Does Human-algorithm Feedback Loop Lead To Error Propagation? Evidence from Zillow's Zestimate *

Runshan Fu[†] New York University

Ginger Zhe Jin [‡] University of Maryland and NBER

Meng Liu [§] Washington U. in St. Louis

Abstract

We study how home sellers and buyers interact with Zillow's Zestimate algorithm throughout the sales cycle of residential properties, with an emphasis on the implications of such interactions. In particular, leveraging Zestimate's algorithm updates as exogenous shocks, we find evidence for a human-algorithm feedback loop: listing and selling outcomes respond significantly to Zestimate, and Zestimate is quickly updated for the focal and comparable homes after a property is listed or sold. This raises a concern that housing market disturbances may propagate and persist because of the feedback loop. However, simulation suggests that disturbances are short-lived and diminish eventually, mainly because all marginal effects across stages of the selling process—though sizable and significant—are less than one. To further validate this insight in the real data, we leverage the COVID-19 pandemic as a natural experiment. We find consistent evidence that the initial disturbances created by the March-2020 declaration of national emergency faded away in a few months. Overall, our results identify the human-algorithm feedback loop in an important real-world setting, but dismiss the concern that such a feedback loop generates persistent error propagation.

Keywords: Algorithm, Human-algorithm feedback loop, Error Propagation, Zillow, Zestimate.

*We thank Emilio Calvano, Bo Cowgill, Andrey Fradkin, Xiang Hui, Zhentong Lu, Nikhil Malik, Daewon Sun, 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.

[†]rf2583@stern.nyu.edu

[‡]ginger@umd.edu

[§]mengl@wustl.edu

1 Introduction

As Machine Learning (ML) algorithms continue to advance and play important roles in human decision making, the *interaction* between algorithms and humans has become a significant aspect of modern society. In many scenarios, humans incorporate the algorithmic output in their decisions; at the same time, algorithms track user decisions and ingest the usage data for algorithmic updates. For example, general search engines rank relevant information based on the browsing history of other users, who 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.

The interactive nature of ML algorithms can create a feedback loop, where the output of an algorithm influences human behavior, which then serves as an input to update the same algorithm. This feedback loop can amplify or suppress certain patterns or behaviors, leading to potentially harmful consequences such as filter bubbles, polarization, and lack of diversity (e.g., Levy [2021]). Previous studies have explored the existence of algorithmic feedback loops and their implications. However, the empirical evidence on each step of an algorithmic feedback loop is still limited. As Cowgill and Tucker [2019] state, "causal inferences about 'algorithmic feedback loop' are inherently difficult for empiricist." They explain the challenges as "identifying a feedback loop requires a researcher not only locate [a quasi-experimental intervention], but also that they must trace intervention as it propagates into codified outcomes as well as future actions and conclusions." In addition, previous discussions about algorithmic feedback loops are mainly in the context of recommender systems in e-commerce or content consumption in social media. However, little is known about algorithmic feedback loops in many other parts of the economy, especially traditional markets that have recently embraced digitization.

In this paper, we study a human-algorithm feedback loop in residential housing markets. We focus on an ML algorithm named "Zestimate", which is developed by Zillow.com, the leading online real estate marketplace company. 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[1] and user-submitted data into Zillow's proprietary formula, also taking into account home facts, location

[1]MLS stands for Multiple Listing Service, a local platform for real estate agents to post and view active property listings.

and market trends." This algorithm produces Zestimates for 104 million homes across the country, which are constantly updated to reflect changing market conditions and new information. As Zillow primarily generates revenue by providing advertising services to real estate agents, mortgage lenders, and other related businesses, it has incentives to provide accurate and informative home value estimates to attract potential buyers and sellers to visit the website. While Zillow may not have access to all the specific local information that market players have, the algorithm, given the computational power, can aggregate and process a vast amount of global information to generate reasonable estimates. 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 housing markets involve significant uncertainty, Zestimates—with reasonable accuracy and 24/7 availability—can potentially influence human decisions and market outcomes such as listing prices and sold prices. These, in turn, are the inputs used by the algorithm to generate Zestimates for other properties, which forms an algorithmic feedback loop. Such a feedback loop is concerning, as it could propagate or sustain disturbances or errors[3] in the system, creating systematic bias in local markets. With a unique dataset that tracks Zestimates and market outcomes over time for properties in Austin, Boston, and Pittsburgh, we have a rare opportunity to observe how Zestimates influence various stages of the home selling process and how Zillow updates Zestimates after listing and sold prices become public information.

Our first objective is to identify and measure the human-Zestimate feedback loop. Specifically, there are four stages in the sales cycle. In the listing stage, a house's pre-listing Zestimate might affect the seller's listing price. After listing, Zestimate may be updated with the new listing price. Near the end of the sales cycle, the pre-sold Zestimate could influence the transaction outcomes (sold or not, days to pending, and sold price). Once the transaction is completed, Zestimates may be updated again with the new transaction information.

The main challenge we face is identifying the causal effect of Zestimate on listing price, sold price, and market outcomes. To address it, we control for various confounders including detailed home characteristics, neighborhood-specific heterogeneity, and market trends at the city-by-round level. However, even with our rich set of controls, there may still be unobservable factors that influence both buyers' and sellers' decisions as well as Zestimates. Consequently, strong correlations can be observed even if Zestimates do not causally affect market outcomes.

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

[3]We use the term "error" to refer to the deviation of an observed or estimated value from the corresponding "true value".

To overcome this challenge, we leverage three algorithm updates during our observation periods. Zillow occasionally updates the Zestimate algorithm to improve its prediction accuracy. The algorithm updates occur unexpectedly for individual users and lead to sudden changes in Zestimate values for most properties. While how the market values a property or specific home characteristic may also change over time, such market-wide shifts are generally gradual and slow. In contrast, the algorithm updates lead to abrupt and significant changes in Zestimate, and therefore creates a source of exogenous variations in Zestimate. However, it is important to note that the algorithm updates do not change Zestimate values randomly. Instead, they aim to improve Zestimate accuracy, hence the Zestimate change brought by the algorithm updates may be correlated with unobserved property characteristics. Therefore, the Zestimate change itself may be endogenous. To address this issue, we construct a set of Instrumental Variables (IVs) by capturing the part of Zestimate changes – caused by Zillow's algorithm updates in April, June, and October 2019 – based on observable factors only. The IVs essentially capture the changes in the Zestimate algorithm's valuation of *observable factors*, and therefore are by construction orthogonal to the unobservables. We then apply the IV-predicted Zestimates to the data periods close to the algorithm updates, controlling for all the observable factors.

With the IVs, we find significant dependence of market outcomes on Zestimates. Specifically, on average, listing prices increase 0.733% for every 1% increase in the property's pre-listing Zestimate. While Zestimate has a sizable influence on the listing price at the margin, its effect on the property's sold price is much more limited — conditional on listing price, a 1% increase in the Zestimate (updated after the property has been listed) leads to a 0.081% higher sold price.

As for the dependence of Zestimates on market outcomes, Zillow explicitly acknowledges that the Zestimate algorithms take listing prices and sold prices into account when generating home value estimates.[4] We provide consistent evidence that Zestimates are updated on listing prices and sold prices immediately after a property is listed or sold. Although the exact algorithm structure may change over time, the fundamental logic of the algorithm remains largely the same. Similar to a human appraisal, the Zestimate algorithm first finds recently on-market properties that are similar to the focal property (i.e., "comparable homes") and then uses their prices as the raw ingredients to produce a home value estimate for the focal property.[5] This algorithm design, which takes market outcomes as inputs, together with the causal impacts of Zestimate on market outcomes, establishes

[4]See https://www.zillow.com/sellers-guide/influencing-your-zestimate/, last accessed on 03/18/2023
[5]See https://www.zillow.com/tech/human-and-machines-home-valuation/, last accessed on 03/18/2023

the user-Zestimate feedback loop in the housing market.

Given the existence of human-Zestimate feedback loop, the second objective of this paper is to examine its market implications. Similar to algorithmic feedback loops in other contexts (such as filter bubbles and echo chambers), the main concern is the potential reinforcement of bias or errors and its contribution to polarization. The interactions between Zestimate and market outcomes may cause disturbances at any stage of the sales cycle to persist and propagate. For example, Zestimate may contain prediction error, and extreme buyer or seller taste may lead to outlier prices. With user-algorithm interactions, these disturbances can persist and, under the comparable-homes model, spread out in local markets. This concern is particularly worrisome as we identify sizable effects in each step of the human-Zestimate interaction.

To examine potential market implications, we first conduct simulations based on the estimated effects. This allows us to understand, in a mechanical way, how disturbances on Zestimates or sold prices evolve over time with user-algorithm interactions. As a stress test, we choose to introduce disturbances that are unrealistically large and widespread. The simulations show that, although the sizable interactive effects seem concerning, they do not cause disturbances to persist or spread for a long time. This is primarily due to two reasons. First, the disturbance's magnitude decreases in each stage of the sales process, because each marginal effect—despite sizeable and significant—is less than one in absolute magnitude. Second, the average mechanism in the Zestimate algorithm[6] limits the influence of disturbances in an individual property, as it allows positive and negative disturbances to cancel out across properties. In reality, disturbances tend to occur at smaller scales for subsets of properties. Moreover, some guardrails in reality, such as Zillow's practice of hiding Zestimate when it is at odds with the listing price, could further reduce the impact of error propagation.

To validate this insight in the real data, we leverage a natural experiment created by the COVID-19 pandemic. The idea is that COVID-19 created a shock of different directions to local housing markets in the weeks after the presidential declaration of national emergency. The shock consists of two parts: changes in buyer taste or market conditions that are real and persistent, and idiosyncratic disturbances that are independent of the real changes. We define a zipcode as a local market and measure the idiosyncratic disturbance for each zipcode in our data. We then compare the market outcomes for properties in zipcodes with positive disturbances with those in zipcodes

[6]The Zestimate algorithm generates home value estimates by taking the average of comparable homes' prices. See https://www.zillow.com/tech/human-and-machines-home-valuation/, last accessed on 03/18/2023.

with negative disturbances. If user-algorithm interactions propagate or sustain disturbances, we should observe widening or persisting gaps in the market outcomes between markets subject to a positive shock at the declaration and markets subject to a negative shock at the declaration. If the market self-corrects as our simulation suggests, we should observe little path dependence and narrowing gaps. With Difference-in-Differences (DiD) regressions, we find that the gaps in market outcomes shrink over time and disappear eventually, which is consistent with our simulation results.

This paper makes several contributions. First, to our knowledge, we are among the first to provide explicit evidence on the human-algorithm feedback loop in an important segment of the economy. While the idea of human-algorithm feedback loop may seem intuitive, documenting its presence and magnitude in different stages is often challenging. Our unique dataset provides a rare opportunity to observe the human-algorithm interactions in various stages of home selling process, and we leverage several algorithm update shocks to identify some key causal links. Second, we discuss the potential market implications of the user-algorithm interactions. We show that the interactive effects—though sizable and significant—do not necessarily lead to persistent propagation of disturbances, as the market self-corrects. Our findings highlight the importance to consider market forces and the context of algorithm use when designing and evaluating algorithms.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the Zestimate algorithm, our data, and the algorithm updates. Section 4 documents user-Zestimate interactions in various stages of home selling. Section 5 explores the potential market implications of user-Zestimate interactions. Section 6 discusses the underlying mechansims for the marginal effects that comprise the feedback loop by examining the neighborhood-level housing heterogeneity. A brief conclusion is offered in Section 7.

2 Literature Review

Our paper is closely related to a fast growing literature on algorithm feedback loop, and its implications such as filter bubbles, echo chambers, polarization of political views, and diversity of content consumptions and product choices. As described before, algorithm feedback loop refers to a process where algorithms—usually in the form of recommender algorithms—affect user behavior and the affected user behavior further feeds into the next iteration of algorithmic recommendations. A common concern about algorithm feedback loop is that it can result in limited user exposure to diverse content or opinions.

The literature offers mixed evidence on these issues, with some studies finding positive evidence on filter bubbles, while others do not find such causal links. 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. Levy [2021] find that Facebook's news feeding algorithm may exacerbate polarization of news consumption by limiting users' exposure to counter-attitudinal news. In contrast, Gentzkow and Shapiro [2011] find that there is "no evidence that the Internet is becoming more segregated over time", and a similar conclusion is reached by Boxell et al. [2017], who find that "polarization has increased the most among the demographic groups least likely to use the Internet and social media"; 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. The definition of diversity can also play a part in this debate. For example, Holtz et al. [2020] find that personalized recommendations on a music streaming platform may reduce within-user consumption diversity while promoting across-user consumption diversity. Lee and Hosanagar [2019] find that recommenders in an e-commerce setting may decrease the aggregate sales diversity but they do not necessarily decrease the within-consumer sales diversity.

