The Public Meeting Paradox: How NIMBY-Dominated Public Meetings Can Enable New Housing*

Allison K. Cuttner[†] Ryan Hübert[‡] B. Pablo Montagnes[§]

February 2024

Abstract

Public meetings to consider new housing proposals often feature visible and vocal opposition from neighboring residents, creating a perception that these meetings impede the growth of the housing supply. We analyze a model where residents can legally challenge a developer's housing proposal. A public meeting serves as a critical tool for developers to identify potential litigants, enabling them to adjust proposals and avoid legal action. Interestingly, developers prefer meetings dominated by opponents since it is easier to identify potentially litigious neighbors. Contrary to common belief, our findings suggest that public meetings dominated by NIMBY opponents can *increase* housing supply by fostering compromise projects. This challenges the prevailing conventional wisdom that unrepresentative meetings significantly restrict housing development. Our analysis instead focuses attention on the threat of litigation as the key driver of the undersupply of housing.

Keywords: formal models, housing politics, NIMBYs, public meetings

Preliminary Draft Prepared for the Political Economy of Housing Conference
The Price School at the University of Southern California

Comments and suggestions welcome.

*Authors are listed in alphabetical order. We thank Ryan Copus and Chris Elmendorf for helpful comments. All errors and omissions are our own.

[†]Assistant Teaching Professor, Department of Quantitative Theory and Methods, Emory University. Email: allison.cuttner@emory.edu.

[‡]Assistant Professor, Department of Political Science, University of California, Davis. Email: rhubert@ucdavis.edu.

[§]Associate Professor, Departments of Political Science and Quantitative Theory and Methods, Emory University. Email: pablo.montagnes@emory.edu.

Local institutions such as planning boards, city government, and school boards often solicit public feedback through local meetings. Understanding the institution of local meetings and the incentives of participants is an important first step to understanding the normative value (or cost) of such institutions and for developing policies to improve outcomes in the communities these meetings purport to serve.

In the context of housing, scholars have provided evidence suggesting that local, public meetings for variances, zoning changes, and other project approvals often lead to costly delays and result in compromise projects (Einstein, Glick, and Palmer 2020). Furthermore, meetings are not usually representative of the broader community that could be affected by the development, and they are often dominated by opponents (Einstein, Palmer, and Glick 2019; Einstein, Glick, and Palmer 2020; Einstein et al. 2023). This invites the question: why can't developers and other interested parties anticipate the opposition and compromise ahead of time? As delays are extremely costly, there is an enormous incentive to propose the optimal project that will be approved without delay. If unrepresentative and undemocratic meetings impose costs on prospective developers, why have meetings at all?

We first develop a model of project proposals by developers absent local planning meetings. Developers make proposals for projects anticipating potential local opposition. Residents in a community are able to observe the project and, if sufficiently motivated, may oppose the project through litigation. Judges resolve cases based on the balance of evidence of harm and the public good and may either approve the developer project or impose costly restrictions in accordance with citizen complaints. The developer must rely on her prior belief about citizen opposition. If developers are sufficiently lawsuit averse, this results in projects that are too small. Citizens take self-interested actions, suing to stop the proposed project if it is sufficiently bad for them. As the threat of these actions shape the choices of developers, costly actions by citizens affect the welfare of other residents.

We then introduce the institution of local meetings. Local meetings come after the initial proposal, but following meetings, developers may alter their proposal. In the model, meetings serve multiple purposes. First, meetings are an opportunity for developers to learn about the intensity and distribution of opposition and support. Motivated by the empirical literature on unrepresentative meeting attendance (Einstein, Glick, and Palmer 2020; Einstein et al. 2023), we will model this as an explicitly biased sample arising from endogenous citizen choices of meeting attendance. Nonetheless, we show that a biased sample is valuable to developers, leading to *more* housing than otherwise would occur. Second, meetings are an administrative avenue of dispute resolution. If a meeting is available for citizens to voice their concerns and they fail to show up and do so, they may not have standing to pursue litigation to delay (or stop) development. Both of these are good for the developer and for increasing the housing supply.

We begin by situating our model in a growing literature on public meetings, citizen opposition to housing development, and anti-development litigation. We define citizen and developer preferences over housing project sizes and first consider a world without meetings where a developer proposes a project and a citizen can sue to stop the development. Lawsuits are costly for both citizens and the developer—despite divergent preferences over the size of a housing project, a developer would rather compromise with citizens then suffer the delays and fees of a suit. With uncertainty over citizen opposition, the process of developing new housing is inefficient. We formalize the value of information and show the conditions under which local meetings improve efficiency and lead to more housing. Our model identifies the heretofore under-explored benefits of unrepresentative public meetings and indicates where policy interventions intending to improve the supply of housing should (lawsuits) and should not (representativeness of meetings) be targeted.

Local Planning Meetings and NIMBYism

We further a substantive literature on housing and NIMBY-ism, which explores the impact of public meetings and litigation on housing supply and development (Einstein, Glick, and Palmer 2020;

Foster and Warren 2022; Brouwer and Trounstine 2024). Despite sky-rocketing housing prices and chronic shortages, developers proposing new housing projects and residents are at odds, even in liberal cities (Brouwer and Trounstine 2024; Ornstein 2023). Much like the model of the strategic interaction between developers and local residents in Foster and Warren (2022), we formalize the incentives of developers, proposing a housing project of variable size, and the citizens who may be opposed to such development. We focus, however, on the institution of a local meeting, even an unrepresentative one, as it can give the developer sufficient information about the opponent citizens to efficiently compromise.

Meetings are important aspects of local political processes, especially housing. Local public meetings, whether city council, planning, or school board, may simply inform citizens in attendance of local government processes and decisions but often include public participation; the opportunity for citizens' public comments– and deliberation between the public and the board– may increase engagement, attendance, and trust in local officials (Collins 2021). In general, meeting attendance and comments are not representative of the community at large and not representative of all who would be impacted by the proposal, such as prospective residents in the case of housing. Planning and zoning meeting attendees tend to be wealthier, whiter, older, and more likely to be male, homeowners, and local voters than the larger communities in which projects are proposed (Einstein, Palmer, and Glick 2019; Einstein, Glick, and Palmer 2020; Einstein et al. 2023; see also Turner and Weninger 2005, on other types of meetings). Meeting comments are dominated by opposition (Einstein, Glick, and Palmer 2020), especially those who would shoulder the concentrated costs of living near a new development (de Benedictis-Kessner and Hankinson 2019; Hankinson and de Benedictis-Kessner 2022; Marble and Nall 2021). Setting aside normative concerns of unrepresentative local "democracy," the lack of participation by neutral and supportive residents (and prospective residents) has been a proffered explanation for the chronic under-supply of housing.

We take an alternative view of unrepresentative local meetings. The fora are not for aggregating the preferences of all impacted by potential development and coming to a procedurally democratic consensus, but for the strongest opponents of the proposal to make themselves, and their preferences, known. Adams (2004) argues that city council and school board meetings in Santa Ana, California, are used not to directly influence governmental decisions, but for citizens to convey information to officials. In the arena of housing, local officials and planning/zoning board members are not the only actors present to garner information about citizen attendees. Meetings offer developers an opportunity to learn more precisely the preferences of a proposed projects' opponents and complete studies and make other adjustments to "mollify neighbors' concerns" (Einstein, Glick, and Palmer 2020, p. 3). This information is vital to developers hoping to complete projects in a timely, cost-effective manner due to the credible threat of citizen lawsuits.

Regardless of a meeting's procedures or outcome, development of new housing can be delayed by opposed citizens through litigation. Judicial "tastes" have become more friendly to anti-development sentiments since the 1970s, making lawsuits serious concerns for developers (Fischel 2004; Glaeser, Gyourko, and Saks 2005). Courts accept "junk science" traffic analyses as evidence sufficient to deny a project entirely or impose costly traffic mitigation expectations on developers (Currans and Stahl 2024). Even lawsuits that are "frivolous" and ultimately unsuccessful force developers to pay legal costs and delay projects; building permits may expire while a case moves through the system (Einstein, Glick, and Palmer 2020, pp. 27, 123). In their case study of the development of St. Aidan's Church in Brookline, Massachusetts, Einstein, Glick, and Palmer (2020) show how impactful anti-development lawsuits can be: after hiring a Canon Lawyer to challenge the Archdiocese's ability to sell the property, seventy-five residents filed a lawsuit that, while eventually settled, delayed the project by a year. The increasing legalization of, and judicial involvement in, urban planning extends outside the United States to municipalities in Latin America, India, and South Africa (Bhan 2009; Sotomayor, Montero, and Ángel-Cabo 2023; Taylor 2020).

