absence episodes found that when more than 75 per cent of shifts worked in the previous seven days were twelve hours or longer, the adjusted odds of a short-term sickness episode rose to 1.28 (Dall'Ora et al. 2019). Neither study records when the absence was notified.[1]

5 A race-condition model of shift recovery

5.1 Setup

A shift starts at time zero. The vacancy becomes known at time $-L$, so L is the notification lead time. A handling buffer b is consumed before the first contact can go out, leaving a usable window of $L - b$. The operation contacts workers from an eligible pool of size n, at a throughput of c attempts per hour. Each contacted worker responds with probability r, at a latency drawn from a distribution with cumulative function F, and accepts with probability a conditional on responding inside the window.

Table 3. Model notation.

Symbol	Meaning
L	Notification lead time: hours between the moment the vacancy becomes known and the shift start.
b	Handling buffer: time consumed before the first contact can go out, and after the last acceptance, so the usable window is $L - b$.
n	Eligible pool: workers who could lawfully and practically work this shift, after credentials, site clearance, hours limits and travel. Not the roster.
c	Contact throughput: attempts that can be initiated per hour.
$m(L, n)$	Effective pool: $\min(n, c(L - b))$. The number of workers who exist as far as this shift is concerned.
F	Cumulative distribution of response latency among workers who ever respond.
r	Share of contacted workers who ever respond.
a	Acceptance probability conditional on responding inside the window.
q	Common-shock probability: the chance that a factor affecting the whole pool at once, weather, a holiday, a plant-wide condition, makes nobody available.

1. That is worth pausing on. The Dall'Ora study attaches absence episodes to individual shifts, which is the finest unit of analysis in any source we verified, and even it does not carry a notification timestamp. The variable is absent at every level of aggregation, not just the national one.

Two structural features distinguish this from the standard treatment. The first is that the number of workers who exist, as far as this shift is concerned, is not n but the number who can actually be contacted inside the window:

$$m(L, n) = \min\{n, c(L - b)\} \tag{1}$$

The second is a correlated component. With probability q, a factor acting on the whole pool at once, weather, a statutory holiday, a plant-wide condition, makes nobody available. The existence of such a component is well established: the canonical call-centre formulation notes that the show-up proportion and the arrival rate may be dependent, for example through weather (Whitt 2006), and a distributionally robust treatment shows that absenteeism is endogenous to the staffing level itself (Jiang and Ryu 2019).

5.2 The probability of fill, and the ceiling

A single contacted worker yields a fill by lead time L with probability

$$p(L) = r \cdot a \cdot F(L - b) \tag{2}$$

Attempts are independent conditional on no common shock, so the parallel-system reliability result gives the probability that at least one succeeds (Lanzafame n.d.), and the void probability of a Poisson process gives the same quantity in continuous time when responses are modelled as a superposed arrival stream (Siegrist n.d.; MIT OpenCourseWare). Combining with the common shock:

$$\mathrm{P}(\text{fill} \mid L, n) = (1 - q) \cdot [\,1 - (1 - p(L))^{m(L,n)}\,] \tag{3}$$

Three consequences follow directly, and none of them depends on the parameter values.

1. *There is a ceiling that depth cannot cross.* The supremum over n is $1 - q$, for every value of L. No roster size reaches it and none exceeds it. The only way to raise the ceiling is to reduce the correlated component, which means predicting the shock rather than outnumbering it.
2. *Lead time enters twice and depth enters once.* L appears inside F, raising the per-worker success probability, and inside m, raising the count of workers who can be reached. n appears only inside m, where it is capped by $c(L - b)$. This asymmetry is the formal content of the paper's title.
3. *In the short-notice regime, depth is irrelevant by construction.* Whenever $c(L - b) < n$, the fill probability does not depend on n at all. An operation whose notice is short enough is operating in a regime where its roster size cannot appear in the answer.

5.3 *Parameterisation*

Every parameter below is either a stated assumption or an assumption anchored to a verified figure from an adjacent domain, and Table 4 states the direction of likely bias in each transfer. We want to be blunt about this: there is no published response-latency distribution for a shift offer, so the numerical results are illustrations of a structure rather than estimates of a quantity.

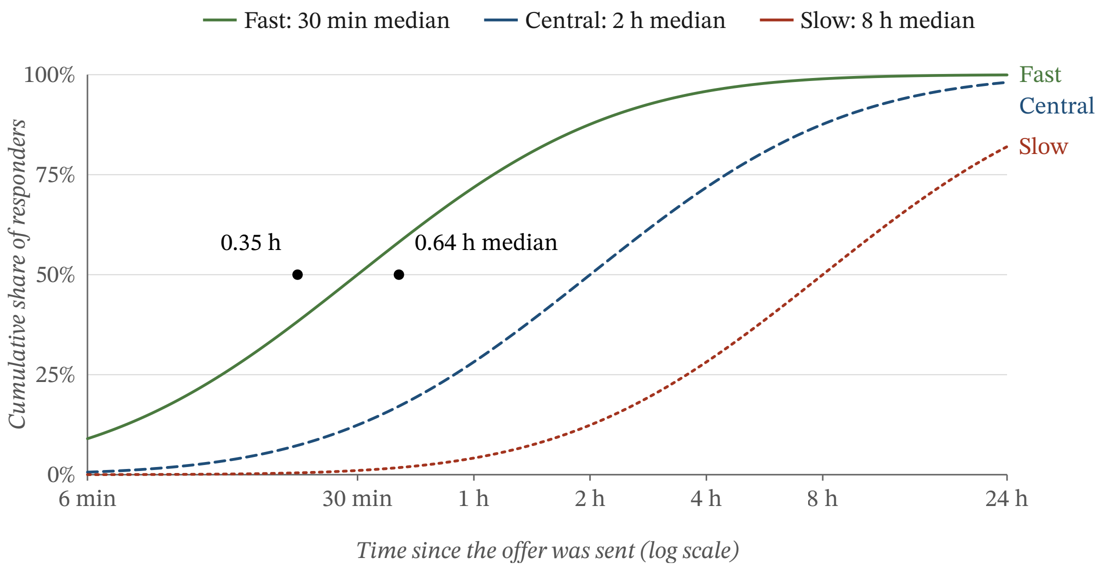

Figure 2. The response-latency distribution, which is the second of the two clocks in the model. The three curves are the parameterisations used throughout: medians of thirty minutes, two hours, and eight hours, with a heavy tail in each case. The two marked points are the only published medians we could verify for a comparable outbound message, both from smoking cessation trials, and both bracketed by the central curve. No published response-latency distribution exists for a shift offer, so these curves are assumptions with verified anchors, not fitted estimates.

Source: Curves computed for this paper (lognormal, log-scale 1.2). Anchor points: Naughton, Riaz and Sutton (2016), *Nicotine & Tobacco Research* 18(5), reporting aggregated median response times of 0.35 hours among 198 pregnant smokers and 0.64 hours among 293 general smokers, with response rates of 61.9 and 67.8 per cent. Peer reviewed. Compatible with Pielot, Church and de Oliveira (2014), which reports notification attendance medians from 3.5 to 27.7 minutes across 6,854 notifications, with a tail past an hour.

Table 4. Parameter values, provenance, and the direction of likely bias. No value in this table is a published measurement of a shift-offer response.

