

C A G E

Working Paper

820/2026
August 2026

The Hidden Geography of Housing Demand Exposes Policy Failure: Evidence from Billions of Housing Searches

Nikhil Datta,
Amrita Kulka,
Jawad Wehbe

ISSN: 2978-0276
Grant number: ES/7504701/1

**UNIVERSITY
OF WARWICK**

Economics



Economic
and Social
Research Council

The Hidden Geography of Housing Demand Exposes Policy Failure: Evidence from Billions of Housing Searches

Nikhil Datta^{†‡}

Amrita Kulka[†]

Jawad Wehbe[†]

August 2026

Abstract

House prices are widely used across the social and environmental sciences to inform crucial business and policy making decisions meant to reduce spatial economic inequalities and future-proof locations at risk from environmental shocks. Yet prices are biased by supply and only observed for the selected properties that transact. This paper is the first to introduce a direct measure of latent location demand, using billions of housing searches in Great Britain between 2019 and 2024. Across three causal applications spanning public health, environmental science and economics, we show that searches reveal demand that transaction data either obscure or cannot measure. The COVID-19 pandemic triggered a “race for space”, but demand for private greenspace returned to pre-pandemic levels before new supply could respond, leaving planners and developers chasing yesterday’s preferences. Local flooding generates temporary dips in demand, suggesting people are myopic. New housing supply developments do not induce increases in location demand, implying new town development policies are misguided. In all three cases, transaction data miss the fundamental behavioural response, leading to suboptimal policy design. Finally, we use the search data to construct a measure of excess housing demand for 235,243 micro neighbourhoods and release it through the WhereToBuild mapping tool.

²Department of Economics and CAGE, University of Warwick

³CEP, London School of Economics

1 Introduction

House prices are a ubiquitous empirical measure across the social and environmental sciences. Environmental and climate researchers use them to price flood risk and environmental quality (Hino and Burke, 2021; Lang et al., 2023; Mamun et al., 2023); geographers and urban scholars use them to quantify neighbourhood change and the value of local accessibility (Aguilar-Velázquez et al., 2024; Gibbons and Machin, 2005); research in finance uses them to value mortgage and corporate collateral (Mian and Sufi, 2011; Cvijanović, 2014); research in wellbeing uses housing costs to construct spatial measures of quality of life (Oswald and Wu, 2010); education researchers use them to infer willingness to pay for school quality (Hu and He, 2024); public-health researchers use them to infer willingness to pay to avoid exposure to infectious disease and other health risks (Liu and Tang, 2021; Qian et al., 2021); and economists use them to value public goods, amenities and disamenities, from clean air and open space to crime and transport access (Chay and Greenstone, 2005; Linden and Rockoff, 2008). Across these applications, the central empirical premise is the same: transacted house prices are interpreted as revealing demand.

However, transacted house prices are an imperfect measure of demand. They are observed only for the small and changing subset of homes that sell, making them sparse at fine spatial scales and difficult to compare over time. Prices also reflect not only where people want to live, but how many homes are available, which properties owners choose to sell and the prices sellers are willing to accept. Moreover, a transaction is typically recorded four to six months after a buyer first begins searching. These limitations can distort policy-relevant estimates and, in places with few or no sales, leave researchers with no usable measure of demand at all.

Our first contribution is to formalise these concerns in a model showing how seller reference dependence, selection into transactions and housing supply responses can cause changes in transacted prices to diverge from the underlying change in demand. In doing so, we provide researchers using transaction prices with a necessary empirical test: estimating the response of transaction volumes in addition to prices using the same research design. A change in volumes indicates that the set of properties, buyers or sellers observed in transactions has changed, so the estimated price response may combine the demand effect with selection or supply adjustment. An unchanged volume response does not prove that prices isolate demand, but it is a necessary condition for that interpretation.

Our second contribution is to introduce a direct measure of demand for locations using data from a major housing-search platform in the UK. Housing searches overcome several of the limitations of transaction prices by providing a high-volume, high-frequency and geographically granular record of expressed location preferences. Crucially, users can search in locations with no available listings, allowing demand to be observed even when no housing is currently available. Our dataset covers 16 billion searches between 2019 and 2024 and includes persistent individual identifiers, measures of engagement such as property views, saves and contacts, and detailed spatial filters including cities, neighbourhoods, train stations, postal-districts and user-defined radii. To ensure that the searches represent actual location demand, we carry out a range of robustness checks using measures of more serious search activity, such as saving a property or contacting the sales or letting broker via the platform. The resulting spatial patterns of demand are robust to these alternative restrictions.

Our main contribution is to demonstrate the power of this measure using three exemplary highly policy relevant settings. We use causal event-study analyses of three distinct shocks, each chosen to highlight a setting in which transaction prices are limited, and could lead to sub-optimal policy conclusions. We study a widespread and persistent public-health shock through the COVID-19 pandemic; short-lived and highly localised environmental shocks through flooding events; and areas with few or no prior transactions which experience large housing-supply shocks. Together, these applications show how housing searches reveal fine-grained changes in location demand that standard transaction data either obscure or cannot measure.

We begin with the setting in which transaction prices should be best able to capture changing location preferences: a relatively persistent shock which should be priced in even though transactions occur with a delay. The COVID-19 pandemic altered working arrangements, residential patterns and urban spatial activity, at least temporarily (Li et al., 2025; Monte et al., 2025; Xu et al., 2023; Gupta et al., 2022). We use housing searches to estimate how preferences for private greenspace changed when pandemic

restrictions sharply reduced people’s access to shared outdoor space.

Our event-study estimates show that, at the peak of the pandemic, an additional $100m^2$ of private greenspace increased searches by approximately 0.8%. The effect emerged gradually and returned to zero by the end of 2023. In contrast, we show that transaction-price effects persist beyond the change observed in searches and are confounded by differential selection into transactions and the subsequent expansion of housing supply in greener areas. Using prices to measure demand would therefore overstate the persistence of preferences for private greenspace, understate the recovery of demand for urban living, and potentially misdirect new housing supply. More generally, when transaction prices adjust through changes in both the number and composition of sales, they can misrepresent the magnitude and duration of changes in underlying preferences and induce costly, spatially misallocated investment.

Next, we turn to flooding events, which generate temporary, and geographically highly localised shocks. A large multidisciplinary literature uses house prices to infer how transaction prices respond to flood and climate risk, but finds mixed effects, ranging from short-run declines (Beltrán et al., 2019) to more persistent responses (Skouralis et al., 2024; Hino and Burke, 2021), including after infrequent floods (Thompson et al., 2026), no pricing into second homes or rental properties Bézy (2025), and even rent increases driven by selective migration (Fan et al., 2025). We study 20 flooding events using housing searches, allowing us to measure the immediate demand response directly and shed light on these conflicting findings.

Our results show a sizeable but temporary decline in housing searches of up to 5% in affected locations during the ten weeks following a flood, relative to unaffected areas within the same TTWA. By contrast, transaction prices show no statistically detectable response. The price estimates are highly imprecise because few properties transact within these narrowly defined areas and time periods, with confidence intervals wide enough to encompass the demand response observed in searches. Transaction volumes also show a suggestive decline around two to three months after flooding, but the estimates are noisy and cannot be distinguished from zero. For policies designed to reduce flood risk, the peak 5% decline in searches rather than the longer-term null effect provides the relevant measure of the behavioural response when households are actively attending to that risk; relying on the absence of a detectable price effect would systematically understate the importance households place on flood protection and resilience. This helps reconcile mixed findings in the price literature and demonstrates the potential for searches to provide real-time evidence on how people respond to climate risk, informing policymakers, planners and insurers responsible for climate adaptation, local investment and risk pricing.

Finally, we apply our new demand measure to a setting relevant to major policy debates in countries facing housing affordability pressures and local economic decline: whether large new developments, including “new towns”, can generate demand for places that currently have little housing or economic activity (Huang et al., 2025; MHCLG, 2026). Transaction data cannot answer this question where no housing, and therefore no transactions, exists before development. Using housing searches, our event study shows that large new-build developments increase searches by only 1.8% in areas with relatively little pre-existing search activity. This corresponds to roughly one to two additional searchers per month for every 100 new housing units developed. This novel result, which cannot be recovered from transaction data, suggests that new housing supply alone generates little additional demand for a location and may therefore be insufficient to deliver the wider objectives associated with new-town development.

Beyond these causal applications, we demonstrate the scalability of housing-search data by combining searches with property listings to construct the “housing gap”, a measure of excess demand for 235,243 micro neighbourhoods across Great Britain. Search data are less readily accessible than transaction records, so we make these estimates freely available through our [WhereToBuild mapping tool](#). The tool reveals substantial variation in excess demand within Local Authorities, the level at which housing targets are administered. This matters because meeting an area-wide target still requires deciding where new housing should be built, and poor spatial allocation may place new supply in neighbourhoods with relatively little demand. The tool helps avoid this by identifying where excess demand is greatest at the hyper-local level.

The remainder of the paper proceeds as follows. Section 2 first examines why transaction prices may fail to recover underlying demand, using a simple model to clarify the roles of supply, seller behaviour and

selection into transactions, before introducing our housing-search measure. Section 3 presents the three causal applications. Section 4 introduces the housing-gap measure and the WhereToBuild mapping tool. Section 5 concludes, and Section 6 details the data construction and empirical methods.

2 A new measure of spatial demand

We introduce a new measure of spatial housing demand based on fine-grained housing search data. To motivate the measure, we first discuss the predominant empirical proxy for demand for places and their characteristics: real estate transaction prices.

2.1 Advantages and disadvantages of using transaction prices

Transaction prices are the dominant empirical measure for inferring preferences over locations and their characteristics. We first discuss why they are useful, before turning to three limitations that motivate our complementary measure based on housing search behaviour.

Transaction prices are a ubiquitous measure of demand for locations and location-specific characteristics. Their popularity reflects both the widespread availability of property transaction data across many countries and their status as a measure of revealed preferences: unlike surveys asking people where they would like to live, transactions capture actual housing choices.

