Before the Factory Opens:

Anticipation and Spatial Adjustment to Announced and Delayed Local Shocks

Harper Stephens

October 2026

Abstract

A factory can move people before it makes anything. I study this channel using Intel's Ohio semiconductor project and a spatial mean field game calibrated to 2021 LODES residences. The agents stand for the region's whole resident labor force, not only the people who might take a job at the plant. I compare workers who know the expected opening date with otherwise identical workers who learn of the project only when it opens. At the original opening horizon, the gap between them equals 46.7 percent of the remaining distance to the long-run equilibrium, measured in total variation. Later delays bend the path without letting the population jump at news dates. A matched two-dimensional extension cuts opening total variation by 16.3 percent but leaves the timing normalization almost untouched. The exercise decomposes an equilibrium by information set; it does not estimate a causal effect.

Keywords: mean field games; spatial equilibrium; forward-looking migration; expectations; local investment.

1 Introduction

A large investment project has three birthdays: the day construction starts, the day production starts, and a much earlier day when people first hear it is coming. Economic adjustment has no obligation to wait for the second. A worker can sign a lease, widen a job search, or move house because a local labor market is expected to improve. When the project slips, the logic runs in reverse, with one asymmetry: the future payoff recedes, but the moving truck has already left.

This paper studies that timing channel through Intel's Ohio One semiconductor project. Intel announced the project on January 21, 2022, promising 3,000 direct Intel jobs, roughly 7,000 construction jobs, and production in 2025; JobsOhio put average direct-job pay at about $135,000 per year (Intel, 2022; JobsOhio, 2022). The production date then moved, more than once. I use those revisions to ask a narrow question: *how much spatial adjustment can occur simply because workers know a local shock is coming before it arrives?*

The object of interest is not a conventional treatment effect. Before opening, the project pays nothing in either of the two counterfactual economies I compare. What differs is continuation information. In the forward-looking economy, workers learn at the announcement that the project will arrive at a maintained future horizon. In the surprise economy, otherwise identical workers ignore the project until the doors open. Both economies start from the same pre-announcement residence distribution and face the same contemporaneous environment before opening. Any pre-opening gap therefore comes from information about the future, plus whatever congestion response the resulting moves set off.

The distinction matters for measurement and for policy. If households start adjusting after the announcement but before operations, a research design that dates the shock at opening will file part of the response under "pre-treatment." A local government that plans housing, roads, and schools around the ribbon-cutting may find that demand arrived ahead of schedule. The economically relevant date is not necessarily the day the machines turn on. It may be the day a future change becomes credible enough to enter current decisions.

Intel Ohio suits this question because its information sequence is unusually public. There is a dated announcement, a stated production horizon, and a series of revisions that push the payoff further out. I do not treat these dates as moments when every worker formed the same belief. They define transparent information regimes for a model experiment. The object is the equilibrium response to a common, dated change in continuation information, not an estimate of heterogeneous subjective expectations.

The model is a spatial mean field game, and its logic fits in a paragraph. A worker weighs the discounted future gain from living closer to the project against the cost of moving now. Every other worker solves the same problem at the same time. Their choices reshape the population distribution, and the distribution determines congestion. In equilibrium, the population path each worker takes as given is the path their choices produce. The Hamilton–Jacobi–Bellman equation solves the worker's problem backward from the future; the Kolmogorov forward equation carries the population forward through time.

The title names a factory, but the model does not follow factory workers. Throughout the paper, "worker" means any member of Central Ohio's resident labor force. The baseline counts every primary job holder in the region, in every industry, and most of them will never work at the plant. The project enters their problem as a premium on nearby locations, and that premium applies to every worker, whether or not they would ever take a job at Intel. The paper tracks where labor lives, not who ends up on the plant's payroll.

I calibrate the model to 2021 Census LODES residence-area counts for primary jobs, the last year before the announcement. In the benchmark, Central Ohio is projected onto the Intel–Columbus axis. The one-dimensional representation is deliberately spare: it keeps the anticipation mechanism visible and allows unusually strict numerical checks. The pre-project distribution is matched exactly rather than approximately. I recover a discrete baseline fundamental under which the smoothed empirical residence distribution is stationary for the same finite-state generator used in the transition. The

no-project counterfactual is therefore flat to machine precision, and the announcement experiment does not inherit drift from a misfit baseline.

The headline result concerns timing. At the original four-year expected-opening horizon, the forward-looking and surprise economies differ by 0.009080 in total variation. The surprise distribution at that date sits 0.019438 from the post-project stationary distribution, which yields a timing normalization of 46.71 percent. In the projected ±5-mile band around the Intel site, the forward-looking economy holds 0.2046 percentage points more residence mass at that horizon. I call these ratios *timing normalizations*. They scale the pre-opening information gap by the transition still left to run. They are not probabilities, and nothing confines them to the unit interval.

The sequence of news matters too. A representative public revision in March 2024 moves the maintained expected-opening horizon from four years to six, and Intel's February 2025 update moves it to nine. At each news date the population distribution is inherited: information can revalue locations instantly, but it cannot teleport households. At the original expected-opening horizon, the sequentially revised path differs from the original-information path by 0.1977 percentage points in the projected band and 0.008414 in total variation.

A two-dimensional extension supplies the main geographic check. It keeps a 90×60 mile rectangle and lets workers move both along and across the Intel–Columbus axis. At matched numerical resolution, adding the second dimension lowers opening total variation from 0.009478 to 0.007935, a 16.3 percent reduction in the amount of reallocation. The timing normalization moves from 46.80 to 46.71 percent. Flattening the map to a line changes how much moves, but hardly changes when.

The paper draws on four literatures. Quantitative spatial economics treats population and economic activity as equilibrium outcomes of wages, amenities, commuting, and mobility (Allen and Arkolakis, 2014; Desmet and Rossi-Hansberg, 2014; Monte, Redding, and Rossi-Hansberg, 2018). Work on local development and announcement effects shows that expectations can be priced in before anything is built (Kline and Moretti, 2014; Cárdenas, Gallego, and Urrutia, 2023; Hyun and Milcheva, 2019). Dynamic spatial models make the timing channel explicit; Bilal and Rossi-Hansberg (2023), for example, study costly forward-looking migration and investment in response to anticipated climate change. Mean field games offer a natural continuous-space representation of forward-looking migration with distributional feedback (Lasry and Lions, 2007; Huang, Malhamé, and Caines, 2006; Achdou and Capuzzo-Dolcetta, 2010), and Bahlali, Boucekkine, and Petit (2026) develop a two-dimensional MFG theory of spatial dynamics.

My contribution is narrower than any of these. I do not claim that anticipation, migration frictions, or spatial mean field games are new. I provide a transparent announcement-to-opening equilibrium decomposition for a concrete local project; make the pre-project empirical distribution exactly stationary under the discrete model actually solved; let public information be revised sequentially; and show that the main timing result survives a matched two-dimensional extension. The exercise is quantitative, and structural in the sense that it solves an equilibrium model. It is not a causal estimate of Intel's realized effect on migration.

Section 2 describes the project, the data, and summary statistics. Section 3 presents the model

and the information counterfactual. Section 4 covers calibration and computation. Section 5 reports the benchmark, delay, and two-dimensional results. Section 6 discusses interpretation and limits. The appendices contain full derivations, the 2D extension, robustness and numerical validation, and data construction.

2 Setting and data

2.1 The information sequence

The January 2022 announcement sets the initial information. Intel said production was expected in 2025, and JobsOhio reported 3,000 direct Intel jobs at an average of about $135,000 (Intel, 2022; JobsOhio, 2022). The benchmark reads the coarse "2025" guidance as an expected horizon of four years after the announcement. That is a modeling convention, not a claim that workers circled a particular operating day on the calendar. Appendix C varies the initial horizon from 3.5 to 4.5 years.

The information changed. I use March 20, 2024 as a representative revision date, tied to reporting that pushed expected Ohio operations later, and February 28, 2025 as the second, when Intel published a further timeline update (Reuters, 2024; Intel, 2025). The maintained expected-opening horizons become six and nine years after the January 2022 announcement. These correspond roughly to the public 2027–28 and 2030–31 regimes, and I treat them as expectations rather than calendar facts.

Order matters here. The January 2022 path is solved using only the information assigned to that regime; later delays are not smuggled back into the initial expectation. When the 2024 revision arrives, the model inherits the population distribution built under the earlier belief and re-solves the continuation problem. The 2025 revision repeats the procedure. The experiment is path dependent by design: workers respond to new information from wherever old information already put them.

The model also separates the expected operating date from construction. Construction can raise local labor demand well before production, but that is not the channel isolated here. In the benchmark, the project premium switches on with the maintained operating regime. Any pre-opening difference between the forward-looking and surprise economies therefore reflects the value of the future operating regime, not an assumed construction boom. Construction employment, supplier entry, and other transitional channels are left out rather than folded into the anticipation effect.

2.2 Residence geography and local labor-market scale

The baseline population comes from the U.S. Census Bureau's LEHD Origin-Destination Employment Statistics (LODES), using residence-area counts for primary jobs (Census Bureau, 2026). The baseline year is 2021, before the announcement. After restricting to the retained Central Ohio geography, the baseline contains 969,829 worker residences. The count covers primary jobs in all industries (LODES segment S000, job type JT01), so the baseline is the region's resident labor force,

not its manufacturing workforce.

For the benchmark model, I project each residence onto the line running from the Intel site toward Columbus. The along-axis domain runs 45 miles on either side of Intel, and observations more than 30 miles off the axis are dropped. The result is a 90-mile state space. The primary local outcome is the mass in the projected ±5-mile band around the Intel site. This is a corridor statistic, not a radius. Appendix B keeps both coordinates and uses a true five-mile circle.

I scale the project shock with Bureau of Labor Statistics QCEW data for Licking County, aggregation level 70 (County, Total Covered) (BLS, 2026). In 2021 the county had 74,622 covered jobs with average annual pay of $47,306. QCEW and LODES count different universes, so I never equate their levels. QCEW has one job: to make the shock scale explicit, by measuring the announced job count against local covered employment and the announced Intel salary against local covered pay.

2.3 Descriptive evolution after the announcement

The post-announcement LODES files are useful as description and useless as causal validation. Table 1 reports the same retained geography for 2021–2023. Total primary-job residences rise, but the two local shares do not: both dip in 2022 and recover in 2023. That is one reason I do not calibrate the model to post-announcement local changes. Plenty of forces besides Intel can move residence patterns over those years, and LODES alone offers no control group that isolates the announcement.

Table 1: Descriptive LODES evolution in the retained geography

Year	Retained residences	Projected ±5-mile share	True five-mile share
2021	969,829	4.528%	0.781%
2022	994,181	4.507%	0.755%
2023	1,014,244	4.599%	0.800%

Notes: LODES primary-job residence counts using the same retained geography as the benchmark. The projected band and true five-mile circle are distinct geographic objects. These rows are descriptive only and are not used to estimate an Intel treatment effect.

The descriptive evidence therefore does a small job well. It shows that the post-announcement data contain no simple mechanical pattern the model is being bent to reproduce. The quantitative exercise instead fixes the pre-announcement geography and asks what the forward-looking mechanism implies under transparent counterfactual information sets. Fitting the model to 2022–2023 movement would blur the line between the mechanism under study and the many contemporaneous forces outside the model, so I keep calibration and post-announcement description separate.

2.4 Summary statistics and the baseline transformation

Table 2 separates three objects that are easy to conflate: observed inputs, the cross-sectional variables on the numerical grid, and maintained numerical and behavioral parameters. Panel A is

measured or announced. Panel B describes the 61 cell-level variables the one-dimensional equilibrium actually uses. Panel C records the reference specification behind the benchmark counterfactual.

