

WHEN SILOED MARKETING TEAMS RAISE ACQUISITION COSTS: THE PRICE OF ANARCHY INSIDE A SINGLE ADVERTISER

Igor Ivitskiy¹

Abstract. Firms that divide customer acquisition among product, regional or channel teams can see the total cost per lead rise even when each team meets its own target. The *subject* of the study is the customer acquisition system of a single advertiser whose teams draw leads from a shared pool of prospects, and the gap between the equilibrium reached by independent teams and the cost-minimizing allocation. The *purpose* of the paper is to specify when this situation forms a congestion game, to quantify its coordination cost and to translate the known remedies of such games into advertising management. *Methodology.* The study applies the theory of atomic splittable congestion games. The model assumes homogeneous, deduplicated leads and fixed quotas, with the volume of each team raising the cost of comparable leads for the others; because Google Ads keeps campaigns within one account from bidding each other up, the cost increase comes from depleting the shared pool rather than from auction prices. A closed-form benchmark is derived analytically and checked by numerical simulation, and two conditional extensions draw on Braess's paradox and on the bullwhip effect. *Results.* The paper introduces the Funnel Coordination Gap, the share of acquisition cost that coordination can avoid. A single team produces no gap and each additional team widens it: in the benchmark, ten teams spend 22% more than a coordinated firm, the avoidable share approaches 25% as teams multiply, and independent teams take up to twice the cost-minimizing volume from the shared pool. Known bounds for this class of games cap the avoidable share at one quarter for identical teams and at one third for teams with different quotas. A personal internal charge per lead, an established toll mechanism, makes the cost-minimizing allocation an equilibrium. Under stated conditions, a costless link between funnel stages can raise total cost by one third, and lower bounds on lead variance multiply across stages unless end-customer conversion data are shared. *Conclusions.* The internal team structure of an advertiser is a measurable source of acquisition cost, and its size depends on the number of teams and on how sharply the cost of the shared pool rises with volume. *Practical implications.* Managers obtain a diagnostic of team structure, an internal price for shared sources, a holdout test for channel removal and a controlled experiment that could falsify the central assumption. *Value / originality.* The paper contributes the conditions under which acquisition inside a firm is a congestion game, a closed-form advertising benchmark placed against exact bounds, and a translation of known pricing mechanisms into management instruments.

Keywords: congestion games, Nash equilibrium, customer acquisition cost, congestion pricing, multichannel advertising, Braess's paradox, bullwhip effect, digital economy.

JEL Classification: M37, C72, L22

1. Introduction

Advertisers rarely buy customers through a single decision maker. Budgets are split among product lines, regions, channels and agencies, and each unit is judged by whether it delivers its quota of leads at an acceptable cost per lead. The units often draw on the same stock of potential customers, who search for the category, see

the brand on the same platforms and submit a form only once. When one unit converts such a prospect, the next comparable lead becomes more expensive for everyone, so every unit can meet its target while the firm pays for coordination it lacks. In the digital economy this loss takes a plain form: advertising spend that buys no additional customers.

¹ Doctor Ads LTD, United Kingdom

E-mail: igor@thedoctorads.com

ORCID: <https://orcid.org/0000-0002-9749-6414>



Research on efficiency in online advertising addresses two other settings. The first is competition between advertisers whose automated bidding agents meet in auctions (Aggarwal et al., 2019; Aggarwal et al., 2023). The second is a single advertiser who controls several channels through per-channel levers, where optimizing only return-on-investment targets can be arbitrarily worse than the global optimum while per-channel budgets suffice to reach it (Deng et al., 2023). Coordination between functions inside the firm has been modelled for marketing and operations (Liu et al., 2015), and the loss caused by independent users of congestible resources is the subject of congestion game theory (Roughgarden & Tardos, 2002; Cominetti et al., 2009). What has not been set out is when customer acquisition inside one advertiser belongs to this class and what the known results mean for its management. Preliminary reports by the author applied the bound for nonatomic routing to funnel teams by analogy (Ivitskiy, 2026a, 2026b, 2026c); the novelty of the present paper is a model built from the costs of the teams themselves, which shows which bound applies to which team structure.