Parameter	Value	Nearest verified evidence	Direction of likely bias
Response latency median, central	2 h	Aggregated medians of 0.35 h and 0.64 h for text-message assessments in cessation trials (Naughton et al. 2016); notification attendance medians of 3.5 to 27.7 minutes across 6,854 notifications (Pielot et al. 2014).	Both anchors come from consented participants with a prior relationship to the sender, which should make them *faster* than a cold shift offer. We therefore set the central median well above both and report a fast case at 0.5 h and a slow case at 8 h.

Parameter	Value	Nearest verified evidence	Direction of likely bias
Ever-responds share (r)	0.55	Response rates of 61.9 and 67.8 per cent in cessation trials, with reminders adding 13.8 and 17.7 points (Naughton et al. 2016); a 20-day web panel response rate of 82 per cent (Pew 2016); single-digit response rates for unsolicited telephone contact to a general population (Pew 2019; Pew 2017).	A shift offer with pay attached to an existing worker should beat cold telephone contact and lose to a consented panel. The value sits between them, closer to the panel.
Acceptance given response (a)	0.65	No verified estimate exists.	*Unanchored.* This is the weakest parameter in the model. It enters multiplicatively with r, so the two cannot be identified separately from fill-rate data alone.
Common-shock probability (q)	0.08	The canonical call-centre formulation treats absence as a random show-up proportion and notes that it may be dependent with demand, for example through weather (Whitt 2006); a distributionally robust formulation shows absenteeism is endogenous to the staffing level itself (Jiang and Ryu 2019).	*Assumed.* The existence of a correlated component is well established in the operations literature; its magnitude in light industrial staffing is not published. It sets the ceiling, so a reader who disagrees should substitute their own value directly.
Sequential contact throughput (c)	8/h	Mean 4.29 calls to first contact among contacted cases, with attempts ranging up to 102 (Olson 2006).	*Assumed.* The mechanism, that attempts consume wall-clock time so effort and lead time are not independent budgets, is verified; the rate is not.
Broadcast throughput	20/h	Texting alongside email raised the share responding within thirty minutes from 6 per cent to 15 (Pew 2016).	*Assumed.* The channel effect on speed is verified; the throughput rate is a modelling choice.
Handling buffer (b)	0.25 h	No verified estimate exists.	*Assumed.* Raising it shifts every curve right and worsens the short-notice case; it does not change the ceiling.
Log-scale of the latency distribution	1.2	Half of notifications viewed within a few minutes and the majority within an hour, implying a heavy right tail (Pielot et al. 2014).	*Assumed.* A thinner tail would make short notice less punishing; a fatter tail would make it worse.

Notes: The purpose of this table is to make the model falsifiable by inspection. Any reader who substitutes their own values for the eight rows above regenerates every figure in section 5, and the qualitative conclusions survive wide variation in all of them except q, which sets the ceiling directly.

5.4 Results

Under the central parameterisation, with sequential voice contact at eight attempts per hour, the probability of filling a shift at thirty minutes' notice is 2.7 per cent, and it is 2.7 per cent whether the eligible pool holds five workers or a hundred and twenty. At one hour it is 34 per cent for any pool of six or more. At two hours it is 84 per cent, and nothing beyond fifteen eligible workers improves it. At four hours and beyond, every pool size of forty or more sits at the ceiling.

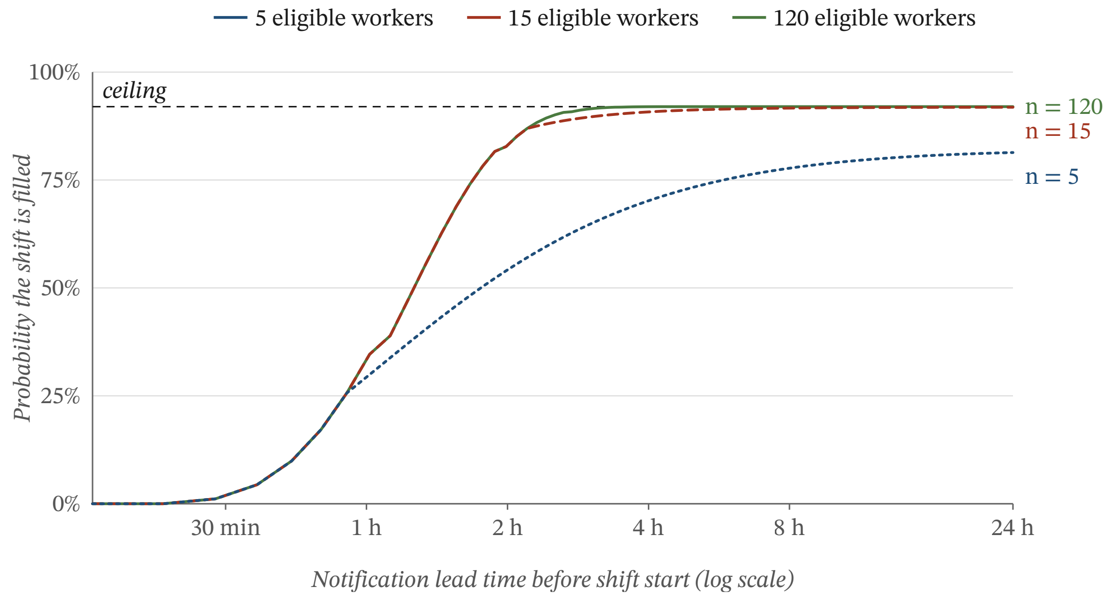

Figure 3. The central result. Below roughly one hour of notice the three curves are indistinguishable: a pool of a hundred and twenty eligible workers performs exactly as a pool of five, because contact throughput caps how many can be reached inside the window and response latency has barely begun. Between about one and four hours, depth matters. Above four hours every pool of fifteen or more is pinned to the ceiling and only the smallest pool is still gaining. Pool depth therefore changes the outcome in one narrow band. Lead time changes it across the whole range, and it is the only variable in the model that reaches the ceiling at all.

Source: Model computed for this paper under the central parameterisation and sequential voice contact at eight attempts per hour. All parameters are stated assumptions; see Table 4 for values, provenance, and the direction of likely bias in each transfer. The dashed ceiling is one minus the common-shock probability and is independent of pool size by construction.

Eligible pool	Notification lead time: 30 min	1 h	2 h	4 h	8 h	24 h
5 workers	3%	29%	54%	70%	78%	81%
15 workers	3%	34%	84%	91%	92%	92%
40 workers	3%	34%	84%	92%	92%	92%
120 workers	3%	34%	84%	92%	92%	92%

Figure 4. The same model as a grid, which makes the asymmetry easier to read across than down. Reading down a column, the gain from multiplying the eligible pool by twenty-four is zero at thirty minutes and zero again beyond four hours. Reading across a row, the gain from lead time is large everywhere and largest for the smallest pool. An operation with five eligible workers and four hours of notice outperforms an operation with a hundred and twenty and one hour.

Source: Model computed for this paper; sequential contact at eight attempts per hour, central response latency, common-shock probability 0.08. Values are percentage probabilities of filling the shift, rounded. This is model output, not measurement.

Read the grid across rather than down and the asymmetry is stark. Multiplying the eligible pool by twenty-four buys nothing at thirty minutes and nothing beyond four hours. It buys something only in the narrow middle band between about one and three hours. Lead time, by contrast, moves the outcome at every point in the range, and it moves it most where the operation is most exposed.

The throughput channel deserves separating out, because it is the part of the argument that a reader is most likely to resist. If all n workers could be contacted simultaneously, short notice would be survivable: an unconstrained pool of forty reaches 41 per cent even at thirty minutes' notice. The constraint is not the pool, it is the rate at which contact can be initiated and confirmed.