At the same time, transaction prices have three main shortcomings as a measure of demand for locations. First, prices are an outcome determined by both demand and supply in an equilibrium. Second, they are observed only when transactions occur. Third, behavioural frictions on the seller side can bias prices away from underlying demand. We now discuss each of these factors in turn.

Transaction prices are the result of the interaction of demand and supply factors. Prices therefore capture how elastic housing supply is (i.e. how responsive new housing supply is to prices), in addition to peoples' preferences for the location and housing unit. As an example, consider the same positive amenity shock, such as cleaner air, occurring in two locations with homogeneous preferences but differently elastic housing supply (Baum-Snow and Han, 2024a), because of differences in construction costs or geographic constraints. In the location with more elastic housing supply, prices will adjust less than in the location with inelastic supply, as the demand shock will be followed by developers building new supply to meet that demand quickly. An analysis based only on price changes could therefore incorrectly suggest that the amenity shock, or preferences for that amenity, differed across the two locations.

Prices also reflect other factors beyond demand for location-specific amenities. At a fundamental level, spatial equilibrium models from the economics literature imply that housing prices capitalise both amenities and local wage levels (Roback, 1982; Rosen, 1979). Positive income shocks can therefore raise housing costs even in lower-cost segments; for example, minimum wage increases have been shown to increase rents (Yamagishi, 2021). Prices also embed development costs (Gyourko and Saiz, 2006), which vary across places because of differences in land prices, labour costs, regulation, and geography (Glaeser et al., 2005; Saiz, 2008).

As a result, places with fewer amenities and lower prices may still face high excess demand relative to their housing stock. Measuring demand using transaction prices therefore requires quasi-experimental settings in which supply adjustments can be ruled out, or approaches that explicitly model both the supply and demand side of the market.

Transaction prices can measure housing demand only when transactions occur. This creates difficulties in areas with few observed sales, whether because of low turnover or because housing has not yet been built, as in potential new towns, rural developments, or urban-fringe sites. In these settings, local price measures are often sparse and noisy, and in some cases demand cannot be estimated from transaction prices at all. The timing of transactions also matters. Sales may occur only after the event of interest, and for short-lived shocks, transaction prices may miss demand responses altogether.

Finally, even when supply conditions are well understood and transaction volumes are high, prices may still fail to capture demand if sellers are reference dependent. A reference-dependent seller evaluates

an offer not only relative to the quality or current market value of the property, but also relative to a reference point, such as the price they paid, a recent valuation, or a price they expected to receive (Andersen et al., 2022). When current market conditions fall below this reference point, sellers may require a higher offer to be willing to sell. Transaction prices therefore reflect not only buyer demand, but also the seller’s reference point.

This creates a problem for studies that use transaction prices to measure changes in housing demand. Consider a simple comparison of price changes in affected and unaffected locations before and after a local demand shock. If demand falls in the affected location, sellers whose expectations remain anchored to earlier, higher prices may reject offers or decide not to sell. The properties that do transact after the shock are therefore not necessarily comparable to those that transacted before it. They may disproportionately involve sellers with lower reference prices or a greater willingness to move, or buyers who place an unusually high value on a particular property. The observed change in transaction prices will then reflect both the direct effect of the demand shock and a change in the composition of the transactions being observed. Appendix B formalises this argument and shows that, when sellers are reference-dependent, changes in transaction prices following a location-specific shock need not provide a clean measure of the change in housing demand.

A useful implication is that studies using transaction prices to infer changes in local demand can examine transaction volumes (which follow directly from transaction prices and therefore require no additional data) using the same empirical design as a minimum diagnostic. In the absence of seller-side frictions or changes in supply, a location-specific demand shock should affect transaction prices without generating a differential change in the probability that a transaction occurs. Re-estimating the specification with transaction volumes as the dependent variable therefore provides a test of this restriction. A significant response in transaction volumes indicates that the set of observed buyers, sellers, and properties has changed, so the corresponding price estimate may combine the underlying demand response with selection into transactions or supply adjustment. A null volume response is not sufficient to establish that transaction prices isolate demand, since offsetting compositional changes may leave aggregate volumes unchanged, but it is a minimum condition for that interpretation. Appendix B.1 shows this result formally.

2.2 Search data

We next develop a method to measure spatial preferences that does not rely on transaction prices. Instead, our measure is based on housing search behaviour, allowing us to capture demand before transactions, or behavioural changes such as moves, occur. This methodology circumvents the limitations inherent to price signals such as confounding effects through supply constraints, low transaction volumes, or behavioural biases like seller reference dependence. We first describe the housing search data used to construct the measure.

To operationalise our measure, we use large-scale data on online housing search activity from Rightmove, Great Britain’s most used housing search platform, covering approximately 10 billion rental and sales searches, linked to the universe of property listings across Great Britain. The data cover user activity from 2019 to 2024 and include persistent, anonymised 32-digit hash identifiers, which allow users to be followed across search sessions, creating a digital trace. Most of our analysis—with the exception of our excess demand results in Section 4—uses searches in the sales market rather than the rental market. This is deliberately done to retain comparability with standard transactions data. Section 4 estimates overall excess demand by micro-neighbourhood and therefore includes information from the rental market to avoid undercounting.

For each search, we observe the complete set of interactions with listed properties, including views, saves, and contacts, as well as the user’s spatial search criteria. Search locations include place names, transport stations, postal areas, postal districts, and full postcodes, and may also include a user-defined radius. These spatial inputs are linked to detailed polygons, allowing each search to be spatially assigned to Output Areas capturing hyper-local, high-frequency human geographic demand. Section 6.2 discusses in detail how we allocate searches to output areas.

In our application to the UK housing market (see Section 4), we also use supply side data, namely the

complete universe of rental and sales listings over the same period. Each listing includes precise geographic coordinates, allowing us to identify the exact set of spatial searches in which it would appear. The data also provide detailed information on prices and other property-level attributes.

We impose several baseline restrictions to ensure that the search data reflect genuine user behaviour. Users must have more than one session, at least one session lasting longer than 60 seconds, at least one recorded search, and fewer than 10,000 logged events. These filters effectively remove automated activity and ensure that our measures capture meaningful preferences for residential locations.

Descriptive information on the data for the 2023-2024 period is presented in Table A1 in the appendix. In the baseline sample, there are 57 million users carrying out 1.8 billion searches. Each user has on average 7 separate sessions, where each session lasts an average of 11 minutes. In total there are 2.2 million listings over the year. In robustness exercises, we restrict the sample to those users that either email a rental or sales broker through the platform, or save a property, to focus on “serious” searchers. The table also presents descriptive statistics for these subsets.

2.3 Using search data to measure demand

As discussed above, measuring residential demand at fine spatial scales requires data that are both geographically granular and not jointly determined by market supply. In principle, researchers could survey individuals directly about where they would like to live, as in Ferreira and Wong (2022). However, obtaining samples large enough to measure demand across many finely defined locations is prohibitively costly.

Search data have several advantages for measuring housing demand. First, they allow demand to be measured at the level of the searched location. These locations can be as fine as individual postcodes but also capture large towns extended by user-specified radii. Second, users can search for a location even when there are no available listings there, or none that satisfy their other filters, such as price or number of bedrooms. We therefore observe demand even when realised choices would necessarily be constrained by housing availability. Third, housing searches are a direct expression of housing demand unlike many demand measures used in the housing and location-choice literature that are either truncated by housing availability; indirect, because they use employment shocks or population growth as proxies for demand; or confounded by additional factors (Baum-Snow and Han, 2024b; Bayer et al., 2007).

Despite these advantages, there are also potential problems with using search data. A central issue in interpreting search data is how searchers incorporate budget constraints. At one extreme, users may pay little attention to affordability when searching. These aspirational searches, for example for high-end properties in highly desirable areas, could overstate demand for locations where users have little realistic prospect of buying or renting. At the other extreme, users may search rigidly within their current budget constraints, excluding locations that are even marginally unaffordable. In that case, search demand should be interpreted as demand at prevailing prices: the market may be able to absorb additional units at current prices, but the measure may not capture how households would substitute across locations in response to small price changes. Search data therefore provide a measure of spatial demand whose interpretation depends on how strongly search behaviour is shaped by affordability constraints.

We are able to rule out the former extreme. Approximately 60% of searches have a maximum price filter when searching, and when restricting to “serious searchers”, e.g. those searchers who contact a sales or letting broker during their search process on the platform, the spatial ranking by OA in excess demand remains broadly unchanged.¹ We are additionally able to plausibly rule out the latter extreme. Users update their maximum price filters during their search process. In the average session a user carries out approximately 5 searches, and they update their price filter with a 25% probability at every search step. Furthermore, the average price of a property that a user views is approximately 12% less than the value of their maximum price filter, conditional on having one. We take this as evidence that the search data predominantly capture preferences conditional on soft budget constraints, but where individuals retain

¹Appendix Table A2 reports pairwise rank–rank correlations of searches across alternative restrictions based on saving properties and contacting agents. The correlations range from 0.95 to 0.98.

the scope to trade off prices, quality of location and property characteristics against one another.

Figure 1 illustrates why housing searches are particularly valuable for measuring demand at fine spatial scales. The left panel maps buying searches (scaled by output area size) one week before Storm Ciara, while the right panel maps the percentage change in searches between one week before and four weeks after the flood. In both panels, black dots denote transactions in the corresponding week and purple outlines identify flooded areas. Most Output Areas record no transaction even before the flood, and transactions become still rarer four weeks afterwards; among the flooded areas, only two post-flood transactions are observed. By contrast, searches are observed across almost all Output Areas and exceed transactions by orders of magnitude, even within these very small neighbourhoods. Search data can therefore reveal geographically detailed changes in demand where transaction data provide too few observations to measure either baseline conditions or post-shock responses reliably. We discuss the corresponding causal change in searches after flooding events in Section 3.2.




Figure 1: Search activity before and after flooding

Note: Panel (a) illustrates a map of the baseline searches per km² in Calderdale in week -1 relative to Storm Ciara, February 2020. Panel (b) shows the % log difference in searches per km² between weeks 4 and -1 relative to the flood; the % changes are winsorised at the 99th percentile for visibility. Black dots denote property transactions in the LA-week. Purple boundaries mark flooded output areas. For details regarding the spatial allocation of floods to OAs, refer to Section 6.4.

