

NBER WORKING PAPER SERIES

HUMAN-ALGORITHM INTERACTIONS:
EVIDENCE FROM ZILLOW.COM

Runshan Fu
Ginger Zhe Jin
Meng Liu

Working Paper 29880
<http://www.nber.org/papers/w29880>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2022

We thank Xiang Hui, Zhentong Lu, Nikhil Malik, and Shunyuan Zhang for constructive comments. None of us have any significant financial interests in Zillow or other housing-related information platforms. The paper only reflects the opinion of the authors, and not the opinion of any organizations. All rights reserved; all errors are ours. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Runshan Fu, Ginger Zhe Jin, and Meng Liu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Human-Algorithm Interactions: Evidence from Zillow.com
Runshan Fu, Ginger Zhe Jin, and Meng Liu
NBER Working Paper No. 29880
March 2022
JEL No. D83,L11,L85,L86

ABSTRACT

Using bi-weekly snapshots of Zillow in three US cities, we document how home sellers and buyers interact with Zillow's Zestimate algorithm during the sales cycle of residential properties. We find that listing and selling outcomes respond significantly to Zestimate, and Zestimate is quickly updated for the focal and comparable houses after a listing or a transaction is completed. The user-Zestimate interactions have mixed implications: on the one hand, listing price depends more on Zestimate if the city does not mandate disclosure of sales information or if the neighborhood is more heterogeneous, suggesting that Zestimate provides valuable information when alternative information is more difficult to obtain; on the other hand, the post-listing update of Zestimate tracks listing price more closely in non-disclosure and heterogeneous neighborhoods, raising the concern that the feedback loop may propagate disturbances in the sales process. However, by leveraging COVID-19 pandemic as a natural experiment, we find no evidence that Zestimate propagates the initial shock from the March-2020 declaration of national emergency, probably because Zestimate has built-in guard rails and users tend to adjust their confidence in Zestimate according to observed market outcomes.

Runshan Fu
Carnegie Mellon University
runshan@cmu.edu

Meng Liu
Washington University in St. Louis
mengl@wustl.edu

Ginger Zhe Jin
University of Maryland
Department of Economics
College Park, MD 20742-7211
and NBER
ginger@umd.edu

1 Introduction

Thanks to artificial intelligence (AI), digital transformation has evolved from accessing and searching digital databases to producing data-driven insights via algorithms. These algorithms have the potential to aggregate decentralized information, improve human decisions, and enhance market efficiency. In the meantime, there are rising concerns that algorithms — often in a black box — may rely on biased inputs, generate erroneous predictions, and mislead human decisions.¹

Underlying the concerns is an implicit assumption of one-directional use: the algorithm makes a recommendation and human users incorporate the algorithmic output in their decisions. In reality, many algorithms, especially those powered by AI, also track how individual users use the algorithm and ingest the usage data for algorithmic updates. For example, general search engines rank relevant information based on the browsing history of other users, which were guided by an earlier version of the same algorithm. E-commerce websites rank products according to historical sales, which depends on how products were presented to previous buyers. Social-media sites recommend what videos to watch, what news to read, and what events to follow based on the likes and dislikes expressed by other users, who made these decisions with the help of the same recommendation algorithm. These two-way interactions could change the quality of algorithmic outputs over time and have profound effects on algorithm-aided human decisions. The goal of this paper is to document the interactive nature of AI-driven algorithms and stimulate a discussion on the potential implications of interactive algorithms.

We use Zillow’s Zestimate algorithm as an example to study interactive algorithms. As defined on Zillow.com, “the Zestimate[®] home valuation model is Zillow’s estimate of a home’s market value. A Zestimate incorporates public, MLS and user-submitted data into Zillow’s proprietary formula, also taking into account home facts, location and market trends.” As of March 10, 2022, Zillow claims that Zestimates have a nationwide median error rate of 1.9% for on-market homes and 6.9% for off-market homes.² Since Zestimates are available for most homes 24/7 but a typical house-selling transaction takes months from listing to sold, we have a rare opportunity to observe how Zestimates influence various stages of house selling and how Zillow updates Zestimates after listing and sold prices become public information.

Based on bi-weekly snapshots of Zillow in Austin, Boston, and Pittsburgh, we find ample

¹See Wall Street Journal 2019/2/13 article “A Crucial Step for Averting AI Disasters”, available at <https://www.wsj.com/articles/a-crucial-step-for-avoiding-ai-disasters-11550069865>.

²See <https://www.zillow.com/z/zestimate/>, last accessed on March 10, 2022.

evidence for user-Zestimate interactions: in the listing stage, the seller’s listing price would go up with an increase in the house’s pre-listing Zestimate. After listing, Zestimate is updated almost immediately, for not only the listed house but also the houses that Zillow deems comparable to the focal house. Towards the end of the sales cycle, transaction outcomes (sold or not, days to pending, days to sold, and sold price) are significantly correlated with its pre-sold Zestimate, and Zestimate is updated once again for the focal and comparable houses once the transaction is completed.

These user-Zestimate interactions are subject to multiple interpretations: First, some unobservable factors may be available to buyers and sellers outside Zillow, but we mistakenly attribute their influence as the impact of Zestimate. To address this concern, we exploit the fact that the Zestimate algorithm incorporates some non-local (i.e., from a different neighborhood) and infrequently shown (on focal home’s Zillow webpage) comparable houses (henceforth “comps”) that are unlikely under the radar of the seller of the focal house (or her real estate agent). Using these non-local and infrequently shown comps to construct an instrument variable (IV) for Zestimate, we continue to find a significant dependence of market outcomes on Zestimate.

In particular, the IV results suggest that listing price would go up 0.66% for every 1% increase in the house’s pre-listing Zestimate. For selling outcomes, we find that a 1% increase in the Zestimate (updated after a property being listed) is, on average, associated with 3.9% fewer days till pending and 0.26% higher sold price, conditional on listing price. We also include a Regression Discontinuity Design (RDD) to focus on cases where the updated Zestimate is within a tiny range around the listing price. Properties that had the Zestimate just above their listing prices have been sold at a faster rate than the properties that had the Zestimate just below the listing prices. These results suggest that, not only do buyers rely on the Zestimate to make decisions, such reliance is but also asymmetric.

As we document the existence of user-Zestimate interaction, the next question is how it influences the housing market. Conceptually, the implications of the user-Zestimate interaction have two sides. On the one hand, it could increase market transparency by providing additional information about property values. The dependence of listing and transaction outcomes on Zestimate may reflect the market’s better use of existing information, because Zestimate makes such information more accessible and more salient to individual buyers and sellers. Our analysis shows that the effects of Zestimate on listing price are highly heterogeneous, where sellers depend more on Zestimate when alternative source of information is more difficult to obtain, such as in a non-disclosure city or in more heterogeneous neighborhoods. This suggests that Zestimate does provide valuable

information to market participants.

On the other hand, the interactions between Zestimate and real-time market information may cause disturbances in any stage of the sales cycle to persist and propagate. For example, Zestimate may contain some prediction error, and some outlier sellers or buyers may have extreme taste relative to market average. If most individual users use Zestimate with confidence, these disturbances can persist and be spread out via user-algorithm interactions. Some data patterns seem to support this concern: for example, we find that Zestimate update is more dependent on listing price in non-disclosure and heterogeneous neighborhoods, reinforcing the co-movement between listing price and Zestimate. Moreover, before a focal property is listed, it only serves as a comparable home for a handful number of properties on Zillow; but once it is listed, it could affect as many as 150 properties via the updated definition of comparable homes. In addition, both sellers and buyers demonstrate asymmetric responses to Zestimate: their reliance on Zestimate is stronger when Zestimate is close to or higher than the property value they may expect from local comparable homes, relative to the cases where Zestimate is lower than the expectation from local comparable homes. All together, these patterns raise the concern that a positive disturbance in the sales cycle (as compared to a negative one) could more easily persist and be further spread out to other houses, contributing to potential housing bubbles.

To test this possibility, we leverage a natural experiment created by COVID-19 to examine whether Zestimate locks in the initial COVID shock from the presidential declaration of national emergency and sets local markets onto different paths. The results suggest no such path dependence in price or Zestimate, and thus do not support the concern of disturbance propagation.

A natural follow-up question is what forces help guard the algorithm from propagating disturbances into a misleading level. Our data suggests two such forces. One is embedded in the Zestimate algorithm: when the confidence interval of a Zestimate is sufficient large, Zillow does not show the Zestimate to the public. The missing value of Zestimate forces the seller and potential buyers to utilize other information. The second force is user adjustment of their confidence in Zestimates. When the Zestimates of same-community comparable houses are far from these houses' actual listing or sold prices, we find that users are less responsive to the Zestimate of the focal house. This suggests that users may adjust their confidence in Zestimate according to how Zestimate aligns with market outcomes. If market outcomes reveal more noise in the Zestimate algorithm, the reduced consumer reaction helps to limit the spread of the disturbances throughout the market.

This paper makes several contributions. First, to our knowledge, we are among the first to provide explicit evidence on the two-way interactions between human users and AI-driven algorithms. While the idea of user-algorithm interactions may seem intuitive, documenting its presence and magnitude is often challenging. To address the challenge, we leverage the long transaction time and granular information updates in the housing market. As detailed in Section 2, we expand a growing literature of data-driven algorithms. Second, we identify the potential positive and negative implications of the human-algorithm interaction, and explore empirical evidence on both sides, including the market’s self-correction mechanisms that could explain the observation that the algorithm did not propagate the initial COVID disturbances. Our findings highlights the importance of considering feedback loops when designing and evaluating algorithms.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3.1 describes the background of Zillow.com and the Zestimate algorithm. Section 3.2 presents our data. Section 4 documents user-Zestimate interactions in various stages of home selling. Section 5 explores the potential implications of user-Zestimate interactions. Section 6 presents two factors that play a role limiting disturbance propagation. A brief conclusion is offered in Section 7.

2 Literature Review

Our work contributes to the growing literature about data-driven algorithms. Most of the literature focuses on evaluating the effect of algorithmic output on the bias and quality of human decision-making: on the positive side, Hoffman et al. [2018] show that managers that reject an algorithm’s recommendation based on job candidates’ job testing performance end up having worse hires on average, suggesting that the algorithm can help to correct human bias or mistake in hiring. Similar results have been shown in the context of bailing [Kleinberg et al., 2018], crowd lending [Fu et al., 2021], and resume screening [Cowgill, 2018]. On the negative side, Ludwig and Mullainathan [2021] summarize the disappointing record of AI in criminal justice, citing evidence for racial and gender bias in algorithms [Angwin et al., 2016, Hamilton, 2019], and lack of better predictability or better performance than human-only decisions [Berk, 2019, Stevenson, 2018, Stevenson and Doleac, 2021]. Such concerns about algorithmic bias and prediction performance are actually shared across many different contexts beyond criminal justice [Lambrecht and Tucker, 2020, 2019, Buolamwini and Gebru, 2018, Chouldechova et al., 2018]. They attribute part of the challenge to human involvement in algorithm building and algorithm use.

There have also been concerns about echo chambers or filter bubbles that the recommendation algorithms behind social medias may limit the information that users encounter or consume online. The empirical evidence is mixed. For example, [Wojcieszak and Mutz \[2009\]](#) find that participants of politics-related online groups, compared with participants of other online groups, were less likely to be exposed to political information that they disagree with. In contrast, [Gentzkow and Shapiro \[2011\]](#) find that “ideological segregation of online news consumption is low in absolute terms” and that there is “no evidence that the Internet is becoming more segregated over time”; [Lambrecht et al. \[2021\]](#) find that YouTube frequently steers users to popular videos that are unrelated to the charity videos that have been previously viewed, demonstrating that firms and institutions are unlikely to benefit from echo chambers. To push it further, we document the *bi-directional* user-Zestimate interactions throughout the cycle of house selling. Similarly, this paper is related to the literature of observational learning [[Banerjee, 1992](#), [Bikhchandani et al., 1992](#), [Smith and Sørensen, 2000](#)] and information cascade [[Anderson and Holt, 1997](#), [Bikhchandani et al., 1998](#), [Çelen and Kariv, 2004](#)], which study how individuals learn from the behavior of others and follow the decisions of the preceding individuals. The algorithm outputs, Zestimate in our context, could reflect multiple sources of information, including but not limited to the behavior of others, and we study how such aggregated and dynamically changing information interacts with participants in housing markets. In addition, we show that such interactions may be an important factor limiting the potential harm of the Zestimate algorithm. These insights could apply to many algorithm-relevant markets beyond housing sales.

The human-algorithm interaction we study in this paper is also related to the literature of systemic risk in financial systems. [Jackson and Pernoud \[2021\]](#) categorizes systemic risk into two types: contagion through network inter-dependencies, and multiple equilibria and self-fulfilling feedback effects. The first refers to the cases where a change in fundamentals move through the financial network [[Rochet and Tirole, 1996](#), [Allen and Gale, 2000](#)], and the second refers to the cases where a shift in beliefs moves the financial system from one equilibrium to another [[Diamond and Dybvig, 1983](#), [Diamond and Rajan, 2011](#), [Bebchuk and Goldstein, 2011](#)]. Similar to the second type of systemic risk, the phenomenon we study also involves shifts in beliefs and changing equilibria. However, we focus more on the external factor that causes belief shift (i.e., an algorithm) and how it interacts with players in the market, rather than multiple equilibria and feedback effects among the players within a system. We also characterize how certain features of an algorithm, such as real-time learning and disclosed individual error rates, influence its interactions with human players.