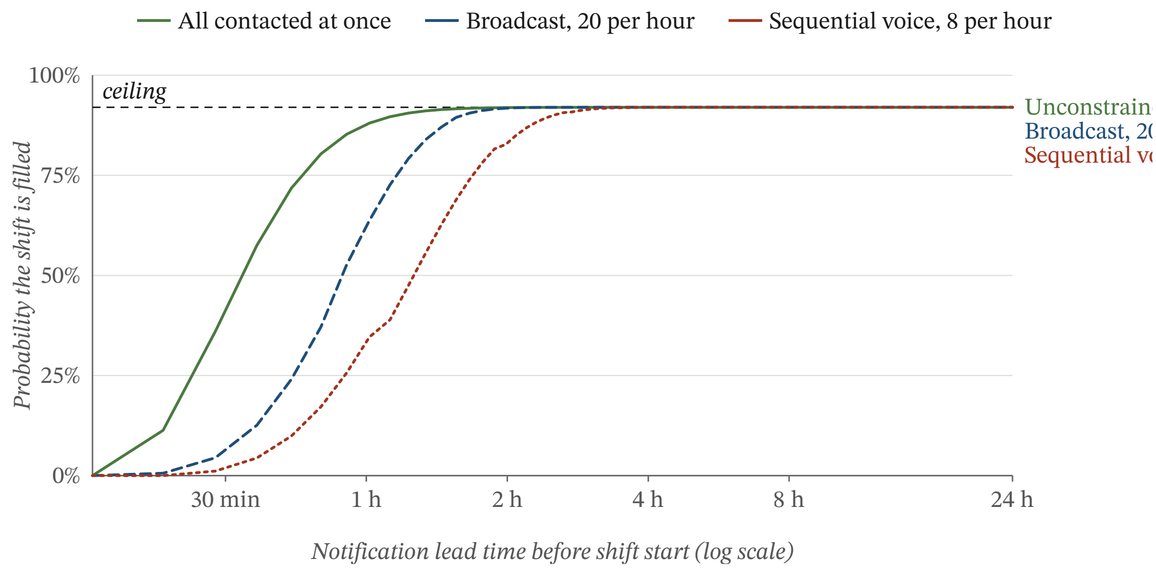

Figure 5. Why lead time enters the model twice. Holding the eligible pool at forty workers, the three curves differ only in how fast contact can be initiated. Under the unconstrained assumption, where all forty are reached at once, short notice is survivable. Under sequential voice contact at eight attempts an hour it is not, because the window bounds how many workers exist as far as the shift is concerned. Lead time therefore acts on the response distribution and on the effective pool simultaneously, while roster size acts on one of them and is capped by the other.

Source: Model computed for this paper. The sequential throughput figure is an assumption, calibrated against the observation that the mean number of calls to first contact among contacted cases in a large survey was 4.29, with attempts ranging up to 102 (Olson, 2006). Peer reviewed. Attempts consume wall-clock time, so effort and lead time are not independent budgets.

This is where the survey methodology literature earns its place in a staffing paper. Attempts consume wall-clock time, so effort and lead time are not independent budgets. In a large study with a 71 per cent response rate, the mean number of calls to first contact among contacted cases was 4.29, with attempts ranging up to 102 (Olson 2006). An operation that responds to short notice by trying harder is spending the same scarce resource it is short of.

6 Independent corroboration of concavity in time

The model's central prediction is that reach is concave in the length of the available window, steeply so at the short end. That prediction has been tested directly, in a different field, in about the cleanest way available.

Holding the population and the contact technology fixed and varying only the field period, a five-day standard survey achieved a 25 per cent response rate while a rigorous survey conducted over twenty-one weeks or more achieved 50 per cent. Seventy-seven of eighty-four comparable items were statistically indistinguishable between the two designs (Keeter et al. 2006). Extending the window roughly thirtyfold doubled reach.

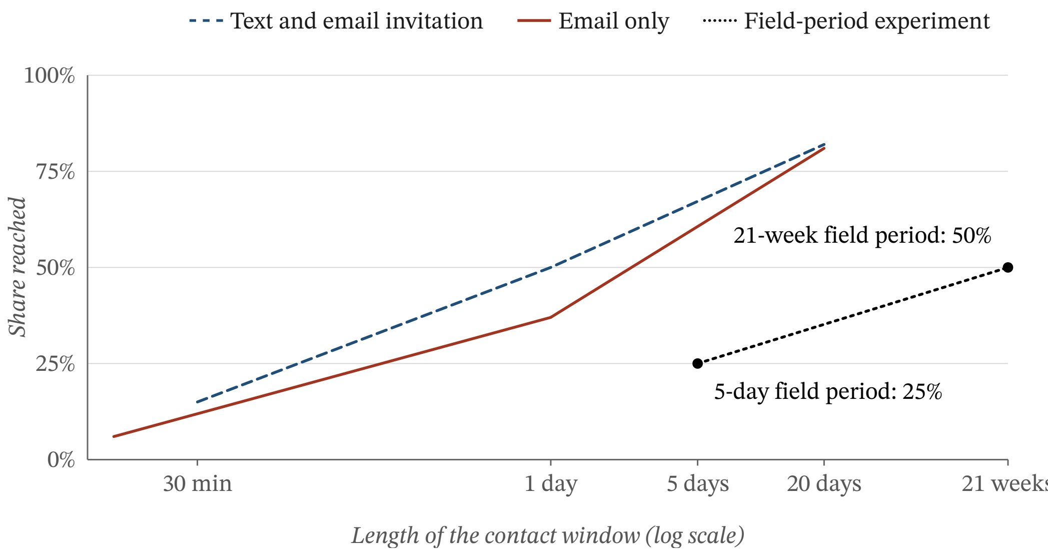

Figure 6. Reach is concave in time, and the concavity has been measured directly in an adjacent field. The two marked points are from a single study that held population and contact technology fixed and varied only the field period: five days produced a 25 per cent response rate and twenty-one weeks produced 50 per cent. Extending the window roughly thirtyfold doubled reach. Read in the other direction, that is the strongest published evidence for the claim this paper makes: compressing the window into hours puts an operation on the steep part of the curve, where the marginal value of lead time is at its maximum and no amount of effort substitutes for it. The two curves show the same shape in a channel experiment where texting alongside email raised the share reached in the first thirty minutes from 6 per cent to 15 and changed the twenty-day total not at all.

Source: Field-period points: Keeter, Kennedy, Dimock, Best and Craighill (2006), *Public Opinion Quarterly* 70(5), comparing a five-day standard survey against a rigorous survey run over twenty-one weeks or more, with 77 of 84 comparable items statistically indistinguishable between designs. Peer reviewed. Channel curves: McGeeney and Yan (2016), Pew Research Center, 2,109 panelists, reporting 15 per cent versus 6 per cent response in the first thirty minutes, 50 versus 37 per cent by the end of the first full day, and 82 versus 81 per cent at twenty days. Research organisation. Curves interpolate between those three reported points in each arm.

Read forwards, that is a result about diminishing returns to patience. Read backwards, it is the strongest published evidence for this paper's claim: if a thirtyfold extension of the window at the long end only doubles reach, then the curve is steep at the short end, and an operation compressed into hours is sitting where the marginal value of an extra hour is at its maximum.