To limit filter bubbles and polarization, a commonly proposed strategy is to recommend users more diverse content.[7] For example, Bail et al. [2018] test the effect of exposing Twitter users to counter-attitudinal views, where they find that such interventions can lead to greater polarization. Chen et al. [2023] find that a recommender designed to provide users with more diverse content may decrease user engagement without increasing the diversity actually consumed, and the effect depends on user engagement on the platform. Di Tella et al. [2021] find that the effect of such strategies depends on whether users start within or outside of the filter bubble. Relatedly, Moehring [2022] demonstrates that a non-personalized news recommendation can alleviate filter bubbles on Reddit.

We contribute to the literature on algorithm feedback loop and filter bubbles in three important ways. First, while much of the literature draws evidence from social media and content consumption, we take these ideas to the housing market, an important part of the economy with high stake transactions. Second, we identify the algorithm feedback loop by making step-by-step causal esti-

[7] We note that there are emerging studies, mostly in the computer science literature, that investigate ways to depolarize from technical perspectives, such as Sun et al. [2019], Krauth et al. [2022], Stray [2021].

mates. To the best of our knowledge, the causal identification of the entire process of algorithm feedback loop is still scarce in the literature. Last, we study the marketwide implications of the feedback loop. We demonstrate how and why the feedback loop in the housing market does not necessarily lead to persistent error propagation in the system.

Our study relates to a broader literature on algorithms' effects on the bias and quality of human decision making. On the one hand, positive evidence is found in a wide range of contexts including hiring (Cowgill [2018], Hoffman et al. [2018]), bailing (Kleinberg et al. [2018]), and crowd lending (Fu et al. [2021]). On the other hand, studies show that algorithms may perform no better than humans (Berk [2019], Stevenson [2018], Stevenson and Doleac [2021]), or even lead to bias and inequality (Angwin et al. [2016], Hamilton [2019]). 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]), where researchers attribute part of the challenge to human involvement in algorithm building and algorithm use. These issues give rise to emerging studies (e.g., Green and Chen [2019]) that investigate the appropriate principles and guidelines to resolve societal problems of "algorithm-in-the-loop." We contribute to this literature by not only providing evidence on Zestimate's effect on housing market outcomes, but also highlighting the two-way interactions between housing market participants and AI algorithms. In addition, the algorithm feedback loop in our context naturally leads to the concern that random disturbances or errors would be propagated or sustained and therefore create systemic bias in local markets. We further investigate the marketwide implications of the feedback loop, and provide rationale and evidence that alleviate such concerns.

Our work also relates to a strand of literature 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]). Huang [2021] further 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. Relatedly, Calder-Wang and Kim [2023] find that the adoption of algorithms in the multifamily rental market enables building managers to set prices more responsively, and markets with higher algorithm usage had lower occupancy rates and higher rents during the post-crisis period. 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 directly related to several recent studies of 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. Our paper provides a new perspective because we study not just Zestimate's effect on market outcomes, but also how the market and Zestimate interact and affect each other as an equilibrium outcome. 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 marketwide implications of the feedback loop.

Finally, the human-algorithm interaction we study in this paper is related to the literature on 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 a belief shift (i.e., an algorithm update) and how it interacts with players in the market, rather than multiple equilibria and feedback effects among the players within a system.

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 could only access market information through an agent, where agents post listings in the local MLS. 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,[8] and Redfin.

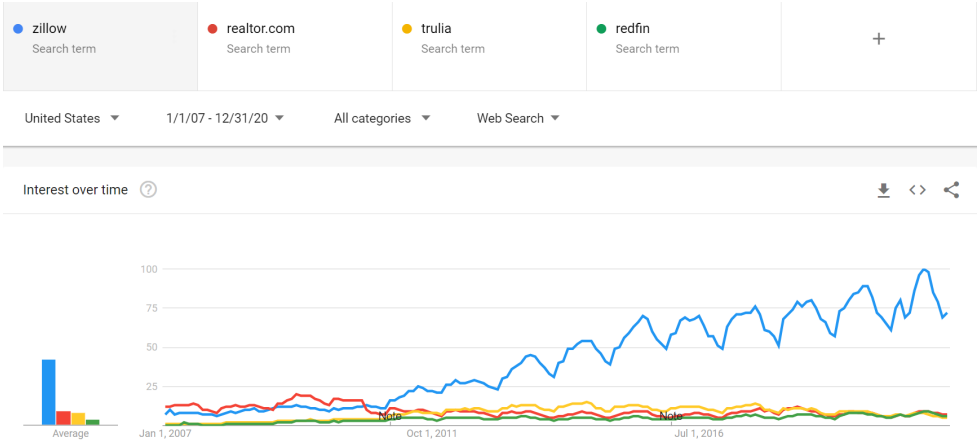
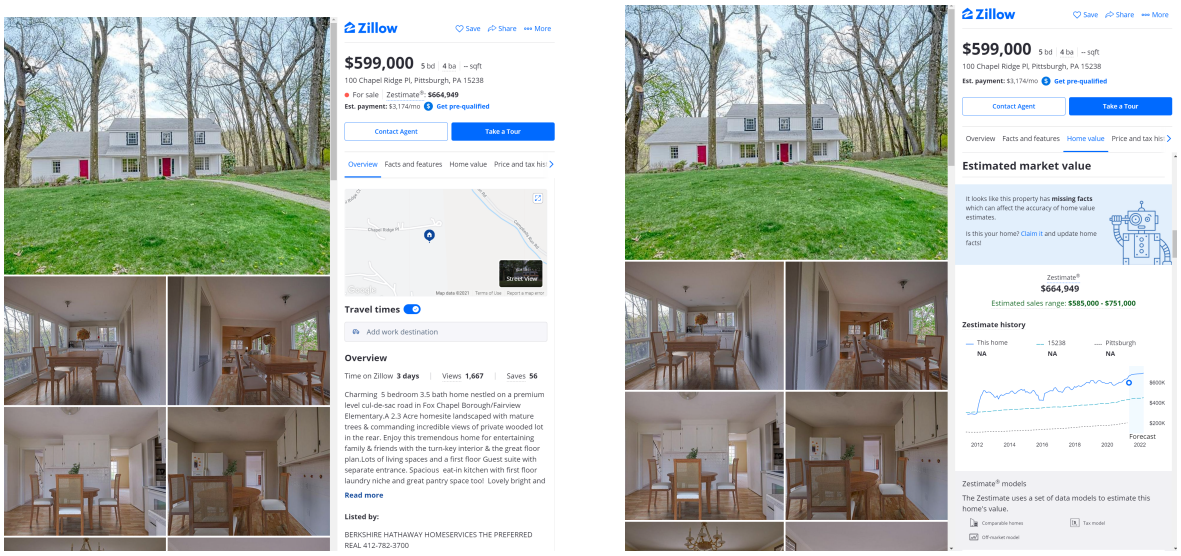


Figure 1: Rise of Digital Real Estate Platforms

In this paper, we focus on Zillow, which has been a dominating player in the market. When users search for listings in a region on Zillow, it returns all active for-sale listings scattered on a map and listed on a panel, 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 2a), where photos of the property will be displayed on the left. On the top right of the listing page, key information about the property is displayed, including listing price, address, number of beds, baths, and square footage. 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.

One of Zillow's most notable features is Zestimate, a home value estimate for individual properties generated by ML algorithms. Zillow primarily profits from providing advertising services

[8] Trulia has been acquired by Zillow in 2014 but remained its website operation



(a) Zillow Listing

(b) Zestimate and Zestimate History

Figure 2: Zillow Listing Page

to real estate agents, mortgage lenders, and other related businesses. As the advertising business crucially depends on organic website traffic, Zillow has incentives to provide accurate and informative Zestimate to encourage page views from potential sellers and buyers. As a main feature of the website, Zestimate is saliently visible on the listing page. As illustrated in Figure 2a, the top of the listing page shows the Zestimate right below the listing price. The listing page also has a separate section dedicated to Zestimate, as showed in Figure 2b. In addition to the point estimate, Zillow also offers an estimated sales price range, which can be viewed as a confidence interval of the 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.[9]

Similar to the practice of real estate agents, the Zestimate algorithm essentially estimates home value via a comparable homes ("comps") model. However, they differ in information access and processing. Real estate agents have detailed knowledge of the local market and usually identify a handful of comparable homes. They use their experience and judgment to adjust prices for comparable homes based on significant differences with the focal properties. On the other hand, the Zestimate algorithm may lack in-depth local information but it excels at processing and aggregating

[9]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 Zestimate 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 that historical time.

vast amounts of data. As a result, it often finds hundreds of comparable homes and takes the average of prices to generate Zestimate for the focal home.[10]

3.2 Data

We collected data from Zillow.com for approximately 800,000 unique residential properties in Austin, Boston, and Pittsburgh.[11] Compared with Census statistics, our sample has a very high coverage of all residential properties in these three cities. The vast majority of these properties are off-market in our sample period and only a tiny share of properties was listed for sale recently. For each property, we collect information on its web page approximately once every two weeks. In each round of data collection, we gather almost all the information available on the webpages with the exception of property photos. Our sample starts from 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 time-varying Zestimate for a given property, which is crucial for our purpose because it 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 years or even decades. Lastly, we observe property-specific characteristics such as address and home facts. In particular, we observe an extensive set of home facts including # of bedrooms, # of bathrooms, square footage, age of the property, type of property (e.g., single family house, condo, etc), detailed housing characteristics (e.g, parking, achitectural style, type of roof, cooling/heating, etc). Table A1 in the appendix lists the full list of home facts we observe and include in various regression analyses.

Our raw data of on-market properties includes 43,474 properties listed between March 26th, 2019 and December 27th, 2020. We apply the following filters to construct our main sample: (1) We remove properties with listing price equal to zero. To prevent extreme values from driving our results, we drop outliers whose listing price is below the 1st percentile or above the 99th percentile, separately for each city. This reduces the sample size to 42,118. (2) We drop outliers whose sold price is below the 1st percentile or above the 99th percentile, separately for each city. The sample

[10] See https://www.zillow.com/tech/human-and-machines-home-valuation/

[11] A 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

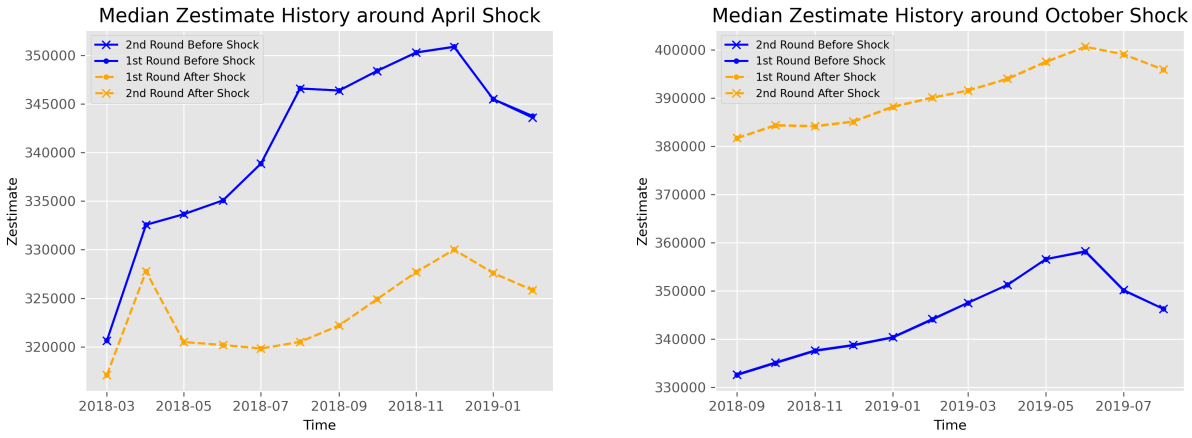
	Austin			Boston			Pittsburgh		
	N	Mean	S.D.	N	Mean	S.D.	N	Mean	S.D.
listprice ($1000)	19042	562.34	408.82	6303	830.55	522.04	13146	253.27	161.11
pending	19042	0.86	0.34	6303	0.80	0.40	13146	0.87	0.33
DaysToPending	16449	23.76	33.22	5034	38.94	43.02	11485	29.99	39.74
sold	–	–	–	6303	0.73	0.44	13146	0.79	0.41
DaysOnMarket	–	–	–	4607	74.39	34.25	10385	71.69	36.42
soldprice ($1000)	–	–	–	4607	722.17	304.21	10385	235.43	140.85
ZestLag1 ($1000)	19042	507.48	404.04	6303	782.18	478.84	13146	223.28	141.91
ZestLead1 ($1000)	17636	543.22	404.39	6232	826.56	510.65	12591	248.78	152.75

is reduced to 41,788. (3) We delete 2,674 properties for which we do not observe Zestimate right before listing. (4) Given that various later analyses require neighborhood fixed effects, we remove 35 neighborhoods with only one listed property in the entire sample period. This reduces the sample to 39,079 properties. (5) We keep properties that are single family houses, multi-family houses, townhouses, and condos. After all the above-mentioned procedures, our main sample contains 38,491 listed properties, where 19,042 are in Austin, 6,303 in Boston, and 13,146 in Pittsburgh.

