Three Essays in Real Estate Finance and Economics

Dissertation

zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften (Dr. rer. pol.)

vorgelegt von
Simon Stehle

Universität Konstanz

Sektion Politik – Recht – Wirtschaft
Fachbereich Wirtschaftswissenschaften
Konstanz, 2021

Tag der mündlichen Prüfung: 17. Juni 2021

1. Referent: Prof. Dr. Marcel Fischer
2. Referent: Prof. Dr. Jens Jackwerth
3. Referent: Prof. Dr. Roland Füss

Danksagung

Diese Dissertation wurde von einer großen Zahl an Personen unterstützt, Betreuer, Kollegen, Familienmitglieder und Freunde. Diese Stelle möchte ich daher nutzen, um allen, die mich in den letzten Jahren betreut, ermutigt und abgelenkt haben, zu danken. Ohne diese Unterstützung wäre diese Arbeit in ihrer Form sicherlich nicht möglich gewesen.

Zunächst möchte ich meinem Doktorvater Professor Marcel Fischer danken, der mir zu jeder Zeit mit Rat und Tat zur Seite stand. Diese Dissertation hat nicht nur von seinen zahlreichen Vorschlägen und Ideen profitiert, sondern auch vom kollegialen Miteinander, das er am Chair of Finance, zum Beispiel durch regelmäßige Lehrstuhlfrühstücke und Veranstaltungen zur Teambildung, geschaffen hat. Dank Marcel Fischer durfte ich früh auf Topkonferenzen fahren, mit exklusiven Daten arbeiten und in teilweise stundenlanger, gemeinsamer Arbeit das Schreiben von wissenschaftlichen Arbeiten lernen. Ich möchte außerdem Herrn Professor Jens Jackwerth für seine stets sehr hilfreichen und konstruktiven Kommentare in den Seminaren in Konstanz und St. Gallen danken, sowie für das Übernehmen des Zweitgutachtens. Professor Roland Füss danke ich für die große Unterstützung durch die umfassende Betreuung, das hilfreiche Feedback, die Einladungen zu Seminarvorträgen in St. Gallen und die Übernahme des Drittgutachtens. Juniorprofessor Daniel Ruf danke ich für die gute Zusammenarbeit als Koautor. Ich danke zudem der Graduate School of Decision Sciences (GSDS) für die finanzielle Unterstützung, sowie Jutta Obenland für ihre große Hilfsbereitschaft.

Weiterhin gilt mein Dank den Mitarbeitern am Lehrstuhl, von denen mich viele über weite Strecken der Promotion begleitet haben. Patrick Hauf, für die stetige Hilfsbereitschaft, aber vor allem für die Freundschaft, die sich während unser gemeinsamen Zeit als Bürokollegen entwickelt hat. Marlene Koch, für die vielen guten Gespräche als wir uns noch das Büro teilen durften und das rege Feedback, insbesondere zu meinem Soloprojekt. Ein großer Dank gilt Ilse Geigges-Marschall, die immer Lösungen für alle Anliegen findet, für die freundschaftlichen und wertvollen Gespräche am Morgen. Außerdem danke ich allen Hiwis für ihre Unterstützung, Angelina, Larissa, Konstantin, Jan, Ingrid, Shend, Adrian, Kristina, Thomas und Manuel.

Mein unschätzbarer Dank gilt meinen Eltern, die mich seit ich denken kann in jeder Lage und Entscheidung mit aller Kraft unterstützt haben. Außerdem meinen beiden Schwestern, Eva und Pia, für die stetige Motivation, die jährlichen Besuche vor Neujahr und jede einzelne wunderbare Ablenkung von der Arbeit. Meiner Tante Cornelia für die positive, motivierende Energie bei den vielen schönen Wiedersehen. Meinen Großeltern, Hella und Herbert, danke ich für ihre fortwährende Unterstützung und die große Herzlichkeit und Freude bei jedem Besuch und Telefonat. Ich denke außerdem an meine wundervollen Großeltern Elisabeth und Franz, die meine Zeit als Doktorand leider nicht mehr erleben konnten.

Meine Promotion wurde von vielen, sehr guten Freunden begleitet. Insbesondere möchte ich Tobi danken für die tiefe Freundschaft, die Abende auf dem Weihnachtsmarkt, das Anglühen und Abgrillen, die Fahrten nach Tschechien und die vielen Spieleabende und Nachmittage, um nur einige Dinge zu nennen. Julie, für die großartige Freundschaft, die fachliche und persönliche Hilfe in allen Lagen, und selbstverständlich für jede einzelne Kaffeepause. Anne, für die wunderbare Freundschaft seit ich denken kann, und die große Hilfe in vielen, wichtigen Momenten. Thomas, für die 10-Uhr Pausen, die vielen lustigen Abende im „Crack House“ und die Geduld als Ski- und Segellehrer. Janni, für die Tunnel beim Kicken und die denkwürdigen Treffen am Grenzzaun während der Pandemie.

Außerdem möchte ich allen Freunden vom „Stammtisch“ danken, für die tollen Abende, ob virtuell oder in Person, Hannah, Caddi, Steffen und Seb. Meinen Freunden aus Schule, Studium und Promotion, Joel, Dennis, Sina, Konzi, Rémi, Tino, Paddy, Ele, Johannes, Michelle, Anke, Julian, Eva, Juli, Leo und Ralf für unzählige schöne Tage und Abende. Der „Donnerstagsgang“, für alle Mensareferate und Quizmittage, Marco, Freddy, Timo, Chris, Franzi, Thomas, Christoph und Michi. Meinen Mitbewohnern in Heidis guter Stube, Alex, Chrissi, Nico, Tabea, Miri und Seli. Außerdem danke ich den Freunden und Kollegen am Fachbereich und der GSDS, Manuel, Carl, Tilli, Stephan, Simon, Timm, Phillip und Mario für alle Pausen und Diskussionen.

Mein besonderer Dank gilt Leo und ihrer Familie für ihre Unterstützung, ob aus der Ferne oder während der vielen Besuche und Reisen. Leo stand mir als Partnerin, Freundin und Mitbewohnerin zu jeder Zeit und in jedem wichtigen Moment zur Seite. Danke für alles, die Reisen von Busan bis Rosemount, das Zuhören und Deine Ideen, die kurzweilige Quarantäne, die Hilfe beim Grafikdesign, und das Auffangen und Aufbauen nach langen, schweren Tagen. Damit hast Du mich nicht nur beim Schreiben dieser Dissertation entscheidend unterstützt, sondern vor allem mein Leben unschätzbar bereichert.

Table of Contents

List of Tables

List of Figures

Summary

The dramatic consequences of the recent US housing market bust illustrate the importance of a better understanding of residential housing markets, for the benefit of not only homeowners, but the real economy as well. Fundamental insights from the financial literature can be transferred to housing markets only to a limited extent, however, as residential real estate differs substantially from traditional asset classes. For instance, in contrast to stocks, homes are infrequently traded, non-divisible, and are simultaneously consumption good and financial asset. These and further features pose eminent challenges to the understanding of determinants and dynamics of price formation in housing markets. This cumulative dissertation contributes to this challenging task in three essays.

The pathbreaking work of Case and Shiller (1989) opened great avenues for real estate research by providing an index methodology that addresses many of the complications, which arise from the special nature of homes as an asset class. Yet, the thereby necessary aggregation of data naturally leads to a loss of information, preventing a deeper understanding of micro-level dynamics. As a result, the housing market literature increasingly focuses on price dynamics on the transaction level. Many, highly local events are documented to influence sales prices of nearby homes, such as foreclosures (e.g., Campbell et al., 2011; Gupta, 2019; Guren and McQuade, 2020), speculative activity (Bayer et al., 2021), and rent decontrol (Autor et al., 2014). While most of the prior literature investigates local price dynamics induced by specific events, the general dynamics of housing transaction prices within local markets received much less attention. This research gap is addressed in Chapter 1, "*Local House Price Comovements*," which is joint work with Prof. Dr. Marcel Fischer from the University of Konstanz and Prof. Dr. Roland Füss from the University of St. Gallen. The manuscript is published in the journal *Real Estate Economics* as Fischer et al. (2021).

In this chapter, we study the transaction level evolution of residential house prices along the spatial and temporal dimension, respectively. Using data on repeat sales of condominiums and apartments on Manhattan Island from 2004 to 2015, we analyze

the local nature of comovements in excess of the ones that stem from aggregate market movements. We document that such excess comovements rapidly decline with increasing spatial distance between properties, but are highly persistent over time. The chapter further analyzes the variation in excess comovements during different states of local and aggregate housing market cycles, respectively. We find that excess comovements are particularly strong in local, underperforming markets, especially when the aggregate market is booming. We argue that excess comovements should continually exist in housing markets, as home sales are price signals that are likely to be incorporated in future nearby transactions. Supporting this hypothesis, we find that excess comovements with publicly recorded observations are stronger compared to such comovements with not yet recorded observations.

In the single-authored Chapter 2, "*How Do Assessed Values Affect Transaction Prices of Homes?*," I analyze effects on transaction prices of homes that arise from another distinctive feature of real estate: periodic, value-based taxation. While the existing literature focuses on the introduction of a property tax scheme (e.g., Bai et al., 2014; Du and Zhang, 2015) or changing tax rates (e.g., Elinder and Persson, 2017; Palmon and Smith, 1998), this chapter analyzes the price impact from the values used as tax base. In many countries worldwide, such as the US, homeowners periodically pay property taxes as a share of their homes' assessed values, which are supposed to reflect the fair market value of a home and are generated by a local assessor. As homes are difficult to price for the various reasons discussed above, assessed values are subject to valuation errors, which is reflected in the documented inequity in tax payments among homeowners (e.g., Hodge et al., 2017). Yet, homes' assessed values contain valuable information about future tax payments and simultaneously constitute updated price estimates, respectively. Despite their significance for homeowners, the effects from potentially misspecified assessed values on the transaction prices of the underlying homes remain largely unexplored, a research gap that is addressed in Chapter 2.

In this chapter, I outline that assessed values should affect the sales prices of the respective homes through two counteracting channels. First, as they are supposed to reflect the market values of the corresponding homes, assessed values are potential reference prices for buyers and sellers, which should lead to a positive effect on transaction prices (anchoring channel). Second, an increased assessed value increases the tax burden of the home, implying a negative effect (tax channel). To separate the causal impact of assessed values on transaction prices from the naturally existing positive correlation, I employ a quasi-experimental framework exploiting publication dates of assessed values as well as geographic variation in the frequencies at which homes are reassessed. My results reveal that the tax channel dominates: An increased assessed value leads to a lower sales price of the respective home. In line with the tax channel, I find that homes associated with higher effective tax rates are affected stronger than homes associated with lower tax rates. This chapter thus documents an additional source of inequity among homeowners emerging from property taxation: Owners of homes with unexpectedly high assessed values do not only pay higher taxes, but also sell their home at a lower price.

Chapter 3, "*Investors in the Housing Market*," which is a joint project with Prof. Dr. Marcel Fischer and Prof. Dr. Roland Füss, as well as Jun-Prof. Daniel Ruf, Ph.D. from the Goethe University Frankfurt, analyzes a wide range of investors that emerges from the particularities of real estate. For instance, the high and positive autocorrelation of house prices, documented for many countries since the work of Case and Shiller (1989), allows for the speculation on momentum, while the option to renovate low-priced homes to sell above market prices allows for "flipping" strategies (Bayer et al., 2020). Besides, the spatial nature allows for external, on average less informed, private investors to enter local housing markets, such as out-of-town, second-home buyers (Chinco and Mayer, 2016) and foreign investors (Cvijanovic and Spaenjers, 2020). Simultaneously, a large share of homes is held by owner-occupiers, who should naturally focus on the consumption role of homes rather than financial gains. The high heterogeneity among different investor groups that emerges with distinct strategies is likely to be associated with a varying exposure to local risk and momentum, respectively, a research question that is addressed in Chapter 3.

In this chapter, we analyze the sizable and systematic differences in annualized capital gains across investor groups in the US housing market. Using nationwide data on more than 21 million repeat sales, we investigate the performance of highly heterogeneous buyers, owner-occupiers, private, short-, and long-term investors. Our results link the observed differences in capital gains to heterogeneous risk-taking. Investor-specific exposure to past local return risk explains a sizable and persistent share of investors' high capital gains. In contrast, neither location choice nor temporal factors on the local or aggregate level, respectively, can help explain investors' outperformance. The results thus indicate that the distinct performances of investors analyzed in the recent literature can be partly explained by the compensation for heterogeneous risk-exposure.

Each of the three chapters in this cumulative dissertation contributes to the understanding of residential house prices on the transaction level. Chapter 1 shows that returns of individual transactions are subject to substantial excess comovements, even within zip-codes, and documents that such comovements vary over the local and aggregate housing market cycle. Chapter 2 shows that assessed values negatively affect the transaction prices of the underlying homes. In Chapter 3, we show that the outperformance of different investor groups in the housing market, documented for capital gains from repeated individual trades of residential homes, can be linked to heterogeneous exposure to local return risk. The results documented in this dissertation thus provide new insights on the formation of housing transaction prices that are valuable for not only homeowners but also private and institutional investors.

References

Autor, D. H., C. J. Palmer, and P. A. Pathak (2014). "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts." *Journal of Political Economy* 122. (3), 661–717.

Bai, C., Q. Li, and M. Ouyang (2014). "Property Taxes and Home Prices: A Tale of Two Cities." *Journal of Econometrics* 180. (1), 1–15.

Bayer, P., C. Geissler, K. Mangum, and J. W. Roberts (2020). "Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market." *Review of Financial Studies* 33. (11), 5212–5247.

Bayer, P., K. Mangum, and J. W. Roberts (2021). "Speculative Fever: Investor Contagion in the Housing Bubble." *American Economic Review* 111. (2), 609–651.

Campbell, J. Y., S. Giglio, and P. Pathak (2011). "Forced Sales and House Prices." *American Economic Review* 101. (5), 2108–2131.

Case, K. E. and R. J. Shiller (1989). "The Efficiency of the Market for Single-Family Homes." *American Economic Review* 79. (1), 125–137.

Chinco, A. and C. Mayer (2016). "Misinformed Speculators and Mispricing in the Housing Market." *Review of Financial Studies* 29. (2), 486–522.

Cvijanovic, D. and C. Spaenjers (2020). "'We'll Always Have Paris': Out-of-Country Buyers in the Housing Market." *Management Science* forthcoming.

Du, Z. and L. Zhang (2015). "Home-Purchase Restriction, Property Tax and Housing Price in China: A Counterfactual Analysis." *Journal of Econometrics* 188. (2), 558–568.

Elinder, M. and L. Persson (2017). "House Price Responses to a National Property Tax Reform." *Journal of Economic Behavior & Organization* 144, 18–39.

Fischer, M., R. Füss, and S. Stehle (2021). "Local House Price Comovements." *Real Estate Economics* 49. (S1), 169–198.

Gupta, A. (2019). "Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults." *Journal of Finance* 74. (5), 2249–2301.

Guren, A. M. and T. J. McQuade (2020). "How Do Foreclosures Exacerbate Housing Downturns?" *Review of Economic Studies* 87. (3), 1331–1364.

Hodge, T. R., D. P. McMillen, G. Sands, and M. Skidmore (2017). "Assessment Inequity in a Declining Housing Market: The Case of Detroit." *Real Estate Economics* 45. (2), 237–258.

Palmon, O. and B. A. Smith (1998). "New Evidence on Property Tax Capitalization." *Journal of Political Economy* 106. (5), 1099–1111.

Zusammenfassung

Die dramatischen Folgen der jüngsten US-Immobilienmarktblase zeigen, wie wichtig ein besseres Verständnis von Häusermärkten ist, sowohl für Wohneigentümer als auch die Realwirtschaft. Ein großer Teil grundlegender Erkenntnisse der Finanzforschung ist jedoch nur in begrenztem Ausmaß auf Wohnungsmärkte übertragbar, da sich Wohnimmobilien erheblich von traditionellen Anlageklassen unterscheiden. So werden Häuser und Wohnungen im Gegensatz zu Aktien selten gehandelt, sind nicht teilbar und sind zeitgleich Konsumgut und Finanzanlage. Diese und weitere Besonderheiten stellen eminente Herausforderungen für das Verständnis sowohl von Determinanten als auch den Dynamiken von Preisbildungen in Wohnimmobilienmärkten dar. Diese Dissertation leistet einen Beitrag zu dieser herausfordernden Aufgabe in drei Aufsätzen.

Die richtungsweisende Arbeit von Case und Shiller (1989) eröffnete eine Vielzahl an Möglichkeiten für die Immobilienforschung, indem sie eine Indexmethodik bereitstellte, welche einige der besonderen Charakteristiken von Immobilien als Anlageklasse adressiert. Die dabei erforderliche Aggregation von Daten führt jedoch zwangsläufig zu einem Informationsverlust, was ein tieferes Verständnis von Preisdynamiken auf Mikroebene verhindert. In den letzten Jahren hat sich daher ein stetig wachsender Strang der Literatur mit der Untersuchung von Hauspreisdynamiken auf Transaktionsebene beschäftigt. Diese dokumentiert einige, hoch lokale Ereignisse, die Preise von nahegelegenen Häusern beeinflussen, wie etwa Zwangsversteigerungen (z.B., Campbell u. a., 2011; Gupta, 2019; Guren und McQuade, 2020), Spekulationsgeschäfte (Bayer u. a., 2021) und das Abschaffen von Mietpreisbremsen (Autor u. a., 2014). Während ein Großteil der bisherigen Forschung lokale Preisdynamiken, die durch spezifische Ereignisse ausgelöst werden, untersucht, wurde der allgemeinen Dynamik von Wohntransaktionspreisen innerhalb lokaler Märkte deutlich weniger Aufmerksamkeit geschenkt. Diese Forschungslücke wird im ersten Kapitel dieser Dissertation, „*Local House Price Comovements*“, adressiert, welches in Zusammenarbeit mit Prof. Dr. Marcel Fischer von der Universität Konstanz und Prof. Dr. Roland Füss von der Universität St. Gallen entstanden und im Journal *Real Estate Economics* als Fischer u. a. (2021) publiziert ist.

In diesem Kapitel untersuchen wir die Entwicklung von Wohnungspreisen jeweils in räumlicher und zeitlicher Dimension. Anhand von Daten über wiederholte Transaktionen („*repeat sales*“) von Eigentumswohnungen und Apartments auf *Manhattan Island* aus den Jahren 2004 bis 2015 analysieren wir die gemeinsame, über die Bewegungen des Gesamtmarktes hinausgehende, lokale Entwicklung von Immobilienpreisen. Wir finden, dass solche Überschusskorrelationen („*excess comovements*“) mit zunehmender räumlicher Distanz zwischen Wohnungen schnell abnehmen, über die Zeit hinweg jedoch sehr beständig sind. Weiterhin analysieren wir Veränderungen dieser Überschusskorrelationen während verschiedener Zustände des lokalen bzw. aggregierten Wohnungsmarktzyklus'. Wir beobachten, dass diese in lokalen, sich unterdurchschnittlich entwickelnden Märkten besonders ausgeprägt sind, insbesondere während eines steigenden Gesamtmarktes. Wir argumentieren, dass Überschusskorrelationen in Wohnungsmärkten fortwährend zu finden sein sollten, da beobachtete Verkäufe Preissignale sind, die sehr wahrscheinlich in zukünftigen Transaktionen aufgenommen werden. Diese Hypothese wird von dem Befund unterstützt, dass Überschusskorrelationen mit öffentlich registrierten Beobachtungen stärker sind als mit solchen, die noch nicht registriert wurden.

In dem von mir eigenständig erarbeiteten Kapitel 2, „*How Do Assessed Values Affect Transaction Prices of Homes?*“, analysiere ich Auswirkungen auf Transaktionspreise von Häusern, die sich aus einer weiteren Besonderheit von Immobilien ergeben: der periodischen, wertbasierten Besteuerung. Während sich die bestehende Literatur auf die Einführung eines Besteuerungssystems für Wohnvermögen (z.B., Bai u. a., 2014; Du und Zhang, 2015) oder sich ändernde Steuersätze innerhalb solcher Systeme (z.B., Elinder und Persson, 2017; Palmon und Smith, 1998) konzentriert, analysiert dieses Kapitel Preisauswirkungen, welche durch die Steuerbasiswerte entstehen. In vielen Ländern der Welt wie z.B. den USA, zahlen Hausbesitzer regelmäßig Wohnvermögenssteuern als Anteil eines Schätzwertes ihres Hauses („*assessed value*“), der den fairen Marktwert eines Objektes widerspiegeln soll und von einem lokalen Gutachter ermittelt wird. Da Häuser aus den bereits angeführten Gründen schwer zu bepreisen sind, unterliegen „*assessed values*“ Bewertungsfehlern, was sich in der von der Literatur dokumentierten Ungerechtigkeit bezüglich Steuerzahlungen widerspiegelt (z.B., Hodge u. a., 2017). Dennoch enthalten die „*assessed values*“ von Häusern wertvolle Informationen über zukünftige Steuerzahlungen und liefern gleichzeitig aktuelle Preisschätzungen für die jeweiligen Objekte. Trotz ihrer signifikanten Bedeutung für Hauseigentümer ist der Einfluss von Schätzwerten auf Transaktionspreise noch weitgehend unerforscht, eine Forschungslücke, die in Kapitel 2 behandelt wird.

In diesem Kapitel stelle ich dar, dass „*assessed values*“ die Verkaufspreise von Immobilien über zwei gegenläufige Kanäle beeinflussen sollten. Erstens, da sie die Marktwerte der entsprechenden Häuser widerspiegeln sollen, sind „*assessed values*“ potenzielle Referenzpreise sowohl für Käufer als auch Verkäufer, was zu einem positiven Effekt auf die Transaktionspreise führen sollte (Ankerkanal). Zweitens führt ein erhöhter „*assessed value*“ zu einer erhöhten Steuerbelastung für die zugrunde liegende Immobilie, was einen negativen Preiseffekt impliziert (Steuerkanal). Um den kausalen Einfluss der „*assessed values*“ auf Transaktionspreise von der generell existierenden, positiven Korrelation zu trennen, verwende ich einen quasi-experimentellen Ansatz, der sowohl die Veröffentlichungszeitpunkte der „*assessed values*“ als auch geographische Variation in den Zeitabständen, in denen Häuser neu bewertet werden, ausnutzt. Meine Ergebnisse zeigen, dass der Steuerkanal dominiert: Ein erhöhter Schätzwert führt zu einem niedrigeren Verkaufspreis des jeweiligen Hauses. In Übereinstimmung mit dem Steuerkanal finde ich, dass Häuser, die mit höheren effektiven Steuersätzen assoziiert sind, stärker betroffen sind als Häuser mit niedrigeren Steuersätzen. Die Ergebnisse in Kapitel 2 dokumentieren somit eine zusätzliche Quelle der Ungerechtigkeit unter Hauseigentümern, die sich aus einer wertbasierten Immobilienbesteuerung ergibt: Eigentümer von Häusern mit unerwartet hohen „*assessed values*“ zahlen nicht nur höhere Steuern, sondern erzielen auch einen niedrigeren Verkaufspreis für ihr Haus.

Kapitel 3, „*Investors in the Housing Market*“, welches in Zusammenarbeit mit Prof. Dr. Marcel Fischer und Prof. Dr. Roland Füss, sowie Jun-Prof. Daniel Ruf, Ph.D. von der Goethe Universität Frankfurt, entstanden ist, analysiert ein breites Spektrum von Investoren das durch die Besonderheiten von Wohnimmobilien geschaffen wird. Zum Beispiel ermöglicht es die hohe und positive Autokorrelation von Hauspreisen, die für viele Länder seit der Arbeit von Case und Shiller (1989) dokumentiert ist, die Spekulation auf steigende Marktpreise, während die Möglichkeit, niedrigpreisige Häuser zu renovieren, um sie über dem Marktpreis zu verkaufen, „flipping“-Strategien gestattet (Bayer u. a., 2020). Außerdem erlaubt die räumliche Beschaffenheit externen, im Durchschnitt weniger informierten, privaten Investoren den Eintritt in die lokalen Wohnungsmärkte, wie z.B. außerstädtischen Käufern von Zweitwohnungen (Chinco und Mayer, 2016) und ausländischen Investoren (Cvijanovic und Spaenjers, 2020). Gleichzeitig wird ein großer Teil der Häuser von den Eigentümern selbst bewohnt, sodass sich einige Käufer mehr auf den Konsumaspekt als auf finanzielle Gewinne fokussieren sollten. Die unterschiedlichen Investitionsstrategien und Ziele sollten mit einer hohen Heterogenität unter den verschiedenen Investorengruppen im Häusermarkt einhergehen. Diese wiederum sollte mit einer unterschiedlichen Exposition gegenüber lokalen Risiken und Preisbewegungen verbunden sein, eine Forschungsfrage, die in Kapitel 3 behandelt wird.

In diesem Kapitel analysieren wir die beträchtlichen und systematischen Unterschiede in annualisierten Kapitalgewinnen zwischen Investorengruppen im US-Amerikanischen Wohnimmobilienmarkt. Unter Verwendung landesweiter Daten zu mehr als 21 Millionen wiederholten Verkäufen einzelner Wohnimmobilien untersuchen wir die Performance von hoch heterogenen Marktteilnehmern: Eigennutzer, private, kurz- und langfristige Investoren. Unsere Ergebnisse bringen die beobachteten Unterschiede in den Kapitalgewinnen mit heterogenem Risikobezug in Verbindung: Investoren-spezifische Exposition gegenüber vergangenem lokalen Renditerisiko erklärt einen erheblichen und persistenten Anteil der hohen Kapitalgewinne von Investoren. Im Gegensatz dazu können weder Standortwahl noch zeitliche Abstimmung der Transaktionen auf lokale oder aggregierte Marktbewegungen zur Erklärung der überdurchschnittlichen Renditen von Investorengruppen beitragen. Diese Ergebnisse deuten darauf hin, dass die in der jüngeren Literatur dokumentierten Performanceunterschiede unter verschiedenen Investoren auch teilweise durch Kompensation für heterogene Risikoexposition erklärt werden können.

Jedes der drei Kapitel in dieser kumulativen Dissertation trägt zum Verständnis von Wohnimmobilienpreisen auf Transaktionsebene bei. Kapitel 1 zeigt, dass die Renditen einzelner Transaktionen selbst innerhalb von Postleitregionen beträchtlichen Überschusskorrelationen unterliegen, und dokumentiert, dass solche Korrelationen über den lokalen und aggregierten Wohnungsmarktzyklus variieren. Kapitel 2 dokumentiert, dass „*assessed values*“ einen negativen Einfluss auf die Transaktionspreise der ihnen zugrunde liegenden Häuser haben. In Kapitel 3 zeigen wir, dass das überdurchschnittliche Abschneiden verschiedener Investorengruppen auf dem Wohnungsmarkt, die für Kapitalgewinne aus wiederholten Transaktionen einzelner Wohnimmobilien dokumentiert wird, mit heterogener Exposition gegenüber lokalem Renditerisiko verbunden ist. Die in dieser Dissertation dokumentierten Ergebnisse liefern somit neue, wertvolle Erkenntnisse für die Formation von Transaktionspreisen in Wohnungsmärkten, sowohl für Eigenheimbesitzer als auch private und institutionelle Investoren.

References

Autor, D. H., C. J. Palmer, and P. A. Pathak (2014). "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts." *Journal of Political Economy* 122. (3), 661–717.

Bai, C., Q. Li, and M. Ouyang (2014). "Property Taxes and Home Prices: A Tale of Two Cities." *Journal of Econometrics* 180. (1), 1–15.

Bayer, P., C. Geissler, K. Mangum, and J. W. Roberts (2020). "Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market." *Review of Financial Studies* 33. (11), 5212–5247.

Bayer, P., K. Mangum, and J. W. Roberts (2021). "Speculative Fever: Investor Contagion in the Housing Bubble." *American Economic Review* 111. (2), 609–651.

Campbell, J. Y., S. Giglio, and P. Pathak (2011). "Forced Sales and House Prices." *American Economic Review* 101. (5), 2108–2131.

Case, K. E. and R. J. Shiller (1989). "The Efficiency of the Market for Single-Family Homes." *American Economic Review* 79. (1), 125–137.

Chinco, A. and C. Mayer (2016). "Misinformed Speculators and Mispricing in the Housing Market." *Review of Financial Studies* 29. (2), 486–522.

Cvijanovic, D. and C. Spaenjers (2020). "'We'll Always Have Paris': Out-of-Country Buyers in the Housing Market." *Management Science* forthcoming.

Du, Z. and L. Zhang (2015). "Home-Purchase Restriction, Property Tax and Housing Price in China: A Counterfactual Analysis." *Journal of Econometrics* 188. (2), 558–568.

Elinder, M. and L. Persson (2017). "House Price Responses to a National Property Tax Reform." *Journal of Economic Behavior & Organization* 144, 18–39.

Fischer, M., R. Füss, and S. Stehle (2021). "Local House Price Comovements." *Real Estate Economics* 49. (S1), 169–198.

Gupta, A. (2019). "Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults." *Journal of Finance* 74. (5), 2249–2301.

Guren, A. M. and T. J. McQuade (2020). "How Do Foreclosures Exacerbate Housing Downturns?" *Review of Economic Studies* 87. (3), 1331–1364.

Hodge, T. R., D. P. McMillen, G. Sands, and M. Skidmore (2017). "Assessment Inequity in a Declining Housing Market: The Case of Detroit." *Real Estate Economics* 45. (2), 237–258.

Palmon, O. and B. A. Smith (1998). "New Evidence on Property Tax Capitalization." *Journal of Political Economy* 106. (5), 1099–1111.

Chapter 1

Local House Price Comovements*

*The authors would like to thank Konstantin Spring for his excellent research assistance. We are grateful for the helpful comments and suggestions from the editor, Wenlan Qian, two anonymous reviewers, Martin Brown, João Cocco, Christian Riis Flor, Günther Franke, Jens Jackwerth, Winfried Koeniger, Frederic Menninger, Daniel Ruf, Nawid Siassi, Bertram Steininger, Yildiray Yildirim, seminar participants at the Universities of Konstanz, Odense, Regensburg, St. Gallen, the EBS Business School, and the NTNU Trondheim, as well as participants at the 2018 AREUEA National Conference in Washington DC, the 2018 EFMA Annual Meeting in Milan, the 2018 Conference on Decision Sciences in Konstanz, the 2018 ReCapNet Conference of the Centre of European Economic Research (ZEW) in Mannheim, the conference on Housing in the 21st Century at the University of Bonn, and the 2020 American Finance Association Annual Meeting in San Diego. An early, preliminary version of this paper existed under the title "Spillover Effects in Residential House Prices".

1.1. Introduction†

The recent boom and bust in house prices dramatically illustrates the need for a better understanding of price dynamics of residential homes. Since the pioneering work of Case and Shiller (1989), it is a well-established fact that returns on national and city-wide house price indices are subject to strong auto- and cross-sectional correlation. These patterns can be explained by comovements and spillovers in residential house prices. Comovements are caused by common underlying factors, such as gradually changing credit conditions (e.g., Amromin et al., 2018; Chambers et al., 2009; Landvoigt et al., 2015). Spillovers, on the other hand, are caused by a trigger, such as gentrification (Guerrieri et al., 2013), or rent decontrol (Autor et al., 2014), that spills over from affected to unaffected properties. Similarly, a foreclosure affects the trading prices of other properties in a neighborhood that do not go through a foreclosure (e.g., Campbell et al., 2011; Gupta, 2019; Guren and McQuade, 2020). Even interior renovations, which are generally unobservable, can raise house prices in close proximity after transaction prices of the renovated homes become public (Szumilo, 2020). Despite all the progress that has been made in exploring specific sources of local house price comovements and spillovers, surprisingly little is known about their general dynamics, for instance with respect to the distance between traded homes, the time between trades, and the dependence on the market cycle.

Beyond specific events, comovements should be largely driven through two main channels. First, comovements can be caused by the information channel. Financing decisions as well as negotiations between buyers and sellers are influenced by different anchors for the value of a home, that typically rely on comparable sales. This valuation approach makes use of sales prices of homes that are similar in terms of location, other characteristics, and past sales date. Important anchors include appraised values by independent parties as well as Zillow's estimates for the marked values of homes, the so-called Zestimates. By affecting anchors for the value of a home, past sales prices should indirectly affect current sales prices. In addition, past sales prices can also directly affect present sales prices by affecting both buyers' and sellers' behavior. For instance, real estate agents should incorporate past sales prices in their offer prices and during price negotiations – for instance, because they do not want to sell at a worse price than their neighbors. Hence, past price changes in the neighborhood should

†This article is published as Fischer et al. (2021) in the journal *Real Estate Economics* and copyrighted, © 2020 American Real Estate and Urban Economics Association. The corresponding "Online Appendix" is placed at the end of this chapter as Appendix 1.B.

affect present trading prices via both buyers' and sellers' incentives to use past sales prices as easily available anchors (Murfin and Pratt, 2019).

Second, homes in the same neighborhood are of similar size and quality, and share common amenities, such as access to schools, recreational areas, shopping facilities, etc. Hence, *ceteris paribus*, homes in the same neighborhood should be better substitutes than more distant ones. When house prices in a given neighborhood increase, a potential buyer's budget constraint is more likely to be binding, thus increasing the incentive to search for cheaper homes in the nearest surrounding. This substitution effect should cause price increases in one neighborhood to also affect close-by neighborhoods.

In this paper, we investigate the micro-level price dynamics of homes in urban areas using repeat sales. We evaluate the order of magnitude, the persistence, and the state dependence of these effects. Specifically, we ask whether outperformance or underperformance relative to other homes in a given zip-code is a locally concentrated and persistent phenomenon. For that purpose, we compute excess returns of homes as the difference between individual and county-wide returns over the respective period. We run regressions to explore excess comovements, i.e., whether excess returns on individual homes can be explained by excess returns of nearby recently traded homes – even after controlling for zip-code-year dynamics, as well as transaction specific, locational, and macroeconomic controls. In these regressions, recently traded homes are grouped by their physical and temporal distance to the traded home. This categorization allows us to explore the dynamics of excess comovements, in particular, whether excess comovements die out with physical and temporal distance.

Our data covers repeat sales on Manhattan Island between 2004 and 2015 from the CoreLogic database. Manhattan Island is well suited for investigating excess comovements in residential house prices in many regards. First, Manhattan is a liquid real estate market. Second, Manhattan is densely populated, implying that new constructions are scarce and unlikely to have major price impacts. Third, down-payments in Manhattan are very high and, over the last years, more than 50% of the condominiums and house sales have consistently closed without mortgage financing.[1] This makes buyers less dependent on the lending policy of banks, reducing trading frictions enormously and turns Manhattan into a highly efficient real estate market. As most of the transactions are conducted by a real estate broker, information travels quickly to the buyers who have appointed an agent. Finally, the exact trading prices

[1] https://www.propertyshark.com/Real-Estate-Reports/2016/12/13/payments-manhattan-now-500k- almost-double-median-sale-price-us/ retrieved on May 18, 2020

for all homes are publicly available on the homepage of the New York City Department of Finance,[2] implying that information is easily available for all market participants.

Consistent with the spillover effects documented in, e.g., Campbell et al. (2011), Guerrieri et al. (2013), and Rossi-Hansberg et al. (2010), we show that excess comovements are strongest in the nearest neighborhood — particularly within the same building — and die out quickly with increasing distance between traded homes. Hence, our results are unlikely to be driven by events that affect larger neighborhoods within a zip-code. For instance, a 6% increase (i.e., one standard deviation) in the annualized excess return of an apartment leads to a 1.3% increase in the expected returns of an apartment located in the same building. This effect decreases by 80% within a 500 feet radius neighborhood. About 21% of past within-building trades are reflected in present trades. For other trades within 500 feet, the order of magnitude decreases to 4% and dies out for distances exceeding 1,000 feet. Our results are robust to controlling for the evolution of house prices on the borough-level, on a monthly basis, as well as zip-code-year based price movements. In extensive robustness checks, we document our results to withstand other model specifications and parameter choices.

In addition to the spatial dimension, performance is also persistent on the temporal dimension, i.e., excess comovements exist over longer time horizons. For example, within the same building, even conditional on most recent excess returns, 6% of past average excess returns from 2 to 2.5 years ago are reflected in excess returns today. Similar to the generally higher level of correlation in markets with falling prices in stock markets (e.g., Ang and Chen, 2002), as well as housing markets on the aggregate level (e.g., Cotter et al., 2015), our results reveal that local excess comovements are stronger in markets with falling prices. Local underperformance is more persistent than local overperformance, particularly in markets with generally appreciating house prices: Within the same building, 36% of negative excess returns are reflected in today's prices during booming periods, in contrast to only 26% during non-booming states of the aggregate market. Our results thus suggest a higher heterogeneity in terms of local house price changes when house prices increase.