Table 2: Data and reference calibration

Panel A. Observed and announced inputs

Quantity	Value	Source / interpretation
Retained 2021 LODES residences	969,829	Primary-job residence counts
Raw projected ±5-mile share	4.528%	2021 LODES
Raw true five-mile share	0.781%	2021 LODES
Licking County covered employment	74,622	2021 QCEW, total covered
Licking County average annual pay	$47,306	2021 QCEW, total covered
Announced direct Intel jobs	3,000	January 2022 information
Announced average Intel salary	$135,000	January 2022 information

Panel B. Cross-sectional grid variables, $N = 61$ cells

Quantity	Mean	SD	P10	Median	P90
Raw cell residence mass (%)	1.639	2.213	0.236	0.564	5.991
Smoothed target cell mass (%)	1.639	1.968	0.300	0.641	5.387
Recovered baseline fundamental	0.532	0.615	-0.347	0.618	1.302
Project premium	0.0405	0.0235	0.0103	0.0404	0.0727

Panel C. Reference specification

Spatial grid	61 cells	Smoothing	3.5 miles
Time step	0.10 years	Project scale	20 miles
Discount rate ρ	0.05	Relocation cost c	1.50
Diffusion variance σ^2	0.15	Congestion scale κ	0.30
Congestion curvature γ	1.00	Opening horizons	4 / 6 / 9 years

Notes: Panel B summarizes the finite-volume cells used in the benchmark model. Cell masses sum to one across the grid. The recovered fundamental and project premium are utility objects, not observed prices. The behavioral and numerical quantities in Panel C are maintained reference values and are varied in Appendix C. QCEW and LODES do not measure the same population and are combined only for shock scaling.

Before solving the equilibrium, the model smooths the 2021 projected residence mass with a 3.5-mile bandwidth. Figure 1 shows why the gap between raw and model geography deserves attention. The raw distribution is jagged, as discrete counts tend to be; smoothing produces the continuous target the PDE discretization needs. Smoothing also shifts local levels, which is why I treat highly local level effects with caution and lean on counterfactual differences and total variation.

3 Model and information counterfactual

The model has one mechanism: because workers look ahead, a location-specific payoff that has not arrived yet can change where they move today. The equations below formalize that sentence. Readers mainly interested in the quantitative experiment can read the HJB as the worker's dynamic optimization problem and the KFE as population bookkeeping; Appendix A derives both step by step.

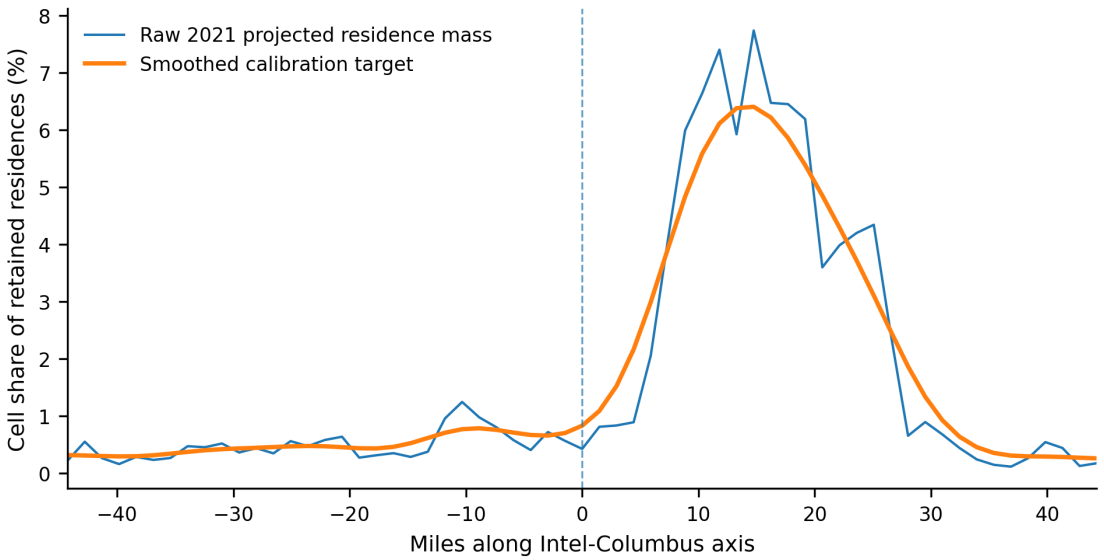


Figure 1: Raw and smoothed 2021 residence mass along the Intel–Columbus axis. The vertical dashed line marks the Intel site.

Notes: Raw residence mass is the 2021 LODES projected baseline after the retained-geography filter. The smoothed series is the exact target used in the benchmark inverse calibration. Both series integrate to one over the retained sample.

3.1 Economic mechanism in words

Start with one worker and two locations that differ mainly in their distance from the future project. Today the plant is not running, so neither location pays an Intel-related dividend. Yet the two locations do not have the same continuation value. Being closer to the future job center makes tomorrow more attractive even when today looks identical.

If moving were free and instant, the model would have almost no transition dynamics. Workers would wait for opening day and then jump. Relocation costs rule that out. Moving fast is expensive, so a worker who expects a valuable future location has a reason to start early and spread the move over time. Discounting pulls the other way: a payoff many years out is worth less than the same payoff next year. The balance between the two determines how soon anticipation shows up.

Population feedback adds a second layer. When many workers want the same places, density rises there and congestion makes those places less appealing. Each worker therefore responds both to the future project and to everyone else's response to it. That interdependence is what makes this a mean field game rather than a pile of independent migration decisions. In equilibrium, individual optimization and aggregate population movement have to agree.

Delays follow from the same logic. When news pushes the expected opening further out, the discounted value of living near Intel falls on the spot. The population cannot jump back on the news date, because residences are a physical state, but optimal future movement can change immediately. The model therefore predicts gradual, partial unwinding of earlier anticipation, not an instant reversal.

3.2 Location, movement, and flow utility

Time is continuous, and the January 2022 announcement is normalized to $t = 0$. In the benchmark, location is $x \in [-1, 1]$, where one unit equals 45 miles along the Intel–Columbus axis and $x = 0$ is the Intel site. A worker chooses purposeful drift a_t, and location follows the reflected diffusion

$$\mathrm{d}X_t = a_t\,\mathrm{d}t + \sigma\,\mathrm{d}W_t + \mathrm{d}K_t, \tag{1}$$

where K_t is the minimal finite-variation regulator that keeps X_t inside $[-1, 1]$. The drift captures deliberate relocation. The diffusion term is a reduced-form stand-in for idiosyncratic movement and omitted heterogeneity. It should not be read literally; households do not follow Brownian paths, however it may look on moving day.

At time t, flow utility at location x is

$$u(t, x, a; m) = b(x) + r(t)\pi(x) - \kappa\big(m(t, x) + \varepsilon\big)^{\gamma} - \frac{c}{2}a^2. \tag{2}$$

The baseline fundamental $b(x)$ summarizes persistent location value before the project: job access, housing, amenities, family ties, and anything else, none of it separately identified. The regime indicator $r(t)$ switches the project payoff on after the maintained opening horizon. Congestion lowers utility where density is high. The quadratic movement cost makes rapid purposeful relocation expensive.

The project premium is smooth and centered on Intel,

$$\pi(x) = \omega \exp\left[-\frac{1}{2}\left(\frac{x}{h}\right)^2\right], \qquad h = 20/45. \tag{3}$$

The 20-mile scale is maintained, not estimated, and the robustness appendix varies it.

3.3 The worker problem and population accounting

Let $V(t, x)$ denote the value of being at x at time t. Dynamic programming gives the Hamilton–Jacobi–Bellman equation

$$-\partial_t V + \rho V = D V_{xx} + \frac{(V_x)^2}{2c} + b(x) + r(t)\pi(x) - \kappa(m + \varepsilon)^{\gamma}, \qquad D = \sigma^2/2, \tag{4}$$

with optimal purposeful drift

$$a^*(t, x) = \frac{V_x(t, x)}{c}. \tag{5}$$

The backward direction carries the economics. A project that does nothing today can still raise the continuation value of nearby locations, which changes V_x and therefore changes movement before opening.

Given those choices, the density of workers evolves according to

$$\partial_t m = D m_{xx} - \partial_x \left(m \frac{V_x}{c} \right). \tag{6}$$

Reflection imposes the boundary conditions

$$V_x(t, \pm 1) = 0, \qquad \left[D m_x - m \frac{V_x}{c} \right]_{x=\pm 1} = 0. \tag{7}$$

The first is the Neumann condition for the worker's value problem; the second says no population leaks out of the domain. Equation (4) asks where a worker wants to go, given the aggregate path. Equation (6) asks where the population goes, given those policies. A mean field equilibrium is a fixed point between the two questions.

3.4 The information experiment

The benchmark compares two economies with the same initial distribution, the same parameters, and the same eventual post-project environment.

(i) In the *forward-looking* economy, workers know about the project at $t = 0$ and expect the regime to switch at the maintained horizon $T_0 = 4$.

(ii) In the *surprise-opening* economy, workers ignore the project before T_0. At T_0 the project is revealed, and from then on the economy faces the same post-opening environment as the forward-looking economy.

For $t < T_0$, both economies have $r(t) = 0$, so the project pays no flow utility in either economy before opening. The only exogenous difference is continuation information. Once policies diverge, population density and congestion can diverge endogenously as well. Appendix A states this decomposition formally.

This is an identification statement *inside the model*. It says what the counterfactual comparison measures: holding fixed the initial distribution, the pre-opening project regime, the behavioral parameters, and the baseline fundamental, the pre-opening gap is generated by different beliefs about the future regime. It does not say that Ohio data would reveal the same gap causally. Actual post-announcement migration also responds to construction, other employers, housing supply, interest rates, demographics, and local policy. Those channels are not differenced out of observed data here. They are simply held outside the structural comparison.

The distinction heads off a common misreading. The forward-looking path is not a forecast of what Ohio "should" have done, and the surprise path is not a statistical control group. They are two solutions to the same equilibrium model under different information sets. Their difference measures the model's anticipation channel, and nothing more.

3.5 How to read the outcome measures

I measure the distance between distributions with total variation,

$$\mathrm{TV}(p, q) = \frac{1}{2} \sum_i |p_i - q_i|. \tag{8}$$

For probability distributions on a common grid, this is the smallest amount of probability mass that must be reassigned to make the two coincide. The timing normalization is

$$R_{TV} = \frac{\mathrm{TV}\Big(p^{\mathrm{FL}}(T_0), p^{\mathrm{UA}}(T_0)\Big)}{\mathrm{TV}(p^{\mathrm{post}}, p^{\mathrm{UA}}(T_0))}. \tag{9}$$

The numerator is the opening-date separation produced by anticipation. The denominator is how far the surprise economy still has to travel at opening to reach the post-project stationary distribution. An analogous local normalization uses mass in the projected ±5-mile band. Both ratios are descriptive normalizations, not structural parameters and not probabilities.

Total variation has the virtue of requiring no neighborhood cutoff. A value of 0.009 means that 0.9 percent of normalized spatial probability mass would have to be reassigned to turn one distribution into the other. It does not mean 0.9 percent of workers move because of the announcement. The simulation does not track individual moves; it works with normalized residence mass, not person-level migration histories. The statistic is a distance between distributions.

The projected-band measure answers a different question: how much extra mass piles up in a narrow corridor around the site? It is easier to picture but more sensitive to smoothing and geometry. I therefore treat total variation as the preferred cross-specification measure and the band as a secondary description. The two-dimensional appendix swaps the projected band for a true geographic circle, one more reason not to compare local levels mechanically across dimensions.