Another strand of literature focuses on algorithms that guide or determine market prices. For example, Airbnb has provided algorithm-based pricing recommendation to hosts but host acceptance is quite limited [Zhang et al., 2021], and Huang [2021] studies how the platform’s pricing algorithms could mitigate pricing frictions. Brown and MacKay [2021] document the pattern of algorithmic pricing among large retailers and discuss how algorithmic price may reshape market competition. Assad et al. [2020] show that algorithmic pricing may have reduced competition in German retail gasoline, while Calvano et al. [2020] find that pricing algorithms consistently learn to charge supracompetitive prices and sustain the high prices by collusive strategies. While algorithm usage is crucial for market outcomes, these studies focus more on the strategic interactions between users through algorithm adoption. In comparison, we focus more on the explicit interaction between users and the Zestimate algorithm, even if every user is atomic and does not engage in strategic interaction with each other.

This paper is also related to the literature on technology use in the housing market, which focuses on examining the effect of emerging technology on housing market outcomes. For example, Carrillo [2012] studies the impact of online listing and advertising, and Buchak et al. [2020] study “iBuyers”, the online platforms that buy and sell residential properties with advanced pricing technology. There are also a few papers studying the Zestimate algorithm. Yu [2020] shows that Zestimate has a positive effect on sold price and people use Zestimate as a summary of information and rely more on it when it is harder to process the information. Lu [2018, 2019] also document the positive impact of Zestimate on sold price. In this paper, we confirm the impact of Zestimate on a broader set of market outcomes, including listing price, selling price and selling time; more importantly, we focus on user-Zestimate interactions and document how user decisions and the algorithm influence each other in a complete home selling cycle. In addition, our data is better suited to study the impact of Zestimate, as we use Zestimates that were actually shown on webpages, while the previous literature largely relies on the “historical Zestimate” graphs that do not necessarily show the exact Zestimates that buyers and sellers observed in the past (see more details in Section 3.2). Malik [2020] uses a theoretical model to show that the feedback loop in the use of algorithms may create a self-fulfilling prophecy. He also establishes necessary primitives for the feedback loop with Zestimate. Our paper complements the theoretical analysis by providing empirical evidence of the feedback loop, characterizing how users and the Zestimate algorithm interact with each other, and documenting evidence for the positive and negative implications of the feedback loop.

3 Context and Data

3.1 Background

The housing market has seen a few digital disrupters in the past two decades. Before the digital age, sellers and buyers can only access market information through an agent, where agents post listings in the local Multiple Listing Service (MLS hereon) system. With the emergence of online real estate companies, housing market information has become more easily accessible. These companies track the MLS listings and make them online, so that buyers can do their own research instead of solely relying on agents. Figure 1 reports Google trends in web search of four leading companies in this market, namely Zillow, Realtor.com, Trulia (acquired by Zillow in 2014 but remained its website operation), and Redfin, where Zillow has been dominating.


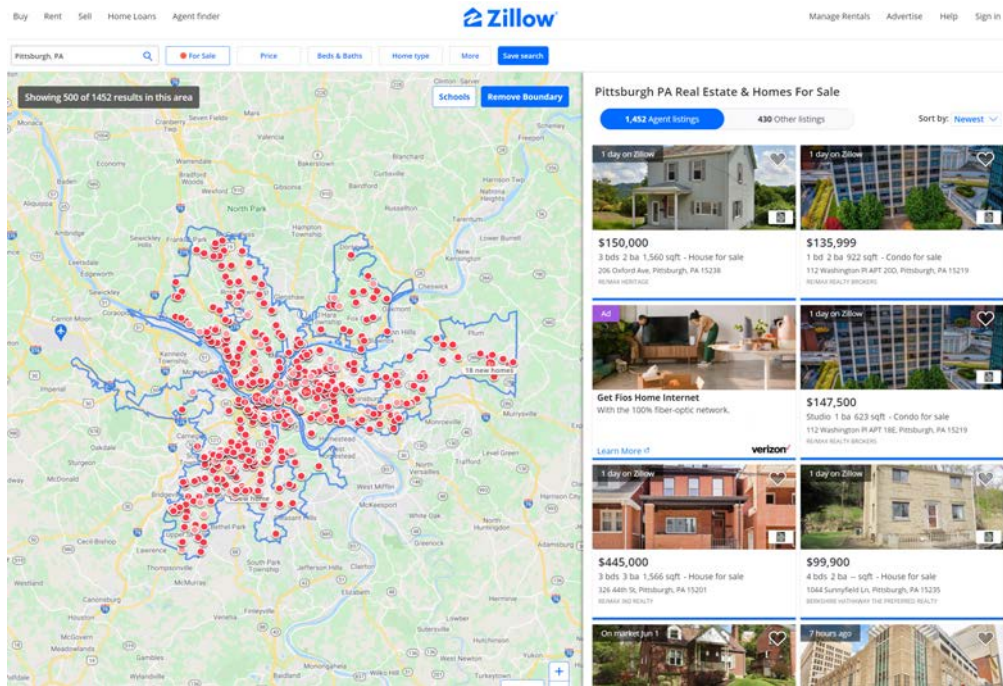
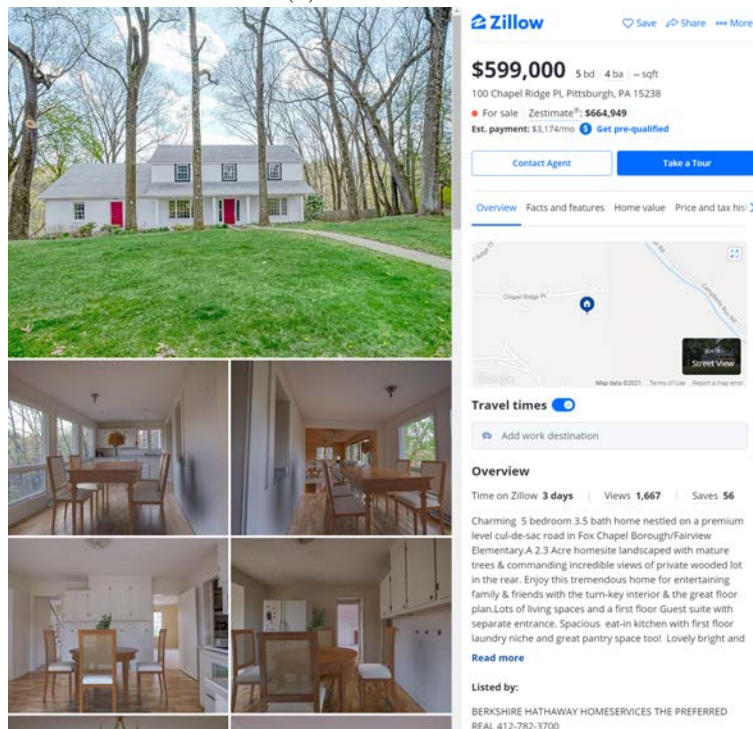


Figure 1: Rise of Digital Real Estate Platforms

Figure 2a shows the typical search interface on Zillow. For example, when users search for listings in Pittsburgh PA, it returns all active for-sale listings scattered on the map (left) and listed on the right, where users can choose the sort order (e.g., by date, by price, etc.). Once users click on a given listing, they will be taken to the listing page (Figure 2b), where photos of the property will be displayed on the left. On the top right of the listing page, key information of the property is displayed, including listing price, address, number of beds, baths, and square footage. It is worth noting that Zestimate is made saliently visible here. As users scroll down, they will see the property's location on the map, the description, the listing agent, detailed property facts, and so on. Zillow makes money by promoting real estate agents, so contact information of advertising agents will be displayed on the property listing page as well.

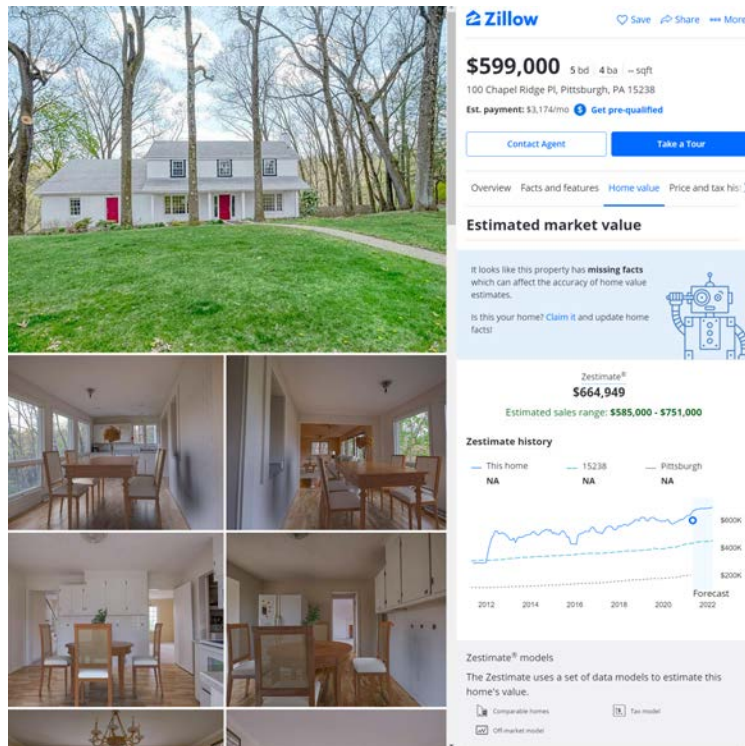


(a) Zillow Search

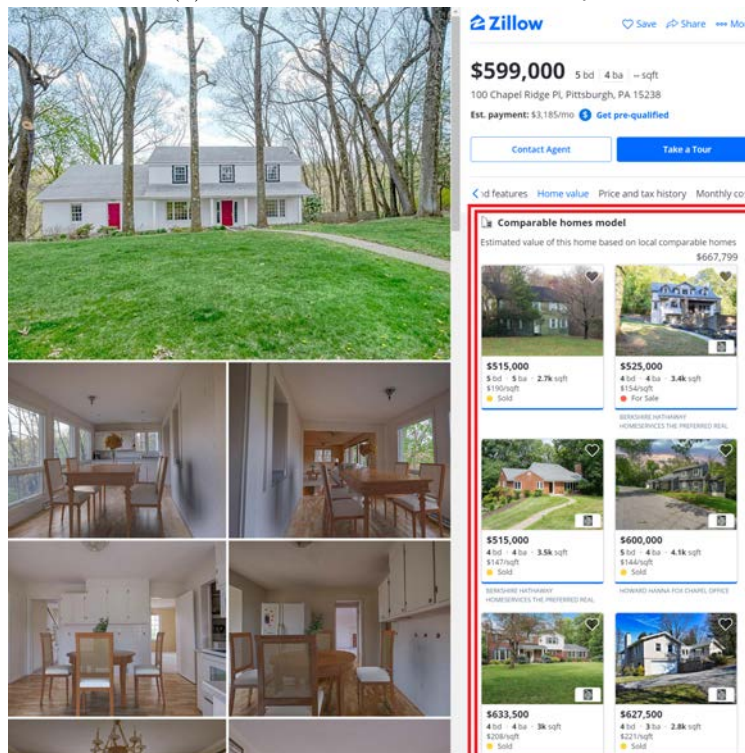


(b) Zillow Listing

Figure 2: Zillow



(a) Zestimate and Zestimate History



(b) Zestimate Comparable Homes Model

Figure 3: Zestimate and the Comparable Homes Model

The listing page has a dedicated section for Zestimate. As illustrated in Figure 3a, it shows the same Zestimate as the number shown at the top of the listing page. Zillow also offers an estimated sales price range, which can be thought of as a confidence interval of the said Zestimate. Besides, Zillow offers “Zestimate history”, where as users move their cursor on the plot they will see Zestimates at given time points, as compared to the average historical Zestimate of the zipcode and of the city. To clarify, the time series are not historical Zestimates *per se*, but Zestimates back-filled by feeding historical data to the most up-to-date algorithm. Therefore, the Zestimate at a given historical time, as shown in the “Zestimate history” graph, can be very different from the actual Zestimate that platform users observed at the same historical time.

As shown in Figure 3a right below the Zestimate history graph, Zillow briefly explains the three different models they use to produce Zestimate. Specifically, they are estimates based on comparable homes, estimates based on property taxes, and estimates generated using the off-market model (it is a model where the algorithm excludes on-market information such as listing price, listing description and days on the market, as if the home were not for sale).

For the comparable homes model, Zillow shows the comps it uses (Figure 3b). Zillow’s algorithm may involve hundreds of comps to build Zestimate of the focal home but usually only five to ten most relevant comps are shown on the listing page³. These comps shown to users (and observed by us) are typically either active listings for sale or properties that were sold in the recent past. Only key information is shown for each comp, including the profile photo, listing/sold price (sometime it is Zestimate for off-market properties), bedrooms, bathrooms, and square footage. Users can click on a given comp and navigate to the page of that comp for more information.

3.2 Data

We scrape data from Zillow.com for approximately 800K residential properties of Austin, Boston, and Pittsburgh.⁴ Compared with Census statistics, our sample has a very high coverage of all residential properties of these three cities. The vast majority of these properties are off-market in our sample period and only a tiny share of properties were listed for sale. For each property, we scrape information on its web page approximately once every two weeks. In each round of scraping, we collect almost all information with the exception of property photos. Our sample starts from

³See <https://www.zillow.com/tech/human-and-machines-home-valuation/>

⁴A property is covered if its city field in the postal address is “Austin”, “Boston” or “Pittsburgh.” By this definition, our sample does not include the properties that locate in the surrounding metropolitan areas but have a different city name in the postal address.