Two further results from the same literature sharpen the point. A channel experiment found that texting alongside email raised the share responding within the first thirty minutes from 6 per cent to 15 per cent, and the share responding by the end of the first full day from 37 per cent to 50 per cent, while leaving the twenty-day totals statistically indistinguishable at 81 and 82 per cent (Pew 2016). Channel changed the speed of reach and not its eventual extent, which is exactly the distinction between filling before a deadline and filling eventually. And a factorial experiment on 17,769 addressable sample members found that the time of day a message was sent did not affect any outcome measured, while the sequencing of contacts within the window shifted mean days to completion by about eight days (Christian et al. 2024). Where you sit in the window matters; what hour of the day you press send does not.

7 What the operations research literature models

The operations research literature on staffing under absence is sophisticated, and it consistently models absence as a quantity resolved at a single decision point rather than as a process with a notification time. The cleanest statement is the canonical call-centre formulation, in which absenteeism enters as a random proportion of scheduled servers who are present, with no time index at all (Whitt 2006).

Table 5. What the operations research literature does and does not model. Absence is treated almost universally as a quantity resolved at a single decision point rather than as a process with a notification time.

Study	What it contributes	How absence enters
Whitt (2006), call centre staffing	The canonical formulation of staffing under absence uncertainty. Absenteeism enters as a random proportion of scheduled servers who are present, with no distributional form imposed and the possibility of dependence with demand noted.	**A scalar with no time index.** This is the cleanest statement of the field's standard treatment.
Maass et al. (2017), nurse staffing	Stochastic programming over patient census and nurse absenteeism with three personnel classes, unit, pool and temporary, and day-to-day census correlation.	A random count resolved at one epoch.
Jiang and Ryu (2019), nurse staffing	Absenteeism is *endogenous*: the number who show up depends partly on the staffing level chosen. Units with higher absenteeism should be pooled.	Endogenous, but still resolved at one epoch. Important qualification for any attempt to treat effort as an independent input.
Slaugh et al. (2018), long-term care	An on-call pool reduces absence-driven staffing costs by 24 per cent while slightly reducing schedule inconsistency; restricted pools outperform unrestricted.	The gain comes from pre-arranging availability, which is a lead-time device rather than a depth device.
Doneda et al. (2026), robust rostering	The one paper whose object is buying lead time. Reserve shifts are on-call duties during which an employee can be called in; the paper asks how accurate absence prediction must be to beat a naive buffer.	As a prediction target, which is the closest the literature comes to the question in this paper.
Yuan (2025), flexible capacity	A multistage staffing strategy indexed by decision horizon: long-term base staffing, short-term agency staffing, real-time allocation.	Implicitly by horizon, which is the right framing, applied to capacity rather than to notification.
Diab et al. (2014), transit extraboard	Four years of driver absence data, 14,305 observations. Mean absence of 9.9 drivers per garage-day-period with a standard deviation of 10.4.	As an overdispersed count. The near-equality of mean and standard deviation is why depth alone is a weak lever: the tail strands you.

Study	What it contributes	How absence enters
Bayliss et al. (2020), airline reserve crew	A probabilistic crew absence and recovery model, matching simulation accuracy at much lower computational cost.	Probabilistically, over reserve deployment timing.
Barrer (1957), queueing with impatience	The origin of the constant-deadline abandonment model, in which a customer whose total waiting time reaches a fixed threshold becomes a lost customer.	Not at all, but this is the correct formal object: an unfilled shift is a job that is lost at a deadline, not delayed past it.

Notes: The gap this table documents is specific. The formal machinery for a fill-before-deadline problem exists and is elementary: superposition of independent point processes gives the aggregate response stream (Siegrist n.d.), the void probability of a Poisson process gives the probability that nothing arrives before the deadline (MIT OpenCourseWare), and the parallel-system reliability result gives the probability that at least one of several independent attempts succeeds (Lanzafame n.d.). Nobody has assembled them into a race between a notification-time distribution and a response-latency distribution, which is what the operating problem actually is.

Two results in that literature point the same way as this paper without saying so. On-call pools reduce absence-driven staffing costs by 24 per cent while slightly reducing schedule inconsistency, and the gain comes from pre-arranging availability rather than from having more people (Slaugh et al. 2018). And the one paper whose object is buying lead time asks how accurate absenteeism prediction must be for a reserve-shift policy to beat a naive buffer, finding that a predict-then-optimise approach outperforms non-data-driven policies even under undemanding accuracy requirements, particularly where skills are interchangeable (Doneda et al. 2026). Both are lead-time devices dressed as capacity devices.

One empirical detail from the transit literature is worth extracting because it explains why depth is a weaker lever than intuition suggests. Across four years and 14,305 garage-day-period observations, mean absence was 9.9 drivers per garage-day-period with a standard deviation of 10.4 (Diab et al. 2014). When the standard deviation approximately equals the mean, the distribution is heavily overdispersed, and a reserve sized to the mean is stranded routinely by the tail. The recommended scenario in that study, a forty per cent adjustment, would cover 92.5 per cent of absences, which is a useful reminder of the order of magnitude at which coverage saturates in practice.

8 Two caveats against the metric itself

8.1 A filled shift is not a neutral outcome

Fill rate treats a covered shift as a success. A cross-sectional study of 8,841 registered nurses across 13,218 acute care shifts found that “the odds of a ‘care left undone’ event are similar for fully staffed shifts with a high temporary nursing staff ratio compared to severely understaffed shifts with no temporary nursing staff,” and that on fully staffed shifts, raising the temporary staff share from zero to ten per cent raised the odds of missed care by six per cent (Senek et al. 2020).

The setting is clinical and the transfer to a warehouse is not automatic. What does transfer is the structural point: a filled shift staffed by someone unfamiliar with the site delivers less than a filled shift staffed by a regular, and fill rate cannot see the difference. The overstatement is worst precisely in the short-notice regime, where the substitute is least likely to have worked the site before.

8.2 Notification timing is not exogenous

The second caveat is more serious for anyone tempted to solve coverage by contacting more people more often. Schedule instability is a strong and robust predictor of turnover. Drawing on panel data from 1,827 hourly workers in retail and food service, exposure to schedule instability predicted turnover, with slightly less than half the relationship mediated by job satisfaction and another quarter by work-family conflict (Choper et al. 2022). The reported associations were substantial: on-call shifts were associated with a 21 per cent higher turnover probability, cancelled shifts 38 per cent higher, and less than one week's notice 35 per cent higher than two weeks or more.

The health evidence points the same way. Across 27,792 workers, exposure to routine schedule instability was associated with psychological distress, poor sleep quality and unhappiness, and the association with instability was "much more strongly" pronounced than the association with low wages (Schneider and Harknett 2019).

The strongest causal evidence closes the loop. A nine-month field experiment across 28 retail stores found that implementing responsible scheduling practices raised store productivity by 5.1 per cent, through a 3.3 per cent sales increase and a 1.8 per cent labour reduction, and identified the mechanism as "motivating additional employee effort and reducing barriers to employees adhering to the scheduled labor plan" (Kesavan et al. 2022).

The reflexivity result

That final clause is the one to hold onto. The mechanism the authors name is not better recovery from vacancies. It is better adherence to the labour plan, which acts on the process that generates the vacancy rather than on the response to it. Read alongside the turnover evidence, which is associational panel data rather than experimental, this suggests that notification lead time is not purely an exogenous constraint on fill rate. If short notice degrades the pool it draws on, then an operation that treats same-day coverage as a contact-volume problem is worsening the input to next quarter's version of the same problem. We state this as a plausible feedback rather than as a demonstrated one, because no study we found closes the loop directly.