To further investigate which channel drives excess comovements, we exploit that after a home is sold, information about the trade is not immediately publicly available. More specifically, our data allows us to distinguish between the sales date at which the home was sold, and the recording date, after which information about the trade is publicly available. Splitting our sample of past transactions into recorded and non-recorded trades, we show that excess comovements are substantially stronger with

[2] http://www1.nyc.gov/site/finance/taxes/property-rolling-sales-data.page retrieved on May 18, 2020

past trades for which information is already publicly available. As the substitution channel would imply no difference in comovements between recorded and non-recorded sales, our results suggest that it is rather the information channel that is driving excess comovements.

Our work contributes to two important strands of literature. First, it contributes to the literature that documents the existence of excess comovements on the index level (e.g., Cohen and Zabel, 2018; Cotter et al., 2015; Kallberg et al., 2014; Landier et al., 2017), by showing that excess comovements also exist in the trading prices of individual homes on the micro level. Second, it contributes to the growing strand of literature investigating that local events, such as gentrification (Guerrieri et al., 2013), urban revitalization (Rossi-Hansberg et al., 2010), air pollution (Chay and Greenstone, 2005), legislative amendment (Autor et al., 2014), unnatural deaths (Bhattacharya et al., 2020), the Low Income Housing Tax Credit (Diamond and McQuade, 2019), and foreclosures (Anenberg and Kung, 2014; Campbell et al., 2011; Gerardi et al., 2015; Gupta, 2019; Guren and McQuade, 2020; Harding et al., 2009) are important drivers of micro-level house price dynamics. In contrast to the work of Rossi-Hansberg et al. (2010), Campbell et al. (2011), Guerrieri et al. (2013), and Szumilo (2020), which focuses on spillovers related to specific events, we document that, potentially driven by the information channel, prices comove even in the absence of specific events.

This paper proceeds as follows: In Section 1.2, we introduce our data as well as our empirical methodology. Section 1.3 presents our results on excess comovements in residential house prices. Section 1.4 documents the robustness of our results. The final section concludes. An Online Appendix provides additional robustness results.

1.2. Data and methodology

Our data is from the CoreLogic database, which covers 99.9% of the US population.[3] We focus on repeat sales in urban areas using data from Manhattan Island, New York City. Manhattan Island is well suited to investigate excess comovements in residential house prices in several regards. First, given that Manhattan is generally perceived as a very attractive place to live, the market for real estate is liquid, and foreclosures are rare.[4] Second, compared to more rural areas, Manhattan Island is densely populated and space for new buildings is therefore extremely scarce. This

[3] https://www.corelogic.com/solutions/university-data-portal.aspx retrieved on May 18, 2020

[4] According to RealtyTrac.com, only one in every 12,410 trades in New York City relates to a foreclosure. See http://www.realtytrac.com/statsandtrends/foreclosuretrends/ny/new-york-county/new-york/ as of April 2019

severely limits the amount of new construction and the price impact of new buildings on existing places. Third, the exact prices of all trades are publicly available at the New York City Department of Finance's homepage. That is, information about actual trading prices of adjacent homes is easily available for all market participants and our results should be less affected by information asymmetries. Our period of investigation covers trades from January 2004 to December 2015. To compute realized returns on these trades, we use past trading prices ranging back to 2000.

1.2.1. Data cleaning

We consider repeat sales of condominiums and apartments in order to compute returns on investments for such places. Initially, our data set of the market on Manhattan Island consists of 43,466 repeat sales, covering the period from January 2004 to December 2015. The dates for the most recent prior sale range back until January 2000, allowing for longer holding periods even at the beginning of the sample. The removal of observations that are not classified as resales (e.g., subdivisions) or for which information about the date of the transaction, the current or most recent preceding sales price (prior sales price) is not available leaves us with 42,301 observations.[5] The removal of duplicates with identical sales prices, prior sales prices, transaction dates, and geographic coordinates leaves our sample with 41,905 observations. Following Landvoigt et al. (2015), we remove speculative trades with holding periods of less than 180 days, leaving us with 41,283 observations. Finally, similar to Campbell et al. (2011), for every year, we remove outliers with current or prior sales prices in the first and 99th percentile, respectively, leading to a data set of 39,771 observations. To account for outliers coming e.g. from data errors or physical changes of the property, we remove observations in the third and 97th percentile of the annualized return distribution.[6] Our final data set then consists of 37,385 observations.

Figure 1.2.1 visualizes the distribution of these observations geographically.[7] Unsurprisingly, there are no trades in park areas, such as in Central Park, or industrial areas, such as 207th Street Train Yard Facility in the northern part of Manhattan Island.

[5]In our data, the characteristics of repeat-sales and non-repeat-sales are remarkably similar, indicating that the removal of non-repeat-sales does not leave us with a non-representative sample. For instance, the average trading price of our repeat-sales is USD 1.41 million (in January 2015 dollars), whereas it is USD 1.42 million for the non-repeat sales. The similarity of both subsamples further suggests that the non-repeat sales do not constitute a systematic, confounding factor in our analysis.

[6]The results are qualitatively robust to the removal of only one or two percent of each tail.

[7]The data for zip-code boundaries was obtained from the US Census Bureau, download link: http://www2.census.gov/geo/tiger/GENZ2015 retrieved on October 30, 2019

Figure 1.2.1
Geographic illustration of the data set

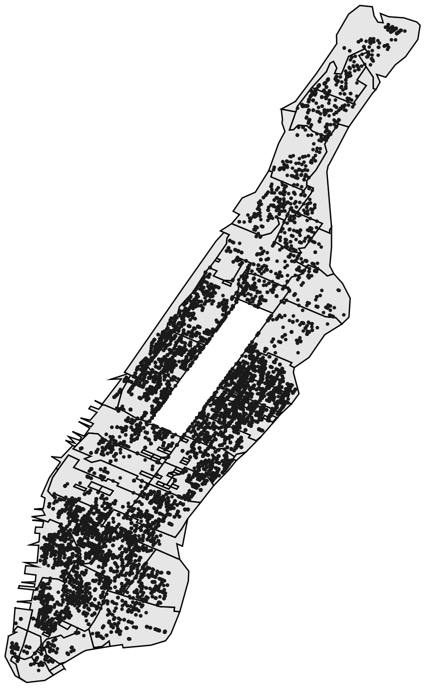

This graph visualizes the geographic distribution of our observations on Manhattan Island. The black lines indicate zip-code boundaries.

Figure 1.2.2 summarizes the evolution of residential house prices in our cleaned data set using a repeat sales index (Case and Shiller, 1989) constructed on a monthly basis for the time period from January 2000 to December 2015. Similar to house prices on the national level, from Figure 1.2.2, Manhattan Island experienced a significant boom during the 2000s, with prices more than doubling from 2000 to 2006. Thereafter, house prices did not show a clear trend until they sharply declined in late 2008 – later than on the national level.[8] This relatively late decline may reflect layoffs in the financial industry. These layoffs did not occur instantly when house prices on the national level started declining, but with a certain delay.

1.2.2. Excess returns

The repeat sales in our data differ along two important dimensions that make a direct comparison of returns difficult. First, the lengths of the time intervals between two trades may differ substantially. Second, returns depend crucially on the phase of

[8] In Subsection 1.3.2, we use these differences in the general evolution of house prices to investigate whether comovements vary with the phase of the housing market cycle.

Figure 1.2.2
Evolution of house prices on Manhattan Island

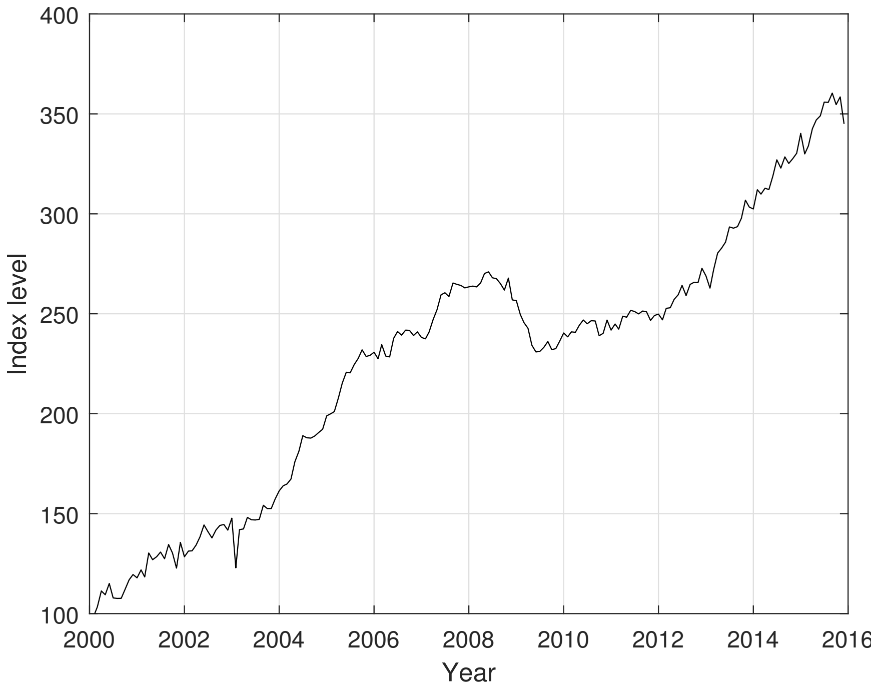

Nominal repeat sales index of Manhattan Island's condominium/apartment market based on our final data set. The index level is normalized to 100 in January 2000.

the housing market cycle. To control for these two effects, we compute annualized market-adjusted excess returns, $r_{t,t-}$, for properties traded at month t and previously traded at month $t-$ as follows:

$$r_{t,t-} = \left(\frac{P_t}{P_{t-}}\right)^{\frac{1}{y(t,t-)}} - \left(\frac{C_t}{C_{t-}}\right)^{\frac{1}{y(t,t-)}}, \tag{1.2.1}$$

in which P_t and P_{t-} denote the present and prior trading prices of the property in months t and $t-$, respectively. $y\,(t, t-)$ is the time distance in years between the two trades, and C_t and C_{t-} denote the index levels of the Manhattan Island repeat-sales price index constructed as in Case and Shiller (1989) from our cleaned data in months t and $t-$, respectively. By subtracting the index return, we remove aggregate effects that systematically affect house prices, such as inflation, seasonal effects, and the phases of the housing market cycle at the moments of the two trading dates, as well as other common specific events, such as the September 11 attacks. Simultaneously, excess returns allow us to distinguish between states with local over- and underperformance, which, as we demonstrate in Subsection 1.3.2, are associated with differential excess comovements.

Likewise, by investigating excess returns over different holding periods, our analysis aims at highlighting the importance of the availability of information for residential house prices. That is, we conjecture that only excess returns that become observable via information about the publicly available trading price, affect excess returns of nearby homes, but not changes in value of untraded homes, for whom the new value does not become public due to the absence of a trade. We further test the role of information frictions more formally by comparing excess comovements at the sales and recording date. We exploit that information about the sales price does not become publicly available prior to the recording date.

1.2.3. Control variables

Our control variables can be broadly split into three different categories: (1) transaction-specific, (2) locational, and (3) macro-financial control variables.

1.2.3.1. Transaction-specific variables

In our data cleaning procedure, we remove transactions with holding periods of less than 180 days, which are likely to be speculative trades. Short holding periods may be targeted at larger renovations during that period, aiming at substantially increasing a property's value. To account for these possible effects, we include mutually exclusive dummy variables for holding periods of less than one and less than two years, respectively.

The results in Landvoigt et al. (2015) document that during the recent housing market boom, housing returns varied substantially between homes in different price segments in a nonlinear fashion. To account for this effect, we control for the log of the inflation-adjusted prior sales price (in January 2015 dollars) as well as its square.

While private investors profit from both their home as a durable consumption good and from house price appreciations, corporations should place higher emphasis on earning higher returns on their investments. To control for these effects, we include two dummies for whether a property is sold or bought by a corporation and a dummy for whether a home is bought to become an owner-occupied home. Transactions in which the buyer is a corporation or the home is bought to serve as an owner-occupied home are already marked in our database. In addition, we construct a dummy variable indicating whether the seller is a corporation or not.[9]

[9] We define a seller as a corporation, if the seller's name contains keywords such as ACQUI, ASSOC, AVENUE, BANK, BOARD, CORP, CREDITOR, EQUIT, ESTATE, FUND, HDFC, HLDGS, HOLDING, HOUSING, HSNG, INC, INVEST, L*L*C, LLC, LP, LTD, OWNER, PARTNER, PLC,

1.2.3.2. Locational variables

The location of a residential home is one of the key factors determining its price (e.g., Can, 1990; Case and Mayer, 1996). To control for possible changes in the pricing of location-specific factors, we control for the view on the Central Park, the waterfront, as well as the walking-distance to these two amenities. We further control for distance to Times Square, the New York Stock Exchange, and the nearest entry to the subway.[10] More specifically, we include a waterfront-view-dummy if a home has a direct view on the water surrounding Manhattan Island; i.e., if the home is separated from water only by a road, a park, or both, but not by a building. We further include a walking-distance dummy, if the city block distance to the waterfront does not exceed 500 feet. To account for easy access to the subway system, we include two dummies: a dummy for very close distances to the nearest entry for city block walking-distances of less than 100 feet and a dummy for close distances of 100 to less than 500 feet.

In a similar fashion, we include mutually exclusive dummies for a view and a walking-distance to Central Park, if the beeline does not exceed 100 feet and the city block walking-distance does not exceed 500 feet, respectively. For Times Square and the New York Stock Exchange we include two dummies for short walking-distance and medium walking-distance if the city-block distance is less than 1,000 feet and 1,000 to less than 2,000 feet, respectively.[11]

Guerrieri et al. (2013) document substantial differences in house price growth across neighborhoods. To account for these differences, we proceed similar to Campbell et al. (2011), who use census-tract-year dummies, and control for zip-code-year fixed effects in the current and the prior year of trade of the home. To attain a reasonable number of observations per zip-code (at least 1,000 observations), we have to cluster a few adjacent zip-codes. A detailed overview over the clustered zip-codes can be found in the Appendix. To control for the impact of liquidity in the local housing market on transaction prices (Caplin and Leahy, 2011), we control for the log of one plus the number of trades in the past 180 days on the zip-code level.

PORTFOLIO, PROP, QUATAR, REALTY, STREET, TRUST, or *LP, in which * signifies blank spaces. A manual comparison of more than 3,000 observations did not indicate any missing words.

[10] The geographic coordinates of the New York subway entries are from NYC Open Data (https://opendata.cityofnewyork.us) retrieved on May 18, 2020

[11] We also investigated a setting in which we excluded our controls for the Central Park, New York Stock Exchange, and Times Square. These results are qualitatively identical to those reported in Table 1.3.1. They can be found in Table 1.B.4 in the Online Appendix.

1.2.3.3. Macro-financial variables

To control for changes in the macroeconomic environment, we include the seasonally-adjusted real growth rate of the GDP relative to the previous quarter with a lag of one period from the US Bureau of Economic Analysis, the seasonally-adjusted monthly growth rate of the unemployment rate in New York City from the Bureau of Labor Statistics, and the percentage change in the average fixed mortgage lending rate from the Federal Housing Finance Board.

Since the pioneering work of Case and Shiller (1989), it is known that residential house prices exhibit a significant degree of autocorrelation. To explain price movements, it is therefore important to control for this persistence. Our analysis focuses on explaining excess returns rather than raw returns, thus removing the systematic autocorrelation.

Table 1.2.1 summarizes key properties of our data. As to be expected, the annualized excess return is not significantly different from zero.[12] The average holding period is only about 5.5 years, indicating that Manhattan is a fairly liquid market for residential homes. About 9% of properties are even resold within up to 2 years, which may, among others, reflect institutional investors' activities that account for about 15% of purchases and 10% of sales. Yet, the majority of trades (53%) still represents sales of owner-occupied places. With an average prior trading price of USD 1.279 million (inflation-adjusted to 2015 prices), prices on Manhattan Island are among the most expensive in the US This high average transaction price suggests that prices should be largely determined by location. In contrast, renovations or a new kitchen should have a lower impact on the trading price, advocating the repeat sales approach. Likewise, the short average holding period provides additional support for the repeat sales approach.

1.2.4. Methodology

The goal of our work is to investigate how recent past excess returns in residential house prices comove with present returns of homes in the neighborhood. We further ask whether the strength of these effects varies with the stage of the housing market cycle and whether excess comovements vanish with temporal distance between two trades. For that purpose, we define K mutually exclusive neighborhoods for each observed trade. We refer to trades with coinciding geographic coordinates, i.e., trades in the same building, as the first-order neighborhood throughout. Additionally, we

[12] The small positive value reflects that the market return constructed using the Case-Shiller methodology weights observations unequally.

Table 1.2.1
Summary statistics

Variable name	Mean	Standard deviation
Annualized excess return	0.01	0.06
Holding period (in years)	5.41	2.67
Liquidity	168.71	100.86
Central Park view	0.03	0.17
Central Park walking	0.04	0.20
Very close subway	0.02	0.15
Close subway	0.19	0.39
Short distance Times Square	0.003	0.06
Medium distance Times Square	0.01	0.09
Short distance NYSE	0.01	0.11
Medium distance NYSE	0.02	0.14
Waterfront view	0.03	0.16
Waterfront walking distance	0.05	0.22
Dummy one year	0.02	0.13
Dummy two years	0.07	0.26
Price (in mio USD)	1.28	1.30
Seller corporation	0.10	0.30
Buyer corporation	0.16	0.36
Owner-occupied	0.53	0.50
Lagged GDP growth	0.005	0.005
Lagged unemployment growth	-0.007	0.02
Lagged interest change * 10,000	-1.94	300.38

This table provides descriptive statistics of the variables used. “Annualized excess returns” are defined in Equation (1.2.1). “Holding period (in years)” is the number of years between two trades of a given residential home. “Liquidity” is the number of sales during the past 180 days in the respective zip-code. “Central Park view” and “Central Park walking” are two dummies indicating whether a home has a view on the Central Park (distance of less than 100 feet beeline) and the city-block distance to the nearest entrance is less than 500 feet, respectively. “Very close subway” and “Close subway” are two mutually exclusive dummies indicating whether the city-block distance to the nearest subway entrance is less than 100 feet or 100 to less than 500 feet, respectively. “Short distance Times Square/NYSE” and “Medium distance Times Square/NYSE” are mutually exclusive dummies for whether the city-block distance to Times Squares/NYSE is less than 1,000 feet or 1,000 to less than 2,000 feet, respectively. “Water front view” is a dummy indicating whether a home has direct view on the water surrounding Manhattan Island. “Waterfront walking distance” is a dummy indicating whether the city-block distance to the waterfront does not exceed 500 feet. “Dummy one year” and “Dummy two years” are indicators for holding periods of one and two years, respectively. “Price (in mio USD)” is the most recent available prior trading price of the home CPI-adjusted to January 2015 dollars. “Seller/Buyer” corporation is a dummy indicating whether the seller/buyer is a corporation. “Owner-occupied” is a dummy indicating whether the buyer is the new inhabitant. “Lagged GDP growth” is the previous quarter’s US GDP growth. “Lagged unemployment growth” is the previous month’s New York City wide unemployment growth rate. “Lagged interest change” is the percentage change of the average fixed mortgage lending rate in the month prior to the sale.

draw $K-1$ circles around each observed trade. We want to end up with the same expected number of observations in each of these $K-1$ circles to make sure that, on average, liquidity is the same in all circles and the average excess returns from all of these neighborhoods are thus estimated with the same precision. We therefore draw the circles such that the area inside each of them is identical.[13]

The first of the $K-1$ circles, also referred to as the second-order neighborhood throughout, is characterized by a maximum distance of 500 feet, roughly corresponding to two blocks.[14] The borders of the other circles, which we refer to as third-, fourth-, fifth-, and sixth-order neighborhoods thus lie at 707, 866, 1,000, and 1,118 feet, respectively, leaving us with on average 5.3 to 6.5 historical trades in the second- to sixth-order neighborhood for every current trade.[15] Figure 1.2.3 visualizes our construction of $K = 6$ neighborhoods for a specific property. For every neighborhood k, we define a neighborhood-specific excess return, $\bar{r}^e_{i,k}$ as the average of the observed excess returns in the T days prior to trade i.

To better illustrate this procedure, consider the following example: In the first-order neighborhood ($k = 1$), a home i that is sold at $t = 0$. Suppose that there are four other observed trades in that neighborhood, realized at $t = -181$, $t = -145$, $t = -25$, and $t = 5$, respectively. In our base case, we consider the past $T = 180$ days, such that when computing $\bar{r}^e_{i,1}$, we only consider the second and the third trade, realized at $t = -145$ and $t = -25$, respectively. The first trade lies too far in the past to be included in $\bar{r}^e_{i,1}$. The fourth trade was not yet realized when i was traded, i.e., this information was not yet available and is therefore not included in $\bar{r}^e_{i,1}$. The average excess return for the first-order neighborhood of observation i, $\bar{r}^e_{i,1}$, is then computed as the average of the two excess returns of trades two and three.

We employ the following regression setup:[16]

[13]The total number of neighboring, recently transacted prices ranges from about 200,000 to 246,000 (each property can be a neighboring transaction multiple times) in the $K-1 = 5$ circles of our empirical analysis, i.e., the numbers of observations among our defined neighborhoods are roughly of equal size.

[14]In Section 1.4, we apply the city-block metric to demonstrate the robustness of our results.

[15]Intuitively, not only the physical distance, measured in feet, may be important, but also the direction, reflecting that neighborhoods in one direction are more similar/attractive than in another direction. It would have been interesting to investigate this issue in more detail. However, with an average number of observations of 5.3 to 6.5 in the neighborhoods, homes are traded too infrequently to split our samples further up. Simultaneously, whether neighborhoods in one direction are more attractive/similar than in another direction may be challenging to decide for some trades. Yet, we do control for a few specific attractive places in the neighborhoods, namely for the water front, the Central Park, and New York Stock Exchange.

[16]Equation (1.2.2) can be easily rewritten in spatial econometrics notation because $\bar{r}^e_{i,k}$ reflects the k-th spatial lag. Nevertheless, under the assumption of homoskedastic error terms, OLS is applicable, since we account for the time-directionality in constructing the spatial weights. In our robustness

Figure 1.2.3
Construction of neighborhoods

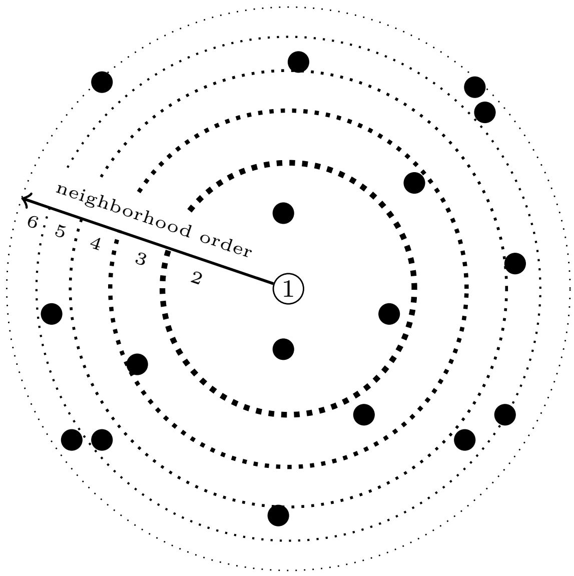

This figure visualizes our construction of neighborhoods. The center symbolizes a trade for a given home. Other trades in the same building are defined as trades in the first-order neighborhood. The dotted circles surrounding the center depict edges of mutually exclusive neighborhoods of orders two to six.

$$r^e_{i,t,t-,z} = \alpha_z + \sum_{k=1}^{K} \rho_k \bar{r}^e_{i,k} + \delta_{a(t),z} - \delta_{a(t-),z} + X_{i,t}\beta + \epsilon_{i,t,t-,z}, \tag{1.2.2}$$

in which $r^e_{i,t,t-,z}$ is the annualized excess return on property i in zip-code z realized between time $t-$ and t, $\delta_{a(t),z}$ and $\delta_{a(t-),z}$ are zip-code-year fixed effect dummies, for which the subscripts $a(t)$ and $a(t-)$ refer to the years in which the trades took place, respectively. For example, if t corresponds to any day of sale in the year 2010, $a(t) = 2010$. $X_{i,t}$ is a vector of control variables, $\epsilon_{i,t,t-,z}$ is a normally-distributed error term, and α_z reflect zip-code level intercepts.

The excess returns, $\bar{r}^e_{i,k}$, from past trades are (deliberately) based on transactions in close areas suggesting that multicollinearity between these returns may be an issue. Empirically, however, multicollinearity between returns in our neighborhoods is not an issue with correlations among the $\bar{r}^e_{i,k}$s ranging from 0.047 to only 0.121.

Our specification has several advantages compared to a spatial autoregressive (SAR) model. First, it avoids the use of a large weighting matrix consisting of more than 37,000 individual properties, as it would be the case with a spatiotemporal

analysis, we also contrast our main findings with those derived from a spatial autoregressive (SAR) model. The results are given in Table 1.B.7 in the Online Appendix.

autoregressive model as, e.g., in Tu et al. (2004). The inclusion of such a highly dimensionally weighting matrix is computationally burdensome and only manageable by either imposing thresholds for the proximity measure or by using a concentrated log likelihood function such as in the specification of a matrix exponential spatial (MESS) model. Second, our specification avoids the use of a log determinant with multiple spatial lags, which would be computationally demanding. Third, the exogenous variables can still be interpreted as marginal effects and not as first round effects without taking into account feedback and spillover effects as in the case of the SAR model. Finally, our model provides a simple way to incorporate the temporal distances as illustrated in Figure 1.3.1 and interaction terms to control for market cycles as in Subsection 1.3.2.

The precision of our estimate for the annualized excess return is generally increasing with the length of the time interval between two trades. Intuitively, when the two trades occur within a relatively short time period, small deviations in observed trading prices of individual properties and short-term fluctuations in the local house price index lead to significant amplifications when being annualized. Hence, annualized excess returns tend to be subject to higher variation when two trades occur within a relatively short time period. To account for this phenomenon in our analysis, we allow the variance of $\epsilon_{i,t,t-,z}$ to depend on the difference D between t and $t-$: $\text{Var}\left(\epsilon_{i,t,t-,z}\right) = \exp\left(\gamma_1 + \gamma_2 D\right)$ where γ_1 and γ_2 are regression-endogenously determined coefficients. We use the exponential function to ensure positivity of the variance in the optimization process.[17]

Our goal is to explore whether recent past excess returns in the neighborhood comove with present excess returns, i.e., whether the ρ_ks are different from zero and, if so, whether price comovements decay with increasing distance, i.e., whether $|\rho_1| > \cdots > |\rho_K|$.

Intuitively, households typically have a preference for a certain location of their homes. This preference could both reflect the neighborhood's facilities, such as good schools, restaurants, shops, as well as social ties, such as other family members or friends living in the neighborhood. A home outside the preferred location is a substitute for the home at the preferred location, because both homes provide households with the same housing services. Comovements in residential house prices should then reflect that households react to price increases for homes in a given neighborhood by purchasing substituting homes in close-by neighborhoods, thus causing price increases in these neighborhoods. The distance between two homes can be interpreted as a proxy for how well two homes can be substituted with each other. More precisely, the smaller

[17] Our estimates are qualitatively robust to the homoskedastic case.

the distance between the two homes, the more a price signal from a previous trade in one home should affect the price of the respective other. The order of magnitude of excess comovements should therefore decrease with increasing distance between homes.

The adjustment of prices due to substitution should not take place instantaneously, as search for houses is time consuming. Thus, the gradual adjustment of prices brings up temporal distance as further dimension of excess comovements in local housing markets. Consequently, the order of magnitude of comovement should not only be affected by the spatial distance between homes, but also by the temporal distance of trades. With increasing distance in the temporal sense, prices of neighboring substitutes become less informative for contemporaneous price movements.

In order to find the best price estimate, i.e., the fair market value, for a home in a given location, agents on both the seller and buyer side have to trade of substitutability (i.e, spatial distance) against timeliness for current market movements (i.e, temporal distance). In other words, if the physical distance between homes is small, agents should be willing to accept a greater time distance for a trade to be used as a reference.

In sum, we test three main predictions. First, positive excess returns from nearby homes should lead to positive excess returns for a given home. Second, the strength of these excess comovements should die out with increasing spatial and temporal distance between traded homes. Third, when physical distance is small, excess comovements should be more persistent in the time dimension.

1.3. Empirical results

For our empirical analysis, we need to determine a few parameters for our model introduced in Subsection 1.2.4. Specifically, we need to choose a meaningful number of distinct neighborhoods. In particular, we want to understand whether excess comovements are strongest in the first-order neighborhood and whether they die out in more distant neighborhoods. We set the number of neighborhoods to $K = 6$, which is a reasonable choice between the longest price-relevant distance and the identical sizes of neighborhoods.[18]

We further need to choose the maximum number of days prior to our trade, T, such that trades on other properties should reasonably have the potential of affecting a home's price. Hence, we consider the persistence of price comovements not only in the

[18]Empirically, it turns out that a larger number of neighborhoods does not further contribute significantly to explaining house prices, while a smaller number does not allow us to fully capture the decay in the comovement magnitude with increasing distance.

spatial, but also in the temporal sense. The choice of T is driven by a tradeoff between two opposing objectives. On the one hand, we want to estimate price comovements as precisely as possible, suggesting that we should use as much past data as possible. On the other hand, the precision can be reduced by using outdated observations that may have little informational content for present prices, among others, because the information is already incorporated in more recent prices. We set $T = 180$ for three main reasons. First, gathering information in the housing market costs more time than for example gathering information about the stock market. Second, finding a buyer for a given home typically takes time. Third, our choice of about half a year provides us with a reasonable number of observations to estimate effects with good precision. We document the robustness of our results to the choice of T in Section 1.4.

1.3.1. Excess comovements

In this section, we provide empirical evidence on the existence and strength of excess comovements in residential housing markets. Table 1.3.1 summarizes the results of five Maximum Likelihood regressions explaining the annualized excess returns of repeat sales relative to trades on Manhattan Island. The first-order neighborhood relates to trades in the same building. Second-, third-, fourth-, fifth-, and sixth-order neighborhoods are less than 500, 500 to 707, 707 to less than 866, 866 to less than 1,000, and 1,000 to less than 1,118 feet. Our choice of distances from the traded homes is motivated by the goal to build neighborhoods of identical sizes in order to end up with similar numbers of traded homes in every neighborhood. Locational controls are our measure for liquidity, as well as dummies indicating Central Park view, Central Park walking distance, a very close subway station, a close subway station, short distance to Times Square, medium distance to Times Square, short distance to the NYSE, medium distance to the NYSE, waterfront view, and waterfront walking distance. Transaction-specific controls include two dummy variables indicating that a resale took place within one year, or between one and two years, respectively, log inflation-adjusted prior sale price and its square, two dummies indicating whether the seller or buyer of the property is a corporation, and a dummy indicating whether the property is owner-occupied. Macro-financial controls are lagged GDP growth, lagged unemployment growth, and lagged percentage interest rate change. Fixed effects are on the zip-code level (zip) or the zip-code-year level (zip-year).[19]

[19]Similarly to Campbell et al. (2011) who use census-tract-year clusters, we cluster standard errors over the zip-code-year level. In Table 1.B.2 in our Online Appendix, we show that our results are robust to alternative ways of clustering, e.g., the two-way clustering of Cameron et al. (2011).

1.3.1.1. Excess comovements and spatial distance

In this section, we put emphasis on the spatial dimension of price comovements. As outlined in Subsection 1.2.4, prices should exhibit excess comovements and these comovements should decay in magnitude as the distance between homes increases. From Table 1.3.1, the coefficients for neighborhoods one to five are all positive and significant. That is, Table 1.3.1 confirms the existence of excess comovements in regular sales. From the first- to the sixth-order neighborhood the coefficients generally decrease, indicating that comovements decrease with increasing distance between two traded homes. Coefficients are monotonically decreasing in space, except for the third-order neighborhood, for which the strength of comovements is of slightly lower magnitude.

For all specifications in Table 1.3.1, the sharpest decline in excess comovements is observed for the transition from the first- to the second-order-neighborhood, for which coefficients drop by at least 63%. This result should be mainly driven through two channels. First, trades within the same building should be among the closest substitutes.[20] Second, within the first-order neighborhood both the transmission of information via informal channels, such as chats among neighbors, but also active search for information, should be most intense.[21]

Local price movements should generally be driven by location-specific events. It is therefore important to control for them. A comparison of columns (1) and (2) in Table 1.3.1 reveals that after including our locational controls and controlling for zip-code fixed effects, the coefficients generally decrease, but remain highly significant for the first five neighborhoods. That is, even after controlling for location-specific events, there is still a strong informational content in excess house price movements in the closest neighborhoods.

However, our results also reveal that the coefficient for the most remote, the sixth-order neighborhood, becomes close to zero and insignificant. In other words, our local controls and zip-code fixed effects already capture local price trends quite well.

[20] To investigate the substitution channel in more detail, we also explored a setting in which we identified co-operatives. Two dwellings within a given co-op tend to be more homogeneous than two random dwellings. Co-ops often span over several buildings or even an entire building block. Hence, co-op homes in adjacent buildings should be better substitutes than non-co-op homes. Our results in Table 1.B.3 in the Online Appendix reveal that excess comovements in second to fifth-order neighborhood from co-op homes are stronger and within-building excess comovements are weaker, since the excess comovements in the adjacent buildings already capture some of the effects.

[21] A further reason for stronger comovements within the same building could be rent stabilization, which applies to entire buildings when built before 1974 containing more than six units. Rent stabilization should have only little effect on price comovements in our work, since we use repeat sales and the stabilization should be priced into the initial purchase.