4 Calibration and computation

4.1 What the data discipline

The calibration uses four kinds of objects, and keeping them apart is half the interpretation. First, the 2021 LODES distribution is observed and fixes the pre-announcement spatial mass the model must support. Second, the Intel announcement and QCEW data pin down the scale of the future local payoff through jobs and relative pay. Third, the baseline fundamental is recovered inversely so that the observed smoothed residence distribution is an equilibrium of the exact discrete model. Fourth, several behavioral and numerical quantities are maintained and stress-tested rather than estimated from these data.

The inverse baseline deserves a closer look. The function $b(x)$ is not an observed amenity index, and it is not the output of a regression of residence density on covariates. It is the reduced-form location value required to rationalize the pre-project distribution under the maintained mobility,

diffusion, discounting, and congestion structure. Economically, it absorbs every persistent reason a worker might prefer one stretch of the corridor to another before Intel: other jobs, housing, amenities, schools, family networks, and any stable local characteristic left unnamed.

The approach trades structural decomposition for a clean transition experiment. Because $b(x)$ is recovered rather than identified piece by piece, the paper cannot say how much of it is "amenities" and how much is "labor-market access." What it can guarantee is that the model generates no unrelated migration before the project arrives. That guarantee matters here, because an anticipation exercise in which the population is already drifting under a no-project counterfactual is measuring something other than anticipation.

The maintained parameters have a different status. Relocation cost, diffusion variance, the discount rate, the congestion parameters, the smoothing bandwidth, and the project-kernel width are reference values, not precise Ohio estimates. They govern the speed and spatial shape of adjustment. Appendix C asks whether the main conclusions survive economically meaningful changes to them. The paper leans hardest on the conclusion that proves most stable across those changes: the timing of adjustment, not an exact neighborhood-level population effect.

4.2 Shock scale and maintained parameters

The reduced-form peak project intensity uses only information available at the 2022 announcement:

$$\omega = \left(\frac{3000}{74622}\right)\left(\frac{135000 - 47306}{47306}\right) = 0.074526. \tag{10}$$

The first factor is announced direct jobs relative to local covered employment. The second is the announced salary premium relative to local covered pay. Their product sets the scale of the location-specific payoff. It is not a wage elasticity, an employment multiplier, or a causal estimate.

The reference values in Panel C of Table 2 are maintained rather than estimated: discount rate $\rho = 0.05$, relocation cost $c = 1.50$, diffusion variance $\sigma^2 = 0.15$, congestion scale $\kappa = 0.30$, and congestion curvature $\gamma = 1$. The benchmark uses 61 spatial cells, a 0.10-year time step, 3.5-mile smoothing, and a 20-mile project kernel. Appendix C varies each economically important choice and reports numerical refinement separately.

4.3 Finite-state approximation

The computation replaces the continuous interval with finite-volume cells. From an interior cell i, a worker chooses a nonnegative rightward intensity a_i^+ and leftward intensity a_i^-. Diffusion adds symmetric transitions between neighbors. The off-diagonal generator entries are

$$q_{i,i+1} = D/\Delta x^2 + a_i^+/\Delta x, \tag{11}$$

$$q_{i,i-1} = D/\Delta x^2 + a_i^-/\Delta x, \tag{12}$$

and each diagonal entry is minus the sum of outward rates, so every row of the generator Q sums to zero.

For a fixed value vector, the optimal directional controls are

$$a_i^+ = \left[\frac{V_{i+1} - V_i}{c\Delta x}\right]_+, \qquad a_i^- = \left[\frac{V_{i-1} - V_i}{c\Delta x}\right]_+. \tag{13}$$

The stationary discrete HJB is

$$\rho V_i = [Q(a)V]_i + b_i + r\pi_i - F_i(p) - C_i(a), \tag{14}$$

where $F_i(p) = \kappa(p_i/\Delta x + \varepsilon)^\gamma$ and $C_i(a) = c[(a_i^+)^2 + (a_i^-)^2]/2$. The forward population equation uses the transpose of the same generator. That shared operator, used for both optimization and bookkeeping, is what makes the inverse calibration below exact.

4.4 Exact discrete inverse baseline

A calibration can fit an observed baseline and still manufacture migration if the baseline is recovered with one operator and simulated with another. The benchmark avoids that trap by making the smoothed 2021 LODES target an equilibrium of the exact discrete generator used in the transition.

Let p^0 denote the smoothed target probabilities. Across each pair of adjacent cells, I choose the value difference so that diffusion and optimal directed movement satisfy detailed balance. This constructs V^0 such that

$$Q(a(V^0))^\top p^0 = 0. \tag{15}$$

The baseline fundamental then follows from the same stationary discrete HJB,

$$b_i = \rho V_i^0 - [Q(a(V^0))V^0]_i + F_i(p^0) + C_i(a(V^0)). \tag{16}$$

The value level is pinned down only up to an additive constant. Controls depend on value differences alone, so the normalization does not touch behavior. Appendix A proves the detailed-balance inverse, normalization invariance, mass conservation, positivity, and invertibility of the policy-evaluation systems.

The resulting baseline is stationary to numerical precision. The TV distance between the smoothed target and the solved pre-project stationary distribution is 7.65×10^{-15}. A no-project path moves only 1.24×10^{-14} in TV by the opening horizon. Both are rounding error next to the economic distances reported below.

4.5 Transition algorithm

Given a candidate population path, the algorithm solves the HJB backward in time and recovers the implied directional controls. It then solves the population equation forward from the empirical initial distribution and moves the candidate path toward that forward response, repeating until the

fixed-point discrepancy falls below 8×10^{-7}. The terminal condition is the stationary post-project value function, and the horizon runs 12 years past the latest maintained opening date so that the transition is close to the post-project stationary distribution when the computation ends.

The benchmark forward-looking and surprise fixed-point errors are 6.96×10^{-7} and 6.98×10^{-7}. Maximum mass errors are about 10^{-14}, and terminal TV distances to the post-project equilibrium are about 10^{-6}. Appendix C reports the full diagnostics and refinement exercises.

5 Results

5.1 Anticipation before opening

Figure 2 is the central experiment. Both paths start from the same pre-announcement distribution. The surprise economy sits still before T_0, because the project appears in neither its current utility nor its continuation value. The forward-looking economy starts moving early, because proximity is already worth something before the plant runs.

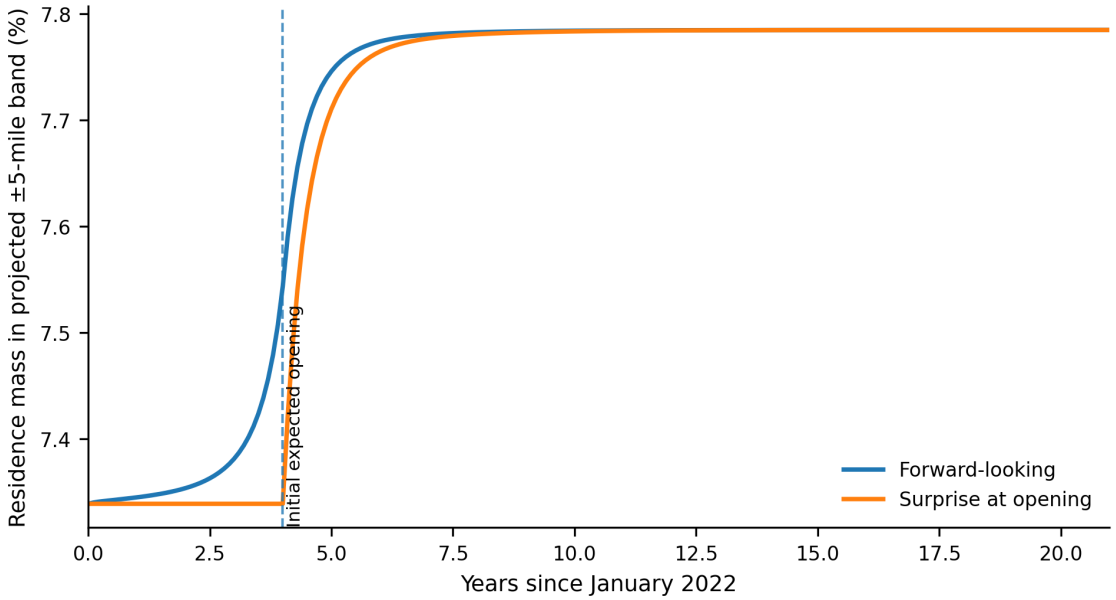


Figure 2: Anticipatory adjustment in the benchmark one-dimensional model. The dashed vertical line marks the original four-year expected-opening horizon.

Notes: Both economies begin from the same stationary 2021 baseline and face the same contemporaneous pre-opening project regime. The blue path incorporates the future project from the January 2022 announcement; the orange path does not incorporate it until opening.

The shape of the paths says as much as the gap between them. The forward-looking path begins to move right after the announcement, but slowly. Early on, the payoff is heavily discounted and there is plenty of time to adjust later. As opening approaches, procrastination gets expensive: less time remains to reposition before the new regime begins. The anticipatory path steepens accordingly. The acceleration is an equilibrium consequence of discounting and costly movement, not a shape imposed on the path.

The surprise path doubles as a diagnostic. It stays at the calibrated baseline before opening because the project is absent from both current utility and continuation value, and because the inverse calibration made that baseline stationary. If the surprise path drifted before opening, the opening gap would mix anticipation with a misspecified starting equilibrium. In the final benchmark, no-project TV drift at the opening horizon is 1.24×10^{-14}, so the visible separation in Figure 2 is not a baseline artifact.

After opening, the surprise economy sets off toward the same post-project equilibrium that the forward-looking economy is already approaching, and the two paths converge. That catch-up is exactly what the experiment is built to isolate: advance information changes the *timing* of the transition, while the destination is common to both economies.

At the opening horizon, the forward-looking economy holds 0.2046 percentage points more mass in the projected ±5-mile band. The full distributions differ by 0.009080 in total variation, or 0.908 percent of normalized population mass. The surprise economy at opening sits 0.019438 from the post-project stationary distribution. Therefore

$$R_{TV} = 0.009080/0.019438 = 0.4671.$$

The analogous projected-band normalization is 45.89 percent.

Table 3 gathers the main quantities and sets the geographic extension beside the benchmark. The local level result is modest and moves with smoothing and spatial representation. The timing statistic holds still. The model is a better guide to *when* adjustment happens than to how many residences end up in a given neighborhood.

Table 3: Main quantitative results

	Benchmark 1D	Matched 1D	Preferred 2D
Spatial resolution	61 cells	121 cells	121×81 cells
Time step	0.10	0.10	0.10
Opening total variation	0.009080	0.009478	0.007935
Long-run TV benchmark	0.019438	—	0.016989
TV timing normalization	46.71%	46.80%	46.71%
Local opening gap (pp)	0.2046[a]	—	0.0400[b]
Local timing normalization	45.89%[a]	—	44.87%[b]

Notes: The matched 1D specification uses the same along-axis resolution and time step as the preferred 2D specification. Moving from matched 1D to 2D reduces opening TV by 16.3 percent. [a]Projected ±5-mile corridor band. [b]True geographic five-mile circle. The two local statistics are not directly comparable; total variation is the common metric.

5.2 Magnitude versus timing

The benchmark makes two kinds of quantitative statement. One is about magnitude: how far apart are the forward-looking and surprise distributions at the original opening horizon? The other is about timing: how large is that gap relative to the transition still needed to reach the post-project

equilibrium? There is no reason for the two to be equally robust. A change in smoothing or in the width of the project premium can shove mass across a narrow local boundary without touching the intertemporal logic that produces early adjustment.