3 Search data reveal demand responses where prices are limited

We use housing searches to estimate causal demand responses to three spatially varying shocks affecting public health, the environment and housing supply, each chosen to highlight a distinct advantage of search data over transaction prices.

First, we study COVID-19 as a widespread and relatively persistent demand shock. In such settings, price responses may be difficult to interpret because shocks can induce selection into transactions and trigger supply responses, both of which complicate the use of prices as a proxy for demand. Second, we examine flooding events as short-lived, highly localised shocks. Transaction-based estimates may fail to detect these demand responses because very few sales occur within the affected area and time period. Estimates based on such small samples are also likely to be noisy, reflecting variation in the characteristics of the properties sold and in the willingness to pay and sell of the particular buyers and sellers who transact. Third, we study a highly localised housing supply shock in areas with little or no pre-existing housing stock. In these settings, transaction prices are censored in the pre-period, making it difficult or impossible to estimate causal effects using sales data. Together, these examples demonstrate how search

data can measure demand in settings where price-based approaches are limited.

We estimate responses to each shock using event-study designs that compare changes in affected locations to changes in appropriate comparison locations around the timing of the shock. In the generalised specification below, we let Y_{irt} denote an outcome in output area i , belonging to region r , and time period t . Depending on the application, Y_{irt} denotes housing searches, transaction prices, transaction volumes or new supply transactions and t can represent year-weeks or year-months. We estimate specifications of the form

$$Y_{irt} = \sum_{\substack{k=\bar{k} \\ k \neq -1}}^{\bar{k}} \beta_k (\text{Treat}_i \times \mathbf{1}\{t - T_i = k\}) + \alpha_i + \gamma_{rt} + \varepsilon_{irt}. \quad (1)$$

Here, Treat_i indicates whether output area i is exposed to the shock, T_i is the timing of exposure, and k indexes time relative to the event. The coefficients β_k trace the evolution of the outcome in treated locations relative to comparison locations, normalised to the period immediately before the shock. The fixed effects α_i absorb time-invariant differences across locations, while γ_{rt} absorb common shocks across locations within the same region, in each time period.

This baseline framework is adapted to each setting. For COVID-19, treatment intensity is continuous: locations differ in their pre-existing amount of private greenspace, and the event time is the onset of the pandemic. For flooding, treatment is defined by whether residential properties in an Output Area are exposed to a flood event. For new-build sites, treatment is the start of construction in locations receiving large new housing developments. The specific outcome definitions, fixed effects, time and region aggregation, and sample construction differ across applications and are described in more detail in section 6.


3.1 Greenspace demand during the COVID-19 Pandemic

We begin by analysing a wide-spread shock, the COVID-19 Pandemic, which substantially changed preferences for space as people temporarily fled dense cities (Gupta et al., 2022; Althoff et al., 2022; Monte et al., 2025; Ramani et al., 2024) and were restricted in their movements through lockdowns leading to a higher valuation of private (green)space (Li et al., 2025; Malik et al., 2023). We compare the impacts of the pandemic outbreak on output-area level searches, transaction price, transaction volumes, and new build transaction volumes in hyper-local neighbourhoods with more versus less average private greenspace. This allows us to measure demand, equilibrium prices, selection changes and new supply of private greenspace respectively.


The resulting event-study estimates are plotted in Figure 2, where the dotted line denotes the first week of March 2020, when lockdown orders came into effect. Panel 2a presents the results for housing searches and shows a clear pattern: searches increased throughout the second quarter of 2020 in OAs with more private greenspace, with the effect peaking from the fourth quarter of 2020 onwards. At the peak, an additional $100m^2$ of median private greenspace increased searches by approximately 0.8%. Given that the average private garden in Great Britain is $188m^2$, the estimates imply that the pandemic increased searches by approximately 1.5% more in an area with an average-sized garden than in an otherwise comparable area with no private greenspace.

The increase persisted for almost two years. Searches began to decline during the summer of 2022 and had returned to their pre-pandemic level by the end of 2023. Overall, the results suggest that the pandemic caused a medium-run increase in people’s preferences for private greenspace that lasted for several years, but ultimately returned to baseline as lockdowns and other pandemic restrictions ended. This is consistent with other findings in the literature documenting a post-pandemic return to cities (Rowe et al., 2023; Monte et al., 2025; Ellen et al., 2026) and helps validate the use of housing-search data to capture changes in residential preferences.


Panel 2b initially displays a broadly similar pattern and magnitude to the search estimates when using transaction prices. Transaction prices increase, though with a lag of approximately six months relative to searches, and reach their peak in 2023. At the peak, an additional $100m^2$ of median private greenspace




(a) Weekly Searches



(b) Monthly Prices



(c) Monthly Number of Transactions



(d) New builds Transactions

Figure 2: Greenspace during Covid-19 Event Study

Note: The above figures present the estimation results of equation 5. Panels (a), (b), (c), and (d) respectively show the % changes in searches, prices, all transactions, and new builds transactions relative to the start of the pandemic. Effects are normalised to relative period -1 for all panels except panel (d) which was normalised to relative period -6 . For details regarding the construction of dataset and the empirical implementation, refer to Section 6.3.

is associated with an approximately 1% larger increase in transaction prices, close to the corresponding increase in searches. The earlier response of searches highlights their ability to capture changes in demand immediately, before they are reflected in completed transactions. However, transaction prices remain above their pre-COVID level at the end of 2024, around one year after searches have returned to baseline. These results are subject to the caveat that direct comparisons between the search and transaction-price estimates should be made with some caution because the underlying samples differ. Searches are observed at high frequency for essentially every OA-week, whereas transaction prices are observed only when a sale occurs and remain highly unbalanced even after aggregation to the OA-month level (only 31.5% of OA-months remain). Differences between the two series may therefore reflect not only differences between demand and equilibrium prices, but also changes in the set of locations and properties for which prices are observed. We show in panel (a) of Figure A.2 that the results are robust to restricting the sample of searchers to serious searchers that either saved a property or made a contact through the platform.

Panels 2c and 2d provide evidence that the price estimates do not capture changes in demand and preferences alone. Interpreting changes in transaction prices as changes in demand relies on housing supply being effectively fixed over the relevant period. Under this assumption, the demand shock should be absorbed primarily through prices, with no systematic differential change in the quantity of housing transacted between areas with more and less private greenspace. However, transaction volumes increase for approximately 16 months after the onset of the pandemic, before falling below their pre-pandemic level and remaining persistently lower for the rest of the sample period. The timing of this subsequent decline closely corresponds to the Stamp Duty Land Tax holiday, which temporarily reduced the tax payable on residential purchases before being phased out in 2021. Importantly, there is no corresponding decline in housing searches when the policy ends. Searches therefore continue to trace underlying demand for private greenspace, while transaction-based outcomes also reflect policy-induced changes in market participation.

Changes in transaction volumes matter for the interpretation of transaction prices because observed prices reflect the willingness to pay of the buyers and the willingness to sell of the sellers who actually transact. The pronounced rise and subsequent fall in transactions therefore indicate that the composition of observed buyers, sellers, and properties changes over the event window. The close correspondence with the Stamp Duty Land Tax holiday further shows that transaction activity was being shaped by a major contemporaneous policy, and that this response differed between areas with more and less private greenspace.

Consistent with the model in Appendix B, the estimated price response may therefore combine changes in underlying demand with differential selection into transactions, rather than identifying the average change in willingness to pay for private greenspace.

Panel 2d points to an additional mechanism. Transactions involving newly built properties increase after the pandemic, indicating a supply response in areas with more private greenspace and a further change in the composition of properties being sold, both of which can affect transacted prices. The former is likely to reduce prices, while the change in composition can go in either direction, depending on the quality of the new houses relative to the rest of the local market. Taken together, these results demonstrate the advantage of searches as a measure of spatial demand. The persistent increase in transaction prices cannot be interpreted straightforwardly as evidence that preferences for private greenspace remained persistently elevated. Instead, prices combine changes in demand with selection into transactions, policy-induced variation in market activity, and adjustments in housing supply.

3.2 Hyper-local demand responses after flooding events

Our second application studies the impact of severe flooding events on location preferences in affected areas. Flooding shocks differ markedly from the COVID-19 pandemic because they are much more spatially concentrated and temporary. For example, Storm Ciara affected the UK over 9–10 February 2020, with more than a month’s rainfall falling in parts of West Yorkshire in around 18 hours and several hundred properties affected by flooding (Met Office, 2020). These two features may make such localised environmental shocks particularly difficult to detect using standard transaction data due to sample size issues. At the same time, understanding how people respond to natural disasters is vital for the design

of climate-adaptation and mitigation policies (Hino and Burke, 2021).

Panel 3a plots the results for searches based on the 20 largest flooding events in England and Wales between 2019-2024, and shows a statistically significant decline of up to 5% during the ten weeks following a flood, after which searches return to their pre-flood level. In contrast to the pandemic, flooding is a temporary shock, but it nevertheless causes a marked short-run change in people’s preferences for affected locations. Skouralis et al. (2024) argue that, because people are myopic about flood risk, these changes in preferences dissipate relatively quickly after a flood. However, Fan et al. (2025) show that even temporary flooding that leaves the overall population unchanged can have persistent consequences by altering the composition of residents, and therefore the risk profile of the local population and the responses required from local government. Short-run changes in housing demand may therefore provide an early indication of the migration and compositional adjustments through which flooding produces longer-run local effects. As before, these results are robust to restricting the sample to serious searchers only (see panel (b) of Figure A.2).


The effects of flooding on prices and transaction volumes are shown in Figures 3b and 3c, respectively. In contrast to searches, neither outcome exhibits a statistically significant response to the flooding events. A central reason is that transactions are comparatively rare at this spatial scale (typically only a handful of output areas are affected in each flooding events). Even after aggregating the data to the monthly level, the estimates are highly imprecise and the confidence intervals widen substantially. Table A3 shows why this is the case: Only 26% of month-year-output area observations in our estimating sample have transactions both before and after the flooding event. At the week-year-output area level, this fraction shrinks to 8%, demonstrating that the analysis cannot be credibly run at this level. Transaction data therefore cannot reveal the high-frequency adjustment that is clearly visible in searches at the weekly level, and which could in principle be measured at an even higher frequency using search data.