Table 1.3.1
Estimation results, base case

Variable name	(1)	(2)	(3)	(4)	(5)
First-order neighborhood	0.27***	0.25***	0.23***	0.23***	0.21***
	(18.3)	(17.3)	(16.6)	(16.4)	(16.7)
Second-order neighborhood	0.10***	0.06***	0.06***	0.06***	0.04***
	(9.7)	(6.1)	(6.0)	(5.7)	(4.0)
Third-order neighborhood	0.07***	0.04***	0.04***	0.04***	0.03**
	(7.1)	(3.9)	(4.2)	(4.0)	(2.9)
Fourth-order neighborhood	0.09***	0.05***	0.05***	0.05***	0.04***
	(9.7)	(6.0)	(5.8)	(5.5)	(4.9)
Fifth-order neighborhood	0.07***	0.04***	0.04**	0.03**	0.02*
	(5.9)	(3.5)	(3.1)	(2.9)	(2.2)
Sixth-order neighborhood	0.05***	0.02	0.02	0.01	0.002
	(4.8)	(1.6)	(1.6)	(1.3)	(0.1)
ln(1+Liquidity)		-0.01***	-0.002***	-0.002**	-0.001
		(-6.7)	(-3.5)	(-2.7)	(-0.8)
Central Park view		0.01**	0.01*	0.01*	0.01*
		(2.6)	(2.4)	(2.4)	(2.5)
Central Park walking distance		0.0004	-0.002	-0.002	-0.001
		(0.2)	(-0.7)	(-0.7)	(-0.2)
Very close subway		0.003	0.003	0.003	0.003
		(1.6)	(1.5)	(1.5)	(1.6)
Close subway		0.003***	0.003***	0.003***	0.003***
		(3.5)	(4.0)	(4.0)	(4.3)
Short distance Times Square		-0.01*	-0.01*	-0.005	-0.01*
		(-2.3)	(-2.0)	(-1.9)	(-2.5)
Medium distance Times Square		0.002	-0.001	-0.001	-0.002
		(0.7)	(-0.4)	(-0.5)	(-0.9)
Short distance NYSE		-0.002	-0.004	-0.004	-0.01*
		(-0.4)	(-1.0)	(-1.0)	(-2.5)
Medium distance NYSE		-0.01***	-0.01***	-0.01***	-0.02***
		(-5.2)	(-5.0)	(-4.7)	(-6.3)
Waterfront view		0.01**	0.004	0.004	0.0005
		(2.8)	(1.4)	(1.4)	(0.2)
Waterfront walking distance		-0.01***	-0.005***	-0.005***	-0.003
		(-5.3)	(-3.5)	(-3.3)	(-1.7)
Dummy one year			0.06***	0.06***	0.05***
			(10.6)	(10.6)	(8.7)
Dummy two years			0.03***	0.03***	0.02***
			(17.4)	(17.3)	(12.9)
ln(Price)			-0.18***	-0.18***	-0.17***
			(-10.9)	(-10.0)	(-11.3)
$\ln(\text{Price})^2/100$			0.62***	0.62***	0.61***
			(10.7)	(9.8)	(11.1)
Seller corporation			0.01***	0.01***	0.01***
			(5.4)	(5.5)	(5.7)
Buyer corporation			0.01***	0.01***	0.01***
			(10.2)	(10.2)	(11.4)
Owner-occupied			-0.001	-0.001*	-0.001*
			(-1.9)	(-2.4)	(-2.4)
Lagged GDP growth				-0.001	-0.05
				(-0.02)	(-0.9)
Lagged unemployment growth				0.1***	0.06**
				(4.5)	(2.6)
Lagged interest change				-0.01	-0.01
				(-0.6)	(-1.3)
Fixed effects	no	zip	zip	zip	zip-year
Akaike criterion	-117,834	-118,571	-120,729	-120,771	-121,762

This table summarizes the results of Maximum Likelihood regressions explaining the annualized excess return of repeat sales relative to trades on Manhattan Island. The first-order neighborhood relates to trades in the same building. Second-, third-, fourth-, fifth-, and sixth-order neighborhoods have distances to the traded home of less than 500 feet, 500 to less than 707 feet, 707 to less than 866 feet, 866 to less than 1,000 feet, 1,000 to less than 1,118 feet, respectively. For further variable descriptions see Table 1.2.1. Fixed effects are on the zip-code level (zip) or the zip-code-year level (zip-year). The figures in parentheses report *t*-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

Furthermore, the insignificance of the coefficient for the sixth-order neighborhood in column (2) points to two conclusions: First, beyond general price trends, the sixth-order neighborhood no longer contains price information. Second, the reduction in the coefficients for the first- to fifth-order neighborhood largely reflects the removal of the location- and zip-code-specific events. The locational controls and the zip-code fixed effects thus should not only capture the general price movement in the sixth-, but also in the first- to fifth-order neighborhoods very well. Changes in our coefficients in the transition from column (2) to (3), in which we add transaction-specific controls, are rather small. In the transition from column (3) to (4), in which we include macro-financial controls, these changes are even smaller, indicating that the excess returns, which our work builds on, already capture the effects of macroeconomic events very well.

Column (5) reports the estimates for our full specification. Compared to the model presented in column (4), we include zip-code-year fixed effects as opposed to zip-code fixed effects. From column (5), 21% of the increase in the annualized excess return in the first-order neighborhood are reflected in future home prices. For example, a one standard deviation increase in the annualized excess return in the first-order neighborhood, i.e., an increase in the annualized excess return by about 6%, leads to an increase in the expected annualized excess return of future home prices of about 1.3%. For an average holding period of about five years, the expected excess return is then about 6.5%. For the second- to fourth-order neighborhoods these effects are around 80% weaker than for the first-order neighborhood. Here, a one standard deviation increase in the second- to fourth-order neighborhood's excess return leads to an increase in the expected future excess return after the typical holding period of five years of 0.9 to 1.2 percent, respectively. With past excess returns of identical signs in the first- to fifth-order neighborhoods, the effects accumulate and expected future excess returns can be even higher. For example, a one standard deviation increase in all five neighborhoods leads to an increase in the expected future excess return of around 11% over five years.

1.3.1.2. Excess comovements and temporal distance

Our results in the previous section document that excess comovements exist in the spatial dimension, but are dying out with increasing distance between traded properties. In this section, we ask whether in addition to the spatial dimension, excess comovements also exist over longer time horizons, i.e., above systematic autocorrelation. Specifically, we investigate whether adding more lagged excess returns from previous periods has additional predictive power for present excess returns, and, if so, whether the predictive power is decaying with increasing temporal distance. More technically, instead of investigating the informational content of only the most recent $T = 180$ days, we analyze the excess comovements of prices from multiple lags of intervals of length T. Accordingly, to be included in lag 1, a neighboring trade should have been settled in the most recent 180 days prior to the respective sale, for lag 2 during the most recent 181 to 360 days, etc. As for all following tables, we only show results for our most advanced specification, including all sets of control variables, as well as zip-code-year fixed effect dummies. To facilitate the interpretation of our results, we present our results graphically in Figure 1.3.1 for five mutually exclusive time lags and the six neighborhoods from our base case (the regression output is reported in Table 1.B.1 in the Online Appendix).

Figure 1.3.1 reveals that excess comovements in residential house prices do not only exist in the spatial dimension, but also over longer time horizons. Lagged excess returns extending beyond the first lag have a strong predictive power for present excess returns – particularly the first-order neighborhood. These effects are decreasing with increasing temporal distance and are dying out completely in all but the first-order neighborhood. Compared to our results from Table 1.3.1 with only one temporal lag of excess returns, the point estimates for the first-order neighborhood are smaller, reflecting that the additional lags are already picking up some of the effects. The persistence of comovements in the first-order neighborhood reveals that when spatial distance between properties is small, excess comovements die out slower in the temporal dimension.

Our results from Figure 1.3.1 suggest implications for the efficiency of local housing markets: Even conditional on the most recent, price movements from greater temporal distances are reflected in excess returns today to a both economically and statistically significant extent. Our results thus indicate that on the local level, information is processed very slowly, suggesting an explanation for the well-documented autocorrelation of residential house prices on the macro level.

1.3.1.3. Test on information frictions

So far, our results reveal the existence of both economically and statistically significant excess comovements in residential house prices. However, it is important to understand in more detail what causes these excess comovements. As we already argued above, information frictions may play an important role in that regard.

Figure 1.3.1
Excess comovements over spatial and temporal distance

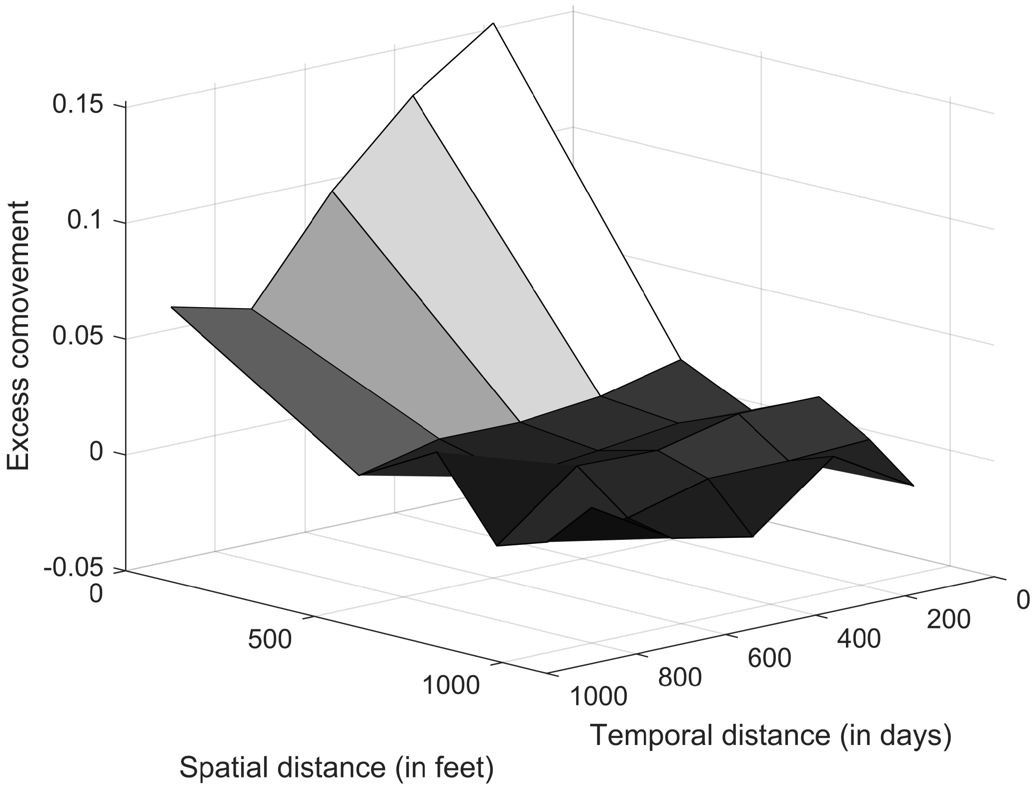

This table visualizes excess comovements over spatial distance, measured in feet, and temporal distance, measured in days. The degree of excess comovements as given on the vertical axis is estimated by extending Equation (1.2.2) by five temporal lags, each of 180 days, such that 6×5 single excess comovement coefficients are estimated. Table 1.B.1 in the Online Appendix provides the exact regression coefficients with corresponding standard errors and significance tests.

To investigate the role of information frictions more formally, in this section we exploit two distinct dates that relate to each individual transaction: the recording date and the contract date. The contract date is the date, at which the contract was signed. The recording date, on the other hand, is the date, at which the information about the transaction is sent to the county's recorder's office. The difference between these two dates is important, because only information that is registered can be made publicly available. Consequently, market participants can only use the information on sales prices after their respective recording dates. That is, the different information available to the public at the sales and the recording date allow us to directly test to which extent information frictions play an important role for excess comovements in residential house prices. More specifically, if the information channel is in play, excess comovements should be weaker prior to the information about the trade becoming public, i.e., before the recording date. That is, the differences between sales and recording dates allow us to ask whether the actual trade on the sales date, or the information about the trade on the recording date are more important for excess comovements in residential house prices.

To test the importance of the information channel more formally, for every trade i, we split the trades over the past $T = 180$ days in each neighborhood into two groups. First, we calculate an average excess return from all observations that have been recorded before home i was sold, i.e., trades that have been publicly available before home i was traded. Second, we calculate an average excess return of homes that have been traded, but not yet recorded. Only in the first set, information about the respective other trades is publicly available. Hence, if the information channel is in play, excess comovements should be stronger in the first set.

Table 1.3.2 reports results from regressions using the two different sets of average excess returns in the neighborhoods. The column marked "Recorded sales" reports results for the first set, in which excess returns are computed using all other past trades in the neighborhoods that have been recorded by the date of a home's sales date. The column "Non-recorded sales" reports results for the second set, in which excess returns are computed using all other past trades in the neighborhoods that have been sold (but not yet recorded) by the date of a home's sales date. From Table 1.3.2, excess comovements are substantially stronger for recorded than non-recorded sales. If information would not play any role, the effects should be of similar size - if ever, effects should be slightly stronger for non-recorded sales, reflecting that the excess returns computed from the non-recorded sales include even more recent trades, namely those trades that have already occurred, but are not yet recorded. Excess comovements for non-recorded sales are largely insignificant, except for the first-order neighborhood, which may reflect that within a building information about sales prices becomes available more quickly via informal channels. Overall, our results in this section reveal that variation in information frictions is important for explaining excess comovements in residential house prices.

1.3.2. Excess comovements over market cycles

Having demonstrated the local nature of excess comovements, our next step is to ask whether the order of magnitude of excess comovements varies with phases of the housing market cycle, i.e., whether the strength of excess comovements and the distance, over which they are measurable, differs between good states with generally

Table 1.3.2
Estimation results, test on information frictions

Variable name	Recorded sales	Non-recorded sales
First-order neighborhood	0.20***	0.09***
	(15.8)	(7.1)
Second-order neighborhood	0.04***	0.005
	(3.9)	(0.5)
Third-order neighborhood	0.03**	0.01
	(3.1)	(1.7)
Fourth-order neighborhood	0.04***	0.02*
	(4.4)	(2.4)
Fifth-order neighborhood	0.02*	-0.02
	(2.0)	(-1.7)
Sixth-order neighborhood	0.01	-0.001
	(0.5)	(-0.1)
Fixed effects	zip-year	
Akaike criterion	-121,796	

This table summarizes the results of a single Maximum Likelihood regression explaining the annualized excess return of repeat sales relative to trades on Manhattan Island. The coefficients describe the comovement of individual trades with average excess returns realized within the past $T = 180$ days in neighborhoods of different orders. For this analysis, the average excess returns from the base case are divided into two categories. The first category are "Recorded sales", i.e., sales that have been reported to the county's recorder's office (and are thus publicly available). The second category contains sales which have already been completed, but have not yet been recorded. The first-order neighborhood relates to trades in the same building. Second-, third-, fourth-, fifth-, and sixth-order neighborhoods have distances to the traded home of less than 500 feet, 500 to less than 707 feet, 707 to less than 866 feet, 866 to less than 1,000 feet, 1,000 to less than 1,118 feet, respectively. The locational, transaction-specific and macro-financial control variables used in this regression are defined in Subsection 1.2.3. Fixed effects are on the zip-code-year level (zip-year). The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, increasing house prices and other stages of the housing market cycle. We differentiate between aggregate (macro) and local (micro) market trends. For the macro perspective, we define boom and non-boom periods ex-post using our price index for Manhattan Island from Figure 1.2.2. According to this index, the boom in the early 2000s ends in October 2005, and house prices start booming again in March 2013. We therefore define the period from November 2005 to February 2013 as the non-boom period and the remaining months as the boom period.[22]

Taking on the local, micro perspective, we define local positive (negative) markets according to the sign of the average excess returns in each neighborhood from the past T days. Consequently, the k-th order neighborhood is above (below) average markets when we observe a positive (negative) average excess return for the past $T = 180$ days. Put differently, we ask whether positive and negative information from past excess returns affect future excess returns asymmetrically.[23] Formally, we extend our empirical application from Equation (1.2.2) to:

$$r^e_{i,t,t-,z} = \alpha_z + \sum_{s=1}^{2}\sum_{k=1}^{K} \rho_{k,s} \times \bar{r}^e_{i,k} \times \mathbb{1}_{t\cap s} + \delta_{a(t),z} - \delta_{a(t-),z} + X_{i,t}\beta + \epsilon_{i,t,t-,z}, \tag{1.3.1}$$

in which the two stages of the cycle are defined by $s \in \{\text{boom,non-boom}\}$ for the macro, and $s \in \{\text{above average, below average}\}$ for the micro trend, and $\mathbb{1}_{t\cap s}$ is an indicator function that equals one if the housing market is in stage s at time t.

Table 1.3.3 summarizes the results of two separate regressions explaining the annualized excess return of repeat sales relative to trades on Manhattan Island in boom and non-boom periods for the macro panel (Panel A: Macro) and positive (negative) trends on the micro level (Panel B: Micro). For both regressions, the full set of control variables as well as zip-code-year fixed effect dummies are used.

From Panel A of Table 1.3.3, coefficients are generally smaller in booming stages of the housing market cycle than in other stages. That is, consistent with Cotter et al. (2015), excess comovements in residential house prices seem to be stronger in markets

[22]Using the publicly available S&P CoreLogic Case-Shiller New York City condominium index (download link https://us.spindices.com/documents/additionalinfo/20170926-589149/589149_cs-condoindices-0926.xls?force_download=true, retrieved on May 18, 2020), we identify a non-boom period between February 2006 and April 2012. Using the S&P Case-Shiller National home price index, we identify a period from March 2006 to March 2012. Similarly, we characterize our non-boom period using a purely liquidity-based approach building on the number of observed trades. In Subsection 1.4.2, we document that our results are robust to all of these alternative specifications.

[23]It is important to note, that a negative excess return does not necessarily mean a loss for the seller. A negative sign only indicates that the performance of the trade was smaller than the performance of the market.

Table 1.3.3
Estimation results, cycle dependencies

	Market trend			
	Panel A: Macro		Panel B: Micro	
Neighborhood	Boom	Non-boom	Above avg.	Below avg.
First-order neigh.	0.19***	0.24***	0.17***	0.32***
	(10.3)	(14.6)	(10.4)	(14.7)
Second-order neigh.	0.04*	0.04***	0.02	0.10***
	(2.3)	(3.4)	(1.6)	(4.3)
Third-order neigh.	0.03*	0.03*	0.03**	0.02
	(2.1)	(2.1)	(2.8)	(0.07)
Fourth-order neigh.	0.04***	0.04***	0.04***	0.05*
	(3.7)	(3.3)	(3.5)	(2.2)
Fifth-order neigh.	0.03	0.02	0.03*	0.01
	(1.7)	(1.6)	(2.2)	(0.5)
Sixth-order neigh.	-0.001	0.003	0.005	-0.01
	(-0.1)	(0.3)	(0.3)	(-0.5)
LR test (p-value)	0.000		0.000	
Akaike criterion	-121,763		-121,819	

This table summarizes Maximum Likelihood regression results on two separate regressions explaining the annualized excess return of repeat sales relative to trades on Manhattan Island. The two regressions depict explanatory power of neighboring excess returns conditional on being in a specific phase of a macro, and a micro cycle. For the macro cycle, we define a boom (January 2004 to October 2005, and March 2015 to December 2015) and a non-boom (November 2005 to February 2013) period. For the micro cycle, we define an above (below) average local market by a positive (negative) average excess return in the past $T = 180$ days. The first-order neighborhood relates to trades in the same building. Second-, third-, fourth-, fifth-, and sixth-order neighborhoods have distances to the traded home of less than 500 feet, 500 to less than 707 feet, 707 to less than 866 feet, 866 to less than 1,000 feet, 1,000 to less than 1,118 feet, respectively. The locational, transaction-specific and macro-financial control variables used for the two regressions are defined in Subsection 1.2.3. Fixed effects are on the zip-code-year level. The Likelihood Ratio test (LR test) is a test of joint equality of neighborhood coefficients, i.e., under the null hypothesis that $\rho_{1,m} = \rho_{1,nm}, \ldots, \rho_{6,m} = \rho_{6,nm}$, where m (nm) denotes being (not being) in micro or macro phase m, respectively. The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

with falling prices. Similarly, from Panel B, coefficients for above-average excess returns are generally smaller than for below-average excess returns. Therefore, excess comovements in residential house prices seem to be stronger for negative deviations from the market than for positive deviations. For instance, about a third of a negative excess return in the first-order neighborhood is reflected in the excess return of an existing trade, whereas only less than 17% of a positive excess return is. The observed asymmetry of local house price comovements can be interpreted as further evidence for the information channel, as the substitution channel would imply rather similar comovements for under- and overperforming areas.

A systematic investigation of local over- and underperformance is challenging, because the number of trades in the same building within a shorter time period is typically small. To investigate the persistence of local over- and underperformance, we therefore ask whether locally, on the census tract level, census tracts with local over- or underperformance are more likely to again over- or underperform in the following years. We find that conditional on a positive average excess return in the past year, the probability for a positive average excess return in the current year is 25% higher after controlling for zip-code fixed effects. An extended analysis with three lags of one year each, again controlling for zip-code fixed effects, leaves us with coefficients of 22.5% for one year, 11.7% for two years, which are both significant at the 1% level, and 3.5%, which is insignificant at the 10%, for three years (standard errors are clustered on the zip-code level). Consequently, local over- and underperformance seem to be systematic for up to two years. For three years, we no longer find a systematic effect. Hence, our results suggest that local over- and underperformance is a systematic phenomenon, which might reflect positive or negative attributes of the neighborhoods.

Having documented differences in how excess returns comove with past ones in positive and negative states of the housing market cycle on both the macro and micro level in Table 1.3.3, we next investigate these two effects jointly. In particular, we want to shed light on whether the generally stronger negative persistence in bad market environments is further amplified by a negative macro trend or not.

Table 1.3.4 presents results from a single regression for four mutually exclusive sets of variables, which we select according to the phase of the housing market cycle on the macro and the micro level as in Table 1.3.3. Irrespective of the stage of the housing market cycle on the macro level, underperformance on the micro level is generally more persistent than local overperformance, particularly in the first- and second-order neighborhoods.

Table 1.3.4
Estimation results, cross-cycle dependencies

Macro level	Boom		Non-boom	
Micro level	Above avg.	Below avg.	Above avg.	Below avg.
First-order neigh.	0.10***	0.36***	0.23***	0.26***
	(4.5)	(12.5)	(11.4)	(8.6)
Second-order neigh.	-0.004	0.12***	0.03*	0.07*
	(-0.2)	(4.0)	(2.2)	(2.3)
Third-order neigh.	0.04*	0.01	0.02	0.02
	(2.2)	(0.4)	(1.8)	(0.7)
Fourth-order neigh.	0.03*	0.07**	0.05**	0.02
	(2.1)	(2.8)	(3.1)	(0.5)
Fifth-order neigh.	0.03	0.03	0.03	-0.01
	(1.5)	(0.9)	(1.8)	(-0.2)
Sixth-order neigh.	-0.01	0.004	0.01	-0.03
	(-0.2)	(0.1)	(0.7)	(-0.9)
Akaike criterion	-121,763			

This table summarizes Maximum Likelihood regression results for a single regression explaining the annualized excess return of repeat sales relative to trades on Manhattan Island. The average past excess returns within neighborhoods are divided according to the current state of the macro and micro cycle in which each transaction was settled. The macro state describes booming and non-booming periods of the Manhattan Island market, where the non-booming period is set to November 2005 to February 2013. Above (below) average returns on the micro level are defined as average excess returns in a neighborhood from the past $T = 180$ days being positive (negative). The first-order neighborhood relates to trades in the same building. Second-, third-, fourth-, fifth-, and sixth-order neighborhoods have distances to the traded home of less than 500 feet, 500 to less than 707 feet, 707 to less than 866 feet, 866 to less than 1,000 feet, 1,000 to less than 1,118 feet, respectively. The locational, transaction-specific and macro-financial control variables used in this regression are defined in Subsection 1.2.3. Fixed effects are on the zip-code-year level (zip-year). The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

Excess comovements are strongest for below-average returns on the micro level in booming stages of the housing market cycle. That is, local underperformance is most persistent when house prices are generally increasing. Simultaneously, excess comovements from above-average local returns in a booming market are relatively weak. In a non-booming stage of the housing market cycle, differences between the strength of comovements induced by local above- and below-average performance are much weaker.

In sum, our results in Table 1.3.4 thus suggest that there is more heterogeneity in terms of the local evolution of house prices in booming stages of the housing market cycle than in non-booming ones.

1.4. Robustness analysis

This section documents the robustness of our key findings with respect to various assumptions. Subsection 1.4.1 provides evidence for our base case parameter setting, in which we do not distinguish between boom and non-boom periods. Subsection 1.4.2 provides results for different definitions of the boom and non-boom periods. Additional robustness results are placed in the Online Appendix. This Online Appendix documents that our results are robust to alternative clustering of standard errors (Table 1.B.2), when nearby homes are part of the same cooperative and thereby better substitutes (Table 1.B.3), to the omission of controls related to Central Park, Times Square, and the New York Stock Exchange, risk-adjusting returns and restricting the neighborhoods to the same zip-code (Table 1.B.4), and accounting for school rezoning, new constructions, and renovations (Table 1.B.5). Furthermore, we investigate whether there is more explanatory power in a neighborhood's excess return if trades in respective lower-order neighborhoods are unavailable (Table 1.B.6), and provide estimates of SAR coefficients for comparison (Table 1.B.7).

1.4.1. Robustness of base case results

With our results in Table 1.3.1, we demonstrate that excess comovements in residential house prices are a highly local phenomenon. In this section, we demonstrate the robustness of our results with regard to four key dimensions and report these results in Table 1.4.1. To simplify the comparison with our base-case results, we repeat the results from Table 1.3.1 in Panel A of Table 1.4.1.

In Panel B of Table 1.4.1, we allow for a different number of past days used to compute average excess returns in the neighborhoods. In our base case parameter

Table 1.4.1
Robustness, base case

Neighborhood order	First	Second	Third	Fourth	Fifth	Sixth
Panel A: Base case						
	0.21***	0.04***	0.03**	0.04***	0.02*	0.002
	(16.7)	(4.0)	(2.9)	(4.9)	(2.2)	(0.1)
Panel B: Varying computation of excess returns in neighborhoods						
T = 240	0.22***	0.05***	0.03*	0.04***	0.03**	-0.003
	(18.6)	(4.2)	(2.4)	(4.9)	(2.8)	(-0.2)
T = 120	0.19***	0.03**	0.03**	0.04***	0.02	-0.02
	(13.6)	(3.1)	(3.2)	(5.3)	(1.8)	(-1.8)
Panel C: Varying neighborhood definitions						
333 feet	0.21***	0.04***	0.02	0.04***	0.02**	0.03**
	(16.7)	(3.7)	(1.6)	(4.1)	(2.6)	(3.2)
0.1, 0.25 miles	0.21***	0.05***	0.07***			
	(16.6)	(5.4)	(3.9)			
NYC neighborhoods	0.21***	0.17***	-0.01			
	(16.5)	(4.7)	(-0.2)			
Panel D: Varying maximum holding period						
Seven years	0.21***	0.03**	0.01	0.03**	0.01	-0.0001
	(16.1)	(2.9)	(1.5)	(3.1)	(0.8)	(-0.01)
Ten years	0.21***	0.04***	0.02**	0.04***	0.02*	0.003
	(16.9)	(4.1)	(2.6)	(3.8)	(2.2)	(0.2)
Panel E: City block and waterfront						
City block metric	0.21***	0.04***	0.04***	0.04***	0.03**	0.01
	(16.6)	(3.8)	(4.4)	(3.8)	(2.8)	(1.4)
Exclude waterfront obs.	0.21***	0.05***	0.03**	0.03***	0.03*	-0.003
	(15.9)	(4.3)	(2.9)	(3.6)	(2.2)	(-0.3)

This table documents the robustness of our key results with respect to various assumptions. Panel B presents results when varying the definition of T, the maximum number of past days used to compute average excess returns in the neighborhood. Panel C presents results for different neighborhood definitions. In the row "333 feet", the second-order neighborhood is defined by a maximum distance of 333 feet. The subsequent neighborhoods are defined such that the area within each neighborhood is the same as in the second-order, yielding borders of 470, 576, 666, and 744 feet. In the row marked "0.1, 0.25 miles", the second- and third-order neighborhoods are defined by maximum distances of 0.1 and 0.25 miles from the traded home, i.e., 528 and 1,320 feet, respectively. "NYC Neighborhoods" depicts results for the case in which the second-order neighborhood is the neighborhood as defined by the City of New York (e.g., Chinatown, Lower East Side, etc.) and the third-order neighborhood are the neighborhoods adjacent to the second-order neighborhood. In Panel D, observations with a holding period of more than seven or ten years, respectively, are excluded. Panel E shows results for a change in the distance measure to the city block metric, and when excluding observations for which the waterfront lies within at least the sixth-order neighborhood (i.e., 1,118 feet). All regressions include the entire set of controls: locational, transaction-specific, and macro-financial. Fixed effects are on the zip-code-year level. The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

setting, we used the past $T = 180$ days, which we consider a good tradeoff between the two opposing goals of having a reasonably larger number of observations and very recent up-to-date observations. In Panel B, we explore the cases in which we set $T = 120$ or $T = 240$ days. Our results for these two cases demonstrate the robustness of our key findings that effects are strongest in the same building, i.e., the first-order neighborhood, remain significant in the second- to fourth-order neighborhood, and fade out for higher-order neighborhoods. Similarly, the point estimates for the strength of comovement in the various neighborhoods are of a very similar order of magnitude.

In Panel C, we vary the definitions of the neighborhoods. In our base case parameter setting, the second-order neighborhood was characterized by a maximum distance from the traded home of not more than 500 feet, roughly corresponding to two blocks. Here, we report results when shrinking this distance measure by two thirds, i.e., to 333 feet. Again, the borders of the higher-order neighborhoods are defined such that the area is the same as in the second-order neighborhood. We also depict results for the case, in which the neighborhoods are defined as in Campbell et al. (2011), i.e., a maximum distance of 0.1 miles, corresponding to 528 feet, for the second-order and 0.25 miles, corresponding to 1,320 feet, for the third-order neighborhood. As in Campbell et al. (2011), we do not account for neighborhoods of a higher order. Finally, we depict results for the case in which the second-order neighborhood is defined by the City of New York (e.g., Chinatown, Lower East Side, etc.) and the third-order neighborhood consists of the corresponding neighborhoods adjacent to the second-order neighborhood.[24]

Our results in Panel C again document the robustness of our key finding that comovements are strongest in the first-order neighborhood. With smaller second- to sixth-order neighborhoods for the former case, results remain significant even in the sixth-order neighborhood, reflecting that the maximum distance of a trade in this neighborhood is 744 feet, corresponding to a trade in the fourth neighborhood in our base-case parameter setting. A more narrow definition of neighborhoods again suffers from the problem of relatively small numbers of historical trades in each of the neighborhoods, which, among others, leads to the coefficient for the third-order neighborhood being insignificant. For instance, the number of historical

[24]Data on the neighborhoods is obtained from New York City Open Data: https://data.cityofnewyork.us/City-Government/Neighborhood-Tabulation-Areas/cpf4-rkhq (retrieved on July 13, 2018), which provides a shape file defining the neighborhoods. The shape file includes a "miscellaneous" area that consists of several dispersed areas, such as parks, cemeteries, etc., that are not related to a particular neighborhood. A few trades in our data fall into this area, but are only a few feet away from the nearest non-"miscellaneous" neighborhood. We assign these observations to the nearest non-"miscellaneous" neighborhood.

trades in this neighborhood decreases by about 65% compared to our base-case parameter setting with wider neighborhoods. Using the definition of New York City neighborhoods yields strong comovements of within-neighborhood excess returns from the first- and second-order neighborhoods, but virtually no connection to adjacent neighborhoods, suggesting that in Manhattan, where adjacent neighborhoods are often very heterogeneous, the preferred locations of households have sharp boundaries. This finding suggests that attention to prices of recently transacted homes is especially attracted within the same neighborhoods. This pattern provides further evidence for the information channel rather than the substitution channel, for which no sharp comovement bounds at neighborhood borders should be observed.

In Panel D, we restrict the maximum holding period to seven and ten years, respectively.[25] Further restricting the maximum holding period to less than seven years leads to such a strong decline in the number of observations that it no longer provides a representative picture of market movements and – due to the lack of this information – predicts largely insignificant effects. Specifically, reducing the maximum holding period to six years removes more than a third of all trades and the information contained in these trades.

In Panel E, we change the distance measure used in the definition of our neighborhoods from the Euclidean to the city-block metric and ask whether our results are affected by excluding observations for which the waterfront lies within at least the sixth-order neighborhood. Intuitively, for such observations, the area covered by higher-order neighborhoods may be smaller than that of smaller-order neighborhoods giving rise to potentially significantly different numbers of past trades in the different neighborhoods. Our results for both cases confirm our key findings that prices comove the strongest in the first-order neighborhood and fade out for the most distant neighborhoods.[26] Further robustness checks can be found in Tables 1.B.4 and 1.B.5 in our Online Appendix.

Overall, our results in this section confirm the robustness of our key findings on house price comovements to various assumptions in our base case parameter setting, in which we do not split the sample into boom and non-boom periods. We next proceed

[25]For shorter holding periods, larger reconstructions and major changes in the neighborhood should be less likely. That is, the repeat-sales approach should yield particularly precise estimates.

[26]We apply the city block metric to proxy commuting distance between properties. This is possible by exploiting the geometric design of Manhattan. We therefore shift the coordinates of the properties in our sample such that the streets approximately align with the lines of longitude and latitude. More precisely, we shift the coordinates (after standardizing) by 35 degrees counterclockwise around the south-east corner of the Central Park. Again, we construct six neighborhoods (of cubic form due to the metric) around each property. To ease comparison with our base case results, each cube encompasses the same area as our base-case circles.

Figure 1.4.1
Identification of non-booming periods

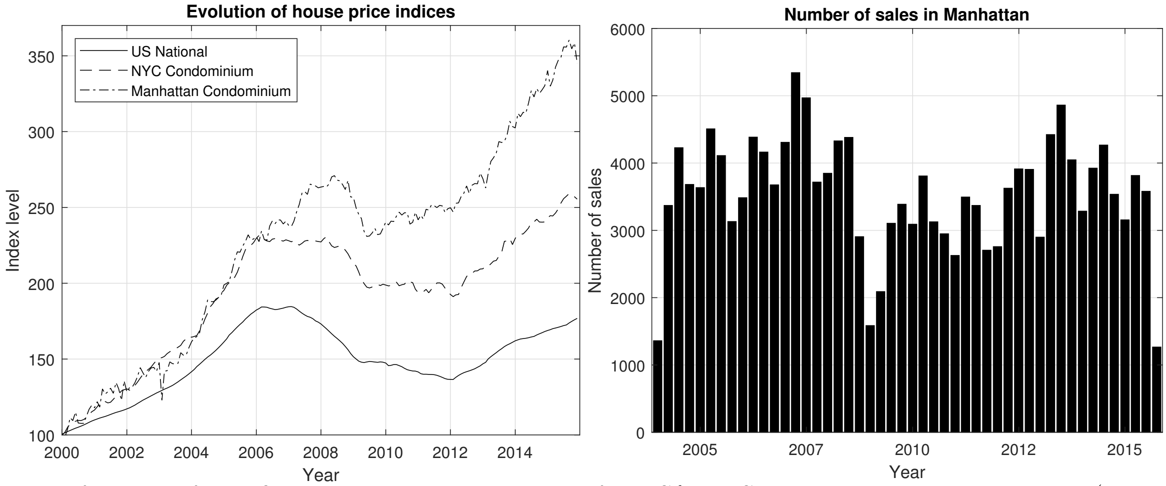

The left panel of this figure depicts the evolution of the S&P US National House Price Index (solid line), the S&P Case-Shiller Condominium Index for New York City (dashed line), and the Manhattan Condominium Index (dotted line) constructed using the methodology of Case and Shiller (1989). Index levels are normalized to 100 in January 2000. The right panel depicts the absolute number of sales of apartments and condominiums on Manhattan Island from the first quarter of 2004 to the fourth quarter of 2015 after removing observations with missing values in sales prices, sales dates, and duplicates.

to demonstrate that our key results on comovements in boom and non-boom periods remain robust when using different criteria to determine these two subperiods.

1.4.2. Robustness of boom versus non-boom

So far, we defined boom and non-boom periods based on our Manhattan Condominium index, constructed using the Case-Shiller methodology (Case and Shiller, 1989). Using this index, our non-boom period lasts from November 2005 to February 2013. We further document that excess comovements in the first-order neighborhood are stronger during non-boom periods and weaker during boom periods and that the other estimates are of comparable magnitude. In this section, we use alternative definitions for the boom and non-boom periods, using different house price indices and a liquidity measure.

The left panel of Figure 1.4.1 depicts the evolution of real house price indices for Manhattan (dotted line), New York (dashed line), and the entire United States (solid line). As before, we define the beginning of a non-boom period as the month in which a previously sharp incline in house prices ends. Likewise, the end of a non-boom period is the month in which a new sharp incline in house prices begins. That is, for the NYC Condominium Index, the non-boom period is March 2006 to April 2012 and for the US National House Price Index, this period is March 2006 to February 2012. The right panel in Figure 1.4.1 depicts the number of sold apartments and condominiums on Manhattan Island after removing observations with missing values in sales prices, sale dates, and duplicates. From this panel, the number of sales declined from 4,381 to 2,906 trades in October 2008 and did not recover systematically before March 2012. As an additional definition for our non-boom period, we therefore use the time period October 2008 to March 2012 as a liquidity-based definition of our non-boom period.

Table 1.4.2 summarizes our results for the different definitions of the non-boom period. For ease of comparison, the results from Table 1.3.3 are repeated in Panel A of Table 1.4.2. Consistent with our key findings from Subsection 1.3.2, our robustness results with different definitions of the non-boom period confirm that during non-boom periods, excess comovements with the first-order neighborhood are stronger. Irrespective of the exact definition of our non-boom period, point estimates for our coefficients are very similar. Our results in Table 1.4.2 thus confirm the finding from Table 1.3.3 that during non-booming periods, excess comovements are stronger for within-building trades. Otherwise, the role of spatial distance is similar for both phases.

1.5. Conclusion

The housing market boom and bust of the early 2000s highlights the importance for a better understanding of the evolution of residential house prices. We contribute to this challenging endeavor by exploring the micro-level evolution of residential house prices, using data from trades on Manhattan Island between 2004 and 2015. In doing so, we also test information frictions more formally by comparing excess comovements at the sales date of a private transaction and the recording date, when the individual sales price becomes publicly available.