Table 3 bears this out. Opening TV is 0.009080 in the 1D benchmark, 0.009478 in the matched higher-resolution 1D model, and 0.007935 in the preferred 2D model. Geography clearly matters for how much mass gets reallocated. The corresponding timing normalizations are 46.71, 46.80, and 46.71 percent. What survives the change in representation is not a level effect. It is the finding that, at the original expected-opening horizon, the information gap is about 47 percent as large as the remaining surprise-to-post-equilibrium TV distance.

For the same reason, I do not convert normalized mass into a count of movers. LODES records residences tied to primary jobs, the model smooths those counts into a continuous density, and total variation summarizes a difference between distributions rather than a set of identified individual moves. Turning 0.009080 into a headcount would claim precision the model does not have. The defensible statement concerns the size of the distributional gap and its timing relative to the model's own long-run benchmark.

5.3 Public delays and inherited geography

The announcement is only the first piece of news. When the expected opening slides further out, continuation values change at once, but the distribution of residences carries over from the previous regime. That state dependence is what separates a dynamic information experiment from a string of unrelated comparative statics.

A delay works through discounting and the option to wait. A worker who expected the project soon had a reason to absorb some relocation cost early. Push the payoff out, and the reason weakens. The optimal response need not undo every earlier move: the worker already lives somewhere new, moving again costs money, and the project is still coming eventually. The model therefore predicts partial unwinding, neither a shrug nor a full retreat.

This gives the revision exercise a different reading from an "early opening versus late opening" comparison. The later regime does not restart from the 2021 baseline. It starts from the distribution the earlier belief created, and that inherited geography is the state variable through which stale information keeps mattering after expectations change.

Figure 3 shows the sequential experiment. Until the first revision, the sequential path coincides with the original 2022-information path. After the 2024 and 2025 revisions, it bends toward the later expected-opening regimes. The surprise path is included as a benchmark for the timing channel.

At the original expected-opening horizon, the sequentially revised path differs from the original-information path by 0.1977 percentage points in the projected band and 0.008414 in total variation. Locally, it has drifted back toward the surprise benchmark: projected-band mass is 7.3460 percent under sequential revisions, against 7.3390 percent under surprise and 7.5437 percent under the original 2022 information. The delays undo most of the local concentration implied by the initial expectation, even though the distribution never jumps at a news date. Both revision segments

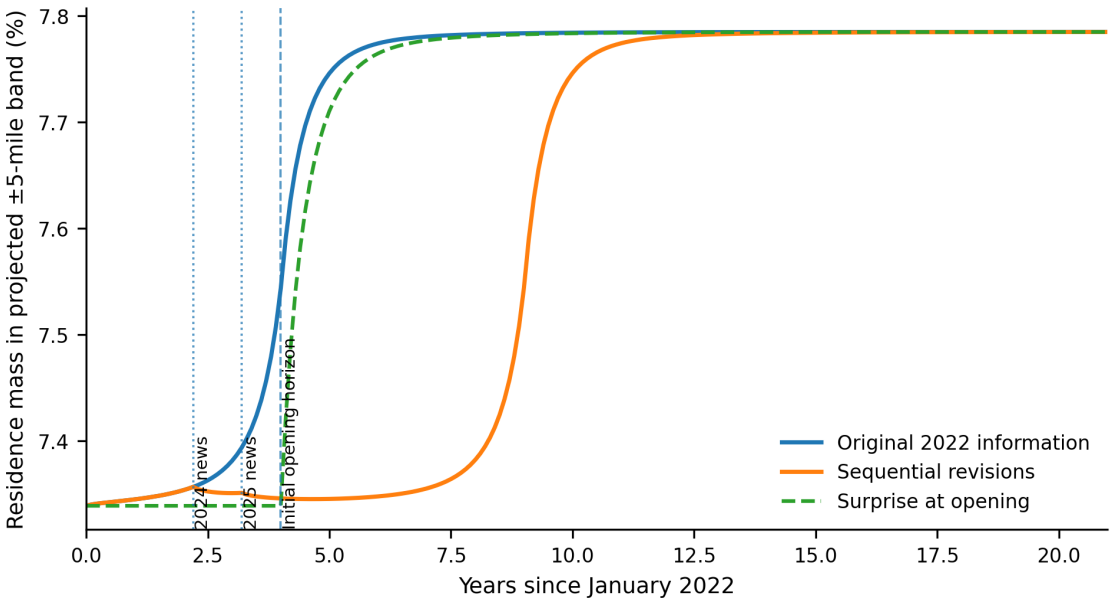


Figure 3: Sequential public revisions. Dotted lines mark the representative 2024 and 2025 news dates; the dashed line marks the original expected-opening horizon.

Notes: The population distribution is continuous at each news date. The value function and optimal controls can change immediately because the continuation information changes.

inherit population exactly in the numerical implementation, with fixed-point errors below 6.3×10^{-7}. Nor does the result hinge on one chosen delay: setting the 2024/2025 expected horizons to 5.5/8.5 or 6.5/9.5 years moves the opening TV in the revision comparison only within 0.008309 to 0.008490.

5.4 Two-dimensional geography

The benchmark lets workers move along one axis only. Appendix B rebuilds the model on the full 90×60 mile retained rectangle. The extension is not decorative: the population does not live on a line. The employment-weighted median absolute distance from the axis is 7.52 miles, and the 90th percentile is 22.63 miles.

The preferred 2D specification uses a 121×81 grid and the same 0.10-year time step as its matched one-dimensional comparator. Opening TV is 0.007935 in 2D against 0.009478 in matched 1D, so the extra direction of movement shrinks the information-induced redistribution by 16.3 percent. The timing normalization is 46.71 percent in 2D and 46.80 percent in matched 1D; the frozen 61-cell benchmark is 46.71 percent. This near equality is the best evidence that the headline timing result is not an artifact of squashing the map onto one axis.

The direction of the magnitude change makes sense. In one dimension, purposeful adjustment can only run along the Intel–Columbus axis. In two, workers can also spread out across it, so the same future project is accommodated through a wider set of moves. The response is less concentrated in the common metric even though the forward-looking incentive arrives at the same moment. Adding geographic freedom changes *where* the equilibrium moves far more than *when* its anticipatory part

begins.

The 2D model also replaces the projected corridor band with a true five-mile circle. In the preferred specification, the forward-looking economy holds 0.0400 percentage points more mass inside that circle at the original opening horizon, and the local timing normalization is 44.87 percent. I treat this level effect as secondary, since highly local quantities bend with smoothing and grid geometry. Total variation remains the cleaner way to compare the one- and two-dimensional models.

5.5 What the sensitivity analysis changes

Appendix C reports the details. The default output contains 40 completed rows: 37 model sensitivity or numerical-refinement specifications and three descriptive LODES checks. Across the model specifications, the TV timing normalization ranges from 44.45 to 48.00 percent, and the projected-band timing normalization from 43.12 to 47.49 percent. These ranges summarize sensitivity; they are not confidence intervals. A separate, optional stress routine holds more aggressive low-diffusion, low-curvature, and high-congestion cases. Those cases are excluded from the reported ranges, and when the current solver fails to complete one, the failure is recorded as a failure rather than as a zero effect.

Level effects move more. Changing movement costs, smoothing, diffusion, or the amplitude and width of the project premium shifts opening TV and the local band gap substantially. The timing normalization stays in a narrow range, apart from the expected finite-step movement when the time step changes. The robustness results tell the same story as the 2D exercise: the model is surer about when adjustment happens than about its exact local size.

6 Discussion and limitations

6.1 Implications for timing and empirical design

The main implication concerns how to date a local economic shock. If a project is announced credibly before it operates, opening day need not be the first economically relevant treatment date. Forward-looking households can respond to expected payoffs while nothing on site has changed. In a reduced-form design, that response could show up as a pre-trend if treatment is dated at opening, or it could shrink the estimated opening effect because part of the adjustment has already happened.

None of this means every announcement should be treated as a realized shock. Credibility, uncertainty, financing, permitting, construction progress, and the reversibility of household decisions all matter. The claim here is narrower: once a future regime is credible enough to enter continuation values, an equilibrium model with costly mobility gives a disciplined reason for adjustment to begin early. The sequential-delay experiment adds the complementary point that later news can redirect that path without resetting the economy to where it started.

The same logic applies to local planning. Housing demand, commuting, and service needs may start shifting while a project is still a construction site, or before it is even that. A policy analysis that compares only the pre-announcement equilibrium with the mature post-opening one skips

the transition that residents actually live through. This paper does not compute welfare, so it cannot say whether earlier adjustment is good or bad. It does show that the transition is itself an equilibrium object, and one large enough to deserve attention.

6.2 What the quantitative exercise identifies

The information experiment answers one counterfactual question: how different is the equilibrium path when workers know about the project before it arrives? It does not estimate how much Intel actually changed observed migration. A causal evaluation would need a design that separates the announcement from other local shocks, endogenous wages, housing supply, construction, supplier entry, interest rates, demographic change, and broader general-equilibrium responses. The 2021–2023 LODES patterns in Table 1 show why the distinction matters: local shares do not move monotonically after the announcement.

Several calibrated objects should be read as model discipline, not empirical estimates. The project amplitude uses announced jobs and relative pay, but it remains a reduced-form utility intensity. The recovered baseline fundamental is whatever location-value function makes the smoothed pre-announcement distribution stationary under the maintained discrete model. It soaks up persistent spatial advantages the model does not represent separately. Exact recovery eliminates artificial baseline drift; it does not turn $b(x)$ into an identified amenity index.

Nor is the timing normalization a causal treatment parameter. It is a ratio of model-implied distances: the opening information gap divided by the remaining distance to the common post-project stationary distribution. Its value lies in comparability across specifications. That it stays put while local magnitudes move suggests the paper's strongest quantitative statement is about the timing structure of the modeled transition, not about how many residences relocate near Intel.

6.3 Scope and omitted margins

The framework is deliberately sparse. Housing supply is fixed, so rising demand does not trigger new construction or rents that feed back through a separate housing market. Wages outside the reduced-form project premium are fixed. Firms, suppliers, and contractors never show up. These omissions can cut either way. More housing could absorb concentration with less crowding, while endogenous wages or supplier entry could spread the shock beyond the radial project kernel. The current model does not sign the net effect.

Workers are identical apart from the diffusion term, so the model does not separate the few thousand people who may work at the plant from the rest of the labor force. Real renters and homeowners face different moving costs; workers in different occupations have different exposure to semiconductor jobs; households differ in income, family structure, and attachment to place. The diffusion term keeps the model from collapsing into a deterministic transport problem, but it is no substitute for observed heterogeneity. A natural extension would add a few worker types with type-specific mobility costs or project exposure while keeping the same information experiment.

Local statistics also depend on geometry and regularization. The benchmark ±5-mile measure

is a projected corridor band; the 2D extension reports a true five-mile circle. Smoothing changes how much mass sits near a sharp geographic boundary. That is why the paper reports the raw LODES shares, shows the smoothed baseline, varies the bandwidth, and prefers total variation when comparing dimensions.

Finally, the benchmark treats each public regime as a commonly understood expected opening horizon. Real beliefs are heterogeneous and may spread probability over several completion dates. Modeling that aggregate uncertainty exactly would require a richer state space, because beliefs about a common project date enter continuation values jointly. I prefer a sequence of transparent deterministic regimes to a hazard approximation that would look more sophisticated while solving the stated problem less exactly.

7 Conclusion

A large project can start rearranging geography before it starts production. In the model, the mechanism is simple: a future payoff raises today's continuation value, and costly relocation turns that into gradual movement. The Intel Ohio announcement and its subsequent delays offer a concrete setting in which to put numbers on it.

In the benchmark, the opening-date gap between the forward-looking and surprise economies is 46.71 percent of the remaining long-run total-variation distance. Sequential delays reshape the path but cannot erase the adjustment that earlier information set in motion. A matched two-dimensional extension reduces predicted reallocation by 16.3 percent and leaves the timing normalization essentially where it was.