The transaction estimates are nevertheless suggestive of a decline around two to three months after a flood, but the effect is too noisy to distinguish from zero. This delayed and imprecisely estimated movement also reinforces the limitations of price-based inference discussed above. If flooding changes which properties, buyers and sellers enter into transactions, any subsequent price response will combine the underlying change in preferences with changes in the composition of observed sales. For example, transactions may become disproportionately concentrated among buyers who are less averse to flood risk and who may also differ systematically in their willingness to pay, access to finance or other characteristics. Together, these results show that short-lived and highly localised shocks can generate clear shifts in preferences that are visible in search behaviour but difficult to detect using either transaction volumes or prices.


3.3 Measuring the demand response after large supply shocks

Third, we consider settings in which demand cannot be measured using transaction data because few or no properties existed in the area before the supply shock. This is particularly relevant for large new-build developments and entirely new settlements, where the absence of pre-treatment transactions makes conventional before-and-after comparisons of prices or sales impossible (Table A5 shows that only 6.5% of week-year-output area observations in our estimation sample experience housing transactions before and after the supply shock). As governments seek to expand housing supply in response to affordability pressures, one increasingly prominent policy is the development of “New Towns”: large new settlements intended to attract residents and stimulate local economic activity and regeneration by providing housing (MHCLG, 2026). Housing-search data allow us to measure how demand for these locations changes as new housing becomes available, providing direct evidence on whether large supply interventions generate interest in places where no housing market previously existed. This offers a way to assess, from the demand side, whether such developments are likely to attract residents and achieve their wider spatial objectives.


Figure 4 presents the resulting event-study estimates for 3,200 new-build developments comprising at least 20 units and completed between 2019 and 2024. Searches increase modestly following construction, with the effect gradually rising and stabilising at approximately 1.8% after 44 weeks. In levels, this corresponds to only around one to two additional searches per Output Area per month relative to pre-treatment search activity, despite the average development adding more than 100 new housing units.



(a) Weekly Searches



(b) Monthly Prices



(c) Monthly Number of Transactions


Figure 3: Floods Event Study

Note: The figures above present the estimation results of equation 6. Panels (a), (b), and (c) respectively show % change in searches, prices, and transactions relative to the onset of the floods. Effects are normalised to relative period -1 . For details regarding the construction of dataset and the empirical implementation, refer to Section 6.4.

Panel (c) in Figure A.2 in the appendix shows that these results are robust to using only serious searchers.

This response is economically negligible: even very large local supply shocks generate only minimal changes in search activity, despite potentially substantial effects on local prices and rents. Housing searches therefore appear to capture relatively stable preferences over locations rather than responding mechanically to the availability of new housing. To our knowledge, this is the first evidence on how location demand responds to large new-build developments in lower-density areas. It begins to fill an important gap in our understanding of whether expanding housing supply in such locations is sufficient to attract additional demand.

Figure 4: Event study of searches around new housing construction



Note: The figures above shows % changes in searches from 26 weeks before to 44 weeks after the start of supply shock. Effects are normalised to week -5 . For details regarding the construction of dataset and the empirical implementation, refer to Section 6.5.

4 The WhereToBuild Mapping Tool

Finally, we demonstrate the scalability of housing-search data as a measure of location preferences by constructing Output Area-level estimates of excess housing demand across Great Britain. Addressing housing affordability has become an increasingly prominent policy objective worldwide, with many governments seeking to expand housing supply substantially. In England, for example, the government has committed to delivering over 1.5 million new homes over the current Parliament (UK Government, 2024). Yet policymakers lack direct, geographically granular measures of demand that can help determine where new housing should be built. We address this gap by developing a first-of-its-kind measure of hyper-local excess housing demand, which we call the “housing gap”, using data on housing searches and listings. We make these estimates publicly available through our *WhereToBuild* mapping tool.

Our measure of excess demand, the housing gap, is constructed as the difference between demand, defined as the number of searchers in an area, and the supply of available listings between May 2023 and April 2024. Larger values indicate a greater imbalance between the number of people searching and the number of properties available, and therefore a larger supply of housing is required to meet demand.

We calculate this measure for 235,243 Output Areas (OAs) for Great Britain. OAs are the smallest census geography in Great Britain and cover 40-250 households. In urban areas, they cover a small neighbourhood, while in rural areas they can be considerably larger. Accordingly, we present our main results in the form of gap per km^2 , rescaling the measure by the OA land area to elicit a like-for-like comparison.

Our methodological approach allocates users and listings across Output Areas using a common weighting procedure. Each user is assigned a total weight of one, distributed across all OAs in which they search, ensuring that each user is counted only once overall. Similarly, each listing is assigned a total weight of one, distributed across all OAs in which it appears. We define the housing gap as the difference between

the weighted number of users searching in an OA and the weighted number of listings available there. Appendix C provides full details of the measure’s construction, including the procedure used to avoid double counting households that search using multiple devices.

4.1 Where do people want to live?

Figure 5 maps the national rank of each OA’s housing gap per square kilometre across Great Britain, with darker colours indicating larger gaps. The largest housing gaps are concentrated in major cities, although there is substantial variation across neighbourhoods within them. In London, for example, parts of boroughs such as Wandsworth rank among the areas with the largest housing gaps in Great Britain. Outside the capital, neighbourhoods in Bristol, Manchester, Salford, Edinburgh, Portsmouth and Sheffield also fall within the top 0.5% of OAs nationally. While high levels of excess demand in central London are well known, the prominence of cities such as Manchester and Bristol in the upper tail of the national distribution is particularly notable.

Figure 6 examines the relationship between the housing gap per km^2 and local housing costs. The left panel compares the sales-market gap with average transaction prices, while the right panel compares the rental-market gap with average rents. Both panels show a positive relationship, but this is concentrated among the most expensive locations—approximately those with average sales prices above GBP 725,000 or average monthly rents above GBP 2,500. Below these thresholds, the relationship is weak, and many lower-priced areas nevertheless exhibit substantial excess demand. This reinforces the argument developed above: prices alone provide an incomplete measure of housing demand and can obscure important variation across locations.

Housing policy is commonly designed around targets and demand assessments defined for relatively large administrative areas. Our hyper-local measure shows that these averages can conceal substantial variation in housing pressure within them. Figure A.1 illustrates this using Liverpool and Cambridge, two large cities in the UK: both have positive housing gaps overall, yet Liverpool contains neighbourhoods ranging from almost no excess demand to more than 2,000 excess searchers per km^2 , while Cambridge is much more spatially uniform. More broadly, 69% of the variation in the housing gap per km^2 outside London occurs within rather than between Local Authorities. Although these estimates are from Great Britain, the implication is general: geographically aggregated measures can obscure where housing demand is concentrated, while search data can inform not only how much housing to build, but where within planning areas new supply is most likely to meet demand.

4.2 Accompanying mapping tool

A practical limitation of housing-search data is that, unlike transaction records in many countries, they are rarely freely available. To make part of the informational value of these data accessible to researchers, policymakers and the public, we release our Output Area-level housing-gap estimates through an open online mapping platform, together with aggregates for Local Authorities. The tool visualises hyper-local housing demand alongside selected housing and demographic indicators and is available at <https://wheretobuild.warwick.ac.uk> (Figure A.3). By making these estimates publicly accessible, the platform allows users to identify where demand is concentrated within broader planning areas, compare neighbourhoods that are obscured by aggregate statistics, and examine where additional housing supply may be most readily absorbed. More broadly, it provides a scalable demonstration of how otherwise costly search data can be converted into an open resource for studying and informing the spatial allocation of housing.

5 Discussion


House prices are routinely used across the physical and social sciences to infer how people value places and respond to spatial shocks. These results then become important inputs into policy decisions ranging from future proofing built environments for climate shocks, deciding insurance premia, and providing new amenities such as parks and schools.

Figure 5: OA Level Map of Gap per km² Rank



Note: The above map shows the gap per km² rank at the OA level for Great Britain. For details on the construction of the gap per km², please refer to Appendix C.

Figure 6: Relationship between average sales/rental prices and the housing gap per km²



Note: The above figures show the scatter plot at the OA level of the sales and rental prices against the gap per km². Price data are collected from Rightmove, which is also used to construct the gap per km². For details regarding the construction of the gap, please refer to Appendix C.

Yet, prices are equilibrium outcomes observed only for properties that transact, and therefore combine demand with housing supply, market institutions, seller behaviour and selection into transactions. These limitations are especially consequential when shocks are short-lived, geographically concentrated, or occur in locations with few or no prior sales. In such settings, transaction prices may provide a delayed, noisy or biased measure of changing preferences.

We show that housing-search data provide a direct measure of spatial demand, available at high frequency, at hyper-local scales and even in places without prior transactions. Across the COVID-19 pandemic, flooding events and large new-build developments, searches reveal changes in preferences that transaction prices either obscure or cannot measure. These are only three examples, but our methodology can be applied widely: Search data allow researchers to observe how people respond to environmental, public-health and structural change before those responses are filtered through market transactions. This is particularly helpful as an early signal for changes in demand or to assess the importance of repeated shocks for human behaviour (for example, repeated flooding), when one shock alone may only cause a temporary deviation in demand but be indicative of more substantive changes in behaviour if repeated. For those continuing to work with transaction prices, we suggest that they additionally run their analyses with transaction volume, as a diagnostic test for the extent to which their results may be influenced by supply factors.

In addition to demonstrating the conceptual importance of search data as a measure of demand, we discuss three novel results: We show that the increased demand for private greenspace during the COVID-19 pandemic returned to pre-pandemic levels earlier than housing transaction prices suggest; that one-time flooding events cause substantive drops in location demand for flooded areas that is temporary and non-detectable in transaction data; and that location demand cannot be induced by new supply of residential homes—a result that cannot be estimated with transaction data due to the lack of transactions prior to the new housing development. These three examples highlight the policy importance of the search measure. They have implications for policies on flood prevention and insurance, providing public versus private greenspace, and the success of new towns in solving the housing affordability crisis. The Where-ToBuild mapping tool demonstrates how these data can also be converted into an open resource for measuring where excess housing demand is concentrated. As search data are rarely as easily available as transactions data, we have aimed to share our results by making this tool publicly accessible. Amongst other things, it can be used to inform policy or industry decisions on housing in Great Britain.