We document that even after controlling for monthly aggregate market movements and zip-code-year based price movements, excess comovements in residential house prices are a highly persistent local phenomenon. The strength of these excess comovements vanishes with the distance between traded homes. In addition to these spatial excess comovements, excess comovements in residential house prices also have a persistent temporal dimension. Unlike in stock markets, house prices seem to adjust slowly to new information, and even price movements from more than two years ago still have a significant impact on present price movements. Moreover, local underperformance is more persistent than local overperformance. This phenomenon is particularly strong when house prices on the aggregate level appreciate.

Table 1.4.2
Robustness, boom versus non-boom

	Boom						Non-boom					
Neighborhood order	First	Second	Third	Fourth	Fifth	Sixth	First	Second	Third	Fourth	Fifth	Sixth
Panel A: Base case												
	0.19***	0.04*	0.03*	0.04***	0.03	-0.001	0.24***	0.04***	0.03*	0.04***	0.02	0.003
	(10.3)	(2.3)	(2.1)	(3.7)	(1.7)	(-0.1)	(14.6)	(3.4)	(2.1)	(3.3)	(1.6)	(0.3)
Panel B: Varying non-boom periods												
03/2006 - 04/2012	0.19***	0.04**	0.04*	0.04***	0.04*	0.01	0.25***	0.04**	0.02	0.04**	0.01	-0.01
	(11.6)	(3.0)	(2.6)	(3.4)	(2.4)	(0.5)	(13.2)	(2.8)	(1.5)	(3.2)	(0.5)	(-0.8)
03/2006 - 02/2012	0.19***	0.04**	0.04**	0.04***	0.04*	0.01	0.25***	0.04**	0.02	0.04**	0.003	-0.01
	(11.8)	(2.9)	(2.6)	(3.6)	(2.5)	(0.5)	(13.4)	(2.9)	(1.3)	(3.2)	(0.2)	(-0.8)
10/2008 - 03/2012	0.20***	0.04**	0.03**	0.04***	0.03*	0.004	0.25***	0.06**	0.02	0.04*	0.02	-0.006
	(14.1)	(3.0)	(2.8)	(3.9)	(2.0)	(0.3)	(9.6)	(3.1)	(1.0)	(2.4)	(1.1)	(-0.4)

This table documents the robustness of our key findings with respect to various ways to define boom and non-boom periods in Panel B. Panel A repeats the results for our base-case parameter setting in which boom and non-boom periods (November 2005 to February 2013) are defined using our Manhattan Condominium index, based on Case and Shiller (1989). The row "03/2006 - 04/2012" depicts results for the case in which the non-boom period is set to its counterpart in the S&P/CS NYC Condominium index, i.e., the time period March 2006 to April 2012; the row "03/2006 - 02/2012" for the case in which the non-boom period is set to its counterpart in the S&P/CS US National Home Price index, i.e., the time period from March 2006 to February 2012. The row "10/2008 - 03/2012" reports results when defining the non-boom period using the liquidity-dry-up-period from Figure 1.4.1, i.e., October 2008 to March 2012. All regressions include the entire set of controls: locational, transaction-specific, and macro-financial. Fixed effects are on the zip-code-year level. The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

References

Amromin, G., J. C. Huang, C. Sialm, and E. Zhong (2018). “Complex Mortgages.” *Review of Finance* 22. (6), 1975–2007.

Anenberg, E. and E. Kung (2014). “Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures.” *American Economic Review* 104. (8), 2527–2551.

Ang, A. and J. Chen (2002). “Asymmetric Correlations of Equity Portfolios.” *Journal of Financial Economics* 63. (3), 443–494.

Autor, D. H., C. J. Palmer, and P. A. Pathak (2014). “Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts.” *Journal of Political Economy* 122. (3), 661–717.

Bhattacharya, U., D. Huang, and K. M. Nielsen (2020). “Spillovers in Prices: The Curious Case of Haunted Houses.” *Review of Finance* forthcoming.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). “Robust Inference With Multiway Clustering.” *Journal of Business and Economic Statistics* 29. (2), 238–249.

Campbell, J. Y., S. Giglio, and P. Pathak (2011). “Forced Sales and House Prices.” *American Economic Review* 101. (5), 2108–2131.

Can, A. (1990). “The Measurement of Neighborhood Dynamics in Urban House Prices.” *Economic Geography* 66. (3), 254–272.

Caplin, A. and J. Leahy (2011). “Trading Frictions and House Price Dynamics.” *Journal of Money, Credit and Banking* 43. (7), 283–303.

Case, K. E. and C. J. Mayer (1996). “Housing Price Dynamics within a Metropolitan Area.” *Regional Science and Urban Economics* 26. (3-4), 387–407.

Case, K. E. and R. J. Shiller (1989). “The Efficiency of the Market for Single-Family Homes.” *American Economic Review* 79. (1), 125–137.

Chambers, M., C. Garriga, and D. E. Schlagenhauf (2009). “Accounting for Changes in the Homeownership Rate.” *International Economic Review* 50. (3), 677–726.

Chay, K. Y. and M. Greenstone (2005). “Does Air Quality Matter? Evidence from the Housing Market.” *Journal of Political Economy* 113. (2), 376–424.

Cohen, J. P. and J. Zabel (2018). “Local House Price Diffusion.” *Real Estate Economics* 48. (3), 710–743.

Cotter, J., S. Gabriel, and R. Roll (2015). “Can Housing Risk Be Diversified? A Cautionary Tale from the Housing Boom and Bust.” *Review of Financial Studies* 38. (3), 913–936.

Diamond, R. and T. McQuade (2019). "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development." *Journal of Political Economy* 127. (3), 1063–1117.

Fischer, M., R. Füss, and S. Stehle (2021). "Local House Price Comovements." *Real Estate Economics* 49. (S1), 169–198.

Gerardi, K., E. Rosenblatt, P. S. Willen, and V. Yao (2015). "Foreclosure Externalities: New Evidence." *Journal of Urban Economics* 87, 42–56.

Guerrieri, V., D. Hartley, and E. Hurst (2013). "Endogenous Gentrification and Housing Price Dynamics." *Journal of Public Economics* 100, 45–60.

Gupta, A. (2019). "Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults." *Journal of Finance* 74. (5), 2249–2301.

Guren, A. M. and T. J. McQuade (2020). "How Do Foreclosures Exacerbate Housing Downturns?" *Review of Economic Studies* 87. (3), 1331–1364.

Harding, J. P., E. Rosenblatt, and V. W. Yao (2009). "The Contagion Effect of Foreclosed Properties." *Journal of Urban Economics* 66. (3), 164–178.

Kallberg, J. G., C. H. Liu, and P. Pasquariello (2014). "On the Price Comovement of US Residential Real Estate Markets." *Real Estate Economics* 42. (1), 71–108.

Landier, A., D. Sraer, and D. Thesmar (2017). "Banking Integration and House Price Co-Movement." *Journal of Financial Economics* 125. (1), 1–25.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). "The Housing Market(s) of San Diego." *American Economic Review* 105. (4), 1371–1407.

Murfin, J. and R. Pratt (2019). "Comparables Pricing." *Review of Financial Studies* 32. (2), 688–737.

Rossi-Hansberg, E., P.-D. Sarte, and R. Owens III (2010). "Housing Externalities." *Journal of Political Economy* 118. (3), 485–535.

Szumilo, N. (2020). "Prices of Peers: Identifying Endogenous Price Effects in the Housing Market." *Economic Journal* forthcoming.

Tu, Y., S.-M. Yu, and H. Sun (2004). "Transaction-Based Office Price Indexes: A Spatiotemporal Modeling Approach." *Real Estate Economics* 32. (2), 297–328.

Appendix 1.A Clustering of zip-codes

The zip-codes have been clustered the following way:

- 10001 & 10011 (Chelsea and Clinton)
- 10002 & 10003 & 10009 (Lower East Side)
- 10004 & 10005 & 10006 & 10007 & 10038 & 10280 & 10282 (Lower Manhattan)
- 10012 & 10013 (Greenwich Village/Lower Manhattan)
- 10017 & 10163 (Gramercy Park and Murray Hill)
- 10018 & 10019 & 10036 & 10129 (Chelsea and Clinton)
- 10023 & 10069 (Upper West Side)
- 10026 & 10027 & 10030 & 10037 & 10039 (Central Harlem)
- 10029 & 10035 & 10128 (East Harlem, 10128 is Upper East)
- 10031 & 10032 & 10033 & 10034 & 10040 (Inwood and Washington Heights)

Appendix 1.B Online appendix

Table 1.B.1
Estimation results, time dimension with six neighborhoods

Neighborhood	lag 1	lag 2	lag 3	lag 4	lag 5
First-order neigh.	0.15***	0.13***	0.09***	0.05***	0.06***
	(14.0)	(11.6)	(10.2)	(4.7)	(5.9)
Second-order neigh.	0.03**	0.02*	0.01	0.01	0.01
	(2.6)	(2.2)	(1.7)	(1.8)	(0.8)
Third-order neigh.	0.01	0.01	0.01	0.01	0.02*
	(1.2)	(1.8)	(1.1)	(0.7)	(2.6)
Fourth-order neigh.	0.03**	0.03**	0.02	0.02*	-0.01
	(2.9)	(2.7)	(1.9)	(2.0)	(-1.2)
Fifth-order neigh.	0.01	0.01	0.01	0.001	-0.002
	(1.1)	(1.3)	(1.1)	(0.1)	(-0.3)
Sixth-order neigh.	-0.003	0.02	-0.01	-0.004	0.02*
	(-0.3)	(1.9)	(-1.1)	(-0.4)	(2.3)
AIC			-122,574		

This table shows results when extending our base case, which includes only one temporal lag, to five ones, as visualized in Figure 1.3.1. For these lags, mutually exclusive time intervals of 180 days are set for which the average excess returns are calculated, e.g., for lag 1 the sales in the most recent 180 days are used, for lag 2 the most recent 360 to 181 days. The regression includes the entire set of controls: locational, transaction-specific, and macro-financial. Fixed effects are on the zip-code-year level. The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

Table 1.B.2
Alternatively clustered standard errors

Neighborhood order	Zip-code-year		Two-way		Zip-code		Year	
	s.e.	*t*-stat	s.e.	*t*-stat	s.e.	*t*-stat	s.e.	*t*-stat
First-order neigh.	0.01	16.7	0.02	10.5	0.02	10.5	0.01	21.0
Second-order neigh.	0.01	4.0	0.01	4.2	0.01	4.6	0.01	4.6
Third-order neigh.	0.01	2.9	0.01	2.9	0.01	2.9	0.01	2.9
Fourth-order neigh.	0.01	4.9	0.01	4.2	0.01	4.2	0.01	4.2
Fifth-order neigh.	0.01	2.2	0.01	2.5	0.01	2.5	0.01	2.5
Sixth-order neigh.	0.01	0.1	0.01	0.1	0.01	0.1	0.01	0.1

This table presents standard errors (s.e.) and corresponding t-statistics (t-stat) for different ways of clustering. The estimated model is our base case, presented in column (5) of Table 1.3.1. Columns "Zip-code-year" serve for ease of comparison and presents results when clustering over zip-code-year level as used throughout the main manuscript. Column "Two way" shows results when applying the two-way clustering by Cameron et al. (2011) with zip-code and year dimension. Columns "Zip-code" and "Year" depict standard errors and t-statistics when clustering over zip-code and year, respectively.

Table 1.B.3
Evidence on substitution: co-operatives

Neighborhood	Base case returns	Co-operative returns
First-order neigh.	0.23***	-0.07*
	(13.0)	(-2.6)
Second-order neigh.	0.03*	0.06***
	(2.5)	(3.6)
Third-order neigh.	0.02	0.04**
	(1.6)	(3.1)
Fourth-order neigh.	0.03***	0.05***
	(3.5)	(3.3)
Fifth-order neigh.	0.02	0.04*
	(1.4)	(2.3)
Sixth-order neigh.	-0.001	0.005
	(-0.1)	(0.3)
AIC	-121,816	

This table documents results for a single regression explaining annualized excess returns of homes. The column "Base case returns" depicts estimates for the average excess returns in each neighborhood as for the base case results in Table 1.3.1. The column "Co-operative returns" shows estimates for the comovement of average excess returns of co-operatives with the subsequent excess returns of single co-operative units. We define a building as co-operative if its land use is flagged as co-operative. All regressions include the entire set of controls: locational, transaction-specific, and macro-financial. Fixed effects are on the zip-code-year level. The figures in parentheses report *t*-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

Table 1.B.4
Further robustness checks, base case

Neighborhood	(1) No CP, TS, NYSE	(2) Risk-adjusted	(3) Zip-code restriction
First-order neigh.	0.21***	0.19***	0.21***
	(16.4)	(15.4)	(16.7)
Second-order neigh.	0.05***	0.03***	0.04***
	(4.3)	(3.5)	(3.7)
Third-order neigh.	0.03**	0.03**	0.03**
	(3.2)	(3.2)	(2.7)
Fourth-order neigh.	0.04***	0.04***	0.04***
	(5.0)	(4.5)	(4.4)
Fifth-order neigh.	0.03*	0.03*	0.03*
	(2.4)	(2.4)	(2.3)
Sixth-order neigh.	0.004	0.003	-0.01
	(0.3)	(0.2)	(-0.6)
AIC	-121,710	89,440	-121,753

This table documents further robustness checks on our key results. Column (1) includes all base case control variables apart from controls related to the Central Park (CP), the Times Square (TS), and the New York Stock Exchange (NYSE). Column (2) presents the results when risk-adjusting the excess returns by the annualized standard deviation of the monthly market return between the respective prior sale date and the second sale date. In column (3), excess returns of neighbors are restricted to be in the same zip-code as the corresponding observation. If not mentioned otherwise, all regressions include the entire set of controls: locational, transaction-specific, and macro-financial. Fixed effects are on the zip-code-year level. The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

Table 1.B.5
Further robustness checks relying on additional data, base case

Neighborhood	(1) New construction	(2) Renovations	(3) Exclude school rezones
First-order neigh.	0.21***	0.21***	0.22***
	(16.7)	(16.6)	(15.1)
Second-order neigh.	0.04***	0.04***	0.05***
	(3.9)	(4.0)	(4.1)
Third-order neigh.	0.03**	0.03**	0.04**
	(2.9)	(2.9)	(3.1)
Fourth-order neigh.	0.04***	0.04***	0.03**
	(4.8)	(4.8)	(3.0)
Fifth-order neigh.	0.02*	0.03*	0.03*
	(2.2)	(2.3)	(2.4)
Sixth-order neigh.	0.002	0.002	-0.004
	(0.1)	(0.1)	(-0.3)
AIC	-121,762	-121,825	-86,992

This table documents further robustness checks on our key results. Column (1) includes six additional dummies indicating a construction related transaction within the respective neighborhood within the past T = 180 days. The constructions are identified using deed data provided by CoreLogic, which include a "new construction" indicator. Column (2) depicts results when including a dummy indicating whether a larger renovation within the building in which a given home is located took place during the holding period. The renovations are identified using the "effective year built" information given in our dataset, which indicates the first year when "the building was assessed with its current components". Column (3) shows results when omitting observations located in areas where school zones were changed. We focus on elementary school zones, since Manhattan allows for city-wide choice for high schools, and enrollment into middle school is at least based on districts, that have been established in the 60s and therefore should not have been subject to changes. To identify changes in school zones, we use data provided by the city of New York, provided via the NYC open data portal. The data ranges back to the school year 2009/2010, which further informs about the most recent prior school zone change. To be sure that no observation is affected by school rezoning, all observations in zones subject to changes before 2009 are dismissed. For all other years, observations are dismissed if, within the holding period of a property, the school zone changes. We further assume that school zone changes are known 6 months prior to the beginning of the school year. For this regression analysis, 26,795 observations remain. All regressions include the entire set of controls: locational, transaction-specific, and macro-financial. Fixed effects are on the zip-code-year level. The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

Table 1.B.6
Robustness, no trades in lower-order neighborhoods

Neighborhood	Base case returns	No neigh. in lower order
First-order neigh.	0.21***	-
	(16.9)	-
Second-order neigh.	0.01	0.06***
	(0.9)	(3.5)
Third-order neigh.	0.03**	0.02
	(2.6)	(0.5)
Fourth-order neigh.	0.04***	0.003
	(4.7)	(0.05)
Fifth-order neigh.	0.03*	-0.07
	(2.2)	(-1.1)
Sixth-order neigh.	-0.0004	0.09
	(-0.04)	(0.8)
AIC	-121756	

This table shows results of a single regression investigating comovements when no trade in the lower-order neighborhoods is available. The left column shows estimates for the neighboring excess returns as in the base case. The right column shows results when including the base case returns and conditioning on no observations available in lower-order neighborhoods. Thus, the regression asks if there is more information in neighboring trades when observations from lower-order neighborhoods are unavailable. All regressions include the entire set of controls: locational, transaction-specific, and macro-financial. Fixed effects are on the zip-code-year level. The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

Table 1.B.7
Robustness, SAR model

Neigh.	First-order	Second-order	Third-order	Fourth-order	Fifth-order	Sixth-order
$\hat{\rho}$	0.21*** (16.8)	0.05*** (4.1)	0.04*** (3.4)	0.05*** (5.3)	0.03** (2.7)	0.005 (0.3)
AIC	-121,685	-120,952	-120,940	-120,957	-120,935	-120,917

This table reports estimates of parameters ρ for six separate spatial autoregressive (SAR) models of the form $Y = \rho WY + X\beta + \epsilon$, where Y is an $n \times 1$ vector that contains chronologically ordered annualized excess returns, X is a $n \times k$ matrix of control variables, and ϵ is a vector of normally distributed nuisance terms. For each regression, we investigate a particular neighborhood and calculate the weighting matrix, W, as follows: In a first step, we assign a one to an entry, if the spatial criterion for the respective neighborhood is met and the trade was executed in the past $T = 180$ days. Otherwise, we assign a zero to the corresponding entry. W is obtained by standardizing all rows to sum up to one. Since Y is chronologically ordered, the determinant, $|I_n - \rho W|$, that has to be calculated when estimating the model, always equals one because W is a triangular matrix with ones on the main diagonal, making the model computationally feasible to being estimated with standard Maximum Likelihood methods. All regressions include the entire set of controls: locational, transaction-specific, and macro-financial. Fixed effects are on the zip-code-year level. The figures in parentheses report t-statistics derived from robust standard errors that are clustered over the zip-code-year level. ***, **, and * denote significance at the 0.1%, 1%, and 5% levels, respectively.

Chapter 2

How Do Assessed Values Affect Transaction Prices of Homes?*

*I am grateful for helpful comments and suggestions from Jan Brueckner, Marcel Fischer, Roland Füss, Patrick Hauf, Jens Jackwerth, Axel Kind, Marlene Koch, Julie Schnaitmann, Steffen Sebastian, Bertram Steininger, Michael Weber, Johannes Zaia, and conference participants at the 2021 AREUEA-ASSA and the IRES Doctoral Symposium 2020, as well as seminar participants at the University of Konstanz and the Royal Institute of Technology Stockholm. An early version of this manuscript existed under the name "The effect of public property valuation". I gratefully acknowledge financial support from the German Research Association (DFG), grant FI2141/5-1.

2.1. Introduction

For many governments worldwide, property taxes are a major source of income. In 2019, for instance, US homeowners paid more than USD 616 billion in property taxes, accounting for almost 40% of state and local tax revenue.[1] To take account of wealth disparities among homeowners, authorities tax each home by a fraction of its official assessed value (AV). AVs constitute an important factor for homeowners, as they are not only central for the calculation of individual tax bills, but also constitute updated market value estimates for the underlying homes. Despite their significant role for homeowners and the sizable redistribution of wealth, it remains largely unexplored how the information contained in AVs is incorporated in homes' transaction prices.

In this paper, I investigate the causal effect of AVs on transaction prices of homes. This effect should be driven by two counteracting channels. On the one hand, once a home's AV is publicly available, it serves as a potential reference price, e.g., during price negotiations. An increasing AV should therefore increase the trading price of the corresponding home (anchoring channel). On the other hand, an unexpected increase in the AV implies a higher future property tax burden, which should in turn negatively affect a home's trading price (tax channel). I identify the dominating channel by employing a novel Difference-in-Differences (DiD) regression setup that exploits AV publication dates and geographic variation in the frequencies in which homes are reassessed. My results document that the tax channel prevails: An unexpectedly higher AV causes a decline in the corresponding transaction price.

My empirical analysis is based on two representative datasets that contain historical tax records and individual transactions of homes, respectively. I focus on single-family homes in New York State for several reasons. First, within each municipality, AVs are published annually at a particular date, allowing me to compare pre- and post-publication periods. Second, reassessment frequencies vary on the municipal level, such that homes sold in tax years without reassessment serve as control group. Third, AVs are not forecasts, but have to reflect home values in the past. Fourth, unlike, e.g., in California, there is no post-transaction adjustment of AVs. Instead, AVs are collectively generated by a local assessor. Together, these characteristics provide the opportunity to employ a DiD setting that allows for causal inference.

The term "assessed value" requires clarification. While local assessors estimate the fair market value of each home, "assessed values" in the most technical sense can

[1] Source: US Census Bureau, Quarterly Summary of State and Local Government Tax Revenue. Last retrieved on March 16, 2021 from https://www.census.gov/data/tables/2019/econ/qtax/historical.Q4.html.

also describe the *share* of the estimated market value that is taxable. In this paper, I avoid this technical definition and use the term "AV" to refer to the full market value estimate provided by the local assessor, which is favorable for several reasons. First, using the full market value preserves cross-municipal comparability, as the taxable share, in contrast, varies over jurisdictions and time. Second, the external validity of this study is enhanced, as market values are not only the nationwide basis for property taxation, but are also similarly employed in many other countries around the world such as Canada or Japan. Third, while the taxable share of market values might be capped in particular areas such as New York City, the fair market value is not, such that its growth rate can inform about the future development of tax payments.[2]

To identify the effect of AVs on transaction prices, a measure other than raw sales prices is necessary. This is the case as homes can be heterogeneously affected through the same channel, even within the same treated municipality. For instance, regardless of which channel prevails, homes with unexpectedly low AVs face opposite price changes compared to homes with unexpectedly high AVs. Consequently, it is possible that AV-induced positive and negative price changes cancel each other, making the investigation of aggregate price changes uninformative for this study. I solve this problem by investigating the relative distance of sales price and AV. In a theory section, I formally show that this distance measure should change in opposing directions under each channel, respectively. In particular, the sales-price-AV distance should decrease under the anchoring channel, as sales prices move towards homes' AVs. In contrast, under the tax channel, the sales-price-AV distance should generally increase. Intuitively, an unexpectedly higher AV deceases the transaction price, which should increase the distance between both quantities.

My results document that AVs negatively affect transaction prices, in line with the tax channel. In particular, the absolute sales-price-AV ratio increases by one percentage point after updated AVs are published. A back-of-the-envelope calculation shows that this result translates to a change of the initial sales price of an average home of about 0.9%, with the effect's sign depending on whether the AV of a given home is unexpectedly high or low, respectively. The results are robust to the inclusion of different local and temporal fixed effects, respectively. Consistent with the tax channel, units associated with higher effective tax rates are stronger affected than units with lower rates. Right after AV publication, no immediate effect is observed, suggesting that when AVs are most salient and up-to-date, anchoring is outweighing the tax channel.

[2]My findings are robust to the exclusion of areas with assessment caps as shown in Section 2.6.

The identification strategy faces multiple challenges that have to be addressed. First, the common trend assumption, crucial for DiD regressions, must be fulfilled. As sale dates are known by day, it is possible to analyze pre-publication trends of treatment and control group, respectively. I find that the conditional means of each group follow similar trends, indicating that the common trend assumption plausibly holds. This conclusion is further supported by similar trends found in several subsamples. I additionally run placebo tests on pre-publication observations, which document no significant difference in pre-treatment trends and coefficient estimates close to zero.

Second, it is possible that buyers and sellers postpone transactions until updated AVs are published to reduce the uncertainty associated with the transaction. This should be of minor concern, as waiting in the housing market is costly, e.g., due to maintenance and opportunity costs, reducing the incentives to hold a property longer than necessary. Additionally, an investigation of transactions in a window around the publication date does not show a jump or a clear trend, and the number of transactions evolve similarly in treatment and control group.

Third, homeowners can challenge the assessment of their home, i.e., the values published at the considered dates are only tentative. To investigate how often AVs are changed after their initial publication, I make use of a dataset from the New York City government on notes sent to homeowners if their AV was changed post-publication. These homes account for only 0.9% of the single-family homes found in my New York City dataset for the following tax year, suggesting that the possibility to contest AVs is of only minor importance for this study.

This work contributes to a growing strand of literature on the effect of property taxation on trading prices. Bai et al. (2014) and Du and Zhang (2015) show that an introduction of a property tax can have a negative or no effect on price growth, exploiting a trial tax in China. Similarly, Elinder and Persson (2017) find only extremely high-valued homes to respond with price declines to an unexpected tax cut in Sweden. Further work, such as Wassmer (1993), Palmon and Smith (1998), Hilber (2017), and Livy (2018), documents that tax rate changes are negatively capitalized in sales prices. While the literature focuses on the tax rate and the introduction of property tax systems, the tax base (here, the AV) received much less attention. This is surprising, given that tax bases can change in different directions within the same district and thus potentially affect prices in opposite directions, even for neighboring homes. This is in sharp contrast to changes in the tax rate, after which prices of all treated homes should adjust in the same direction collectively. I thus contribute to the taxation literature by first, showing that the tax base itself affects trading prices and second, by uncovering an additional taxation effect that is heterogeneous even within the same treated location.

I further contribute to the literature on inequity in property taxation by uncovering price distortion as an additional source of inequitable outcomes. So far, other work such as Goolsby (1997), Allen and Dare (2002), Sirmans et al. (2008), and Hodge et al. (2017) focusing on inequity in tax payments. I extend this strand of literature by showing that misspecified assessments are even capable of distorting sales prices themselves. As my results indicate that AVs negatively impact sales prices, I uncover a double-punishment for homeowners with unjustified high AVs. Not only do they have to pay an excessive amount of taxes, they also suffer from a price discount when selling their home.

By investigating AVs as a potential anchor for buyers and sellers of homes, I contribute to the literature that underpins the importance of reference quantities in the housing market. So far, several anchoring phenomena have been investigated. Northcraft and Neale (1987) document that even professional real estate agents adjust their appraisals towards a randomized listing price. Genesove and Mayer (2001) show that homeowners consider the initial purchase price as reference point when they are selling their home. Andersen et al. (2021) estimate a structural model of listing decisions and identify the nominal purchase price as a reference point for homeowners. Fischer et al. (2021) document that realized returns of homes traded in close neighborhoods have an increased predictive power for future prices once they are publicly recorded. Similarly, Bailey et al. (2018) show that individuals rely on the house price growth experienced by distant friends when making their buy or rent decision.

In the context of anchoring on AVs, Jones (2020) shows that homeowners confronted with an increase in their AV have a higher propensity to contest their home's assessment, which can be linked to loss aversion. Considering the AVs as an anchor for valuations instead, the evidence provided by the current literature is mixed. Cypher and Hansz (2003) do not find anchoring on assessed values in an experimental setting. In contrast, Levy et al. (2016) find homeowners in New Zealand to be influenced by values that are used for property taxation, but are not necessarily market value estimates. While these studies investigate AVs primarily as anchor, this study expands this view by studying the interplay between tax and anchoring channel in a quasi-experimental framework.

The remainder of this paper is structured as follows. In the next Section 2.2, I motivate the channels through which AVs should influence trading prices within a theoretical framework. Section 2.3 describes the empirical approach followed in this paper. The data used and the validity of the methodology applied is discussed in Section 2.4. Results and robustness checks are presented in Sections 2.5 and 2.6, respectively. Section 2.7 concludes.

2.2. Theoretical considerations

In this section, I first briefly present crucial features of the New York State property tax practices to set the practical foundations for the model presented afterwards. This model first, theoretically motivates both anchoring and tax channel, respectively, and second, illustrates both channels' implications for the causal impact of AVs on transaction prices.

2.2.1. Some preliminaries

In each municipality in New York State, AVs are published annually by the local assessor at a particular date T. The frequencies in which homes are reassessed are determined by the municipal governments from annual reassessment to once in several decades. If homes are reassessed in a given period, the updated AVs become available to all market participants at T, e.g., on a public webpage. In contrast, if the municipality is not reassessing homes in the given period, the known AVs from the prior year are published.[3] The existence of pre- and post-publication periods as well as years without reassessment in some municipalities yield the key ingredients for the DiD analysis that is described in Section 2.3.

Apart from AVs, taxes paid by homeowners depend on multiple factors, such as the local tax rate, the local assessment ratio, individual exemptions, and the total municipal tax levy, such that the effective tax rate, τ, can vary substantially across individual homeowners, even within the same jurisdiction. These include the local tax rate, the local assessment ratio, local budgeting, as well as individual exemptions, such that effective tax rates can vary substantially even within the same jurisdiction.[4]

[3] In some municipality-tax-year combinations, AVs are adjusted collectively by the same factor, e.g., to adjust for inflation or general market movements. I exclude such observations in the analysis as they can be neither assigned to treatment nor control group.

[4] Additionally, there are two policies that are limiting the increase of tax payments. First, in New York City and Nassau county, the taxable share of the fair market values can maximally increase by 6% per year or 20% over five years, which might limit the effect of the tax channel. In the robustness Section 2.6, I show that my results are robust to the exclusion of these areas. Second, in most areas, such as counties and cities outside NYC, the annual increase of the districts' total tax levy is capped by the minimum of either 2% or the CPI inflation rate. Tax levy limits do not rule out significant individual changes in tax payments, however. For instance, a substantial AV increase relative to

An extended description of the New York State property tax system, based on the state's property tax calendar,[5] is provided in Appendix 2.A. The following section provides a theoretical model that illustrates the effect ov AVs on transaction prices of homes.

2.2.2. A simple model

I consider a model similar to the closed-form framework in Landvoigt et al. (2015), extended by a second period and property taxation, but with divisible housing. A representative household maximizes lifetime utility, V, over two periods by choosing between (numéraire) consumption c_t $(t = 1, 2)$ and units of (divisible) housing stock $n > 0$.

In the first period, the household buys a home at a price that is determined by $p(n) = \dot{p}n$, in which $\dot{p} > 0$ is the price per unit of a home. In the second period, the household pays property taxes as a fraction of the home's known AV, $\tau AV(n)$, in which $AV(n) = \hat{p}n$ with $\hat{p} > 0$ is again linearly increasing in n to reflect that larger homes tend to be assessed at a higher value. The per-unit transaction price, $\dot{p}$, is likely to diverge from the per-unit AV, $\hat{p}$, as the local assessor is uncertain about the true model and can thus only provide a best estimate for $\dot{p}$. The parameter τ is the household's effective tax rate.

For simplicity, further assume that the household is a home buyer and holds initial wealth, $W > 0$, that is used to pay for the home, property taxes, as well as consumption in both periods. Assuming a log-additive utility function and that wealth can be frictionlessly transferred to the second period, the household solves the optimization problem

$$\max_{c_1,c_2,n} \quad V(c_1, c_2, n) = \ln(c_1) + \beta \ln(c_2) + (1 + \beta)\theta \ln(n) \tag{2.2.1}$$

$$s.t. \quad W = c_1 + c_2 + n(\dot{p} + \tau\hat{p}), \tag{2.2.2}$$

in which $0 < \beta < 1$ is a time preference parameter and $\theta > 0$ determines the importance of housing relative to regular consumption. Optimizing for c_1, c_2, and n, the first order conditions imply

other properties in the municipality would still lead to a stark increase in tax payments, even if the total tax levy remained constant.

[5]The official tax calendar can be found at the New York State Department of Taxation and Finance website: https://www.tax.ny.gov/pit/property/learn/proptaxcal.htm last retrieved on March 16, 2021.

$$\frac{1}{c_1} = \beta \frac{1}{c_2} \qquad \text{and} \tag{2.2.3}$$

$$\frac{(1+\beta)\theta c_1}{n} = \dot{p} + \tau \hat{p}. \tag{2.2.4}$$

Equation (2.2.4) shows that, in addition to the closed-form model of Landvoigt et al. (2015), the marginal rate of substitution between lifetime housing utility and consumption depends not only on the marginal house price at size n, $\dot{p}$, but also on the marginal tax payment at n, $\tau\hat{p}$. Increasing marginal tax payments through either higher τ or $\hat{p}$ therefore imply a higher willingness of the household to substitute housing with consumption. Solving the first-order conditions and assuming market clearing at fixed housing supply $\bar{n}$, I find the equilibrium per-unit transaction price given by

$$\dot{p}^o = \frac{W\theta}{(1+\theta)\bar{n}} - \tau\hat{p}. \tag{2.2.5}$$

It follows from Equation (2.2.5) that the market clearing price of one unit of a home is decreasing in the per-unit AV, $\hat{p}$, illustrating the tax channel: A ceteris paribus higher AV decreases the sales price of a home through an increased tax burden.

The rational choice model described so far does not include anchoring. This is the case as anchoring itself is not rational. For instance, Northcraft and Neale (1987) document that even professional real estate agents adjust their appraisals towards randomly assigned ask prices. Similarly, Black and Diaz III (1996) document random adjustments of ask prices to influence offering prices as well as final transaction prices in an experimental setting.[6] A rational agent would simply adjust the optimal choice with respect to the purchase price to account for the heuristic bias. The anchoring-adjusted price function of a rational agent would then simply coincide with the optimal decision, $\dot{p}^o$. In consequence, anchoring must result in a deviation from the optimal choice, unless the AV is a perfect prediction of the sales price before anchoring.

With anchoring, the final price per unit (and thus the overall price paid for a home) can be written as a linear combination of the per-unit price from Equation (2.2.5) and the per-unit AV (see, e.g., Gibbs and Kulish, 2017), given as

$$\dot{p}^{anch} = (1-\alpha)\left(\frac{W\theta}{(1+\theta)\bar{n}} - \tau\hat{p}\right) + \alpha\hat{p}, \tag{2.2.6}$$

[6] In the pioneering work of Tversky and Kahneman (1974), subjects are influenced in their judgment by a (seemingly) random wheel of fortune, illustrating the irrationality of the anchoring heuristic.

in which $0 \leq \alpha \leq 1$ is the equilibrium degree of anchoring. If $\alpha = 0$, $\dot{p}^{anch}$ corresponds to the rational choice, $\dot{p}^{o}$, and if $\alpha = 1$, the transaction price corresponds to the AV.

The effect of each channel on the transaction price can be now illustrated with the first derivative with respect to the per-unit AV, $\hat{p}$,

$$\frac{\partial \dot{p}^{anch}}{\partial \hat{p}} = -(1-\alpha)\tau + \alpha, \tag{2.2.7}$$

which indicates that the tax channel dominates, i.e., an increasing AV decreases the transaction price, if $\alpha - (1-\alpha)\tau < 0$, and the anchoring channel dominates, i.e., an increasing AV leads to an increasing sales price, if $\alpha - (1-\alpha)\tau > 0$.

The order of magnitude of which a channel is dominating should be further influenced by (i) the drivers of the degree of anchoring, i.e., the determinants of α and (ii) the individual effective tax rate τ. Accordingly, it should hold that first, the tax channel is most pronounced for units associated with high effective tax rates. Second, anchoring should be most dominant when AVs are most salient and up-to-date, i.e., right after publication. Both of these hypotheses are tested in Section 2.5.2.

2.2.3. Measuring the impact of AVs on transaction prices

So far, the model illustrated that whether a sales price is positively or negatively influenced by the AV depends on two factors, first, whether tax or anchoring channel dominates and, second, whether the AV is relatively high or low. Empirically, the latter factor results in an issue that needs to be addressed, since individual homes within a treated municipality are affected heterogeneously. Intuitively, while over- and undervalued homes should be affected in opposite directions regardless which channel dominates, the aggregate effect on prices at the treatment level might well be zero. In consequence, DiD regressions investigating changes in nominal prices are uninformative for the causal effect of AVs on transaction prices.[7]

I solve this issue by investigating the absolute distance between AV and sales price P, given as

$$DA(n) = \left| P(n) - AV(n) \right| = n \left| \dot{p}^{anch} - \hat{p} \right|. \tag{2.2.8}$$

[7] Another approach would be to investigate subsamples of over- and undervalued homes separately. Following this approach, however, requires knowledge of additional factors that are difficult to identify, such as which homes are actually under- or overvalued and the expectations of market participants. Additionally, due to the associated selection process, the construction of a suitable control group would be rather difficult.