Anti-development citizen litigation can take many forms, but environmental lawsuits are common, especially in California under the California Environmental Quality Act (CEQA). Environmental suits have both been maligned as special interest-driven allowances for private individuals and organizations to (attempt to) enforce laws without political accountability and lauded as a vital component of environmental administration, especially in regions and under presidencies with anti-environmental ideological leanings (Adelman and Glicksman 2019). While the requirements for standing in state court may vary from state to state, the right to sue is often contingent on the plaintiff exhausting all administrative remedies prior to filing (Lahav 2016). If lawsuits are, most simply, a peaceful means of resolving disputes, other administrative proceedings that may similarly resolve the dispute must be utilized prior to proceeding to court (Lahav 2016). Public meetings fill this role of an administrative avenue of dispute resolution: if a public meeting regarding a new development is held, a potential litigant must attempt to resolve their dispute with the developer (or municipality allowing the development) there. An anti-development lawsuit may immediately be thrown out if the plaintiff did not exhaust all administrative remedies– did not attend the meeting.

In addition to limiting potential citizen litigation by forcing prospective plaintiffs to use an administrative proceeding before suing, housing meetings transmit information to developers. As a public forum in which the most opposed neighbors voice their concerns, meeting attendance can indicate citizens' willingness to pursue litigation against developers (Einstein, Glick, and Palmer 2020, p. 122). Meetings attended by the opposition allow developers to identify who the potentially litigious neighbors are and the demands they have for the project. Many housing and local democracy advocates argue that making meetings less costly to attend (in time and location) and conducting wider outreach to bring more than immediate neighbors to meetings can make planning and zoning meetings more representative of both proponents and opponents of housing development (Einstein, Glick, and Palmer 2020; see also Cole and Caputo 1984, on meetings). However, making meetings more representative of the larger community, as well as including prospective residents, would dilute the information on their staunchest opponents that the developers need to make efficient compromises.

In the purview of formal theory, our work speaks to a political economy literature on information aggregation (Ali et al. 2008) and meetings as tools of preference aggregation (Aragonès and

Sánchez-Pagés 2009; Collins 2021; Osborne, Rosenthal, and Turner 2000), but we take significant departures by embedding meetings as a source of information in a proposer-litigation setting rather than the direct mechanism for decision making. This generates qualitatively different predictions about who attends meetings and the value of the meeting itself. The inefficiency of uncertainty in our model is similar to canonical work in international relations on bargaining and war. Two sides prefer a negotiated settlement to fighting, as war is costly; however, with private information on their own power and willingness to fight, wars with deadweight loss can still occur rationally (Fearon 1995). Much like war, lawsuits, with their associated legal fees and delays, are costly and if the developer has information about the opposed citizens' willingness to sue, an efficient compromise can be made.

A Model of Citizen Lawsuits

We model the interactions between a developer (D, pronoun "she") who has chosen to develop a new housing site that has $n > 1$ neighbors, each of whom we index by $i \in \{1, 2, ..., n\}$ (pronoun "he" or "they"). The developer chooses the size of the project, $s_D \geq 0$. One or more neighbors can decide to sue the developer to try to reduce the size of the developer's project to a court-imposed project size of $s_0 \geq 0$, which is strictly smaller than the developer's proposal of s_D.

Preferences over housing We will denote the final size of the project as s, which can be either s_D (the developer's proposal) or s_0 (the court-imposed project size). We assume each neighbor i gets a benefit from new housing, which is increasing in s and captures the idea that new housing is economically beneficial for the community in which it is built (Li 2021; Asquith, Mast, and Reed 2023; Mast 2023). However, we also assume each neighbor is negatively impacted by new housing and suffers a cost that is also increasing in s, capturing the idea that new housing could increase traffic, introduce new shadows, change the character of the neighborhood, or impact adjacent property values.

To capture this trade-off, we represent each neighbor i's preferences over a project s with the following payoff function:

$$v_i(s, \alpha_i) = bs - \alpha_i s^2$$

where $b \geq 0$ and $\alpha_i > 0$. This utility function is strictly concave, single-peaked and symmetric around ideal point $\hat{s}_i(\alpha_i) \equiv b/2\alpha_i \geq 0$. While neighbor i may prefer *some* housing (i.e., when $b > 0$ so that $\hat{s}_i > 0$), the costs created by new housing generally mean that he prefers relatively smaller projects than the developer (or a non-neighboring resident of the community) would prefer.

Importantly, we assume each neighbor faces a potentially different trade-off between the benefit and cost of new housing. We formalize this by allowing α_i to vary across neighbors, capturing the idea that some neighbors are closer to a development site and may face more severe costs, or that some neighbors are simply more prone to be NIMBYs than others. While heterogeneity among neighbors is important for our core arguments, we formalize this notion in the simplest manner possible, assuming that each neighbor i's cost can take one of two values: $\alpha_i \in \{\underline{\alpha}, \overline{\alpha}\}$, where $\overline{\alpha} > \underline{\alpha} > 0$.[1] We will often refer to a neighbor i with $\alpha_i = \overline{\alpha}$ as a "high cost neighbor" or simply "high type" and a neighbor i with $\alpha_i = \underline{\alpha}$ as a "low cost neighbor" or simply "low type."

Each neighbor i knows their own cost, but this is private information. Formally, we assume that each α_i is independently drawn from a commonly known distribution where $\Pr(\alpha_i = \underline{\alpha}) = \rho$ for all i. Substantively, the probability $1 - \rho$ captures the prevalence of extreme NIMBYs in the community around the proposed development site. It can roughly be understood as a measure of the community's NIMBYism.

We assume the developer strictly prefers larger projects up to a point where they become too costly. Accordingly, we assume her payoff from a development of size s is $v_D(s)$, where v_D is a

1. In the analysis below, the discrete type space for the neighbors induces a discrete action space for the developer since it allows to eliminate a wide range of project sizes that are strictly dominated.

single-peaked function that is weakly concave and has an ideal point that is strictly higher than all of the neighbors, $\hat{s}_D > \hat{s}_i$ for all i.

In our analysis, we will limit our attention to situations in which a successful lawsuit against a developer by a neighbor would reduce the size of the project substantially—to a level that is lower than each neighbor's ideal point. Practically speaking, this reduces the number of cases we have to consider below.[2] This assumption encompasses the situation in which a successful lawsuit blocks a project all together, $s_0 = 0$, either directly or by altering the project so much that it is no longer economically feasible for the developer.

> **Assumption 1** (court-imposed project size)**.** The court-imposed project size is sufficiently small so that all neighbors weakly prefer a higher project size than the court-imposed project size. Formally, $s_0 \leq \hat{s}_i(\alpha_i)$ for all i.

Of course, since the developer prefers more housing than each neighbor, it follows that $s_0 < \hat{s}_D$.

Lawsuits We model lawsuits in a simple and tractable manner. Each neighbor independently decides whether to sue the developer, and we will let $\tau_i \in [0, 1]$ indicate the probability that player i sues. If the developer wins all of the lawsuits, then she implements her project of size s_D. However, if she loses any one of the lawsuits, the court reduces the size of the project from the developer's proposal of s_D to the default value of $s_0 < s_D$. Lawsuits always *reduce* the size of projects if they are successful, capturing the idea that a developer cannot be legally compelled (via a lawsuit) to produce larger projects than she wants.

