

Too Vague to Forecast, Too Smart to Pay for One: When Narrowing the Demand Box Buys Nothing

Athaa Asyfahani Nurhermandhiya*, Dermawan Wibisono, and Ima Fatima

School of Business and Management, Institut Teknologi Bandung, Bandung, Indonesia

Abstract—Demand that can only be bracketed is the normal condition of a facility siting study, and minimax regret is the standard way to turn a bracket into a site. What it will not tell you is whether the bracket did any work, since closed-form conditions for when an uncertainty set stops changing a decision exist only where the problem separates coordinatewise, so the reflex is to buy a sharper forecast and hope it earns its fee. We reparametrise the worst-case pairwise margin by box width; because freight geometry fixes the adversarial vertex against each rival, the margin is affine in width and changes sign at most once, so the binding frontier is closed form, a minimum of ratios over rival sites, which we evaluate on parametrised line and star networks. On both families that minimum collapses to the adjacent rival whenever the fixed-cost field is affine in position, leaving frontiers written in demand mass alone. Positions, gaps, spoke lengths and the freight rate all cancel, so distance settles which site wins and says nothing about when the box begins to bind, which is awkward, because distance is the one quantity a siting study measures carefully. Nonnegative floors also cap any derived box at total central mass, a ceiling a posited box does not carry. The frontier therefore prices the forecast: below it, narrowing the box buys nothing; above it, better data buys assurance but not a different warehouse; higher still, it changes the site.

Keywords—facility location, minimax regret, interval uncertainty, partial identification, value of information

I. INTRODUCTION

A firm choosing where to build a warehouse must commit to one site while demand at the nodes it serves is known only within a range. In many industrial settings that range is what the evidence licenses rather than a modelling convenience: forward commitments establish floors, authorisation caps establish ceilings, and no defensible argument collapses the resulting interval to a point forecast. The input to the siting decision is then a box of demand vectors, every one of which the evidence admits.

A set-valued input requires a criterion to collapse it to a recommendation. Minimax regret is the canonical choice for interval data [1], and location theory has developed it thoroughly. Given a box, the criterion returns a site. What it does not report is whether the box mattered - whether the same site would have been

selected had the intervals been narrower, or had demand been known exactly. That question is answered one instance at a time, by solving the problem again.

The distinction carries managerial content, because measurement is costly. If the recommended site is invariant across the box, narrowing the box changes nothing the decision can use, and investment in demand intelligence buys assurance rather than a different warehouse. If the box binds, the same investment buys a different warehouse. A manager deciding whether to commission a demand study is asking which of the two situations he is in, and would prefer to know before commissioning it.

This paper characterises the boundary between them. We show that the worst-case advantage of the midpoint-optimal site over any rival is affine in the width of the box once freight geometry is held fixed, so that its sign changes at most once. The width at which it changes is the binding frontier, available in closed form as a minimum of ratios taken over rival sites. We evaluate the frontier for two parametrised families of synthetic networks, a line and a star, obtain a closed form in the demand primitives for each, and classify the families into a regime in which the box binds the site choice and a regime in which freight geometry identifies the site regardless of the box. We then read the frontier as a price for demand measurement.

Two limits are declared at the outset. The regret machinery is not ours and is cited as known; the box is taken as given, and the measurement model that derives it from operational primitives on a real instance is the subject of a companion working paper. The instances here are synthetic and demonstrative, parametrised to exhibit the regimes the theorem predicts, and nothing below supports a claim about how often those regimes arise in practice.

II. LITERATURE REVIEW

A. Interval minimax-regret location

Locating a facility when demand is known only to lie in an interval box is a mature problem, and minimax regret is its canonical criterion [1]. The criterion selects the site minimising the largest shortfall against the site one would have chosen with hindsight, which makes it computable from the interval endpoints alone and explains why location theory adopted it so quickly. The work that followed extended its reach: from the p-center with uncertain demand weights [2] to the 1-median on a

*Corresponding Author: athaa.asyfahani@gmail.com

tree [3], to trees and general networks [4], [5], to location-allocation [6], and most recently to two-stage capacitated siting with budgeted sets [7]. Each extension enlarges the class of problems the criterion can handle. None changes its position on the box, which is that the box is an input. A wider box yields a different recommendation, and the criterion reports the recommendation without reporting whether the width produced it.