Importantly, the true AV (or the true per-unit price, $\hat{p}$) that is subtracted is known after publication, but not before. Thus, when later constructing the dataset, the AV that is matched to the sales price is always the one published at the tentative roll date closest to the transaction date.

In the empirical analysis, I make use of a standardized version of DA, as further outlined in Section 2.3. For simplicity, I illustrate the effects from anchoring and tax channel on the absolute measure DA first. It is then straightforward to show that the same advantages of DA also hold for the relative measure.

In the remainder of this section, I show that DA should change in opposing directions for anchoring and tax channel, respectively. That is, relative to the control group, DA should increase for all units in treated municipalities if the tax channel dominates, and decrease for all units in treated municipalities if the anchoring channel dominates. For the sake of clarity, the following two sections separately investigate each channel's effect on DA.

2.2.3.1. Anchoring channel

In this section, I show that the anchoring channel reduces the absolute distance between sales price and AV, DA. Intuitively, this is the case as anchoring moves the transaction price towards the AV, reducing the distance between both quantities, DA. This intuition is formally outlined below.

For the sake of simplification, define $\dot{p}^* := \frac{W\theta}{(1+\theta)\bar{n}}$. Then, the change in the absolute sales-price-AV difference from pre- to post AV publication, given that only the anchoring channel is present, can be described with

$$DA^{pre}(n) - DA^{post}(n) = n\left|\dot{p}^* - \hat{p}\right| - n\left|(1-\alpha)\dot{p}^* + \alpha\hat{p} - \hat{p}\right| = \alpha n|\dot{p} - \hat{p}|. \quad (2.2.9)$$

Note that anchoring is not possible pre-publication since $\hat{p}$ is not yet known. From Equation (2.2.9), it holds that $DA^{pre} \geq DA^{post}$, as $n > 0$ as well as $\alpha \geq 0$. Notably, anchoring strictly reduces DA for all homes, whether over- or undervalued, given that $\hat{p}$ does not perfectly predict $\dot{p}^*$.

2.2.3.2. Tax channel

The tax channel, in contrast, generally implies an opposite, increasing effect on DA. Intuitively, an unexpectedly higher AV leads to a lower sales price, such that the distance between both quantities, DA increases. In the opposite case, an unexpectedly lower AV increases the transaction price, again increasing DA. The formal conditions under which this intuition holds are derived in the following.

As future tax payments should play a role for homeowners in any case, I assume that they form expectations about their future tax burden, $n\tau E[\hat{p}]$. Hence, before updated AVs are published, households purchase their home under the expected future AV, and replace expectations with the true value post-publication. Thus, when considering the tax channel isolated from anchoring, it holds that

$$DA^{pre}(n) - DA^{post}(n) = n\left|\dot{p}^* - \tau E[\hat{p}] - \hat{p}\right| - n\left|\dot{p}^* - \tau\hat{p} - \hat{p}\right|. \tag{2.2.10}$$

In Appendix 2.B.1, I show that $DA^{pre}(n) \leq DA^{post}(n)$, i.e., that the tax channel has an increasing effect on DA relative to the pre-publication value, if one of the sufficient conditions

$$\hat{p} \leq E[\hat{p}] \leq \dot{p}^{pre} + \frac{|\epsilon|}{2} \tag{2.2.11}$$

$$\hat{p} \geq E[\hat{p}] \geq \dot{p}^{pre} - \frac{|\epsilon|}{2} \tag{2.2.12}$$

holds, in which ϵ is the valuation error, i.e., the difference between $\dot{p}^{pre} = \dot{p}^* - \tau E[\hat{p}]$ and $\dot{p}^{post} = \dot{p}^* - \tau\hat{p}$. This shows that when the household's expectations about the AV, $E[\hat{p}]$, lie between the actual AV and the pre-publication price, the tax channel has an increasing effect on DA, regardless whether a home is over- or undervalued. Furthermore, as DA as well as the change in DA for treated and control group, respectively, are observable, it is not necessary to know the expectations of households.

2.3. Methodology

The goal of this paper is to analyze the effects of value-based property taxation on transaction prices. To be able to causally interpret the results, I run DiD regressions that compare pre- and post-publication transactions of homes in municipalities publishing updated and previously unknown AVs, with homes in municipalities that did not reassess homes, thus publishing the already known AVs from the previous year.

As my dataset contains multiple years of data for a large amount of jurisdictions, I define municipality-tax-year clusters c by employing symmetric time-windows of ± 180 days around each publication date T_c. [8] I do not consider longer time-spans to avoid overlaps between treated and control clusters.

[8] The results are robust for alternative time-windows of 90, 120, and 150 days, respectively, as shown in Table 2.5.4.

In the previous section, I proposed a simple, absolute measure, DA, that is changing homogeneously for under- and overvalued homes and that is moving in opposing directions for tax and anchoring channel, respectively. In the empirical application, I use a standardized version of DA, to prevent that higher-priced homes drive the regression results. Leaving the theoretical framework from Section 2.2, I calculate for each home i, transacted within municipality-tax-year cluster c, the absolute ratio between P_{ic} and AV_{ic} as

$$DAR_{ic} = \left| \frac{DA_{ic}}{AV_{ic}} \right| = \frac{|P_{ic} - AV_{ic}|}{AV_{ic}}. \tag{2.3.1}$$

The conditions derived for DA_{ic} in theory Section 2.2 hold for DAR_{ic} as well, as AV_{ic} is strictly larger than zero and is fixed for each home i within a given tax year. That is, post-publication, the tax channel should lead to an increase in DAR_{ic} for treated municipalities, while the anchoring channel should lead to a decrease in DAR_{ic} for treated units. Intuitively, after AVs are published, the anchoring channel moves prices towards these values, reducing the sales-price-AV distance, whereas the tax channel drives prices away from AVs, increasing the distance between both quantities. Note that the AV that is matched to the sales price is always the one published at T_c. Thus, the empirical analysis compares the relationship of sales prices to assessed values, which stem from the same assessment year within a cluster c. Under the condition that the control group follows the common trend assumption, the treatment effect on transaction prices thus stems from the knowledge about the true AV.

It is necessary to standardize by AV_{ic} rather than P_{ic} in Equation (2.3.1), as AV_{ic} remains constant before and after T by definition, whereas the transaction price should be, as illustrated above, affected by anchoring and tax considerations. Consequently, dividing DA_{ic} by the transaction price, P_{ic}, instead of AV_{ic} would yield an unstable and endogenous measure. It is important to note that investigating the absolute value of a relative measure makes the implicit assumption that an overvaluation of 50% can be treated equally to an undervaluation of the same amount. While this assumption should be reasonable for most ratios, it becomes less plausible for larger deviations, e.g., an overvaluation of 99% should be generally less extreme than an undervaluation of 99% due to the natural lower bound of -100%. Therefore, in the robustness section, I first show that my results quantitatively hold when investigating the non-standardized measure DA_{ic}, and second, that they are robust to setting a conservative upper bound for DAR_{ic}.

I define a dummy variable $Treat_{ic}$ that equals one, if home i is sold within a municipality-tax-year cluster c in which homes have been reassessed collectively, i.e., a new AV is available for all respective homes, and zero otherwise. I continue on following the standard DiD framework by defining a dummy variable $Post_{ic}$ that equals one if home i was sold after T_c, and zero otherwise. Consequently, I run regressions of the form

$$DAR_{itc} = \beta_0 + \beta_1 Treat_{ic} + \beta_2 Post_{itc} + \gamma Treat_{ic} \times Post_{itc} + \delta_c + \nu_t + \epsilon_{itc}, \qquad (2.3.2)$$

in which DAR_{itc}, as defined in Equation (2.3.1), is the absolute ratio between P_i and AV_i, β_0 is an intercept, δ_c and ν_t denote location-tax-year (e.g, zip-code-tax-year) and temporal (e.g., year-quarter) fixed effects, respectively, $Treat_{ic} \times Post_{itc}$ is the interaction between the dummies identifying treatment and post groups, respectively, and ϵ_{itc} is a nuisance term. The coefficient of interest γ measures the treatment effect on DAR_{itc} from the updated AVs relative to the control group. As illustrated in Section 2.2, if γ is positive, the tax channel is dominating. In contrast, the anchoring channel is dominating if γ is negative. Before presenting the empirical results, I introduce the data used and discuss the identification in the following section.

2.4. Data and identification

The first part of this section briefly describes the data cleaning process, how the datasets are merged, and how treated and non-treated units are identified. After illustrating the cleaned dataset, I discuss the validity of my identification strategy.

2.4.1. Data cleaning

I merge property transactions of single-family homes with historical tax records from 2007 to 2017 that contain assessments for up to eleven years per property. Both datasets are obtained from the data vendor *CoreLogic*, who provides a US property record coverage of more than 99%.[9] I focus on New York State properties for several reasons. First, the municipality-dependent publication dates T_c are easily accessible through the municipal profile webpage provided by the state government.[10] Second, the revaluation frequencies across municipalities differ, allowing for the construction of a control group. Third, AVs are published annually, providing a distinct point at which the new information is available. Fourth, the state includes the largest US city as well as more rural areas, strengthening the external validity of the results.

[9]As reported by the data provider at https://www.corelogic.com/solutions/university-data-portal.aspx, last retrieved on March 16, 2021.

[10]To be found at http://orps1.orpts.ny.gov/cfapps/MuniPro/, last retrieved on March 16, 2021.

I start out by cleaning the transaction dataset by following DeFusco et al. (2020) in dismissing all observations that are not classified as "arms length", have a missing sales price, are associated with a foreclosure, or ones identified as duplicate.[11] Additionally, I remove all transactions without an assessor's parcel number (APN), which is used to merge the transactions with the tax records. Afterwards, I follow Bollerslev et al. (2016) and set sharp nominal bounds for the transaction prices. Based on municipality and sale date, I identify for each of the remaining homes the corresponding municipality-tax-year cluster c. Based on this classification, I then match the corresponding AV_{ic} to the respective home i.

In addition to the contemporaneous AV that is used to construct DAR_{ic}, I match the one-year AV lag for two reasons. First, to filter out observations with unusually large valuation changes, e.g., induced by substantial renovations. Second, to employ a data-driven way to identify treatment and control clusters c. I do this by investigating whether at least 75% of the remaining observations have the same one-year AV return. To rule out that small deviations confound the classification, I round the returns to the third digit before doing so. If at least 75% of observations in a municipality-tax-year cluster c have a one-year AV return of zero, I assign all observations in c control unit status and dismiss all observations in this c with a return different from zero. If at least 75% of returns within a particular c have the same return, but the return is different from zero, e.g., because the local assessor market-adjusted AVs in the given year, I dismiss all observations in this cluster c, as they can neither be considered treated nor control units. The remaining observations are then assigned to the treatment group.

After filtering out extreme values of DAR_{ic} and DA_{ic} based on sample quantiles, respectively, I conduct a final step by dismissing clusters c with only few observations (less than 100), as treated units appear to be more often located in larger municipalities, such that comparability of treatment and control group is increased. The final dataset consists of 193,494 observations.

Table 2.4.1 provides summary statistics of key variables used in the later analyses. Panel A reports statistics for the treated units, Panel B for units sold within control clusters. From the total 193,494 observations, about 59% are treated units. For them, the average sales price is about USD 383,000, while the mean AV is slightly lower with about USD 381,000. In contrast, the control group contains units that tend to be comparably lower priced with an average sales price of USD 255,000. For both groups, a larger share of transactions are realized after the publication date, with 55% for the treated and 59% for the control units. A potential reason is that most municipalities

[11] A more detailed description of the data cleaning procedure is provided in Appendix 2.C.

publish AVs at the beginning of the year and in spring, and turnover is typically highest in summer (e.g., Ngai and Tenreyro, 2014). For the treatment group, the average increase of the AV is 1%, reflecting the general increase of house prices over the sample period, but also that the sample includes the post-bust period from 2007 to 2010. With a standard deviation of about 9%, substantial changes in the AV from one year to another appear to be rather common. The effective tax rate is on average higher for the control group. One reason for this observation could be the necessity to increase tax rates in lower-priced areas, to generate sufficient governmental income.

Figure 2.4.1 illustrates the geographic distribution of transactions within the state of New York, separated by treatment status. Panel A shows transactions for treated units, and Panel B for control ones. Unsurprisingly, properties in New York City, found in the southeast of the state, are all treated units, as here, homes are reassessed annually. Other cities, such as Buffalo, located in the western end of the state, reassess less frequently, which is why homes from Buffalo can be found in both groups. Note that as at least hundred observations are required per municipality-tax-year cluster to increase comparability between treated and control sample, there are fewer observations found in rural areas. The comparison of Panels A and B illustrates why the control group should match the treated units well: there are many overlaps among cities and towns, such that the AVs of both groups should be similarly precise and distributed, as the local assessors should remain more or less constant over time.

2.4.2. Identification

The key underlying assumption of the DiD analysis is that the trend of the dependent variable DAR_{ic} would have been the same for treated and control group in the absence of treatment, i.e., the one of a common trend. If this assumption holds, the control units can be used to infer about the counterfactual outcome of the treated units. In other words, the homes in the control group inform about how the dependent variable DAR_{ic} of the treated homes would have developed if AVs had not been updated.

To check the validity of this assumption, I plot pre-publication trends of the dependent variable DAR_{ic} for both groups in Panels A and B of Figure 2.4.2. Each panel shows conditional means of DAR_{ic} for two periods prior to publication, one for 180 to 91 days, and one for 90 to 1 day before publication of AVs. Panel A shows trends conditional on the base case controls zip-code-tax-year and year-quarter fixed effect dummies for both groups, respectively. DAR_{ic} is upward sloping for both groups at a highly similar magnitude. Likewise, Panel B shows the trends of DAR_{ic} for both groups when conditioning on separate sets of municipality and tax-year level fixed

Table 2.4.1
Summary statistics

Panel A: Treated units						
	Obs.	Mean	Std.	1st Q.	Median	3rd Q.
Sales price	113,701	382,521	307,676	184,000	342,000	485,000
Assessed value (AV)	113,701	380,698	295,568	191,863	346,000	485,000
Abs. price-AV ratio (DAR)	113,701	0.174	0.158	0.059	0.129	0.240
Abs. price-AV diff. (DA)	113,701	61,688	71,564	15,000	37,500	81,209
One-year return AV	113,701	0.010	0.087	-0.038	0	0.046
Post	113,701	0.548	0.498	0	1	1
Crisis (2007-2010)	113,701	0.226	0.418	0	0	0
Effective tax rate seller	93,903	0.017	0.011	0.008	0.014	0.025
Effective tax rate buyer	88,671	0.019	0.011	0.010	0.016	0.027
Panel B: Control units						
	Obs.	Mean	Std.	1st Q.	Median	3rd Q.
Sales price	79,793	255,205	206,511	121,000	203,000	333,900
Assessed value (AV)	79,793	247,902	193,296	118,600	207,000	325,000
Abs. price-AV ratio (DAR)	79,793	0.174	0.162	0.059	0.125	0.238
Abs. price-AV diff. (DA)	79,793	40,613	48,977	10,000	23,000	51,545
One-year return AV	79,793	0	0	0	0	0
Post	79,793	0.594	0.491	0	1	1
Crisis (2007-2010)	79,793	0.231	0.421	0	0	0
Effective tax rate seller	72,236	0.026	0.010	0.021	0.027	0.032
Effective tax rate buyer	58,253	0.028	0.010	0.023	0.029	0.034

This table provides summary statistics for the variables used in the empirical application. Panel A shows summary statistics for the homes located in treated municipality-tax-year-clusters c (i.e., municipalities in which homes have been reassessed in a given year surrounding the respective AV publication date). Panel B shows the same statistics for the control units, i.e., sales of homes within municipality-tax-year clusters that did not reassess homes. "Sales price" denotes the nominal transaction price, and the "Assessed value (AV)" is the publicly available estimated market value provided by the local assessor, which is known if the corresponding unit was sold before the publication date T_c, and known afterwards. "Abs. price-AV ratio" is the absolute value of the ratio between nominal sales price and AV (DAR). "Abs. price-AV diff." is the absolute value of the difference between nominal sales price and AV (DA). "One-year return AV" is the relative change in AV with respect to the prior year. "Post" is a dummy indicating whether a property was sold before or after the publication of AVs. "Crisis" is a dummy that equals one if the property value was published between 2007 and 2010, and zero otherwise. The "Effective tax rates" are defined as tax amount paid within a particular year, divided by a given AV. For the effective tax rate of the seller (buyer), the tax amount in the year prior to (after) the respective tax year is used. The AV used to derive the seller's (buyer's) tax rate is the one from the prior (contemporaneous) tax year.

Figure 2.4.1
Geographic dispersion of transactions

Panel A: Treated units

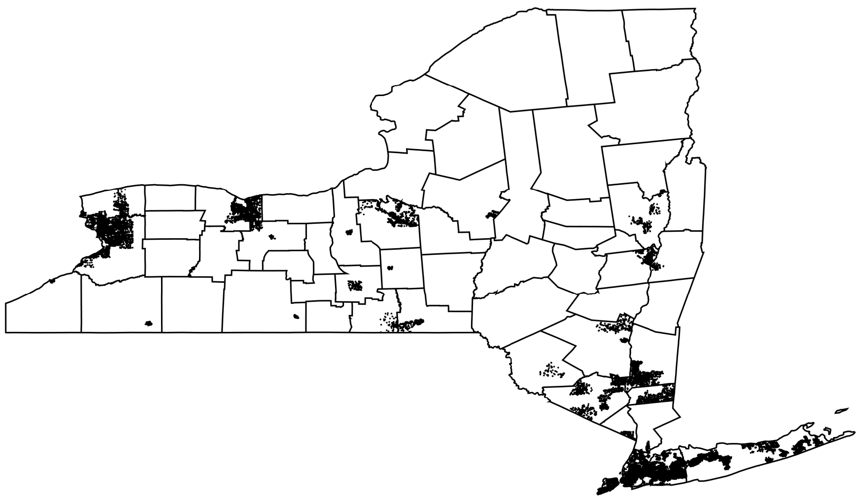

Panel B: Control units

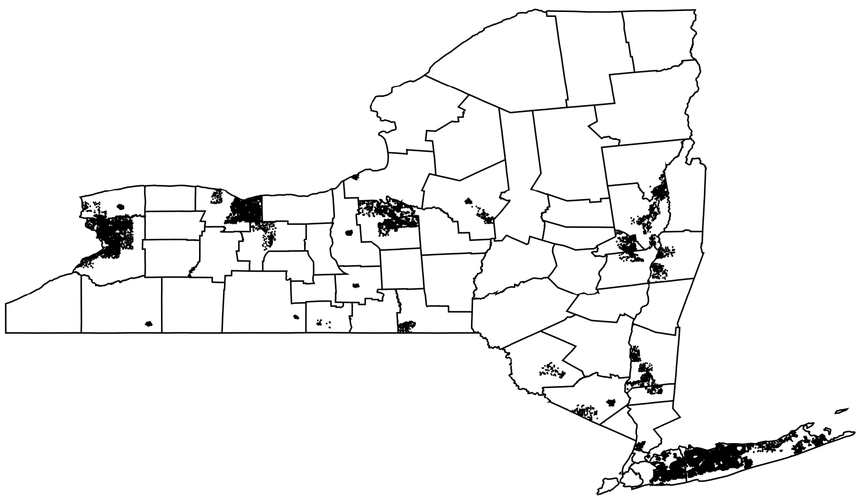

Panels A and B show the geographic dispersion of observations in the cleaned dataset within New York State. Note that observations from smaller areas are excluded to increase comparability between treatment and control group, as smaller municipalities tend to reassess less often than larger municipalities. Panel A shows the distribution of observations assigned to the treatment group, i.e., sales that took place in a time window of ± 180 days around the publication of updated AVs. Panel B shows the geographic distribution of homes in the control group, i.e., ones that were transacted ± 180 days around AV publication in tax-years in which the AVs were not updated (i.e., coincide with last year's AV). The solid lines indicate county and state borders. County and state border data is provided by the US Census Bureau.

Figure 2.4.2
Trends of DAR_{ic} before treatment

Panel A: Base case controls *Panel B: Simple controls*

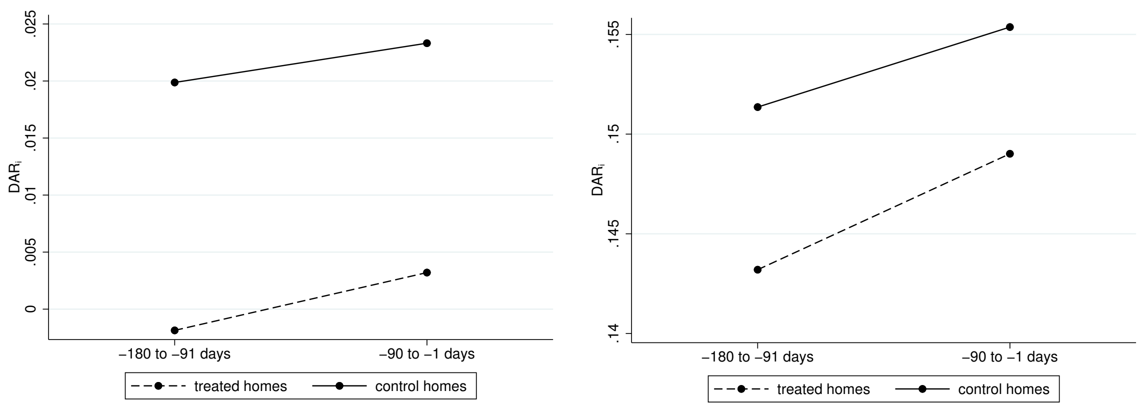

This figure shows the pre-publication development of the average absolute sales-price-AV ratio, DAR, defined in Equation (2.3.1), for both treated (dashed line) and control units (solid line). The points displayed indicate means for two subperiods of 90 days each. Panel A depicts trends conditional on zip-code-tax-year and year-quarter fixed effect dummies, respectively. The means in Panel B are conditional on separate sets of municipality as well as tax-year fixed effect dummies, respectively.

effect dummies, respectively. Even under the much simpler controls, both means are upward sloping at a similar magnitude, indicating that the common trend assumption holds reasonably well, and homes in control municipalities can thus be used to infer about the counterfactual outcome for the treated homes. Yet, the comparison of Panels A and B indicates that more fine-grained controls are helpful in correcting for unobserved heterogeneity over locations and time. Note that the absolute level of each point is uninformative as the conditional means displayed are estimated relative to a base category. This does not constitute a problem, as only the trend needs to be common in DiD regressions. Additional pre-publication trends for subsamples are presented in Section 2.5.2, e.g., for samples based on different effective tax rates. Several (placebo) tests on differences in pre-treatment trends are presented in Section 2.6.

A further potential issue is that homeowners wait for publication until they sell their homes. Figure 2.4.3 shows the frequency of transactions relative to AV publication for both treated (Panel A) and non-treated units (Panel B). For the treatment group, no change in the turnover trend is visible around day zero, indicating no immediate effect of AV publication on liquidity. With increasing amount of days after the publication date, however, there is an increase in observation frequency. Similarly, an increase in observations is observed for the control group, as shown in Panel B. For the control units, the increase is even higher, suggesting that AVs do not play a role

Figure 2.4.3
Amount of sales relative to the event date

Panel A: Treated units *Panel B: Control units*

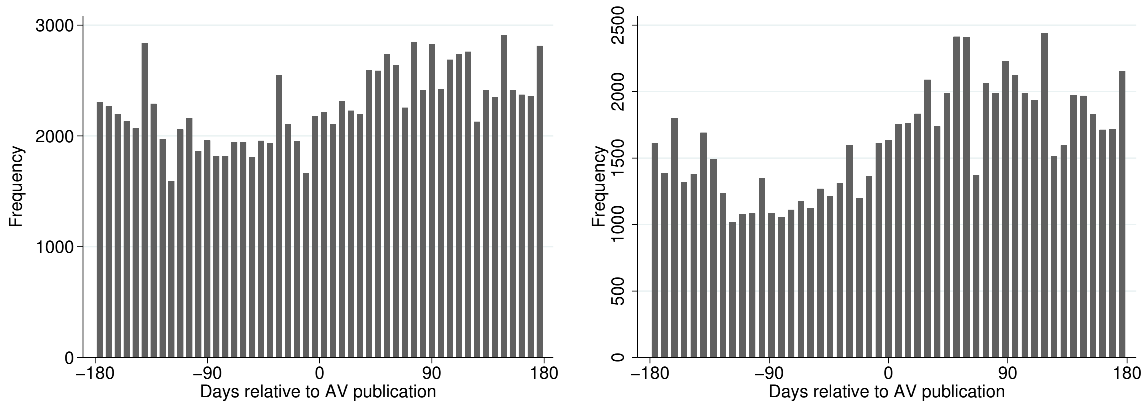

The histograms displayed in Panels A and B show the number of observations in the dataset relative to AV publication in a time window of ±180 days. Panel A shows the frequency for treated units, i.e., units for which the AV is revalued and is known by market participants only after day zero (T_c). Panel B shows the same variable for control units, i.e., ones for which the (old and new) AV is not revalued and thus known before and after day zero.

when considering the timing of a sale. A potential reason for the observed increase over time for both groups is likely to be that AVs are mostly published in spring, and turnover is typically highest in summer. Additionally, waiting for AV updates should not be a concern in the housing market, as it is costly due to factors such as interest payments and opportunity costs.

Another concern that needs to be discussed is that homeowners are able to challenge their assessment once they received notice of their updated AV. To investigate how frequently AVs are changed, I utilize data from the New York City government, which provides such information for the year of 2016.[12] Investigating the subsample for single-family homes, I find that only about 6,000 AVs have been updated after the tentative roll date. This accounts for only about 0.9% of the about 697,000 homes that are listed in the following year's tax records in the *CoreLogic* dataset, suggesting that the possibility to contest public estimates is affecting only a small share of observations and should therefore be of minor relevance.

Given that successful contests of tentative AVs are rather rare, the estimates published at T_c can be viewed as quasi-fixed, thus mitigating further endogeneity concerns. Taking further into account that the estimates, derived by using sales prices of comparable properties, are not meant to be forecasts but have to reflect home

[12] Available at https://data.cityofnewyork.us/City-Government/Revised-Notice-of-Property-Value-RNOPV-/8vgb-zm6e, last retrieved on August 17, 2020.

values at a particular date in the past (the "valuation date") should further support the causal interpretation of the results.

Finally, spillovers across municipalities should not be a concern, as homes are individually affected, depending on whether the updated AV is relatively high or low. That is, within a municipality, some homes can be positively, and some negatively affected. As this implies no shift in the aggregate price level into a particular direction within a municipality-tax-year cluster, no spillover effects across units due to AV publication should be expected. Having discussed the validity of the identification strategy, the following section presents the empirical results.

2.5. Results

The aim of this work is to show the impact of AVs on sales prices. To identify this effect, I investigate DAR_{ic}, the absolute ratio between the sales price of a home i and the corresponding AV, and exploit the timing of publication at time T_c in a time window of ±180 days. The sign of this effect is ex-ante not clear. As illustrated in Section 2.2, the tax channel should lead to an increase in DAR_{ic} through an induced change in tax payments. In contrast, anchoring should reduce DAR_{ic}. Using municipality-tax-year clusters in which homes are not reassessed as a control group, I run DiD regressions following Equation (2.3.2).

2.5.1. Base case

Table 2.5.1 shows results for OLS regressions with DAR_{ic}, the absolute sales-price-AV ratio, as dependent variable. The causal effect of AVs on transaction prices that is measured by DAR_{ic} is given by the interaction between dummy variables "Treatment" and "Post". Standard errors are clustered over counties and are shown in parentheses below.[13] Column (1) shows estimates when only municipality fixed effects are included. The effect of the interaction between treatment and post-publication dummy is estimated at 1.1%. The coefficient is statistically significant at the 0.1% level and indicates that knowledge about new AVs is influencing DAR_{ic} positively, in line with the tax channel. That is, an increased assessed value decreases the transaction price of the respective home.

Column (2) includes municipality-tax-year fixed effects, thus additionally accounting for temporal variation within local districts. The result is significant at the

[13] An inspection showed that the base case results are robust to alternative clustering of standard errors, such as on municipality or zip-code level, as well as Cameron et al. (2011) two-way clustering with county and tax year.

0.1% level and suggests a positive causal effect of AVs on DAR_{ic} of 1.1 percentage points. Using the more fine-grained zip-code-tax-year fixed effects and additionally year-quarter fixed effects as shown in columns (3) and (4), respectively, slightly reduces the base case estimate to one percentage point. The common pre-publication trend when using the base case specification in column (4) is shown in Panel A of Figure 2.4.2. Together, the positive estimates indicate that the tax channel is dominating and increasing the absolute sales-price-AV ratio by about one percentage point.

To provide an economic interpretation of these results, I do a back-of-the-envelope calculation making use of the relationship

$$\Delta P_{ic} = \frac{\gamma AV_{ic}}{P^c_{ic}}, \tag{2.5.1}$$

in which P^c_{ic} is the price under control conditions, and ΔP_{ic} the relative change in the sales price between treatment and control state. The coefficient γ is the treatment effect, which is estimated to be 1% in the base case. The simplifying assumption that allows this closed-form solution is that both transaction prices, under treated and control conditions, are larger than the AV, as shown in Appendix 2.B.2.[14]

Applying Equation (2.5.1) to the sample means of the treatment group, it is possible to derive an effect on transaction prices. Plugging-in the average AV of the treatment group (about USD 383,000), and using the average difference between sales price and the AV of USD 62,000 (which must be added to the sales price such that the underlying assumption is fulfilled), the change in P due to treatment is 0.9%. This indicates that value-based property taxation leads to economically significant price distortions of about one percent of the non-treatment sales price.

2.5.2. Evidence for the underlying channels

The results presented in the prior section indicate that the tax channel is dominating. The purpose of this section is to first, present supporting evidence that this is indeed the case, and second, to investigate the interplay between tax and anchoring channel in more detail.

If the tax channel is indeed in play, transactions that involve units associated with higher effective tax rates should be affected stronger, as illustrated in Equation (2.2.7). To test this prediction, I analyze subsamples selected on the basis of the effective tax

[14]This assumption is in line with the sufficient condition (2.2.12) from Section 2.2, which implies an increase in the sales price post-publication as well. Doing the same calculations, but assuming instead that the transaction prices are smaller than the AV, a decline in the transaction price is implied. Again, this decline is in line with the corresponding condition (2.2.11).

Table 2.5.1
Base case results

	(1)	(2)	(3)	(4)
Treatment	-0.009**			-0.024***
	(0.003)			(0.006)
Post	0.005**	0.005***	-0.003	-0.005
	(0.001)	(0.001)	(0.002)	(0.003)
Treatment × Post	0.011***	0.011***	0.010***	0.010***
	(0.002)	(0.002)	(0.002)	(0.002)
Municipality fixed effects	X	-	-	-
Municipality tax-year fixed effects	-	X	-	-
Zip-code-tax-year fixed effects	-	-	X	X
Year-quarter fixed effects	-	-	-	X
Adj. R-sq.	0.054	0.070	0.070	0.097
Observations	193,494	193,494	193,494	193,494

This table provides coefficient estimates on four separate Difference-in-Differences regressions based on Equation (2.3.2), with DAR_{ic}, the absolute sales-price-AV ratio, as dependent variable. "Treatment" is a dummy that equals one if the unit was sold within a municipality-tax-year cluster in which AVs were updated, and zero otherwise. "Post" is a dummy that equals one if the unit was sold after the AV was known, and zero otherwise. "×" denotes an interaction between two variables. "X" indicates that the set of control variables is used, whereas "-" indicates that the set remained unused for the particular model. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

rates of buyers and sellers, which can differ for several reasons. First, property tax rates change over time. Second, different households can have different exemptions on their payments (e.g., veterans or senior citizens). Third, the annual increase in payments might be capped.

I define the effective tax rate as tax payments made during a particular year, divided by a given AV. The ratio therefore indicates how many cents the homeowner has to pay in taxes per dollar AV. As it remains unclear in the data whether buyer or seller paid the tax amount in the year of the underlying transaction, I make use of the tax payments made for a given home in the years before or after the respective AV was published, respectively. Accordingly, the seller's effective tax rate is defined as the prior year's tax payments divided by the one-year AV lag. For the buyer, I use the forward lag of tax payments, but divide by the current AV to avoid the additional loss of a larger amount of data. As this is done for all observations in the sample collectively, and as most municipalities do not reassess on an annual basis, this definition should not significantly affect the results, while simultaneously preserving the size of the dataset.

Based on the sample medians of buyer and seller effective tax rates, respectively, I define a high and a low tax subsample. As a high tax paying individual might sell to a low tax buyer (i.e., with some tax exemptions), I make use of a combined measure that indicates whether a property is associated with a higher or lower tax rate. Accordingly, the high tax sample includes observations for which both buyers' and sellers' effective tax rates lie above their respective sample median. I proceed similarly for the low tax subsample, but check whether tax rates are below the same medians, respectively.

Table 2.5.2 presents results on DiD regressions for the two subsamples. Importantly, the treatment dummy cannot be displayed as it is absorbed in the fixed effects. The corresponding common trends are shown in Panels A and B of Figure 2.5.1, indicating that the common trend assumption holds well even for different subsamples. The first two columns of Table 2.5.2 show results when separating sellers by high and low tax rates. The table shows that the estimated treatment effect is at about 1.4% for the high tax sample and 40% lower for the low tax subsample. The coefficient for the low tax is only borderline significant, further indicating that the low tax group is less affected through the tax channel.

Having underpinned the tax channel as driving force of the impact of AVs on transaction prices, I further analyze how this impact varies across states of the housing market cycle. The dataset used for AVs starts in 2007, right when the housing bubble began to burst. This raises the question whether the subsequent years of turmoil that followed affected the impact of AVs on trading prices as well. For instance, when prices are more volatile, an anchor such as the AV could be more relevant for market participants, as it remains unclear where prices are heading. In contrast, the volatility of AVs itself should increase and therefore be perceived as less trustworthy for buyers and sellers, thus lowering their relevance for future tax payments. In contrast, when the economy is struggling and marginal consumption of households is high, saving taxes could become a more important issue. Taking these considerations together, it remains unclear how the situation of the economy and the financial system in the sample years 2007 to 2010 influence the effect of AVs on trading prices.

Table 2.5.3 shows results for two separate DiD regressions based on subsamples related to the year in which the respective AVs were published. The corresponding common trends are shown in Panels C and D of Figure 2.5.1. The graphs show that for both subperiods, the pre-publication trends of treatment and control group are

Table 2.5.2
Treatment effect by effective property tax rates

	High tax	Low tax
Post	0.000	-0.008*
	(0.002)	(0.003)
Treatment × Post	0.014***	0.010*
	(0.003)	(0.004)
Zip-code-tax-year fixed effects	X	X
Year-quarter fixed effects	X	X
Adj. R-sq.	0.101	0.119
Observations	48,328	59,939

This table provides coefficient estimates on four separate Difference-in-Differences regressions based on Equation (2.3.2), with DAR_{ic}, the absolute sales-price-AV ratio, as dependent variable. The subsamples are selected according to the effective tax rates of the buyer and seller, respectively. The buyer (seller) effective tax rate is derived as tax amount paid in the year after (before) the underlying transaction, divided by corresponding AV of the previous (contemporaneous) tax year, respectively. The "High tax" ("Low tax") subsample is based on observations with effective tax rates of both buyer and seller above (below) the respective sample medians. "Treatment" is a dummy that equals one if the unit was sold within a municipality-tax-year cluster in which AVs were updated, and zero otherwise. Note that the non-interacting "Treatment" dummy is absorbed by the fixed effects. "Post" is a dummy that equals one if the unit was sold after the AV was known, and zero otherwise. "×" denotes an interaction between two variables. "X" indicates that the set of control variables is used. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

remarkably similar, indicating that the common trend assumption holds well even for different states of the housing market cycle.

The first column of Table 2.5.3, depicting results for the 2007-2010 crisis subsample, shows that the coefficient estimate is very close to zero and insignificant. This could imply that with volatile prices, homeowners value the AV as an anchor, outweighing tax considerations. In contrast, the coefficient for the 2010-2017 sample is positive and at about 1.3%, slightly larger than for the base case scenario. Hence, the results indicate that in upward trending periods, market participants emphasize on tax considerations.