9 Notice in the field, and what regulation does to it

How much notice is there in practice? The distribution is bimodal rather than centred, which matters because an operation does not face an average notice period. In a service-sector sample of 5,364 workers, 31 per cent received less than one week's notice, 28 per cent between one and two weeks, and 41 per cent two weeks or more (Schneider et al. 2023). In a national early-career sample of 3,739 wage and salaried workers, 38 per cent knew their schedule one week or less in advance, rising to 41 per cent among hourly workers, while an almost identical 39 per cent of hourly workers knew it four or more weeks ahead (Lambert et al. 2014; Golden 2015).

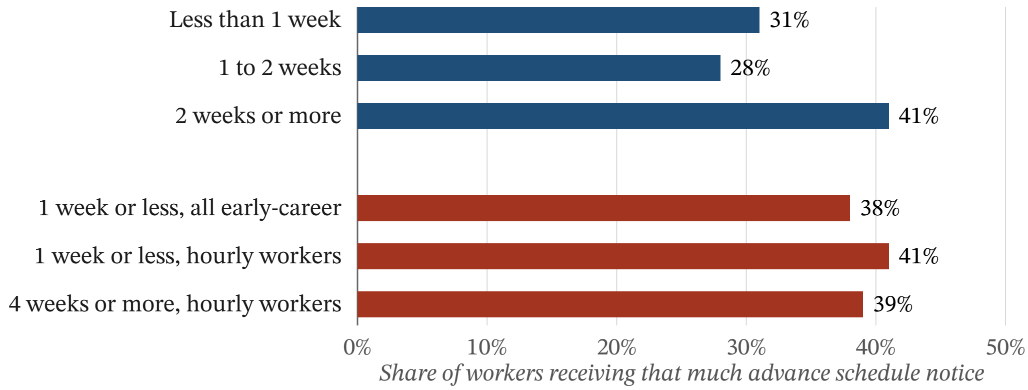

Figure 7. Advance schedule notice is bimodal, not centred. The upper group is a service-sector sample; the lower group is a national early-career sample in which roughly two in five hourly workers know their schedule a week or less ahead and roughly the same share of hourly workers know it four or more weeks ahead, with little mass in between. That shape matters for the model: an operation does not face an average notice period, it faces a mixture of a comfortable regime and a very short one, and its overall fill rate is a weighted average of two very different points on the curves in Figure 3.

Source: Upper group: Schneider, Harknett, Gailliot and Mopsick (2023), *Working in the Service Sector in Michigan*, The Shift Project, 5,364 workers surveyed July 2017 to December 2022. Research centre publication. Lower group: Lambert, Fugiel and Henly (2014), analysing the National Longitudinal Survey of Youth 1997 round 15, N = 3,739 wage and salaried workers aged 26 to 32 Federal longitudinal survey., with the same figures reported in Golden (2015). The two groups are different populations and the categories are not directly comparable; they are shown together because both establish the bimodality.

The thin middle has a modelling implication. An operation's realised fill rate is a weighted average over two very different points on the curves in Figure 3, and a change in that weighting will move the headline number without any change in capability. This is a second reason unconditional fill-rate comparison misleads.

Regulation compresses the short tail without eliminating it. Two years after Seattle's secure scheduling ordinance, the share of covered workers receiving two weeks' notice or more had risen from 46.3 to 57.0 per cent, and last-minute changes without pay had fallen from 69.5 to 49.6 per cent. On-call shifts moved from 25.2 to 21.2 per cent and unpaid cancellations from 14.5 to 13.7 per cent, but neither of those two impact estimates was statistically significant (Harknett et al. 2021).

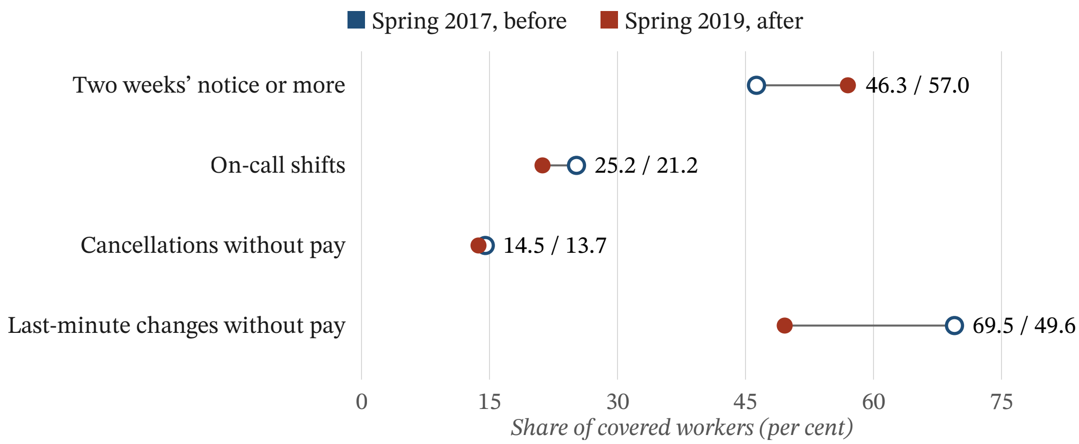

Figure 8. What regulation does to the notice distribution, and what it does not do. Two years after a well-enforced predictable-scheduling ordinance, the share of covered workers receiving two weeks' notice or more had risen from 46 per cent to 57 per cent. That is a real effect and it leaves roughly two in five covered shifts still inside the short-notice regime. Of the four impact estimates, only the gains on two weeks' notice and on last-minute changes were statistically significant; the reductions in on-call shifts and in unpaid cancellations were not. Predictable-scheduling law compresses the tail of the notice distribution; it does not remove the same-day problem, which is why the model in section 5 remains the binding description of the operating environment even under regulation.

Source: Harknett, Schneider and Irwin (2021), *Seattle's Secure Scheduling Ordinance: Year 2 Worker Impact Report*, Seattle Office of City Auditor. Municipal audit office report. Comparing spring 2017 with spring 2019: Seattle baseline 759 covered workers and follow-up 441, against comparison-city samples of 5,394 and 2,619. Reported impact estimates were +10.1 percentage points on two-week notice, −6.2 on on-call shifts, −12.5 on last-minute changes, and −3.2 on cancellations. The on-call and cancellation estimates were not statistically significant.

Those are real effects and they leave roughly two in five covered workers still inside the short-notice regime. Predictable-scheduling law is not a solution to the problem this paper describes; it is a partial redistribution of the notice distribution, and the same-day regime survives it.

10 A measurement protocol

The outcome measure is inherited rather than invented. The Canadian light industrial benchmark study now in the field defines same-day coverage as, of the shifts "that lost their worker inside 12 hours of the start time, the share that were covered before the shift started," and defines a filled shift as one "covered by a confirmed worker who attended" (Kordis Research n.d.). We adopt both. The twelve-hour boundary is the operative one for this paper, because it is the regime in which the curves of Figure 3 are steep, and because a coverage measure that counts a filled shift the worker did not attend would understate exactly the quality problem set out in section 8.1.

The model becomes a measurement with five fields, all of which an operation that dispatches by phone or message already generates implicitly.

Table 6. Minimum viable event schema. Five fields turn the model in section 5 into a measurement, and an operation that already dispatches by phone or message generates all five implicitly.