Table 1 reports the summary statistics of 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. Listings in Pittsburgh are most likely to go to the pending sale stage (87%), and the probability of pending sale for Austin and Boston are 86% and 80%, respectively. Austin properties have the shortest days to pending while Boston properties take the longest time to go pending. Properties that go pending are not necessarily sold eventually, for various reasons such as buyers' financing plan does not go through. Therefore, we observe a lower probability of sale than the probability of pending sale. Days on market are defined as the days from listing date to closing date, which is by definition longer than days to pending, given that the time between pending and final closing is mainly for home inspection, appraisal, repair, loan approval, etc. Note that we have non-trivial missing data on "sold" and "DaysOnMarket" for Austin properties (because Texas does not mandate their public disclosure), so we do not report the summary statistics of these variables for Austin. For properties that are eventually sold, we observe the sold price, which is on average lower than the listing price. 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.

3.3 Algorithm updates

One important feature of our data is that it covers three major Zestimate algorithm updates in 2019. Zillow occasionally updates the Zestimate algorithm to make it more accurate. These algorithm updates include changes in the algorithm architecture and incorporation of new input features. Usually, they cannot be anticipated but can significantly change Zestimate values for most of the properties in a short period of time.



(a) (b)

Figure 3: Zestimate Algorithm Updates in 2019

The three major Zestimate algorithm updates we identified happened in April (Round 3), June (Round 8), and October (Round 17), all in 2019. The April and October updates are reflected in the changes of Zestimate history. As mentioned before, Zillow shows the Zestimate history as their best estimate of property values in the past, based on historical data and current algorithm. The values in the Zestimate history usually remain unchanged, as the historical data is fixed and the algorithm is mostly stable. However, when there are major algorithm updates, Zillow may recalculate historical Zestimates to reflect its improvement.[12] In our dataset, we have multiple snapshots of Zestimate history for each property, collected throughout the observational period. Comparing these snapshots, we find that the Zestimate history remained unchanged most of the time; significant changes for most properties, however, were observed in April and October.

The April algorithm update occurred during Round 3 of scraping. Therefore, some properties had their Round 3 data collected before the update, while others had their data collected after. To illustrate the update, Figure 3a plots the median Zestimate history for the two rounds before and

[12]See the Q&A section on https://www.zillow.com/z/zestimate/, last accessed on 5/22/2023.

two rounds after the update.[13] We can see that the historical Zestimate values changed significantly after the update. Furthermore, the historical Zestimate values stayed the same within the before and after periods, as the lines for the two rounds before and after the update overlap completely. In fact, without algorithm updates, Zestimate history values for the vast majority of properties remain the same across different rounds of data collection.

Similarly, Figure 3b shows the median Zestimate history in the four rounds around the October algorithm update, which occurred during Round 17.[14] We see the same patterns.

The June algorithm update is confirmed by Zillow's official announcement.[15] We do not observe significant changes in Zestimate history around it, probably because Zillow implemented the algorithm updates without retrospectively changing historical values.[16]

We leverage these algorithm updates to construct IVs to identify the impacts of Zestimates on market outcomes, which we explain in the next section.

4 Identifying the Human-algorithm Feedback Loop

To fix ideas, consider the following process by which the housing market unfolds for a typical transaction (Figure 4): (1) the property is off-market and has a Zestimate; (2) at some point, the seller decides to sell the home and determines a listing price with possible consultation of Zestimate; (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 makes buying decisions after seeing the listing price and the updated Zestimate (as well as touring the home and placing a bid); and (5) lastly, Zestimate updates once again after the transaction is completed and transaction outcomes are observed by Zillow.

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 listing. Admittedly, such Zestimate could also affect whether and when to list the property. We shy away from these decisions because they depend on many idiosyncratic factors we do not observe, for example the home owner may

[13] For properties whose Round 3 data was collected before the update, the four rounds from "the 2nd Round before shock" to "the 2nd Round after shock" are Round 2-5 respectively; for properties whose Round 3 data was collected after the update, these four rounds are Round 1-4 respectively.

[14] For properties whose Round 17 data was collected before the update, the four rounds from "the 2nd Round before shock" to "the 2nd Round after shock" are Round 16-19 respectively; for properties whose Round 17 data was collected after the update, these four rounds are Round 15-18 respectively.

[15] See https://www.zillow.com/tech/introducing-a-new-and-improved-zestimate-algorithm/, last visited 03/14/2023.

[16] Our main results are robust when we remove Round 8 update.

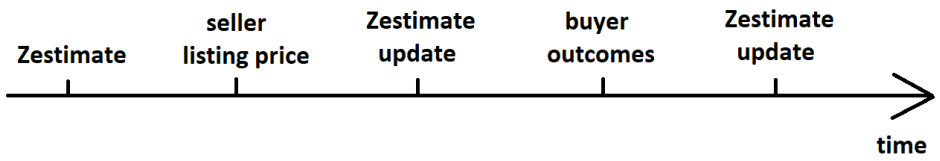


Figure 4: Sequence of Events

need to relocate for job, schooling, or other personal reasons. After listing, we study how Zestimate updates based on listing price. Then we aim to understand if home buyers are affected by the Zestimate which is already updated upon the 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

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 and relevant covariates that may affect the listing price. Specifically, we estimate:

$$ln(ListingPrice)_i = \beta ln(ZestLag1)_i + X_i\Theta + Nbhd_FE + CityRound_FE + \epsilon_i, \tag{1}$$

where *ZestLag1* is the Zestimate in the round prior to listing. X_i denotes a rich set of property characteristics as listed in Appendix Table A1. We also control for neighborhood fixed effects and city-by-round fixed effects to account for unobserved market conditions. A remaining identification challenge is that *ZestLag1* may still capture some relevant information that the seller observes but we do not, and thus an OLS regression of Equation 1 may lead to a biased estimate of β.

To identify the true β, we exploit the Zestimate algorithm updates to construct instrumental variables for *ZestLag1*. In Section 3.3, we showed that there are three major Zestimate algorithm updates in our observational period. These algorithm updates are unlikely to be anticipated by individual market participants, as Zillow does not have a fixed updating schedule and usually does not announce the major algorithm updates.[17] Meanwhile, the algorithm updates lead to abrupt and

[17]The June 2019 algorithmic update was an exception, which Zillow announced *after* its implementation. In a robustness check, we remove the June 2019 update from the analysis and find qualitatively similar results. Nevertheless, we contend that the June 2019 update is unlikely to have been anticipated by market participants, as there was

significant changes in Zestimate values (see Section 3.3). One may argue that how the market values a property or specific home characteristics keeps changing over time, and the algorithm updates could merely reflect the ongoing market shifts. However, the market-wide shifts are generally gradual and slow, especially when compared to the sudden and significant changes caused by the algorithm updates. In the immediate short-term surrounding these algorithm updates, it is reasonable to suspect that the housing market remains relatively stable in terms of valuing these observables, while the algorithm updates lead to sharp, discontinuous changes in Zestimate.

Given the above reasoning, we believe that the algorithm updates create a source of exogenous variation in Zestimate. However, the algorithm updates do not change Zestimate value randomly. Instead, they aim to improve Zestimate accuracy, which could make the Zestimate change endogenous. For example, with an algorithm update, Zestimate could start to incorporate new features generated by the analysis of property photos, such as whether a kitchen has granite countertops. As a result, the changes in Zestimate brought by the algorithm update would capture these image features, which are unobservable to us but could influence listing price. In other words, the Zestimate change brought by the algorithm updates may be correlated with ϵ_i in Equation 1. Therefore, directly taking these changes in Zestimate as exogenous variation is problematic.

To address this issue, we construct instrumental variables which capture the part of Zestimate algorithm updates that are associated with only *observable* home characteristics. In particular, we model *ZestLag1* as a function of home characteristics as well as those characteristics interacting with three dummies indicating the periods post the April, June, and October algorithm updates respectively. This way, the predicted *ZestLag1* reflects how each of the three algorithm updates triggers a different mapping from *observable* home characteristics to Zestimate, and avoids the potential influence of the unobservables on Zestimate change. The exact first-stage equation is:

$$ln(ZestLag1)_i = \lambda_0 X_i + \sum_{k=1}^{3} \lambda_k X_i Post_k + Nbhd_FE + CityRound_FE + \eta_i, \tag{2}$$

where $Post_k$ equals 1 if *ZestLag1* is queried after the k-th algorithm update; 0 otherwise.

With our IV construction, the key identifying assumption is that individual market participants do not *change* how they value the *observable* home characteristics in the same way *and* at precisely the same time as the algorithm update. Given that 1) the algorithm updates cannot be anticipated, and 2) the algorithm updates bring abrupt and significant changes in Zestimate while the housing

no prior announcement.

market remains relatively stable, we believe this assumption is reasonable. Under this assumption, the transient change in Zestimate—as a function of observable home characteristics— will only affect outcomes through Zestimate.

The two-stage least squares (2SLS) estimation results are reported in Table 2. Recall that our identification assumption is most valid in the periods close to the algorithm updates, thus our main specification limits the sample to properties listed in 2019. Column (1) reports the IV result of Equation 1. The estimated β is 0.733, meaning that a 1% increase in the lagged Zestimate leads to a 0.733% increase in the listing price, everything else held the same. Columns (2) and (3) report the results on the house and condo subsamples, respectively, where house includes single family houses, multi-family houses, and townhouses. The effect of Zestimate is positive and strong across both housing types.

Table 2: Zestimate's Effect on Listing Price (2SLS)

	(1)	(2)	(3)	(4) Short	(5) Short	(6) Short
D.V.: log(Listing Price)	All	Houses	Condos	All	Houses	Condos
log(ZestLag1)	0.733***	0.702***	0.717***	0.579***	0.562***	0.682***
	(0.043)	(0.048)	(0.047)	(0.048)	(0.053)	(0.053)
neighborhood FE	YES	YES	YES	YES	YES	YES
city-round FE	YES	YES	YES	YES	YES	YES
Observations	17,364	14,583	2,754	10,492	8,833	1,636
R-squared	0.659	0.561	0.821	0.665	0.572	0.820

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

In Columns (4), (5), and (6), we repeat the same IV regressions by further reducing the sample to two rounds before and two rounds after each algorithm update. That is, we only focus on approximately one month before each algorithm update and one month after, so that our identifying assumption is even more likely to hold. We find that the effect of Zestimate on listing price remains positive and significant, albeit the effect size is somewhat lower in the "short" sample. In Appendix Table A2, we show robust results by removing the June algorithm update (which amounts to dropping the two rounds before and the two rounds after Round 8), as Round 8 is the only algorithm update confirmed by Zillow's official announcement but does not lead to any changes in Zestimate history. We find qualitatively similar results.

We also report OLS estimation results of Equation 1 in Appendix Table A3, using the same

samples for the IV regressions. For the 2019 sample, we find that OLS estimates are smaller than their IV counterparts except for condos, suggesting that Zestimate before listing may be negatively correlated with unobservables of property value. This can happen, for example, when a property that looks bad on paper (thus low Zestimate) has undergone a renovation that is unobserved by Zillow before the property is listed. While one may argue that renovation is just one example, the reasoning extends to various types of value-adding work that sellers often do before listing their property for sale, such as improving curb appeal, fresh wall paint, minor repair, etc.[18] This is more likely to happen in the pre-listing stage, as the information available to Zillow is usually limited and potentially outdated for off-market properties. That said, when looking at the shorter sample period, we do not find significant differences in OLS and IV estimates.

4.2 Zestimate Update after Listing

By definition, Zestimate is an algorithmic estimate that is updated dynamically based on real-time information Zillow has access to. Usually, for a given property, the biggest update in Zestimate takes place immediately after the property goes on the market, when seller's agent provides a major update on the property in the local MLS. Because listing price contains seller's valuable private information about the property, it is one of the strongest signals that Zillow can leverage for an on-market property. As Zillow states in their official explanation,[19] "The Zestimate takes into account: ... On-market data such as listing price, description, comparable homes in the area and days on the market...." In this section, we provide consistent evidence for this causal chain.

Figure 5 plots the average Zestimate-to-listing-price ratio six rounds before and six rounds after listing, where we benchmark all listed properties against their own listing round (denoted by 0). We can see a clear discontinuity in the time series — the Zestimate-to-listing price ratio jumps up right after the listing price is posted. In addition, we find that the confidence interval of Zestimate shrinks substantially after listing, suggesting that the updated Zestimate follows the listing price more closely than the lagged Zestimate. There seems to be a mild increase in Zestimate-to-listing price ratio in the listing round (Round 0). 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

[18] https://www.realtor.com/advice/sell/home-selling-checklist-things-to-do-before-selling/

[19] See https://www.zillow.com/sellers-guide/influencing-your-zestimate/, last accessed 03/18/2023.

and should be considered as a mixture of Round -1 and Round 1.