Table 1: Summary Statistics

	Austin			Boston			Pittsburgh		
	N	Mean	S.D.	N	Mean	S.D.	N	Mean	S.D.
listprice (\$1000)	17817	565.43	412.52	7160	887.60	640.06	14430	251.98	161.99
sold	17817	0.64	0.48	7160	0.74	0.44	14430	0.82	0.39
DaysOnMarket	11433	68.68	70.13	5278	93.79	79.78	11780	93.92	83.87
DaysToPending	16075	39.93	71.91	5740	54.15	78.30	13083	52.48	83.90
soldprice (\$1000)	1552	532.62	372.32	5274	799.81	512.56	11723	239.81	404.51
ZestLag1 (\$1000)	17817	511.34	412.10	7160	836.31	592.87	14430	223.99	143.23
ZestLead1 (\$1000)	15834	540.48	404.52	7066	880.26	620.14	13615	248.68	153.47
age	17817	30.79	21.94	7160	93.43	43.26	14430	75.84	30.43
bathrooms	17817	2.67	1.10	7160	1.97	1.05	14430	2.16	0.94
bedrooms	17817	3.26	1.03	7160	2.97	2.00	14430	3.08	0.99
floorsize	17817	2132.43	1084.66	7160	1622.90	2262.68	14430	1672.48	727.46
condo	17817	0.11	0.31	7160	0.60	0.49	14430	0.06	0.24

March 2019 and ends in March 2021. We have a fairly balanced panel where each property is tracked for about 60 rounds in this two-year period.

A unique feature of our data is that we observe various time-varying information about a property, such as Zestimate and comparable homes that are used to construct Zestimate. It is crucial for our purpose because this is the information that market participants observe and act upon. We are able to track the status of on-market properties with outcomes such as listing price, price change, days on the market, sold price, listing agent, etc. For all properties, we observe their past sales history and tax records that usually go back in years or even decades. Lastly, we observe property-specific characteristics such as address and home facts.

Our raw data of on-market properties includes 44,327 properties listed between March 25th, 2019 and March 7th, 2021. We apply the following filters to construct our main sample: (1) We remove listings for which we do not observe key home facts including age, bedrooms, bathrooms, and floor size. This removes 3,445 properties from the sample. (2) We remove properties with listing price less than \$10,000. To prevent extreme values from driving our results, we drop outliers whose listing price is below 1st percentile or above 99th percentile, separately for each city. This reduces the sample size to 39,631. (3) We delete 193 properties for which we do not observe Zestimate right before listing. (4) Given that various later analyses require neighborhood fixed effects, we remove 31 neighborhoods with only one listed property in the entire sample period. After all the above-mentioned procedures, our main sample contains 39,407 listed properties, where 17,817 are in Austin, 7,160 are in Boston, and 14,430 are in Pittsburgh.

Table 1 reports the summary statistics of the key variables of listed properties, separately

for Austin, Boston, and Pittsburgh. Residential properties are on average the most expensive in Boston and least expensive in Pittsburgh. We observe that listings in Pittsburgh have the highest probability of sales (82%), while only 64% of Austin listings were sold. Days on market are defined as the days from listing date to closing date. It is approximately the same for Boston and Pittsburgh properties (94 days on average) but it is considerably shorter for Austin homes (69 days). A similar pattern holds for days to pending, which is the days from listing date to the date when the property enters the pending stage. Due to construction, it is shorter than days on market and is plausibly a better measure of time on market, given that the time between pending and final closing is mainly for home inspection, appraisal, repair, loan approval, etc, which do not necessarily reflect the popularity of the property. Because Austin is a non-disclosure city where home sales price are not required to be made public, we only observe sales price for 1,552 out of 11,433 sales. However, we have almost perfect coverage of sales price for Boston and Pittsburgh (5,274 out of 5,278 and 11,723 out of 11,780, respectively). Zestimate before the round of listing, *ZestLag1*, is lower than listing price on average, for all three cities. Zestimate after the listing round, *ZestLead1*, is greater than *ZestLag1* but lower than listing price on average. Note that for some properties Zestimate will disappear after the listing round and we will leverage the time-varying coverage in later analyses. Finally, we summarize the key housing characteristics, namely age, bathrooms, bedrooms, floor size (in square footage), and whether the property is a condo or a house⁵. The statistics suggest that Austin homes are on average younger and bigger than properties in Boston and Pittsburgh. The majority of Boston properties in our sample are condos while Austin and Pittsburgh homes are predominately single family houses, multi-family houses, or townhouses.

4 Characterization of Zestimate’s Role in the Housing Market

To fix ideas, let us consider the following process by which the housing market unfolds for a typical transaction. Figure 4 illustrates the sequence of events: (1) The property is off-market and has a Zestimate. (2) Then at some point, the seller decides to sell the home by consulting the Zestimate and decides on a listing price. (3) Immediately after the home is listed, Zestimate will get updated based on the listing price, as well as the updated home facts provided in the listing. (4) The demand side comes in and makes buying decisions after seeing the listing price and the updated Zestimate (as well as touring the home and placing a bid). (5) Lastly, Zestimate updates once again after the

⁵We categorize all apartments and condos as condos, and group all single family houses, townhouses, and multi-family houses as houses.

transaction is done and transaction outcomes are observed by Zillow.




Figure 4: Sequence of Events

Given this, we first study whether sellers use Zestimate to price their homes. Note that the Zestimate that can affect listing price is the Zestimate prior to the listing. Then, we study how Zestimate updates based on listing price. After that, we aim to understand if home buyers are affected by Zestimate which is already updated on listing price. Lastly, we study how Zestimate continues to update with information on market outcomes of the property. Our goal is to characterize how sellers, buyers, and Zestimate interact with each other throughout the process.

4.1 Seller Response to Zestimate

In this section, we aim to understand the causal effect of Zestimate on the seller’s pricing decision. We specify the logged listing price as a linear function of the logged Zestimate in the round before listing. We also control for relevant covariates that may affect the listing price. Specifically, we estimate the following regression equation:

$$\ln(ListingPrice)_i = \beta \ln(ZestLag1)_i + X_i \Theta + \epsilon_i, \quad (1)$$

where *ZestLag1* is the Zestimate in the round prior to listing. We use X_i to denote the set of relevant property characteristics including # of bedrooms, # of bathrooms, square footage, age of the property, whether the property is condo or not.⁶ We also include in X_i two measures of local comparable homes value, namely the average recent sales price of comps in the same neighborhood and the average recent sales price of comps within one mile of the focal property. These variables are controlled for in the regression because they could be a direct input for the seller’s pricing

⁶In the regression, we control for the logarithm of square footage because it is a skewed measure. Age is also skewed, but a tiny fraction of properties have 0 age (i.e., new construction), so we use $\log(\text{Age}+1)$ instead.

decision if the seller or the seller’s agent finds it from the MLS or other non-Zillow sources. We further control for neighborhood-month fixed effects to account for unobserved market conditions. A remaining identification challenge is that *ZestLag1* may still capture some non-Zillow information that the seller observes but we do not, and thus an OLS regression of Equation 1 will lead to a biased estimate of β . To identify the causal effect, we instrument *ZestLag1* with the average Zestimate of non-local *and* infrequently shown comparable homes of the focal property. Specifically, the comps used for IV construction need to satisfy two conditions: (1) “*non-local*” — they are not from the neighborhood of the focal property; (2) “*infrequently shown*” — they are only shown on Zillow as comps of the focal property for no more than two times out of the six snapshots we scrape. We now explain the validity of this instrument.

Listing agents usually prepare a comparative market analysis report to help sellers price their homes, by aggregating past sales information from their local MLS. This report typically consists of three to five recently sold properties that are similar in key aspects, such as location and size. While we do not find a one-size-fits-all rule in terms of the location of the comparable homes agents select, there seems to be a consensus that comparable homes should be ideally in the same neighborhood or school district, and they should be as close to the focal property as possible (for example, within a mile).⁷

For a given focal property, Zillow usually shows five to ten comps on the listing page at a given time. This list of comps can change over time, possibly because Zillow gives new listings and newly sold properties more weight as comps. In Figure 5, we plot the distribution of the number of unique comps shown on Zillow for a focal property in the six rounds before listing. For an average property, about twenty comps are used and shown by Zillow in the approximately 3-month period before listing.

In addition, the comps vary in their distance to the focal property and the frequency they were shown on the focal property’s listing page. We find that Zillow uses both local and non-local comps, where we define local as the same neighborhood as the focal property. In fact, we find that a vast majority of properties in our sample do have non-local comps, although local comps are more common. The variation in distance and localness of comps suggest that Zillow’s comps algorithm may take a more “global” perspective than that of an average agent, to the extent that agents are

⁷See <https://www.realtor.com/advice/sell/comparative-market-analysis-explained/>, <https://www.zillow.com/home-buying-guide/comparative-market-analysis/>, <https://www.rocketmortgage.com/learn/comparative-market-analysis>, <https://www.investopedia.com/terms/c/comparative-market-analysis.asp>.

more likely to restrict comps to be within the same neighborhood as the focal property.

Furthermore, we observe that some comps tend to stay as comps of the focal property over time while others only appear briefly. In Figure 6, we provide a breakdown of comps by their frequency of appearance in our data. For example, 16% represents comps that only appear once in the 6-round period and 83% represents comps that were observed five times out of six. We see that there are many comps that only appear once or twice, for both local and non-local comps.




Figure 5: The Number of Unique Comparable Homes

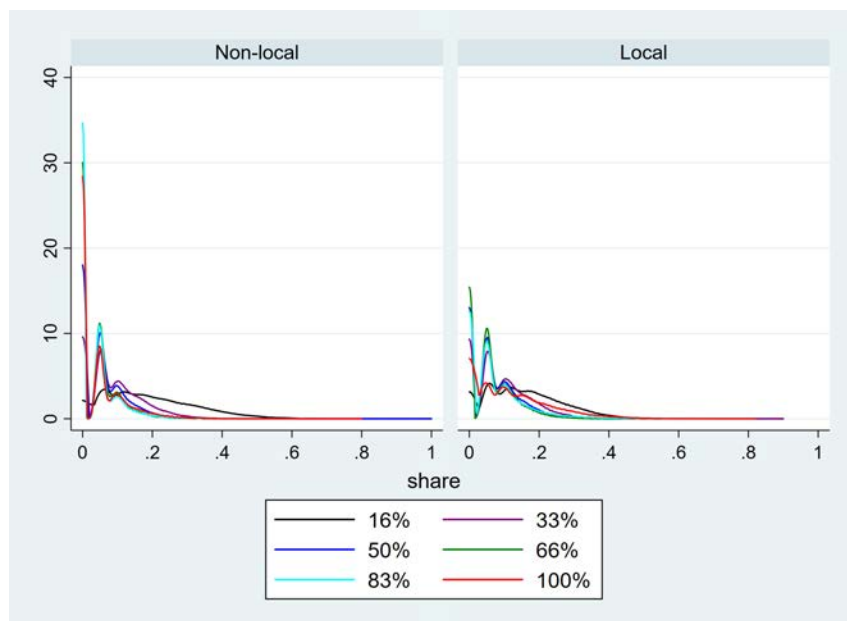


Figure 6: Localness and Frequency of Comparable Homes

Given these patterns, we use the average Zestimate of non-local and infrequent comps as an instrument, under the assumption that they are outside of the direct information set of the seller

or the seller’s agent when pricing the home. We believe that the instrument satisfies the exclusion restriction for the following reasons: First, Zillow’s algorithm takes a more global perspective at comps whereas seller and their agents usually focus on local comps. To the extent that location is one of the most important factors in selecting comps, homes in a different neighborhood or zip code are not likely to be considered by humans as valid comps. Machine learning algorithms, however, may not have clear-cut restrictions like humans do. Secondly, we need to restrict the comps to be the ones infrequently shown on Zillow, because sellers and their agents may consult Zillow before making pricing decisions, which leads to a direct correlation between non-local comps and focal property’s price. If the seller observes that a non-local home is listed as a comp repeatedly, she is likely to visit the comp’s Zillow page to get more detail and use that information to guide her pricing. Therefore, we use infrequently-shown non-local homes under the assumption that sellers have very limited exposure to these comps.⁸

We find a positive effect of Zestimate on listing price. The regression results are reported in Table 2. In Column (1), we report the IV result of Equation 1 that shows an estimated elasticity of 0.66, meaning that a one percent increase in lagged Zestimate leads to a 0.66% increase in listing price, everything else held the same. The positive coefficients of *logRecentSoldPriceComp1m* (logged average recent sales price of comps within one mile) and *logRecentSoldPriceLComp* (logged average recent sales price of comps in the same neighborhood) suggest that local comps matter but they do not affect the listing price as strongly as the property’s own Zestimate does. In addition, we observe that listing price is positively correlated with features such as bedrooms, bathrooms, and floor size, and negatively correlated with age and the condo dummy.