We assume that each neighbor who decides to sue wins their case with probability ω, where all lawsuits are independent draws from the distribution over potential case outcomes. Then, the

2. Note to reader: we are pretty sure we can relax this assumption without changing our results, and will do so for the final draft. Of particular interest will be the case where $\hat{s}_i = 0 < s_0$, so that neighbor i wants *no* development and the court-imposed project size is a compromise between what the neighbor wants and what the developer wants.

probability that any lawsuit is successful at reducing the developer's proposal is given by:[3]

$$\gamma(\boldsymbol{\tau}) = 1 - \prod_{i=1}^{n}(1 - \omega\tau_i)$$

where $\boldsymbol{\tau} = (\tau_1, \tau_2, ..., \tau_n)$ is a vector indicating each neighbor's probability of suing.

Lawsuits are costly to both the suing neighbors and the developer. If neighbor i files a lawsuit, he pays a cost $\kappa_i > 0$. For now, we assume this cost equal for all neighbors, so that $\kappa_i = \kappa_P$ for all i.[4] If any neighbor files a lawsuit, then the developer pays a cost $\kappa_D > 0$. These costs can represent both the direct costs of litigation (e.g., attorneys, court fees, etc.) as well as opportunity costs (e.g., resources diverted away from developing other projects).[5]

Sequence and solution concept Summarizing the discussion above, we can now write out the game sequence.

1. The developer makes a proposal of size $s_D \geq 0$.

2. For each neighbor i, their cost α_i is independently drawn and privately revealed to that neighbor.

3. Each neighbor i simultaneously decides whether to sue the developer, $\tau_i \in [0, 1]$.

4. The outcome of each lawsuit is drawn, where with probability $1 - \omega$ the developer wins and probability ω the developer loses.

5. If the developer is not sued or if she wins all lawsuits, then project s_D is implemented; otherwise project s_0 is implemented.

---

3. An alternative way to model lawsuits is to assume that the court merges lawsuits into one consolidated lawsuit that the developer loses with probability $\gamma(k)$, where k is the number of litigants who sue, and γ increases in k.

4. When we introduce public meetings to our model, this cost may vary across neighbors. See Assumption 6.

5. We assume the developer pays a single fixed cost if at least one lawsuit is filed. Our results would be qualitatively similar if we instead assumed the developer paid this cost for each lawsuit filed, but the analysis would be more cumbersome and require more notation.

6. Each player receives a total payoff equal to their payoff from the final project size s, less the cost of a lawsuit (if they were involved in one).

An equilibrium of this game will consist of a strategy profile in which each player is playing sequentially rational strategies, and beliefs that are updated via Bayes' rule when possible. Throughout our analysis, we will make assumptions about players' beliefs that will simplify our analysis, and we will be explicit about these assumptions.

Potential Coordination Among Neighbors

A key feature of our model is that there are multiple neighbors, each of whom makes their own decisions about how to respond to the developer's proposal. This raises the possibility that neighbors may consider other neighbors' actions when making their own decisions. These coordination motives can cut in two directions.

First, and perhaps most interestingly, each neighbor may have an incentive to free ride on the effort of other neighbors. If, for example, neighbor i believes several other neighbors are planning to sue, then this will disincentivize i from paying a cost to file their own lawsuit. On the other hand, coordination among neighbors could be attractive insofar as it enables them to either increase their probability of winning or share the costs of litigation. Because of these two competing dynamics, it is not obvious whether coordination motives would increase or decrease neighbors' willingness to sue a developer.

At a general level, while the coordination dynamics among neighbors are surely important, we largely set them aside in our current analysis since our primary focus is on the developer's uncertainty about the extent of local resistance to her proposed project. Indeed, coordination is complicated to model and we will assume away the possibility of collective and/or "class action" lawsuits. However, since the incentive to free ride potentially cuts against some of our core arguments, we will parameterize this incentive in a tractable way by making the following behavioral assumption about the neighbors in our model.

> **Assumption 2** (behavioral neighbors)**.** Each neighbor i believes that every other neighbor $j \neq i$ will sue with probability $\tilde{\tau} \in [0, 1]$. This belief is fixed and need not be correct in equilibrium.

The parameter $\tilde{\tau}$ gives us a flexible and simple way to incorporate each neighbor's incentive to free-ride. As $\tilde{\tau}$ increases, the coordination problem worsens. But we can also "shut down" coordination motives all together by simply setting $\tilde{\tau} = 0$. Except for this particular assumption, elsewhere in our model we will make the standard rationalist assumption that players form correct beliefs via Bayes' rule based on other players' equilibrium best responses (when possible).

To make our notation more concise going forward, we will let $\tilde{\gamma}_i^{\tau_i}(n)$ indicate neighbor i's belief about the probability of a successful lawsuit when he sues with probability τ_i and he believes all other $n-1$ neighbors $j \neq i$ sue with probability $\tau_j = \tilde{\tau}$. Then:

$$\tilde{\gamma}_i^{\tau_i}(n) = 1 - (1 - \omega\tau_i)(1 - \omega\tilde{\tau})^{n-1}$$

Suing the Developer

Given Assumption 2, a neighbor i finds it weakly better to sue ($\tau_i = 1$) than not sue ($\tau_i = 0$) if

$$\underbrace{(1 - \tilde{\gamma}_i^1(n))v_i(s_D, \alpha_i) + \tilde{\gamma}_i^1(n)v_i(s_0, \alpha_i) - \kappa_P}_{\text{expected utility from suing}} \geq \underbrace{(1 - \tilde{\gamma}_i^0(n))v_i(s_D, \alpha_i) + \tilde{\gamma}_i^0(n)v_i(s_0, \alpha_i)}_{\text{expected utility from not suing}}$$

Rearranging this condition and noting that $\tilde{\gamma}_i^1(n) - \tilde{\gamma}_i^0(n) = \omega(1 - \omega\tilde{\tau})^{n-1}$, the condition simplifies to:

$$\underbrace{(v_i(s_0, \alpha_i) - v_i(s_D, \alpha_i))}_{\substack{\text{Net value to } i \text{ from a successful} \\ \text{lawsuit reducing the project size}}} \geq \underbrace{\frac{\kappa_P}{\omega(1 - \omega\tilde{\tau})^{n-1}}}_{\substack{\text{Cost of litigation weighted} \\ \text{by increased probability} \\ \text{of successful lawsuit}}} \qquad (1)$$

This condition demonstrates that when deciding whether or not to file a lawsuit, a neighbor i must consider three factors: the net value of a successful lawsuit, $v_i(s_0, \alpha_i) - v_i(s_D, \alpha_i)$, the increase in the probability that a lawsuit will be successful if he files his own lawsuit, $\omega(1 - \omega\tilde{\tau})^{n-1} = \tilde{\gamma}_i^1(n) - \tilde{\gamma}_i^0(n)$, and the cost of the lawsuit, κ_P. In the next result, we show that neighbor i is weakly disincentivized to sue if the developer's proposal is not too far from the neighbor's ideal point.

> **Lemma 1.** For a neighbor i, there exist a threshold, $\tilde{s}(\alpha_i) > 0$ such that i will not sue for all $s_D \in [s_0, \tilde{s}(\alpha_i))$, is indifferent between suing and not suing if $s_D = \tilde{s}(\alpha_i)$, and will sue if $s_D > \tilde{s}(\alpha_i)$.

Proof of Lemma 1. This follows from condition (1). The right hand side of this condition is strictly positive and constant in s since $\kappa_P > 0$. The left hand side is a bit more complicated to characterize since v_i is non-monotonic in s. First, for $s_D < \hat{s}_i(\alpha_i)$, v_i is strictly increasing and the left hand side is therefore strictly decreasing. The left hand side is weakly positive at $s_D = 0$, and strictly negative at $s_D = \hat{s}_i(\alpha_i)$, implying there exists some threshold $s'_D \in [0, \hat{s}_i(\alpha_i))$ such that (1) holds for all $s_D \leq s'_D$. Moreover, note that from the condition $s_0 > s'_D$, indicating that any lawsuit for a proposal $s_D < s'_D$ would *increase* the developer's proposal. Since we rule that out by assuming $s_D \geq s_0$, then there are no proposals $s_D \leq s'_D$ where a lawsuit occurs. Second, for $s_D > \hat{s}_i(\alpha_i)$, v_i is strictly decreasing and the left hand side is therefore strictly increasing with no upper bound. Moreover, at $s_D = \hat{s}_i(\alpha_i)$, the left hand side is negative, implying there exists some threshold $\tilde{s}(\alpha_i) > \hat{s}_i(\alpha_i)$ such that (1) holds for all $s_D \geq \tilde{s}(\alpha_i)$. For a neighbor i, these two cases imply that it will be weakly optimal for him not to sue if $s_D \in [s_0, \tilde{s}(\alpha_i)]$. □

Note that when condition (1) binds, neighbor i is indifferent between suing and not suing. If neighbor i sues when indifferent, this creates an open set problem for the developer, and there can be no equilibrium in which i sues when indifferent. Going forward, we will assume that when a developer's proposal makes a neighbor indifferent between suing and not suing, he does not sue.