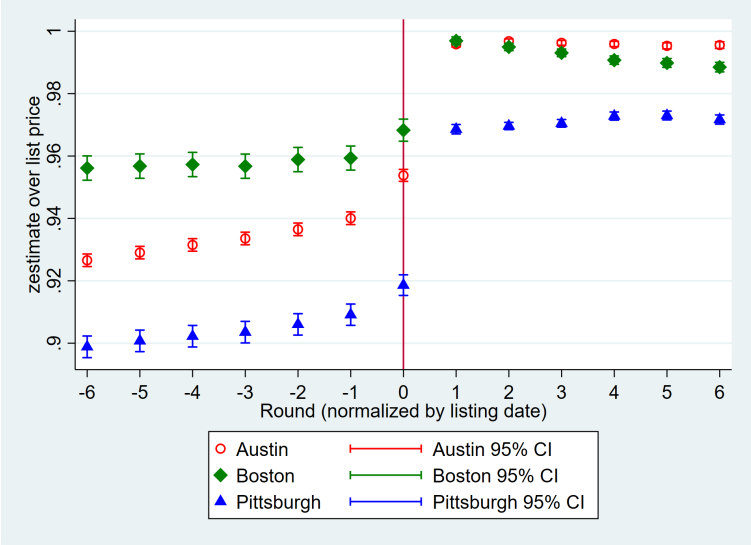


Figure 5: Zestimate Update After Listing

An alternative explanation for the patterns in Figure 5 is that Zillow and the seller may observe new market information 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 collection is bi-weekly, we instead focus on a subsample that consists of (a) the crawling was done 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, controlling for the home characteristics, as well as neighborhood and city-round fixed effects. 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 of $loglistprice \times After$.

4.3 The Effect of Zestimate on Market Outcomes

We now evaluate the relationship between Zestimate and market outcomes besides listing price. Specifically, we focus on whether the property is pending sale, logged days from listing to pending,

Table 3: Zestimate Update on Listing price

D.V.: log(Zestimate)	(1)
loglistprice	0.539***
	(0.013)
loglistprice×after	0.026**
	(0.013)
After	-0.293
	(0.196)
Observations	5,510
R-squared	0.940
neighborhood FE	YES
city-round FE	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. We control for the full list of home facts as listed in Table A1 and each of the variable interacting with "after". We do not report individual estimates due to space limit.

whether property is sold, logged days from listing to closing,[20] and logged sold price. Because we do not observe most of these outcome variables for Austin properties, we focus on Boston and Pittsburgh for this analysis. Note that once the property is listed, Zestimate gets updated based on the listing price, so we use the udpated Zestimate ($ZestLead1$) to represent the Zestimate observed by the demand side. We estimate:

$$Y_i = \beta_1 ln(ZestLead1)_i + \beta_2 ln(listprice)_i + X_i\Theta + Nbhd_FE + CityRound_FE + \epsilon_i, \quad (3)$$

where $ZestLead1_i$ is the Zestimate after listing and $listprice_i$ is the listing price of the property. X_i denotes the set of property characteristics as before. We further control for neighborhood fixed effects and city-by-round fixed effects.

Given that our primary focus is on the causal effect of *ZestLead1* on market outcomes, we instrument $ZestLead1_i$ while taking $listprice_i$ in the regression as given.[21] We adopt the same IVs as in Section 4.2, i.e., home characteristics interacting the post dummies, where the post dummies indicate whether $ZestLead1_i$ is after one of the three algorithm updates.

To the extent that eager and serious buyers check Zillow on a daily basis and stay alerted about

[20] A tiny fraction of active properties have zero days to pending or zero days to closing, so we define logged days to pending as log(days to pending+1) and logged days to closing as log(days to closing+1).

[21] If our purpose was to understand the causal effects of both *listprice* and *ZestLead1*, then we would need two separate sets of 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.

the market, the *ZestLead1* that buyers see and act upon is plausibly the Zestimate immediately after the property is listed and then updated by Zillow. However, we are not able to capture this "immediate after" *ZestLead1* for every listing, because the information for each property is collected approximately every two weeks. So for some listings, their *ZestLead1* might be too late into the process, e.g., 10 days after the property is listed. This is problematic because it is possible an outcome already affected by market participants' actions, for example, the property may have been pending since Day 3 and ZestLead1 is only observed on Day 10. To overcome this measurement problem, we apply a few filters to our sample: (1) We limit the scraping of *ZestLead1* to occur within X days of listing, where X is set to 3 for our main specification but is tested at 1, 2, 4, and 5 days for robustness. (2) We exclude cases where a property enters a pending status before *ZestLead1* is collected. For example, if a property becomes pending on Day 2 after listing and ZestLead1 is collected on Day 3, that listing is omitted. Furthermore, observations from Rounds 2 and 3 are lost due to the unavailability of data collection dates (this information is only available starting in Round 4), resulting in IV estimation for market outcomes based on two algorithm updates (R8 and R17) instead of three. These filters leave the sample to be much smaller than the one used for listing price analysis.

Table 4 reports the IV regression results. Column (1) suggests that the probability of pending sale is positively affected by Zestimate but the effect is insignificant. Column (2) shows that days to pending is negatively affected by Zestimate and the effect is also weak. As shown in Column (3) and (4), the effects of Zestimate on the probability of sale and days on market are both small and negative. They are statistically insignificant due to large standard errors. While all estimates in Column (1)-(4) are statistically insignificant because of the small sample coupled with the IV estimation and many fixed effects, they mostly have signs (except for sold) indicating that higher Zestimate leads to easier sales and shorter sale duration. In Column (5), we observe a strong, positive effect of Zestimate on the sold price. The coefficient estimate of 0.081 suggests that a 1% increase in Zestimate translates to a 0.081% higher sold price. Furthermore, the correlations between listing price and market outcomes are intuitive: listing price is negatively correlated with pending probability and sale probability, while positively correlated with days till pending, days on market, and sold price. Please note that the coefficients of listing price cannot be interpreted as causal, as there could be unobserved confounders for listing price and market outcomes.

In Columns (6)-(10), we repeat the analysis on the shorter sample (i.e., approximately one month before and one month after each algorithm update). With half of an already small sample,

Table 4: Zestimate's Effect on Market Outcomes

	(1) pending	(2) log(DaysToPending)	(3) sold	(4) log(DaysOnMarket)	(5) log(soldprice)
logZestLead1	0.122	-0.136	-0.030	-0.029	0.081**
	(0.142)	(0.447)	(0.164)	(0.183)	(0.039)
loglistprice	-0.233**	0.587*	-0.119	0.231*	0.939***
	(0.101)	(0.324)	(0.118)	(0.135)	(0.029)
nbhd FE	YES	YES	YES	YES	YES
city_round FE	YES	YES	YES	YES	YES
Observations	1,491	1,203	1,491	1,059	1,059
R-squared	0.117	0.104	0.100	0.130	0.944
	(6) Short pending	(7) Short log(DaysToPending)	(8) Short sold	(9) Short log(DaysOnMarket)	(10) Short log(soldprice)
logZestLead1	-0.094	-0.117	-0.097	-0.073	0.097**
	(0.144)	(0.465)	(0.163)	(0.213)	(0.045)
loglistprice	-0.089	0.503	-0.042	0.252	0.979***
	(0.114)	(0.383)	(0.128)	(0.180)	(0.038)
nbhd FE	YES	YES	YES	YES	YES
city_round FE	YES	YES	YES	YES	YES
Observations	787	636	787	556	556
R-squared	0.184	0.177	0.201	0.212	0.949

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

we observe even noisier estimates for pending, days to pending, sold, and days on market, compared to those of the main specification. Nonetheless, we find a strong positive effect of Zestimate on the sold price, and the size of the effect is of similar magnitude as in Column (5) (0.097 vs. 0.081). Overall, these patterns suggest that a higher Zestimate is a strong cause of a higher sold price.

We report the OLS estimates of Equation 3 in Table A4, where the OLS estimate on *log(soldprice)* is quite close to the IV estimates with overlapping standard errors. Furthermore, the estimates for variables associated with sale probability and sale duration are noisy even in OLS estimation. This noise helps explain why we observe limited significance for these variables in the IV regression.

We conduct two sets of robustness checks. First, we exclude the Round 8 algorithm update and focus only on the Round 17 update for IV construction. The results, shown in Appendix Table A5, demonstrate a decrease in the size and statistical significance of the effect on sold price, while the effect on *pending* becomes borderline significant. These findings align with the main results in Table 4 across various outcomes. Second, we test alternative caps for the number of days between

listing and the scraping date of ZestLead1. In the main analysis, the cap is set at 3 days. However, we also examine caps of 1, 2, 4, and 5 days, as presented in Table A6. The estimates indicate that pending, days to pending, sold, and days on the market remain mostly insignificant, except when the cap is set to 5 days, revealing a significant decrease in days to pending and days on the market due to the Zestimate. Notably, the estimate on sold price remains strong and stable in size across all specifications.

4.4 Zestimate Update on Market Outcomes

Zillow acknowledges that it takes into account past sales (i.e., sold price) in producing Zestimate.[22] Moreover, Zestimate is an estimate of a home's market value, which is usually measured by the final sold price. In other words, sold price is often viewed as the "ground truth" for Zestimate. Therefore, it is expected that Zestimate will closely align with the sold price once a property is sold. Figure 6 plots the average Zestimate-to-sold price ratio 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 updates upon the sold price, which is likely taken by the algorithm as ground truth to converge to and train its later versions with.

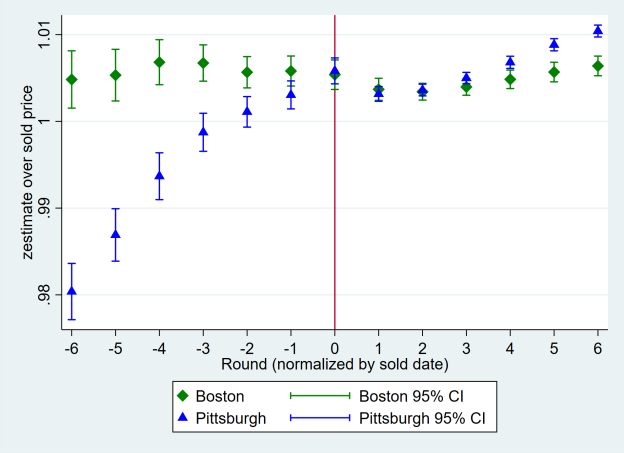


Figure 6: Zestimate Updates on Sold Price

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

[22] See https://www.zillow.com/sellers-guide/influencing-your-zestimate/, last accessed on 03/18/2023.

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 Does the Feedback Loop Cause Error Propogation?

The user-Zestimate feedback loop documented in the previous section may lead to error propagation. For example, if the initial Zestimate has a positive error, this can make the listing price higher than it would have been without the Zestimate. Observing the high listing price, the algorithm judges that the property value is indeed high, and thus generates a high updated Zestimate. This updated Zestimate leads to a higher sold price, which feeds another round of Zestimate increase, thereby locking in the initial error and propagating it to comparable properties and eventually to the entire market. In fact, these chain reactions can propagate any disturbance in any stage of the sales process, potentially creating bubbles and hampering the long-run efficiency of the housing market.

In this section, we explore the marketwide implications of the feedback loop. Quantifying these implications is challenging, due to their dynamic nature as well as many other factors that may affect the dynamics. Given this challenge, we first conduct simulations in order to understand how the interactions work together in a mechanical way. We then present empirical evidence that supports the simulation results, leveraging the onset of the COVID-19 pandemic.

5.1 Simulation

To understand how the feedback loop may propagate disturbances, we conduct a series of simulations. While we make simplifying assumptions, the simulation has several features that lend credibility to the results. First, we calibrate our simulation model with parameter estimates from the real data, so it gives a flavor of the real housing market. Second, stages in our simulation model map our empirical findings in the first part of the paper and the parameters used are based on causal identification. Lastly, the simulation can describe the evolution of market outcomes and highlight the underlying mechanisms.

We start by simulating 10,000 properties. We assume these properties are comparable to each other. Their baseline sold price is 100 and baseline listing price is 102 (i.e., on average sellers choose to start slightly higher than true market value). The baseline sold price and listing price are the

prices that would be determined in the absence of Zestimate. They can be viewed as natural market outcomes without the influence of Zestimate. We use discretized time and each period represents one week.

The initial Zestimate is set to be 100 for all the properties, i.e, Zestimate is a perfect estimate of the sold price in the first period. In each period, with probability 0.05%, an off-market property (i.e., a property that has never been listed before) is listed on market, and its listing price is the baseline listing price plus the effect of Zestimate on listing price.

If a property is listed in the previous period, then its Zestimate in the current period will be updated based on its listing price. This Zestimate update for on-market properties only happens in the first period after the property is listed. That is, Zestimate remains unchanged for all subsequent periods until the property is sold.

In each period, after potential Zestimate updates, an on-market property is sold with certain probability. We allow this probability to be higher when listing price is lower, everything else held constant. If a property is sold, its sold price is determined by the baseline sold price and the effect of Zestimate on sold price. The parameters in the simulation are calibrated with our estimated effects and statistics in our data; the detailed specifications of listing price, updated Zestimate, the probability of being sold in a period, and sold price are presented in the Appendix B.1.