Finally, although search intensity is not itself a monetary measure of value, search data can be given additional structure to recover willingness to pay. Search shares across locations, or the underlying location choices made by individual users, can be embedded within standard discrete-choice models in which prices enter alongside local amenities and other attributes. Under the usual assumptions, the

resulting estimates translate changes in search behaviour into marginal willingness to pay for location characteristics and spatial shocks. Search data can therefore do more than reveal where and when demand changes: they can also provide the micro-foundations for direct valuation which we leave for future work.

6 Methodology

6.1 Additional data sources

We combine the search data with information from several additional datasets to construct our measures and empirical samples. Unless otherwise stated, our unit of observation is the Output Area (OA), the smallest census geography in Great Britain. OAs generally contain between 40 and 250 households, with approximately 125 households on average.

National House Building Council (NHBC): We use NHBC data, covering approximately 70% of new-build housing in Great Britain, to identify and geolocate housing-supply shocks used in Section 3.3. We focus on development sites comprising at least 20 units for which construction began in 2019 or later. We repair invalid project geometries and combine all polygons belonging to the same development site into a single spatial unit.

HM Land Registry (HMLR): To compare our search-based demand measure with conventional housing-market outcomes (throughout Section 3), we use property-transaction records from HM Land Registry. These data cover England and Wales. We restrict the sample to transactions from 2019-2024 to align it with the time period used in our search data.

Ordnance Survey: We use AddressBase to identify residential buildings and the National Geographic Database to measure land use and construct our measure of private greenspace (used in Section 3.1).

Flood data: In Section 3.2, we use flood outlines for England and Wales from the Environment Agency and DataMapWales respectively which cover the history of floods in the both regions and the areas they cover.

6.2 Calculating demand from searches

We construct a weekly panel of property searches at the Output Area (OA) level. Searches on the property platform are recorded against a search location, which may be a named place (for example, “Bethnal Green”), a postcode area or district (for example, “SE1”), or a railway station (for example, “Manchester Piccadilly”). Users may also specify a radius around the selected location. Each location–radius combination therefore defines a search polygon that may intersect multiple OAs.

Let j index users, i index OAs, p index search polygons, and t index year-weeks. Let S_{jpt} denote the number of non-bot searches made by user j in polygon p during week t .

To allocate searches from polygons to OAs, define the share of polygon p assigned to OA i as

$$\alpha_{ip} = \frac{\text{Area}(A_i \cap P_p)}{\text{Area}(P_p)}, \quad (2)$$

where A_i denotes the geographical boundary of OA i and P_p denotes the boundary of search polygon p . Thus, an OA receives a larger share of a polygon’s searches when it accounts for a larger proportion of the polygon’s area. Since the OAs partition the relevant geographical area, the allocation shares for each polygon sum to one:

$$\sum_i \alpha_{ip} = 1. \quad (3)$$

The total number of searches allocated to OA i in week t is then

$$S_{it} = \sum_j \sum_p \alpha_{ip} S_{jpt}. \quad (4)$$

6.3 Methodology details for COVID-19 event study

6.3.1 Dataset Construction

To construct our sample, we filter building coordinates from AddressBase Premium to houses, semi-detached houses, terraced houses, and flats, excluding buildings removed before 2019 within England and Wales. We spatially join the remaining points to their building polygons from OS Open Zoomstack and keep the relevant residential polygons.

We construct private greenspace from de-duplicated land parcels from the OS National Geographic Database. For each parcel, we subtract the area occupied by residential building footprints to identify the remaining private greenspace. We divide this area, measured in m^2 , by the number of non-flat residential units associated with the parcel to obtain private greenspace per dwelling. Parcels containing only flats are assigned a value of zero as greenspace is shared with others in the building.

Where a land parcel intersects more than one Output Area (OA), we allocate its private greenspace measure across OAs in proportion to the share of the parcel lying within each OA. We then construct the OA-level measure as the median private greenspace per dwelling unit across the parcels allocated to each OA.

To compare our search-based demand measure with conventional housing-market outcomes, we use transaction data from HMLR. We assign each transaction to an OA using its postcode and aggregate the data to the OA-month level. For transaction volumes, OA-months with no sales are assigned a value of zero. For transaction prices, we calculate the mean sale price among transactions occurring within each OA-month; prices are therefore unobserved in OA-months with no transactions. Appendix Table A4 documents the sparsity of the transaction data and shows that fewer than 9% of OAs record transactions in both the pre- and post-March 2020 periods at the year-week-OA level. We merge the resulting monthly transaction counts and average prices with the greenspace and search measures to construct the regression sample.

6.3.2 Empirical Implementation

We begin by estimating the effect of the pandemic on demand for private green space using our search measure as the dependent variable. We then re-estimate the same specification using the number of transactions, the number of transactions filtered to new builds only, and average prices as outcomes. New builds follow the same construction as transactions, restricted to newly built properties only. The regression specification common to all five outcomes is:

$$Y_{i,t} = \sum_{\substack{k=\bar{k} \\ k \notin \{-1, -6\}}}^{\bar{k}} \beta_k (\text{Treat}_i \times \mathbb{1}\{t - T = k\}) + \gamma_{TTWA,t} + \delta_{LSOA,s(t)} + \theta_i + \varepsilon_{i,t} \quad (5)$$

where $Y_{i,t}$ is the outcome of interest for Output Area (OA) i in period t . A period is a week in the weekly specifications (searches) and a month in the monthly specifications (prices, number of transactions, number of new-build transactions). Treat_i is a *continuous* treatment-intensity measure equal to the median private green space (in $100m^2$) in OA i . The term $t - T$ denotes event time relative to the shock period T (the last week of February 2020, or February 2020, in the weekly and monthly specifications respectively).

$\gamma_{TTWA,t}$ are TTWA-by-calendar-time fixed effects, where t is year-month in the monthly specifications and year-week in the weekly ones. This soaks up shocks to the TTWA at particular points in time. $\delta_{LSOA,s(t)}$ are LSOA by week or month fixed effects depending on specification, where $s(t)$ is month of the year (e.g. January, February, etc) in the monthly specifications, or week of the year (e.g. week 52) in the weekly specifications. δ accounts for recurring seasonality by month or week of the year. θ_i are OA fixed effects. Standard errors are clustered at the OA level.

We estimate four versions of equation (5), which differ in the outcome, its transformation, and the frequency at which observations are defined. We use the following outcome variables:

1. log buying searches (weekly);
2. transactions, Inverse Hyperbolic Sine-transformed (monthly);
3. new-build transactions, Inverse Hyperbolic Sine-transformed (monthly);
4. log average prices (monthly).

Transactions are inverse-hyperbolic-sine (IHS) transformed to accommodate OA-period combinations containing zero transactions; the log specifications are estimated on the strictly positive support. Weekly specifications use the event-time window $(\underline{k}, \bar{k}) = (-60, 201)$ and monthly specifications use $(-14, 46)$. In all specifications $k = -1$ is the omitted reference period, with the exception of the new-build specification, where we normalise to $k = -6$ to account for potential expectations around construction commencement.

For the IHS outcomes, as the raw coefficient is not a semi-elasticity, we transform the raw coefficient as recommended in [Bellemare and Wichman \(2020\)](#) using:

$$\tilde{\beta}_k = \beta_k \cdot \frac{\sqrt{\bar{y}^2 + 1}}{\bar{y}},$$

where $\bar{y}$ is the mean of the *untransformed* outcome in the estimation sample. In all specifications, the reported coefficients correspond to a 100 m² increase in median private green space and expressed in percent for ease of interpretation.

6.4 Methodology details for flooding event study

6.4.1 Dataset construction

To construct the estimation sample for our analysis of floods, we first use the OS Classification Scheme to identify residential buildings and assign each building to its corresponding Output Area (OA). We then spatially intersect these buildings with the 20 largest flood events in England and Wales, ranked by the size of the affected area, that occurred during the search-data period from January 2019 to May 2024. An OA is classified as treated for a given event if at least one residential building within it was affected by the flood.²

For each flood, the comparison group consists of OAs in the same Travel to Work Area (TTWA) as the treated OAs that were not exposed to any of the 20 selected flood events. This produces geographically proximate comparison areas that remained untreated throughout the study period. Due to the staggered nature of the natural experiment, we construct a separate dataset for each flood comprised of treated and untreated OAs for that flooding event, and then stack these event-specific samples. This structure takes into account that only never flooded output areas serve as control groups for a set of output areas treated in a given flood. Following the procedures described in Section 6.3, we calculate total housing searches at the OA-week and OA-month levels and construct monthly measures of transaction volumes and average transaction prices. We then merge these outcomes with the stacked flood sample to create the regression dataset.

6.4.2 Empirical Implementation

We begin by estimating the causal effect of floods on housing demand using our search measure as the dependent variable. We then re-estimate the same specification using the number of transactions and average prices as outcomes. The regression specification common to all three outcomes is:

$$Y_{i,s,t} = \sum_{\substack{k=\underline{k} \\ k \neq -1}}^{\bar{k}} \beta_k (\text{Treat}_{is} \times \mathbb{1}\{t - T_s = k\}) + \gamma_{TTWA,t} + \alpha_{i,s} + \varepsilon_{i,s,t} \quad (6)$$

²Defining treatment using affected residential buildings, rather than the intersection with the OA boundary alone, prevents OAs from being classified as treated when flooding affected only non-residential land.

where $Y_{i,s,t}$ is the outcome of interest for Output Area (OA) i , flood event s , and period t ; a period is a week in the weekly specifications and a month in the monthly specifications. $\text{Treat}_{i,s}$ is an indicator equal to one if OA i is exposed to flood event s . The term $t - T_s$ denotes event time (the number of periods relative to flood s , which occurs at T_s), so $k = 0$ is the period of the flood. $\gamma_{TTWA,t}$ are TTWA-by-period fixed effects and $\alpha_{i,s}$ are OA-by-stack fixed effects. Standard errors are clustered at the OA-stack level. We estimate three versions of equation (6), which differ in the outcome, its transformation, and the frequency at which observations are defined. We consider the following three outcome variables:

1. log buying searches (weekly);
2. transactions, IHS-transformed (monthly);
3. log average prices (monthly).