To investigate post-publication dynamics of anchoring and tax channel, respectively, I study subsamples according different time windows around the publication date. Intuitively, if anchoring is present, it should be strongest right after homeowners and other market participants become aware of the updated AVs. Over time, AVs become more and more outdated, such that the tax channel should become increasingly dominant.

Table 2.5.4 shows results on four DiD regressions based on subsamples of time windows with 60, 90, 120, and 150 days before and after the publication date, re-

Table 2.5.3
Treatment effect by subperiods

	Crisis (2007-2010)	2011-2017
Post	-0.005	-0.004
	(0.005)	(0.002)
Treatment × Post	-0.003	0.013***
	(0.003)	(0.002)
Zip-code-tax-year fixed effects	X	X
Year-quarter fixed effects	X	X
Adj. R-sq.	0.113	0.092
Observations	44,140	149,354

This table provides coefficient estimates on three separate Difference-in-Differences regressions based on Equation (2.3.2), with DAR_{ic}, the absolute sales-price-AV ratio, as dependent variable. "Treatment" is a dummy that equals one if the unit was sold within a municipality-tax-year cluster in which AVs were updated, and zero otherwise. Note that the non-interacting "Treatment" dummy is absorbed by the fixed effects. "Post" is a dummy that equals one if the unit was sold after the AV was known, and zero otherwise. The column names indicate period according to the publication year on which the regressions are based. "×" denotes an interaction between two variables. "X" indicates that the set of control variables is used. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

spectively. The first column, which analyzes a time window of ± 60 days around T_c shows that there is no significant effect observed immediately after publication, and the coefficient estimate is close to zero. As outlined above, this result could suggest that the anchoring channel offsets the tax channel right after the publication date, at which the AV should be most salient to homeowners.

When investigating 90 days around the publication date, the estimate becomes significant at the 1% level and closer to the base case estimate, suggesting that anchoring is reduced when AVs become less salient and more outdated over time. Expanding the time window to 150 days pre- and post-publication, the coefficient estimate is again significant at the 1% as documented in the base case analysis using a time window of 180 days. Table 2.5.4 thus indicates that first, the base case results are robust to reducing the time-window by half, and second, it takes time until the information contained in the new AVs is capitalized into transaction prices, potentially due to the offsetting anchoring channel. The base case effect of 1% is almost reached after 150 days. This suggests that after about 150 days, the tax channel reaches its maximum impact, while anchoring is only present when AVs are most salient and timely.

Figure 2.5.1
Trends of DAR_{ic} before treatment by subsamples

Panel A: High tax sellers

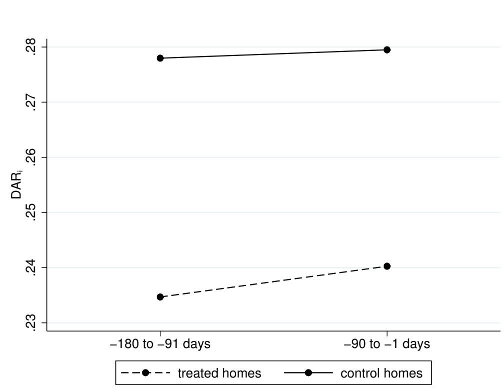

Panel B: Low tax sellers

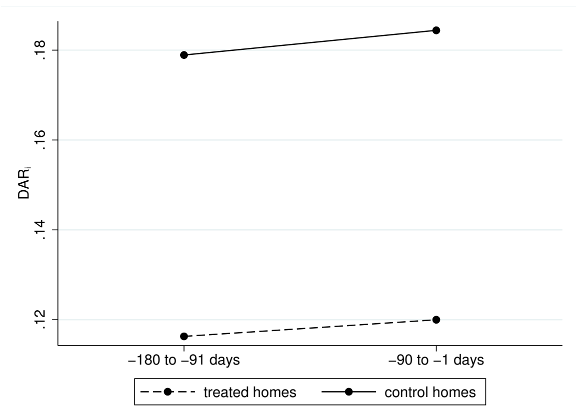

Panel C: 2007-2010

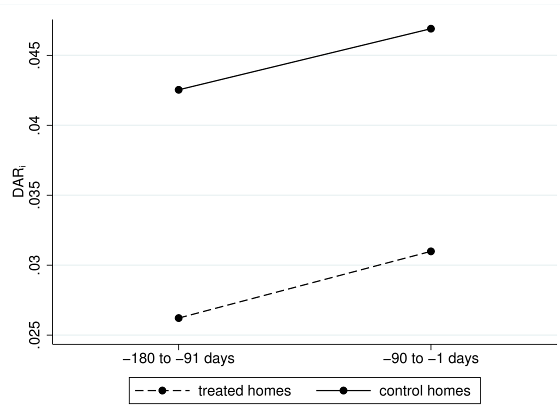

Panel D: 2011-2017

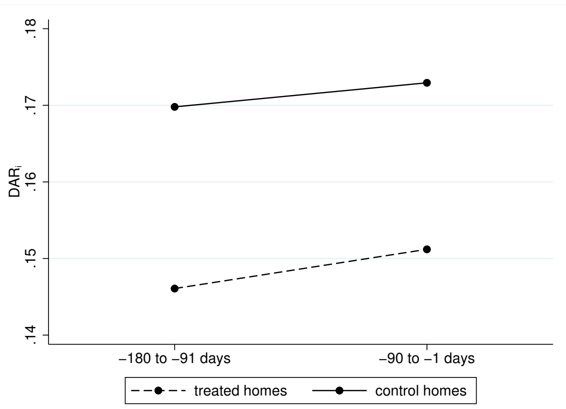

This figure shows the pre-publication development of the absolute sales-price-AV ratio, DAR_{ic}, defined in Equation (2.3.1), for both treated (dashed line) and control units (solid line). Panel A (B) shows pre-publication trends for observations for which the effective tax rates of seller and buyer is both above (below) the respective sample medians. The buyer (seller) effective tax rate is derived as tax amount paid in the year after (before) the underlying transaction, divided by corresponding AV of the previous (contemporaneous) tax year, respectively. Panel C shows pre-publication trends for homes with AVs published during the crisis period of 2007 to 2010. Panel D shows pre-publication trends for the period of recovery from 2011 to 2017.

2.6. Robustness

The purpose of this section is to analyze whether the results presented in Section 2.5 are robust with respect to several choices made throughout the paper. Table 2.6.1 summarizes all robustness checks. In Panel A, results with respect to alternative models and subsamples are presented. For the sake of brevity, only the interaction coefficient is presented. All models use zip-code-tax-year and year-quarter fixed effects,

Table 2.5.4
Results for different time windows

	60 days	90 days	120 days	150 days
Post	-0.000	-0.004	-0.005	-0.005
	(0.002)	(0.003)	(0.003)	(0.003)
Treatment× Post	0.001	0.007*	0.008**	0.009***
	(0.002)	(0.002)	(0.003)	(0.002)
Zip-code-tax-year fixed effects	X	X	X	X
Year-quarter fixed effects	X	X	X	X
Adj. R-sq.	0.105	0.102	0.098	0.096
Observations	64,805	95,427	127,529	160,488

This table provides coefficient estimates on three separate Difference-in-Differences regressions based on Equation (2.3.2), with DAR_{ic}, the absolute sales-price-AV ratio, as dependent variable. "Treatment" is a dummy that equals one if the unit was sold within a municipality-tax-year cluster in which AVs were updated, and zero otherwise. Note that the non-interacting "Treatment" dummy is absorbed by the fixed effects. "Post" is a dummy that equals one if the unit was sold after the AV was known, and zero otherwise. The column names indicate the time window relative to the publication date used for the regression: "60", "90", "120", and "150 days" before and after T_c. "×" denotes an interaction between variables. "X" indicates that the set of control variables is used. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

respectively, as in the base case. Panel B provides several placebo tests to check for differences in common trends pre-publication.

About 19% of properties analyzed in the paper are located in New York City (NYC). This might be an issue as first, NYC reassesses homes annually, which is not the case for most of the other municipalities in the sample. Second, NYC is the largest city in the US and thus might behave differently from other, less urbanized regions in the sample. I therefore check whether the base case results hold when excluding NYC properties. As shown in Panel A of Table 2.6.1, the coefficient estimate of the interaction term is at 0.8%, similar to the base case. This underpins the external validity of the results: The tax channel remains at a similar order of magnitude, whether properties are located in highly urbanized areas or in comparably rural ones. As NYC and Nassau county impose limits on how much the taxable share of the AV can increase, I check whether my results hold when excluding observations in these locations. The point estimate of 0.7% that is shown in the second row indicates that this appears to be the case.

When investigating only larger municipalities, defined by having at least 1,000 observations per municipality-tax-year cluster, the treatment effect is still highly significant and positive. The same is true for subsamples divided by the median nominal sales price (about USD 280,000) in the sample. The treatment effect is slightly higher for above-median sales prices with a coefficient estimate difference of 0.3%. Together, these results document that the tax channel is consistently dominating in different market segments.

As mentioned earlier in the paper, using the absolute value of a relative measure, such as the base case variable DAR, might be associated with some drawbacks, as undervaluations of, e.g., 90% are valued the same as overvaluations of the same magnitude, even though such undervaluations can be considered to be much more extreme. I therefore limit DAR to be less than 50% to see whether my base case results still hold. With a significant coefficient of 0.9%, this appears to be the case. I further investigate the unstandardized measure DA. The coefficient is positive and highly significant as well, suggesting that my results are robust with respect to standardization.

Panel B of Table 2.6.1 depicts regression results for three placebo tests using a subsample of pre-publication observations. The placebo date is set to 90 days prior to the actual publication, thus symmetrically dividing the subsample period. The idea is to check for the existence of differences in pre-publication trends between treatment and control units. Models (1)-(3) in Panel B indicate no significant difference between treatment and control group in the dependent variable after the pseudo-publication date. For the base case scenario, shown in column (3), the coefficient estimate is very close to zero, further supporting the assumption of common trends.

2.7. Conclusion

Property taxes are commonly based on assessed values (AVs) of homes which can only be estimated and thus are prone to valuation errors. This paper analyzes whether AVs themselves affect trading prices of homes. Theoretically, the impact of AVs on transaction prices should be driven by two opposing channels. As AVs constitute salient reference prices for market participants, an anchoring channel indicates a positive causal effect of AVs on transaction prices. At the same time, an unexpected increase in the AV implies a higher future property tax burden associated with the underlying home. Thus, a tax channel should lead to a negative effect of AVs on sales prices.

I analyze this ambiguous relationship with a Difference-in-Differences framework that exploits AV publication dates and variation in reassessment frequencies. The results document a robust negative causal impact of AVs on transaction prices, in line with the tax channel. With an AV-induced change of about 0.9% in the transaction

Table 2.6.1
Robustness checks

Panel A: Alternative models and subsamples			
	Treatment × Post	Observations	Adj. R-sq.
Without NYC	0.008***	156,195	0.096
	(0.002)		
Without NYC & Nassau	0.007**	123,646	0.097
	(0.002)		
Larger municipalities only	0.011**	67,212	0.091
	(0.003)		
Below median sales price	0.008**	95,927	0.167
	(0.003)		
Above median sales price	0.011***	97,567	0.117
	(0.002)		
$DAR < 0.5$	0.007***	183,876	0.079
	(0.002)		
DA as dependent variable	0.085***	193,494	0.302
	(0.013)		

Panel B: Placebo tests			
	(1)	(2)	(3)
Post (placebo)	0.005*	0.005*	0.004
	(0.002)	(0.002)	(0.003)
Treatment × Post (placebo)	0.005	0.004	0.002
	(0.003)	(0.003)	(0.003)
Location-tax-year fixed effects	municipality	zip-code	zip-code
Year-quarter fixed effects	-	-	X
Adj. R-sq.	0.073	0.107	0.108
Observations	83,755	83,755	83,755

This table documents the results of multiple robustness checks. Panel A shows estimates for the interaction term of treatment group and post-AV publication ("Treatment × Post") that informs about the causal effect of AV publication on transaction prices, measured by the change in the dependent variable DAR, the absolute value of the sales-price-AV ratio. Each row of Panel A shows results for a separate regression. The specification is indicated in the first column. "Without NYC" indicates that all observations placed in New York City have been left out, and "Without NYC & Nassau" indicates that all observations placed in New York City or Nassau county are excluded, respectively. "$DAR < 0.5$" indicates an upper bound for the dependent variable DAR of fifty percent. For model "Larger municipalities only", only municipality-tax-year clusters with more than 1,000 observations have been included. "Below (Above) median sales price" investigates subsamples of transactions with sales prices of less than (equal or more than) the median nominal sample sales price (about USD 280,000). "DA as dependent variable" indicates that the (unstandardized) absolute distance between sales price and AV was used as dependent variable. For all regressions in Panel A, the base case scenario with zip-code-tax-year and year-quarter fixed effect dummies, respectively, is employed. Each column in Panel B shows a placebo regression with varying control variables. All regressions investigate observations before AV publication. The "Post (placebo)" dummy was defined to split the pre-publication date into two sales with equal time window (90 days each). "X" indicates that the set of control variables is used, whereas "-" indicates that the controls are not used for the specification. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

price of the average home, the results of this paper are not only statistically robust, but also of economic significance. Supporting the tax channel as driver of my findings, homes associated with high tax rates are affected stronger than homes related to lower tax rates. The results documented in this work are robust with respect to different subsamples and fixed effects. The findings of this work have implications for the redistribution of wealth through the property tax system: Imperfect valuation does not only lead to inequitable tax payments, but also adversely distorts transaction prices of homes.

References

Allen, M. and W. Dare (2002). "Identifying Determinants of Horizontal Property Tax Inequity: Evidence from Florida." *Journal of Real Estate Research* 24. (2), 153–164.

Andersen, S., C. Badarinza, L. Liu, J. Marx, and T. Ramadorai (2021). "Reference Dependence in the Housing Market." Available at SSRN: https://ssrn.com/abstract=3396506.

Bai, C., Q. Li, and M. Ouyang (2014). "Property Taxes and Home Prices: A Tale of Two Cities." *Journal of Econometrics* 180. (1), 1–15.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). "The Economic Effects of Social Networks: Evidence from the Housing Market." *Journal of Political Economy* 126. (6), 2224–2276.

Black, R. T. and J. Diaz III (1996). "The Use of Information Versus Asking Price in the Real Property Negotiation Process." *Journal of Property Research* 13. (4), 287–297.

Bollerslev, T., A. J. Patton, and W. Wang (2016). "Daily House Price Indices: Construction, Modeling, and Longer-Run Predictions." *Journal of Applied Econometrics* 31. (6), 1005–1025.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). "Robust Inference With Multiway Clustering." *Journal of Business and Economic Statistics* 29. (2), 238–249.

Cypher, M. and J. A. Hansz (2003). "Does Assessed Value Influence Market Value Judgments?" *Journal of Property Research* 20. (4), 305–318.

DeFusco, A., C. Nathanson, and E. Zwick (2020). *Speculative Dynamics of Prices and Volume*. Working Paper 23449. National Bureau of Economic Research.

Du, Z. and L. Zhang (2015). "Home-Purchase Restriction, Property Tax and Housing Price in China: A Counterfactual Analysis." *Journal of Econometrics* 188. (2), 558–568.

Elinder, M. and L. Persson (2017). "House Price Responses to a National Property Tax Reform." *Journal of Economic Behavior & Organization* 144, 18–39.

Fischer, M., R. Füss, and S. Stehle (2021). "Local House Price Comovements." *Real Estate Economics* 49. (S1), 169–198.

Genesove, D. and C. Mayer (2001). "Loss Aversion and Seller Behavior: Evidence From the Housing Market." *Quarterly Journal of Economics* 116. (4), 1233–1260.

Gibbs, C. G. and M. Kulish (2017). "Disinflations in a Model of Imperfectly Anchored Expectations." *European Economic Review* 100, 157–174.

Goolsby, W. (1997). “Assessment Error in the Valuation of Owner-Occupied Housing.” *Journal of Real Estate Research* 13. (1), 33–45.

Hilber, C. A. (2017). “The Economic Implications of House Price Capitalization: A Synthesis.” *Real Estate Economics* 45. (2), 301–339.

Hodge, T. R., D. P. McMillen, G. Sands, and M. Skidmore (2017). “Assessment Inequity in a Declining Housing Market: The Case of Detroit.” *Real Estate Economics* 45. (2), 237–258.

Jones, P. (2020). “Loss Aversion and Property Tax Avoidance.” Available at SSRN: https://ssrn.com/abstract=3511751.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). “The Housing Market(s) of San Diego.” *American Economic Review* 105. (4), 1371–1407.

Levy, D., Z. Dong, and J. Young (2016). “Unintended Consequences: The Use of Property Tax Valuations as Guide Prices in Wellington, New Zealand.” *Housing Studies* 31. (5), 578–597.

Livy, M. R. (2018). “Intra-School District Capitalization of Property Tax Rates.” *Journal of Housing Economics* 41, 227–236.

McMillen, D. P. (2013). “The Effect of Appeals on Assessment Ratio Distributions: Some Nonparametric Approaches.” *Real Estate Economics* 41. (1), 165–191.

Ngai, L. R. and S. Tenreyro (2014). “Hot and Cold Seasons in the Housing Market.” *American Economic Review* 104. (12), 3991–4026.

Northcraft, G. and M. Neale (1987). “Experts, Amateurs, and Real Estate: An Anchoring-and-Adjustment Perspective on Property Pricing Decisions.” *Organizational Behavior and Human Decision Processes* 39. (1), 84–97.

Palmon, O. and B. A. Smith (1998). “New Evidence on Property Tax Capitalization.” *Journal of Political Economy* 106. (5), 1099–1111.

Sirmans, S., D. Gatzlaff, and D. Macpherson (2008). “Horizontal and Vertical Inequity in Real Property Taxation.” *Journal of Real Estate Literature* 16. (2), 167–180.

Tversky, A. and D. Kahneman (1974). “Judgment Under Uncertainty: Heuristics and Biases.” *Science* 185. (4157), 1124–1131.

Wassmer, R. W. (1993). “Property Taxation, Property Base, and Property Value: An Empirical Test of the “New View”.” *National Tax Journal* 46. (2), 135–159.

Appendix 2.A The property tax system in New York State

This section gives an extended overview of the New York State property tax system. Figure 2.A.1 illustrates the sequence of important dates according to the official tax calendar that is followed by all municipalities in the state.[15]

In each jurisdiction, AVs are published annually at the "tentative roll date", T, as shown in the center of Figure 2.A.1. Beginning at T, the new information contained in the updated AVs cannot only be used by homeowners, but also by other market participants, as property assessments are publicly available. The most common tentative roll date is May 1, but other dates are used as well, such as January 15 in New York City. Although AVs are published every year, homes are not necessarily reassessed at the same frequency. The frequencies are decided by municipal governments and vary substantially. For instance, homes in New York City are reassessed annually, while some smaller municipalities, such as Westerlo, did not reassess since 1974.[16] In some years, AVs either remain unchanged or are collectively adjusted by the same percentage, e.g., to correct for inflation. As discussed in Section 2.4.1, I exclude municipality tax-year clusters in which AVs have been collectively adjusted by the same percentage, as the corresponding observations can be neither assigned to treatment nor control group.

Prior to AV publication at T, there are two dates that are relevant for this work. The first is the "valuation date". If a municipality reassesses homes for the upcoming tax year, the AVs have to reflect home values at this particular day. AVs are thus not meant as a forecast, but reflect pre-publication prices. Valuation of residential real estate is typically done with a sales comparison approach. AVs are thus based on past sales prices of units that are similar to the home that is assessed. The second relevant date before T is the "taxable status date". The AV has to be based on the condition of the considered home at this particular date. Hence, homes whose condition changes after this date (e.g., through a fire), should be severely mispriced by authorities. In the data cleaning process, I therefore remove observations with extreme one-year AV returns.

At T, AVs are not final. Until the "grievance day", typically four weeks after T, homeowners have the chance to contest their home's assessment. This is a potential concern for the empirical analysis, since the values are not necessarily fixed at T, raising

[15] The link to the New York State Department of Taxation and Finance website is given in Section 2.2.1.

[16] As reported by the New York State Department of Taxation and Finance on https://www.tax.ny.gov/pdf/ORPTS/recent-reassessments.pdf, last retrieved on March 15, 2021.

Figure 2.A.1
The New York State property tax calendar

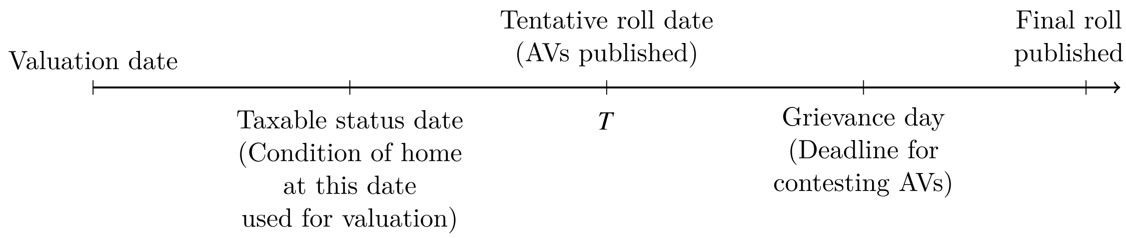

This figure visualizes the timeline of the assessment process in New York State for a representative municipality in a given year. Market value estimates have to reflect the homes' values at the "valuation date". Homes have to be valued according to the condition of the home at the "taxable status date". AVs are published at the "tentative roll date" T. AVs can be challenged until the "grievance date". The "final roll" is published at the beginning of the new tax year.

endogeneity concerns. In Section 2.4.2, I show that successful contests are relatively rare in a dataset from New York City, affecting less than 1% of all single-family homes, and thus are likely to be of minor concern. AV contests that have not been accepted are not necessarily final, but homeowners have to apply for judicial review, making further changes rather difficult.

The "final roll" is published at the beginning of the new tax year and contains the finalized AVs. Until the next reassessment, which fraction of AVs is taxed depends on local and individual factors. These include the local tax rate, the local assessment ratio, local budgeting, as well as individual exemptions, such that effective tax rates can vary substantially even within the same jurisdiction. Additionally, there are two policies that are limiting the increase of tax payments. First, in New York City (NYC) and Nassau county, assessed values can maximally increase by 6% per year or 20% over five years. Second, changes in the total tax levy (i.e., how much taxes a given entity can collect in total) are capped in most jurisdictions (e.g., counties outside NYC and independent school districts) by the minimum of either two percent or the CPI inflation rate. Both caps should not be an issue for this paper. First, In the robustness section, I show that my results are robust to the exclusion of observations in NYC and Nassau. Second, tax levy limits do not rule out substantial individual changes in tax payments. For instance, an AV increase relative to other properties in the municipality would lead to an increase in tax payments even if the total tax levy remained constant.

Appendix 2.B Derivations

2.B.1 Derivations for the tax channel

The tax channel increases the absolute sales-price-AV distance, DA, post AV publication if

$$
\begin{aligned}
& DA^{pre}(\bar{n}) \leq DA^{post}(\bar{n}) && \text{(A1)}\\
\Leftrightarrow \quad & \bar{n}\left|\dot{p}^{*} - \tau E[\hat{p}] - \hat{p}\right| \leq \bar{n}\left|\dot{p}^{*} - \tau \hat{p} - \hat{p}\right| && \text{(A2)}\\
\Leftrightarrow \quad & \left|\dot{p}^{*} - \tau E[\hat{p}] - \hat{p}\right| \leq \left|\dot{p}^{*} - \tau \hat{p} - \hat{p}\right|, && \text{(A3)}
\end{aligned}
$$

holds. The last equivalence follows as the total housing supply is larger than zero, i.e., $\bar{n} > 0$. The expected tax payments, $\tau E[\hat{p}]$, are replaced with the true value, $\tau \hat{p}$, once the AV is published (post). Using $(|a| \leq b) \quad \Leftrightarrow \quad (-b \leq a \leq b)$, inequality (A1) holds if

$$
\begin{aligned}
-\left|\dot{p}^{*} - \tau \hat{p} - \hat{p}\right| &\leq \dot{p}^{*} - \tau E[\hat{p}] - \hat{p} \qquad \text{and} && \text{(A4)}\\
\left|\dot{p}^{*} - \tau \hat{p} - \hat{p}\right| &\geq \dot{p}^{*} - \tau E[\hat{p}] - \hat{p}. && \text{(A5)}
\end{aligned}
$$

It is now to show under which conditions inequalities (A4) and (A5) hold. Starting with (A4), as $(|a| \geq b) \quad \Leftrightarrow \quad (a \leq -b) \quad \text{or} \quad (a \geq b)$, it must hold that either

$$
\begin{aligned}
\dot{p}^{*} - \tau \hat{p} - \hat{p} &\geq \hat{p} - \dot{p}^{*} + \tau E[\hat{p}] \qquad \text{or} && \text{(A6)}\\
\dot{p}^{*} - \tau \hat{p} - \hat{p} &\leq \dot{p}^{*} - \tau E[\hat{p}] - \hat{p}. && \text{(A7)}
\end{aligned}
$$

Defining $\dot{p}^{pre} := \dot{p}^{*} - \tau E[\hat{p}]$ and $\dot{p}^{post} := \dot{p}^{*} - \tau \hat{p}$, it is easy to show that it follows from inequality (A6) that

$$
\frac{\dot{p}^{pre} + \dot{p}^{post}}{2} \geq \hat{p}, \tag{A8}
$$

and from (A7) that

$$
E[\hat{p}] \leq \hat{p}. \tag{A9}
$$

Proceeding with similar calculations and arguments for , it can be shown that (A5) is true if either

$$
\frac{\dot{p}^{pre} + \dot{p}^{post}}{2} \leq \hat{p} \qquad \text{or} \tag{A10}
$$

$$E[\hat{p}] \geq \hat{p} \tag{A11}$$

holds. This means that (A1) is fulfilled when one out of the two inequalities (A8) and (A9) holds together with either (A10) or (A11). This leaves four combinations to check. Inequalities (A8) and (A10) hold together if $\hat{p} = E[\hat{p}]$. The same is true for inequalities (A9) and (A11). From these two pairs, it follows that DA remains unchanged when expectations match the outcome. Now, consider (A8) and (A11). From (A11), it follows that $\dot{p}^{post} \geq \dot{p}^{pre} \Leftrightarrow \epsilon > 0$ and thus rewriting (A8) yields

$$\frac{2\dot{p}^{pre} + \epsilon}{2} \geq \hat{p}, \tag{A12}$$

in which $\epsilon := \dot{p}^{post} - \dot{p}^{pre}$ is the difference between expected AV and actual AV times the effective tax rate, it follows that (A1) holds if

$$\hat{p} \leq E[\hat{p}] \leq \dot{p}^{pre} + \frac{|\epsilon|}{2}. \tag{A13}$$

Proceeding with similar calculations and arguments (A10) and (A9), it follows that (A1) holds if

$$\hat{p} \geq E[\hat{p}] \geq \dot{p}^{pre} - \frac{|\epsilon|}{2}. \tag{A14}$$

Hence, inequalities (A13) and (A14) are sufficient conditions for (A1).

2.B.2 Derivations for the back-of-the-envelope calculations

Defining the price after treatment as P^t and the price under control conditions as P_c, per definition, the estimated treatment effect γ is the difference in DAR between treatment and control condition, given as

$$\left|\frac{P^t - AV}{AV}\right| - \left|\frac{P^c - AV}{AV}\right| = \gamma. \tag{A15}$$

Note that the AV is the same for both scenarios, mirroring the empirical application in which the contemporaneous AV is used to derive the absolute sales-price-AV ratio. For simplicity, assume that P^t and P^c are larger than AV. Then, it follows from Equation (A15) that

$$P^t - AV = P^c - AV + \gamma AV, \tag{A16}$$

which can be rearranged to

$$\Delta P_{ic} = \frac{P^t - P^c}{P^c} = \frac{\gamma AV}{P^c}. \tag{A17}$$

To see that these calculations fit to the model presented in Section 2.2, note that the assumption made that P^t and P^c are smaller than the AV is consistent with condition (2.2.12), which implies that the sales price is increasing as well. Similarly, when assuming that P^t and P^c are smaller than AV instead, it follows that

$$\frac{P^t - P^c}{P^c} = -\frac{\gamma AV}{P^c}, \tag{A18}$$

which is again in line with condition (2.2.11).

Appendix 2.C Data handling

I merge the tax record database with the housing transactions as follows. For each municipality tax-year cluster c, I define a symmetric time-window of 360 days around the publication date T_c. For all transactions that are observed for each c, I identify the particular property-specific AV that was published at the corresponding date T_c. Thus, the AV that is matched to a transaction i was unknown if i was sold in the 180 days before T_c, and known if i was sold in the 180 days afterwards instead. The merging process described is applied to all transactions that remain after the cleaning procedure that is described below.

I focus on single-family homes as they are associated with a unique assessor's parcel number (APN) which is used as a key for the matching process. Starting out with 1,253,756 arms-length transactions from December 2006 (the earliest possible date to identify a matching AV for) to December 2017 with non-missing sales price and available APN, I follow several steps of DeFusco et al. (2020). First, I remove all foreclosure related transactions (1,184,960 remain). Second, I remove duplicates according to two criteria. If names of buyers and sellers, as well as sales prices are identical for more than one observation, I keep only the earliest recorded transaction. If there are multiple transactions for a particular property at the same day, all but one transaction is dismissed (1,129,881 remain). Similar to Bollerslev et al. (2016), I set fixed bounds on the nominal values of USD 5,000 and USD 10,000,000. Thus the high right-skewness of the price distribution is mitigated, and scaling is more appropriate for the regression analysis. This leaves me with a cleaned transaction sample of 1,126,596 observations.

The cleaned transaction data is matched to the publication dates obtained from the Municipal Profile webpage of the New York State government. The general publication date is May 1, but some municipalities choose to deviate from this date, e.g., to match other budget related purposes. I use the Statewide Information System (SWIS) codes that uniquely identify municipalities to match the respective publication dates to the observations.[17] Thus, observations for which the SWIS code is missing cannot be used and must be dismissed (1,125,338 observations remain). Finally, the transactions are matched to the tax records. Here, each transaction is assigned to the up-to-date AV that is (yet) unknown if the property transacted before publication (i.e., in the time

[17]For New York City, I use the county FIPS codes. For most of the remaining counties, the SWIS code is part of the assessor's parcel number. For the few counties that do not follow this coding procedure, SWIS codes are given in another variable of the dataset that is identified manually.

window -180 to -1 days), and known afterwards (day 0 to 180). The matched dataset contains 747,254 observations.

For the analysis, the one-year AV lag is necessary for two reasons. First, to identify the control group, which requires that homes have not been reassessed in the given tax year. Second, extreme changes in valuation are possible as the condition of a home could have changed dramatically, e.g., through fire damage. Accordingly, I dismiss observations for which the prior year's AV is either missing or for which the one-year AV return exceeds the bounds of the second and 98th percentile. I chose the second and 98th percentile, as trimming based on the first and 99th percentiles did leave extreme outliers in the sample. After these steps, 516,165 observations remain.

I identify treatment and control group in a data driven way. For each of the remaining observations, I calculate the one-year AV return and round it up to the third decimal place to correct for small differences in returns. I identify three different types of municipality-tax-year clusters based on how many observations within a cluster have the same return. First, clusters for which the mode return makes up less than 75% of returns is assigned treatment status. Second, clusters for which the return of at least 75% of observations equals zero is assigned control status. Third, clusters in which the mode return makes up more than 75% of observations, but the mode return is different from zero, is assigned the market- or inflation-adjusted status.

Based on these definitions, I dismiss observations in the control group with returns different from zero, as these are likely to be caused by unusually large physical changes, and AV adjustments are likely to be expected for such properties. I further dismiss all observations in the market- or inflation-adjusted group as they are neither fully treatment nor control group. Additionally, observations in municipalities (SWIS codes) that never reassess homes in the sample are also dismissed to increase the comparability of the control group. After these steps, 315,402 observations remain. Once the observations from the market-or inflation-adjusted group are removed, I deal with extreme assessment ratios which are commonly observed in housing markets (e.g., McMillen, 2013). To rule out that my findings are driven by such outlier properties, I dismiss trades of homes for which either the ratio of sales price and AV or the absolute distance between both quantities is extreme (again 2nd and 98th percentile). The latter is done for both treated and control group separately to account for disparities in the price distribution. After this step, 288,546 observations remain.

To increase comparability between control and treatment group, I conduct a final step. As treated municipalities are typically larger than the control group ones, I adjust both groups such that at least 100 observations per municipality-tax-year cluster are required. The final dataset consists of 193,494 observations.

Chapter 3

Investors in the Housing Market*

*We are grateful for comments and suggestions from Alexander Braun, Martin Brown, Jens Jackwerth, Axel Kind, Winfried Koeniger, Steffen Sebastian, and seminar participants at the Universities of Konstanz, Regensburg, and St. Gallen. Marcel Fischer and Simon Stehle gratefully acknowledge financial support from the German Research Association (DFG), grant FI2141/5-1.

3.1. Introduction

Residential real estate differs from traditional asset classes, such as stocks or bonds, in multiple ways. For instance, while it is notoriously difficult to predict future stock prices, housing markets are subject to strong and positive momentum (Case and Shiller, 1989). Additionally, homes are non-divisible and comparably expensive, making housing risk difficult to diversify. Simultaneously, the housing market is occupied by a wide range of buyers with contrasting strategies and targets, varying in, e.g., investment horizon, risk-taking, and whether a home is primarily bought as consumption good or financial asset. Consequently, buyers' and sellers' ties to local momentum and housing risk are likely to be highly divergent, suggesting that the heterogeneous exposure to risk and expected returns is crucial in understanding investment behavior in these markets.

In this paper, we investigate the role of investor-specific exposure to local risk and expected returns for the respective performances in capital gains. Our analyses are based on four prevailing categories of market participants that are likely to be heterogeneous in their main objectives as well as their underlying strategies. Our first category consists of owner-occupiers, whose primary interest should lie in consumption rather than financial yields. As consumption implies selection based on preferred location and taste, the compensation for engaging in riskier trades should be of secondary importance for this group. In contrast, we expect investors in the housing market to seek high rents through increased risk-taking. We analyze private investors, who might buy close-by homes as future savings for retirement, but were also active during the last housing boom (e.g., Bayer et al., 2020, 2021), suggesting high dependence on momentum and increased risk exposure. We further analyze short-term institutional investors (ST IIs), who are likely to speculate on high appreciation rates and should thus be highly exposed to local momentum and risk. Our final group, long-term institutional investors (LT IIs), should be more likely to trade on stable dividends in form of rental income, which implies a relatively lower exposure to momentum and risk, respectively.

We analyze capital gains that stem from more than 21 million repeat sales of US residential real estate and document persistent and sizable differences among the four groups of market participants. Our first finding documents that these differences can be neither explained by location choice, nor timing on the aggregate or local level, respectively. In contrast, the heterogeneous exposure of investors' capital gains to past local return risk explains a sizable share of capital gain disparities. Moreover, after controlling for several factors, such as documented construction activities, holding period, and county-level fixed effects, the outperformance of private investors and ST IIs can even be entirely explained by heterogeneous return risk exposure. In contrast, even though investors' exposure to local momentum is highly heterogeneous, it can explain only little of the performance differences. Our results thus document that the systematically higher capital gains realized by investors can be linked to the higher risk-taking associated with their strategies.

Our analyses are based on the generalized portfolio sorts model proposed by Hoechle et al. (2020) that we convey to our housing market setting. Being able to perfectly replicate the Jensen et al. (1972) portfolio sorts approach in a single regression, the methodology allows us to simultaneously estimate and test the relative performance of investor groups compared to owner-occupiers, as well as the exposure of each group to one or more factors. Hence, the procedure allows to run multiple tests at the same time and to include control variables as well as fixed effects that might affect both, factor exposure and relative performance. For instance, our regressions show that the relative alphas of private and short-term institutional investors become insignificant after including local return risk as a factor. The results are robust to instrumenting volatility similar to Han (2013), alternative definitions of "short-term", and different subsamples. In contrast, the relative alphas of investors remain large and significant after including multiple other factors such as local or nationwide momentum.