The empirical and numerical design exists to make that timing statement interpretable. The baseline is the last pre-announcement LODES geography; it is exactly stationary under the generator used in the transition; the surprise economy stays flat before opening; and the 1D and 2D models are compared at matched resolution. None of this makes the exercise causal. It does close several mechanical routes by which a model could manufacture an anticipation effect on its own.

The conclusion, then, is not that some particular number of households must move before a factory opens. It is that, in a forward-looking spatial equilibrium with costly mobility, a substantial share of the distributional response can happen while the plant is still a promise. For projects whose completion dates keep moving, the path of public expectations is part of the shock.

A Proofs and mathematical details

This appendix collects the derivations and finite-dimensional proofs used in the paper. The order follows the computation: the continuous worker problem, population accounting, the discrete control problem, the Markov generator, the inverse baseline, and finally the information and two-dimensional extensions. It does *not* claim a new global existence-and-uniqueness theorem for the continuous mean field game. Standard MFG existence and numerical-analysis results are available in the literature, including Lasry and Lions (2007), Achdou and Capuzzo-Dolcetta (2010), and, for the spatial setting, Bahlali, Boucekkine, and Petit (2026). The paper's computational claims concern the finite-dimensional system it actually solves.

A.1 Discrete system used in the computation

The numerical proofs below concern a finite-state continuous-time Markov process on neighboring spatial cells. For a value vector V, the directional controls are a_i^+ and a_i^-, the controlled generator is $Q(a)$, and the forward population vector is p. The stationary discrete equilibrium for a fixed project regime is the pair (V, p) satisfying

$$\rho V_i = [Q(a(V))V]_i + b_i + r\pi_i - F_i(p) - C_i(a(V)), \tag{A1}$$

$$0 = Q(a(V))^\top p, \qquad \mathbf{1}^\top p = 1, \qquad p \geq 0. \tag{A2}$$

The dynamic implementation uses the same generator in the backward policy problem and the forward population equation. The proofs establish the algebraic properties that implementation needs; they do not assert a global existence theorem for the nonlinear continuous-space MFG.

A.2 From the worker's control problem to the HJB

Proposition A.1 (Optimal continuous drift and HJB)**.** *Suppose V is sufficiently smooth and a worker at (t, x) chooses an adapted drift process a_s to maximize expected discounted utility*

$$V(t,x) = \sup_a \mathbb{E}_{t,x}\left[\int_t^T e^{-\rho(s-t)} u(s, X_s, a_s; m)\,\mathrm{d}s + e^{-\rho(T-t)} G(X_T)\right],$$

subject to the reflected process $\mathrm{d}X_s = a_s\,\mathrm{d}s + \sigma\,\mathrm{d}W_s + \mathrm{d}K_s$, where u is given by equation (2)*. Then, in the interior, the pointwise optimal drift is $a^* = V_x/c$ and the HJB is equation* (4)*; reflection supplies the Neumann boundary condition $V_x = 0$ on $\partial\Omega$.*

Proof. Fix an interior state x and a small increment $\Delta t > 0$. Over an interval before the reflected process reaches the boundary, the regulator K is constant. By the dynamic programming principle, today's optimal value equals expected utility over the next short interval plus the discounted

continuation value at $t + \Delta t$:

$$V(t,x) = \sup_a \mathbb{E}\left[\int_t^{t+\Delta t} e^{-\rho(s-t)} u(s, X_s, a_s; m)\,\mathrm{d}s + e^{-\rho\Delta t} V(t+\Delta t, X_{t+\Delta t})\right].$$

Over a vanishing interval, hold the control fixed at a constant a. The flow integral is

$$u(t,x,a;m)\Delta t + o(\Delta t).$$

Itô's formula gives

$$\mathbb{E}[V(t+\Delta t, X_{t+\Delta t})] = V(t,x) + \Delta t\left(V_t + aV_x + \frac{\sigma^2}{2}V_{xx}\right) + o(\Delta t).$$

Also,

$$e^{-\rho\Delta t} = 1 - \rho\Delta t + o(\Delta t).$$

Substituting the three expansions into the Bellman equation, cancelling $V(t,x)$ from both sides, dividing by Δt, and letting $\Delta t \downarrow 0$ yields

$$0 = \max_a \left\{u(t,x,a;m) + V_t + aV_x + \frac{\sigma^2}{2}V_{xx} - \rho V\right\}.$$

Now insert

$$u = b + r\pi - \kappa(m+\varepsilon)^\gamma - \frac{c}{2}a^2.$$

The terms that depend on a are

$$\phi(a) = aV_x - \frac{c}{2}a^2.$$

Because $c > 0$, $\phi''(a) = -c < 0$, so ϕ is strictly concave and has a unique maximizer. The first-order condition is

$$\phi'(a) = V_x - ca = 0,$$

which gives

$$a^* = \frac{V_x}{c}.$$

Evaluating the maximized expression,

$$a^*V_x - \frac{c}{2}(a^*)^2 = \frac{V_x^2}{c} - \frac{V_x^2}{2c} = \frac{V_x^2}{2c}.$$

Therefore

$$0 = V_t + DV_{xx} + \frac{V_x^2}{2c} + b + r\pi - \kappa(m+\varepsilon)^\gamma - \rho V,$$

where $D = \sigma^2/2$. Rearranging gives equation (4) in the interior. For a reflected diffusion, the generator acts on functions satisfying the Neumann condition $V_x = 0$ at the endpoints; equivalently,

the boundary local-time term in Itô's formula vanishes for functions in the generator domain. This supplies equation (7) for the value problem. □

A.3 From individual motion to the population equation

Proposition A.2 (Kolmogorov forward equation)**.** *Let workers follow the feedback drift $a^*(t,x) = V_x(t,x)/c$. If $m(t,x)$ is the density of worker locations and the boundary has zero outward flux, then m satisfies equation* (6)*.*

Proof. For any smooth test function φ, the infinitesimal generator of the worker diffusion is

$$\mathcal{L}\varphi = a^*\varphi_x + D\varphi_{xx}.$$

Taking expectations and writing them as integrals against the density,

$$\frac{\mathrm{d}}{\mathrm{d}t}\int \varphi(x)m(t,x)\,\mathrm{d}x = \int \left(a^*\varphi_x + D\varphi_{xx}\right) m\,\mathrm{d}x.$$

Integrate the drift term by parts once:

$$\int a^* m\varphi_x\,\mathrm{d}x = \left[a^* m\varphi\right]_{\partial\Omega} - \int \varphi\,\partial_x(a^* m)\,\mathrm{d}x.$$

Integrate the diffusion term by parts twice:

$$\int Dm\varphi_{xx}\,\mathrm{d}x = \left[Dm\varphi_x - Dm_x\varphi\right]_{\partial\Omega} + \int \varphi Dm_{xx}\,\mathrm{d}x.$$

Reflecting boundaries impose zero net population flux,

$$a^* m - Dm_x = 0 \quad \text{on } \partial\Omega.$$

The generator domain for the reflected diffusion uses test functions with $\varphi_x = 0$ at the endpoints, so the combined boundary contribution

$$\left[(a^* m - Dm_x)\varphi + Dm\varphi_x\right]_{\partial\Omega}$$

vanishes. Hence

$$\frac{\mathrm{d}}{\mathrm{d}t}\int \varphi m\,\mathrm{d}x = \int \varphi\left[-\partial_x(a^* m) + Dm_{xx}\right]\,\mathrm{d}x.$$

Because this identity holds for every smooth test function φ, the density satisfies

$$\partial_t m = Dm_{xx} - \partial_x(a^* m).$$

Substituting $a^* = V_x/c$ gives equation (6). □

A.4 Optimal directional controls in the finite-state model

Proposition A.3 (Optimal directional controls)**.** *At an interior cell i, let $a_i^+, a_i^- \geq 0$ denote right and left movement intensities. Suppose the contribution of those controls to the discrete HJB is*

$$\frac{a_i^+}{\Delta x}(V_{i+1} - V_i) + \frac{a_i^-}{\Delta x}(V_{i-1} - V_i) - \frac{c}{2}\left[(a_i^+)^2 + (a_i^-)^2\right].$$

Then the unique maximizers are

$$a_i^+ = \left[\frac{V_{i+1} - V_i}{c\Delta x}\right]_+, \qquad a_i^- = \left[\frac{V_{i-1} - V_i}{c\Delta x}\right]_+.$$

At a reflecting boundary, the outward control is omitted.

Proof. The objective is the sum of a function of a_i^+ and a function of a_i^-, so the two controls can be maximized separately.

For the right control, define

$$\phi_+(a) = \frac{a}{\Delta x}(V_{i+1} - V_i) - \frac{c}{2}a^2, \qquad a \geq 0.$$

Its second derivative is $\phi_+''(a) = -c < 0$, so it is strictly concave. The unconstrained first-order condition is

$$\phi_+'(a) = \frac{V_{i+1} - V_i}{\Delta x} - ca = 0,$$

which gives

$$\widehat{a}_i^+ = \frac{V_{i+1} - V_i}{c\Delta x}.$$

If $\widehat{a}_i^+ \geq 0$, strict concavity makes it the unique constrained maximizer. If $\widehat{a}_i^+ < 0$, then $\phi_+'(0) < 0$, the concave function is decreasing at the left endpoint of the feasible set, and the unique constrained maximizer is $a_i^+ = 0$. Both cases are written compactly as

$$a_i^+ = \left[\frac{V_{i+1} - V_i}{c\Delta x}\right]_+.$$

The same argument with $V_{i-1} - V_i$ gives the formula for a_i^-. At a boundary there is no cell outside the domain, so the outward transition is removed rather than optimized. □

A.5 Generator properties, mass conservation, and positivity

Proposition A.4 (Generator properties)**.** *Let Q have nonnegative off-diagonal entries and diagonal entries*

$$q_{ii} = -\sum_{j \neq i} q_{ij}.$$

Then $Q\mathbf{1} = 0$. The continuous-time forward equation $\dot{p} = Q^\top p$ preserves total mass.

Proof. Fix a row i. By construction,

$$\sum_j q_{ij} = q_{ii} + \sum_{j \neq i} q_{ij} = -\sum_{j \neq i} q_{ij} + \sum_{j \neq i} q_{ij} = 0.$$

This holds for every row, so $Q\mathbf{1} = 0$.

Now let $M(t) = \mathbf{1}^\top p(t)$ denote total probability mass. Differentiate:

$$\dot{M}(t) = \mathbf{1}^\top \dot{p}(t) = \mathbf{1}^\top Q^\top p(t) = (Q\mathbf{1})^\top p(t) = 0.$$

Hence $M(t)$ is constant. If $\mathbf{1}^\top p(0) = 1$, then $\mathbf{1}^\top p(t) = 1$ for all t. □

Proposition A.5 (Positivity of the continuous-time Markov evolution)**.** *If $p(0) \geq 0$ componentwise and Q is a finite-state Markov generator, then $p(t) = e^{tQ^\top} p(0) \geq 0$ for every $t \geq 0$.*

Proof. Choose $\lambda > \max_i(-q_{ii})$ (and hence $\lambda > 0$). Define

$$P = I + \frac{Q^\top}{\lambda}.$$

For $i \neq j$, the entry $P_{ij} = q_{ji}/\lambda$ is nonnegative because the off-diagonal entries of Q are nonnegative. For a diagonal entry,

$$P_{ii} = 1 + \frac{q_{ii}}{\lambda} \geq 0$$

by the choice of λ. Also, because $Q\mathbf{1} = 0$, every column of $Q^\top$ sums to zero, so every column of P sums to one.