In Table 2 Column (2) - (4), we report IV results on subsamples by city. The effect of Zestimate on listing price is positive and strong across cities. The effect size varies and is the largest in Austin and smallest in Pittsburgh. We then split the sample into houses and condos (Columns 5 and 6), where we find that the effect is present for both housing types. We also report OLS estimation results of Equation 1 in Appendix Table A1, using the same samples for the IV regressions. Across all specifications except the one on condos, we find that OLS estimates are smaller than their IV counterparts, suggesting a downward bias from unobservables. In appendix Table A2, we report the

⁸An alternative way to define the IV is to use the distance between comps and the focal property. We believe localness is a more reasonable choice because (1) agents typically select comps based on neighborhoods or zipcodes instead of exact distance, (2) the same cutoff of distance could mean very different degrees of comparability in different cities, and this is likely one reason why there is no commonly agreed-upon guideline in terms of using distance to define comps. Nonetheless, we use a 3-mile cutoff as an alternative way to define non-local comps in a robustness check, where we observe qualitatively similar results.

Table 2: Zestimate’s Effect on Listing Price

D.V.: logged(Listing Price)	(1) All	(2) Austin	(3) Boston	(4) Pittsburgh	(5) Houses	(6) Condos
logZestLag1	0.659*** (0.018)	0.878*** (0.026)	0.639*** (0.047)	0.545*** (0.028)	0.649*** (0.021)	0.525*** (0.043)
logRecentSoldPriceComp1m	0.046*** (0.014)	0.028 (0.018)	0.058*** (0.022)	0.083*** (0.026)	0.038** (0.017)	0.119*** (0.024)
logRecentSoldPriceLComp	0.042*** (0.015)	-0.019 (0.021)	0.048* (0.028)	0.052* (0.027)	0.041** (0.018)	0.048* (0.025)
bedrooms	0.015*** (0.002)	0.006** (0.003)	0.019*** (0.003)	0.018*** (0.003)	0.018*** (0.002)	0.016*** (0.006)
bathrooms	0.055*** (0.003)	0.012*** (0.003)	0.040*** (0.004)	0.093*** (0.005)	0.059*** (0.003)	0.046*** (0.007)
logFloorsize	0.087*** (0.007)	0.090*** (0.011)	0.130*** (0.016)	0.062*** (0.011)	0.077*** (0.008)	0.173*** (0.018)
log(Age+1)	-0.016*** (0.002)	-0.031*** (0.003)	-0.006 (0.004)	-0.018*** (0.006)	-0.019*** (0.003)	-0.008** (0.004)
condo	-0.061*** (0.007)	-0.077*** (0.009)	-0.004 (0.009)	-0.043*** (0.016)		
Observations	24,622	9,000	4,583	11,039	20,519	3,418
R-squared	0.691	0.778	0.821	0.605	0.605	0.822
nbhd-by-month FE	YES	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

first stage results for the IV regressions, which confirm a strong correlation between the instrument and the endogenous variable for all specifications.

We use two alternative measures of the IV to provide robustness. First, we use the average sales price, instead of average Zestimate, of those non-local and infrequent comps that were sold recently. Table A3 reports the regression results that are qualitatively similar to the main results reported in Table 2. Secondly, we use distance to define non-local properties as an alternative measure. Specifically, we use the average Zestimate of infrequently shown comps that are at least 3 miles away from the focal property. As Table A4 shows, the estimated effect of Zestimate on listing price is qualitatively similar to the results in Table 2.

4.2 Zestimate Update after Listing

Zestimate is an algorithmic estimate that is dynamically changing over time based on real-time information that Zillow has access to. We observe that usually the biggest update in Zestimate for a given property takes place immediately after the property goes on the market. This is the time when seller’s agent provides a major update on the property in the local MLS.

Figure 7 plots the average of Zestimate-listing-price ratio for six rounds before the listing and six rounds after, where we benchmark all listed properties against their own listing round (denoted by 0). As shown in this graph, for all three cities, Zestimate is lower than listing price in both the rounds before listing and the rounds after. In addition, we find a clear discontinuity in the time series — it appears that immediately after listing price is posted, Zestimate is updated. In particular, the updated Zestimate increases significantly as indicated by the rising Zestimate-price ratio. Lastly, we find that the confidence interval of Zestimate shrinks after listing, suggesting that the updated Zestimate tends to follow listing price more closely than does the lagged Zestimate. In addition, there seems to be a mild increase in Zestimate-price ratio in Round 0, the listing round, and this is an artifact of how the sample is collected. Specifically, it takes approximately two weeks for our scraping algorithm to complete one round of queries. For the round where a given property goes on market, its Zestimate query of that round may be scraped before or after the specific listing date. Therefore, Round 0 contains both cases and should be considered as a mixture of Round -1 and Round 1.

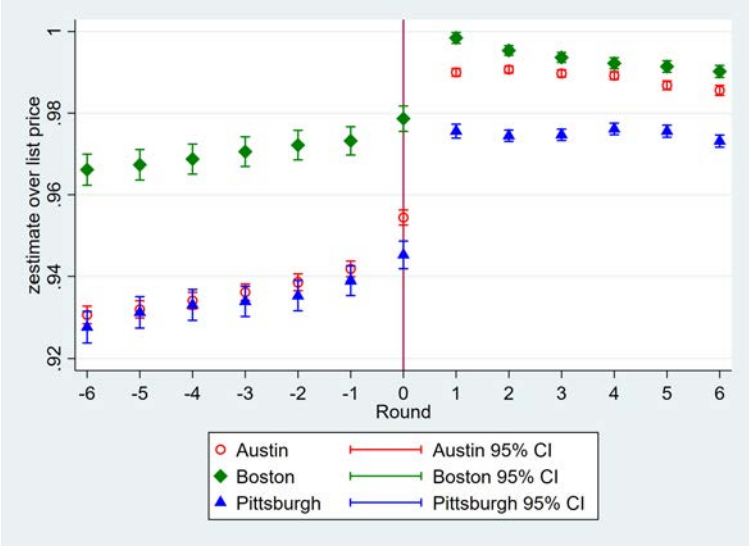


Figure 7: Zestimate Update After Listing

An alternative explanation for the patterns in Figure 7 is that Zillow and the seller may observe new market conditions simultaneously and the Zestimate update simply reflects Zillow’s independent observation of these market changes. Intuitively, if we were able to observe Zestimate on the daily basis, we could use the timing of Zestimate update to address this problem, assuming that the aforementioned market conditions change somewhat continuously. Given that our data crawling is bi-weekly, we instead focus on a subsample that consists of (a) the crawling was done

Table 3: Zestimate Update on Listing price

D.V.: $\log(\text{ZestLead1})$	(1)
$\log\text{listprice}$	0.574*** (0.011)
$\log\text{listprice} \times \text{After}$	0.025*** (0.007)
After	-0.294*** (0.094)
bedrooms	-0.005 (0.003)
bathrooms	0.016*** (0.004)
$\log\text{Floorsize}$	0.255*** (0.011)
$\log(\text{Age}+1)$	-0.017*** (0.004)
condo	-0.027** (0.011)
Observations	5,367
R-squared	0.937
Neighborhood FE	YES
Year-month FE	YES

Notes: Standard errors in parentheses.
 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

one day after the property was listed and (b) the crawling was done one day before the property was listed. We then compare how the correlation between the updated Zestimate and listing price differs between these two cases. If the unobserved factor affects the updated Zestimate and listing price simultaneously, there should be little difference between correlations in (a) and (b). However, if Zestimate is updated based on listing price, then we should expect a stronger correlation in (a) than in (b). As shown in Table 3, Zestimate has a larger correlation with price for cases in (a) than for cases in (b), as evidenced by the positive coefficient estimate of $\log\text{listprice}$ times *After*.

Why does Zestimate so closely follow the listing price? One obvious reason is that listing price contains valuable private information about the property, which makes listing price possibly the strongest signal of the property value that Zillow can leverage for an on-market property. As we discuss in Section 5, the tight co-movement of Zestimate and listing price may give rise to disturbance propagation. For example, a random noise in Zestimate before listing may affect listing price, which is later on propagated to the updated Zestimate, causing the final sales outcomes to deviate from what they should be. Section 6 will further discuss the factors that may limit the scope of disturbance propagation and help the market to self correct.

4.3 Zestimate and Other Market Outcomes

We now evaluate the relationship between Zestimate and other market outcomes besides listing price. Specifically, we focus on the dummy variable of whether property is sold or not, logged days from listing to pending, logged days from listing to closing⁹, and logged sold price (where sold price is observable). Note that once the property is listed, Zestimate gets updated based on listing price, so we use the updated Zestimate (*ZestLead1*) in these analyses to represent the Zestimate that is observed by the demand side post listing. We estimate the following regression equation:

$$Y_i = \beta_1 \ln(\text{ZestLead1})_i + \beta_2 \ln(\text{listprice})_i + X_i \Theta + \epsilon_i, \quad (2)$$

where *ZestLead1* is the Zestimate in the round after listing and *listprice* is the listing price of the property. We use X_i to denote the set of relevant property characteristics including # of bedrooms, # of bathrooms, square footage, age of the property, whether the property is condo or not. We also include in X_i two measures of local comps value, namely the average recent sales price of comps in the same neighborhood and the average recent sales price of comps within one mile of the focal property. These variables are controlled for because they may capture the part of property value that is observable to market participants but unobservable to us, for example community attributes common among local comparable homes. We further control for neighborhood-month fixed effects.

Given that our primary focus is on the causal effect of *ZestLead1* on market outcomes, we choose to instrument *ZestLead1* while taking *listprice* in the regression as given.¹⁰ We use the same instrument as in Section 4.2, i.e., average Zestimate of infrequent non-local comparable homes prior to the listing of the focal home.

Table 4 reports the IV regression results. Estimates in Column (1) suggest that the sales probability is positively affected by Zestimate but the effect is weak. Column (2) shows that days to pending is negatively affected by Zestimate. In particular, a 1% increase in Zestimate leads to a 3.9% decline in days till pending. A similar but smaller effect is observed for days to closing (Column 3), where a 1% increase in Zestimate means a 1.9% shorter days on market. Lastly, we observe a positive effect of Zestimate on the sold price (Column 4), suggesting that a 1% increase

⁹A tiny fraction of active properties have zero days to pending or zero days to closing, so we define logged days to pending as $\log(\text{days to pending}+1)$ and logged days to closing as $\log(\text{days to closing}+1)$.

¹⁰If our purpose was to understand the causal effects of both *listprice* and *ZestLead1*, then we would need two separate instruments. When we take listing price as given, the coefficient of *ZestLead1* identifies how the market responds to a random change in *ZestLead1* conditional on whatever information that is already embodied in the listing price.

Table 4: Zestimate and Other Market Outcomes

	(1) sold	(2) log(DaysToPending+1)	(3) log(DaysOnMarket+1)	(4) logsoldprice
logZestLead1	0.227 (0.152)	-3.904*** (0.624)	-1.893*** (0.334)	0.260*** (0.076)
loglistprice	-0.269** (0.121)	3.824*** (0.502)	1.819*** (0.265)	0.756*** (0.060)
logRecentSoldPriceComp1m	0.001 (0.024)	0.063 (0.091)	-0.011 (0.049)	-0.008 (0.011)
logRecentSoldPriceLComp	0.051* (0.028)	-0.026 (0.107)	0.045 (0.059)	0.005 (0.013)
bedrooms	-0.010*** (0.003)	-0.019* (0.012)	0.004 (0.006)	0.003** (0.001)
bathrooms	-0.002 (0.004)	0.046*** (0.017)	0.023** (0.009)	0.000 (0.002)
logFloorsize	-0.021* (0.013)	0.315*** (0.047)	0.214*** (0.028)	-0.034*** (0.006)
log(Age+1)	0.001 (0.004)	-0.022 (0.017)	-0.002 (0.010)	-0.001 (0.002)
condo	-0.078*** (0.012)	0.175*** (0.044)	0.101*** (0.024)	-0.001 (0.005)
Observations	22,772	20,149	16,711	11,982
R-squared	0.004	-0.038	-0.041	0.904
nbhd-by-month FE	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

in Zestimate translates to 0.26% higher sold price. Furthermore, the correlation between listing price and market outcomes are intuitive: listing price is negatively correlated with sales probability and positively correlated with days till pending, days on market, and sold price. Overall, these empirical patterns suggest that a higher Zestimate is a strong cause of shorter time on market and a higher sold price.

We now turn to a Regression Discontinuity Design (RDD) to provide further evidence for the causal effect of *ZestLead1* on market outcomes. Consider an example from the perspective of home buyers, where hypothetically two houses are identical except for *ZestLead1*. In particular, *ZestLead1* is slightly lower than listing price in one case and slightly above in the other. If this crossover is random at the margin, then any differences in market outcomes can be interpreted by home buyers perceive these *ZestLead1* as different quality signals. In particular, Zillow may give buyers more confidence for the just-above cases than the just below cases.

To mimic this hypothetical ideal in econometrics, we focus on cases where *ZestLead1* is within a tiny range around listing price. We find that the distribution of *ZestLead1* around listing price is

not uniform, with more cases where $ZestLead1$ is equal or greater than listing price than below it. Given this, We limit the RDD sample to the observations where $ZestLead1$ is above listing price up to 0.1% or below listing price no more than 1%. This means that for a home listed at \$400,000, $ZestLead1$ is at most \$400 above or \$4000 below. We verify that the just-above and just-below listings are balanced on a number of observables including listing price, lagged Zestimate, and home facts. The balance check result is reported in Appendix Table A5.