The key quantity of our analysis consists of annualized capital gains in excess to the risk-free rate. We focus on this measure rather than rental income for several reasons. First, capital gains are of interest to all four types of investors. Even owner-occupiers should have high interest in homes with high appreciation rates, as they profit from increased consumption (Campbell and Cocco, 2007), entrepreneurship (Corradin and Popov, 2015), higher local economic growth, and (Loutskina and Strahan, 2015), as well as from a hedge against future housing consumption risk (Han, 2013). Second, capital gains can be directly observed for all trades. In contrast, rents must either be imputed for the largest share of homeowners, as owner-occupants are by far the largest group, or used at the aggregate level, in which case they can be addressed with fixed effects. Third, rents tend to be stronger regulated than prices (e.g., rent controls in New York City), and regulations vary across jurisdictions, making them difficult to analyze in a nationwide setting. Additionally, in efficient markets, expected rental income should be priced in the purchase transaction. Thus, the comparatively stable rental income is further addressed with our repeat sales approach and additionally captured by local fixed effects. Lastly, we control for the existence of construction related documents associated with a transaction to mitigate the impact of larger renovations, as well as for new constructions to account for land speculation (Nathanson and Zwick, 2018).

Our work provides a unifying framework to analyze and compare performance and risk exposure of different types of investors that have been investigated in the prior literature. For instance, similar to our category of private investors, second-home buyers have been analyzed by Chinco and Mayer (2016) and Cvijanovic and Spaenjers (2020), focusing on out-of-town and foreign buyers, respectively. Bayer et al. (2021) document investor contagion from experienced investors towards new, inexperienced ones who then performed worse than their professional counterparts. Speculators, typically engaging in short-term activity, have been extensively analyzed in the literature (e.g., Bayer et al., 2020; Fu and Qian, 2014; Fu et al., 2016). For instance, short-term speculators drive up trading volume, as documented by DeFusco et al. (2020). Institutional investors further sped up the recovery of housing markets after the recent bust, as shown by Lambie-Hanson et al. (2020). Concerning long-term investors that should be interested in rental income, Mills et al. (2019) analyze buyers of single-family homes that operate on a larger scale to securitize the respective payment streams.

We contribute to the literature on investors in the housing market by showing that a sizable share of the systematic differences in capital gains across prevailing investor groups can be linked to heterogeneous risk exposure. Our results thus suggest that the strategies of investors analyzed in the prior literature are associated with a change in risk-taking, which in turn results in increased compensation. Simultaneously, we show that trading on momentum can be beneficial in general, but can explain only little in the disparities across investor groups. Our analyses not only focus on (selected) MSAs but on the entire US mainland, for a period of more than two decades. Thus, we can shed light on the effects of location choice, timing on aggregate as well as local markets, associated risks and returns. Our results thus not only have implications for the most densely populated areas, but also for more rural regions. Moreover, our results are not only relevant for a particular investor group, but for a wide range of homeowners in the US.

By highlighting the importance of local return risk exposure for the understanding of investors' outperformance of owner-occupiers, we contribute to a growing strand of literature that analyzes the relationship of risk and return in the housing market. Han (2010) shows that incentives to buy a home as hedge against future housing risk can influence the demand for homes in local markets, uncovering a heterogeneous impact of risk on local markets. These hedging incentives are further affecting the risk-return relationship across metropolitan statistical areas (MSAs). As documented by Han (2013), hedging incentives against future housing risk can even lead to a negative risk-return relationship in some areas. Similarly, Peng and Thibodeau (2017) find that, on average, idiosyncratic risk is not compensated with higher appreciation rates. We further contribute to this strand of literature by relating local risk and return to the activity of multiple classes of investors as well as their underlying performance in capital gains. Our results document that even within the same markets, investors' risk exposure is heterogeneous, which can additionally help explain the ambiguous relationship of risk and return in housing markets.

The remainder of this paper is structured as follows. In Section 3.2, we describe our dataset, the associated cleaning process, and how we derive our key variables. Our methodology and results are presented in Section 3.3. Section 3.4 provides robustness checks. The final section concludes.

3.2. Data

In the following section, we present our data on individual transactions of residential real estate, along with a detailed description of the data cleaning process. In Section 3.2.2, we explain how investor identities are defined and assigned to each observation. Our key performance measure, annualized excess capital gains from repeated transactions of homes, as well as our measures for local return risk and local momentum, are described subsequently in Section 3.2.3. Additional variables and further data used are described in Section 3.2.4. Finally, we present summary statistics of the resulting dataset in Section 3.2.5.

3.2.1. Housing transactions

We obtain data on US property transactions from data vendor *CoreLogic*. The dataset is representative, covering more than 99% of US properties.[1] We start our data cleaning process with 99,757,949 arms-length transactions of residential real estate with non-missing sales price and complete sale date from 1995 to 2017. To remove duplicate transactions, we follow DeFusco et al. (2020). First, if there is more than one transaction with coinciding buyer and seller names as well as the same sales price, we only keep the transaction with the earliest recording date. Second, if the same property is transacted multiple times a day, we keep only one of the given transactions

[1] The property record coverage is given by the data vendor on https://www.corelogic.com/solutions/university-data-portal.aspx. Last retrieved on March 16, 2021.

(85,966,973 remain). To account for extreme prices that are likely to be associated with data errors, we follow Bollerslev et al. (2016) by setting fixed nominal bounds of USD 5,000 and USD 100,000,000 and removing all transactions that fall outside the resulting interval (85,641,056 remain).

Out of the remaining transactions, we are able to form 33,280,346 repeat sales. Again following Bollerslev et al. (2016), we exclude observations with extreme characteristics by first, dismissing capital gains of less than -50% or more than 100% per year, and second, removing transactions with a holding period of less than 180 days (29,748,815 remain). Afterwards, we assign an investor identity to each repeat sale, as described in detail in Section 3.2.2. Observations for which the identity remains unknown, e.g., ones with missing address of a private homeowner, are dismissed (26,217,238 remain).

Our cleaning process is finalized by two steps. After analyzing the distribution of annualized capital gains in the remaining sample, we dismiss observations in the first and last percentile (25,692,895 remain). Lastly, we again follow the literature (e.g., DeFusco et al., 2020) in removing all foreclosure or real estate owned (REO) related repeat sales. We conduct this step at the final stage for two reasons. First, removing these observations before matching transactions of the same home to repeat sale pairs would lead to "false" transaction pairs and thus wrong assignment of investor identities. For instance, consider a home is transacted three times during the sample period with the second transaction being a foreclosure. An early omission of the foreclosure would lead to a misleading match of first and third transaction and to the wrong assignment of investor identity to the third sale. Second, including foreclosure related transactions when determining capital gain outliers helps to identify the most extreme values more precisely. Intuitively, as sales prices associated with a foreclosure are typically subject to a large discount (e.g., Anenberg and Kung, 2014; Campbell et al., 2011; Gerardi et al., 2015), transactions not flagged as foreclosures with similar or lower return can be more plausibly identified as irregular. The final sample contains 21,178,869 observations. The data cleaning process is summarized in Table 3.A.1 in the Appendix.

3.2.2. Defining investor types

We define mutually exclusive investor types that should have a high propensity to differ in their underlying strategy. Our first group encompasses owner-occupiers, i.e., buyers that live in the home underlying the transaction. Similar to recent literature on speculation in housing markets, we define a separate, private category that consists of second-home buyers (e.g., Chinco and Mayer, 2016; Cvijanovic and Spaenjers, 2020; Gao et al., 2020). We further divide institutional investors in two categories. First, short-term operating professionals, who should mainly be interested in capital gains, and long-term institutional investors, who should primarily focus on steady rental income such as the large-scale single-family home buyers analyzed in Mills et al. (2019). Technically, we assign one of the four investor identities to transactions as follows.

We first distinguish between corporate and private identities using an identifier provided in the *CoreLogic* dataset. To divide private buyers in owner-occupiers and investors, we compare homeowners' mailing addresses with the ones of the corresponding properties following DeFusco et al. (2020). As primary indicator for owner-occupancy status, we check if house numbers of both addresses match. If at least one of the two house numbers is missing, we compare street names instead. If the respective criterion indicates a match, owner-occupancy status is assigned, and private investor identity otherwise. If an observation is not providing sufficient information, we dismiss the corresponding datapoint.

Institutional investors are divided into either ST or LT II according to the respective holding period of the underlying transaction. Recent literature shows that speculation is associated with very short holding periods. For instance, Bayer et al. (2020), show that 90% of large-scale flippers resell homes within four years after the purchase. We follow DeFusco et al. (2020) and define a holding period of less or equal three years as short-term, and assign long-term status otherwise.[2]

3.2.3. Construction of key variables

In this paper, we are interested in comparing capital gains across investor types. However, capital growth rates themselves are difficult to compare in the cross-section, as they are first, realized over different holding periods, and second, originated between different points in time. Thus, we make two adjustments similar to Fischer et al. (2021). To account for varying holding periods, we annualize each growth rate using the respective holding period. Additionally, we correct for time-varying opportunity costs by subtracting the annualized return of an investment in the risk-free rate over

[2] In the robustness Section 3.4, we show that our base case results are robust to alternative cut-offs at two or four years, respectively.

the same period. Accordingly, we derive annualized excess capital gains, $r_{it_1t_2}$, of home i, bought at t_1 and sold at t_2 as

$$r_{it_1t_2} = \left(\frac{P_{it_2}}{P_{it_1}}\right)^{\frac{1}{y(t_1,t_2)}} - \left(R^f_{t_1t_2}\right)^{\frac{1}{y(t_1,t_2)}}, \tag{3.2.1}$$

in which P_{it} is the nominal sales price at time t, $y(t_1, t_2)$ is the holding period in years, and $R^f_{t_1,t_2}$ is the gross return of the risk-free rate from t_1 to t_2.

We then derive measures for local returns to infer about local price expectations (momentum) and for local return risk to measure the local uncertainty about capital gains. We capture local price dynamics with the Case and Shiller (1989) methodology, which requires the repeat sales structure that is given in our data.[3]

To measure to which extent investors' outperformance can be linked to heterogeneous risk exposure, we propose a simple measure that directly relies on the local dispersion of returns. Local return risk is derived as the standard deviation of all annualized capital gains that are realized in a given period in a given location. To account for aggregate changes in return volatility, we standardize the local measure by the nationwide standard deviation. Our return risk measure thus indicates whether the local dispersion of returns is relatively high or low compared to the nationwide dispersion. Simultaneously, the standardization has the advantage to account for the systematic change in holding periods over the sample period that arises with our repeat sales structure. As we estimate return volatility directly from the underlying data, our procedure does not rely on prior index aggregation or assumptions on the autoregressive structure as, e.g., in GARCH models (e.g., Han, 2013). Formally, our measure for local return risk, Vol_{ct}, is derived as

$$Vol_{ct} = \frac{\sigma_{ct}}{\sigma_t}, \tag{3.2.2}$$

in which σ_{ct} is the standard deviation of annualized returns realized in county c at time t, and the nationwide standard deviation of returns at time t is defined as σ_t.

To achieve a reasonable tradeoff between estimation precision and cross-sectional as well as temporal variation, both price indices and local return risk are measured at the county-quarter level. Otherwise, the usage of more fine-grained frequencies

[3]When estimating price indices, we face the problem of missing observations in in-between periods for some counties. To make use of the maximum amount of observations possible, we solve this issue by omitting the respective intervals in the estimation process and assigning missing values to these periods in the resulting time-series. To further increase the precision of our estimates, we use data on transactions ranging back to 1980 wherever possible, as well as the transactions with unknown identity, such that indices are estimated before dismissing these observations in the cleaning process.

or areas could lead to the loss of more rural areas in the sample, which would bias our results towards large MSAs. Once local return risk and house price returns are derived and matched to purchase and sale date of each return, respectively, we remove outliers by trimming both variables at the 0.5 and 99.5 percent level, respectively.

3.2.4. Control variables

We investigate channels other than local risk and expected returns by making use of additional data sources and deriving additional variables with our housing dataset. To account for conditions at both purchase and sale date, we match each of these additional variables to both dates of each transaction, respectively. We further analyze changes in these variables over the respective holding periods to investigate how investors' capital gains move along with each of these factors, respectively.

Based on our housing dataset, we calculate two variables, again at the county-quarter level. To have a proxy of liquidity for the local market, we measure turnover as the sum of all individual transactions in a given period. Local credit conditions are measured with the average primary loan-to-value ratio (LTV), derived as the average mortgage amount divided by the average sales price, excluding individual ratios larger than one in the mean calculation, similar to Fuster and Vickery (2015). Intuitively, the LTV measures the importance of local credit supply for the transactions in a given period. Again, to increase the precision of our estimates, we use the dataset before dismissing transactions with unknown identity when deriving both variables.

Local economic fundamentals are analyzed using panel data on the county-year level. We use information on median income, the unemployment rate, and population. Additionally, we use data on land areas to calculate population densities. We obtain population and land area (measured in square miles) from the US Census Bureau. Median income is provided by the Small Area Income and Poverty Estimates Release of the US Census, obtained from the FRED database of the Federal Reserve Bank of St. Louis. The unemployment rate is provided by the US Bureau of Labor Statistics, again obtained from FRED.

In order to analyze macro-financial (risk) factors, we use monthly time series. To investigate investors' sensitivity with respect to the aggregate housing market, we use monthly data of the S&P Case-Shiller US National Home Price Index (not seasonally adjusted) and the average 30-year mortgage interest rate which we both obtain from

FRED. We further attain the Fama-French stock market factor as well as data on the risk-free rate from Kenneth French's webpage.[4]

We further control for several factors on the individual level. First, when investigating capital gains, larger renovations can confound the true performance, as observed gains might be driven upwards from additional investments rather than general price appreciation. We therefore construct a dummy that indicates existing construction related documents (e.g., a construction mortgage) associated with the purchase. When deriving this variable, we do not include new constructions. Instead, to exclude that land speculation (Nathanson and Zwick, 2018) is driving our results, we generate an additional dummy indicating that our data provider classifies the initial purchase as new construction. We further control for unobserved heterogeneity related to the holding period. For instance, shorter holding periods are documented to be generally associated with speculation thus possibly implying higher returns (Bayer et al., 2020). To account for a potentially non-linear relationship, we generate dummies based on holding period sample deciles that we include in our regressions.

3.2.5. Descriptive statistics

Panel A of Table 3.2.1 provides summary statistics for the full sample, and Panel B summarizes excess capital gains by investor group. Panel A shows that the sample average of annualized excess capital gains is 4.2%, reflecting the general increase in market prices over the sample period. This finding is further reflected in the comparison of nominal purchase and sale prices, for which the average increases from about USD 250,000 to USD 310,000. A comparison with the median indicates the typical right-skewness of the house price distribution. Almost one percent of purchases is associated with larger observed construction activity. The number of observations of our local expected return (momentum) and risk measures is reduced compared to our overall sample for two reasons. First, we account for extreme values by trimming observations out of the 0.5 to 99.5 percent interval. Second, as we use lagged versions of these variables in our analyses to cope with potential simultaneity issues, we report summary statistics of the time-shifted variables for the sake of consistency.

Panel B summarizes our performance measure, annualized excess capital gains, by investor group. With almost 74%, owner-occupiers make up the largest share of market participants in our sample, while short-and long-term investors are responsible for about 2% of trades in our sample, respectively. With about 3.5%, owner-occupiers

[4] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/f-f_factors.html, last retrieved on February 3, 2021.

perform the worst among the groups. Yet, the corresponding standard deviation indicates that, as a group, owner-occupiers faced the lowest variation in excess capital gains among the investors. The risk-return structure holds for all subgroups: A higher group-average is associated with a higher standard deviation. The highest gains (15.8%), but also the highest dispersion (24.2%) is realized by ST IIs. Although a large portion of this high performance should arise from activities associated with shorter holding periods such as renovations, this finding already suggests that risk-taking, reflected in the varying dispersion of returns might contribute to the different performances across investors.

Figure 3.2.1 illustrates that the differences in annualized excess capital gains do not only hold cross-sectionally, but are also persistent over time. In line with Panel B of Table 3.2.1, owner-occupiers, depicted by the black solid line, persistently realize the lowest capital gains since 1996. In contrast, short-term institutional investors perform best since 1998. Private investors performed better than LT II during the last boom, but the pattern reversed when the bubble bursted, again suggesting heterogeneity in risk-taking across groups. The persistence of the performance differences over time further underlines that the investors generally follow similar strategies over time. Having presented descriptive evidence of our data, we present more formal analyses on the documented performance differences in the following section, with particular focus on local return risk and momentum.

3.3. Empirical results

In this section, we present our results on the differences in annualized excess capital gains across investor groups. In Section 3.3.1, we first analyze whether the documented outperformance of investor groups persists after including various sets of control variables. We then present our results on the exposure of capital gains with respect to local risk and expected returns (momentum) in Section 3.3.2. As we document that local return risk is a sizable factor in explaining disparities across investors' capital gains, we further analyze its role in the purchase decision for market participants in Section 3.3.3, and its relationship to local fundamentals in Section 3.3.4.

3.3.1. Systematic differences in capital gains

While the last section presented descriptive evidence on the systematic differences across investor groups, we provide more formal analyses in this section. Table 3.3.1 analyzes the average annualized capital gains of private investors, ST IIs, and LT

Table 3.2.1
Summary statistics

Panel A: Full sample				
	Observations	Mean	Std.	Median
Ann. excess capital gain	21,178,869	0.042	0.121	0.019
Purchase price	21,178,869	253,182	486,727	179,037
Sale price	21,178,869	313,509	677,930	220,000
Construction	21,178,869	0.008	0.088	0
New construction	21,178,869	0.158	0.365	0
Holding period	21,178,869	5.750	4.005	4.712
County unemployment	21,170,934	5.341	2.144	4.900
County income	21,156,784	49,813	12,767	47,215
County pop. density	21,178,517	1,679	4,890	685.966
County primary LTV	21,178,545	0.638	0.168	0.682
County turnover	21,178,869	4,021	5,169	2,019
County momentum	20,825,036	0.017	0.039	0.019
County return risk	19,760,762	0.950	0.264	0.931
Panel B: Annualized excess capital gains by investor				
	Observations (%)	Mean	Std.	Median
Owner-occupiers	15,660,837 (73.9)	0.035	0.105	0.017
Private investors	4,590,293 (21.7)	0.055	0.144	0.023
ST II	435,175 (2.1)	0.158	0.242	0.120
LT II	492,564 (2.3)	0.052	0.139	0.024

Panel A of this table shows summary statistics for the full sample. “Ann. excess capital gain” is the annualized capital gain of a repeat sale minus the return of the risk-free rate over the respective holding period. “Purchase (sale) price” is the nominal transaction price corresponding to the first (second) sale of the respective repeat sales pair. “Construction” is a dummy that indicates that the transaction was associated with documents indicating construction such as a construction mortgage, excluding new constructions. “New construction” is a dummy that indicates that the purchase transaction was the one of a newly constructed home. “Holding period” is the period between purchase and sale date given in years. The “County” variables indicate the conditions of the local market related to each transactions’ purchase period. “County unemployment” is the county-year level unemployment rate in percentage points, “County income” is the county-year level median income,“County pop. density” is the population density on the county-year level, derived as number of inhabitants per square-mile of land, “County primary LTV” is the county-quarter level average of the mortgage amount divided by the average sales price, and “County turnover” is the sum of all transactions in a given quarter in a given county, including observations with unknown identity. “County momentum” is the one-quarter lag of our estimated county-level index returns. “County return risk” is the one-quarter lag of the local annualized return dispersion relative to nationwide dispersion. Panel B shows mean, standard deviation, and median of annualized excess capital gains separated by investor type, as indicated by the row names. “Owner-occupiers” are private individuals that live in the underlying home. “Private investors” are private individuals who do not live in the home purchased. “ST II” indicates the property was purchased by an institutional investor (II) with a short-term (ST) holding period, i.e., three years or less, “LT II” are institutional holders with a long-term (LT) holding period, i.e., more than three years. The numbers in parentheses next to the absolute amount of observations per investor group indicate the relative frequency of each group in the sample.

Figure 3.2.1
Capital gains in the US residential housing market

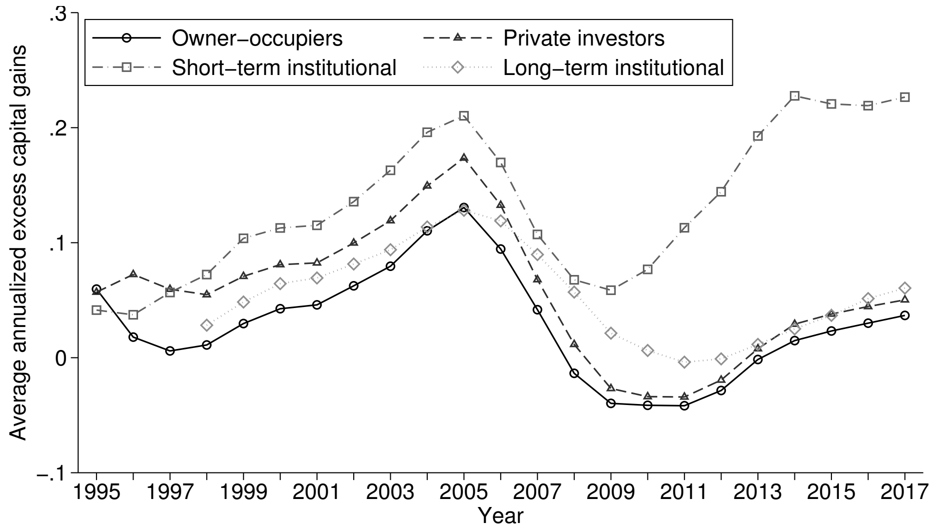

This figure plots average annualized excess capital gains realized by repeated trades of individual homes. The graph is based on more than 21 million repeat sales of US residential homes traded between the years 1995 and 2017. Each datapoint shows the average of annualized excess capital gains realized at the given year by a particular group of sellers. Investor identity is assigned according to information associated with each individual trade. The solid line (circled datapoints) shows the development for owner-occupiers, i.e. households that live in the underlying home. The dashed line (triangle datapoints) depicts capital gains of private investors that are defined as private individuals who do not live in the traded home. The dashed-dotted line (cubed datapoints) shows the return development for short-term institutional investors (ST IIs), i.e., trades made by institutional investors with a holding period of three or less years. The dotted line (diamond datapoints) shows capital gain dynamics for long-term institutional investors (LT IIs), i.e., institutional sellers with a holding period of more than three years.

IIs relative to owner-occupiers in a regression framework. We regress annualized excess capital gains on dummies indicating the respective investor type, as well as different sets of control variables. As for all performance analyses in the paper, we use owner-occupiers as base category. The results presented in column (1) indicates that on average, ST IIs outperform owner-occupiers by more than 12 percentage points per year, while private investors and LT IIs perform 2 and 1.6 percentage points better, respectively.

As a large portion of the outperformance documented in column (1) might stem from speculative activity such as larger renovations, we add dummy variables related to observed construction activities, newly built homes, or decile bins of different holding periods for the regressions shown in column (2). We find a sizable reduction in capital gains for ST IIs relative to owner-occupiers, but also for private investors, which might be explained by short-term value creation due to renovations or real estate developers selling new built properties in the market (e.g., Gao et al., 2020; Nathanson and Zwick, 2018). By controlling for the holding period, we also filter out unobserved variation related to price increases which are driven by general short-term activity such as house flipping strategies (e.g., Bayer et al., 2020). Furthermore, we show that all investor groups consistently outperform owner-occupiers when comparing them in the same holding period bin.

Columns (3) and (4) include county-level fixed effects as well as dummies to capture the timing of purchase and sale at the annual and quarterly level, respectively, to further analyze the impact of aggregate price dynamics on investors' performance. The disparities in realized capital gains remain positive and statistically significant. This indicates that choice of preferable locations or buying and selling at particularly beneficial points in time contribute only little to the remaining outperformance of investor groups. Consequently, we document systematic performance differences in capital gains between investor groups not only across counties, but also within local markets and time periods.

As indicated in column (5), interacting location with both purchase and sale quarters to county-quarter fixed effects can explain up to 46% of the variation in annualized excess capital gains. Despite this sizable explanatory power, the differences across investors persist. The pattern even remains after the inclusion of zip-code-quarter fixed effects, as shown in column (6). The results of Table 3.3.1 thus indicate that the outperformance of investor groups relative to owner-occupiers is largely driven by investor-specific factors. We therefore continue our analysis by investigating investor-specific exposure to local and macro-financial variables, respectively.

In Table 3.3.2, we analyze the average outperformance of each investor after accounting for the investor-specific sensitivity to several variables. We include the S&P Case-Shiller US National Home Price Index to analyze the impact of the sizable housing market momentum on investor performance (e.g., Case and Shiller, 1989). Additionally, we investigate returns of the stock market and the 30-year mortgage rate to investigate whether changes in the aggregate economy impact investors differently. We also control for local market fundamentals, such as income, unemployment, the population density, and the turnover as a proxy for liquidity in the market. We also analyze investors' sensitivity to the average LTV in the local market to capture the potential price impact of a generally less restrictive credit supply, which might be translated in lower down payment requirements (Anenberg et al., 2019; Mian and Sufi, 2019).

Table 3.3.1
Average differences in capital gains

	(1)	(2)	(3)	(4)	(5)	(6)
Private inv.	0.020***	0.014***	0.018***	0.019***	0.019***	0.016***
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
ST II	0.123***	0.048***	0.060***	0.060***	0.067***	0.055***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
LT II	0.016***	0.030***	0.032***	0.032***	0.029***	0.027***
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Construction	-	X	X	X	X	X
New constr.	-	X	X	X	X	X
Holding pd.	-	X	X	X	X	X
Local FE	-	-	county	county	-	-
Purchase FE	-	-	year	quarter	-	-
Sale FE	-	-	year	quarter	-	-
L×T FE	-	-	-	-	cty-qtr	zip-qtr
Adj. R-sq.	0.024	0.135	0.324	0.328	0.467	0.519
Observations	21,178,869	21,178,869	21,178,869	21,178,869	21,178,869	21,087,407

This table shows regression results with annualized excess capital gains as dependent variable. The displayed coefficients are estimates for group-specific dummies. The reference category are owner-occupiers, i.e., private individuals that live in the home purchased during the holding period. "Private inv.", "ST II", "LT II" are dummy coefficients indicating that the corresponding trade was made by a private investor, short-term institutional investor or long-term institutional investor, respectively. Private investors are defined as individuals that do not live in the underlying homes. A transaction is classified as "short-term" if the corresponding holding period is three years or less, and "long-term" otherwise. "Construction" indicates that the purchase transaction has associated documents indicating construction activity, excluding new constructions. "New construction" indicates that the purchase is identified as a transaction of a newly built home. "Holding period" stands for dummies based on holding period sample deciles. "Local FE" indicates local fixed effect dummies on the level given in the respective column. "Purchase FE" and "Sale FE" indicate that fixed effect dummies corresponding to the respective purchase or sale date on yearly or quarterly level are included in the model, respectively. "L×T" indicate that interactions of local and time fixed effect dummies are used, respectively. Here, "cty-qtr" denotes county-quarter, and "zip-qtr" denotes zip-code quarter fixed effects, respectively. For all local-time fixed effects, interactions of both purchase and sale quarter are used. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

In each regression reported in Panel A of Table 3.3.2, we interact the returns of the variable given in the first column with each buyer identity. As each repeat sale is exposed to each factor at the purchase as well as the sale date, respectively, we include realizations of the given variable at both points in time. For all of the models analyzed in Panel A, the estimated coefficients remain positive at the same order of magnitude and highly significant. This indicates that neither of the tested variables can help explain the disparities in investor performance. Our results therefore indicate that the investor groups do not systematically invest in heterogeneous market segments where they are exposed to different (e.g., increasing vs. decreasing) market characteristics for which they are rewarded with a higher return compensation.

For the results shown in Panel B, we use the same variables as in Panel A, but transform them to annualized changes over the respective holding periods instead, similar to the holding period factor models proposed by Peng (2016) and Peng and Zhang (2019). Again, the outperformance of investors remains at the same order of magnitude and highly significant for all specifications tested. The results of Panels A and B of Table 3.3.2 thus indicate that the capital gain disparities are highly systematic and are likely to be unrelated to local economic fundamentals as well as macro-financial factors. Yet, our analysis so far neglected the potential impact of local risk and momentum, respectively, which is analyzed in more detail in the following sections.

3.3.2. Results on local risk and return

In this section, we test whether the performance of different investor types can be explained by local return risk and momentum. As each group should have a distinct investment strategy, the relationship of investor returns with respect to local risk and return might be crucial in understanding these differences in appreciation returns. Our goal is to evaluate the investor-specific performance along two dimensions. First, we estimate the group-specific exposure to both local return risk and momentum, respectively, to infer whether the risk-return relationship is indeed heterogeneous. Second, we study whether this heterogeneity can help explain the observed differences in capital gains.

To answer both questions simultaneously, we proceed similar to the portfolio sorts approach proposed by Hoechle et al. (2020). The methodology has the advantage that first, exposures and relative performances can be estimated in a single regression, and second, that the specification allows for the inclusion of fixed effects and other controls, in contrast to traditional portfolio sorts approaches. In the context of housing

Table 3.3.2
Relative performance after investor-specific exposure

Panel A: Return at purchase and sale date								
	Private inv.		ST II		LT II		Adj. R-sq.	Observations
National index	0.013***	(0.001)	0.045***	(0.003)	0.039***	(0.002)	0.188	21,160,538
Stock market	0.013***	(0.001)	0.048***	(0.003)	0.031***	(0.001)	0.162	21,160,538
30-year mortgage	0.013***	(0.001)	0.049***	(0.003)	0.031***	(0.001)	0.163	21,160,538
County income	0.014***	(0.001)	0.048***	(0.003)	0.036***	(0.001)	0.200	20,492,659
County unemployment	0.013***	(0.001)	0.049***	(0.003)	0.035***	(0.001)	0.188	20,486,270
County density	0.011***	(0.002)	0.070***	(0.004)	0.032***	(0.002)	0.168	20,496,123
County turnover	0.013***	(0.001)	0.050***	(0.003)	0.030***	(0.001)	0.162	21,156,300
County LTV	0.012***	(0.001)	0.052***	(0.003)	0.028***	(0.001)	0.166	20,348,816
Panel B: Annualized changes over the holding period								
	Private inv.		ST II		LT II		Adj. R-sq.	Observations
National index	0.016***	(0.001)	0.076***	(0.004)	0.031***	(0.001)	0.332	20,967,087
Stock market	0.009***	(0.002)	0.049***	(0.003)	0.034***	(0.002)	0.171	20,967,094
30-year mortgage	0.016***	(0.001)	0.045***	(0.003)	0.046***	(0.002)	0.175	20,967,139
County income	0.007***	(0.002)	0.065***	(0.003)	0.015***	(0.003)	0.187	21,156,784
County unemployment	0.013***	(0.001)	0.051***	(0.003)	0.031***	(0.001)	0.230	21,170,761
County density	0.010***	(0.002)	0.071***	(0.004)	0.027***	(0.002)	0.173	21,178,517
County turnover	0.015***	(0.001)	0.056***	(0.003)	0.027***	(0.001)	0.193	21,178,869
County LTV	0.013***	(0.001)	0.052***	(0.003)	0.027***	(0.001)	0.172	20,444,399

This table shows regression results with annualized excess capital gains as dependent variable. The displayed coefficients are estimates for group-specific dummies. The reference category are owner-occupiers. For each regression, we interact the variable given in the first column with dummies for each buyer type, respectively. In each regression, the realization of the underlying variable at purchase and sale date is used, respectively. “National index” is the monthly S&P Case-Shiller US National Home Price Index, “Stock market” is the monthly Fama-French stock market factor, “30-year mortgage” is the monthly average 30-year mortgage interest rate, “County income” is the county’s median income in a given year, “County unemployment” is the county’s unemployment rate in the given year, “County density” is the county’s population per square mile in a given year. “County turnover” is the quarterly number of transactions in a given period and “County LTV” is the county’s quarterly average mortgage amount divided by the average sales price (see Section 3.2.4). “Private inv.”, “ST II”, “LT II” are dummy coefficients indicating that the corresponding trade was made by a private investor, short-term institutional investor or long-term institutional investor, respectively. Private investors are defined as individuals that do not live in the underlying homes. A transaction is classified as “short-term” if the corresponding holding period is three years or less, and “long-term” otherwise. For Panel A, we interact each group with the given variable at purchase and sale period, respectively. For Panel B, annualized changes over the given holding period are interacted with group dummies. The annualized variables are trimmed at the 0.5 and 99.5 percent level. All regressions include the same set of control variables: Dummies for construction, new construction, and holding period sample deciles, as well as county-level fixed effects. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

markets, the latter reason crucially allows to control for unobserved heterogeneity across different areas. Consequently, we explain the average annualized excess capital gains realized between purchase date t_1 and sale date t_2 by an investor of group k for property i that is located in county c, $r_{ict_1t_2}$, with regressions of the form

$$r_{ict_1t_2} = (\boldsymbol{\alpha_i} \otimes \boldsymbol{X_{ct-1}})\,\boldsymbol{\theta} + \nu_c + Z_i\beta + \epsilon_{ict_1t_2}. \tag{3.3.1}$$

Vector $\boldsymbol{\alpha_i} = \begin{bmatrix} 1 & \alpha_i^{(PI)} & \alpha_i^{(ST)} & \alpha_i^{(LT)} \end{bmatrix}$ contains dummy variables of the investor type identified for repeat sale i, $\boldsymbol{X_{ct}} = \begin{bmatrix} 1 & \boldsymbol{Q_{ct_1}} & \boldsymbol{Q_{ct_2}} \end{bmatrix}$ includes a set of regressors, with $\boldsymbol{Q_{ct_1}}$ containing variables related to the purchase date of i, and $\boldsymbol{Q_{ct_2}}$ to the respective sale date. In our baseline specifications, the vectors of factors, $\boldsymbol{Q}$, contain either local return risk, Vol_{ct}, or local momentum, $\boldsymbol{Ret_{ct}}$, respectively. The explanatory variables are included with a one-period lag to avoid simultaneity concerns and a look-ahead bias that might arise, as we use our LHS variables to measure each variable.[5]

Vector $\boldsymbol{\theta}$ contains the regression coefficients. The estimates for the interactions of $\alpha_i^{(PI)}$, $\alpha_i^{(ST)}$, and $\alpha_i^{(LT)}$ with the "intercept" in $\boldsymbol{X_{ct}}$ measure the factor-adjusted performance of each investor relative to the reference group of owner-occupiers ("relative alpha"). An insignificant coefficient estimate associated with the corresponding investor type therefore indicates that the outperformance relative to the reference group can be explained by the set of (local) factors included in the regression. Parameter ν_c captures unobserved county heterogeneity, such as land supply restrictions (Saiz, 2010), and ensures the comparability of the performance between investor types within a local market. Matrix Z_i includes additional confounding variables associated with transaction i such as (new) construction and holding period decile dummies. In the following, we separately test whether local return risk and momentum can contribute to explain the performance differences.

In Table 3.3.3, we first study investor exposure to local market momentum as a predictor for capital gains. Panel A shows estimates for investor-specific exposure at purchase and sale date, respectively, conditional on the full set of controls. The coefficients are positive and statistically significant for all investor groups at the time of purchase as well as the time of resale. The findings are consistent with, e.g., Landvoigt (2017), who shows that households' growth expectations are based on past market returns. Yet, comparing coefficient size across investor types, the loadings are indeed heterogeneous. For instance, ST IIs load strongest on momentum, while the long-term strategy of LT IIs appears to depend much less on current house price movements.

[5] In Section 3.4, we proceed similar to Han (2013) and instrument contemporaneous return risk with up to two of its lags. Our results are robust to these alternative specifications.

Panel B of Table 3.3.3 analyzes whether this heterogeneity in exposure is translated into the documented capital gain disparities. The panel shows results on investors' performance relative to owner-occupiers after the inclusion of local momentum exposure and multiple controls. For ease of comparison, we show the baseline results conditional on control variables in column (1). The results shown in columns (2)-(4) indicate that, despite sizable heterogeneity in momentum exposure, the outperformance of investor groups remains at the same order of magnitude, whether including momentum at the purchase quarter, sale quarter, or both, respectively. This indicates that although trading on momentum can generally be beneficial, it appears to contribute only little to the outperformance of investors.

While Table 3.3.3 focuses on local return momentum on the returns of investor types, the implications of local return risk for investor performance have been neglected so far. In Table 3.3.4, we therefore investigate the exposure of annualized excess capital gains to past local return risk. Again, we regress annualized excess capital gains on the investor type dummies and the corresponding interaction terms with one-quarter lagged local return risk. Panel A shows investors' estimated exposure to local return risk relative to the purchase and sale quarter, respectively. The purchase loadings are similar and positive for owner-occupiers, private investors, and LT IIs. The positive coefficients suggest a general compensation for risk-taking when purchasing the home. Similarly, the sale date exposure is negative, suggesting that selling during riskier times is generally associated with a discount. This reflects the intuition of the positive purchase date coefficients, as the new buyer again wants to be compensated for engaging in higher risk. The coefficient of sale date risk-exposure of owner-occupiers is of higher magnitude than for private investors and LT IIs, suggesting that the risk-discount is higher for individuals with a primary motive of consumption. The exposure of ST IIs is remarkably different to the other investors, with a large, positive exposure at the sale date. This indicates that the trading strategy of ST IIs is paying off most in highly volatile markets, potentially reflecting engagement speculative activity.