The aim of the paper is to specify the conditions under which the acquisition system of a single advertiser forms an atomic splittable congestion game, to quantify the coordination cost of independent teams and to translate known remedies into management instruments. The tasks are to formulate the model and the mechanism through which one team raises the costs of another, to derive a closed-form benchmark and compare it with known bounds, and to interpret the toll mechanism as an internal price of leads, adding two conditional extensions. The methodology is theoretical modelling with analytical derivation and numerical verification. The analysis is static, assumes locally affine costs, contains no empirical estimates and claims no new theorems. Section 2 presents the theoretical basis and the model, Section 3 the results and Section 4 the discussion, followed by the conclusions.

2. Theoretical Basis and Model

A congestion game describes players who share resources whose cost grows with the total load placed on them; its inefficiency is measured by the price of anarchy, the ratio between the total cost of the worst equilibrium and the minimum

achievable cost. With affine costs the price of anarchy is at most $4/3$ when each player controls an infinitesimal share of the load (Roughgarden & Tardos, 2002). In the atomic splittable version each player controls a fixed, divisible volume (Orda et al., 1993) and bears only part of the congestion it creates. For K symmetric players with affine costs the price of anarchy is at most $4K^2/((K+1)(3K-1))$, a tight bound that tends to $4/3$, while for K players of general form the upper bound is $(3K+1)/(2K+2)$, which exceeds $4/3$ once K is larger than five (Cominetti et al., 2009); the limit $3/2$ is the exact worst case (Roughgarden & Schoppmann, 2015). If a coordinator may charge different prices to different players, charging each player the part of the marginal cost that it does not bear makes the cost-minimizing allocation an equilibrium (Cominetti et al., 2009), and anonymous tolls with the same effect exist (Swamy, 2012; Yang & Zhang, 2008). Losses from decentralized decisions have also been quantified in supply chains (Perakis & Roels, 2007), and budget allocation based on attribution carries caveats that can mislead the decision maker (Danaher & van Heerde, 2018).

A firm acquires a volume of comparable leads during a period, normalized to one, through K teams; team i must deliver a quota d_i , and the quotas sum to one. Leads come from sources $e = 1, \dots, m$, such as a pool of prospects reached through several campaigns or an alternative channel that other teams do not use. Let x_{ie} be the volume that team i obtains from source e and f_e the total volume that all teams obtain from it. The average cost of a comparable lead is

$$c_e(f_e) = a_e f_e + b_e, \quad (1)$$

where $a_e \geq 0$ is the sensitivity of the cost to the total volume taken from source e and $b_e \geq 0$ is the cost at negligible volume. The cost reported by team i and the total cost of the firm are

$$C_i = \sum_{e=1}^m x_{ie} c_e(f_e), \quad C = \sum_{e=1}^m f_e c_e(f_e), \quad (2)$$

where $\sum_e x_{ie} = d_i$, $x_{ie} \geq 0$ and $f_e = \sum_i x_{ie}$. An allocation in which no team can lower its own cost by moving part of its quota is a Nash equilibrium x^{NE} , and the allocation that minimizes C is x^{OPT} . The *Funnel Coordination Gap* is the share of the cost of the firm under independent teams that coordination can avoid:

$$G = 1 - \frac{C(x^{OPT})}{C(x^{NE})} = 1 - \frac{1}{PoA}, \quad (3)$$

where $PoA = C(x^{NE})/C(x^{OPT})$ is the price of anarchy of the team structure of the firm.

The model rests on four assumptions. First, leads are homogeneous and deduplicated: a prospect converts once, for whichever team reaches it first, and no team can restrict itself to prospects that others have not approached. Second, the average cost of a comparable lead rises with the total volume the firm has already taken from the shared pool. Diminishing response to advertising, with a mean short-term elasticity of 0.12 across 751 estimates (Sethuraman et al., 2011), makes this plausible for total volume, but that evidence does not establish the cross-team part of the assumption, which Section 4 proposes to test. Third, quotas are fixed for the period and costs are locally affine. Fourth, each team minimizes its own cost. The interaction does not run through auction prices inside one account, because Google Ads gives priority to the campaign with the highest Ad Rank, so that campaigns in the same account do not bid each other up (Google, n.d.); if the audiences of the teams do not overlap, the game disappears. Under these assumptions the acquisition system is an atomic splittable congestion game (Orda et al., 1993; Cominetti et al., 2009).