How is Zestimate determined for off-market properties? The actual Zestimate algorithm mainly uses the listing price and the sold price of comparable properties. The listing price of an on-market home affects Zestimates of off-market homes starting in the period after focal property's listing. Similarly, the sold price of an on-market home affects Zestimates of off-market homes starting in the period after the focal home is sold. To align with Zestimate's methodology, we first convert the listing prices to equivalent sold prices by multiplying the ratio of the baseline sold price to the baseline listing price. This is because Zestimate uses the sold price as the anchor for determining market value. Then we take the average of all the adjusted listing prices and the actual sold prices in the previous period to be the Zestimate for all off-market properties. Once its Zestimate is updated, any off-market property can be listed and goes through the selling process as described earlier.

As mentioned above, the initial Zestimate is set to be the baseline sold price for all the properties. We then simulate market outcomes and Zestimate according to the above process for the first 20 periods. We use these periods for sanity check, as listing prices and sold prices should remain at the baseline level when the initial Zestimate is perfectly accurate. In other words, the system itself

should not create any bias or disturbance. Additionally, the first 20 periods allow us to generate a stable market environment with a certain number of on-market properties and sold properties.

We introduce a disturbance on Zestimate in the 21st period. We add a random value drawn from a normal distribution $\mathcal{N}(20, 10^2)$ to each property's Zestimate. This disturbance changes the Zestimate of all the properties and leads to a 20% increase on average. Note that in reality, disturbances in Zestimate tend to occur at much smaller scales for only a proportion of properties. Here, we intentionally set the disturbance in Zestimate to be unrealistically large and widespread as a stress test. After adding the disturbance, we continue to simulate market outcomes and subsequent Zestimates for 80 periods. We repeat the simulation for 10,000 times and calculate the average listing price and the average sold price in each period along with the 95 percent confidence interval.

Figure 7 shows the results. We see that in the first 20 periods, listing prices and sold prices indeed maintain at the baseline level, confirming that the system itself does not generate bias or disturbance. In period 21, both the average listing price and the average sold price surge, reflecting the impact of the disturbance on Zestimate. The increase in average listing price is larger than the increase in average sold price, because listing price is much more responsive to Zestimate in our estimates (0.733 vs 0.081).

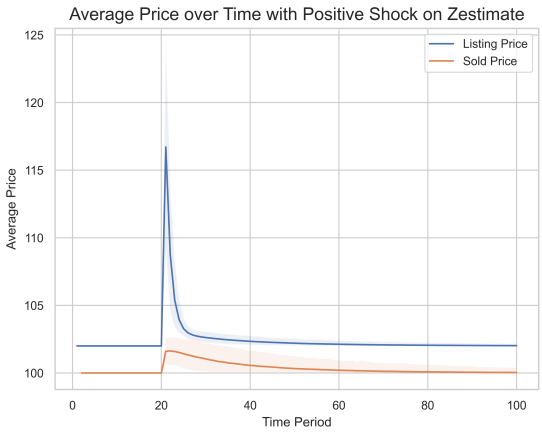


Figure 7: Simulation results for the disturbance on Zestimate

Interestingly, the positive disturbance does not persist for a long time. Instead, both listing price and sold price drop quickly and return to the baseline levels after about 60 rounds. In our simulation, listing and sold prices of on-market and recently sold properties determine the Zestimate for off-market properties. This, in turn, influences off-market properties' listing prices

and sold prices (if listed and sold). Thus, the impact of the disturbance can spread out to all other properties. However, note that the estimated coefficients of the interactive effects, albeit positive, are all less than 1. This means that the initial disturbance will reduce in strength each time it is passed to the next element (listing price, sold price, or Zestimate). As a result, the disturbance quickly fades away as market forces correct it in each stage of the selling cycle.

In the real world, another limiting factor is that Zestimate for off-market properties is calculated as an average. The effect of disturbance in one property can spread to other properties, but is diluted when averaged across multiple comparable properties. When taking the average, positive and negative disturbances on individual properties also cancel out. Therefore, only market-level disturbances, such as the 20% average increase in our initial simulation, will significantly affect subsequent Zestimate and market outcomes for some periods before they disappear eventually.[23]

Note that the shrinkage of the ripple effect is not limited to disturbances on Zestimate, but rather applies to disturbance in any stage of the selling process. In Appendix B.2, we show similar simulation results when we introduce a disturbance on sold price. Additionally, in reality, there could be multiple guardrails that further reduce the impact of disturbances. For example, we find that Zillow is more likely to remove Zestimate of a property when it finds that Zestimate is too far away from the listing price. We provide a detailed discussion on this in Appendix B.3.

These simulations demonstrate that, even with unrealistically large and widespread disturbances, the interactive effects can still limit the persistence or propagation of disturbances in a pure mechanical way. In other words, the estimated effects in human-algorithm loop are unlikely to contribute to persistent price bubbles. Next, we further validate this insight in the real data by leveraging the shock of the COVID-19 pandemic.

5.2 Empirical Evidence from COVID Shock

To empirically test the implications of the feedback loop and verify whether it will sustain or propagate errors, we first need to identify some errors. While there are many potential errors in the housing market, it is challenging to separate them from the intrinsic value of a property, as the "true value" of a property is usually unobserved to researchers.

The ideal setting is to have a replica of housing markets and a sizable shock happens in a random subset of the markets. Unfortunately, we do not have such an experimental setting. Instead, we

[23] We also tried with mean-zero disturbances in our simulations, and they did not lead to significant changes in average prices.

turn to the natural experiment created by the COVID-19 pandemic. The idea is that COVID-19 created temporary disturbances to the local housing markets in the initial weeks since the presidential declaration of national emergency.[24] 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 prices, we can measure the size of the shock to the local market. If Zestimate plays an important role in error propagation, we should observe markets subject to a positive shock at the declaration go up further afterwards, compared to markets subject to a negative shock at the declaration. If, however, the market corrects itself as our simulation suggests, such path dependence should fade away in the data.

We define a zipcode as a local market and first measure the COVID disturbance for each zipcode. This allows us to have enough number of local markets for comparison while maintaining a sufficient number of properties in each local market. We identify the properties that were listed between February 14, 2020 and March 13, 2020 and went into pending after March 13, 2020. These properties were caught by the surprise of COVID. For each of these "surprised" properties, we construct their sold price to listing price ratio (henceforth sold-to-listing ratio).

The initial COVID-19 shock to the demand side would be reflected in the sold-to-listing ratio. However, cross-sectionally, this ratio could reflect many fundamental differences across properties and zipcodes even without COVID-19, such as market thickness and bargaining power of either side of the market. Moreover, the COVID-19 shock consists of two parts. The first part is the real and persistent changes in buyer taste or market conditions. The second part is the idiosyncratic errors to the market that are independent of the real changes. It is the second part that we want to capture, as the first part will change the intrinsic value of properties instead of causing a one-time disturbance.

To separate the two parts, a natural idea is to compare the zipcode fixed effects in sold-to-listing ratio before and after the COVID-19 shock. The difference in the fixed effects should capture the average change in sold-to-listing ratio in a zipcode brought by the COVID-19 shock. However, this difference may still reflect persistent changes in buyer taste or market conditions caused by the COVID-19 (e.g., a preference for larger homes in a less densely populated area during the pandemic), which we want to exclude. Therefore, we take a more flexible approach, which accounts

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

for zipcode characteristics, to measure the COVID disturbances.

In particular, we first regress sold-to-listing ratio on property characteristics, zipcode characteristics, and month fixed effects on the sample of the properties caught in surprise by COVID:

$$SoL_{ij} = X_{ij}\Theta + C_j\Gamma + Month_FE + \epsilon_{ij}. \tag{4}$$

SoL_{ij} is the sold-to-listing ratio for property i in zipcode j; X_{ij} is property characteristics, which include the full list of home facts; C_j is zipcode characteristics, which include the average property characteristics (e.g., number of bedrooms, number of bathrooms, floor size, and age) in the zipcode, as well as population density. We take the residuals from this regression and group them at zipcode level. This zipcode level average residual, denoted as $\bar{r}_j^{\text{COVID}}$, captures the part of average sold-to-listing ratio in each zipcode after accounting for property- and zipcode-level observables.

To further separate the impact of COVID on sold-to-listing ratio from the fundamental difference in sold-to-listing ratio across zipcodes, we rerun the above regression on all the properties sold in 2019. Since we run the two regressions separately, the real and persistent changes caused by COVID are partially captured by the changes in the regression coefficients. Like before, we take the residuals and group them at zipcode level. This zipcode level average residual in 2019 is denoted as $\bar{r}_j^{2019}$. We then calculate the difference between the average residual during the COVID shock period and the average residual in 2019:

$$\Delta\bar{r}_j = \bar{r}_j^{\text{COVID}} - \bar{r}_j^{2019}. \tag{5}$$

This difference captures the change in the average sold-to-listing after accounting for property features, zipcode features, month fixed effects, and potential change in preference for features, and thus is a measure of the idiosyncratic disturbance caused by COVID in the zipcode.[25]

To ensure meaningful calculation of the average residual, we remove zipcodes whose number of sold properties in 2019 or in the COVID shock period is in the bottom 25 percentile. The mean of the change in the average residual ($\Delta\bar{r}_j$) is -0.002 and the standard deviation is 0.035. We define zipcodes with $\Delta\bar{r}_j > 0.01$ as subject to a positive shock and those with $\Delta\bar{r}_j < -0.01$ as subject to a negative shock. We drop zipcodes with a change between -0.01 and 0.01 to ensure the disturbance is large enough to detect changes in market outcomes.[26]

[25] This measure could still capture some of the persistent changes and the intrinsic features in a zipcode, but these factors should strengthen our results, as discussed in the end of this section.

[26] Here, the goal is to find disturbances that are sufficiently large to enable meaningful observation of their dynamics.

After identifying zipcodes with positive or negative idiosyncratic disturbances because of COVID-19, we turn to examine how the markets respond to these disturbances. We take a Difference-in-Differences (DiD) approach, where the shock is the COVID-19 declaration. In constructing the sample, we start with all the properties in these zipcodes, and drop those listed in 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. Properties listed between February 14, 2020 and March 13, 2020 are isolated to construct the disturbance measure, and are excluded from the DiD regression sample.

Recall that simulation suggests the disturbances would lead to changes in market outcomes, but they should diminish over time. To test the dynamics, we divide the post period into 5 subperiods, where each subperiod spans 2 months.[27] Correspondingly, we divide the 10 months before the shock into 5 subperiods.[28] We adopt the following DiD specification:

$$Y_i = \sum_{s=1}^{5} \lambda_s \; Positive_i \times Post_i^s + \sum_{s=1}^{5} \gamma_s \; Positive_i \times Pre_i^s + X_i\Theta + Nbhd_FE + CityMonth_FE + \epsilon_i, \tag{6}$$

where Y_i is the outcome variable for property i, including logged $ZestLag1$, logged listing price, logged $ZestLead1$, logged days to pending, and logged sold price. We use the dummy variable $Post_i^s$ to denote the property being listed in the s-th period after the declaration, and the dummy variable Pre_i^s to denote the property being listed in the s-th period before the declaration. The dummy variable $Positive_i$ indicates whether the property is in a zipcode subjective to positive disturbance. X_i is property characteristics. We further control for neighborhood fixed effects and city-month fixed effects. All standard errors are clustered by zipcode. In short, the baseline group is the properties in zipcodes with negative COVID disturbances. We then compare the market outcomes of properties in zipcodes with positive COVID disturbances to this baseline group.

Table 5 reports the DiD regression results. Before the COVID shock, there is minimal difference in the outcome variables between the two groups,[29] suggesting that they are comparable. If the

We also tried the analysis with smaller cutoffs (-0.005 and 0.005) or without dropping any zipcodes, and obtained results that were qualitatively similar, though with smaller effect sizes and weaker significance levels.

[27] The 5 periods are: Mach-April(post1); May-June(post2); July-August(post3); September-October(post4); November-December(post5); all in 2020.

[28] The 5 periods are: January-February 2020(pre1); November-December 2019(pre2); September-October 2019(pre3); July-August 2019(pre4); May-June 2019(pre5).

[29] With the exception of the difference in log(ZestLead) in the 4th pre-period (July-August 2019) being significant.