Field	Definition	Why it is required
`t_known`	Timestamp at which the vacancy became known to the operation: the call-out, the no-show determination, or the client's order.	This is the variable that does not exist in any published dataset. Everything else in the model is conditional on it.
`t_start`	Scheduled shift start.	`t_start` − `t_known` is the lead time L, and it is the paper's independent variable.
`eligible_n`	Count of workers who could lawfully and practically take the shift at `t_known`, after credentials, site clearance, hours limits and travel.	Separates pool depth from roster size. Fill-rate benchmarks that use roster size are measuring the wrong denominator.
`attempts[]`	Per attempt: worker, channel, dispatch timestamp, delivery timestamp, response timestamp, outcome.	Yields the response-latency distribution F and the ever-responds share r directly, and separates dispatch from delivery, which a broadcast channel conflates.
`filled_by`, `t_fill`, `worked_to_completion`	Which attempt produced the fill, when, and whether the replacement actually worked the shift.	The last field is not pedantry. A filled shift that is then abandoned is a worse outcome than an unfilled one, and fill rate as normally reported cannot see the difference.

Notes: The estimand is the probability of fill before shift start as a function of lead time, with requisition and client-site fixed effects to absorb the possibility that short-notice shifts are also unattractive shifts. The identifying variation is within-site variation in `t_known`, which arises for reasons unrelated to the shift itself, such as when in the preceding day a worker reports illness. A single quarter of one mid-sized operation's dispatch log would be the first dataset in the published record capable of estimating it.

Two design notes. The dispatch-to-delivery distinction in the attempts array is not fussiness: a broadcast channel conflates them, and a queued message that never arrives is not contact. And the estimand should be a probability of fill before shift start rather than an eventual fill rate, because the two differ exactly in the way the channel experiment showed, where texting changed speed without changing the twenty-day total (Pew 2016).

11 Discussion and limitations

The practical implications are uncomfortable for the way this industry sells and measures itself.

Unconditional fill rate is not comparable across operations. Two operations with identical capability will report different fill rates if their clients differ in when they call. The comparable metric is fill probability conditional on lead time, which requires the timestamp in Table 6 and which nobody currently keeps.

Roster growth is the wrong response to a short-notice problem. In the regime where $c(L - b) < n$, the fill probability does not contain n. Recruiting more workers to solve a forty-minute call-out is spending on a term that has dropped out of the equation.

The highest-return interventions act on the clock. Anything that moves notice earlier: predicting absence from prior absence (Gohar et al. 2021), pre-committing availability through reserve or on-call arrangements (Slaugh et al. 2018; Doneda et al. 2026), or reducing the vacancy-generating process through schedule stability (Kesavan et al. 2022), operates on the variable that enters the model twice.

Limitations

- *The model is illustrative, not estimated.* Not one of its parameters is a published measurement of a shift-offer response. Table 4 states each value and the direction of likely bias, and readers who substitute their own values regenerate every figure. The qualitative conclusions survive wide variation in all parameters except q, which sets the ceiling directly.
- *Every latency parameter is transferred across domains.* The verified response-latency evidence comes from smoking cessation trials, web survey panels, random-digit-dial telephone surveys, and mobile phone notification studies (Naughton et al. 2016; Pielot et al. 2014; Pew 2016; Olson 2006). A shift offer with pay attached plausibly beats a cessation-trial message on motivation and loses to a consented panel on obligation. The transfer is unavoidable and it is stated rather than hidden.
- *The independence assumption is doing real work.* Conditional on no common shock, contacted workers are treated as independent. If responses are positively correlated for reasons the common-shock term does not capture, for example because eligible workers share a commute or a household, the fill probability is lower than modelled at every point.
- *The notification distribution is assumed, not measured.* This is the paper's own central gap. We model outcomes conditional on L rather than integrating over its distribution, precisely because that distribution is unpublished.
- *The Canadian absence figures are dated.* The detailed composition we cite has a 2011 reference year. The series continues to 2025 in table 14-10-0191-01, but its values do not render to an automated reader and we did not retrieve them, so we report the older figures rather than approximate the newer ones (Statistics Canada n.d.).
- *The predictable-scheduling evidence is US service sector, not Canadian light industrial.* The turnover and health associations (Choper et al. 2022; Schneider and Harknett 2019) and the ordinance evaluation (Harknett et al. 2021) come from retail and food service in the United States. We could not verify a Canadian equivalent.
- *The quality caveat in section 8.1 is clinical.* The missed-care finding (Senek et al. 2020) is from acute nursing, where the consequences of unfamiliarity are unusually visible. Whether the same magnitude holds in a warehouse is unknown.

- *The vendor-class exclusion cuts both ways.* We declined to cite any staffing fill-rate benchmark because none had a stated methodology. That keeps the paper clean and it also means we offer no empirical baseline against which to compare the model output.

12 Conclusion

An unfilled shift is not a delayed job. It is a job that is lost at a deadline, which is a formally different object, and the operations literature has had the right model for it since 1957 (Barrer 1957). What has not been done is to write shift recovery as a race between the distribution of when the vacancy becomes known and the distribution of how quickly workers respond. Doing so produces a result that is hard to see from inside the industry's own metrics: lead time enters the fill probability twice, once through response latency and once through the number of workers who can be reached at all, while pool depth enters once and is bounded above by a correlated component that no roster crosses.

The practical version is that in the short-notice regime a large roster and a small one perform identically, and that in the comfortable regime they also perform identically. Pool depth matters in a narrow middle band. Notice matters everywhere. And if the association between short notice and turnover is causal, an operation that meets a timing problem with contact volume is trading a solved shift today for a thinner roster next quarter.

None of this is measured, and the reason is that a single timestamp is not recorded. No statistical agency in Canada or the United States collects the moment a vacancy becomes known, and no published fill rate is conditioned on it. One quarter of one mid-sized operation's dispatch log, with the five fields in Table 6, would be the first dataset in the published record capable of answering the question. Until then the honest description of the industry's headline metric is that it measures the behaviour of its clients at least as much as the capability of its operators.

Disclosure

Funding and interest

This paper was prepared at Kordis Research. Kordis sells software to temporary staffing agencies, and has a direct commercial interest in the conclusion that notification timing and response speed dominate coverage outcomes. Readers should weigh the argument accordingly. The model is presented with every parameter tabulated and its bias direction stated so that a sceptical reader can rebuild it on their own assumptions. No external funding was received and no client data was used.

What is derived rather than reported

Figures 2 through 5 and the numbers in Table 4 and Figure 4 are output of a model built for this paper. The model is specified in section 5, its parameters are tabulated with provenance in Table 4, and none of them is a published measurement of a shift-offer response. Figure 6 interpolates between reported points in each arm; only the marked and stated values are published. Figures 1, 7 and 8 contain published values only.

Typographic note

Set in STIX Two Text. There are no em dashes anywhere in the series. En dashes appear only in numeric ranges, page ranges, and quoted source titles.

References

Barrer, D. Y. 1957. “Queuing with impatient customers and ordered service.” *Operations Research* 5 (5): 650–656. https://ideas.repec.org/a/inm/oropre/v5y1957i5p650-656.html.

Bayliss, C., G. De Maere, J. A. Atkin, and M. Paelinck. 2020. “Scheduling airline reserve crew using a probabilistic crew absence and recovery model.” *Journal of the Operational Research Society* 71 (4): 543–565. https://nottingham-repository.worktribe.com/output/1458047.

Chartered Institute of Personnel and Development, with Simplyhealth. 2023. *Health and Wellbeing at Work*, September. https://www.cipd.org/globalassets/media/knowledge/knowledge-hub/reports/2023-pdfs/8436-health-and-wellbeing-report-2023.pdf.