For the transactions event studies, we assign a value of zero to OA-period observations in which no transaction occurred; however, due to the low frequency of transactions at the weekly level and the fact that prices are only observed if a transaction occurs, both transactions and prices are estimated only at the monthly level as in the greenspace analysis. Transactions are IHS transformed to accommodate OA-period combinations containing zero transactions. Weekly specifications use the event-time window $(\bar{k}, \bar{k}) = (-24, 52)$ and monthly specifications use $(-6, 12)$, i.e. half a year before and a year after the flood; in both cases $k = -1$ is the omitted reference period and $k = 0$ is the period of the flood. As explained in Section 6.3, the IHS coefficients are transformed and all reported coefficients are rescaled to be interpretable as percent.

6.5 Methodology details for New Build Demand Shock

6.5.1 Dataset Construction

We use the NHBC data set to identify significant new-build sites. The data is organised into construction projects across Great Britain and development sites within these projects. We keep only sites that are active or have been completed. With that, our cleaned initial data set consists of approximately 141,000 sites and 3.2 million builds. We are only interested in sizable new-build sites which we define as 20 new units or more. We define the shock week as the first week of construction of the first build within a site and restrict our sample to sites that had the first build starting post-2019 corresponding with the start of the search data. We also restrict the dates of the shocks to ensure we have a complete pre- and post-treatment window of search data for the analysis, resulting in sites with start dates after June 2019 and before May 2023. Our final sample used consists of 3,200 sites that meet these criteria. We then assign the intersecting OAs for each site, and create the corresponding controls for a site through the following algorithm:

1. Across all OAs in Great Britain, we keep OAs that are in MSOAs that never intersected a treated OA.
2. From the surviving OAs, we assign each site its control OAs such that the latter lie within the same TTWA as the intersecting treated OAs.

Lastly, we stack all treated and control OAs, creating the final data set with around 265 million observations.

We note here, that this analysis is not possible using property transactions. In our sample of new-build sites, the average number of transactions per OA over the 6 months prior to the supply shock is only 5. Therefore, this evidence presents novel evidence on the extent to which new supply can induce demand — a mechanism that is often implicitly assumed by policymakers to reflect reality.

6.5.2 Empirical Implementation

Using the constructed dataset, Figure 4 is produced by running the following regression specification and clustering the standard errors at the OA-site level:

$$Y_{s,o,t} = \sum_{\substack{k=-26 \\ k \neq -5}}^{44} \beta_k (\text{Treat}_{s,o} \times \mathbb{1}\{t - T_s = k\}) + \gamma_{s,o} + \lambda_{s,t} + \varepsilon_{s,o,t} \quad (7)$$

where $Y_{s,o,t} = \ln(\text{Searches}_{s,o,t})$ is the log number of buying searches in site stack s , OA o , and year-week t . Treat_{so} is an indicator equal to one if OA o in stack s is a treated OA, and zero otherwise. The term $t - T_s$ denotes event time (the number of weeks relative to the shock in stack s , which occurs at year-week T_s), so $k = 0$ is the week of the shock. The effects are normalised to $k = -5$, the omitted reference period. $\gamma_{s,o}$ are OA-by-stack fixed effects and $\lambda_{s,t}$ are year-week-by-stack fixed effects.


References

- Aguilar-Velázquez, D., I. Rivera Islas, G. Romero Tecua, and A. Valenzuela-Aguilera (2024). Gentrification and access to housing in Mexico City during 2000 to 2022. *Proceedings of the National Academy of Sciences* 121(10), e2314455121.
- Althoff, L., F. Eckert, S. Ganapati, and C. Walsh (2022). The geography of remote work. *Regional Science and Urban Economics* 93, 103770.
- Andersen, S., C. Badarinza, L. Liu, J. Marx, and T. Ramadorai (2022). Reference dependence in the housing market. *American Economic Review* 112(10), 3398–3440.
- Baum-Snow, N. and L. Han (2024a). The microgeography of housing supply. *Journal of Political Economy* 132(6), 1897–1946.
- Baum-Snow, N. and L. Han (2024b). The microgeography of housing supply. *Journal of Political Economy* 132(6), 1897–1946.
- Bayer, P., F. Ferreira, and R. McMillan (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of political economy* 115(4), 588–638.
- Bellemare, M. F. and C. J. Wichman (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics* 82(1), 50–61.
- Beltrán, A., D. Maddison, and R. Elliott (2019). The impact of flooding on property prices: A repeat-sales approach. *Journal of Environmental Economics and Management* 95, 62–86.
- Bézy, T. (2025). Real estate wealth inequality and exposure to natural disasters.
- Chay, K. Y. and M. Greenstone (2005). Does air quality matter? evidence from the housing market. *Journal of Political Economy* 113(2), 376–424.
- Cvijanović, D. (2014). Real estate prices and firm capital structure. *The Review of Financial Studies* 27(9), 2690–2735.
- Ellen, I. G., A. Kulka, and H.-L. Liao (2026). High-resolution mobility data reveals a stalling of racial integration in new york city during the covid-19 pandemic. *Available at SSRN* 6344178.
- Fan, Y., Q. Gao, Y. E. Sitoh, and W. X. Wan (2025). Post-flood selective migration interacts with media sentiment and income effects. *Nature Climate Change* 15(6), 619–626.
- Ferreira, F. V. and M. Wong (2022). Neighborhood choice after covid: the role of rents, amenities, and work-from-home. Technical report, National Bureau of Economic Research.
- Gibbons, S. and S. Machin (2005). Valuing rail access using transport innovations. *Journal of Urban Economics* 57(1), 148–169.
- Glaeser, E. L., J. Gyourko, and R. E. Saks (2005). Why have housing prices gone up? *American Economic Review* 95(2), 329–333.
- Gupta, A., V. Mittal, J. Peeters, and S. Van Nieuwerburgh (2022). Flattening the curve: Pandemic-induced revaluation of urban real estate. *Journal of Financial Economics* 146(2), 594–636.
- Gyourko, J. and A. Saiz (2006). Construction costs and the supply of housing structure. *Journal of regional Science* 46(4), 661–680.
- Hino, M. and M. Burke (2021). The effect of information about climate risk on property values. *Proceedings of the National Academy of Sciences* 118(17), e2003374118.
- Hu, L. and S. He (2024). Entrance opportunity vs. academic performance: unpacking the nonlinear capitalization effects of multidimensional school qualities on housing sales and rental prices. *Humanities and Social Sciences Communications* 11, 1254.
- Huang, X., H. Yan, and Z. Zhang (2025). Economic impacts of capital city relocation in myanmar. *Nature Cities* 2(4), 299–303.

- Lang, C., J. VanCeylon, and A. W. Ando (2023). Distribution of capitalized benefits from land conservation. *Proceedings of the National Academy of Sciences* 120(18), e2215262120.
- Li, H., W. Xiao, Y. Zhou, and I. K. Moise (2025). COVID-19 impacts on residential preferences for parks and green-blue spaces: A case study from Miami-Dade county, Florida. *Cities* 157, 105628.
- Linden, L. L. and J. E. Rockoff (2008). Estimates of the impact of crime risk on property values from megan’s laws. *American Economic Review* 98(3), 1103–1127.
- Liu, Y. and Y. Tang (2021). Epidemic shocks and housing price responses: Evidence from china’s urban residential communities. *Regional Science and Urban Economics* 89, 103695.
- Malik, K., S. Kim, and B. J. Cultice (2023). The impact of remote work on green space values in regional housing markets. *Journal of Housing Economics* 62, 101967.
- Mamun, S., A. Castillo-Castillo, K. Swedberg, J. Zhang, K. J. Boyle, D. Cardoso, C. L. Kling, C. Nolte, M. Papenfus, D. Phaneuf, and S. Polasky (2023). Valuing water quality in the United States using a national dataset on property values. *Proceedings of the National Academy of Sciences* 120(15), e2210417120.
- Met Office (2020). Storm ciara. Technical report, Met Office.
- MHCLG (2026, March). New towns draft programme. Consultation. Accessed 18 July 2026.
- Mian, A. and A. Sufi (2011). House prices, home equity–based borrowing, and the us household leverage crisis. *American economic review* 101(5), 2132–2156.
- Monte, F., C. Porcher, and E. Rossi-Hansberg (2025, November). Remote work and city structure. Working paper, Georgetown University and University of Chicago.
- Oswald, A. J. and S. Wu (2010). Objective confirmation of subjective measures of human well-being: Evidence from the U.S.A. *Science* 327(5965), 576–579.
- Qian, X., S. Qiu, and G. Zhang (2021). The impact of covid-19 on housing price: Evidence from china. *Finance research letters* 43, 101944.
- Ramani, A., J. Alcedo, and N. Bloom (2024). How working from home reshapes cities. *Proceedings of the National Academy of Sciences* 121(45), e2408930121.
- Roback, J. (1982). Wages, rents, and the quality of life. *Journal of political Economy* 90(6), 1257–1278.
- Rosen, S. (1979). Wages-Based Indexes of Urban Quality of Life. In P. Mieszkowski and M. Straszheim (Eds.), *Current Issues in Urban Economics*, pp. 74–104. Baltimore: Johns Hopkins University Press.
- Rowe, F., A. Calafiore, D. Arribas-Bel, K. Samardzhiev, and M. Fleischmann (2023). Urban exodus? Understanding human mobility in Britain during the COVID-19 pandemic using Meta-Facebook data. *Population, Space and Place* 29(1), e2637.
- Saiz, A. (2008). On local housing supply elasticity. *Available at SSRN* 1193422.
- Skouralis, A., N. Lux, and M. Andrew (2024). Does flood risk affect property prices? Evidence from a property-level flood score. *Journal of Housing Economics* 66, 102027.
- Thompson, J. J., R. L. Wilby, J. K. Hillier, R. Connell, and N. Gunn (2026). Flood-affected homes in England lose value and are owned longer than those in unaffected areas. *Communications Earth & Environment*. Volume and article number to confirm on publication.
- UK Government (2024). Plan for change. Government policy programme. Accessed 19 July 2026.
- Xu, Y., L. E. Olmos, D. Mateo, A. Hernando, X. Yang, and M. C. González (2023). Urban dynamics through the lens of human mobility. *Nature computational science* 3(7), 611–620.
- Yamagishi, A. (2021). Minimum wages and housing rents: Theory and evidence. *Regional Science and Urban Economics* 87, 103649.