B. Critical set sizes and inverse robustness

Closed-form critical sizes do exist, but where the problem separates. Chassein and Goerigk obtain the largest interval set preserving regret-optimality in explicit form for the unconstrained combinatorial problem [8], whose optimum decouples coordinatewise; a companion result makes the nominal solution regret-optimal at every set size for selection [9]. Once combinatorial structure couples the coordinates, their answer is a mixed-integer program solved per instance. The closed form is bought by separability. The frontier derived below is not: it carries a minimum over candidate sites, which structure resolves rather than removes, and which fails to resolve as soon as the fixed-cost field is not affine in position.

The same work supplies a caution. Under regular interval sets a solution can lose regret-optimality and regain it on arbitrarily many disjoint sizes [8], so a single threshold is not generic under min-max regret. The mechanism is that their regret is a maximum over an exponential comparator set, hence convex and piecewise linear, and two such curves cross repeatedly. The difference is where the uncertainty sits. They box objective coefficients directly, so the adversarial scenario moves with the solution. We box demand primitives; each site's cost is a known functional of the same box, and the adversarial vertex against a given competitor is fixed by freight geometry rather than by the decision. The margin is then affine in width, which Section III establishes rather than assumes. Their companion paper integrates regret across sizes into a single compromise solution [9]; we characterise the critical size instead of averaging over it.

C. Adjacent literatures

The closest question outside this line is the stability radius: how far may demand move before a given site loses its standing? Carrizosa and Nickel maximise the smallest perturbation that pushes a planar site past an admissible cost budget [10], and Pham et al. bound the radius for the median on a tree [11]. The radius is measured against an exogenous budget, and what it establishes is that one site remains acceptable, not that the choice between sites is undisturbed. Our frontier answers the second question, which is why an instance can sit far inside a stability radius and still have its site decided by the width of the box.

Three further literatures touch the width of an uncertainty set without asking when it binds. The price of robustness trades conservatism against cost along a curve [12] that need not even be monotone for capacitated siting [13]; convergence results describe the limit as the set shrinks [14] rather than the width at which the decision switches; and value-of-information calculations identify parameters that repay learning [15] but require a prior the interval frame declines to supply. The methodological cousin lies outside location theory altogether, in instance space analysis [16], which maps

performance onto instance features to expose where methods succeed and fail. We borrow that ambition and pursue it analytically. The framing of a decision identified only up to a set is Manski's [17].

D. Evidence versus characterisation

What this literature does not settle is when the demand box matters to the siting decision at all, as a function of the instance. The absence is one of characterisation rather than of evidence. McGarvey and Thorsen come closest [7]: in roughly a third of their smallest instances, a solution computed without robustness protection already attains the minimax-regret optimum, and their experimental design relates cost structure to the conditions under which the criterion bites. That is evidence that a non-binding region exists and is patterned. It stops short of locating the boundary, and of expressing it in the primitives of a siting instance. We are aware of no closed-form condition, in those primitives, separating boxes that bind the site choice from those leaving it freight-identified, and no classification of instances into the two regimes.

III. METHODOLOGY

A. Cost model and the identified set

A finite set of candidate sites X serves nodes $i = 1, \dots, n$. A demand vector d assigns $d_i \geq 0$ to each node; the rate $c_i(x) \geq 0$ is the cost of serving node i from site x , and $f(x) \geq 0$ is an optional site fixed cost. Total cost is given by (1).

$$C(x, d) = f(x) + \sum_i c_i(x) d_i \quad (1)$$

The identified set is a hyperrectangle $D = \prod_i [\ell_i, u_i]$, each side licensed by the evidence available for that node. We reparametrise it by midpoint and half-width, $m_i = (\ell_i + u_i)/2$ and $w_i = (u_i - \ell_i)/2$, and index its size by a single scalar. Fix a width profile ω with $\omega_i \geq 0$ and $\sum_i \omega_i = 1$, and set $w = t\omega$. The dial t is total half-width, measured in demand units; ω fixes how that width is distributed across nodes. Widths scale about a fixed midpoint, so m is held while t varies. Nonnegative floors require $t\omega_i \leq m_i$ at every node, a constraint we return to in IV-C.