3. Results

3.1 Closed-form benchmark

Consider two sources: a shared pool with cost $c_s(F) = aF + b_s$, where $a > 0$, and an alternative channel with constant cost $b > b_s$, used by K symmetric teams with quotas $1/K$. Let F be the volume the firm takes from the pool, $s = (b - b_s)/a$ the price advantage of the pool in units of its price sensitivity and $\beta = b/a$ the scaled cost of the alternative channel, with $0 < s \leq 1$ and $\beta \geq s$. The cost of team i is $x_i(aF + b_s) + (1/K - x_i)b$; its derivative is $a(F + x_i - s)$ and its second derivative is $2a > 0$, so in a symmetric equilibrium $x_i = s/(K+1)$. The cost of the firm, $a(F^2 - sF + \beta)$, is minimized at half the price advantage, hence

$$F^{NE} = \frac{Ks}{K+1}, \quad F^{OPT} = \frac{s}{2}, \quad (4)$$

where F^{NE} and F^{OPT} are the volumes taken from the pool at the equilibrium and at the cost-

minimizing allocation. Substituting (4) into the cost of the firm and using (3) gives

$$G(K; s, \beta) = 1 - \frac{\beta - s^2/4}{\beta - Ks^2/(K+1)^2}, \quad (5)$$

where $G(K; s, \beta)$ is the share of acquisition cost that coordination can avoid.

A single team produces no gap, because $K/(K+1)^2 = 1/4$ at $K=1$. The gap widens with every additional team and tends to $s^2/(4\beta)$, which equals 25% when the pool has no base cost and $s=1$; independent teams take $2K/(K+1)$ times the cost-minimizing volume from the pool. For $s=\beta=1$, two teams leave a gap of 3.6%, five teams 12.9%, ten teams 18.2%, which means 22.3% more spending than under coordination, and fifty teams 23.5%. Against the exact bound for symmetric teams, the benchmark reaches 56% of the permitted excess cost at $K=2$ and 88% at $K=10$. Best-response dynamics reproduced (5) to within 10^{-4} , and among 400 random symmetric games with affine costs the largest price of anarchy was 1.210.

3.2 Bounds and internal charge

Table 1 converts the known bounds into the largest possible gap for any configuration that satisfies the assumptions of the model.

The bound for teams of general form is an upper limit for finite K and is not claimed to be attained. The gap arises because a team that takes one more lead from source e raises the cost of the $f_e - x_{ie}$ leads of the other teams by a_e each. Applying the price discrimination mechanism of the congestion game literature (Cominetti et al., 2009), the internal charge per lead is

$$\tau_{ie} = a_e(f_e^{OPT} - x_{ie}^{OPT}), \quad (6)$$

where x_{ie}^{OPT} is the volume assigned to team i from source e in a cost-minimizing allocation and f_e^{OPT} is the total volume from that source. With this charge the marginal cost of a team at the optimum equals the marginal cost of the firm, $2a_e f_e^{OPT} + b_e$, so the optimum satisfies the equilibrium condition of every team; in 100 random games the equilibrium under the charges matched the optimum to within 10^{-12} . For symmetric teams the charge equals $(b - b_s)(K - 1)/(2K)$ in the benchmark, which is 45% of the price advantage of the pool with ten teams. Setting budgets by source from the centre reaches the same allocation, consistent with the




Figure 1. Funnel Coordination Gap by the number of symmetric teams K for $s = 1, 0.5$ and 0.25 with $\beta = s$, and the exact upper bound for symmetric teams with affine costs (dash-dot line); the dotted line marks 25%.