A Additional Figures & Tables

Figure A.1: Differences in Variation within LAs




Note: The above maps show the gap per km² at the OA level across several LAs. The gap per km² variable is constructed from Rightmove data. For details regarding the construction of the variable, see Appendix C.


Table A1: Descriptive Information on Rightmove Data (2023-2024)

Metric	Buying	Letting	Combined
Distinct Users (All)	57,514,595	18,368,066	69,569,246
Distinct Users (emailed)	1,317,158	2,087,639	3,366,794
Distinct Users (saved)	1,903,819	1,074,503	2,867,241
Distinct Users (emailed and saved)	189,178	354,035	542,373
Total Searches (All)	1,816,107,042	394,831,293	2,210,938,335
Total Searches (emailed)	203,216,347	129,108,179	361,886,253
Total Searches (saved)	203,249,830	64,707,167	280,777,704
Total Searches (emailed and saved)	32,582,814	29,438,985	67,633,484
Listings	2,242,668	1,294,453	3,537,121
Unique Search Locations	381,457	270,659	391,454
Avg Time Per Session (mins)	11	10	11
Avg Searches Per Session	4	4	4
Avg Sessions Per User	7	5	7
Avg Searches Per User	32	21	32


Note: The above table shows Rightmove users and session descriptive measures between 2023-2024 in the sales and letting markets. User-level statistics are segmented into all non-bot users, users who contacted an agent, users who saved a listing, and users who both saved a listing and contacted an agent. Session-level statistics are calculated using all non-bot sessions.



(a) Covid Shock



(b) Floods Shock




(c) Supply Shock

Figure A.2: Saved or Contacted User Sample

Note: The above figures present the results of using serious searches as an outcome variable. Serious searches are constructed by restricting focus to users who ever saved a property or contacted an agent on the platform. Panels (a), (b), and (c) respectively show the results upon re-estimating equations 5, 6, 7.

Figure A.3: Mapping tool



Note: This Figure shows a screenshot from our online mapping tool that visualises our new measure of housing demand—the housing gap—introduced in this paper and can be found [here](#).

Table A2: Rank–rank correlations between user restriction groups

Restriction 1	Restriction 2	Rank–Rank Correlation (ρ)
All non-bot users	Saved a property	0.981
All non-bot users	Contacted estate agent	0.970
All non-bot users	Saved and contacted estate agent	0.946
Contacted estate agent	Saved and contacted estate agent	0.970
Contacted estate agent	Saved a property	0.948
Saved and contacted estate agent	Saved a property	0.966

Note: Table shows the rank-rank correlation of buying searches at the output-area level for various combinations of samples. Restriction 1 and 2 refer to the restrictions applied to create a given sample. All samples are restricted to non-bot users.

Table A3: Transaction activity in the flood regression data

Statistic	Flood sample
Week-Year-OA with transaction	8.8%
Week-Year-OA with transaction both pre and post flood	8.0%
Month-Year-OA with transaction	29.1%
Month-Year-OA with transaction both pre and post flood	26.3%
Year-OA with transaction	73.2%
Year-OA with transaction both pre and post flood	64.1%

Note: Flood statistics are calculated using a 24-week pre-treatment and 52-week post-treatment window for each flood event, matching the transaction event-study window.

Table A4: Transaction activity in the greenspace regression data

Statistic	Greenspace sample
Week-Year-OA with transaction	8.9%
Week-Year-OA with transaction both pre and post Covid	8.8%
Month-Year-OA with transaction	31.5%
Month-Year-OA with transaction both pre and post Covid	31%
Year-OA with transaction	94.4%
Year-OA with transaction both pre and post Covid	92.1%

Note: Greenspace statistics 60-week pre-Covid and 201-week post-Covid window, matching the covid event-study window.

Table A5: Transaction activity in the supply shock data

Statistic	Supply shock sample
Week-Year-OA with transaction	6.6%
Week-Year-OA with transaction both pre and post supply shock	6.5%
Month-Year-OA with transaction	23%
Month-Year-OA with transaction both pre and post supply shock	22.9%
Year-OA with transaction	67.6%
Year-OA with transaction both pre and post supply shock	67.1%

Note: Supply shock statistics 26-week pre-supply shock and 44-week post-shock window, matching the supply shock event-study window.

B A model of seller reference dependence

Consider a housing market with properties indexed by k , buyers indexed by i , locations indexed by l , and time indexed by t .

Latent property quality is given by

$$q_{klt} = \alpha_k + \lambda_l + \xi_t + shock_{lt},$$

where α_k is a property fixed effect and λ_l is the location component capturing local demand or amenity conditions, ξ_t is the time-varying component of local demand.

Buyer i 's willingness to pay for property j is

$$W_{iklt} = \mu_i + q_{klt} + \varepsilon_{iklt},$$

where μ_i is a buyer fixed effect and ε_{iklt} is an idiosyncratic buyer-property shock.

Seller j 's reservation price is

$$R_{jklt} = m_j + q_{klt} + \phi(p_j^{ref} - q_{klt}) + \eta_{jklt},$$

where m_j captures moving costs or outside options, p_j^{ref} is a reference price, $\phi(\cdot)$ captures reference dependence or loss aversion with $\phi'(\cdot) > 0$, and η_{jklt} is a seller-property-specific shock. The term $\phi(p_j^{ref} - q_{klt})$ implies that when current latent values fall sufficiently below the seller's reference point, sellers require a higher reservation price to transact. This can be seen using

$$\frac{dR_{jklt}}{dq} = 1 - \phi'(p_j^{ref} - q_{klt})$$

A transaction occurs if and only if buyer willingness to pay exceeds the seller reservation price:

$$D_{ijklt} = \mathbf{1}\{W_{iklt} \geq R_{jklt}\}.$$

Substituting for buyer willingness to pay and seller reservation prices yields

$$D_{ijklt} = \mathbf{1}\left\{\mu_i - m_j + \varepsilon_{iklt} - \eta_{jklt} \geq \phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - shock_{lt})\right\}.$$

Hence the transaction decision depends on whether the net surplus from trade exceeds a seller-side reference-dependent cutoff,

$$c_{jklt} = \phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - shock_{lt}).$$

Because

$$\frac{\partial c_{jklt}}{\partial shock_{lt}} = -\phi'(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - shock_{lt}) < 0,$$

negative local demand shocks raise the cutoff for trade. Observed transactions therefore become increasingly selected toward buyer-seller matches with high values of

$$\mu_i - m_j + \varepsilon_{iklt} - \eta_{jklt}.$$

Observed transaction prices are determined through bargaining between buyers and sellers:

$$p_{ijklt} = \tau W_{iklt} + (1 - \tau) R_{jklt}, \quad \tau \in [0, 1].$$

Substituting in for buyer willingness to pay and seller reservation prices yields

$$p_{ijklt} = q_{klt} + \tau \mu_i + (1 - \tau) m_j + (1 - \tau) \phi(p_j^{ref} - q_{klt}) + \tau \varepsilon_{iklt} + (1 - \tau) \eta_{jklt}.$$

Substituting further for latent quality,

$$p_{ijklt} = \alpha_k + \lambda_l + \xi_t + \tau \mu_i + (1 - \tau) m_j + (1 - \tau) \phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - shock_{lt}) + \tau \varepsilon_{iklt} + (1 - \tau) \eta_{jklt} + shock_{lt}$$

Suppose the econometrician estimates a difference-in-differences specification.

The econometrician estimates

$$p_{ijklt} = \theta_l + \kappa_t + \delta (\text{Treat}_l \times \text{Post}_t) + u_{ijklt},$$

using observed transaction prices only.

The difference-in-differences estimand is therefore

$$\begin{aligned} \delta = & E[p_{ijklt}(1) - p_{ijklt}(0) \mid D_{ijklt} = 1, \text{Treat}_l = 1] \\ & - E[p_{ijklt}(1) - p_{ijklt}(0) \mid D_{ijklt} = 1, \text{Treat}_l = 0]. \end{aligned} \quad (\text{B.1})$$

where $p_{ijklt}(1)$ and $p_{ijklt}(0)$ denote equilibrium transaction prices post- and pre-treatment respectively. Realising that the conditional expectation of μ , ϵ , m and η varies with the cutoff yields:

$$\begin{aligned} E[p_{ijklt}(s) \mid D_{ijklt} = 1, \text{Treat}_l = 1] = & E[\alpha_k + \lambda_l + \xi_t(s) + shock_{lt}(s) + \tau\mu_i(s) + (1 - \tau)m_j(s) \\ & + (1 - \tau)\phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t(s) - shock_{lt}(s)) \\ & + \tau\epsilon_{iklt}(s) + (1 - \tau)\eta_{jklt}(s) \mid D_{ijklt} = 1, \text{Treat}_l = 1] \end{aligned}$$

where $s \in \{0, 1\}$.