Table 5: Zestimate and Other Market Outcomes — RDD

	(1)	(2)	(3)	(4)
	sold	log(DaysToPending+1)	log(DaysOnMarket+1)	logsoldprice
Just above	-0.004	-0.174***	-0.086**	0.003
	(0.021)	(0.067)	(0.034)	(0.019)
bedrooms	-0.017	-0.017	-0.003	-0.005
	(0.011)	(0.036)	(0.019)	(0.011)
bathrooms	-0.024	-0.058	0.002	0.120***
	(0.014)	(0.051)	(0.026)	(0.015)
logFloorsize	0.050	0.323***	0.154**	0.350***
	(0.037)	(0.123)	(0.063)	(0.039)
log(Age+1)	-0.021	0.009	0.011	-0.002
	(0.014)	(0.046)	(0.024)	(0.019)
Observations	1,873	1,764	1,476	816
R-squared	0.080	0.054	0.116	0.850
city FE	YES	YES	YES	YES
Year-month FE	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 5 shows results by regressing market outcomes on the “Just above” dummy, controlling for various home facts. Given the sample size, we are no longer able to control for granular fixed effects at the neighborhood-month level. Instead, we allow for two separate fixed effects at the city level and year-month level. In Column (1), we observe no statistically significant effect on the probability of sale. In Column (2) and (3), We do find that the just-above cases can be sold at a faster rate than just-below cases. In particular, the just-above cases need 17.4% less time to go to pending and 8.6% less time to get closed. Lastly, as Column (4) suggests, we do not observe a strong difference in sold price. Note that the sample is reduced considerably when the dependent variable is sold price because (1) not all listings end up being sold, and (2) Austin properties need to be removed from the sample because Austin is a non-disclosure city and most properties choose to not disclose their sold price. Overall, we find evidence that the demand side responds to just-above cases more favorably than just-below cases. These findings imply that the demand side may

have incomplete information and therefore need to leverage Zestimate as an information source in addition to other signals in the market.

We can think of two explanations for the asymmetry. First, when Zestimate is above listing price, little other information is available to justify why and how much the Zestimate is higher than listing price, so buyers count more on the Zestimate as a guidance. In contrast, when the Zestimate is below listing price, the seller has strong incentives to explain why he/she chooses to set a high listing price, which helps to divert buyer attention from Zestimate. Alternatively, most buyers may have some investment motive in house shopping, which could generate an upward confirmation bias. In particular, when the Zestimate is higher than listing price, it paints a rosy picture about the property, minimizing the buyer’s incentive to investigate further about the true property value; but when the Zestimate is below listing price, buyers are compelled to do more homework to justify their buying behavior.

4.4 Zestimate Update on Market Outcomes

Lastly, we provide evidence that Zestimate updates again once the market clears for the focal property. Figure 8 plots the average Zestimate-sold-price ratio in the six rounds before and six rounds after closing. We see that Zestimate converges to the sold price as the property goes towards closing. In addition, Zestimate’s deviation from the sold price shrinks after the home is sold, as reflected by the smaller confidence interval. We take this as evidence that Zestimate tends to update based on the sold price.

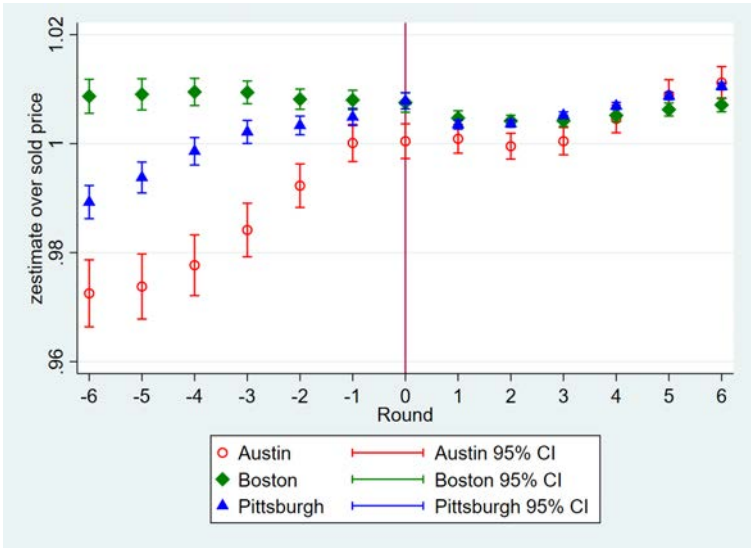


Figure 8: Zestimate Updates on Sold Price

In summary, we document a clear, sequential pattern of interactions between Zestimate and user-driven market outcomes. Specifically, the initial Zestimate affects the listing price, which Zestimate uses as an input to update its next version; then the demand side takes into account the updated Zestimate when making the purchase decision, generating a sold price that feeds another round of Zestimate update.

5 Marketwide Implications

As we present evidence on the interactions between Zestimate and human users in the previous section, we now move to discuss and analyze the potential implications that such user-Zestimate interactions may bring to the long-run efficiency of the housing market. Since Zestimate is a salient aggregation of public information, the benefit of Zestimate to home sellers and buyers depends on the availability and efficiency of alternative sources of information. That is, when it is more costly for sellers and buyers to acquire the information by other means, we expect them to demonstrate a greater dependence on Zestimate.

Such dependence could have mixed effects on the overall market efficiency: on the one hand, it may bring attention to important but otherwise obscure information, so that sellers and buyers can overcome the existing information frictions and make a more informed decision. This could help housing price to converge to the true property value. On the other hand, if Zestimate incorporates noise but some sellers or buyers use it with ultimate confidence, they may be misled by the Zestimate. Their twisted decisions feed into the next update of the Zestimate algorithm, which in turn affects the comps, the comps of comps, etc. In theory, these chain reactions could propagate any disturbance in any stage of the sales process, which has a potential to stimulate bubbles and hamper the long-run efficiency of the housing market.

It is challenging to quantify these positive and negative implications, partly because alternative information is hard to observe, partly because the implications are dynamic by definition but many other factors may affect the dynamics simultaneously. Given this challenge, we look for clues that may point to the presence of these implications and try to identify factors that may affect the size of the implications, even if we cannot pin down their absolute magnitudes.

5.1 Value Added of Zestimate

We rely on three variables to proxy the costs of finding alternative sources of information. The first is regulatory difference in sales price disclosure. In the US, a dozen states do not require sellers to disclose the historical sold price of residential properties. Sellers in these states can choose to disclose their sales price but they are not required to and in fact many choose not to. In these non-disclosure states, Zillow does not have access to all sales prices, nor does the general public. Typically, one must go through a licensed realtor to get information on non-disclosed sales price. The common reason given by the supporters of the non-disclosure law is to protect home owner’s privacy, under the assumption that people do not like others to know how much they paid or how much they sold their home for. Our data includes one non-disclosure city, Austin, and two disclosure cities, Boston and Pittsburgh.

To examine how the effect of Zestimate on listing price varies by city, we estimate Equation 1 by including a interaction term between the dummy for Austin properties and $\log ZestLag1$. Estimates reported in Column (1) of Table 6 suggest that sellers seem to depend more on Zestimate in Austin than in two other cities. Specifically, the elasticity is 14.7% (0.095 above 0.646) higher in Austin than the other cities. Zestimate in Austin is plausibly less accurate than in Boston and Pittsburgh because Zillow has less and lower-quality information on sales price. However, the higher dependence of listing price on Zestimate implies that it is even more difficult for Austin sellers to do their own research with scarce data available to them, and therefore they need to rely on Zestimate more than sellers in other cities do.

One may argue that Austin may differ from Boston and Pittsburgh in many dimensions, so the above results are not necessarily driven by the difference in sales price disclosure. To address this concern, our second measure of information cost exploits within-city variations. In particular, sellers often form belief of property values by looking at nearby houses, and this process is more difficult if the neighborhood is more heterogeneous in housing characteristics. Empirically, for each city, we split neighborhoods into low, medium, and high-heterogeneity, where neighborhood heterogeneity is measured by the within-neighborhood variation in Zestimate, beds, baths, floor size, and house age. We restrict the sample to Boston and Pittsburgh only, so that neighborhood heterogeneity is not driven by disclosure law as is the case in Column (1). Table 6 Column (2) shows that listing price of high-heterogeneity neighborhoods has a higher dependence on Zestimate than that of low- and medium-heterogeneity neighborhoods. This is consistent with the hypothesis

Table 6: Zestimate’s Effect on Listing Price: Heterogeneous Effects

D.V.: log(List Price)	(1) Austin vs. other cities	(2) Nbhd. heterogeneity	(3) ZestLag1 confidence interval
logZestLag1	0.646*** (0.018)	0.532*** (0.026)	0.577*** (0.025)
logZestLag1×Austin	0.095*** (0.013)		
logZestLag1×heteroHi		0.070*** (0.016)	
logZestLag1×heteroLo		0.014 (0.019)	
logZestLag1×CI.Hi			0.022*** (0.006)
logRecentSoldPriceComp1m	0.045*** (0.014)	0.071*** (0.018)	0.053*** (0.019)
logRecentSoldPriceLComp	0.039*** (0.015)	0.063*** (0.020)	0.062*** (0.020)
bedrooms	0.017*** (0.002)	0.019*** (0.002)	0.018*** (0.002)
bathrooms	0.054*** (0.003)	0.074*** (0.003)	0.073*** (0.003)
logFloorsize	0.082*** (0.007)	0.083*** (0.009)	0.087*** (0.009)
log(Age+1)	-0.014*** (0.003)	-0.009** (0.004)	-0.010*** (0.004)
condo	-0.054*** (0.007)	-0.022** (0.009)	-0.028*** (0.009)
CI.Hi			-0.287*** (0.082)
Observations	24,622	15,622	15,224
R-squared	0.692	0.672	0.676
nbhd-by-month FE	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

that information acquisition is more difficult for sellers in more heterogeneous neighborhoods, which warrants a higher dependence on Zestimate.

The third measure of information cost is the confidence interval (CI) of Zestimate, which by definition, reflects how confident Zillow is when it uses public information to create Zestimate. In particular, for each *ZestLag1*, we observe the associated Zestimate CI in percentage terms. For example, a 5% upper bound and a 5% lower bound around the Zestimate together indicate a Zestimate CI of 10. This measure is positive for all properties by construction, and it is skewed with a mean of 19.07, a median of 16, and a standard deviation of 12.43. Based on the median, we group properties into high-CI and low-CI and use the dummy “CI.Hi” as the corresponding

indicator. Evidence shown in Column (3) of Table 6 reports a higher effect of Zestimate on listing price for properties with higher CI.

In summary, we identify strong heterogeneous effects of Zestimate on listing price, which indicate that sellers depend more on Zestimate when the underlying information is more costly to collect from alternative sources. These findings suggest that Zestimate may improve market efficiency by making harder-to-find information more salient and available to housing market participants.

5.2 Possible disturbance propagation

While Zestimate may improve market transparency, the way it closely tracks real-time market information such as listing and sold prices is also concerning because it may cause market disturbances to persist and propagate.

To examine this possibility, we first document how post-listing Zestimate tracks listing price in places where Zillow has less or lower-quality information. Similar to analyses in Section 4.2, we test these implications using variations in non-disclosure, neighborhood heterogeneity, and Zestimate’s confidence interval. For example, we regress the updated Zestimate on lagged Zestimate and its interaction with Austin, plus the listing price and its interaction with Austin. The aim of this analysis is to evaluate how the correlation between lead Zestimate and listing price and lagged Zestimate vary across different scenarios.

Results in Table 7 Column (1) suggest that the updated Zestimate follows both the listing price and the lagged Zestimate, but the correlation with listing price is much greater (0.717 vs. 0.254). Column (2) reports evidence that in Austin, the updated Zestimate is less correlated with lagged Zestimate and more correlated with listing price, compared to other cities. We get similar results for heterogeneous neighborhoods (Column 3).

These results suggest that in places where there is less public information or less comparability across nearby houses, Zestimates may depend more on listing price. Note that our earlier results show that these are the same markets where sellers rely more on the lagged Zestimate. Therefore, these findings suggest a stronger co-movement between Zestimate and listing price in these places. That being said, when we test the heterogeneity across Zestimate confidence interval (Column 4), we do find an opposite pattern — the updated Zestimate is less correlated with listing price and more correlated with lagged Zestimate if the focal property has a greater confidence interval in its Zestimate. This suggests that the feedback between Zestimate and listing price does not always reinforce each other in both directions.

Table 7: Zestimate Closely Tracks Listing Price

D.V.: log(ZestLead1)	(1) All	(2) Austin vs. other cities	(3) Nbhd. heterogeneity	(4) ZestLag1 confidence interval
loglistprice	0.717*** (0.003)	0.711*** (0.003)	0.713*** (0.007)	0.757*** (0.006)
logZestLag1	0.254*** (0.003)	0.263*** (0.003)	0.276*** (0.007)	0.216*** (0.006)
loglistprice × Austin		0.025*** (0.006)		
logZestLag1 × Austin		-0.035*** (0.006)		
loglistprice × heteroHi			0.020** (0.009)	
loglistprice × heteroLo			-0.032*** (0.009)	
logZestLag1 × heteroHi			-0.036*** (0.009)	
logZestLag1 × heteroLo			-0.004 (0.009)	
loglistprice × CI.Hi				-0.070*** (0.007)
logZestLag1 × CI.Hi				0.073*** (0.007)
bedrooms	0.000 (0.001)	0.000 (0.001)	0.001 (0.001)	0.000 (0.001)
bathrooms	0.002* (0.001)	0.002** (0.001)	0.003*** (0.001)	0.003** (0.001)
logFloorsize	0.016*** (0.002)	0.016*** (0.002)	0.017*** (0.003)	0.016*** (0.003)
logAge	-0.001 (0.001)	-0.001 (0.001)	-0.000 (0.001)	-0.000 (0.001)
condo	-0.011*** (0.002)	-0.011*** (0.002)	-0.008** (0.003)	-0.008** (0.003)
CI.Hi				-0.027 (0.030)
Observations	35,502	35,502	20,032	19,747
R-squared	0.983	0.983	0.984	0.984
nbhd-by-month FE	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Moving beyond the listing stage, Figure 9 plots the ratio of Zestimate to sold price in six rounds before and six rounds after the listing round (i.e., Round 0). We observe two interesting patterns here: First, Zestimates before listing (i.e., Round -6 to Round -1) tend to underpredict the sold price as the Zestimate-sold-price ratio is predominately below 1 for all cities and all rounds. This is reasonable because Zestimate before listing is based on incomplete information. For example, there

may be a recent renovation with the property that was not captured by lagged Zestimate. Second, in the six rounds post listing (Rounds 1 to 6), the Zestimate-sold-price ratio is significantly above 1 for almost all cities and almost all rounds, suggesting that when Zestimate follows the listing price too closely, it may generate erroneous predictions of the market value of the property.