Table 1

Upper bounds on the price of anarchy and on the Funnel Coordination Gap, affine costs

Team structure	Bound on PoA	Largest G for $K = 2 / 5 / 10 / \text{unlimited}$
Single team	1	0 / 0 / 0 / 0
Symmetric teams	$4K^2 / ((K+1)(3K-1))$	6.3% / 16.0% / 20.3% / 25.0%
Teams of general form	$(3K+1) / (2K+2)$	14.3% / 25.0% / 29.0% / 33.3%

Source: compiled from bounds in the literature (Cominetti et al., 2009; Roughgarden & Schoppmann, 2015)

finding that per-channel budgets suffice for the global optimum (Deng et al., 2023).

3.3 Conditional extensions

If many small independent optimizers, such as campaign-level bidding algorithms, route prospects through funnel stages, the nonatomic limit applies. Let one unit of prospects move from the pool to conversion through stage A, with step costs x and 1, or stage B, with step costs 1 and x , where x is the volume on the step. The flow splits equally at a cost of 1.5 per unit. Adding a costless link from A to B, for example retargeting that moves prospects from the awareness activity of one team into the consideration activity of another, draws all flow onto the path through A and B, and every path then costs 2.0, one third more (Braess, 1968). Under reasonable assumptions such a paradox is about as likely to occur as not (Steinberg & Zangwill, 1983), and detecting it in a general network is NP-hard (Roughgarden,

2006), so the effect of a chosen change should be computed within the model and checked with a holdout test.

If each stage forecasts with a moving average over p periods of the requests it receives from the stage closer to the customer and sets its own request with an order-up-to rule after a processing lag L , variance travels from sales toward advertising (Lee et al., 1997). Requests are treated as planning quantities in units of end-customer demand and, as in the source model, may be negative (Chen et al., 1999). For one stage facing autoregressive demand

$$\frac{\text{Var}(q)}{\text{Var}(D)} \geq 1 + \left(\frac{2L}{p} + \frac{2L^2}{p^2} \right) (1 - \rho^p), \quad (7)$$

where q is the request of the stage, D the demand it observes and ρ the autocorrelation of demand (Chen et al., 2000). For uncorrelated demand and k stages with lags L_j , the Lead Variance Amplification $\text{LVA}_k = \text{Var}(q_k) / \text{Var}(D)$ satisfies

$$LVA_k \geq \prod_{j=1}^k \left(1 + \frac{2L_j}{p} + \frac{2L_j^2}{p^2} \right) \quad (8)$$

when each stage forecasts from the requests of its neighbour, and

$$LVA_k \geq 1 + \frac{2 \sum_j L_j}{p} + \frac{2 \left(\sum_j L_j \right)^2}{p^2} \quad (9)$$

when every stage forecasts from end-customer demand (Chen et al., 1999). With $p = 4$ and $L = 1$ the bound for five stages is 11.331 under (8) and 6.625 under (9): shared data slow the growth of variance but do not remove amplification.

4. Discussion

The model does not apply where a central planner already fixes how much each team takes from each source, or where the audiences of the teams do not overlap; the size of the overlap is the first quantity a manager has to measure. Decentralization also has benefits, since independent teams test sources that a planner would not try and limit the damage of central errors, and coordination has its own costs. The Funnel Coordination Gap is therefore an upper limit on what coordination can recover, and coordination pays only when that limit exceeds its cost. The same externality can be read as overuse of a common-pool resource or as a problem of team incentives; the congestion game formulation is used because it provides exact bounds and a pricing mechanism with known properties.

Three boundary conditions apply. The interaction runs through depletion of a shared pool, not through auction prices. Leads must be of comparable quality, since a team that systematically attracts better prospects breaks the common cost structure. The game is static and costs are affine only locally, so the parameters a and b drift in repeated auctions with learning bidding algorithms, and (5) holds only while the linear approximation is adequate.