Then, substituting in the terms of equation (B.1), we obtain:

$$\begin{aligned} E[p_{ijklt}(1) - p_{ijklt}(0) \mid D_{ijklt} = 1, \text{Treat}_l = 1] = & E[\Delta(\xi_t) + \Delta shock_{lt} + \tau\Delta\mu_i + (1 - \tau)\Delta m_j \\ & + (1 - \tau)\Delta\phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - shock_{lt}) \\ & + \tau\Delta\epsilon_{iklt} + (1 - \tau)\Delta\eta_{jklt} \mid D_{ijklt} = 1, \text{Treat}_l = 1] \end{aligned}$$

where

$$\Delta\phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - shock_{lt}) \equiv \phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_1 - shock_{l,1}) - \phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_0).$$

Simplifying then gives:

$$\begin{aligned} E[p_{ijklt}(1) - p_{ijklt}(0) \mid D_{ijklt} = 1, \text{Treat}_l = 1] = & \Delta(\xi_t) + \Delta shock_{lt} + (1 - \tau)\Delta\phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - \\ & shock_{lt}) + E[\tau\Delta\mu_i + (1 - \tau)\Delta m_j + \tau\Delta\epsilon_{iklt} \\ & + (1 - \tau)\Delta\eta_{jklt} \mid D_{ijklt} = 1, \text{Treat}_l = 1] \end{aligned}$$

Similarly for the untreated

$$E[p_{ijklt}(1) - p_{ijklt}(0) \mid D_{ijklt} = 1, \text{Treat}_l = 0] = \Delta(\xi_t) + (1 - \tau)\Delta'\phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t)$$

where

$$\Delta'\phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t) \equiv \phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_1) - \phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_0).$$

Substituting in the equation for δ yields:

$$\begin{aligned} \delta = & \underbrace{\Delta shock_{lt}}_{\text{direct demand effect}} + (1 - \tau) \underbrace{\left[\Delta\phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - shock_{lt}) - \Delta'\phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t) \right]}_{\text{seller reference-dependence effect}} \\ & + \underbrace{E[\tau\Delta\mu_i \mid D_{ijklt} = 1, \text{Treat}_l = 1]}_{\text{buyer-composition selection}} + \underbrace{E[(1 - \tau)\Delta m_j \mid D_{ijklt} = 1, \text{Treat}_l = 1]}_{\text{seller-composition selection}} \\ & + \underbrace{E[\tau\Delta\epsilon_{iklt} + (1 - \tau)\Delta\eta_{jklt} \mid D_{ijklt} = 1, \text{Treat}_l = 1]}_{\text{buyer-seller match selection}}. \end{aligned} \quad (\text{B.2})$$

The direct demand effect captures the treatment-induced change in the latent value of housing in treated locations. The seller reference-dependence effect captures the additional change in observed transaction prices arising because the shock changes sellers' reservation prices relative to their reference points, net of the corresponding change in untreated locations.

The remaining terms capture selection into observed transactions. Buyer-composition selection reflects changes in the underlying willingness to pay of buyers who transact, while seller-composition selection reflects changes in the moving costs or outside options of sellers who transact. Buyer-seller match selection captures changes in the idiosyncratic quality of observed matches, with transactions becoming concentrated among buyers with relatively high property-specific valuations and sellers with relatively low property-specific reservation-price shocks.

B.1 Transaction Volumes as a Robustness Check

The model provides a direct robustness check for studies using transaction prices to measure changes in local demand. Recall that a transaction occurs when

$$D_{ijklt} = \mathbf{1} \left\{ \mu_i - m_j + \varepsilon_{iklt} - \eta_{jkl t} \geq \phi(p_j^{ref} - \alpha_k - \lambda_l - \xi_t - shock_{lt}) \right\}.$$

In the absence of seller reference dependence, such that $\phi(\cdot) = 0$, the transaction condition becomes

$$D_{ijklt} = \mathbf{1} \{ \mu_i - m_j + \varepsilon_{iklt} - \eta_{jkl t} \geq 0 \}.$$

The local demand shock no longer enters the transaction decision. This is because the shock raises buyer willingness to pay and the seller reservation price by the same amount. It therefore changes the equilibrium transaction price, but leaves the surplus from trade and the probability of transaction unchanged.

Suppose the econometrician estimates the corresponding difference-in-differences specification for transaction volumes, where the latter is defined as:

$$V_{lt} \equiv \sum_{i,j,k} D_{ijklt},$$

so the regression would be,

$$V_{lt} = \theta_l^V + \kappa_t^V + \delta^V (\text{Treat}_l \times \text{Post}_t) + \nu_{lt}.$$

The difference-in-differences estimand is

$$\delta^V = E[V_{lt}(1) - V_{lt}(0) \mid \text{Treat}_l = 1] - E[V_{lt}(1) - V_{lt}(0) \mid \text{Treat}_l = 0]$$

When $\phi(\cdot) = 0$, the transaction decision is unaffected by the local demand shock, so that

$$D_{ijklt}(1) = D_{ijklt}(0),$$

and since

$$E[V_{lt}(1)] = \sum_{i,j,k} E[D_{ijklt}(1)] = \sum_{i,j,k} E[D_{ijklt}(0)] = E[V_{lt}(0)],$$

it follows then,

$$\delta^V = 0.$$

By contrast, with seller reference dependence,

$$\frac{\partial D_{ijklt}}{\partial shock_{lt}} \neq 0$$

for buyer-seller matches close to the transaction threshold. A negative demand shock raises the reference-dependent cutoff and causes some matches that would previously have transacted to exit the observed transaction sample. The corresponding transaction-price estimate is then conditional on a changing set of buyers, sellers and buyer-seller matches.

The model therefore implies that transaction volumes should be examined alongside transaction prices. A differential change in transaction volumes rejects the restriction implied by the frictionless model and provides evidence that the price estimate may combine the underlying demand response with selection into transactions. A zero transaction-volume effect is a necessary, though not sufficient, condition for interpreting the transaction-price response as a clean measure of the demand shock.

C Housing Gap

C.0.1 Searcher Calculations

First, we focus on allocation searches to output areas. Let g represent a grid, p a polygon, i a user, n a search number, and m a listing. The calculation relies on the following principles:

- Each user should only count once, therefore for any user, summing over all the OAs in the UK should give a value of 1, for that user.
- Each listing should only count once, therefore for any listing, summing over all the OAs in the UK should give a value of 1, for that listing.

Note in this context g is an OA, while p is a search location on the property platform. As such p may be a location (“Bethnal Green”), a postal area/district (“SE1”), a train station (“Manchester Piccadilly”), and may also have a radius attached.

Users and listings therefore need to be allocated to OAs. This needs to take into account two factors. Users can search across multiple polygons (and listings can appear in multiple polygons), and polygons contain multiple OAs. For searcher calculations, a user i search n is equal to 1 for polygon p if they searched that polygon, and 0 for all others

$$S_{ipn} = \mathbb{1}[i \text{ searched } p] \quad (\text{B.3})$$

Thus, summing over all searches gives a count of searches for each polygon for user i

$$S_{ip} = \sum_n S_{ipn} \quad (\text{B.4})$$

In the first step, we need to allocate the searcher across all the polygons they searched in such that it sums to 1. Hence, let ω_p be some set of weights.

$$\hat{S}_{ip} = \frac{S_{ip}\omega_p}{\sum_p S_{ip}\omega_p} \quad (\text{B.5})$$

We use a weight of $\omega_p = 1$ such that each polygon gets an equal weight.

Once users are allocated across polygons, users still need to be allocated to OAs within those polygons. All OAs receive equal weighting within the polygon. Therefore large polygons give each grid within it a smaller share and vice versa.

$$S_{igp} = \hat{S}_{ip} * \frac{g \in p}{\sum_{g \in p} 1} \quad (\text{B.6})$$

To then get the user allocation to a grid, we sum over all polygons:

$$S_{ig} = \sum_p S_{igp} \quad (\text{B.7})$$

To get total grid level searches we sum over users i

$$\mathcal{S}_g = \sum_i S_{ig} \quad (\text{B.8})$$

C.0.2 Listing Calculations

We count each unique property listing only once, in a similar vein to the approach with searchers. We use a similar allocation procedure as we do for searchers to ensure consistency.

A listing m is equal to 1 for polygon p if that listing appears in that polygon p :

$$L_{mp} = \mathbb{1}[m \in p] \quad (\text{B.9})$$

First, we need to allocate the listing across all the polygons it falls in such that it sums to 1. Let ω_p be some set of weights:

$$\hat{L}_{mp} = \frac{L_{mp}\omega_p}{\sum_p L_{mp}\omega_p} \quad (\text{B.10})$$

Again we use equal weighting ($\omega_p = 1$ for all p). Now, to allocate listings within polygons to OAs, we give each grid equal weighting within the polygon:

$$L_{m gp} = \hat{L}_{mp} * \frac{g \in p}{\sum_{g \in p}} \quad (\text{B.11})$$

To then get the listing allocation to a grid, we sum over all polygons:

$$L_{mg} = \sum_p L_{m gp} \quad (\text{B.12})$$

To get total grid level listings we sum over listings m :

$$\mathcal{L}_g = \sum_m L_{mg} \quad (\text{B.13})$$

C.0.3 Combining searchers and listings to calculate the housing gap

The gap is defined as:

$$\mathcal{G}_g = \mathcal{S}_g - \mathcal{L}_g \quad (\text{B.14})$$

and thus the gap per km^2 is:

$$\mathcal{G}^{\text{km}}_g = \mathcal{G}_g / \text{Area}_g \quad (\text{B.15})$$

We multiply searchers by a scaling factor α for two main reasons. One, it is highly likely that different members within a household are searching simultaneously for their next property, and will use multiple devices (e.g. laptop, phone, tablet). Additionally, a given user may clear their cookies, which, in the absence of a user being logged-in, is how we track individual users. These adjustments do not affect the spatial pattern of results. As such this gives us:

$$\mathcal{G}_g = \alpha \mathcal{S}_g - \mathcal{L}_g \quad (\text{B.16})$$

We use an $\alpha = 0.125$, which in turn allows for two users in a single household, using two devices, who refresh their cookies once during the year. Tests reveal that the rank-rank correlation of $\mathcal{G}^{\text{km}}_g$ is still very high at this level, while reducing the gap to what is seen as a more realistic level. Specifically, the Spearman rank correlation between baseline and scaled $\mathcal{G}^{\text{km}}_g$ is 0.97, indicating that the spatial ordering of areas is largely invariant to the choice of α .