Rewrite

$$Q^\top = \lambda(P - I).$$

Because P commutes with I,

$$e^{tQ^\top} = e^{-\lambda t} e^{\lambda t P} = e^{-\lambda t} \sum_{k=0}^{\infty} \frac{(\lambda t)^k}{k!} P^k.$$

Each P^k is componentwise nonnegative, and every scalar coefficient in the series is nonnegative, so $e^{tQ^\top}$ is componentwise nonnegative. Multiplying it by $p(0) \geq 0$ yields $p(t) \geq 0$. □

Corollary A.1 (Implicit forward step preserves mass and positivity)**.** *The numerical forward step*

$$(I - \Delta t Q^\top) p^{n+1} = p^n, \qquad \Delta t > 0,$$

preserves total mass. If $p^n \geq 0$, then $p^{n+1} \geq 0$.

Proof. Premultiply the discrete equation by $\mathbf{1}^\top$:

$$\mathbf{1}^\top p^{n+1} - \Delta t\, \mathbf{1}^\top Q^\top p^{n+1} = \mathbf{1}^\top p^n.$$

Since $\mathbf{1}^\top Q^\top = (Q\mathbf{1})^\top = 0$, total mass is unchanged.

For positivity, use the resolvent identity

$$(I - \Delta t Q^\top)^{-1} = \int_0^\infty e^{-s} e^{s\Delta t Q^\top} \, \mathrm{d}s.$$

The integral is well defined because a finite-state Markov generator has spectral bound at most zero. By Proposition A.5, $e^{s\Delta t Q^\top}$ is nonnegative for every $s \geq 0$, so the integrand is a nonnegative matrix and so is the inverse. Hence

$$p^{n+1} = (I - \Delta t Q^\top)^{-1} p^n \geq 0$$

whenever $p^n \geq 0$. □

A.6 Well-posed policy evaluation

Proposition A.6 (Invertibility of the policy-evaluation systems)**.** *For any finite-state generator Q and $\rho > 0$, the matrix $\rho I - Q$ is nonsingular. For any $\Delta t > 0$, the dynamic HJB matrix $(1/\Delta t + \rho)I - Q$ is also nonsingular.*

Proof. Every eigenvalue λ of a Markov generator has nonpositive real part. Gershgorin's theorem gives a direct argument. Row i has center $q_{ii} \leq 0$ and radius

$$\sum_{j \neq i} |q_{ij}| = \sum_{j \neq i} q_{ij} = -q_{ii}.$$

Every Gershgorin disc therefore lies in the closed left half-plane and touches the imaginary axis only at zero, so $\Re(\lambda) \leq 0$ for every eigenvalue of Q.

If $\rho I - Q$ were singular, ρ would be an eigenvalue of Q. But $\rho > 0$ has positive real part, a contradiction. Hence $\rho I - Q$ is nonsingular.

The same argument applies to $(1/\Delta t + \rho)I - Q$, because $1/\Delta t + \rho > 0$. □

A.7 Exact one-dimensional inverse baseline

Proposition A.7 (Detailed-balance inverse in one dimension)**.** *Let $p_i^0 > 0$ be a target probability distribution on a one-dimensional nearest-neighbor grid. Let $D > 0$. Define value differences across each face by*

$$V_{i+1}^0 - V_i^0 = cD \frac{p_{i+1}^0 - p_i^0}{\min\{p_i^0, p_{i+1}^0\}}. \tag{A3}$$

Choose any additive constant for V^0, and construct the optimal directional controls from Proposition A.3. Then every adjacent pair satisfies detailed balance and

$$Q(a(V^0))^\top p^0 = 0.$$

Proof. Consider the face between cells i and $i+1$. Write

$$d = \frac{D}{\Delta x^2}$$

for the symmetric diffusion transition rate.

Case 1: $p^0_{i+1} \geq p^0_i$. Equation (A3) gives

$$V^0_{i+1} - V^0_i = cD\frac{p^0_{i+1} - p^0_i}{p^0_i} \geq 0.$$

So the rightward control from cell i is active and the leftward control from cell $i+1$ is zero:

$$a^+_i = \frac{V^0_{i+1} - V^0_i}{c\Delta x} = \frac{D(p^0_{i+1} - p^0_i)}{p^0_i \Delta x}, \qquad a^-_{i+1} = 0.$$

The total transition rate from i to $i+1$ is

$$q_{i,i+1} = d + \frac{a^+_i}{\Delta x} = \frac{D}{\Delta x^2} + \frac{D(p^0_{i+1} - p^0_i)}{p^0_i \Delta x^2} = \frac{D}{\Delta x^2}\frac{p^0_{i+1}}{p^0_i}.$$

The reverse rate is diffusion alone:

$$q_{i+1,i} = d = \frac{D}{\Delta x^2}.$$

Multiplying by the corresponding probabilities gives

$$p^0_i q_{i,i+1} = \frac{D}{\Delta x^2} p^0_{i+1} = p^0_{i+1} q_{i+1,i},$$

so net probability flux across the face is zero.

Case 2: $p^0_{i+1} < p^0_i$. Equation (A3) is now negative, with denominator p^0_{i+1}. The rightward control from i is zero, and the leftward control from $i+1$ is

$$a^-_{i+1} = \frac{V^0_i - V^0_{i+1}}{c\Delta x} = \frac{D(p^0_i - p^0_{i+1})}{p^0_{i+1} \Delta x}.$$

Hence

$$q_{i,i+1} = \frac{D}{\Delta x^2},$$

and

$$q_{i+1,i} = \frac{D}{\Delta x^2} + \frac{D(p^0_i - p^0_{i+1})}{p^0_{i+1} \Delta x^2} = \frac{D}{\Delta x^2}\frac{p^0_i}{p^0_{i+1}}.$$

Again,

$$p^0_i q_{i,i+1} = p^0_{i+1} q_{i+1,i}.$$

Detailed balance therefore holds across every interior face. Reflecting boundaries have no outward

face. Summing the zero net fluxes into and out of each cell gives

$$Q(a(V^0))^\top p^0 = 0.$$

The additive constant in V^0 affects no value difference, control, or generator entry. □

Proposition A.8 (Recovered baseline fundamental)**.** *Let p^0 and V^0 satisfy Proposition A.7. Define*

$$F_i(p^0) = \kappa \left(p_i^0 / \Delta x + \varepsilon \right)^\gamma$$

and let $C_i(a)$ be the directional quadratic control cost. If

$$b_i = \rho V_i^0 - [Q(a(V^0))V^0]_i + F_i(p^0) + C_i(a(V^0)), \tag{A4}$$

then (V^0, p^0) satisfies both the stationary no-project discrete HJB and the stationary KFE exactly.

Proof. The KFE condition follows directly from Proposition A.7:

$$Q(a(V^0))^\top p^0 = 0.$$

For the HJB, the stationary no-project equation under the optimal directional controls is

$$\rho V_i = [Q(a(V))V]_i + b_i - F_i(p) - C_i(a(V)).$$

Set $V = V^0$ and $p = p^0$, and substitute b_i from equation (A4):

$$\begin{aligned} [Q(a(V^0))V^0]_i + b_i - F_i(p^0) - C_i(a(V^0)) &= [QV^0]_i + \rho V_i^0 - [QV^0]_i + F_i(p^0) + C_i(a) \\ &\quad - F_i(p^0) - C_i(a) \\ &= \rho V_i^0. \end{aligned}$$

The HJB holds cell by cell. Both stationary equilibrium conditions therefore hold for exactly the generator used in the transition model. □

Corollary A.2 (Constant-shift normalization)**.** *Suppose (V, p, b) satisfies the stationary discrete HJB and KFE for a fixed project regime. For any constant δ, define*

$$\widetilde{b} = b + \delta \mathbf{1}, \qquad \widetilde{V} = V + \frac{\delta}{\rho} \mathbf{1}.$$

Then $(\widetilde{V}, p, \widetilde{b})$ has the same optimal controls and the same generator, and it satisfies the same stationary equilibrium conditions with the shifted fundamental.

Proof. Every directional control depends on a neighboring value difference. Adding the same

constant δ/ρ to every component of V leaves all differences unchanged, so

$$a(\widetilde{V}) = a(V) \quad \text{and} \quad Q(a(\widetilde{V})) = Q(a(V)).$$

Because a Markov generator annihilates constants,

$$Q\widetilde{V} = QV + \frac{\delta}{\rho} Q\mathbf{1} = QV.$$

The left side of the shifted HJB is

$$\rho\widetilde{V} = \rho V + \delta\mathbf{1}.$$

The right side gains exactly $\delta\mathbf{1}$ through $\widetilde{b} = b + \delta\mathbf{1}$, while QV, congestion, and control cost are unchanged. The HJB therefore still holds. The KFE is untouched because the generator and p are untouched. This justifies picking a convenient level for the recovered baseline fundamental without changing behavior. □

Corollary A.3 (A constant no-project equilibrium path)**.** *If the no-project economy begins at the stationary pair (V^0, p^0) from Proposition A.8, then the constant population path $p(t) = p^0$ together with the stationary control $a(V^0)$ is a no-project equilibrium path.*

Proof. Under the stationary control, Proposition A.8 gives

$$Q(a(V^0))^\top p^0 = 0.$$

The constant function $p(t) = p^0$ therefore satisfies the forward equation, since its time derivative is zero. By the same proposition, V^0 satisfies the no-project HJB, so the control remains individually optimal against the constant population path. The pair is a mean field equilibrium path.

The corollary establishes that this constant path exists; it does not claim that every nonlinear MFG path is unique. In the computation, the forward–backward fixed-point solver recovers the constant path with no-project TV drift of 1.24×10^{-14} at the opening horizon. □

A.8 The information decomposition

Proposition A.9 (Information-channel decomposition)**.** *For $t < T_0$, the forward-looking and surprise-opening economies share the same exogenous contemporaneous project regime, baseline fundamental, movement technology, diffusion, and congestion function. Their pre-opening paths diverge only because of different continuation information. Endogenous congestion can subsequently amplify or damp that information-induced divergence.*

Proof. For every $t < T_0$, both experiments set the project regime indicator to zero:

$$r^{\mathrm{FL}}(t) = r^{\mathrm{UA}}(t) = 0.$$

Neither economy receives the contemporaneous project premium $\pi(x)$ before opening. They also share $b(x)$, c, σ, κ, γ, and the initial population distribution.

The backward problems differ because the forward-looking economy knows that $r(t)$ will switch to one at T_0, while the surprise economy leaves that future regime out of its pre-opening continuation problem. Value gradients can therefore differ before T_0, and with them the optimal movement controls.

Once controls differ, population distributions can differ. Because congestion depends on the endogenous population, current congestion terms may then differ too. That later congestion difference is not an independent contemporaneous treatment; it is an equilibrium consequence of the initial information perturbation. The full pre-opening path difference is therefore the equilibrium information-channel decomposition used in the paper. □

A.9 Sequential public news

Proposition A.10 (Population continuity at a revision date)**.** *Suppose public information changes at time τ, altering the continuation problem and therefore the control generator from $Q^-(t)$ to $Q^+(t)$. If the forward equation contains no impulse at τ and the generators are locally bounded, then the population distribution is continuous at τ:*

$$p(\tau^+) = p(\tau^-).$$

The value function may change because the information set changes.