The reference site x^* is a midpoint optimizer, $x^* \in \operatorname{argmin}_x C(x, m)$: the site one would select by ignoring the box and taking its centre at face value. Everything below asks whether the box overturns that selection.

B. Inherited machinery, cited as known

Three results are used and not re-derived. First, the worst case against a fixed pair of sites is attained at a vertex of the box, so the inner infimum over D reduces to a signed sum in closed form rather than an enumeration [1]. Second, the regret of the reference site is $R(x^*) = \max(0, -\sigma)$, where σ is the minimum over rivals of the worst-case pairwise advantage margin. Third, $\sigma \geq 0$ holds if and only if x^* is optimal at every point of the box, so a nonnegative margin establishes necessary optimality rather than merely scoring the selection [18]. What follows is one reparametrisation of these margins together with its consequences; no new regret machinery is introduced.

C. The midpoint-exposure identity

Fix x^* and a rival y . Write $b_i(y) = c_i(y) - c_i(x^*)$ for the freight differential at node i , positive where x^* is advantaged, and $\Delta f(y) = f(y) - f(x^*)$. The worst-case margin is the infimum of $C(y, d) - C(x^*, d)$ over $d \in D$, which vertex reduction evaluates as $\Delta f(y) + \sum_i [b_i^+ \ell_i - b_i^- u_i]$, the fixed-cost term riding outside the infimum because it does not depend on d . Substituting $\ell_i = m_i - w_i$ and $u_i = m_i + w_i$, node i contributes $b_i^+(m_i - w_i) - b_i^-(m_i + w_i) = b_i m_i - |b_i| w_i$, using only the identities $a = a^+ - a^-$ and $|a| = a^+ + a^-$. Summing over nodes gives (2).

$$s(y) = A(y) - E(y) \quad (2)$$

with midpoint margin $A(y) = \Delta f(y) + \sum_i b_i(y) m_i$ and exposure $E(y) = \sum_i |b_i(y)| w_i$. Here $A(y) = C(y, m) - C(x^*, m)$ is the advantage of x^* at the centre of the box, and $E(y)$ is the adversary's best use of the box against the pair. The per-node step exhibits the mechanism: the worst-case corner moves every coordinate by its full half-width in the adverse direction, so each node contributes $|b_i| w_i$ whichever side of the differential it sits on. Because x^* minimises cost at m , $A(y) \geq 0$ for every rival. This is the only place the choice of reference site is used, and it is what excludes a vacuous frontier.

D. Affineness in width

On the width ray the exposure is linear in the dial, $E(y; t) = t\kappa(y)$, where $\kappa(y) = \sum_i |b_i(y)| w_i$ is the exposure per unit width. Each pairwise margin is therefore affine and nonincreasing in t , and the stability margin is a pointwise minimum of finitely many such functions, as in (3).

$$\sigma(t) = \min_{y \neq x^*} [A(y) - t\kappa(y)] \quad (3)$$

A minimum of affine decreasing functions is concave and decreasing, so σ changes sign at most once and the set of widths at which dominance holds is an initial interval. This is what licenses the word regime, and it is where the present setting departs from variable-sized uncertainty in the sense of [8]. There, regret is a maximum over comparators and is convex and piecewise linear, so optimality can be lost and regained on arbitrarily many disjoint sizes. Here the coefficients $A(y)$ and $\kappa(y)$ are fixed by freight geometry and do not move as the box grows, so the lower envelope is an envelope of lines and each rival binds on one interval.

E. Two synthetic families

The frontier is evaluated on two parametrised families. In the line family, nodes sit at positions $q_1 < \dots < q_n$, sites are collocated with nodes, and freight is proportional to distance, $c_i(x_k) = r|q_k - q_i|$; the midpoint optimizer is the weighted median of m . In the star family, a central site is joined to leaves at spoke lengths ρ_k , with $c_i(x_0) = r\rho_i$ and $c_i(x_k) = r(\rho_k + \rho_i)$ for $i \neq k$; the centre is midpoint-optimal precisely when no single leaf holds half of the central mass.