Table 5: COVID Shock

	(1) log(ZestLag1)	(2) log(ListingPrice)	(3) log(ZestLead1)	(4) sold	(5) log(DaysToPending)	(6) log(SoldPrice)
positive×post1	0.056*	0.016	0.073**	0.010	-0.291**	-0.015
	(0.032)	(0.032)	(0.029)	(0.030)	(0.137)	(0.029)
positive×post2	0.048**	0.043**	0.052**	-0.011	0.136	0.047**
	(0.023)	(0.018)	(0.020)	(0.027)	(0.119)	(0.017)
positive×post3	0.045*	0.049**	0.054**	-0.009	0.043	0.027
	(0.024)	(0.022)	(0.023)	(0.025)	(0.090)	(0.024)
positive×post4	0.021	0.027	0.024	-0.013	-0.005	0.017
	(0.025)	(0.025)	(0.021)	(0.024)	(0.111)	(0.023)
positive×post5	0.040	0.023	0.035	-0.044	0.027	-0.027
	(0.032)	(0.029)	(0.027)	(0.039)	(0.128)	(0.028)
positive×pre5	0.009	-0.005	-0.008	-0.024	0.068	-0.010
	(0.033)	(0.018)	(0.020)	(0.023)	(0.089)	(0.018)
positive×pre4	0.009	0.035	0.042**	-0.033	0.030	0.022
	(0.030)	(0.021)	(0.016)	(0.022)	(0.095)	(0.021)
positive×pre3	0.013	0.002	0.026	0.033	0.015	-0.024
	(0.028)	(0.024)	(0.020)	(0.030)	(0.089)	(0.021)
positive×pre2	0.023	-0.011	-0.014	0.046	-0.081	-0.012
	(0.035)	(0.025)	(0.024)	(0.037)	(0.156)	(0.023)
positive×pre1	0.029	0.001	0.009	0.019	0.046	-0.001
	(0.031)	(0.034)	(0.024)	(0.034)	(0.188)	(0.033)
Home facts	YES	YES	YES	YES	YES	YES
nbhd FE	YES	YES	YES	YES	YES	YES
city-month FE	YES	YES	YES	YES	YES	YES
Observations	12,111	13,948	13,187	13,948	11,858	10,720
R-squared	0.875	0.858	0.880	0.133	0.127	0.855

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

COVID disturbances we construct are meaningful and the demand-side disturbances do affect Zestimates and market outcomes, we should observe significant differences in the outcome variables after the COVID shock. We can see in Table 5 that the differences in all the outcomes variables except the probability of being sold become significant within two periods after the COVID shock. These results are highly consistent with our findings in Section 4 that Zestimate positively affects listing price and sold price, and has no significant impact on the probability of being sold.

If error propagation exists and drives path dependence, we shall expect the differences in Zestimate and market outcomes to grow or at least persist over time. However, the results show that

Along with the insignificant but relatively large coefficients in this period for log(listing price) and log(sold price), this may suggest some temporary change during this period.



Figure 8: The Coefficient Plots of the COVID Shock on Prices

all the differences start to shrink a while after the shock and the statistical significance disappears after four periods. We also plot the coefficients in the regressions of listing prices and sold prices in Figure 8, and the patterns are similar to the simulation results in the previous subsection. Note that for the purpose of illustration, our simulation chooses unrealistically large disturbances. Here, the COVID disturbances are much smaller in scale. In addition, many other factors could affect the market outcomes and the dynamics in reality, which our simulations cannot capture. Nonetheless, this empirical analysis leveraging the COVID shock shows similar results, which provide supporting evidence that the disturbances diminish over time and do not necessarily lead to path dependence.

It is possible that our measure of the idiosyncratic disturbances ($\Delta\bar{r}_j$) is imperfect. It might capture some of the real and persistent changes in buyer taste or market conditions caused by the COVID shock, as well as some intrinsic differences across zipcodes. However, these factors should actually strengthen our results, as the differences in outcomes driven by persistent change or intrinsic differences should not diminish over time. In other words, if our measure captures these factors, we would expect to see persistent differences, which is not what we observe.

In summary, we leverage the idiosyncratic disturbances caused by the COVID shock and evaluate how these disturbances affect the market evolution over time. Our results show that the demand-side disturbances do carry over to Zestimates and market outcomes (consistent with our findings in Section 4). However, these effects diminish over time rather than propagate or persist, in line with our simulation results in Section 5.1.

6 Exploring within-neighborhood Housing Heterogeneity

Our Covid-led natural experiment suggests that error propagation is largely absent in this context. Recall that the simulation analysis provides a straightforward explanation for this absence: the influence of Zestimate on listing and sold prices is less than one. More specifically, while Zestimate has a significant impact on listing prices, its effect on the sold price is much smaller, which places a significant constraint on error propagation. However, it is important to delve deeper into understanding what explains the limited marginal effects in the first place. In this aim, we study the heterogeneous effects of Zestimate on the listing price and the sold price, across neighborhoods of varying uniqueness of housing characteristics. This approach serves two crucial purposes. First, it enables us to gain a deeper understanding of the absence of error propagation beyond the mechanical factors explained in our simulation analysis. Second, it bolsters the causal arguments regarding the feedback loop with intuitive and consistent evidence.

For individual sellers and buyers in the housing market, estimating the market value of a home involves acquiring and processing a substantial amount of relevant information. It is arguably easier to estimate home values in neighborhoods with more homogeneous homes, such as a large development of similar properties, than in neighborhoods with more heterogeneous homes, where each home is unique and finding comparable homes is challenging. If neighborhoods with heterogeneous homes create greater information friction for sellers and buyers, they may place a higher value on Zestimate's accuracy in these areas. However, Zestimate may also be less accurate in these neighborhoods due to the same reason, which implies that Zestimate may have more incentive to closely follow human decisions in these markets. Given this, we may expect a stronger human-algorithm feedback loop in more heterogeneous neighborhoods. However, it is also possible that buyers in heterogeneous neighborhoods rationally account for the lower Zestimate accuracy and the stronger Zestimate-listing price interdependence, and therefore trust Zestimate less, which in turn may make the feedback looser in these neighborhoods.

To test these hypotheses apart, we constructed a measure of housing heterogeneity at the neighborhood level by the variation in key housing characteristics, including the number of bedrooms, number of bathrooms, floor size, property age, and Zestimate.[30] This measure provides insights into the level of housing heterogeneity in a given neighborhood. Given that this measure is pos-

[30]Specifically, for each neighborhood, we compute the standard deviation and the mean of number of beds, number of baths, floor size, property age, and Zestimate. We then compute the coefficient of variation (CV, i.e., the standard deviation divided by mean) at the neighborhood level for each of these factors. We then get the average across these CV's so different neighborhoods can be comparable.

itively skewed,[31] we take its logarithm and interact it with key variables in various stages of the home selling process.

Our analysis shows that, in the listing stage, the listing price is more influenced by Zestimate in more heterogeneous neighborhoods, as demonstrated by the positive coefficient estimate of the interaction term in Table A7. Furthermore, after a property is listed, we observe that Zestimate is more correlated with the listing price and less correlated with its previous version (ZestLag1) in more heterogeneous neighborhoods, as indicated in our results in Table A8. In line with our analysis at other stages, we examined how buyers' responses to Zestimate are affected by housing heterogeneity. Our findings in Table A9 suggest that buyers are influenced by Zestimate to a lesser extent in more heterogeneous neighborhoods, as indicated by the negative coefficient estimate of the interaction term for the sold price. We also ran simulations using the estimated effects in neighborhoods with low (5 percentile) and high (95 percentile) heterogeneity. The results show similar patterns of dissipating disturbance as before. In high-heterogeneity neighborhoods, the convergence to baseline prices is faster due to the significantly lower reliance on Zestimate in sold price.

One concern is that neighborhood housing heterogeneity might be associated with other factors, potentially undermining the interpretation of the above results as reflecting heterogeneous effects of housing heterogeneity. For instance, a confounding variable could be neighborhood housing prices, as richer neighborhoods might have more unique and less comparable homes. To address this concern, we divided neighborhoods into deciles based on their average Zestimate in 2019 for all properties (listed or not), separately for each city. We found no significant correlation between neighborhood housing diversity and price levels. Additionally, we controlled for neighborhood price levels interacted with key variables into the regression analysis, and the interaction with neighborhood housing heterogeneity remained statistically significant. The results are reported in Table A10, A11, and A12.

These findings collectively imply that in more heterogeneous neighborhoods: (1) sellers may place greater trust in Zestimate, (2) Zestimate aligns more closely with sellers' listing prices, but (3) buyers tend to rely less on Zestimate. These dynamics indicate that sellers in heterogeneous neighborhoods rely more on Zestimate, which strengthens the human-algorithm feedback loop. However, buyers in these neighborhoods trust Zestimate less, limiting the strength of the feedback

[31] The log of this neighborhood-specific scalar has a mean of 0.036 and a standard deviation of 1.029, with a 10th percentile at -0.965 and 90th percentile at 1.652.

loop. The lower reliance on Zestimate in sold prices in heterogeneous neighborhoods can have multiple explanations. For example, buyers in these neighborhoods may put more efforts in searching offline, and therefore perceive Zestimate as an inferior signal compared to other sources of information. Alternatively, heterogeneous neighborhoods may attract buyers with more diverse tastes who naturally assign less weight to Zestimate. Moreover, buyers in these neighborhoods may be aware that sellers rely more on Zestimate and that Zestimate more closely tracks listing prices, leading them to place less trust in it. Regardless of the reasons behind buyers' lower reliance on Zestimate in more heterogeneous neighborhoods, this observation translates into a counterforce to the stronger human-algorithm interaction in the earlier stages.

In summary, our findings suggest that sellers and buyers use various information to guide their decision-making, which acts as a safeguard to the feedback loop and contributes to the market's self-correction.

7 Conclusion

Using a dataset that tracks market outcomes and Zestimates over time for properties in three U.S. cities, we find a strong co-movement between Zestimate and market outcomes in the residential housing market. Specifically, Zestimate leads the market by affecting the decision making of both sellers and buyers; it also follows the market by tracking listing price and market outcomes closely. The sizable effects in this algorithmic feedback loop raise a concern that 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.

Our simulation suggests that the algorithm does not cause errors to persist or spread out, as the disturbance's magnitude weakens in each stage as it ripples through the home sales cycle, eventually dissipating. The analysis that leverages a natural experiment created by the COVID-19 shock provides consistent evidence, as we see the gap in the market outcomes between properties in the zipcodes subject to positive disturbances and those in the zipcodes subject to negative disturbances all decrease over time and disappear eventually.

As ML algorithms becomes more prevalent in our lives, human-algorithm interaction is intensifying. Algorithms take real world data as inputs, and when human decision is influenced by algorithm outputs, it can create a self-reinforcing cycle. However, market forces and algorithm designs can also limit the propagation of disturbances in the feedback loop. It is therefore important

to consider the context of algorithm use when designing and evaluating algorithms. We hope our work would encourage more research in this fast-growing area.

Our results offer a few policy implications. First, human-algorithm interactions exist in not only online content consumption and social media but also in traditional high-stake markets which calls for more research. Second, our results suggest that Zestimate's ability to manipulate the market seems low. Third, part of our analysis is limited by the fact that many Austin properties miss sold price and other sold outcomes because Texas does not mandate their disclosure to the public. Algorithm creators and the general public may benefit from more data sharing, so that the benefits of the algorithm can be improved and any disturbance propagation, if persistent, can be better studied and prevented.

Our research is subject to a number of 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 collected 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 that we ignore the potential impact of Zestimate on the selection of properties going on the market. Finally, while human-algorithm feedback loop is common in many other settings, such as e-commerce and social media, 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 human-algorithm feedback loops in other markets.

References

Franklin Allen and Douglas Gale. Financial contagion. *Journal of political economy*, 108(1):1–33, 2000.

Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias. propublica (2016). *URL: https://www. propublica. org/article/machine-bias-risk-asses sments-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.

Christopher A Bail, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. Exposure

to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221, 2018.

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.

Levi Boxell, Matthew Gentzkow, and Jesse M Shapiro. Greater internet use is not associated with faster growth in political polarization among us demographic groups. *Proceedings of the National Academy of Sciences*, 114(40):10612–10617, 2017.

Zach Y Brown and Alexander MacKay. Competition in pricing algorithms. Technical report, National Bureau of Economic Research, 2021.

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.

Sophie Calder-Wang and Gi Heung Kim. Coordinated vs efficient prices: The impact of algorithmic pricing on multifamily rental markets. *Available at SSRN*, 2023.

Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–97, 2020.

Guangying Chen, Tat Chan, Dennis Zhang, Senmao Liu, and Yuxiang Wu. The effects of diversity in algorithmic recommendations on digital content consumption: A field experiment. *Available at SSRN 4365121*, 2023.

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.

Bo Cowgill and Catherine E Tucker. Economics, fairness and algorithmic bias. *preparation for: Journal of Economic Perspectives*, 2019.

Rafael Di Tella, Ramiro H Gálvez, and Ernesto Schargrodsky. Does social media cause polarization? evidence from access to twitter echo chambers during the 2019 argentine presidential debate. Technical report, National Bureau of Economic Research, 2021.

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.

Ben Green and Yiling Chen. The principles and limits of algorithm-in-the-loop decision making. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–24, 2019.

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.

David Holtz, Ben Carterette, Praveen Chandar, Zahra Nazari, Henriette Cramer, and Sinan Aral. The engagement-diversity connection: Evidence from a field experiment on spotify. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 75–76, 2020.

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.

Karl Krauth, Yixin Wang, and Michael I Jordan. Breaking feedback loops in recommender systems with causal inference. *arXiv preprint arXiv:2207.01616*, 2022.

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.

Dokyun Lee and Kartik Hosanagar. How do recommender systems affect sales diversity? a cross-category investigation via randomized field experiment. *Information Systems Research*, 30(1): 239–259, 2019.

Ro'ee Levy. Social media, news consumption, and polarization: Evidence from a field experiment. *American economic review*, 111(3):831–870, 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.

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.