Choper, J., D. Schneider, and K. Harknett. 2022. “Uncertain time: Precarious schedules and job turnover in the US service sector.” *ILR Review* 75 (5): 1177–1210. https://doi.org/10.1177/00197939211048484.

Christian, L. M., H. Sun, Z. Slowinski, C. Hansen, and M. McRoy. 2024. “When to text? How the timing of text message contacts in mixed-mode surveys impacts response.” *Journal of Survey Statistics and Methodology* 12 (3): 674–696. https://doi.org/10.1093/jssam/smae014.

Dall’Ora, C., J. Ball, O. Redfern, A. Recio-Saucedo, A. Maruotti, P. Meredith, and P. Griffiths. 2019. “Are long nursing shifts on hospital wards associated with sickness absence?” *Journal of Nursing Management* 27 (1): 19–26. https://researchportal.port.ac.uk/en/publications/are-long-nursing-shifts-on-hospital-wards-associated-with-sicknes.

Diab, E. I., R. A. Wasfi, and A. M. El-Geneidy. 2014. “Extraboard team sizing: An analysis of short unscheduled absences among regular transit drivers.” *Transportation Research Part A*. https://tram.mcgill.ca/Research/Publications/Extra_board_team_sizing.pdf.

Doneda, M., P. Smet, G. Carello, E. Lanzarone, and G. Vanden Berghe. 2026. “Robust personnel rostering: How accurate should absenteeism predictions be?” *Journal of Scheduling* 29 (1): 67–82. https://doi.org/10.1007/s10951-025-00858-9.

Gardner, P., K. Blackwell, and S. Young. 2025. *Bi-annual nurse staffing review*. Board paper 018.25. Shrewsbury and Telford Hospital NHS Trust, January 16. https://www.sath.nhs.uk/wp-content/uploads/2025/01/018.25-Bi-annual-Nurse-Staffing-Review.pdf.

Gohar, B., M. Larivière, N. Lightfoot, C. Larivière, E. Wenghofer, and B. Nowrouzi-kia. 2021. “Demographic, lifestyle, and physical health predictors of sickness absenteeism in nursing: A meta-analysis.” *Safety and Health at Work* 12: 536–543. https://doi.org/10.1016/j.shaw.2021.07.006.

Golden, L. 2015. *Irregular work scheduling and its consequences*. Briefing Paper No. 394. Economic Policy Institute, April 9. https://www.epi.org/publication/irregular-work-scheduling-and-its-consequences/.

Harknett, K., D. Schneider, and V. Irwin. 2021. *Seattle's Secure Scheduling Ordinance: Year 2 worker impact report*. Seattle Office of City Auditor, February. https://shift.hks.harvard.edu/wp-content/uploads/2021/02/Seattle-Year-2-Evaluation.pdf.

Harrison, D. A., and J. J. Martocchio. 1998. "Time for absenteeism: A 20-year review of origins, offshoots, and outcomes." *Journal of Management* 24 (3): 305–350. https://doi.org/10.1177/014920639802400303.

Jiang, R., and M. Ryu. 2019. *Nurse staffing under absenteeism: A distributionally robust optimization approach* [Preprint]. Optimization Online. https://optimization-online.org/2019/09/7381/.

Keeter, S., C. Kennedy, M. Dimock, J. Best, and P. Craighill. 2006. "Gauging the impact of growing nonresponse on estimates from a national RDD telephone survey." *Public Opinion Quarterly* 70 (5): 759–779. https://doi.org/10.1093/poq/nfl035.

Keeter, S., N. Hatley, C. Kennedy, and A. Lau. 2017. *What low response rates mean for telephone surveys*. Pew Research Center, May 15. https://www.pewresearch.org/methods/2017/05/15/what-low-response-rates-mean-for-telephone-surveys/.

Kennedy, C., and H. Hartig. 2019. *Response rates in telephone surveys have resumed their decline*. Pew Research Center, February 27. https://www.pewresearch.org/short-reads/2019/02/27/response-rates-in-telephone-surveys-have-resumed-their-decline/.

Kesavan, S., S. J. Lambert, J. C. Williams, and P. K. Pendem. 2022. "Doing well by doing good: Improving retail store performance with responsible scheduling practices at the Gap, Inc." *Management Science* 68 (11): 7818–7836. https://ideas.repec.org/a/inm/ormnsc/v68y2022i11p7818-7836.html.

Kordis Research. n.d. *Canadian temp staffing operations benchmarks (2026)*. Study protocol and pre-registered definitions. https://kordisai.com/research/canadian-light-industrial-staffing-benchmarks-2026.

Lambert, S. J., P. J. Fugiel, and J. R. Henly. 2014. *Precarious work schedules among early-career employees in the US: A national snapshot*. University of Chicago, EINet, August 27. https://bpb-us-w2.wpmucdn.com/voices.uchicago.edu/dist/5/1068/files/2018/05/lambert.fugiel.henly_.precarious_work_schedules.august2014_0-298fz5i.pdf.

Lanzafame, R. n.d. *Risk and Reliability for Engineers*, chapter 4: System reliability. Delft University of Technology. https://teachbooks.tudelft.nl/risk-reliability/reliability-system/overview.html.

Maass, K. L., B. Liu, M. S. Daskin, M. Duck, Z. Wang, R. Mwenesi, and H. Schapiro. 2017. "Incorporating nurse absenteeism into staffing with demand uncertainty." *Health Care Management Science* 20 (1): 141–155. https://doi.org/10.1007/s10729-015-9345-z.

Margheritti, S., L. Corthésy-Blondin, S. Vila Masse, and A. Negrini. 2025. "Work-related psychosocial risk and protective factors influencing workplace sickness absence: A systematic literature review." *Journal of Occupational Rehabilitation* 36 (1): 1–22. https://doi.org/10.1007/s10926-024-10265-0.

McGeeney, K., and H. Y. Yan. 2016. *Text message notification for web surveys*. Pew Research Center, September 7. https://www.pewresearch.org/methods/2016/09/07/text-message-notification-for-web-surveys/.

Naughton, F., M. Riaz, and S. Sutton. 2016. "Response parameters for SMS text message assessments among pregnant and general smokers participating in SMS cessation trials." *Nicotine & Tobacco Research* 18 (5): 1210–1214. https://doi.org/10.1093/ntr/ntv266.

Olson, K. 2006. “Survey participation, nonresponse bias, measurement error bias, and total bias.” *Public Opinion Quarterly* 70 (5): 737–758. https://doi.org/10.1093/poq/nfl038.

Pielot, M., K. Church, and R. de Oliveira. 2014. “An in-situ study of mobile phone notifications.” In *Proceedings of MobileHCI ’14*. ACM. https://pielot.org/pubs/Pielot2014-MobileHCI-Notifications.pdf.

Schneider, D., and K. Harknett. 2019. “Consequences of routine work-schedule instability for worker health and well-being.” *American Sociological Review* 84 (1): 82–114. https://doi.org/10.1177/0003122418823184.

Schneider, D., K. Harknett, A. Gailliot, and J. Mopsick. 2023. *Working in the service sector in Michigan*. The Shift Project, Harvard Kennedy School, February. https://shift.hks.harvard.edu/wp-content/uploads/2023/02/MI_Brief_2022_Update_Layout_Final.pdf.