Figure 9: Zestimate-Sold Price Ratio Before and After Listing

Intuitively, disturbance propagation is not problematic if the market is able to self correct. For example, market can self correct if disturbances eventually cancel out each other — for example, if Zestimate introduces positive errors to some properties and negative errors to others, then the market as a whole can still converge to the efficiency benchmark, provided that market participants respond to positive and negative noises similarly.

We show evidence that sellers in fact have asymmetric responses to Zestimate. In Figure 10, we divide all properties into three groups based on their values. For each value group, we plot the linear fit of logged listing price on logged $ZestLag1$, separately for three cases: Zest Medium, Zest Low, and Zest High. We use local comparable homes (i.e. comparable homes within one mile) to construct seller’s expected fair market value of their own property, so when the lagged Zestimate is significantly below ($< 80\%$) that perceived fair market value, we categorize that case as “Zest Low”. It is “Zest High” if the lagged Zestimate is significantly greater than ($> 120\%$) the perceived fair market value. Everything in between is categorized as “Zest Medium”. As shown in Figure 10, seller’s reliance on lagged Zestimate is the strongest when it is closer to their expected value of the property. The dependence is weaker when Zestimate is higher and weakest when Zestimate is lower. This suggests that sellers find Zestimate most believable when it aligns with their expectation; when

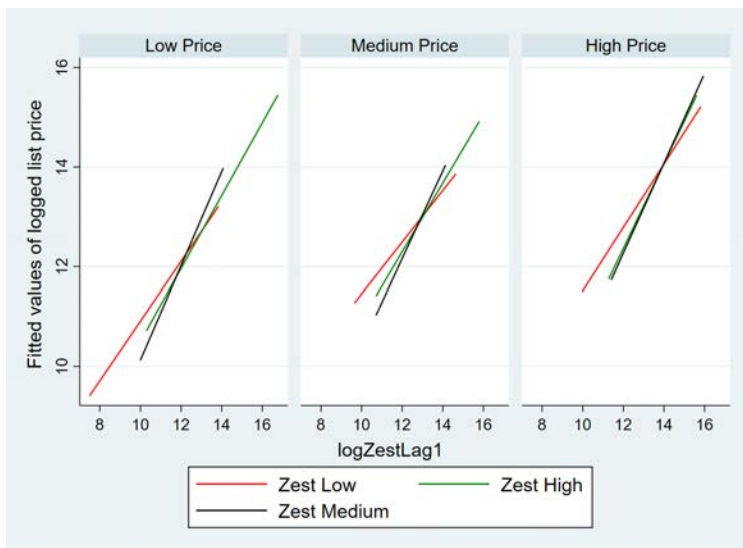


Figure 10: Asymmetry in Seller Response to $ZestLag1$

it does not, they put less trust in it when it is below their own valuation than when it is above.

Similarly, buyers also exhibit asymmetric responses to $ZestLead1$. In Figure 11, we plot the linear fit of logged sold price on logged $ZestLead1$, separately for Zest Medium, Zest Low, and Zest High. Here we define “Zest Low” (“Zest High”) if a property’s $ZestLead1$ is significantly lower (higher) than the listing price. We observe that, across all property-value types, there is a smaller correlation between $ZestLead1$ and logged sold price when $ZestLead1$ is significantly below the listing price than when it is not. Therefore, both sellers and buyers are more likely to agree with Zestimate when it is close to or higher than their evaluation benchmark (i.e., local comps for sellers and the listing price for buyers).

These asymmetries are concerning, because noises reinforced more in the positive direction may contribute to a housing market bubble. On top of that, the way that Zillow defines comparable homes could compound the effect. Figure 12a counts the number of properties who have the focal property as a comparable home, in the six rounds before and six rounds after the listing round. It shows that the focal property can only affect a limited number of properties when it is off market, but once listed, the focal property can affect as many as 150 properties. Therefore, any error in the focal property’s value can be amplified by this large comparable homes network. In particular, Figure 12b shows that the homes affected after listing are on average within one mile from the focal property, suggesting that the noise is more likely to affect local properties although distant properties may also be affected indirectly (through the comps of comps, for example).




Figure 11: Asymmetry in Buyer Response to *ZestLead1*

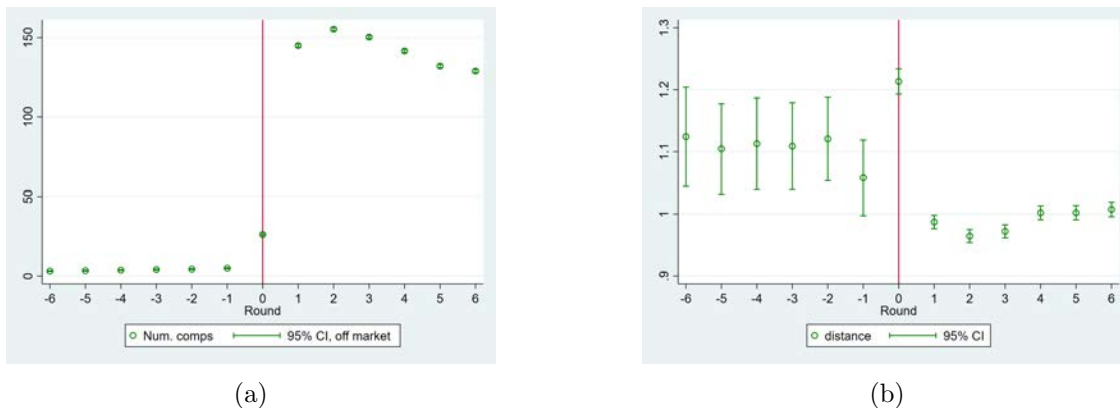


Figure 12: Homes Affected by the Focal Property, Before and After Listing

5.3 Does Zestimate Propagate the Initial Covid-19 Shock?

So far, empirical patterns point to a mixture of potential implications of Zestimate: on the one hand, users rely more on Zestimate when it is more costly to access the same information from alternative sources, implying potential efficiency gains from the improved transparency because of Zestimate; on the other hand, Zestimate update also tracks listing price more closely when information from alternative sources is more costly, implying a feedback loop that could propagate disturbances in any stage of the sales process.

To causally tease out the net effect of these two forces, an ideal setting is to have a replica of housing markets where Zestimate is only available in a random subset of the markets. In reality, we do not have such experimental or quasi-experimental settings, because: (1) Zillow’s launch in

2006 was nationwide and there was no staggered entry to leverage; (2) Zestimate availability at the property level does vary within the same geographical market, but such a study may require matching of properties which is apparently subject to critique in that properties are unique in nature and matching on observables cannot cleanly construct counterfactuals. Besides, matching-based analysis may suffer from spillover bias given that matched properties are inherently competitors.

Given the lack of clean lab settings, we turn to the natural experiment created by COVID-19. The assumption is that COVID-19 created a shock to the local housing market in the initial weeks since the presidential declaration of national emergency¹¹. If we focus on the properties that had already been active on the market at the time of the declaration, their listing price was preset. For them, the sudden direct shock from the declaration is from other market participants, likely on the demand side. From their final sales price, we can measure the size of the shock to the local market. If Zestimate plays an important role in disturbance propagation, we should observe markets subject to a positive shock at the declaration to go up further afterwards, and markets subject to a negative shock at the declaration to go down further afterwards. If, however, the market self corrects, such path dependence may not be present in the data.

To construct a measure of the neighborhood-level COVID shock, We first identify the properties that were listed between January 1, 2020 and March 13, 2020 and went into pending after March 13, 2020 but before May 1, 2020. We take these properties as the ones that were caught in surprise by COVID. For each of these “surprise” properties, we construct their sold price to listing price ratio (henceforth sold-to-listing ratio), under the assumption that this ratio reflects the initial COVID-19 shock to the demand side. However, cross-sectionally, this ratio could also reflect many fundamental differences across neighborhoods even without COVID-19, and thus we need to normalize it by the historical sold-to-listing ratio within the same neighborhood. To do so, for each neighborhood, we identify all properties sold in 2019 and construct their sold-to-listing ratio. Then we collapse these two sold-to-listing ratios — one for the surprise properties and one for those sold in 2019 — to the neighborhood level by taking the median across properties. The difference between these two ratios tells us the direction and size of the initial COVID shock.

Accordingly, we sort all neighborhoods into three terciles (positive shock, minimal shock, and negative shock), depending on whether the median sold-to-listing ratio of the surprise properties is considerably higher than, similar to, or considerably lower than the neighborhood’s median ratio in

¹¹See <https://www.federalregister.gov/documents/2020/03/18/2020-05794/declaring-a-national-emergency-concerning-the-novel-coronavirus-disease-covid-19-outbreak>

2019. More specifically, an average positive-shock neighborhood has seen its median sold-to-listing ratio of the surprise properties go up 0.055 from its 2019 median¹², while an average negative shock neighborhood saw its median sold-to-listing ratio of the surprise properties go down 0.112 from its 2019 median.¹³ In comparison, the shock in the minimal-shock neighborhoods is bounded between -0.035 and +0.016.

In constructing the sample, we remove condos from the data because condos are often investment properties which may be affected by COVID in ways different than laid out above. We also drop the last three months of the data (namely Jan, Feb, Mar 2021) out of concern of data censoring — some properties that were listed close to the end of the sample period may have missing values of market outcomes by the end, and these properties are likely to be low-quality listings which gives rise to censoring bias. The data between January 1, 2020 to April 30, 2020 is isolated to construct the dummies of positive, minimal, and negative shock. This interim period is excluded from the regression sample, as we aim to identify how neighborhoods respond to the shock by comparing their market outcomes strictly before (i.e., before January 1, 2020) and strictly after (i.e., after April 30, 2020). We adopt the following Difference-in-Differences specification:

$$Y_{it} = \lambda_1 Positive_i \times Post_t + \lambda_2 Negative_i \times Post_t + \delta_i + \eta_t + \epsilon_{it}, \quad (3)$$

where Y_{it} is the outcome variable for neighborhood i at month t , including logged $ZestLag1$, logged listing price, sales probability, logged days to pending, logged days on market¹⁴, and logged sold price. We use the dummy variable $Post_t$ to denote months after the declaration, and $Positive$ and $Negative$ stand for neighborhoods subject to Positive and Negative COVID shocks, respectively. We control for neighborhood and month fixed effects with δ_i and η_t . Standard errors are clustered at the neighborhood level. Constructed this way, the baseline group is the neighborhoods with Minimal COVID shock. If disturbance propagation exists and drives path dependence, we expect λ_1 to be positive and λ_2 to be negative. To test the assumption of parallel trends, we plot the time series of logged $ZestLag1$, logged listing price, sales probability, logged days to pending, logged days on market, and logged sold price in Appendix Figure A1, which shows parallel trends across the three shock groups for all outcome variables.

Table 8 reports the regression results. For all price- or value-related variables, we do not

¹²The full range is from 0.017 to 0.216, all going up.

¹³The full range is from 0.036 to 0.372, all going down.

¹⁴Logged days to pending is defined as $\log(1+\text{days to pending})$ and logged days to market is defined as $\log(1+\text{days to market})$, because a small sample of the data (< 1%) has zero days to pending or zero days to market.

Table 8: COVID Shock

	(1)	(2)	(3)	(4)	(5)	(6)
	logZestLag1	loglistprice	sold	log(DaysToPending +1)	log(DaysOnMarket +1)	logsoldprice
Positive \times Post	0.010 (0.036)	-0.021 (0.028)	-0.037 (0.026)	0.071 (0.101)	0.084 (0.058)	-0.028 (0.039)
Negative \times Post	0.045 (0.040)	-0.026 (0.025)	-0.022 (0.034)	0.315*** (0.119)	0.180*** (0.061)	-0.013 (0.037)
Observations	1,435	1,435	1,435	1,428	1,391	1,234
R-squared	0.918	0.919	0.419	0.408	0.420	0.887
neighborhood FE	YES	YES	YES	YES	YES	YES
month FE	YES	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

find any significant coefficient estimates for variables of interest. For timing of sales (logged days to pending and logged days on market), properties that locate in the neighborhoods subject to Negative COVID shock need a longer time to sell after the shock, although probability of sales and sold price do not change significantly. In combination, these results do not suggest a strong path dependence set forth by the initial COVID shock, especially not in terms of price or property value.