Alex Moehring. News feeds and user engagement: Evidence from the reddit news tab. *Available at SSRN 4193945*, 2022.

Jean-Charles Rochet and Jean Tirole. Interbank lending and systemic risk. *Journal of Money, credit and Banking*, 28(4):733–762, 1996.

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.

Jonathan Stray. Designing recommender systems to depolarize. *arXiv preprint arXiv:2107.04953*, 2021.

Wenlong Sun, Sami Khenissi, Olfa Nasraoui, and Patrick Shafto. Debiasing the human-recommender system feedback loop in collaborative filtering. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 645–651, 2019.

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.

Shunyuan 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 Tables

Table A1: Detailed Home Characteristics

Variable	Note
Number of bedrooms	
Number of full bathrooms	
Number of half bathrooms	
Total number of rooms	
Total number of units	
Number of stories	
Solar potential	
Has video	
Number of photos	
Age	year listed - year constructed
Floor area (sqft)	
Lot size (acres)	
Has virtual tour	
Property type	single family, multi family, townhouse, condo
Roof type	builtup, composition, metal, shingle, slate, tile, other
Parking type	attached, carport, detached, garage, street, none, other
Flooring	carpet, wood, tile, laminate, concrete, vinyl, stone, other
Heating	gas, forced air, electric, heat pump, radiant, baseboard, none, other
Cooling	central, electric, wall, AC, none, other
Patio type	deck, patio, porch
Other interior features	fireplace, ceiling fan, storage, vaulted ceiling, furnished
Exterior material	composition, cement concrete, brick, stucco, stone, vinyl, shingle, wood products, wood, metal, other
Architectural Style	bungalow, cape cod, colonial, contemporary, craftsman, detached, french, Georgian, house, loft, modern, ranch, Spanish, split level, stilt home, Tudor, victorian, lowrise, midrise, highrise, undefined

Table A2: Zestimate's Effect on Listing Price — Assume No Round 8 Algorithm Update

	(1)	(2)	(3)	(4)	(5)	(6)
				Short	Short	Short
D.V.: log(Listing Price)	All	Houses	Condos	All	Houses	Condos
logZestLag1	0.741***	0.705***	0.718***	0.579***	0.561***	0.718***
	(0.049)	(0.057)	(0.059)	(0.061)	(0.064)	(0.070)
Observations	17,364	14,583	2,754	6,428	5,355	1,043
R-squared	0.658	0.560	0.821	0.677	0.582	0.819
nbhd FE	YES	YES	YES	YES	YES	YES
city-round FE	YES	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A3: Zestimate's Effect on Listing Price — OLS results

D.V.: log(Listing Price)	(1) All	(2) Houses	(3) Condos	(4) Short All	(5) Short Houses	(6) Short Condos
logZestLag1	0.639***	0.585***	0.775***	0.641***	0.588***	0.721***
	(0.007)	(0.008)	(0.016)	(0.009)	(0.011)	(0.022)
Observations	17,364	14,583	2,754	10,492	8,833	1,636
R-squared	0.899	0.892	0.953	0.898	0.892	0.952
neighborhood FE	YES	YES	YES	YES	YES	YES
city-round FE	YES	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A4: Zestimate and Market Outcomes — OLS Estimates

	(1) pending	(2) log(DaysToPending)	(3) sold	(4) log(DaysOnMarket)	(5) log(soldprice)
logZestLead1	0.094*	-0.355*	0.096*	-0.081	0.078***
	(0.049)	(0.183)	(0.057)	(0.080)	(0.017)
loglistprice	-0.215***	0.728***	-0.202***	0.265***	0.940***
	(0.051)	(0.191)	(0.059)	(0.085)	(0.018)
nbhd FE	YES	YES	YES	YES	YES
city_round FE	YES	YES	YES	YES	YES
Observations	1,491	1,203	1,491	1,059	1,059
R-squared	0.247	0.226	0.234	0.237	0.987
	(6) Short pending	(7) Short log(DaysToPending)	(8) Short sold	(9) Short log(DaysOnMarket)	(10) Short log(soldprice)
logZestLead1	0.096	-0.190	0.104	0.056	0.102***
	(0.072)	(0.259)	(0.081)	(0.119)	(0.025)
loglistprice	-0.215***	0.551*	-0.176**	0.165	0.976***
	(0.077)	(0.283)	(0.087)	(0.135)	(0.028)
nbhd FE	YES	YES	YES	YES	YES
city_round FE	YES	YES	YES	YES	YES
Observations	787	636	787	556	556
R-squared	0.327	0.341	0.356	0.342	0.989

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A5: Zestimate's Effect on market outcomes — Assume No Round 8 Algorithm Update

	(1) pending	(2) log(DaysToPending)	(3) sold	(4) log(DaysOnMarket)	(5) log(soldprice)
logZestLead1	0.389*	-0.146	0.290	-0.050	0.054
	(0.219)	(0.743)	(0.252)	(0.316)	(0.068)
loglistprice	-0.409***	0.594	-0.329*	0.245	0.956***
	(0.150)	(0.501)	(0.172)	(0.214)	(0.046)
Observations	1,491	1,203	1,491	1,059	1,059
R-squared	0.091	0.105	0.095	0.130	0.944
neighborhood FE	YES	YES	YES	YES	YES
city-round FE	YES	YES	YES	YES	YES

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A6: Zestimate's Effect on Market Outcomes
— Alternative Caps on Days between Listing and Scraping Date of ZestLead1

				1 Day	
	(1)	(2)	(3)	(4)	(5)
	pending	log(DaysToPending)	sold	log(DaysOnMarket)	log(soldprice)
logZestLead1	0.116	-0.229	0.139	-0.213	0.087**
	(0.114)	(0.409)	(0.132)	(0.184)	(0.041)
loglistprice	-0.227**	0.574*	-0.194*	0.318**	0.942***
	(0.090)	(0.344)	(0.104)	(0.157)	(0.035)
Observations	526	423	526	381	381
R-squared	0.285	0.252	0.259	0.292	0.949
				2 Days	
	(6)	(7)	(8)	(9)	(10)
	pending	log(DaysToPending)	sold	log(DaysOnMarket)	log(soldprice)
logZestLead1	0.097	-0.231	-0.035	-0.116	0.078**
	(0.124)	(0.425)	(0.144)	(0.172)	(0.037)
loglistprice	-0.200**	0.774**	-0.110	0.348***	0.944***
	(0.089)	(0.318)	(0.104)	(0.133)	(0.029)
Observations	1,015	824	1,015	726	726
R-squared	0.133	0.137	0.128	0.174	0.941
				4 Days	
	(6)	(7)	(8)	(9)	(10)
	pending	log(DaysToPending)	sold	log(DaysOnMarket)	log(soldprice)
logZestLead1	0.270*	-0.636	0.120	-0.276	0.111**
	(0.158)	(0.499)	(0.180)	(0.206)	(0.044)
loglistprice	-0.346***	0.975***	-0.240*	0.386**	0.915***
	(0.114)	(0.364)	(0.130)	(0.152)	(0.033)
Observations	1,970	1,586	1,970	1,407	1,407
R-squared	0.089	0.091	0.094	0.101	0.945
				5 Days	
	(6)	(7)	(8)	(9)	(10)
	pending	log(DaysToPending)	sold	log(DaysOnMarket)	log(soldprice)
logZestLead1	0.050	-1.375***	-0.005	-0.457**	0.181***
	(0.163)	(0.501)	(0.186)	(0.213)	(0.046)
loglistprice	-0.208*	1.511***	-0.163	0.530***	0.860***
	(0.120)	(0.368)	(0.137)	(0.156)	(0.034)
Observations	2,395	1,921	2,395	1,697	1,697
R-squared	0.086	0.052	0.082	0.079	0.944

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include neighborhood fixed effects, city-by-round fixed effects, and the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A7: Zestimate's Heterogeneous Effects on Listing Price (2SLS) — Neighborhood Heterogeneity

	(1)	(2)
	Full sample	Short sample
D.V.: log(Listing Price)		
log(ZestLag1)	0.760***	0.708***
	(0.028)	(0.030)
log(ZestLag1)×log(het)	0.016**	0.019**
	(0.007)	(0.009)
neighborhood FE	YES	YES
city-round FE	YES	YES
Observations	17,364	10,492
R-squared	0.657	0.665

Notes: We use "het" to denote the housing heterogeneity at the neighborhood level. Specifically, for each neighborhood, we compute the standard deviation and the mean of number of beds, number of baths, floor size, property age, and Zestimate. We then compute the coefficient of variation (CV, i.e., the standard deviation divided by mean) at the neighborhood level for each of these factors. We then get the average across these CV's so different neighborhoods can be comparable. The variable "het" is positively skewed, and the log of "het" has a mean of 0.036 and a standard deviation of 1.029, with a 10th percentile at -0.965 and 90th percentile at 1.652. We use "PriceLevel" to denote the housing price level of each neighborhood, measured by the neighborhood average Zestimate of all properties (listed or not) in 2019, which is then converted into deciles by each city. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A8: Zestimate Tracks Listing Price — Neighborhood Heterogeneity

	(1)	(2)
D.V.: log(ZestLead1)		
log(listprice)	0.631***	0.636***
	(0.005)	(0.005)
log(listprice)×log(het)		0.026***
		(0.004)
log(ZestLag1)	0.339***	0.336***
	(0.005)	(0.005)
log(ZestLag1)×log(het)		-0.020***
		(0.005)
neighborhood FE	YES	YES
city-round FE	YES	YES
Observations	13,995	13,995
R-squared	0.975	0.975

Notes: We use "het" to denote the housing heterogeneity at the neighborhood level. Specifically, for each neighborhood, we compute the standard deviation and the mean of number of beds, number of baths, floor size, property age, and Zestimate. We then compute the coefficient of variation (CV, i.e., the standard deviation divided by mean) at the neighborhood level for each of these factors. We then get the average across these CV's so different neighborhoods can be comparable. The variable "het" is positively skewed, and the log of "het" has a mean of 0.036 and a standard deviation of 1.029, with a 10th percentile at -0.965 and 90th percentile at 1.652. We use "PriceLevel" to denote the housing price level of each neighborhood, measured by the neighborhood average Zestimate of all properties (listed or not) in 2019, which is then converted into deciles by each city. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A9: Zestimate's Heterogeneous Effect on Market Outcomes — Neighborhood Heterogeneity

	(1) pending	(2) log(DaysToPending)	(3) sold	(4) log(DaysOnMarket)	(5) log(soldprice)
log(ZestLead1)	0.178*	-0.150	0.104	0.035	0.079**
	(0.095)	(0.333)	(0.110)	(0.145)	(0.031)
log(ZestLead1)×log(het)	-0.014	0.082	-0.045	-0.039	-0.020**
	(0.025)	(0.095)	(0.029)	(0.044)	(0.009)
log(listprice)	-0.268***	0.595**	-0.200**	0.191*	0.940***
	(0.075)	(0.262)	(0.086)	(0.115)	(0.025)
nbhd FE	YES	YES	YES	YES	YES
city_round FE	YES	YES	YES	YES	YES
Observations	1,491	1,203	1,491	1,059	1,059
R-squared	0.115	0.107	0.105	0.127	0.944
	(6) Short pending	(7) Short log(DaysToPending)	(8) Short sold	(9) Short log(DaysOnMarket)	(10) Short log(soldprice)
log(ZestLead1)	0.032	-0.151	0.014	-0.146	0.084**
	(0.109)	(0.370)	(0.123)	(0.167)	(0.035)
log(ZestLead1)×log(het)	0.002	0.099	-0.052	-0.009	-0.036**
	(0.033)	(0.130)	(0.037)	(0.067)	(0.014)
log(listprice)	-0.173*	0.525	-0.108	0.301*	0.987***
	(0.095)	(0.333)	(0.107)	(0.156)	(0.033)
nbhd FE	YES	YES	YES	YES	YES
city_round FE	YES	YES	YES	YES	YES
Observations	787	636	787	556	556
R-squared	0.193	0.180	0.211	0.209	0.950

Notes: We use "het" to denote the housing heterogeneity at the neighborhood level. Specifically, for each neighborhood, we compute the standard deviation and the mean of number of beds, number of baths, floor size, property age, and Zestimate. We then compute the coefficient of variation (CV, i.e., the standard deviation divided by mean) at the neighborhood level for each of these factors. We then get the average across these CV's so different neighborhoods can be comparable. The variable "het" is positively skewed, and the log of "het" has a mean of 0.036 and a standard deviation of 1.029, with a 10th percentile at -0.965 and 90th percentile at 1.652. We use "PriceLevel" to denote the housing price level of each neighborhood, measured by the neighborhood average Zestimate of all properties (listed or not) in 2019, which is then converted into deciles by each city. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A10: Zestimate's Heterogeneous Effects on Listing Price (2SLS) — Neighborhood Heterogeneity vs Neighborhood Housing Price