In Figure 1, we plot a neighbor i's utility over project sizes, illustrating the range of proposals for which he is unwilling to file a lawsuit, given condition (1) above.

Figure 1: We plot neighbor i's utility over project sizes to illustrate his strategic calculus over when to sue the developer, given that the court-imposed project size is s_0. Note that the default size must be smaller than the developer's proposal, so $s < s_0$ is not feasible, and accordingly greyed out.

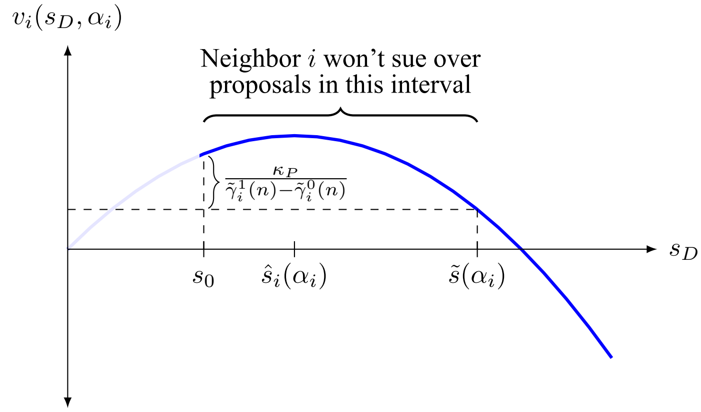

The costlier a neighbor finds new housing, the smaller this interval is. In Lemma A.1 in the Online Appendix, we demonstrate that a neighbor with lower costs is willing to accept a larger housing development project than a neighbor with higher costs. Going forward, let $s_1 \equiv \tilde{s}(\overline{\alpha})$ and $s_2 \equiv \tilde{s}(\underline{\alpha})$.

Setting aside the possibility that the developer makes a proposal smaller than the default value s_0 (which would be ruled out by Assumption 1), the thresholds we defined above allow us to characterize which neighbors will find it optimal to sue for which kinds of proposals.

Lemma 2. Let s_D be the developer's proposal, and assume $s_D > s_0$. Then,

- If $s_D \in (s_0, s_1]$, none of the neighbors sue.
- If $s_D \in (s_1, s_2]$, all neighbors with high costs ($\overline{\alpha}$) sue.
- If $s_D \in (s_2, \infty)$, all neighbors sue.

Proof of Lemma 2. In text. □

The Developer's Proposal

When she is developing her proposal, the developer does not (yet) know if a lawsuit will occur. For a proposal s_D, let $\tilde{\gamma}_D(s_D)$ be her belief that there will be a successful lawsuit, and $\tilde{\sigma}_D(s_D)$ be her belief that there will be any lawsuit (successful or not). Then, her ex ante expected utility from making a proposal s_D is:

$$U_D(s_D) = (1 - \tilde{\gamma}_D(s_D))v_D(s_D) + \tilde{\gamma}_D(s_D)v_D(s_0) - \tilde{\sigma}_D(s_D)\kappa_D$$

We are most interested in situations where there is substantial conflict between the developer and the neighbors over the ideal size of the project. We will make the following assumption, which ensures the developer's ideal project size is always large enough that every neighbor would find it optimal to sue if the developer were to propose her ideal project, $\hat{s}_D$.

> **Assumption 3** (preference conflict)**.** The developer's ideal project size is sufficiently large that every neighbor would prefer to sue than accept a project of that size. Formally, $\hat{s}_D > s_2 > s_1$.

With this assumption we can make the following statement about the developer's induced choice set.

Lemma 3. In equilibrium, the developer chooses a project size $s_D \in \{s_1, s_2, \hat{s}_D\}$.

Proof. Given Assumption 3 and the fact that v_D is single-peaked around $\hat{s}_D$, then v_D is strictly increasing for all $s \in [0, \hat{s}_D)$. Then, for any generic right-closed interval S_k such that $S_k \subseteq [0, \hat{s}_D)$ and the probability of a (successful) lawsuit is fixed, it is strictly dominated for the developer to choose any $s_D < \max S_k$. Finally, since Lemma 2 defines three intervals in which each the probability of a lawsuit is constant, if $s_D > s_0$, then it is strictly dominated for the developer to choose any s_D such that $s_D \notin \{s_1, s_2, \hat{s}_D\}$. Moreover, note that for all $s_D \leq s_0$, $v_D(s_D) < v_D(s_0)$, and there is a weakly greater probability of a lawsuit, costing κ_D. □

Lemma 3 means that the developer's decision boils down to choosing a project size based on her expectation about the number of lawsuits it will provoke. There are three scenarios to consider.

- **No compromise.** First, consider if the developer proposes her ideal project. We refer to this as "no compromise." By Assumption 3, it follows that both types sue. The probability of a lawsuit is $\tilde{\sigma}_D(\hat{s}_D) = 1$, and the probability of a success is $\tilde{\gamma}_D(\hat{s}_D) = 1 - (1-\omega)^n$. In this case, her expected payoff is:

$$U_D(\hat{s}_D) = (1-\omega)^n v_D(\hat{s}_D) + (1-(1-\omega)^n)v_D(s_0) - \kappa_D$$

- **Partial compromise.** Next, consider if she proposes a project that satisfies the low types, but not the high types. We refer to this as "partial compromise" that generates a "medium" sized project, s_2. Then the probability of a lawsuit is $\tilde{\sigma}_D(s_2) = 1 - \rho^n$ and the probability of a successful lawsuit is $\tilde{\gamma}_D(s_2) = 1 - (1-\omega(1-\rho))^n$. In this case, her expected payoff is:

$$U_D(s_2) = (1-\omega(1-\rho))^n v_D(s_2) + (1-(1-\omega(1-\rho))^n)v_D(s_0) - (1-\rho^n)\kappa_D$$

- **Full compromise.** Finally, consider if she proposes a project that satisfies both the low types and the high types. We refer to this as "full compromise" that generates a "small" sized project, s_1. In this situation, there is no lawsuit, so $\tilde{\sigma}_D(s_1) = \tilde{\gamma}_D(s_1) = 0$. In this case, her expected payoff is:

$$U_D(s_1) = v_D(s_1)$$

Whether the developer chooses a project smaller project size to avoid a lawsuit depends on both the cost of a lawsuit and how she evaluates the lottery associated with a lawsuit. Intuitively, if lawsuits are especially costly or the expected project size induced by a lawsuit is sufficiently bad, then the developer will find the full compromise appealing.

Lemma 4. Given Assumption 1, the developer always prefers to offer a full compromise, $s_D = s_1$ if:

(i) $\kappa_D \geq \overline{\kappa}_D$ where $\overline{\kappa}_D > 0$ is defined in the proof; or

(ii) $s_1 > (1-\omega(1-\rho))^n s_2 + (1-(1-\omega(1-\rho))^n)s_0$.

Moreover, there are regions of the parameter space where the condition in part (ii) holds.

Proof of Lemma 4. First note that as $\kappa_D \to \infty$, s_1 is strictly preferred to s_2 or $\hat{s}_D$ since $U_D(s_1)$ is a finite real number while $U_D(s_2)$ and $U_D(\hat{s}_D)$ approach $-\infty$. Moreover, $U_D(s_2)$ and $U_D(\hat{s}_D)$ are both strictly decreasing in κ_D. This implies that there exists some $\tilde{\kappa}_D$ such that $U_D(s_1) > \max\{U_D(s_2), U_D(\hat{s}_D)\}$ for all $\kappa_D > \tilde{\kappa}_D$. Of course, this $\tilde{\kappa}_D$ is not constrained to be strictly positive, if for example, the developer prefers the full compromise for all $\kappa_D > 0$. So, define $\hat{\kappa}_D = \max\{\tilde{\kappa}_D, 0\}$.