Each family is dialled by the shape of m , the dispersion of the freight field (spacing coefficient of variation for lines, spoke coefficient of variation for stars), the number of nodes, and the width profile ω . The fixed-cost field is taken affine in position: $f(x_k) = f_0 + g q_k$ for lines and $f(x_k) = f_0 + g \rho_k$ for stars. That restriction is not cosmetic: it is exactly the class on which the collapse of

IV-B holds, and IV-B reports where it fails. These instances are demonstrative. They are parametrised to exhibit the regimes the theorem predicts, and no claim is made about the frequency with which real networks occupy them.

IV. FINDINGS AND DISCUSSION

A. The binding frontier

For each rival, $s(y; t) \geq 0$ holds if and only if $t\kappa(y) \leq A(y)$. Because $A(y) \geq 0$, the condition holds at every width when $\kappa(y) = 0$, a rival that differs from x^* in fixed cost alone being unable to use the box at all; otherwise it holds exactly when $t \leq A(y)/\kappa(y)$. Write $\tau(y) = A(y)/\kappa(y)$ for the binding ratio of rival y , with $\tau(y) = \infty$ when $\kappa(y) = 0$. Dominance survives at width t precisely when every rival ratio exceeds it, and intersecting the per-rival intervals gives the frontier (4).

$$t^* = \min_{y \neq x^*} \tau(y) = \min_{y \neq x^*} \frac{A(y)}{\kappa(y)} \quad (4)$$

The frontier is closed form in the primitives: one ratio per rival, with no enumeration, no optimisation and no sampling. Each $\tau(y)$ is the exact width at which rival y first overtakes x^* somewhere in the box, and t^* is the earliest such overtaking. Below it the box leaves the midpoint optimizer undominated without selecting anything, which we call Regime I; above it the recommendation switches somewhere inside the box and the identified set does genuine work, Regime II. Scaling all costs, $c \mapsto \lambda c$ and $f \mapsto \lambda f$, multiplies numerator and denominator alike, so t^* is a shape property of the cost field rather than a scale property. This licenses the dimensionless width $\theta = t/M$, where M is the total central mass.

B. Closed forms and the collapse

Evaluate (4) on the line family against the adjacent rival x_{k+1} , writing h_k for the gap. The freight differential is rh_k at every node at or left of k and $-rh_k$ at every node to the right, so $A = rh_k(M_{\leq k} - M_{> k})$ and $\kappa = rh_k$. The rate and the gap cancel exactly, leaving $\tau(x_{k+1}) = M_{\leq k} - M_{> k}$; the mirror rival gives $\tau(x_{k-1}) = M_{\geq k} - M_{< k}$. Distant rivals need not be checked. For f affine in position the τ -argmin is always adjacent to x^* : margins telescope along the path and are therefore additive, exposure is subadditive by the triangle inequality, and the median inequality bounds any distant ratio below by the smallest step ratio on the path, while the step ratios increase away from the weighted median. Minimising over the two adjacent rivals then gives (5),

$$t_{\text{line}}^* = m_{x^*} - |(S_{<} - S_{>}) + g/r| \quad (5)$$

where $S_{<}$ and $S_{>}$ are the central mass strictly left and strictly right of x^* . The star is direct: against leaf k the differential is $r\rho_k$ at every other node and $-r\rho_k$ at k itself, so $A(x_k) = r\rho_k(M - 2m_k)$ and $\kappa(x_k) = r\rho_k$, giving $\tau(x_k) = M - 2m_k$ and hence (6).

$$t_{\text{star}}^* = g/r + M - 2m_{\max} \quad (6)$$

The star frontier reproduces an adversarial-majority reading from first principles: the box binds exactly when it is wide enough for some single leaf to hold a majority

of demand. The collapse behind (5), however, is bought by the affine restriction of III-E and not by structure. Once f is not affine in position, additivity of the margin fails, distant rivals bind, and the minimum in (4) must be taken as written. This is where the present result differs from a closed form obtained by separability [8]: the minimum over sites is doing work, and the two families resolve it rather than remove it.