Proof. For $t \neq \tau$, the finite-dimensional population state satisfies

$$\dot{p}(t) = Q(t)^\top p(t).$$

Integrate over $[\tau - \eta, \tau + \eta]$:

$$p(\tau + \eta) - p(\tau - \eta) = \int_{\tau-\eta}^{\tau+\eta} Q(s)^\top p(s)\,\mathrm{d}s.$$

Because the state space is finite, $p(s)$ is bounded. By assumption, $Q(s)$ is locally bounded on either side of the news date. There is therefore a finite constant M such that

$$\|Q(s)^\top p(s)\| \leq M$$

for sufficiently small η. Hence

$$\|p(\tau + \eta) - p(\tau - \eta)\| \leq 2M\eta.$$

Letting $\eta \downarrow 0$ gives

$$p(\tau^+) = p(\tau^-).$$

No analogous requirement forces the value function to be continuous across an unanticipated information change, because the conditional distribution of future regimes shifts at the news date. The implementation therefore inherits the population exactly and re-solves the backward problem under the new continuation information. The reported inheritance error is zero in both revision segments. □

A.10 Dimensional congestion scaling

Proposition A.11 (Dimensional congestion rescaling)**.** *Let the normalized two-dimensional domain have perpendicular width W. If a one-dimensional density $m_{1D}(x)$ is spread uniformly across the perpendicular coordinate, so that*

$$m_{2D}(x, y) = \frac{m_{1D}(x)}{W},$$

then choosing

$$\kappa_{2D} = \kappa_{1D} W^{\gamma}$$

makes the unregularized congestion term κm^{γ} identical in one and two dimensions under this uniform benchmark. More generally, the regularized term $\kappa(m + \varepsilon)^{\gamma}$ is exactly preserved if $\varepsilon_{2D} = \varepsilon_{1D}/W$.

Proof. For the unregularized term,

$$\begin{aligned} \kappa_{2D} m_{2D}^{\gamma} &= \kappa_{1D} W^{\gamma} \left(\frac{m_{1D}}{W} \right)^{\gamma} \\ &= \kappa_{1D} m_{1D}^{\gamma}. \end{aligned}$$

For the regularized term, also set $\varepsilon_{2D} = \varepsilon_{1D}/W$. Then

$$\begin{aligned} \kappa_{2D}(m_{2D} + \varepsilon_{2D})^{\gamma} &= \kappa_{1D} W^{\gamma} \left(\frac{m_{1D} + \varepsilon_{1D}}{W} \right)^{\gamma} \\ &= \kappa_{1D}(m_{1D} + \varepsilon_{1D})^{\gamma}. \end{aligned}$$

The scaling thus removes the mechanical change in congestion that would otherwise come purely from measuring the same population per unit area instead of per unit length. □

A.11 Two-dimensional discrete inverse conditional on the active set

Proposition A.12 (Two-dimensional graph inverse)**.** *Fix a strictly positive target probability vector p on the two-dimensional nearest-neighbor grid and fix, for every undirected edge (u, v), which endpoint has the higher value. Under that fixed active-edge ordering, the stationary KFE is linear in the unknown value vector V. If V solves the weighted graph-Poisson system used in Appendix B, and the implied value ordering agrees with the fixed active set, then the controlled generator satisfies*

$$Q(V)^{\top} p = 0.$$

Recovering b from the discrete stationary HJB then makes (V, p) an exact stationary discrete equilibrium.

Proof. Consider a neighboring pair (u, v) with edge length h. Suppose first that the fixed active set says $V_v \geq V_u$. Then only the control from u toward v is active on that edge, and the directed rate contribution generated by the control is proportional to

$$\frac{V_v - V_u}{ch^2}.$$

Multiplying by the probability at the origin cell shows that the control contribution to net flux is linear in V with weight p_u/h^2. If instead $V_u > V_v$, only the reverse control is active and the weight is p_v/h^2. Once the sign of every neighboring value difference is fixed, then, all control-flow terms in $Q(V)^\top p$ are linear in V.

The diffusion contribution is already linear: it is a graph Laplacian applied to p, scaled by D. Moving the diffusion terms to the right-hand side gives a weighted graph-Poisson system of the form

$$A_{\mathcal{A}} V = cD\, L_h p,$$

where the active set $\mathcal{A}$ determines whether an edge receives weight p_u/h^2 or p_v/h^2. The Neumann structure determines V only up to an additive constant, so the implementation replaces one redundant node equation with the probability-weighted gauge

$$p^\top V = 0.$$

The omitted KFE row is indeed redundant. Because every controlled generator satisfies $Q\mathbf{1} = 0$,

$$\mathbf{1}^\top Q(V)^\top p = (Q(V)\mathbf{1})^\top p = 0,$$

so the components of the KFE residual sum to zero. If the other $n - 1$ node equations hold, the omitted component must be zero as well. Solving the gauged system therefore makes diffusion and directed control fluxes cancel at every node, which is exactly

$$Q(V)^\top p = 0.$$

The active-set algorithm then recomputes the signs of all neighboring value differences. If they agree with the active set used to build $A_{\mathcal{A}}$, the linearization is self-consistent and the KFE identity holds for the actual controlled generator.

Finally, define

$$b = \rho V - Q(V)V + F(p) + C(a(V)).$$

Substituting into the stationary discrete HJB yields an identity, by the same algebra as in Proposition A.8. Thus (V, p) satisfies both stationary equations. The preferred 2D computation verifies a KFE residual of 1.09×10^{-14} and an HJB residual of 2.36×10^{-12}. □

B Two-dimensional geographic extension

The benchmark collapses Central Ohio onto the Intel–Columbus axis. This appendix asks whether that shortcut drives the timing result. The 2D model keeps both the along-axis and perpendicular coordinates on the same retained 90×60 mile rectangle.

B.1 Geography and equations

Let x denote distance along the Intel–Columbus axis divided by 45 miles, and y perpendicular distance divided by 45 miles. Then

$$(x, y) \in \Omega = [-1, 1] \times [-2/3, 2/3],$$

with Intel at $(0, 0)$. The preferred specification uses a 121×81 finite-volume grid and $\Delta t = 0.10$, giving cells of roughly 0.74×0.74 miles.

The continuous equations are

$$-\partial_t V + \rho V = D\Delta V + \frac{\|\nabla V\|^2}{2c} + b(x, y) + r(t)\pi(x, y) - \kappa_{2D}(m + \varepsilon)^\gamma, \tag{B1}$$

$$\partial_t m = D\Delta m - \nabla \cdot \left(m \frac{\nabla V}{c} \right), \tag{B2}$$

with reflecting boundaries,

$$\nabla V \cdot n = 0, \qquad \left(D\nabla m - m\frac{\nabla V}{c} \right) \cdot n = 0 \quad \text{on } \partial\Omega.$$

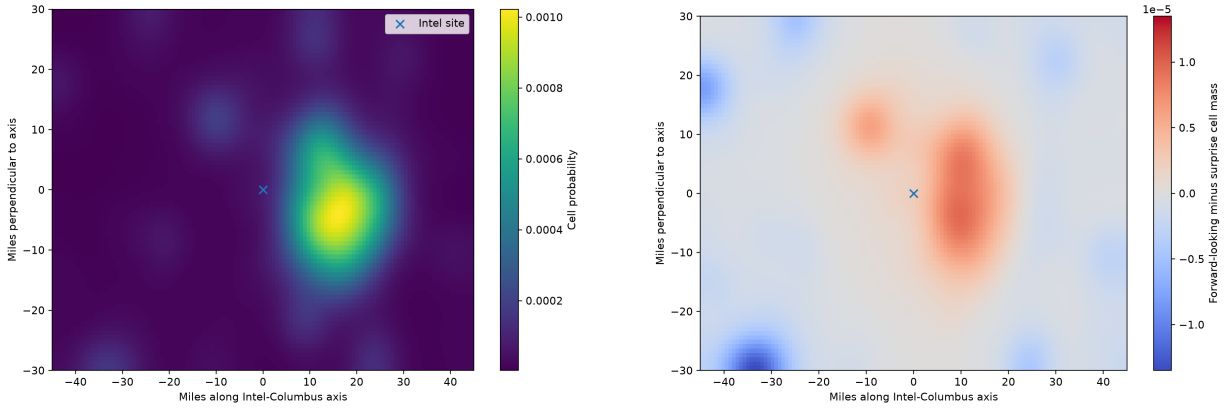
The project kernel is radial,

$$\pi(x, y) = \omega \exp\left[-\frac{x^2 + y^2}{2h^2} \right], \qquad h = 20/45, \tag{B3}$$

and the local outcome is mass inside the true geographic five-mile circle.

Density carries different physical units in one and two dimensions. To avoid conjuring a dimensional effect out of units alone, I rescale the 2D congestion coefficient. If the same one-dimensional density were spread uniformly over perpendicular width $W = 4/3$, then $m_{2D} = m_{1D}/W$, and setting $\kappa_{2D} = \kappa_{1D} W^\gamma$ preserves the unregularized congestion term κm^γ exactly. With $\gamma = 1$, $\kappa_{2D} = 0.40$. The code keeps the same numerical floor $\varepsilon = 10^{-10}$ in both dimensions rather than rescaling it. In the preferred $\gamma = 1$ case this adds only a spatially constant term of order 10^{-11} under the uniform benchmark, which shifts value levels but not controls and is absorbed by the inverse-baseline normalization. Proposition A.11 states the exact scaling result.

The same 3.5-mile smoothing rule applies in both dimensions, and the baseline is again recovered with the exact discrete generator. In the preferred 2D run, the target-to-stationary TV error is 5.87×10^{-15} and the no-project TV drift at the opening horizon is 1.56×10^{-14}.



(a) Smoothed 2021 baseline.

(b) Forward-looking minus surprise mass at opening.

Figure B1: Two-dimensional geography in the preferred 121×81 specification. The marker denotes the Intel site.

Notes: Panel (a) is the smoothed LODES baseline. Panel (b) shows the cell-by-cell difference between the forward-looking and surprise economies at the original expected-opening horizon. Positive and negative regions integrate to zero because total population mass is conserved.

B.2 Matched dimensional comparison and refinement

Table B1 reports the matched comparison and the refinement sequence. The comparison that matters is the 121-cell 1D model against the 121×81 2D model at $\Delta t = 0.10$. Adding the second direction of movement lowers opening TV by 16.3 percent but moves the timing normalization by only 0.09 percentage points.

Table B1: Two-dimensional comparison and numerical refinement

Specification	Opening TV	TV timing	5-mile gap (pp)	5-mile timing
Matched 1D, 121 cells, $\Delta t = .10$	0.009478	46.80%	—	—
2D 61×41, $\Delta t = .20$	0.007290	45.04%	0.0341	41.83%
2D 91×61, $\Delta t = .20$	0.007548	45.18%	0.0360	41.82%
2D 121×81, $\Delta t = .20$	0.007659	45.08%	0.0373	41.85%
2D 91×61, $\Delta t = .10$	0.007825	46.83%	0.0386	44.82%
2D 91×61, $\Delta t = .05$	0.007984	47.78%	0.0402	46.65%
2D 121×81, $\Delta t = .10$	0.007935	46.71%	0.0400	44.87%

Notes: The 5-mile statistic is a true geographic circle in 2D and is therefore not directly comparable to the projected 1D band. Spatial refinement is tight for the TV timing statistic. Time-step refinement moves the level somewhat more; the matched 1D and 2D timing ratios move together, which is the relevant dimensional check.

The sequential information experiment survives the move to two dimensions. The 2024 and 2025 segments inherit the distribution exactly, with fixed-point errors of 5.77×10^{-7} and 5.23×10^{-7}. At the original expected-opening horizon, the revised-information path differs from the original-information path by 0.0382 percentage points inside the true five-mile circle and 0.007258 in total variation.