We now prove part (ii). The first claim in part (ii) is easy to show. Given that v_D is weakly concave, it follows that if $s_1 > (1-\omega(1-\rho))^n s_2 + (1-(1-\omega(1-\rho))^n)s_0$, then

$$U_D(s_1) = v_D(s_1) > \underbrace{(1-\omega(1-\rho))^n v_D(s_2) + (1-(1-\omega(1-\rho))^n)v_D(s_0)}_{U_D(s_2)+(1-\rho^n)\kappa_D} > U_D(s_2)$$

We now prove the second claim in part (ii) by showing an example. First, note the condition is:

$$\begin{aligned}
&\frac{b}{2\overline{\alpha}} + \frac{1}{2\overline{\alpha}}\sqrt{b^2 + 4\overline{\alpha}\left(\frac{\kappa_P}{\omega(1-\tilde{\tau}\omega)^{n-1}}\right) - 4\overline{\alpha}v_i(s_0,\overline{\alpha})} > \\
&\qquad (1-\omega(1-\rho))^n\left(\frac{b}{2\underline{\alpha}} + \frac{1}{2\underline{\alpha}}\sqrt{b^2 + 4\underline{\alpha}\left(\frac{\kappa_P}{\omega(1-\tilde{\tau}\omega)^{n-1}}\right) - 4\underline{\alpha}v_i(s_0,\underline{\alpha})}\right) \\
&\qquad + (1-(1-\omega(1-\rho))^n)s_0
\end{aligned}$$

Suppose if a lawsuit is successful, the project doesn't get built, $s_0 = 0$. This then reduces to:

$$\overline{\alpha} < \left[\frac{1}{(1-\omega(1-\rho))^n}\right]\underline{\alpha}$$

Since the term in the square brackets is strictly greater than one, then there exist values of $\overline{\alpha} > \underline{\alpha}$ such that the condition holds. □

The previous result is an existence result insofar as it demonstrates that it is possible that the developer will prefer to propose a full compromise instead of risking a lawsuit. Because we are interested in examining how developer's aversion to lawsuits shapes their incentives around meetings, we will assume going forward that when facing the risk of a lawsuit, the developer prefers a full compromise that forecloses a lawsuit. We will refer to a developer like this as "lawsuit averse."

Assumption 4 (developer lawsuit aversion)**.** The developer is lawsuit averse. Formally, $U_D(s_1) > \max\{U_D(s_2), U_D(\hat{s}_D)\}$.

Finally, we can characterize the equilibrium of the model of citizen lawsuits satisfying the scope conditions we've identified in our assumptions above.

Proposition 1. In the model of citizen lawsuits in which the court-imposed project size is very small (Assumption 1), neighbors don't fully account for other neighbors' lawsuits (Assumption 2), and the developer wants both very large projects (Assumption 3) and is averse to lawsuits (Assumption 4), then there is a unique equilibrium in which the developer proposes a full compromise project, $s_D^* = s_1$ and a neighbor i uses the following lawsuit strategy:

$$\tau_i^*(s_D, \alpha_i) = \begin{cases} 0 & \text{if } s_D \in [s_0, \tilde{s}(\alpha_i)] \\ 1 & \text{otherwise} \end{cases}$$

Proof of Proposition 1. In text. □

The crucial assumption for what follows is the last one—Assumption 4—which ensures that the developer would rather avoid lawsuits. If, on the other hand, there are circumstances in which she would be willing to face a lawsuit over a larger project, then the informational value of meetings will be lessened.

The Value of Information

Lawsuits are both costly to and risky for a developer. When she is lawsuit averse, she makes a full compromise and proposes s_1. In effect, she presumes there will be strong opposition, and proposes a substantially smaller project in line with this presumption. Maybe she would be better off if she were able to learn about the extent of opposition *before* making her proposal? If she did, she would sometimes learn there is actually only mild opposition and she could propose a larger project. Learning about potential opposition in advance would therefore give her an expected utility of:

$$\underbrace{v_D(s_2)\rho^n}_{\substack{\text{Learns there are}\\\text{only low types}}} + \underbrace{v_D(s_1)(1-\rho^n)}_{\substack{\text{Learns there are}\\\text{some high types}}}$$

Then, the value of information to the developer would be

$$\underbrace{v_D(s_2)\rho^n + v_D(s_1)(1-\rho^n)}_{\substack{\text{utility with}\\\text{information}}} - \underbrace{v_D(s_1)}_{\substack{\text{utility without}\\\text{information}}} = \rho^n(v_D(s_2) - v_D(s_1))$$

Her value from information is therefore decreasing in n and increasing in ρ. In very large communities (large n), there is a higher probability of a high type, making the value of information lower. But, if the probability of high types is fairly low (high ρ), then the value of information is higher. Finally, if the high types have really high costs ("extreme NIMBYs") or the low types have really low cots, then $s_2 - s_1$ will be larger and the value of information is higher.

In addition to the developer being better off with information about potential opposition, access to this information also increases the expected project size since she can perfectly calibrate her proposal to avoid lawsuits:

$$E(s) = s_2\rho^n + s_1(1-\rho^n) > s_1$$

We have demonstrated, at least at an abstract level, that information is valuable to the developer and increases the supply of housing. But how does she acquire information? And is the information she acquires good information? In the next section, we develop a model of planning meetings that explore these questions in detail.

Planning Meetings

In the previous section, we characterized the conditions under which the developer is better off knowing the degree of potential opposition among the neighbors. Given that lawsuits generate an inefficiency, the developer will often find it valuable to know in advance whether she is likely to be sued.

In many jurisdictions in the United States, local governments require public meetings to review many (if not all) housing proposals made by developers. These meetings provide a forum for any member of the community (such as negatively impacted neighbors) to provide comment on the proposals under consideration. A public meeting thus creates an opportunity for the developer to learn whether there are neighbors in the community that strongly oppose her project.

However, two issues complicate this story. First, which neighbors actually have an incentive to show up to these meetings? And second, how much does a developer actually learn learn from these meetings? We now extend our baseline model to include the possibility of a planning meeting in order to explore these issues.

Suppose that the local government adopts an ordinance requiring a public meeting to review every housing proposal made by a developer. The developer now has an opportunity to revise her initial proposal in response to the feedback in the meeting, but we assume she can only revise it down.[6] As before, we denote the post-meeting (i.e., final) proposal by s_D. We now denote the developer's pre-meeting (i.e., provisional) proposal by s_M, where $s_D \leq s_M$.

Each neighbor decides whether to attend the public meeting, $a_i \in \{0, 1\}$, where $a_i = 1$ indicates attendance. We will let $a = \sum_{i \in R} a_i$ be the total number of meeting attendees. In our analysis of planning meetings, we will simplify and assume each neighbor i of type $\alpha_i \in \{\underline{\alpha}, \overline{\alpha}\}$ uses the same pure strategy when deciding to attend. Moreover, when we consider deviations to establish the existence of equilibria, we will consider only deviations by type, not by individual neighbors. The

6. If she tried to revise it up, for example, she would have to attend another meeting, starting the whole process over again.

latter assumption presents the hardest case for our analysis since deviations become more profitable for neighbors (and thus more tempting).

Assumption 5 (type-level strategies)**.** In every equilibrium, all neighbors of a type use the same pure strategy. Moreover, when considering potential deviations from equilibrium behavior, we will assume type-level deviations.

While meetings provide an opportunity for the developer to learn about potential opposition to her proposal, they are not perfectly informative. For example, time limits on public meetings may prevent some strong opponents from publicly registering their objections. More generally, the information provided at meetings may be fairly "noisy," which can prevent the developer from perfectly learning the types of all the attendees. We will model this idea in a simple manner, assuming that at the meeting the types of only m attendees are publicly revealed (to both the developer and all neighbors). If the number of attendees exceeds m, then the each of the attendees has an equal probability of being one of the m attendees whose types are publicly revealed.