This pattern is confirmed in the raw data of Zestimate. In Figure 13, we plot the trends of Zestimate aggregated across *all* properties (i.e., on market and off market) in Positive, Minimal, and Negative COVID shock neighborhoods. Specifically, we compute the median Zestimate across properties of a given neighborhood in a given round, and then take the average across neighborhoods of the same COVID shock group in that month. No evidence suggests that Zestimate after COVID goes up the most for Positive-shock neighborhoods and goes down the most for Negative-shock neighborhoods, relative to Minimal-shock neighborhoods. Above all, evidence tends to support the null of no path dependence on the initial COVID shock, suggesting that the risk of disturbance propagation in the Zestimate algorithm, if it exists, does not manifest itself in the form of price divergence based on the initial COVID shock.

6 Mechanisms that Limit Disturbance Propagation

Since no strong evidence supports disturbance propagation in light of the initial COVID shock, a natural question is what factors help to prevent the market from propagating the shock via the Zestimate algorithm. In this section, we describe two such factors.

First, we find evidence that Zillow is more likely to remove Zestimate when Zestimate is more

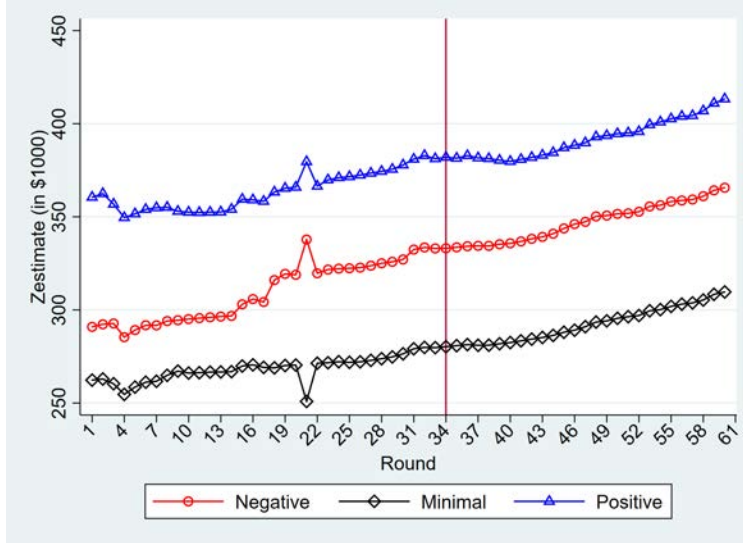


Figure 13: Zestimate by Covid Shock

at odds with the listing price, which may mitigate disturbance propagation. In Figure 14, for each city, we plot the share of properties with non-missing Zestimates in the six rounds before and six rounds after the listing round, separately by three cases — Zest Low, Zest Medium, and Zest High. For each property, we compute the average Zestimate in the six rounds before listing and compare it with the listing price. If this measure is considerably lower than listing price, this property is considered to be “Zest Low”. “Zest Medium” and “Zest High” are defined similarly. We observe that across cities, Zestimate in general becomes less available for “Zest Low” and “Zest High” than “Zest Medium”, suggesting that Zillow is more likely to make Zestimate disappear when they find it in less agreement with the listing price. In particular, we find that Zestimate is much more likely to disappear when Zestimate is lower than listing price than the other two cases in Austin, which can help limit disturbance propagation given the sellers’ asymmetric response to Zestimate. We also note that Zestimate removal is more frequent when Zestimate is higher than the listing price in Pittsburgh than when it is lower than the listing price, but the difference and magnitude is not as significant as that in Austin.

Next, we provide evidence that sellers are less affected by Zestimate when Zestimate is at odds with the information they observe. We assume that sellers closely monitor the recent sales of homes in their local neighborhoods prior to deciding on the listing price. If the sold prices of sellers’ local comparable homes are rather different from these homes’ Zestimates, then the seller may consider Zestimate to be less accurate and thus depend less on Zestimate. For each listing, we compute the absolute percentage deviation of the average Zestimate of local comparable homes in the 6 rounds

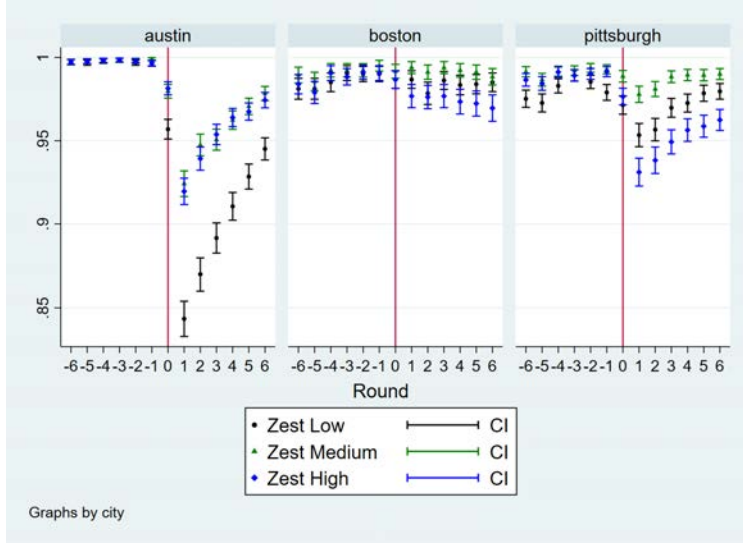


Figure 14: Zestimate Availability

before the focal property’s listing from the average sold price of these comparable homes during the same time period. This percentage deviation is highly skewed, with a median of 5.3%, a mean of 11.8%, and standard deviation of 56.1%. Given this, we transform this variable into deciles and estimate Equation 1 by including this decile variable and its interaction with $\log ZestLag1$. Our hypothesis is that in cases with a higher deviation of Zestimate from the sold price of local comps, the seller is less likely to use Zestimate in pricing her own home.

Results are reported in Table 9, where we use $dec.comps_Zest_soldP_diff$ to denote the above-mentioned Zestimate-sold price disagreement. For the estimation, We use the same variable, i.e., the average Zestimate of infrequent, non-local comparable homes to instrument for the endogenous $\log ZestLag1$ and its interaction with the disagreement variable to instrument for the interaction term. We estimate a statistically significant coefficient estimate of the interaction term at -0.002. This indicates that when the disagreement increases by 1 decile, the effect of Zestimate on listing price decreases by 0.002. Therefore, evidence supports the hypothesis that sellers are less affected by Zestimate when Zestimate is in less agreement with other sources of information.

Lastly, we adopt a similar approach to evaluate if buyers are less affected by Zestimate when they find that the lagged Zestimate is more different from the listing price. Specifically, for each property, we construct the absolute percentage deviation of lagged Zestimate from the listing price. We then transform this variable into deciles given that it is highly skewed and name the transformed variable as $dec_ZestLag_listprice_diff$. We then estimate Equation 2 by allowing for the interaction term between $dec_ZestLag_listprice_diff$ and $\log ZestLead1$. The objective is to test if the demand

Table 9: Smaller Effect of Zestimate on Listing Price When Zestimate Disagrees with Sold Price of Local Comps

D.V.: log(List Price)	(1)
logZestLag1	0.668*** (0.019)
dec_comps.Zest_soldP_diff×logZestLag1	-0.002** (0.001)
dec_comps.Zest_soldP_diff	0.024** (0.011)
logRecentSoldPriceComp1m	0.047*** (0.014)
logRecentSoldPriceLComp	0.046*** (0.015)
bedrooms	0.016*** (0.002)
bathrooms	0.054*** (0.003)
logFloorsize	0.086*** (0.007)
log(Age+1)	-0.016*** (0.002)
condo	-0.060*** (0.007)
Observations	24,620
R-squared	0.691
nbhd-by-month FE	YES

Notes: Standard errors in parentheses.
 *** p<0.01, ** p<0.05, * p<0.1

side reacts to the updated Zestimate differently when the historical Zestimate of the property (i.e., lagged Zestimate) is far from the listing price. Similar to what we did with the supply side above, we continue to use the IV approach so that the coefficient estimates can be interpreted as causal effects.

Regression outputs in Table 10 suggest that buyers are less affected by Zestimate when they find the lagged Zestimate was further away from the listing price. Specifically, we find that the updated Zestimate after listing affects the the sales probability, days to pending, days on market, and the sold price to a lesser extent.

In summary, we find that Zillow actively manages the prediction accuracy of Zestimate by pruning cases where Zestimate disagrees greatly with other types of information. At the same time, both supply and demand sides seem to use various sources of information in their decision making and they rely less on Zestimate if Zestimate is found off.

Table 10: Smaller Effect of Zestimate on Market Outcomes When Lagged Zestimate Disagrees with Listing Price

	(1) sold	(2) log(DaysToPending +1)	(3) log(DaysOnMarket +1)	(4) logsoldprice
logZestLead1	0.167 (0.155)	-3.595*** (0.632)	-1.813*** (0.343)	0.227*** (0.078)
dec_ZestLag_listprice_diff×logZestLead1	-0.004*** (0.001)	0.018*** (0.006)	0.006** (0.003)	-0.001* (0.001)
loglistprice	-0.183 (0.122)	3.407*** (0.503)	1.700*** (0.269)	0.795*** (0.061)
dec_ZestLag_listprice_diff	0.043** (0.019)	-0.195*** (0.071)	-0.072* (0.039)	0.010 (0.008)
logRecentSoldPriceComp1m	-0.001 (0.024)	0.074 (0.090)	-0.007 (0.049)	-0.009 (0.011)
logRecentSoldPriceLComp	0.048* (0.028)	-0.020 (0.106)	0.047 (0.059)	0.004 (0.013)
bedrooms	-0.010*** (0.003)	-0.022* (0.012)	0.003 (0.006)	0.003** (0.001)
bathrooms	-0.002 (0.004)	0.043** (0.017)	0.022** (0.009)	0.001 (0.002)
logFloorsize	-0.018 (0.013)	0.304*** (0.048)	0.209*** (0.028)	-0.031*** (0.006)
log(Age+1)	0.003 (0.004)	-0.030* (0.017)	-0.005 (0.009)	-0.000 (0.002)
condo	-0.081*** (0.012)	0.186*** (0.043)	0.104*** (0.024)	-0.003 (0.005)
Observations	22,772	20,149	16,711	11,982
R-squared	0.009	-0.017	-0.028	0.905
nbhd-by-month FE	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

7 Conclusion

Using bi-weekly snapshots of Zillow in three US cities, we find a strong co-movement between Zestimate and the residential housing market. In particular, Zestimate leads the market by affecting the decision making of both sellers and buyers; it also follows the market by tracking listing price and other market outcomes closely. These two-way interactions have mixed implications for market efficiency: on the one hand, Zestimate may enhance efficiency because its 24/7 presence can empower market participants with useful, easy-to-get information, especially in the places that are costly to access alternative sources of information. On the other hand, the Zestimate algorithm has the potential to ingest market disturbance at any stage of the sales cycle, and propagate it over time and across properties.

Although our study on the initial COVID shock does not reveal strong evidence in support of disturbance propagation, the COVID shock is only one of many disturbances that may occur in the sales cycle. We cannot rule out the possibility that an opposite market shock may have happened later on and cancel out the initial COVID shock. In that case, Zestimate could propagate both, but it is not detectable in our research design. It is also possible that Zestimate has little effect propagating big shocks like COVID, because the propagation may trigger the guard rails embedded in the algorithm or raise enough doubt in user confidence in Zestimate. When the shock is small or gradual, these limiting factors may not be as effective. More future research is definitely needed to better understand the risk of user-algorithm interactions.

Our research is subject to a number of other caveats. Due to technical constraint, we can only track Zillow in three US cities for two years. Whether our findings are representative of the whole US remains to be seen. In addition, we do not have any insider information as to how Zillow constructs or evolves its Zestimate algorithm. Our understanding of the algorithm is based on public information of Zillow as well as our inference from the scraped data. Furthermore, because we do not observe the private information that drives a seller's decision to go on the market, we take whether a property is on the market as exogenously given, and focus on market outcomes since the start of listing. This implies we ignore the potential impact of Zestimate on the selection of going on the market. Finally, while user-algorithm interaction is common on many e-commerce websites, it is conceivable that users treat residential properties differently, because housing sales involve a large amount of money and many information frictions are unique to the housing market. We hope our findings in the housing market would encourage future research of interactive algorithms in other markets.

References


- Franklin Allen and Douglas Gale. Financial contagion. *Journal of political economy*, 108(1):1–33, 2000.
- Lisa R Anderson and Charles A Holt. Information cascades in the laboratory. *The American economic review*, pages 847–862, 1997.
- Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias. *propublica* (2016). URL: <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>, 2016.
- Stephanie Assad, Robert Clark, Daniel Ershov, and Lei Xu. Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market. 2020.
- Abhijit V Banerjee. A simple model of herd behavior. *The quarterly journal of economics*, 107(3): 797–817, 1992.
- Lucian A Bebchuk and Itay Goldstein. Self-fulfilling credit market freezes. *The Review of Financial Studies*, 24(11):3519–3555, 2011.
- Richard Berk. Accuracy and fairness for juvenile justice risk assessments. *Journal of Empirical Legal Studies*, 16(1):175–194, 2019.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, 100(5):992–1026, 1992.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of economic perspectives*, 12(3):151–170, 1998.
- Zach Y Brown and Alexander MacKay. Competition in pricing algorithms. Technical report, National Bureau of Economic Research, 2021.
- Greg Buchak, Gregor Matvos, Tomasz Piskorski, and Amit Seru. Why is intermediating houses so difficult? evidence from ibuyers. Technical report, National Bureau of Economic Research, 2020.
- Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018.
- Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–97, 2020.
- Paul E Carrillo. An empirical stationary equilibrium search model of the housing market. *International Economic Review*, 53(1):203–234, 2012.
- Boğaçhan Çelen and Shachar Kariv. Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review*, 94(3):484–498, 2004.
- Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*, pages 134–148. PMLR, 2018.