Senek, M., S. Robertson, T. Ryan, R. King, E. Wood, and A. Tod. 2020. “The association between care left undone and temporary nursing staff ratios in acute settings.” *BMC Health Services Research* 20: 637. https://doi.org/10.1186/s12913-020-05493-y.

Siegrist, K. n.d. *Probability, Mathematical Statistics, and Stochastic Processes*, section 14.5: Thinning and superposition. Statistics LibreTexts. https://stats.libretexts.org/Bookshelves/Probability_Theory/Probability_Mathematical_Statistics_and_Stochastic_Processes_(Siegrist)/14:_The_Poisson_Process/14.05:_Thinning_and_Superpositon.

Slaugh, V. W., A. A. Scheller-Wolf, and S. R. Tayur. 2018. “Consistent staffing for long-term care through on-call pools.” *Production and Operations Management* 27 (12): 2144–2161. https://ideas.repec.org/a/bla/popmgt/v27y2018i12p2144-2161.html.

Society for Human Resource Management and Kronos. 2014. *Total financial impact of employee absences across the United States, China, Australia, Europe, India and Mexico*. https://www.shrm.org/content/dam/en/shrm/topics-tools/news/employee-relations/Total-Financial-Impact-of-Employee-Absences-Report.pdf.

Sohoni, M. G., E. L. Johnson, and T. G. Bailey. 2006. “Operational airline reserve crew planning.” *Journal of Scheduling* 9 (3): 203–221. https://doi.org/10.1007/s10951-006-6778-8.

Statistics Canada, Labour Statistics Division. 2012. *Work Absence Rates 2011*. Catalogue 71-211-X. https://www150.statcan.gc.ca/n1/pub/71-211-x/71-211-x2012000-eng.pdf.

Statistics Canada. n.d. *Quality of Employment in Canada: Night work, 2022*. Catalogue 14-28-0001. https://www150.statcan.gc.ca/n1/pub/14-28-0001/2023001/article/00008-eng.htm.

Statistics Canada. n.d. *Work absence of full-time employees by industry, annual*. Table 14-10-0191-01, series 1987 to 2025. https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1410019101.

Statistics Canada. n.d. *Work Absence Rates: data quality, concepts and methodology* (survey documentation 3701_D68_T25). https://www.statcan.gc.ca/en/statistical-programs/document/3701_D68_T25.

Tsitsiklis, J., and P. Jaillet. 2018. *Introduction to Probability*, RES.6-012, Lecture 22: The Poisson process. MIT OpenCourseWare. https://ocw.mit.edu/courses/res-6-012-introduction-to-probability-spring-2018/1f25a3930276f573b7be8cce5243f885_MITRES_6_012S18_L22AS.pdf.

US Bureau of Labor Statistics. n.d. *Table 47. Absences from work of employed full-time wage and salary workers by occupation and industry*. Current Population Survey, annual averages. https://www.bls.gov/cps/cpsaat47.pdf.

Whitt, W. 2006. “Staffing a call center with uncertain arrival rate and absenteeism.” *Production and Operations Management* 15 (1): 88–102. https://www.columbia.edu/~ww2040/POM2006.pdf.

Yuan, Y. 2025. “Managing flexible capacity in service systems with worker shortages.” *Manufacturing & Service Operations Management* 27 (3): 808–824. https://ideas.repec.org/a/inm/ormsom/v27y2025i3p808-824.html.

Appendix A. Source classification

Every source cited in this paper, with the class assigned to it on retrieval. Full details are in the reference list.

Table A1. Source classification.

Source	Class
Barrer, D. Y. 1957	Peer-reviewed
Bayliss, C., G. De Maere, J. A. Atkin, and M. Paelinck 2020	Peer-reviewed
Chartered Institute of Personnel and Development, with Simplyhealth 2023	Industry body
Choper, J., D. Schneider, and K. Harknett 2022	Peer-reviewed
Christian, L. M., H. Sun, Z. Slowinski, C. Hansen, and M. McRoy 2024	Peer-reviewed
Dall’Ora, C., J. Ball, O. Redfern, A. Recio-Saucedo, A. Maruotti, P. Meredith, and P. Griffiths 2019	Peer-reviewed
Diab, E. I., R. A. Wasfi, and A. M. El-Geneidy 2014	Peer-reviewed
Doneda, M., P. Smet, G. Carello, E. Lanzarone, and G. Vanden Berghe 2026	Peer-reviewed
Gardner, P., K. Blackwell, and S. Young 2025	Government department or agency
Gohar, B., M. Larivière, N. Lightfoot, C. Larivière, E. Wenghofer, and B. Nowrouzi-kia 2021	Peer-reviewed
Golden, L. 2015	Research centre
Harknett, K., D. Schneider, and V. Irwin 2021	Government department or agency
Harrison, D. A., and J. J. Martocchio 1998	Peer-reviewed
Jiang, R., and M. Ryu 2019	Peer-reviewed
Keeter, S., C. Kennedy, M. Dimock, J. Best, and P. Craighill 2006	Peer-reviewed
Keeter, S., N. Hatley, C. Kennedy, and A. Lau 2017	Research centre
Kennedy, C., and H. Hartig 2019	Research centre
Kesavan, S., S. J. Lambert, J. C. Williams, and P. K. Pendem 2022	Peer-reviewed
Kordis Research n.d.	Author’s own research programme
Lambert, S. J., P. J. Fugiel, and J. R. Henly 2014	Research centre
Lanzafame, R. n.d.	Monograph or courseware
Maass, K. L., B. Liu, M. S. Daskin, M. Duck, Z. Wang, R. Mwenesi, and H. Schapiro 2017	Peer-reviewed
Margheritti, S., L. Corthésy-Blondin, S. Vila Masse, and A. Negrini 2025	Peer-reviewed
McGeeney, K., and H. Y. Yan 2016	Research centre
Naughton, F., M. Riaz, and S. Sutton 2016	Peer-reviewed

Source	Class
Olson, K. 2006	Peer-reviewed
Pielot, M., K. Church, and R. de Oliveira 2014	Peer-reviewed
Schneider, D., and K. Harknett 2019	Peer-reviewed
Schneider, D., K. Harknett, A. Gailliot, and J. Mopsick 2023	Research centre
Senek, M., S. Robertson, T. Ryan, R. King, E. Wood, and A. Tod 2020	Peer-reviewed
Siegrist, K. n.d.	Monograph or courseware
Slaugh, V. W., A. A. Scheller-Wolf, and S. R. Tayur 2018	Peer-reviewed
Society for Human Resource Management and Kronos 2014	Industry body
Sohoni, M. G., E. L. Johnson, and T. G. Bailey 2006	Peer-reviewed
Statistics Canada n.d.	National statistical agency
Statistics Canada n.d.	National statistical agency
Statistics Canada n.d.	National statistical agency
Statistics Canada, Labour Statistics Division 2012	National statistical agency
Tsitsiklis, J., and P. Jaillet 2018	Monograph or courseware
US Bureau of Labor Statistics n.d.	National statistical agency
Whitt, W. 2006	Peer-reviewed
Yuan, Y. 2025	Peer-reviewed

Notes: Assigned on retrieval and not revised. Where a source appears in more than one paper in this series it carries the same class in each.

Kordis Research Working Paper Series · KRS-2026-03 · Version 1.0, August 2026 · Patrick Underwood
Not peer reviewed. Cite as: Underwood, P. (2026). *Notification Lead Time and the Recoverability of Unfilled Shifts* (Kordis Research Working Paper KRS-2026-03).