Given a neighborhood of size n, and a meeting constraint of m, and the attendance strategy used by each type, $a_i(\underline{\alpha})$ and $a_i(\overline{\alpha})$, then the probability of observing all low types in a meeting is:[7]

$$\lambda(n, m, a_i(\underline{\alpha}), a_i(\overline{\alpha})) = \underbrace{\rho^n a_i(\underline{\alpha})}_{\substack{\text{Probability all} \\ \text{neighbors are} \\ \text{low types and} \\ \text{in attendance}}} + \underbrace{\sum_{z=m}^{n-1} \binom{n-1}{z} \rho^z (1-\rho)^{n-1-z} \left(\frac{z a_i(\underline{\alpha})}{z a_i(\underline{\alpha}) + (n-z) a_i(\overline{\alpha})} \right)^m}_{\substack{\text{Probability of mistakenly observing all low types at} \\ \text{the meeting when there is a high type in attendance}}}$$

Each attendee pays a cost $\kappa_M > 0$ to attend.[8] Showing up gives the developer more information about what types exist in the community. In particular, the developer can learn if she is facing only

7. For example, suppose there is a neighborhood of size $n = 3$ where all neighbors attend the meeting, but only $m = 2$ neighbors are observed by the developer. Then,

$$\lambda(3, 2, 1, 1) = \left[\binom{3}{2} \rho^2 (1-\rho)^1 \left(\frac{2}{3} \right)^2 \right] + \left[\binom{3}{3} \rho^3 (1-\rho)^0 \left(\frac{3}{3} \right)^2 \right] = \frac{(4-\rho)\rho^2}{3}.$$

8. Since the meeting is about developer's proposal, she must show up. Accordingly, her cost is fixed and plays no role in the analysis. For clarity, we omit it.

low types. What do residents get? In many jurisdictions, it is more difficult to win a lawsuit against a developer if objections are not initially raised in a public meeting. This legal principle is known as "exhaustion of administrative remedies." We will keep things simple and make the following assumption about the value of meeting attendance for future litigation.

Assumption 6 (exhaustion of administrative remedies)**.** A neighbor i cannot sue the developer over her proposed project s_M if he does not attend the public meeting on that proposed project. Formally, for neighbor i, $\kappa_i \to \infty$ if $a_i = 0$.

In many jurisdictions, members of the public may also get a chance to persuade local officials—such as a planning board—to administratively reject a proposal, allowing neighbors to prevail against a developer without going to court. While this is obviously important (and a promising direction for future research), we set it aside here and treat the planning meeting solely as a forum for public comment.

Each neighbor's calculation about whether to attend the meeting will depend on what kind of (potentially revised) project he thinks will result from the meeting. In particular, he will consider what types of neighbors are likely to attend the meeting, and how the developer will respond to those types. We will make the following assumption about the preferences of the low cost neighbors, which ensures that (1) there is "sufficient" alignment between their interests and the interests of the high cost neighbors, and (2) they prefer more housing than what the court would impose if a neighbor wins a lawsuit.

Assumption 7. The low cost neighbors strictly prefer project size s_1 over any larger project $s' > s_1$ and over the court-imposed size of s_0.

Assumption 7 is sufficient (although not necessary)[9] for preventing a scenario in which low cost neighbors and high cost neighbors use the planning meeting to compete over the project size. Without this assumption, the low cost neighbors may have an incentive to attend the meeting to

9. Note to reader: for a future draft, we hope to characterize necessary and sufficient conditions, but for our purposes, this is somewhat of a technicality. We simply seek to ensure that there is sufficient preference alignment between high and low types that.

"shout down" the high cost neighbors, and potentially undermining the developer's ability to learn about opposition to her proposed project. While this is an interesting case, it is beyond the scope of our present analysis. As a result, our model is best understood as a model of meetings that are dominated by opponents of the developer's proposal.

Given that the strategic calculus will be somewhat complex, we will make a technical—albeit reasonable—set of assumptions about the developer's beliefs off the equilibrium path.

> **Assumption 8** (off path beliefs)**.** We make the following assumptions about the developer's beliefs off the equilibrium path:
>
> (i) If any high type is observed at a meeting, then with probability 1, the developer believes there is a high type that could sue, and correctly anticipates that high type's strategic behavior in the lawsuit subgame.
>
> (ii) If no high types are observed off the equilibrium path, and this is unexpected for the developer, then the developer has a belief $\tilde{\eta} \in [0, 1]$ that there is at least one high type at the meeting who was not one of the m attendees that was observed.

Part (i) of this assumption ensures that the developer does not ignore (or forget) that she observed high types in a meeting when they are observed off the equilibrium path. Part (ii) parameterizes the developer's belief in situations where she expected no low types were going to attend the meeting but observes them anyway.

The developer's choice of her initial proposal, s_M will determine which neighbors show up to the meeting to complain. Her proposal could provoke all neighbors to attend, no neighbors to attend or just one type of neighbor to attend. Which neighbors attend determine how informative the meeting is for the developer.

Note to reader: In this draft, we consider a situation in which the developer is potentially able to learn whether there are high types in the community. In particular, for now we only study an equilibrium where the high cost neighbors attend the meeting and low cost neighbors do not. We believe this equilibrium demonstrates the most novel aspects of meetings in our setting: they serve to screen types and provide information to the developer. Of note, they do so *even when they are* *unrepresentative*. In later drafts, we will explore other equilibria, such as equilibria with representative meeting attendance, including questions of existence.[10]

Unrepresentative Planning Meetings

The most informative meeting, from the developer's perspective, are meetings where she anticipates only high types show up. If none do, then she can safely conclude that there actually aren't any high cost neighbors that are willing to sue her over her project. We now examine (1) whether an equilibrium exists where only high types show up to the meeting, and (2) what kinds of initial proposals provoke this kind of attendance.

If only high types attend the planning meeting, then there is either no attendance (implying all neighbors are low types), or all attendees the developer observes at the meeting are high types. So, if the developer observes any meeting attendance, she observes only high types and given Assumption 4, she revises her project down from s_M to s_1.

To induce high types to show up, they have to be convinced that their attendance will be worth it. Formally, the developer needs to propose an s_M that is large enough so that the following holds:

$$v_i(s_1, \overline{\alpha}) - \kappa_M \geq v_i(s_M, \overline{\alpha})$$

Given that $\kappa_M > 0$, and since v_i is declining in s_M, then there exists some proposal $s_M = \underline{s} > s_1$ where the condition binds. This is the minimum project size the developer could propose that would provoke the high cost neighbors to attend; although they also do so for any s_M that is strictly larger. If none do in the equilibrium, she gets to implement this project, which is better than if she preemptively compromised to s_1.

Suppose the developer makes a proposal of size $s_M \geq \underline{s}$ that induces the high types to attend (by

10. Note to reader: Our preliminary work suggests that the meetings where the developer preemptively compromises so that no neighbors attend the meeting might exist. However, meetings where all attend (i.e., "representative meetings") are dominated from the developer's perspective by the kinds of unrepresentative meetings that we study here.

satisfying the condition above). Then the key question for our analysis is whether it is sequentially rational for the low types not to attend. In the Online Appendix, we walk through the analysis in some detail, which we now summarize.

If a low cost neighbor does not go to the meeting, he knows that either project s_M or project s_1 will be implemented, depending on whether there are any high types who showed up to the meeting. Of course, if the low types would *prefer* a much larger project (e.g., s_2), then they might have an incentive to attend to try to convince the developer that only low types live in the neighborhood. But Assumption 7 rules this out and ensures that the low types would prefer smaller projects (e.g., s_1) since they're relatively aligned with the high types. Given that it's costly to attend the meeting, and the high types' attendance already induces a smaller proposal, there is no incentive for the low types to attend the meeting. If they did, they'd have to pay the cost of attendance, plus there is a chance that they cause the developer to wrongly conclude there are no high types, ensuring a larger project. In short, there is no upside for a low type to attend.