- Bo Cowgill. Bias and productivity in humans and algorithms: Theory and evidence from resume screening. *Columbia Business School, Columbia University*, 29, 2018.
- Douglas W Diamond and Philip H Dybvig. Bank runs, deposit insurance, and liquidity. *Journal of political economy*, 91(3):401–419, 1983.
- Douglas W Diamond and Raghuram G Rajan. Fear of fire sales, illiquidity seeking, and credit freezes. *The Quarterly Journal of Economics*, 126(2):557–591, 2011.
- Runshan Fu, Yan Huang, and Param Vir Singh. Crowds, lending, machine, and bias. *Information Systems Research*, 32(1):72–92, 2021.
- Matthew Gentzkow and Jesse M Shapiro. Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839, 2011.
- Melissa Hamilton. The sexist algorithm. *Behavioral sciences & the law*, 37(2):145–157, 2019.
- Mitchell Hoffman, Lisa B Kahn, and Danielle Li. Discretion in hiring. *The Quarterly Journal of Economics*, 133(2):765–800, 2018.
- Yufeng Huang. Seller-pricing frictions and platform remedies. 2021.
- Matthew O Jackson and Agathe Pernoud. Systemic risk in financial networks: A survey. *Annual Review of Economics*, 13:171–202, 2021.
- Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. Human decisions and machine predictions. *The quarterly journal of economics*, 133(1):237–293, 2018.
- Anja Lambrecht and Catherine Tucker. Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management science*, 65(7):2966–2981, 2019.
- Anja Lambrecht and Catherine E Tucker. Apparent algorithmic discrimination and real-time algorithmic learning. *Available at SSRN 3570076*, 2020.
- Anja Lambrecht, Ananya Sen, Catherine E Tucker, and Caroline Wiertz. Algorithmic recommendations and earned media: Investigating product echo chambers on youtube. *Available at SSRN 3951425*, 2021.
- Guangli Lu. Learning from online appraisal information and housing prices. *Available at SSRN 3489522*, 2018.
- Guangli Lu. How machine learning mitigates racial bias in the us housing market. *Available at SSRN 3489519*, 2019.
- Jens Ludwig and Sendhil Mullainathan. Fragile algorithms and fallible decision-makers: Lessons from the justice system. *Journal of Economic Perspectives*, 35(4):71–96, 2021.
- Nikhil Malik. Does machine learning amplify pricing errors in housing market?: Economics of ml feedback loops. *Economics of ML Feedback Loops (September 18, 2020)*, 2020.
- Jean-Charles Rochet and Jean Tirole. Interbank lending and systemic risk. *Journal of Money, credit and Banking*, 28(4):733–762, 1996.

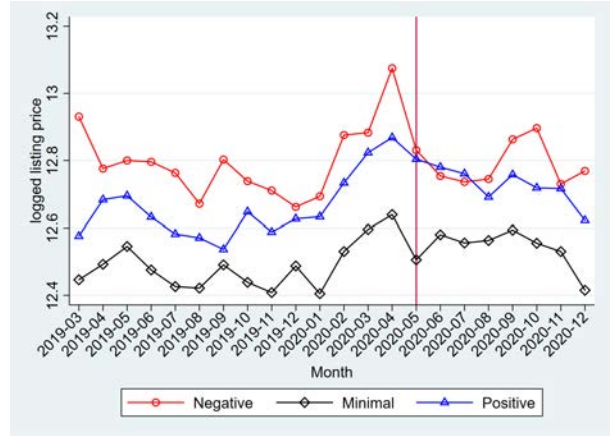
- Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.
- Megan Stevenson. Assessing risk assessment in action. *Minn. Law Review*, 103:303, 2018.
- Megan T Stevenson and Jennifer L Doleac. Algorithmic risk assessment in the hands of humans. *Available at SSRN 3489440*, 2021.
- Magdalena E Wojcieszak and Diana C Mutz. Online groups and political discourse: Do online discussion spaces facilitate exposure to political disagreement? *Journal of communication*, 59(1):40–56, 2009.
- Shuyi Yu. Algorithmic outputs as information source: The effects of zestimates on home prices and racial bias in the housing market. *Available at SSRN 3584896*, 2020.
- Shunyuang Zhang, Nitin Mehta, Param Vir Singh, and Kannan Srinivasan. Frontiers: Can an artificial intelligence algorithm mitigate racial economic inequality? an analysis in the context of airbnb. *Marketing Science*, 40(5):813–820, 2021.

Appendix


A Figures



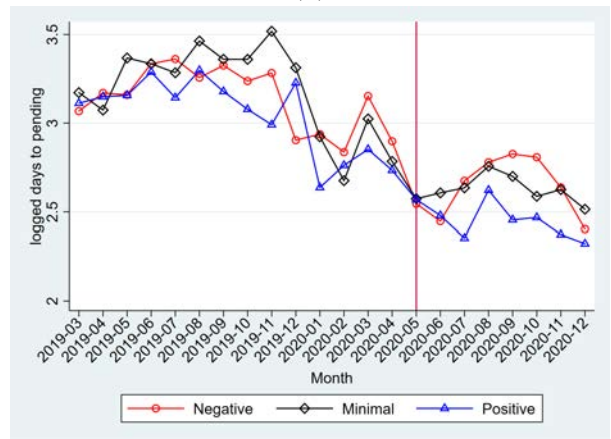
(a)




(b)




(c)



(d)



(e)



(f)

Figure A1: Paralell Trends — Covid Shock

B Tables

Table A1: Zestimate's Effect on Listing Price — OLS results

D.V.: log(List Price)	(1) All	(2) Austin	(3) Boston	(4) Pittsburgh	(5) Houses	(6) Condos
logZestLag1	0.501*** (0.008)	0.704*** (0.013)	0.520*** (0.017)	0.397*** (0.014)	0.469*** (0.010)	0.524*** (0.018)
logRecentSoldPriceComp1m	0.087*** (0.013)	0.048*** (0.018)	0.084*** (0.019)	0.137*** (0.025)	0.083*** (0.016)	0.120*** (0.022)
logRecentSoldPriceLComp	0.104*** (0.014)	0.044** (0.019)	0.098*** (0.021)	0.101*** (0.025)	0.114*** (0.017)	0.049** (0.023)
bedrooms	0.016*** (0.002)	0.007*** (0.003)	0.018*** (0.003)	0.020*** (0.003)	0.019*** (0.002)	0.016*** (0.006)
bathrooms	0.062*** (0.002)	0.019*** (0.003)	0.044*** (0.004)	0.102*** (0.004)	0.068*** (0.003)	0.046*** (0.007)
logFloorsize	0.114*** (0.006)	0.132*** (0.009)	0.157*** (0.012)	0.083*** (0.010)	0.105*** (0.007)	0.174*** (0.015)
log(Age+1)	-0.018*** (0.002)	-0.031*** (0.003)	-0.006* (0.004)	-0.026*** (0.006)	-0.022*** (0.003)	-0.008** (0.004)
condo	-0.065*** (0.006)	-0.099*** (0.008)	-0.003 (0.009)	-0.044*** (0.016)		
Observations	24,622	9,000	4,583	11,039	20,519	3,418
R-squared	0.923	0.928	0.881	0.818	0.917	0.951
nbhd-by-month FE	YES	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A2: Zestimate's Effect on Listing Price — IV First Stage

D.V.: log(ZestLag1)	(1) All	(2) Austin	(3) Boston	(4) Pittsburgh	(5) Houses	(6) Condos
log(IV)	0.336*** (0.004)	0.300*** (0.006)	0.204*** (0.008)	0.422*** (0.008)	0.349*** (0.005)	0.249*** (0.011)
logRecentSoldPriceComp1m	0.097*** (0.009)	0.066*** (0.014)	0.080*** (0.017)	0.152*** (0.016)	0.109*** (0.011)	0.098*** (0.023)
logRecentSoldPriceLComp	0.338*** (0.010)	0.271*** (0.014)	0.414*** (0.017)	0.264*** (0.016)	0.325*** (0.011)	0.304*** (0.022)
bedrooms	0.000 (0.001)	0.002 (0.002)	-0.009*** (0.002)	0.009*** (0.002)	0.002* (0.001)	0.018*** (0.006)
bathrooms	0.030*** (0.002)	0.026*** (0.002)	0.021*** (0.004)	0.034*** (0.003)	0.030*** (0.002)	0.031*** (0.007)
logFloorsize	0.110*** (0.004)	0.170*** (0.007)	0.189*** (0.010)	0.072*** (0.007)	0.096*** (0.005)	0.210*** (0.014)
log(Age+1)	-0.013*** (0.002)	-0.003 (0.002)	-0.004 (0.003)	-0.037*** (0.004)	-0.016*** (0.002)	-0.004 (0.004)
condo	-0.009* (0.005)	-0.089*** (0.006)	0.016** (0.008)	0.000 (0.010)		
Observations	24,622	9,000	4,583	11,039	20,519	3,418
R-squared	0.961	0.950	0.904	0.923	0.961	0.953
nbhd-by-month FE	YES	YES	YES	YES	YES	YES
F-stat (IV=0)	6109.92	2516.07	616.63	3159.28	5158.19	537.04

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A3: Zestimate's Effect on Listing Price — Alt. IV (1)

D.V.: log(List Price)	(1) All	(2) Austin	(3) Boston	(4) Pittsburgh	(5) Houses	(6) Condos
logZestLag1	0.631*** (0.025)	0.875*** (0.047)	0.653*** (0.060)	0.566*** (0.032)	0.625*** (0.029)	0.531*** (0.057)
logRecentSoldPriceComp1m	0.054*** (0.016)	0.050* (0.027)	0.053** (0.023)	0.073*** (0.027)	0.046** (0.020)	0.116*** (0.027)
logRecentSoldPriceLComp	0.047** (0.018)	-0.040 (0.033)	0.045 (0.032)	0.045* (0.027)	0.043* (0.022)	0.045 (0.028)
bedrooms	0.017*** (0.002)	0.010** (0.004)	0.019*** (0.003)	0.018*** (0.004)	0.020*** (0.002)	0.016** (0.007)
bathrooms	0.061*** (0.003)	0.014*** (0.005)	0.039*** (0.005)	0.090*** (0.005)	0.067*** (0.004)	0.051*** (0.008)
logFloorsize	0.083*** (0.008)	0.074*** (0.015)	0.127*** (0.018)	0.061*** (0.011)	0.069*** (0.009)	0.167*** (0.021)
log(Age+1)	-0.013*** (0.003)	-0.030*** (0.004)	-0.006 (0.004)	-0.017*** (0.006)	-0.014*** (0.004)	-0.007* (0.004)
condo	-0.052*** (0.008)	-0.086*** (0.014)	-0.006 (0.009)	-0.040** (0.016)		
Observations	19,391	4,170	4,439	10,782	15,748	3,021
R-squared	0.682	0.784	0.821	0.603	0.592	0.820
nbhd-by-month FE	YES	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A4: Zestimate's Effect on Listing Price — Alt. IV (2)

D.V.: log(List Price)	(1) All	(2) Austin	(3) Boston	(4) Pittsburgh	(5) Houses	(6) Condos
logZestLag1	0.705*** (0.027)	0.808*** (0.032)	0.761*** (0.087)	0.630*** (0.046)	0.703*** (0.030)	0.623*** (0.063)
logRecentSoldPriceComp1m	0.056*** (0.016)	0.033 (0.020)	0.058* (0.030)	0.082*** (0.029)	0.058*** (0.018)	0.147*** (0.035)
logRecentSoldPriceLComp	0.031 (0.021)	0.038 (0.025)	-0.035 (0.052)	0.036 (0.036)	0.024 (0.023)	-0.037 (0.042)
bedrooms	0.016*** (0.002)	-0.000 (0.003)	0.027*** (0.003)	0.018*** (0.004)	0.015*** (0.002)	0.018** (0.008)
bathrooms	0.041*** (0.003)	0.011*** (0.004)	0.024*** (0.006)	0.068*** (0.005)	0.045*** (0.003)	0.032*** (0.009)
logFloorsize	0.067*** (0.008)	0.106*** (0.014)	0.091*** (0.021)	0.039*** (0.013)	0.062*** (0.009)	0.105*** (0.024)
log(Age+1)	-0.013*** (0.003)	-0.019*** (0.003)	-0.009* (0.005)	-0.012* (0.006)	-0.012*** (0.003)	-0.015*** (0.005)
condo	-0.066*** (0.007)	-0.067*** (0.010)	-0.031*** (0.012)	-0.053*** (0.016)		
Observations	17,657	6,699	2,472	8,486	15,331	1,673
R-squared	0.716	0.810	0.811	0.647	0.651	0.816
nbhd-by-month FE	YES	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A5: RDD Balance Check

	(1) loglistprice	(2) logZestLag1	(3) bedrooms	(4) bathrooms	(5) logFloorsize	(6) log(Age+1)
above	-0.002 (0.017)	-0.006 (0.017)	-0.011 (0.055)	0.001 (0.047)	-0.019 (0.019)	0.048 (0.038)
Observations	1,873	1,873	1,873	1,873	1,873	1,873
R-squared	0.677	0.672	0.182	0.115	0.172	0.401
city FE	YES	YES	YES	YES	YES	YES
year-month FE	YES	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1