C. Dispersion invariance

Neither closed form contains a freight quantity. Positions, gaps, spoke lengths and the rate r have all cancelled, and the only survivor is g/r , the ratio of fixed to variable cost. Nor does spacing re-enter through the reference site, since the weighted median index depends on m alone. The frontier is therefore invariant to the dispersion of the freight field. Freight geometry determines which site is optimal; it never determines the width at which the box begins to bind. In dimensionless form θ^* is a function of the mass shape and g/r only, and for the star it is $\theta^* = g/(rM) + 1 - 2m_{\max}/M$. A phase diagram of width against freight dispersion is exactly flat. That is a finding about the two families, not a defect of the projection. Affinity is not derived here: it is imposed on both families in III-E, and it is what licenses the collapse of (4) to the single adjacent ratio from which the freight quantities cancel. Fixed costs thresholded at a capacity break, or convex in position, destroy that collapse, since margins no longer telescope, distant rivals can attain the minimum, and the min over rivals must be taken as written. The invariance is then local rather than global: it holds on the region of a family where one rival binds, and the frontier stays closed form, but its reading as a statistic of demand mass alone is no longer available across the family.

D. The feasibility ceiling

Because the box is derived rather than posited, its floors must be nonnegative, so $t\omega_i \leq m_i$ at every node and the admissible width is capped at $t_{\text{feas}} = \min_i m_i/\omega_i$. Since ω is a probability vector, $\min_i (m_i/\omega_i) \leq \sum_i \omega_i (m_i/\omega_i) = M$, so that $t_{\text{feas}} \leq M$ in every instance of either family. The frontier obeys the same bound, $t^* \leq m_{x^*} \leq M$ on the line with $g = 0$. Frontier and ceiling are thus denominated in the same currency, total central mass, and their ordering decides the regime an instance occupies. An uncertainty set that is not derived from primitives carries no such ceiling, which is the sense in which derivedness, rather than interval-valuedness alone, does work here.

E. What measurement buys

The frontier prices demand measurement in three tiers rather than two. A manager weighing a study is asking which tier he is in. Below t^* the recommended site is optimal everywhere in the box, so whatever a sharper forecast returns, the warehouse goes where it was going. Above t^* dominance fails, but lost dominance is not a moved warehouse: there is a band in which regret at x^* is strictly positive while x^* remains the regret argmin. On the star the band is provable. At $t = t^*$ regret at the centre is zero by definition, while leaf h carries regret at least $2r\rho_h t^*$, strictly positive whenever $t^* > 0$; the centre therefore wins strictly at the frontier and, by continuity, beyond it. Within that band measurement buys assurance, not a different warehouse: worth its fee when the site must be defended to a board or a

regulator, not when the only question is where to build. Only above the switch width does measurement change the site.

The three tiers are not equally reachable. In some instances the switch width exceeds t_{feas} , so no admissible box reaches the third tier at fixed m and the middle tier covers the whole feasible range. Such a manager can be told something stronger than that his intervals are slack: no valid demand box, however wide, moves the warehouse.

F. Limitations

The instances are synthetic and demonstrative. The paper reports which regimes exist and what separates them, not how often they arise. The frontier (4) assumes nothing about topology, but the closed forms (5) and (6) and the collapse behind them do: each is family-specific, and the collapse additionally requires f affine in position. The measurement model that derives the box from operational primitives belongs to the companion paper. A computational sweep of the two families, and a check of whether (5) can be recovered by a coordinatewise argument in the sense of [8], are deferred to the extended version, as is any non-synthetic instance, the operator network motivating this work not being available at the resolution the frontier requires.

V. CONCLUSION

The frontier answers a question managers ask and the location literature has not: whether the demand box is doing anything. What makes the answer cheap is that the adversarial vertex is fixed by freight geometry, so the margin is affine in width and the binding width is a minimum of ratios rather than a search.

What the two families then reveal is a division of labour that neither the regret literature nor the stability-radius literature would predict. Freight geometry is what a siting study actually measures, and it settles which site wins. It has nothing to say about when the box matters: positions, gaps, spoke lengths and the rate all cancel, and the binding width is a statistic of demand mass alone. The dispersion of the cost field, in other words, is the wrong thing to look at when deciding whether to buy better demand data, and it is exactly the thing a siting analyst has closest to hand.