The developer understand this thought process and because she is better off with larger project sizes, in the unrepresentative meeting equilibrium, she'll propose a project $s_M \geq \underline{s}$ that induces the high types to show up to the meeting (and reveal themselves, if they exist), but that is not so high that it *also* provokes the low cost neighbors to show up. What is this level? In the Online Appendix, we show that it depends on the assumption we make about the developer's off path belief. This is a somewhat techincal point, but we demonstrate that there always exists a threshold s_M^{high} that just makes the low cost neighbor indifferent between attending and not attending, but is sufficiently high to induce attendance by high types. Importantly, it is always strictly greater that s_1 and sometimes it is even strictly higher than s_2.

Proposition 2. There is an unrepresentative meeting equilibrium in which all low cost neighbors do not attend and all high cost neighbors do attend. Formally for a neighbor i,

$$a_i^*(\alpha_i) = \begin{cases} 0 & \text{if } \alpha_i = \underline{\alpha} \\ 1 & \text{if } \alpha_i = \overline{\alpha} \end{cases}$$

In this equilibrium, the developer proposes $s_M = s_M^{\text{high}} > s_1$ and revises downward to $s_D = s_1$ if and only if she observes a high type at the meeting or holds off equilibrium path beliefs that there high types at the meeting who were not observed. Moreover, for some off-path beliefs, the developer's proposal is strictly higher than s_2: $s_M^{\text{high}} > s_2$.

Proof of Proposition 2. In text and in Online Appendix A.1. □

Consider the developer's payoff in this equilibrium. From an ex ante perspective, she gets the following lottery:

$$\rho^n v_D(s_M^{\text{high}}) + (1 - \rho^n)v_D(s_1)$$

Notice that the value of information in this context is strictly higher than it is in our baseline model without planning meetings, in which she proposes s_1 and receives $v_D(s_1)$. Moreover, when $s_M^{\text{high}} > s_2$, then unrepresentative meetings are actually better for the developer than the full information benchmark from above. This is because the developer uses the meeting to learn about opposition *and* it forecloses lawsuits by low cost neighbors who do not attend.

While this unrepresentative meeting is good for the developer (relative to a world with no meetings), the key issue for policy is that these meetings actually enable more housing. Again, a lawsuit averse developer would, in the absence of information, preemptively compromise and propose small projects (s_1) to prevent lawsuits. Unrepresentative meetings reveal to the developer whether there are high cost neighbors who are very likely to sue—when there aren't, she can propose bigger projects. Formally, the expected project size with unrepresentative meetings is larger than s_1

$$E(s) = \rho^n s_M^{\text{high}} + (1 - \rho^n)s_1$$

Somewhat subtly, unrepresentative meetings are better for the developer because what she really wants to know is if there are high cost neighbors. If the low type neighbors were to show up to the meeting as well, they'd make it more difficult for the developer to observe whether there are high cost neighbors. Given the increased risk of a lawsuit in a "representative" meeting, the developer is better off preemptively compromising down to a smaller project size. In this sense, unrepresentative meetings dominated by the strongest opponents enable more housing.

Conclusion

By situating public meetings in the full context of developer-citizen disputes, which includes the credible threat of citizen lawsuits, we show the informational value of unrepresentative public meetings for creating a more efficient development process and increasing the supply of housing. For policy makers and housing advocates that are looking for institutional solutions to the housing crisis, our model suggests that changing public meeting structures to improve attendance or make meetings more representative of the larger community may not yield more housing. Our perspective, that meetings are an informational channel by which developers can avoid the inefficiency of citizen lawsuits, points toward a different set of policy interventions that may be more fruitful for housing supply. In particular, changes to lawsuit procedures, whether through citizen costs or the reversion project size, would be impactful.

While our model addresses two institutional environs (lawsuits with and without meetings) as well as a variety of community sizes and project sizes, our analysis is somewhat limited. Meetings as a forum for conflict between neighbors– in which those less opposed or even supportive of a proposal attend a meeting to drown out the voices of the opposition– is outside the current scope. We have also set aside the meeting administrators' (i.e. the planning board) incentives and potential proposal rejection in order to focus on the conflict between developers and citizens. In future work, we will consider more aspects of meetings than we have so far. In addition to being an information channel between citizens and developers and an administrative remedy, public meetings provide an opportunity for like-minded citizens to meet one another and coordinate. Meetings may allow opposition citizens to coordinate against a developer, perhaps sharing the costs of a lawsuit or improving their chances of winning by working together; from the developers' perspective, meetings would be a double-edged sword, where the information gained creates greater efficiency, but at the cost of creating a stronger legal opponent.

Beyond the context of housing, our model speaks to a larger class of problems in which citizen action is asymmetric: opponents can veto a project (through litigation), but cannot mandate a larger project. Optimal institutional design that gives interested parties' some say in the outcome but the proposal can only be vetoed, not expanded or mandated, is a difficult but important venture.

References

Adams, Brian. 2004. "Public meetings and the democratic process." *Public Administration Review* 64 (1): 43–54.

Adelman, David E, and Robert L Glicksman. 2019. "The limits of citizen environmental litigation." *Natural Resources & Environment* 33 (4): 17–21.

Aragonès, Enriqueta, and Santiago Sánchez-Pagés. 2009. "A theory of participatory democracy based on the real case of Porto Alegre." *European Economic Review* 53 (1): 56–72.

Asquith, Brian J., Evan Mast, and Davin Reed. 2023. "Local Effects of Large New Apartment Buildings in Low-Income Areas." *The Review of Economics and Statistics* 105 (2): 359–375. https://doi.org/10.1162/rest_a_01055.

Bhan, Gautam. 2009. ""This is no longer the city I once knew". Evictions, the urban poor and the right to the city in millennial Delhi." *Environment and urbanization* 21 (1): 127–142.

Brouwer, NR, and Jessica Trounstine. 2024. "NIMBYs, YIMBYs, and the Politics of Land Use in American Cities." *Annual Review of Political Science* 27.

Cole, Richard L, and David A Caputo. 1984. "The public hearing as an effective citizen participation mechanism: A case study of the general revenue sharing program." *American Political Science Review* 78 (2): 404–416.

Collins, Jonathan E. 2021. "Does the meeting style matter? The effects of exposure to participatory and deliberative school board meetings." *American Political Science Review* 115 (3): 790–804.

Currans, Kristina M, and Kenneth A Stahl. 2024. "Are Traffic Studies "Junk Science" That Don't Belong in Court?" *Journal of the American Planning Association,* 1–9.

de Benedictis-Kessner, Justin, and Michael Hankinson. 2019. "Concentrated burdens: How self-interest and partisanship shape opinion on opioid treatment policy." *American Political Science Review* 113 (4): 1078–1084.

Einstein, Katherine Levine, David Glick, Luisa Godinez Puig, and Maxwell Palmer. 2023. "Still muted: The limited participatory democracy of Zoom public meetings." *Urban Affairs Review* 59 (4): 1279–1291.

Einstein, Katherine Levine, David M. Glick, and Maxwell Palmer. 2020. *Neighborhood Defenders: Participatory Politics and America's Housing Crisis.* New York: Cambridge University Press.

Einstein, Katherine Levine, Maxwell Palmer, and David M Glick. 2019. "Who participates in local government? Evidence from meeting minutes." *Perspectives on politics* 17 (1): 28–46.

Fearon, James D. 1995. "Rationalist explanations for war." *International organization* 49 (3): 379–414.

Fischel, William A. 2004. "An economic history of zoning and a cure for its exclusionary effects." *Urban Studies* 41 (2): 317–340.

Foster, David, and Joseph Warren. 2022. "The NIMBY problem." *Journal of Theoretical Politics* 34 (1): 145–172.

Glaeser, Edward L, Joseph Gyourko, and Raven E Saks. 2005. "Why have housing prices gone up?" *American Economic Review* 95 (2): 329–333.

Hankinson, Michael, and Justin de Benedictis-Kessner. 2022. "How Self-Interest and Symbolic Politics Shape the Effectiveness of Compensation for Nearby Housing Development."

Lahav, Alexandra D. 2016. "The Roles of Litigation in American Democracy." *Emory LJ* 65:1657.

Li, Xiaodi. 2021. "Do new housing units in your backyard raise your rents?" *Journal of Economic Geography* 22 (6): 1309–1352. https://doi.org/10.1093/jeg/lbab034.