The preferred 2D run passes every required numerical check. Baseline HJB and KFE residuals

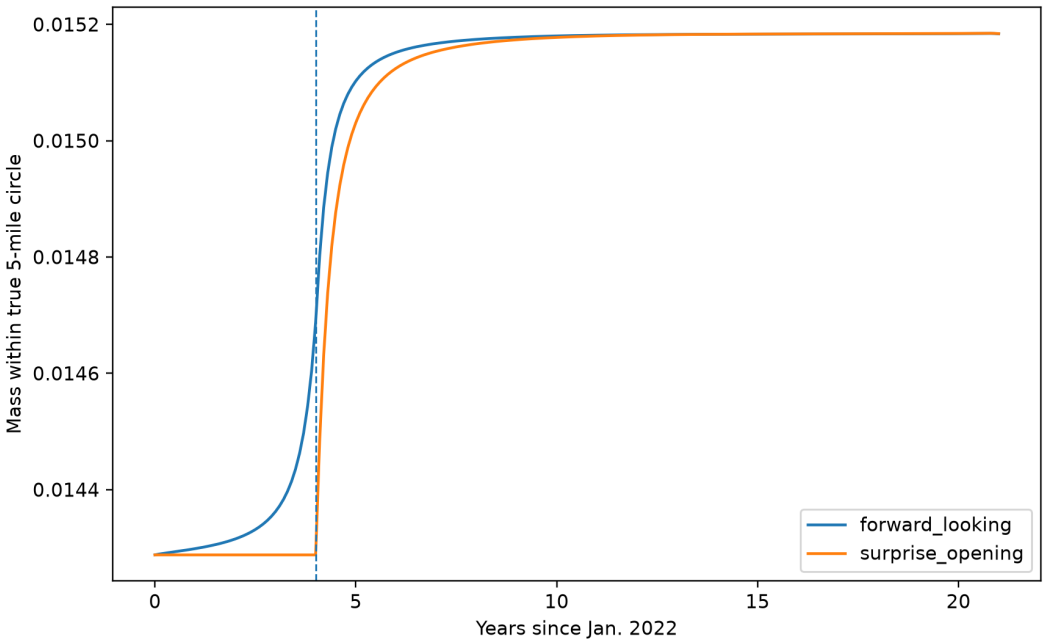


Figure B2: Forward-looking and surprise paths for mass inside the true five-mile circle in the preferred 2D specification.

are 2.36×10^{-12} and 1.09×10^{-14}; the post-project KFE residual is 1.80×10^{-7}; forward-looking and surprise fixed-point errors are below 7.3×10^{-7}; and the maximum mass error is 1.37×10^{-14}.

C Robustness and numerical validation

The robustness exercises are kept apart from the main results on purpose. Their job is to find out which conclusions lean on maintained modeling or numerical choices. The reported ranges condition on successful numerical solutions and are not confidence intervals.

The default robustness output contains 40 completed rows: 37 model sensitivity or numerical-refinement specifications and three descriptive LODES checks. Table C1 summarizes the 37 model specifications rather than listing every row. A separate, optional stress routine contains $\sigma^2 = 0.10$, $\gamma = 0.5$, and $\kappa = 2\kappa_0$. Those harsher cases stay outside the reported ranges unless they are explicitly run and converge, and the code records any numerical failure as a failure rather than as a zero effect. The clearest pattern in the reported suite is that the amount of reallocation moves much more than the TV timing normalization.

Table C1: Robustness summary

Variation	Successful cases	Opening TV range	TV timing range
Structural parameters	8/8	0.006307–0.011406	45.71–47.26%
Smoothing bandwidth	4/4	0.008397–0.011949	46.53–47.50%
Project amplitude and width	9/9	0.003057–0.015132	46.11–47.13%
Spatial grid	4/4	0.008760–0.009424	46.71–46.81%
Time step	3/3	0.008640–0.009330	44.45–48.00%
Terminal tail	3/3	0.009080–0.009080	46.71–46.71%
Initial opening horizon	3/3	0.009052–0.009097	46.57–46.80%
Sequential revision horizons	3/3	0.008309–0.008490[a]	—

Notes: Structural cases vary relocation cost, diffusion, discounting, congestion curvature, and congestion scale. Smoothing varies from 2 to 7.5 miles; project amplitude varies from 0.5 to 1.5 times the benchmark and width from 10 to 30 miles; N_x ranges from 41 to 101; Δt ranges from 0.05 to 0.20 years. [a]For sequential revisions, the reported range is TV between the revised and original-information paths at the initial expected-opening horizon rather than the benchmark forward-looking/surprise TV.

Time discretization is the most visible numerical sensitivity. At $\Delta t = 0.20, 0.10, 0.05$, the 1D TV timing normalization is 44.45, 46.71, and 48.00 percent. The paper therefore reports finite-step sensitivity rather than claiming exact convergence in time. Spatial refinement is tighter: across N_x from 41 to 101, the timing normalization moves by less than a tenth of a percentage point once the grid reaches the 61-cell benchmark.

Highly local level effects are touchier. Widening the smoothing bandwidth from 2 to 7.5 miles raises the projected-band opening gap, and changing the project amplitude or width moves opening TV substantially. These exercises are why the paper emphasizes timing, and total variation as the common cross-dimensional statistic.

C.1 Benchmark numerical diagnostics

The numerical errors sit orders of magnitude below the economic distances reported in the paper. In particular, the baseline target is stationary under the implemented generator and the no-project

Table C2: Numerical diagnostics for the preferred 1D and 2D specifications

Diagnostic	1D benchmark	2D preferred
Baseline HJB residual	6.11×10^{-13}	2.36×10^{-12}
Baseline KFE residual	9.29×10^{-14}	1.09×10^{-14}
Post-project HJB residual	6.29×10^{-13}	3.20×10^{-12}
Post-project KFE residual	7.62×10^{-8}	1.80×10^{-7}
Forward-looking fixed-point error	6.96×10^{-7}	6.86×10^{-7}
Surprise fixed-point error	6.98×10^{-7}	7.29×10^{-7}
Baseline target–stationary TV	7.65×10^{-15}	5.87×10^{-15}
No-project TV at opening	1.24×10^{-14}	1.56×10^{-14}
Maximum mass error	4.60×10^{-14}	1.37×10^{-14}
Terminal TV to post equilibrium	2.16×10^{-6}	1.31×10^{-4}
All required checks	Passed	Passed

path is essentially flat. That is the diagnostic needed to rule out phantom pre-opening migration from an inconsistent baseline calibration.

D Data construction and measurement

D.1 LODES residence mass

The empirical baseline uses LODES residence-area counts for primary jobs. Residence block centroids are converted to coordinates relative to the Intel site and the Intel–Columbus axis. The along-axis coordinate is the orthogonal projection onto that axis; the perpendicular coordinate is the signed residual distance. The 1D sample keeps residences within 30 miles of the axis and within 45 miles along it. The 2D model keeps the same rectangle and both coordinates.

The 2021 distribution is the calibration baseline because it predates the announcement. The 2022 and 2023 LODES extracts are descriptive checks only and play no role in estimating a causal effect. Their retained counts and raw local shares appear in Table 1 in the main text, so the empirical context is on the table before any model result is.

D.2 Smoothing, local measurement, and boundaries

The baseline projected mass is smoothed with a 3.5-mile Gaussian bandwidth. The implementation uses even reflection at the spatial boundaries, consistent with the reflecting-boundary PDE, so density does not leak out at the edges. The smoothed probabilities are then normalized to sum to one.

The benchmark local statistic uses fractional cell overlap with the projected band $|x| \leq 5/45$, rather than counting a cell only when its center falls inside. This keeps the reported local share from shifting mechanically when the grid is refined. The true five-mile circle in two dimensions follows the same principle: each boundary cell is weighted by the fraction of its area inside the circle. A uniform-density area test has error 6.68×10^{-13} on the preferred 2D grid.

D.3 QCEW scaling and information timing

QCEW aggregation level 70 (County, Total Covered) supplies the Licking County employment and pay benchmark in equation (10). The model does not claim that QCEW covered employment equals the LODES primary-job residence universe. The two datasets meet only to make the reduced-form project scale explicit and reproducible.

The initial shock uses the 3,000 direct jobs and $135,000 average salary reported in January 2022. Later payroll or employment commitments are not back-cast into the 2022 information set. Each public revision date is placed at the first time-grid point on or after the event, so no revision ever affects behavior before it is public.

D.4 Replication structure

The accompanying scripts separate the benchmark model, the robustness exercises, and the 2D extension. The benchmark script retrieves and builds the public data inputs, calibrates the exact discrete baseline, solves the forward-looking, surprise, no-project, and sequential-revision paths, and

writes the numerical diagnostics reported in the paper. The robustness script reruns the maintained default alternatives and records any failure explicitly; the optional stress routine sits outside the reported default ranges. The 2D script repeats the inverse calibration and information experiment on the two-dimensional grid and computes the matched 1D comparator.

References

Achdou, Yves, and Italo Capuzzo-Dolcetta. 2010. "Mean Field Games: Numerical Methods." *SIAM Journal on Numerical Analysis* 48(3): 1136–1162.

Allen, Treb, and Costas Arkolakis. 2014. "Trade and the Topography of the Spatial Economy." *Quarterly Journal of Economics* 129(3): 1085–1140.

Bahlali, Mohamed, Raouf Boucekkine, and Quentin Petit. 2026. "When Lasry-Lions Meet Krugman: A Mean-Field Game Theory of Spatial Dynamics." *Journal of Economic Theory* 236: 106217.

Bilal, Adrien, and Esteban Rossi-Hansberg. 2023. "Anticipating Climate Change Across the United States." NBER Working Paper 31323, revised October 2025.

Bureau of Labor Statistics. 2026. "Quarterly Census of Employment and Wages: Aggregation Level Titles." https://www.bls.gov/cew/classifications/aggregation/agg-level-titles.htm.

Cárdenas, Jeisson, Juan Miguel Gallego, and María Alejandra Urrutia. 2023. "Announcement of the First Metro Line and Its Impact on Housing Prices in Bogotá." *Case Studies on Transport Policy* 11: 100941.

U.S. Census Bureau. 2026. "LEHD Origin-Destination Employment Statistics Data." https://lehd.ces.census.gov/data/.

Desmet, Klaus, and Esteban Rossi-Hansberg. 2014. "Spatial Development." *American Economic Review* 104(4): 1211–1243.

Huang, Minyi, Roland P. Malhamé, and Peter E. Caines. 2006. "Large Population Stochastic Dynamic Games: Closed-Loop McKean-Vlasov Systems and the Nash Certainty Equivalence Principle." *Communications in Information and Systems* 6(3): 221–252.

Hyun, Dongwoo, and Stanimira Milcheva. 2019. "Spatio-Temporal Effects of an Urban Development Announcement and Its Cancellation on House Prices: A Quasi-Natural Experiment." *Journal of Housing Economics* 43: 23–36.

Intel. 2022. "Intel Announces Next US Site with Landmark Investment in Ohio." January 21. https://www.intc.com/news-events/press-releases/detail/1521/intel-announces-next-us-site-with-landmark-investment-in.

Intel. 2025. "Ohio One Construction Timeline Update." February 28. https://newsroom.intel.com/corporate/ohio-one-construction-timeline-update.

JobsOhio. 2022. "Intel to Bring Advanced Semiconductor Manufacturing to Ohio." January 21. https://www.jobsohio.com/newsroom/news-press/intel-chooses-ohio-for-chip-manufacturing.

Kline, Patrick, and Enrico Moretti. 2014. "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs." *Annual Review of Economics* 6: 629–662.

Lasry, Jean-Michel, and Pierre-Louis Lions. 2007. "Mean Field Games." *Japanese Journal of Mathematics* 2: 229–260.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg. 2018. "Commuting, Migration, and Local Employment Elasticities." *American Economic Review* 108(12): 3855–3890.

Reuters. 2024. "Intel Prepares $100 Billion Spending Spree Across Four U.S. States." March 20. https://www.reuters.com/technology/intel-prepares-100-bln-spending-spree-across-four-us-states-2024-03-20/.