That cuts against the intuition that a network with wildly uneven freight costs is the one most exposed to demand uncertainty. It is not, at least not here. Whether the invariance survives outside the affine class is the first thing worth testing, since the same condition governs the collapse to the adjacent rival; if it fails, the frontier remains closed form but its reading as a pure demand statistic does not. The line frontier also invites a direct comparison with the stability radius of the weighted median, which would settle how much of the result is new and how much is an old quantity seen from a different side.

REFERENCES

- [1] P. Kouvelis and G. Yu, Robust Discrete Optimization and Its Applications.

- Dordrecht, The Netherlands: Kluwer Academic, 1997.
- [2] I. Averbakh and O. Berman, "Minimax regret p-center location on a network with demand uncertainty," *Location Science*, vol. 5, no. 4, pp. 247–254, 1997.
 - [3] B. Chen and C.-S. Lin, "Minmax-regret robust 1-median location on a tree," *Networks*, vol. 31, no. 2, pp. 93–103, 1998.
 - [4] I. Averbakh and O. Berman, "Algorithms for the robust 1-center problem on a tree," *Eur. J. Oper. Res.*, vol. 123, no. 2, pp. 292–302, 2000.
 - [5] I. Averbakh and O. Berman, "Minimax regret median location on a network under uncertainty," *INFORMS J. Comput.*, vol. 12, no. 2, pp. 104–110, 2000.
 - [6] E. Conde, "Minimax regret location–allocation problem on a network under uncertainty," *Eur. J. Oper. Res.*, vol. 179, no. 3, pp. 1025–1039, 2007.
 - [7] R. G. McGarvey and A. Thorsen, "Minimizing maximum regret in two-stage facility location with site preparation costs," *Eur. J. Oper. Res.*, 2026, in press.
 - [8] A. Chassein and M. Goerigk, "Variable-sized uncertainty and inverse problems in robust optimization," *Eur. J. Oper. Res.*, vol. 264, no. 1, pp. 17–28, 2018.
 - [9] A. Chassein and M. Goerigk, "Compromise solutions for robust combinatorial optimization with variable-sized uncertainty," *Eur. J. Oper. Res.*, vol. 269, no. 2, pp. 544–555, 2018.
 - [10] E. Carrizosa and S. Nickel, "Robust facility location," *Math. Methods Oper. Res.*, vol. 58, no. 2, pp. 331–349, 2003.
 - [11] V. H. Pham, N. Nguyen, M.-Q. Cao, H.-P. Dang, and T.-L. Tran, "Stability radius and an upgrading model of median location on trees," *Asia-Pacific J. Oper. Res.*, vol. 43, no. 2, 2026.
 - [12] D. Bertsimas and M. Sim, "The price of robustness," *Oper. Res.*, vol. 52, no. 1, pp. 35–53, 2004.
 - [13] R. Baldacci, M. Caserta, E. Traversi, and R. Wolfler Calvo, "Robustness of solutions to the capacitated facility location problem with uncertain demand," *Optim. Lett.*, vol. 16, no. 9, pp. 2545–2560, 2022.
 - [14] G. P. Crespi, D. Kuroiwa, and M. Rocca, "Robust optimization: Sensitivity to uncertainty in scalar and vector cases, with applications," *Oper. Res. Perspect.*, vol. 5, pp. 113–119, 2018.
 - [15] A. Zitrou, T. Bedford, and A. Daneshkhah, "Robustness of maintenance decisions: Uncertainty modelling and value of information," *Rel. Eng. Syst. Saf.*, vol. 109, pp. 89–103, 2013.
 - [16] K. Smith-Miles and M. A. Muñoz, "Instance space analysis for algorithm testing: Methodology and software tools," *ACM Comput. Surv.*, vol. 55, no. 12, art. 255, pp. 1–31, 2023.
 - [17] C. F. Manski, "Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice," *J. Econometrics*, vol. 95, no. 2, pp. 415–442, 2000.
 - [18] M. Inuiguchi and M. Sakawa, "Possible and necessary optimality tests in possibilistic linear programming problems," *Fuzzy Sets Syst.*, vol. 67, no. 1, pp. 29–46, 1994.