Marble, William, and Clayton Nall. 2021. "Where self-interest trumps ideology: liberal homeowners and local opposition to housing development." *The Journal of Politics* 83 (4): 1747–1763.

Mast, Evan. 2023. "JUE Insight: The effect of new market-rate housing construction on the low-income housing market." *Journal of Urban Economics* 133:103383. https://doi.org/https://doi.org/10.1016/j.jue.2021.103383.

Ornstein, Joseph T. 2023. "Zone defense: Why liberal cities build too few homes." *Journal of Theoretical Politics* 35 (4): 310–323.

Osborne, Martin J, Jeffrey S Rosenthal, and Matthew A Turner. 2000. "Meetings with costly participation." *American Economic Review* 90 (4): 927–943.

Sotomayor, Luisa, Sergio Montero, and Natalia Ángel-Cabo. 2023. "Mobilizing legal expertise in and against cities: urban planning amidst increased legal action in Bogotá." *Urban Geography* 44 (3): 447–469.

Taylor, Whitney K. 2020. "On the social construction of legal grievances: Evidence from Colombia and South Africa." *Comparative Political Studies* 53 (8): 1326–1356.

Turner, Matthew, and Quinn Weninger. 2005. "Meetings with costly participation: an empirical analysis." *The Review of economic studies* 72 (1): 247–268.

Online Appendix for

“The Public Meeting Paradox: How NIMBY-Dominated Public Meetings Can Enable New Housing”

Allison K. Cuttner, Emory University
Ryan Hübert, UC Davis
B. Pablo Montagnes, Emory University

February 2024

For Online Publication

Contents

A Proofs and Additional Formal Results

Lemma A.1. Given Assumption 1, then $\tilde{s}(\overline{\alpha}) < \tilde{s}(\underline{\alpha})$.

Proof of Lemma A.1. Let $s_0 < \hat{s}_i(\alpha_i)$, and define $s_0'(\alpha_i) = 2\hat{s}_i(\alpha_i) - s_0$. Since v_i is strictly concave, $|\hat{s}_i(\overline{\alpha}) - s_0| < |\hat{s}_i(\overline{\alpha}) - s_0|$, and $\kappa_P/(\tilde{\gamma}_i^1(n) - \tilde{\gamma}_i^0(n))$ is constant in s (using Assumption 2), then it follows that $|\tilde{s}(\overline{\alpha}) - s_0'(\overline{\alpha})| < |\tilde{s}(\underline{\alpha}) - s_0'(\underline{\alpha})|$. Moreover, $s_0'(\overline{\alpha}) < s_0'(\underline{\alpha})$ since $\hat{s}_i(\overline{\alpha}) < \hat{s}_i(\underline{\alpha})$, and then:

$$\tilde{s}(\overline{\alpha}) = s_0'(\overline{\alpha}) + |\tilde{s}(\overline{\alpha}) - s_0'(\overline{\alpha})| < s_0'(\underline{\alpha}) + |\tilde{s}(\underline{\alpha}) - s_0'(\underline{\alpha})| = \tilde{s}(\underline{\alpha})$$

We have thus shown that $\tilde{s}(\overline{\alpha}) < \tilde{s}(\underline{\alpha})$. □

A.1 Low Types' Incentives in the Only-High Meeting Equilibrium

Now consider the strategic calculations of a representative neighbor i who is a low type. He gets the following payoff from not attending:

$$\underbrace{\rho^{n-1} v_i(s_M, \underline{\alpha})}_{\substack{\text{All other neighbors} \\ \text{are low types} \\ \text{(no attendance)}}} + \underbrace{(1-\rho^{n-1}) v_i(s_1, \underline{\alpha})}_{\substack{\text{At least one other neighbor} \\ \text{is a high type}}}$$

However, what happens if low types deviate and attend the meeting?

In this case, all n neighbors attend and $m \leq n$ are randomly chosen for the developer to observe. Since observing low types at the meeting is off the equilibrium path, perfect Bayesian equilibrium does not pin down the developer's belief in that situation. By Assumption 8, if she observes any high types, she (correctly) believes there is a high type at the meeting who is eligible to sue. And, if she observes all low types, she believes there is a high type in the meeting with probability $\tilde{\eta} \in [0, 1]$.

There are two scenarios in which the developer observes only low types. First, all neighbors could be low types, and without the deviation, there would be no meeting attendance. This happens with probability ρ^{n-1}. In this situation, the deviating low types get a payoff that depends on the developer's off path belief η.

Second, the neighborhood could have a mix of low and high types (probability $1 - \rho^{n-1}$), but with the deviation, only low types are randomly chosen to be observed. A deviating low type i believes this happens with probability:

$$\tilde{\lambda}_i(n,m) = \sum_{z=m-1}^{n-2} \binom{n-2}{z} \rho^z (1-\rho)^{n-2-z} \left(\frac{z+1}{n}\right)^m$$

Which is simply the probability that there are between m and n low types, and only low types were chosen.

For a low type, $l_i = 1$, and considering this deviation, the belief becomes: Recall that by Assumption 2, each neighbor i believes there will be a successful lawsuit with probability $\tilde{\gamma}_i^{\tau_i}(n)$. Moreover, since each neighbor's cost is private information, a neighbor i only learns the types of the m attendees whose types are revealed.[1] Let $s_2' \equiv \min\{s_2, s_M\}$ Then, the expected payoff to a low type neighbor i for deviating and attending the meeting is:

$$\begin{aligned}
&\rho^{n-1}(\tilde{\eta} v_i(s_1,\underline{\alpha}) + (1-\tilde{\eta}) v_i(s_2,\underline{\alpha})) \\
&\quad + \tilde{\lambda}_i(n,m)(\tilde{\eta} v_i(s_1,\underline{\alpha}) + (1-\tilde{\eta})((1-\tilde{\gamma}_i^0(n)) v_i(s_2',\underline{\alpha}) + \tilde{\gamma}_i^0(n) v_i(s_0,\underline{\alpha}))) \\
&\quad + (1-\tilde{\lambda}_i(n,m) - \rho^{n-1}) v_i(s_1,\underline{\alpha}) \\
&\quad - \kappa_M
\end{aligned}$$

Note that when $\tilde{\eta} = 1$, this collapses to

$$\underbrace{\rho^{n-1} v_i(s_1,\underline{\alpha}) + \tilde{\lambda}_i(n,m) v_i(s_1,\underline{\alpha}) + (1-\tilde{\lambda}_i(n,m) - \rho^{n-1}) v_i(s_1,\underline{\alpha})}_{v_i(s_1,\underline{\alpha})} - \kappa_M$$

Since $s_M > s_1$, and if s_M is just set to induce the high type to attend, she's clearly better off not attending.

When $\tilde{\eta} = 0$, this collapses to

$$\rho^{n-1} v_i(s_2,\underline{\alpha}) + \tilde{\lambda}_i(n,m)((1-\tilde{\gamma}_i^0(n)) v_i(s_2',\underline{\alpha}) + \tilde{\gamma}_i^0(n) v_i(s_0,\underline{\alpha})) + (1-\tilde{\lambda}_i(n,m) - \rho^{n-1}) v_i(s_1,\underline{\alpha}) - \kappa_M$$

If $s_M = s_2$, then the first term is equal to the first term in his utility from not attending, and the remainder (except the cost of attendance) is a lottery over s_2, s_1 and s_0. By Assumption 7, he clearly prefers s_1 to this lottery. He'd be worse off attending than not attending. Still, however, he pays a cost κ_M to attend, so the developer can set some $s_M > s_2$ and induce the high types to attend and the low types not to attend. Then, the s_M that induces non-attendance for the low types depends on $\tilde{\eta}$. More specifically, there exists a threshold $s_M^{\text{high}}(\tilde{\eta})$ that is decreasing in $\tilde{\eta}$ from some maximum level $s_M^{\text{high}}(0) > s_2$ to the level that just induces the high types to attend, $s_M^{\text{high}}(1) > s_1$.

1. It is possible that a given neighbor i ends up knowing the types of $m+1$ neighbors if she is not among the m whose types are publicly revealed. This won't dramatically change our analysis, and we thus ignore it to make our results more parsimonious.