D.V.: log(Listing Price)	(1) Full sample	(2) Short sample
log(ZestLag1)	0.660***	0.653***
	(0.033)	(0.038)
log(ZestLag1)×log(het)	0.017**	0.019**
	(0.007)	(0.009)
log(ZestLag1)×PriceLevel	0.020***	0.014***
	(0.003)	(0.004)
neighborhood FE	YES	YES
city-round FE	YES	YES
Observations	17,364	10,492
R-squared	0.659	0.666

Notes: We use "het" to denote the housing heterogeneity at the neighborhood level. Specifically, for each neighborhood, we compute the standard deviation and the mean of number of beds, number of baths, floor size, property age, and Zestimate. We then compute the coefficient of variation (CV, i.e., the standard deviation divided by mean) at the neighborhood level for each of these factors. We then get the average across these CV's so different neighborhoods can be comparable. The variable "het" is positively skewed, and the log of "het" has a mean of 0.036 and a standard deviation of 1.029, with a 10th percentile at -0.965 and 90th percentile at 1.652. We use "PriceLevel" to denote the housing price level of each neighborhood, measured by the neighborhood average Zestimate of all properties (listed or not) in 2019, which is then converted into deciles by each city. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A11: Zestimate Tracks Listing Price — Neighborhood Heterogeneity vs Neighborhood Housing Price

D.V.: log(ZestLead1)	
log(listprice)	0.647***
	(0.011)
log(listprice)×log(het)	0.027***
	(0.004)
log(listprice)×PriceLevel	-0.002
	(0.002)
log(ZestLag1)	0.331***
	(0.011)
log(ZestLag1)×log(het)	-0.020***
	(0.005)
log(ZestLag1)×PriceLevel	0.001
	(0.002)
neighborhood FE	YES
city-round FE	YES
Observations	13,995
R-squared	0.975

Notes: We use "het" to denote the housing heterogeneity at the neighborhood level. Specifically, for each neighborhood, we compute the standard deviation and the mean of number of beds, number of baths, floor size, property age, and Zestimate. We then compute the coefficient of variation (CV, i.e., the standard deviation divided by mean) at the neighborhood level for each of these factors. We then get the average across these CV's so different neighborhoods can be comparable. The variable "het" is positively skewed, and the log of "het" has a mean of 0.036 and a standard deviation of 1.029, with a 10th percentile at -0.965 and 90th percentile at 1.652. We use "PriceLevel" to denote the housing price level of each neighborhood, measured by the neighborhood average Zestimate of all properties (listed or not) in 2019, which is then converted into deciles by each city. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

Table A12: Zestimate's Heterogeneous Effect on Market Outcomes — Neighborhood Heterogeneity vs Neighborhood Housing Price

	(1) pending	(2) log(DaysToPending)	(3) sold	(4) log(DaysOnMarket)	(5) log(soldprice)
log(ZestLead1)	0.237**	-0.057	0.246*	-0.142	0.045
	(0.112)	(0.411)	(0.129)	(0.183)	(0.039)
log(ZestLead1)×log(het)	-0.011	0.137	-0.044	-0.038	-0.024***
	(0.025)	(0.095)	(0.029)	(0.043)	(0.009)
log(ZestLead1)×PriceLevel	-0.007	-0.052	-0.010	0.010	0.008*
	(0.014)	(0.050)	(0.016)	(0.022)	(0.005)
log(listprice)	-0.279***	0.733***	-0.253***	0.264**	0.928***
	(0.067)	(0.244)	(0.078)	(0.106)	(0.023)
nbhd FE	YES	YES	YES	YES	YES
city_round FE	YES	YES	YES	YES	YES
Observations	1,491	1,203	1,491	1,059	1,059
R-squared	0.114	0.109	0.104	0.130	0.945
	(6) Short pending	(7) Short log(DaysToPending)	(8) Short sold	(9) Short log(DaysOnMarket)	(10) Short log(soldprice)
log(ZestLead1)	0.067	-0.455	0.235	-0.482**	0.034
	(0.142)	(0.504)	(0.160)	(0.231)	(0.048)
log(ZestLead1)×log(het)	-0.021	0.102	-0.065*	-0.042	-0.042***
	(0.033)	(0.133)	(0.037)	(0.067)	(0.014)
log(ZestLead1)×PriceLevel	0.004	0.024	-0.026	0.079**	0.009
	(0.019)	(0.070)	(0.021)	(0.032)	(0.007)
log(listprice)	-0.210**	0.630**	-0.145	0.193	0.982***
	(0.089)	(0.310)	(0.100)	(0.144)	(0.030)
nbhd FE	YES	YES	YES	YES	YES
city_round FE	YES	YES	YES	YES	YES
Observations	787	636	787	556	556
R-squared	0.193	0.181	0.215	0.226	0.950

Notes: We use "het" to denote the housing heterogeneity at the neighborhood level. Specifically, for each neighborhood, we compute the standard deviation and the mean of number of beds, number of baths, floor size, property age, and Zestimate. We then compute the coefficient of variation (CV, i.e., the standard deviation divided by mean) at the neighborhood level for each of these factors. We then get the average across these CV's so different neighborhoods can be comparable. The variable "het" is positively skewed, and the log of "het" has a mean of 0.036 and a standard deviation of 1.029, with a 10th percentile at -0.965 and 90th percentile at 1.652. We use "PriceLevel" to denote the housing price level of each neighborhood, measured by the neighborhood average Zestimate of all properties (listed or not) in 2019, which is then converted into deciles by each city. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Across all specifications, we include the full list of home facts as listed in Table A1 to the regressions without reporting those estimates due to space limit.

B Simulation

B.1 Setup Specifications

In our simulations, the initial listing price of a property is the baseline listing price plus the effect of Zestimate on listing price:

$$ListingPrice_i = BaseListingPrice + \frac{\mathrm{d}ListingPrice}{\mathrm{d}ZestLag1} \cdot (Zest_{it} - BaseSoldPrice) \tag{7}$$

where $Zest_{it}$ is the Zestimate of the property i in period t (the listing period), and $\frac{\mathrm{d}ListingPrice}{\mathrm{d}ZestLag1} = 0.733$ based on our previous estimation.

If a property is listed in the previous period, then its Zestimate in the current period will be updated based on its listing price:

$$Zest_{it} = Zest_{i(t-1)} + \frac{\mathrm{d}Zest}{\mathrm{d}ListingPrice} \cdot (ListingPrice_i - BaseListingPrice), \tag{8}$$

where $Zest_{i(t-1)}$ is the Zestimate of property i in period $t-1$, and $\frac{\mathrm{d}Zest}{\mathrm{d}ListingPrice} = 0.565$ based on our previous estimation. For subsequent periods when the property is on market, Zestimate remains unchanged:

$$Zest_{it} = Zest_{i(t-1)} \tag{9}$$

The probability of an on-market property being sold in a period is

$$\Pr(sold)_{it} = \frac{0.11}{1 + \exp(0.1 \cdot (ListingPrice_i - BaseListingPrice))}. \tag{10}$$

This means that a property is easier to sell when listing price is lower, everything else held constant.

If a property is sold, its sold price is determined by the baseline sold price and the effect of Zestimate on sold price:

$$SoldPrice_i = BaseSoldPrice + \frac{\mathrm{d}SoldPrice}{\mathrm{d}ZestLead1} \cdot (Zest_{it} - BaseSoldPrice), \tag{11}$$

where $\frac{\mathrm{d}SoldPrice}{\mathrm{d}ZestLead1} = 0.081$ based on our previous estimation.

B.2 The Case of Disturbance in Sold Price

To show that the shrinkage of the ripple effect is not limited to disturbances on Zestimate, we simulate a disturbance on sold price. Similar to the case presented in the main paper, we introduce the disturbance in the 21st period, but instead of adding a shock to Zestimate, we add a random value drawn from $\mathcal{N}(20, 10^2)$ to the sold price of all the properties sold in that period. After adding the disturbance on the sold prices, we continue to simulate market outcomes and subsequent Zestimates for 80 periods. We again repeat the simulation for 10,000 times, and Figure A1 shows the average listing price and the average sold price along with the 95 percent confidence interval.

The average sold price increases by 20% in period 21, reflecting the disturbance on sold price. This directly affects the Zestimate of off-market properties, resulting in a surge in listing prices in

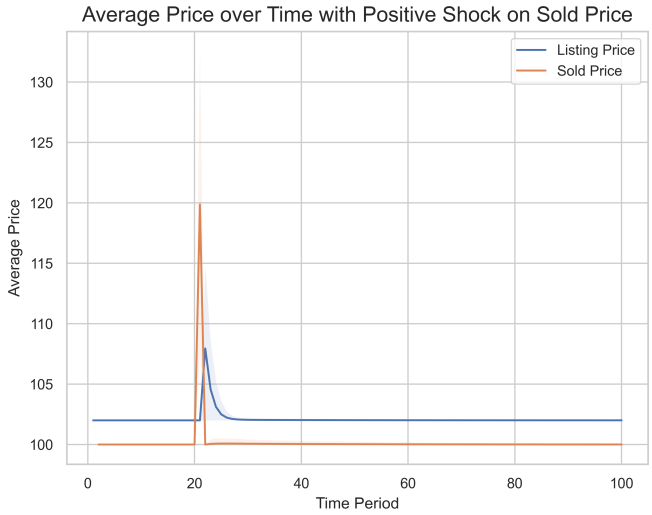


Figure A1: Simulation results for the disturbance on sold price

period 22. The increase in average listing price is smaller than the increase in average sold price because the Zestimate of off-market properties also depends on listing prices in period 21, which are not affected by the disturbance. This means that the impact of the sold price disturbance is reduced when passed to the Zestimate. Furthermore, the disturbance diminishes when passed from Zetsimate to listing price.

The sold price in period 22 returns to the baseline level, as properties sold during this period are not yet affected by the disturbance. Subsequent periods see the disturbance passed through Zestimate and listing prices, affecting sold prices. After the initial increase, both average listing price and average sold price quickly drop and return to baseline levels after 20 periods. Again, the disturbance on sold price does not persist or propagate for a long time.

B.3 Potential Guardrails

There are many guardrails in reality that could further reduce the impact of disturbances. For example, in our simulations, the Zestimate of off-market properties is based on only the listing prices and sold prices in the previous period. The real algorithm usually tracks listing prices and sold prices over a relatively long period (e.g., the past 6 months). Therefore, Zestimate is less affected by a short-term disturbance.

Another potential guardrail is Zillow's practice of hiding Zestimte when it is at odds with the listing price. For each property in our dataset, we compute the average Zestimate in the six rounds before listing and compare it with the listing price. If the average Zestimate is considerably lower than the listing price, then this property is considered to be "Zest Low." We define "Zest Medium" and "Zest High" similarly. Figure A2 shows that Zestimate, in general, becomes less available for "Zest Low" and "Zest High" than "Zest Medium", suggesting that Zillow is more likely to hide the Zestimate when they find it less agreeable with the listing price.

To illustrate the limiting effect of guardrails, we run another simulation, identical to the first case (i.e., introducing disturbance on Zestimate) except that the Zestimate for an on-market property will be removed if the absolute value of its difference from the listing price exceeded 20. Once the Zestimate is removed, it has no impact on sold price. We compare the market outcomes in this

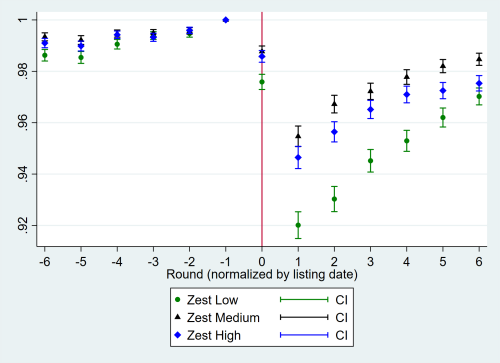


Figure A2: Zestimate Availability

simulation to the first case (as the baseline). Figure A3 shows the results.

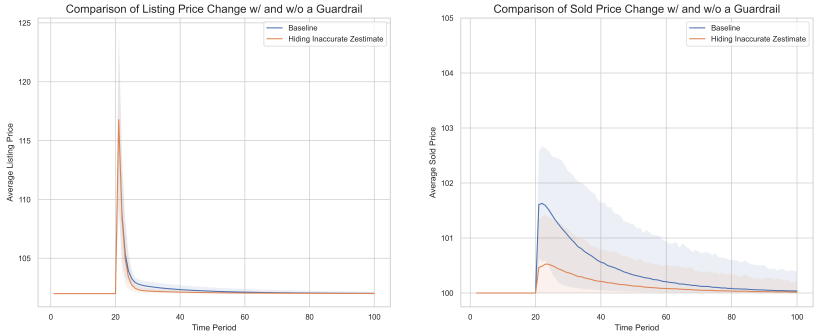


Figure A3: Comparison of market outcomes with and without hiding Zestimate (same disturbance on Zestimate)

The left panel plots the average listing price and the right panel plots the average sold price. The baseline case is the same as the results shown in Figure 7. We observe that when the wildly inaccurate Zestimate is hidden, the effect of Zestimate on sold price is significantly reduced, and the convergence in listing price is faster. This demonstrates that guardrails can reduce the impact of disturbances and speed up the correction.