The central assumption can be tested directly. In a controlled experiment the firm raises the volume of team j by a fixed proportion in randomly assigned periods or regions, keeps the budget and targets of team i unchanged, and measures the cost of a deduplicated, quality-matched lead for team i . The test should fix in advance the smallest cost increase that matters for the firm: an

estimate whose confidence interval excludes that effect rejects the practically relevant version of the assumption for the pool, while an interval wide enough to include both zero and that effect leaves the question open. Where the assumption holds, a manager can estimate the audience overlap, the slope a from the relation between volume and lead cost, and the number of teams, and then compute the gap from (5), or numerically when quotas differ, and the internal charge from (6). All of these steps require conversion data that are not corrupted.

5. Conclusions

The paper established the conditions under which customer acquisition inside a single advertiser forms an atomic splittable congestion game: homogeneous, deduplicated leads drawn from a shared pool, fixed quotas and a cost of comparable leads that rises with the total volume taken by all teams. Under these conditions the known results for this class of games apply to the organization of the firm itself. For symmetric teams choosing between a shared pool and an uncontested channel the Funnel Coordination Gap has a closed form: it is zero for one team, grows with every additional team, reaches 18.2% for ten teams in the benchmark, which is 22.3% more spending, and approaches 25%, while known bounds cap the avoidable share at one quarter for symmetric teams and one third for teams of general form. A personal internal charge per lead makes the cost-minimizing allocation an equilibrium without central control of campaigns. Under additional conditions, adding a costless link between funnel stages can raise total cost by one third, and lower bounds on lead variance multiply across stages unless end-customer data are shared.

Further research should test the central assumption in controlled experiments, estimate the gap on data from firms with several acquisition teams, extend the model to repeated auctions with learning bidding algorithms, and allow for teams that differ in lead quality and access to sources.

Declarations. The author is the founder of an advertising agency; no client data were used and the research received no external funding. Scripts for the numerical checks and the figure are available from the author. Generative AI tools

(Claude, Anthropic; Codex, OpenAI) assisted with drafting, language editing, verification of bibliographic metadata and code for numerical

checks; all claims, derivations, citations and wording were verified by the author, who bears full responsibility for the paper.

References:

- Aggarwal, G., Badanidiyuru, A., & Mehta, A. (2019). Autobidding with constraints. In I. Caragiannis, V. Mirrokni, & E. Nikolova (Eds.), *Web and Internet economics* (pp. 17–30). Springer International Publishing. https://doi.org/10.1007/978-3-030-35389-6_2
- Aggarwal, G., Perlröth, A., & Zhao, J. (2023). Multi-channel auction design in the autobidding world. In *Proceedings of the 24th ACM conference on economics and computation* (p. 21). ACM. <https://doi.org/10.1145/3580507.3597707>
- Braess, D. (1968). Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung Operations Research – Recherche Opérationnelle*, 12(1), 258–268. <https://doi.org/10.1007/bf01918335>
- Chen, F., Drezner, Z., Ryan, J. K., & Simchi-Levi, D. (1999). The bullwhip effect: Managerial insights on the impact of forecasting and information on variability in a supply chain. In S. Tayur, R. Ganeshan, & M. Magazine (Eds.), *Quantitative models for supply chain management* (pp. 417–439). Springer US. https://doi.org/10.1007/978-1-4615-4949-9_14
- Chen, F., Drezner, Z., Ryan, J. K., & Simchi-Levi, D. (2000). Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Science*, 46(3), 436–443. <https://doi.org/10.1287/mnsc.46.3.436.12069>
- Cominetti, R., Correa, J. R., & Stier-Moses, N. E. (2009). The impact of oligopolistic competition in networks. *Operations Research*, 57(6), 1421–1437. <https://doi.org/10.1287/opre.1080.0653>
- Danaher, P. J., & van Heerde, H. J. (2018). Delusion in attribution: Caveats in using attribution for multimedia budget allocation. *Journal of Marketing Research*, 55(5), 667–685. <https://doi.org/10.1177/0022243718802845>
- Deng, Y., Golrezaei, N., Jaillet, P., Liang, J. C. N., & Mirrokni, V. (2023). Multi-channel autobidding with budget and ROI constraints. In *Proceedings of the 40th International Conference on Machine Learning* (Vol. 202, pp. 7617–7644). PMLR. <https://proceedings.mlr.press/v202/deng23c.html>
- Google. (n.d.). *How Performance Max interacts with other campaigns in your account*. Google Ads Help. <https://support.google.com/google-ads/answer/13810170>
- Ivitskiy, I. (2026a). Braess's paradox in multi-channel advertising: When adding channels reduces conversion performance. In *VII International Scientific Conference "Period of Transformation Processes in World Science: Tasks and Challenges"* (Ukraine, Poltava, May 22, 2026) (pp. 76–78). TOV UKRLOHOS Hrup. <https://doi.org/10.62731/mcnd-22.05.2026.002>
- Ivitskiy, I. (2026b). The bullwhip effect in lead generation funnels: Demand variance amplification under decentralized processing. In *XI International Scientific Conference "Problems and Prospects of Implementing Interdisciplinary Scientific Achievements"* (Ukraine, Kamianets-Podilskyi, May 29, 2026) (pp. 56–59). TOV UKRLOHOS Hrup. <https://doi.org/10.62731/mcnd-29.05.2026.003>
- Ivitskiy, I. (2026c). The price of anarchy in siloed funnel management: Nash equilibrium bounds on decentralized conversion optimization. In *VI International Scientific Conference "Innovative Science: Searching for Answers to Modern Challenges"* (Ukraine, Chernivtsi, May 15, 2026) (pp. 93–96). TOV UKRLOHOS Hrup. <https://doi.org/10.62731/mcnd-15.05.2026.002>
- Lee, H. L., Padmanabhan, V., & Whang, S. (1997). Information distortion in a supply chain: The bullwhip effect. *Management Science*, 43(4), 546–558. <https://doi.org/10.1287/mnsc.43.4.546>
- Liu, G., Zhang, J., & Tang, W. (2015). Strategic transfer pricing in a marketing-operations interface with quality level and advertising dependent goodwill. *Omega*, 56, 1–15. <https://doi.org/10.1016/j.omega.2015.01.004>
- Orda, A., Rom, R., & Shimkin, N. (1993). Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking*, 1(5), 510–521. <https://doi.org/10.1109/90.251910>
- Perakis, G., & Roels, G. (2007). The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Science*, 53(8), 1249–1268. <https://doi.org/10.1287/mnsc.1060.0656>
- Roughgarden, T. (2006). On the severity of Braess's paradox: Designing networks for selfish users is hard. *Journal of Computer and System Sciences*, 72(5), 922–953. <https://doi.org/10.1016/j.jcss.2005.05.009>
- Roughgarden, T., & Schoppmann, F. (2015). Local smoothness and the price of anarchy in splittable congestion games. *Journal of Economic Theory*, 156, 317–342. <https://doi.org/10.1016/j.jet.2014.04.005>

- Roughgarden, T., & Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM*, 49(2), 236–259. <https://doi.org/10.1145/506147.506153>
- Sethuraman, R., Tellis, G. J., & Briesch, R. A. (2011). How well does advertising work? Generalizations from meta-analysis of brand advertising elasticities. *Journal of Marketing Research*, 48(3), 457–471. <https://doi.org/10.1509/jmkr.48.3.457>
- Steinberg, R., & Zangwill, W. I. (1983). The prevalence of Braess' paradox. *Transportation Science*, 17(3), 301–318. <https://doi.org/10.1287/trsc.17.3.301>
- Swamy, C. (2012). The effectiveness of Stackelberg strategies and tolls for network congestion games. *ACM Transactions on Algorithms*, 8(4), 1–19. <https://doi.org/10.1145/2344422.2344426>
- Yang, H., & Zhang, X. (2008). Existence of anonymous link tolls for system optimum on networks with mixed equilibrium behaviors. *Transportation Research Part B: Methodological*, 42(2), 99–112. <https://doi.org/10.1016/j.trb.2007.07.001>

Received on: 18th of August, 2026

Accepted on: 25th of September, 2026

Published on: 14th of October, 2026