In contrast to any other channel tested in this work so far, the exposure to return risk can explain a sizable share in the outperformance of investors relative to owner-occupiers, as illustrated in Panel B of Table 3.3.4. Again, the results after controls are displayed in column (1) for ease of comparability. As shown in column (3), the exposure to local return risk can fully explain the remaining outperformance of investors. Controlling additionally for purchase level activity in column (4), thus accounting for the riskiness existing prior to the respective investor's market entry,

Table 3.3.3
Investor performance and local momentum

Panel A: Local momentum exposure by investor type

	Owner-occupier	Private investor	ST II	LT II
Purchase	0.110***	0.082***	0.409***	0.029**
	(0.012)	(0.012)	(0.033)	(0.011)
Sale	0.651***	0.623***	0.807***	0.267***
	(0.044)	(0.040)	(0.048)	(0.018)
Construction		X		
New construction		X		
Holding period		X		
County FE		X		
Adj. R-sq.		0.206		
Observations		20,656,495		

Panel B: Relative investor alphas under local momentum exposure

	(1)	(2)	(3)	(4)
Private inv.	0.013***	0.013***	0.015***	0.015***
	(0.001)	(0.001)	(0.001)	(0.001)
ST II	0.049***	0.044***	0.052***	0.048***
	(0.003)	(0.003)	(0.003)	(0.003)
LT II	0.031***	0.031***	0.037***	0.037***
	(0.001)	(0.001)	(0.001)	(0.001)
Local momentum	-	purchase	sale	purchase/sale
Construction	X	X	X	X
New construction	X	X	X	X
Holding period	X	X	X	X
County FE	X	X	X	X
Adj. R-sq.	0.161	0.164	0.203	0.206
Observations	21,178,869	20,825,036	20,868,406	20,656,495

This table shows results for regressions with annualized excess capital gains of individual repeat sales as dependent variable. Panel A shows estimates for the exposure of investors with respect to local momentum, i.e., the one-quarter lag of the county-level index return corresponding to the purchase and sale date of each repeat sale, respectively. The regression includes investor dummies. Panel B reports the average capital gains of investors relative to owner-occupiers, after including different sets of one-quarter lags of "Local momentum", i.e., the quarterly return of the county-level index. Each set of lags is included either relative to the purchase quarter, the sale quarter, or both, respectively, as indicated in the row "Local momentum". "Private inv.", "ST II", "LT II" are dummy coefficients indicating that the corresponding trade was made by a private investor, short-term institutional investor or long-term institutional investor, respectively. Private investors are defined as individuals that do not live in the underlying homes. A transaction is classified as "short-term" if the corresponding holding period is three years or less, and "long-term" otherwise. "Construction" indicates that the purchase transaction has associated documents indicating construction activity, excluding new constructions. "New construction" indicates that the purchase is identified as a transaction of a newly built home. "Holding period" stands for dummies based on holding period sample deciles. "County FE" stands for county-level fixed effect dummies. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

leaves only a slight and borderline significant outperformance of LT IIs. While this outperformance appears the be unstable regarding our robustness checks, it might well indicate that the long-term strategy pays off due to lower risks taken relative to their performance.

Taking the results of Tables 3.3.3 and 3.3.4 together, we find that, although loadings on local momentum are strong and heterogeneous across groups, they cannot explain the return discrepancies. Accounting for local risk, in contrast, leads to a similar performance of private investors and ST IIs compared to owner-occupiers. Conditional on local risk, only LT IIs still outperform owner-occupiers, although a larger amount of their performance can be explained by local return risk as well. This pattern holds when including both, local momentum and risk in a single regression, as shown in Table 3.A.2 in the Appendix.

3.3.3. The purchase decision

In the last section, we illustrated that the exposure to local return risk can help explain differences in returns across buyer groups. The exposure of investors is likely to be related to the selection process underlying the purchase decision, e.g., choosing either markets associated with relatively high or low return risk. This decision should be associated with the investor's underlying strategy, and is potentially impactful. For instance, Table 3.3.4 documents that purchase and sale date related loadings differ in size, suggesting that even constant return risk might impact the buyer's performance. In this section, we thus investigate potential determinants for market selection across investor groups, in particular local momentum and return risk.

We do so by forming indicator variables that equal one if a transaction was made by a particular group of investors and zero otherwise. Afterwards, we run probit regressions to investigate the relationship of investor group activity with respect to local variables such as income or unemployment at the time of purchase. Accordingly, we model the conditional probability that buyer i belongs to group k to buy a home in county c at time $T = t_1$ as

$$P(\textit{Buyer } i \textit{ belongs to group } k | X_{ct_1}, T = t_1) = \Phi(X_{ct_1}\beta + \nu_{t_1}), \tag{3.3.2}$$

in which $\Phi()$ is the normal cumulative probability distribution, X a set of explanatory variables, ν_{t_1} purchase quarter fixed effect dummies, β a vector of parameters, and the set of groups is $k \in$ {owner-occupier, private investor, ST II, LT II}. For each group of market participants, we run a separate regression. In the regressors X, we include contemporaneous and lagged versions of local momentum to determine who is

Table 3.3.4
Investor performance and local return risk

Panel A: Local risk exposure by investor type

	Owner-occupier	Private investor	ST II	LT II
Purchase	0.028***	0.026***	-0.006	0.028***
	(0.008)	(0.007)	(0.006)	(0.005)
Sale	-0.060***	-0.044***	0.041**	-0.044***
	(0.016)	(0.012)	(0.012)	(0.012)
Construction		X		
New construction		X		
Holding period		X		
County FE		X		
Adj. R-sq.		0.175		
Observations		19,626,314		

Panel B: Relative investor alphas under local return risk exposure

	(1)	(2)	(3)	(4)
Private inv.	0.013***	0.011**	-0.001	0.001
	(0.001)	(0.004)	(0.006)	(0.005)
ST II	0.049***	0.049***	-0.017	-0.011
	(0.003)	(0.008)	(0.016)	(0.015)
LT II	0.031***	0.023***	0.011	0.014*
	(0.001)	(0.005)	(0.006)	(0.006)
Local return risk	-	purchase	sale	purchase/sale
Construction	X	X	X	X
New construction	X	X	X	X
Holding period	X	X	X	X
County FE	X	X	X	X
Adj. R-sq.	0.161	0.170	0.166	0.175
Observations	21,178,869	19,760,762	20,931,159	19,626,314

This table shows results for regressions with annualized excess capital gains of individual repeat sales as dependent variable. Panel A shows estimates for the exposure of investors with respect to local return risk, i.e., the one-quarter lag of the county-level dispersion of annualized capital gains relative to nationwide dispersion in the same quarter. The regression includes investor dummies. Panel B reports the average capital gains of investors relative to owner-occupiers, after including different sets of one-quarter lags of "Local momentum", i.e., the quarterly return of the county-level index. Each set of lags is included either relative to the purchase quarter, the sale quarter, or both, respectively, as indicated in the row "Local return risk". "Private inv.", "ST II", "LT II" are dummy coefficients indicating that the corresponding trade was made by a private investor, short-term institutional investor or long-term institutional investor, respectively. Private investors are defined as individuals that do not live in the underlying homes. A transaction is classified as "short-term" if the corresponding holding period is three years or less, and "long-term" otherwise. "Construction" indicates that the purchase transaction has associated documents indicating construction activity, excluding new constructions. "New construction" indicates that the purchase is identified as a transaction of a newly built home. "Holding period" stands for dummies based on holding period sample deciles. "County FE" stands for county-level fixed effect dummies. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

active at which state of the local housing market cycle, and potentially trading on autocorrelation. Similarly, we investigate which investor groups buy in higher-risk locations, and whether past information on risk matters was well.

Table 3.3.5 depicts results of four probit regressions that model the probability that a given investor group made a particular trade. Columns (1)-(4) separately model the probability that a purchase was made by an owner-occupier, private investor, ST II, or LT II. All regressions include fixed effect dummies on the quarter-year level to filter out macroeconomic determinants. If not indicated otherwise, the regressors refer to the contemporaneous period of purchase. All variables are measured at the county level.

The first column of Table 3.3.5 investigates the activity of owner-occupiers relative to other investor groups. The regression coefficients show that owner-occupiers tend to invest in counties with lower return risk, as indicated by the negative and significant coefficient estimate for contemporaneous return risk. Even conditional on contemporaneous volatility, the first lag is significant, suggesting that higher experienced volatility tends to reduce owner-occupier activity. The coefficients for the index returns are negative as well, suggesting that owner-occupiers tend to buy in local markets with falling prices relative to the other investor groups. Furthermore, purchase activity tends to be higher when unemployment is lower. This could reflect that owner-occupiers like to move closer to their workplace, or that homes become more affordable to a wider range of individuals when employment is high. Lastly, a predictor for owner-occupier activity is the average primary loan-to-value ratio which could indicate that owner-occupiers profit from easy access to credit, or are willing to take out larger loans to buy a home in the preferred location.

As owner-occupiers make up about 75% of the total observations, the results for the investors have to be largely interpreted relative to owner-occupiers. All three investor groups, private investors, ST IIs, and LT IIs, tend to invest in counties subject to comparably high return risk. This could reflect a general rent-seeking pattern, in which investors aim for higher expected capital gains by engaging in riskier trades. In contrast, the importance of local momentum varies across investors. While private investors are mostly active in booming markets, as indicated by the large and highly significant coefficients, ST IIs tend to invest in falling markets, which could be an explanation for the high returns of ST IIs in the post-bust period of recovery. The activity of LT IIs is not significantly related to momentum, suggesting that current price movements lose in relevance for long-term strategies of professionals. This result is intuitive, as short-term deviations from the general house price trend should not matter if homes are held for longer periods. This finding indicates that LT IIs are indeed planning to hold their properties for longer periods when buying the property, underpinning our cut-off choice of three years. Other than that, institutional investors choose similar locations, characterized by higher population density, lower income, and lower unemployment. This could suggest that investors tend to buy homes in gentrifying neighborhoods.

In summary, all investors, especially private investors and ST IIs, tend to invest in markets with relatively high return risk. When it comes to local momentum, however, strategies appear to be different. Private investors seem to be active mostly in booming markets, while ST IIs tend to buy in falling ones. In contrast to all other groups, LT IIs do not vary their activity with local price trends.

3.3.4. Where is local return risk high?

While the previous sections documented that local return risk can help explain a significant share of the differences in capital gains across investor groups, the drivers of local return risk remain largely unknown. To analyze which determinants are associated with higher risk, we run panel regressions on the county-quarter level. In addition to the variables investigated in the paper, we generate a new one to capture how the activity of ST IIs is related to return risk, conditional on local fundamentals. To do so, we calculate the share of sales made by ST IIs in a given market at a given period. To correct for the systematically changing holding period over time due to the repeat sales structure, we standardize this share by the nationwide share of short-term sales.

Table 3.3.6 shows results for three panel regressions that explain our local return risk measure. The model in column (1) shows that higher housing returns are associated with higher return risk, which holds when including county fixed effects as well as year fixed effects. These results support the finding from Table 3.3.4 that buyers are compensated for taking higher risk associated with the purchase. The relationship of income and return risk is generally negative. One explanation could be that lower income households are more sensitive to credit restrictions, which led to higher appreciation rates for low priced homes during last boom's credit expansions, as shown in Landvoigt et al. (2015).

The share of ST IIs in a market is positively related to local return risk as well. This reflects the finding that ST IIs are more active in markets with higher risk from Table 3.3.5. Additionally, we find that this result is robust to the inclusion of year fixed effects as well as county fixed effects, as shown in columns (2) and (3). This finding

Table 3.3.5
Local variables and the purchase decision of investor groups

	OO	PI	ST II	LT II
County return risk	-0.280***	0.249***	0.202***	0.123***
	(0.057)	(0.061)	(0.021)	(0.025)
County return risk (t-1)	-0.156**	0.138*	0.124***	0.071***
	(0.053)	(0.057)	(0.016)	(0.018)
County momentum	-0.856**	1.004**	-0.360***	-0.158
	(0.296)	(0.315)	(0.090)	(0.103)
County momentum (t-1)	-0.941***	1.038***	-0.213*	-0.007
	(0.271)	(0.288)	(0.085)	(0.099)
Log county pop. density	0.004	-0.021	0.062***	0.036***
	(0.026)	(0.027)	(0.007)	(0.010)
Log county income	0.157	-0.080	-0.250***	-0.331***
	(0.148)	(0.162)	(0.055)	(0.056)
County unemployment	-0.043**	0.050**	-0.015**	-0.009
	(0.015)	(0.016)	(0.005)	(0.006)
Log county turnover	0.026	-0.024	-0.010	-0.013
	(0.024)	(0.025)	(0.010)	(0.013)
County primary LTV	1.227***	-1.073***	-0.582***	-0.759***
	(0.148)	(0.135)	(0.069)	(0.083)
Purchase quarter FE	X	X	X	X
Pseudo R-sq.	0.042	0.035	0.065	0.029
Observations	19,428,196	19,428,196	19,428,196	19,094,592

This table shows results of four probit regressions that explain whether a particular investment was made by one of our four investor groups, owner-occupiers (OOs), private investors (PIs), short-term institutional (ST II), or long-term institutional investors (LT IIs), respectively. Owner-occupants are private individuals that live in the home underlying the transaction. Private investors are defined as individuals that do not live in the respective homes. A transaction is classified as "short-term" if the corresponding holding period is three years or less, and "long-term" otherwise. "County return risk" indicates the county-quarter level return risk at the purchase quarter of the underlying transaction. Return risk is measured as local relative to nationwide dispersion of annualized capital gains within the same quarter. "County momentum" indicates the county-quarter level index return at the purchase quarter of the underlying transaction. The term "t-1" indicates a one-quarter lag with respect to the purchase quarter. "Log county pop. density" is the natural log of the county's population density, derived as number of inhabitants at the purchase year per square-mile of land. "Log county income" is the natural log of the median county income at the purchase year, "County unemployment" is the county's unemployment rate in percentage points at the purchase year, "Log county turnover" the natural log of the sum of total transactions in the purchase quarter, and "County primary LTV" is the average mortgage amount divided by the average sales price in a given county in a given quarter, excluding LTVs larger than one. All regressions include purchase year-quarter fixed effect dummies. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

suggests that higher activity of ST IIs implies higher return risk, within counties. One reason for this could be that ST IIs themselves are driving the volatility they profit from, although Table 3.3.6 provides correlative evidence only. The positive coefficient for turnover in the county fixed effect models shown in columns (2) and (3) indicates that markets with increased trading volume are subject to higher return risk. Together with the finding on ST IIs, our results are in line with the finding of DeFusco et al. (2020), who argue that short-term speculators increase trading volume, and destabilize prices in local markets.

3.4. Robustness

In Section 3.3, we documented that investors' outperformance with respect to annualized excess capital gains can be explained by heterogeneous return risk exposure. The goal of this section is to provide evidence that this result is robust with respect to alternative model specifications and subsamples. Table 3.4.1 summarizes our robustness checks.

Panel A of Table 3.4.1 reports estimates for the relative alphas of investors based on alternative versions of variables that are used in our study throughout. First, we test for issues with simultaneity and look ahead bias by instrumenting contemporaneous return risk with one and two of its lags, respectively, following Han (2013). Using both of these instruments, we find that local return risk is still able to fully explain the outperformance of private investors as well as ST IIs. Even the performance dummy of LT IIs becomes insignificant, although the coefficient estimate remains positive and above one percentage point, underpinning the sensitivity of the outperformance of LT IIs to model specifications.

We then provide alternative definitions for "short-term". In the base case versions of our regressions, we specified three years of holding period as a cut-off, following DeFusco et al. (2020). As this might be considered a rather subjective choice, we test whether our results still hold when using two or four years as a breakpoint instead. The results show that when using two years as cut-off, the outperformance of all three investors remain insignificant, while using four years leads to a significant outperformance of LT IIs. Thus, our key result that local return risk explains private investor and ST II outperformance is robust to alternative cut-off values. Given that the outperformance of LT IIs seems to emerge with a breakpoint between two and four years, our choice of three years appears to be a balanced choice. Lastly, the results presented in the final column show that measuring return risk on the yearly level does not alter our results.

Table 3.3.6
Panel regressions explaining local return risk

	(1)	(2)	(3)
Housing market return	0.295***	0.226***	0.317***
	(0.023)	(0.021)	(0.022)
Adj. share ST II	0.138***	0.097***	0.087***
	(0.012)	(0.013)	(0.012)
Log income	-0.222***	-0.187***	0.106**
	(0.015)	(0.023)	(0.040)
Unemployment rate	-0.005***	-0.005***	0.004*
	(0.001)	(0.001)	(0.002)
Log population density	0.001	-0.106**	-0.017
	(0.004)	(0.036)	(0.035)
Log turnover	-0.009**	0.036***	0.038***
	(0.003)	(0.004)	(0.004)
Primary LTV	-0.032*	-0.089***	-0.094***
	(0.013)	(0.013)	(0.013)
County FE	-	X	X
Year FE	-	-	X
Adj. R-sq.	0.076	0.035	0.051
Observations	64,522	64,522	64,522

This table shows results for panel regressions on county-quarter level with local return risk as dependent variable. Each explanatory variable is included on the contemporaneous quarter, respectively. "Housing market return" is the quarterly county Case and Shiller (1989) index return, estimated with our housing data. "Adj. share ST II" is the share of selling short-term investors in a given market relative to the nationwide share of short-term sellers. "Log income" is the natural log of the median county income, "Unemployment" is the county's unemployment rate, "Log county density" is the natural log of the county's population density, measured as number of inhabitants per square mile of land, "Log turnover" the natural log of the sum of total transactions in the given period including observations with unknown identity, and "County primary LTV" is the average mortgage amount divided by the average sales price in a given county-quarter combination. "FE" stands for fixed effect dummies. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

Our full sample considers repeat sales generated between 1995 and 2017. Due to the starting period only very short holding periods can be found right at the beginning of our sample. To ease this restriction, we investigate two different subperiods that allow for a “burn-in” period of five years. Panel B shows the results when investigating transactions with sale dates between 2000 and 2015 only, but allowing purchases to range back until 1995. For this subperiod, the coefficients still remain insignificant. We additionally investigate a subsample of transactions with sales between 2000 and 2010 for two reasons. First, by investigating a period of 11 years while allowing for a 5-year burn-in period, we compare observations with more similar holding periods. Second, the period covers the most recent housing boom and bust, such that the period with the highest risk but also highest rewards is analyzed. The results show that during this period, none of the investors outperformed owner-occupiers after adjusting for local return risk. In contrast, the adjusted performance of ST IIs is even negative and borderline significant, suggesting that ST IIs took high risks during the boom that were overproportionally punished during the bust.

Finally, we restrict all groups to have holding period of three or less years, thus comparing only short-term investors. We do so to underscore that our holding period controls do perform well in removing short- and long-term related activities. Our results in the last column of Panel B of Table 3.4.1 show that when comparing short-term investors only, investors’ outperformance relative to owner-occupiers still remains insignificant. In sum, our robustness checks provide supporting evidence for local return risk as an important variable in explaining the systematic differences in capital gains across investor groups.

3.5. Conclusion

In this paper, we explore the systematic outperformance of investor groups relative to owner-occupiers. These investors are likely to differ in their underlying strategies and primary goals. While owner-occupiers should be primarily interested in drawing utility from housing consumption, investors should primarily hold financial interests in their homes. We thus expect investors to require compensation for engaging in riskier trades. Although private investors are likely to invest in close-by homes and save for retirement, their activity increased during the last housing boom, suggesting tendencies to speculate on momentum. Short-term institutional investors (ST IIs), in turn, are likely to engage in speculative activity, aiming for rapid appreciation rates for their investments. Long-term institutional investors (LT IIs), in contrast, should rather trade for steady dividends from rental income. Varying in their primary

Table 3.4.1
Robustness checks

Panel A: IV and alternative definitions of variables

	IV 1 lag	IV 2 lags	ST 2 yr	ST 4 yr	Yearly risk
Private inv.	-0.002	-0.002	0.001	0.001	-0.005
	(0.007)	(0.007)	(0.005)	(0.005)	(0.005)
ST II	-0.022	-0.029	-0.016	-0.009	-0.005
	(0.019)	(0.019)	(0.017)	(0.014)	(0.013)
LT II	0.011	0.013	0.010	0.019^{**}	0.009
	(0.008)	(0.008)	(0.007)	(0.006)	(0.006)

Panel B: Subsamples

	2000-2015	2000-2010	HP < 3 years
Private inv.	-0.005	-0.012	0.011
	(0.006)	(0.007)	(0.008)
ST II	-0.026	-0.027^{*}	0.014
	(0.014)	(0.013)	(0.015)
LT II	0.012	0.001	
	(0.007)	(0.011)	

This table shows results of OLS regressions with annualized excess capital gains as dependent variable. The displayed coefficients are estimates for group-specific dummies. The reference group consists of owner-occupiers, i.e., private individuals that live in the home purchased during the holding period. "Private inv.", "ST II", "LT II" are dummy coefficients indicating that the corresponding trade was made by a private investor, short-term institutional investor or long-term institutional investor, respectively. Private investors are defined as individuals that do not live in the underlying homes. A transaction is classified as "short-term" if the corresponding holding period is three years or less, and "long-term" otherwise. If not indicated otherwise, the regressions for which results are displayed in this table include group-specific exposure to the one-quarter lag of county return risk relative to the purchase and sale date, respectively. All regressions include controls for "Construction", "New construction", and "Holding period" deciles, as well as county-level fixed effect dummies. Panel A shows regression results when using the one-period lag of return risk as instrument for the contemporaneous risk, the one and two period lags as instrument similar to Han (2013), altering the definition of ST to 2 and 4 years, respectively, and by using the one-year lag of local return risk instead of the quarterly measure that is used in the base case regressions. Panel B shows three regressions on subsamples, using transactions realized with sales between 2000-2015 and 2000-2010, respectively, as well as one using transactions with a holding period of less than three years only. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

motivation and their underlying trading strategies, investors and owner-occupiers should be heterogeneously exposed to risk and momentum, affecting their respective performance.

Using nationwide data on transactions of US residential real estate, we document sizable and persistent differences in annualized excess capital gains among buyers in the housing market. We find that the heterogeneous exposure of investors to past local return risk explains a sizable share of performance differences across investors. Conditional on controls such as documented construction activity, holding period, and county fixed effects, the outperformance of private investors and ST II is even completely explained after accounting for heterogeneous risk exposure. In contrast, neither selection of particular areas nor timing with respect to the aggregate or local housing market states can explain the performance differences among market participants. Our results on the heterogeneous risk exposure are robust to instrumenting local return risk with up to two of its lags, alternative definitions of "short-term", different subperiods, and local fixed effects. Our results indicate that a sizable share of investor outperformance does not stem from the ability to select good locations or to time the market, but from the higher exposure to local risk.

References

Anenberg, E., A. Hizmo, E. Kung, and R. Mollow (2019). "Measuring Mortgage Credit Availability: A Frontier Estimation Approach." *Journal of Applied Econometrics* 34. (6), 865–882.

Anenberg, E. and E. Kung (2014). "Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures." *American Economic Review* 104. (8), 2527–2551.

Bayer, P., C. Geissler, K. Mangum, and J. W. Roberts (2020). "Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market." *Review of Financial Studies* 33. (11), 5212–5247.

Bayer, P., K. Mangum, and J. W. Roberts (2021). "Speculative Fever: Investor Contagion in the Housing Bubble." *American Economic Review* 111. (2), 609–651.

Bollerslev, T., A. J. Patton, and W. Wang (2016). "Daily House Price Indices: Construction, Modeling, and Longer-Run Predictions." *Journal of Applied Econometrics* 31. (6), 1005–1025.

Campbell, J. Y. and J. F. Cocco (2007). "How Do House Prices Affect Consumption? Evidence from Micro Data." *Journal of Monetary Economics* 54. (3), 591–621.

Campbell, J. Y., S. Giglio, and P. Pathak (2011). "Forced Sales and House Prices." *American Economic Review* 101. (5), 2108–2131.

Case, K. E. and R. J. Shiller (1989). "The Efficiency of the Market for Single-Family Homes." *American Economic Review* 79. (1), 125–137.

Chinco, A. and C. Mayer (2016). "Misinformed Speculators and Mispricing in the Housing Market." *Review of Financial Studies* 29. (2), 486–522.

Corradin, S. and A. Popov (2015). "House Prices, Home Equity Borrowing, and Entrepreneurship." *Review of Financial Studies* 28. (8), 2399–2428.

Cvijanovic, D. and C. Spaenjers (2020). "'We'll Always Have Paris': Out-of-Country Buyers in the Housing Market." *Management Science* forthcoming.

DeFusco, A., C. Nathanson, and E. Zwick (2020). *Speculative Dynamics of Prices and Volume*. Working Paper 23449. National Bureau of Economic Research.

Fischer, M., R. Füss, and S. Stehle (2021). "Local House Price Comovements." *Real Estate Economics* 49. (S1), 169–198.

Fu, Y. and W. Qian (2014). "Speculators and Price Overreaction in the Housing Market." *Real Estate Economics* 42. (4), 977–1007.

Fu, Y., W. Qian, and B. Yeung (2016). "Speculative Investors and Transactions Tax: Evidence from the Housing Market." *Management Science* 62. (11), 3254–3270.

Fuster, A. and J. Vickery (2015). "Securitization and the Fixed-Rate Mortgage." *Review of Financial Studies* 28. (1), 176–211.

Gao, Z., M. Sockin, and W. Xiong (2020). "Economic Consequences of Housing Speculation." *Review of Fiancial Studies* 33. (11), 5248–5287.

Gerardi, K., E. Rosenblatt, P. S. Willen, and V. Yao (2015). "Foreclosure Externalities: New Evidence." *Journal of Urban Economics* 87, 42–56.

Han, L. (2010). "The Effects of Price Risk on Housing Demand: Empirical Evidence from US Markets." *Review of Financial Studies* 23. (11), 3889–3928.

— (2013). "Understanding the Puzzling Risk-Return Relationship for Housing." *Review of Financial Studies* 26. (4), 877–928.

Hoechle, D., M. Schmid, and H. Zimmermann (2020). "Does Unobservable Heterogeneity Matter for Portfolio-Based Asset Pricing Tests?" Available at SSRN: https://ssrn.com/abstract=3569485.

Jensen, M. C., F. Black, and M. S. Scholes (1972). "The Capital Asset Pricing Model: Some Empirical Tests." In: *Studies in the Theory of Capital Markets*. Ed. by Jensen, M. C. Praeger Publishers Inc.

Lambie-Hanson, L., W. Li, and M. Slonkosky (2020). *Leaving Households Behind: Institutional Investors and the U.S. Housing Recovery*. Working Paper 19-1. Federal Reserve Bank of Philadelphia.

Landvoigt, T. (2017). "Housing Demand During the Boom: The Role of Expectations and Credit Constraints." *Review of Financial Studies* 30. (6), 1865–1902.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). "The Housing Market(s) of San Diego." *American Economic Review* 105. (4), 1371–1407.

Loutskina, E. and P. E. Strahan (2015). "Financial Integration, Housing, and Economic Volatility." *Journal of Financial Economics* 115. (1), 25–41.

Mian, A. and A. Sufi (2019). *Credit Supply and Housing Speculation*. Working Paper 24823. National Bureau of Economic Research.

Mills, J., R. Molloy, and R. Zarutskie (2019). "Large-Scale Buy-to-Rent Investors in the Single-Family Housing Market: The Emergence of a New Asset Class." *Real Estate Economics* 47. (2), 399–430.

Nathanson, C. G. and E. Zwick (2018). "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market." *Journal of Finance* 73. (6), 2587–2633.

Peng, L. (2016). "The Risk and Return of Commercial Real Estate: A Property Level Analysis." *Real Estate Economics* 44. (3), 555–583.

Peng, L. and T. G. Thibodeau (2017). "Idiosyncratic Risk of House Prices: Evidence from 26 Million Home Sales." *Real Estate Economics* 45. (2), 340–375.

Peng, L. and L. Zhang (2019). “House Prices and Systematic Risk: Evidence from Microdata.” *Real Estate Economics* forthcoming.

Saiz, A. (2010). “The Geographic Determinants of Housing Supply.” *Quarterly Journal of Economics* 125. (3), 1253–1296.

Appendix 3.A Additional tables

Table 3.A.1
Data cleaning process

	Dismissed observations	Remaining observations
Starting observations	-	99,757,949
Duplicates	13,790,976	85,966,973
Extreme prices	325,917	85,641,056
Resulting repeat sales	-	33,280,346
Holing period < 180 days	2,544,614	30,735,732
Extreme annualized capital gains	986,917	29,748,815
Unknown identity	3,531,577	26,217,238
Full sample extreme capital gains	524,344	25,692,894
REO/foreclosure related	4,514,025	21,178,869

This table shows the amount of observations lost in each cleaning step of the housing data. The steps are followed in descending order. The upper panel shows numbers for individual transactions, the lower panel for the resulting repeat sales. The starting observations are transactions of residential homes purchased from 1995 to 2017 that are declared "arms length" by the data provider, have a non-missing sales price and a full date including day, month, and year of sale. Duplicates are removed according to DeFusco et al. (2020). Extreme prices (annualized capital gains) are defined as outside the USD 5,000 to USD 100,000,000 (-50%, +100%) interval as in Bollerslev et al. (2016). Observations with "Unknown identity" are repeat sales to which no investor type could be assigned due to missing information on situs or mailing address. "Full sample extreme capital gains" are defined as exceeding the first and last percentile of the return distribution in the full sample at the given cleaning step. "REO/foreclosure related" are repeat sales for which either purchase or sale transaction was associated with a foreclosure or foreclosure related transfer of the underlying home, respectively.

Table 3.A.2
Investor performance, local return risk and local momentum

	(1)	(2)	(3)	(4)
Private inv.	0.013***	0.011**	0.005	0.006
	(0.001)	(0.004)	(0.005)	(0.005)
ST II	0.049***	0.045***	-0.015	-0.010
	(0.003)	(0.008)	(0.015)	(0.015)
LT II	0.031***	0.023***	0.016**	0.019***
	(0.001)	(0.005)	(0.006)	(0.006)
Local risk/return	-	purchase	sale	purchase & sale
Construction	X	X	X	X
New construction	X	X	X	X
Holding period	X	X	X	X
County FE	X	X	X	X
Adj. R-sq.	0.161	0.172	0.209	0.222
Observations	21,178,869	19,611,315	20,680,566	19,369,779

This table shows results for regressions with annualized excess capital gains of individual repeat sales as dependent variable. The coefficients displayed indicate average capital gains of investors relative to owner-occupiers, after including different sets of one-quarter lags of "Local momentum", i.e., the quarterly return of the county-level index, and "Local risk", i.e., the dispersion of county-level annualized capital gains relative to the nationwide dispersion of returns within the same quarter. Each set of lags is included either relative to the purchase quarter, the sale quarter, or both, respectively, as indicated in the row "Local risk/momentum". "Private inv.", "ST II", "LT II" are dummy coefficients indicating that the corresponding trade was made by a private investor, short-term institutional investor or long-term institutional investor, respectively. Private investors are defined as individuals that do not live in the underlying homes. A transaction is classified as "short-term" if the corresponding holding period is three years or less, and "long-term" otherwise. "Construction" indicates that the purchase transaction has associated documents indicating construction activity, excluding new constructions. "New construction" indicates that the purchase is identified as a transaction of a newly built home. "Holding period" stands for dummies based on holding period sample deciles. "County FE" stands for county-level fixed effect dummies. Standard errors are clustered over counties. *,**, and *** indicate significance on the 5%, 1%, and 0.1% level, respectively.

Complete References

Allen, M. and W. Dare (2002). “Identifying Determinants of Horizontal Property Tax Inequity: Evidence from Florida.” *Journal of Real Estate Research* 24. (2), 153–164.

Amromin, G., J. C. Huang, C. Sialm, and E. Zhong (2018). “Complex Mortgages.” *Review of Finance* 22. (6), 1975–2007.

Andersen, S., C. Badarinza, L. Liu, J. Marx, and T. Ramadorai (2021). “Reference Dependence in the Housing Market.” Available at SSRN: https://ssrn.com/abstract=3396506.

Anenberg, E., A. Hizmo, E. Kung, and R. Mollow (2019). “Measuring Mortgage Credit Availability: A Frontier Estimation Approach.” *Journal of Applied Econometrics* 34. (6), 865–882.

Anenberg, E. and E. Kung (2014). “Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures.” *American Economic Review* 104. (8), 2527–2551.

Ang, A. and J. Chen (2002). “Asymmetric Correlations of Equity Portfolios.” *Journal of Financial Economics* 63. (3), 443–494.

Autor, D. H., C. J. Palmer, and P. A. Pathak (2014). “Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts.” *Journal of Political Economy* 122. (3), 661–717.

Bai, C., Q. Li, and M. Ouyang (2014). “Property Taxes and Home Prices: A Tale of Two Cities.” *Journal of Econometrics* 180. (1), 1–15.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). “The Economic Effects of Social Networks: Evidence from the Housing Market.” *Journal of Political Economy* 126. (6), 2224–2276.

Bayer, P., C. Geissler, K. Mangum, and J. W. Roberts (2020). “Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market.” *Review of Financial Studies* 33. (11), 5212–5247.

Bayer, P., K. Mangum, and J. W. Roberts (2021). “Speculative Fever: Investor Contagion in the Housing Bubble.” *American Economic Review* 111. (2), 609–651.

Bhattacharya, U., D. Huang, and K. M. Nielsen (2020). “Spillovers in Prices: The Curious Case of Haunted Houses.” *Review of Finance* forthcoming.

Black, R. T. and J. Diaz III (1996). “The Use of Information Versus Asking Price in the Real Property Negotiation Process.” *Journal of Property Research* 13. (4), 287–297.

Bollerslev, T., A. J. Patton, and W. Wang (2016). “Daily House Price Indices: Construction, Modeling, and Longer-Run Predictions.” *Journal of Applied Econometrics* 31. (6), 1005–1025.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). “Robust Inference With Multiway Clustering.” *Journal of Business and Economic Statistics* 29. (2), 238–249.

Campbell, J. Y. and J. F. Cocco (2007). “How Do House Prices Affect Consumption? Evidence from Micro Data.” *Journal of Monetary Economics* 54. (3), 591–621.

Campbell, J. Y., S. Giglio, and P. Pathak (2011). “Forced Sales and House Prices.” *American Economic Review* 101. (5), 2108–2131.

Can, A. (1990). “The Measurement of Neighborhood Dynamics in Urban House Prices.” *Economic Geography* 66. (3), 254–272.

Caplin, A. and J. Leahy (2011). “Trading Frictions and House Price Dynamics.” *Journal of Money, Credit and Banking* 43. (7), 283–303.

Case, K. E. and C. J. Mayer (1996). “Housing Price Dynamics within a Metropolitan Area.” *Regional Science and Urban Economics* 26. (3-4), 387–407.

Case, K. E. and R. J. Shiller (1989). “The Efficiency of the Market for Single-Family Homes.” *American Economic Review* 79. (1), 125–137.

Chambers, M., C. Garriga, and D. E. Schlagenhauf (2009). “Accounting for Changes in the Homeownership Rate.” *International Economic Review* 50. (3), 677–726.

Chay, K. Y. and M. Greenstone (2005). “Does Air Quality Matter? Evidence from the Housing Market.” *Journal of Political Economy* 113. (2), 376–424.

Chinco, A. and C. Mayer (2016). “Misinformed Speculators and Mispricing in the Housing Market.” *Review of Financial Studies* 29. (2), 486–522.

Cohen, J. P. and J. Zabel (2018). “Local House Price Diffusion.” *Real Estate Economics* 48. (3), 710–743.

Corradin, S. and A. Popov (2015). “House Prices, Home Equity Borrowing, and Entrepreneurship.” *Review of Financial Studies* 28. (8), 2399–2428.

Cotter, J., S. Gabriel, and R. Roll (2015). “Can Housing Risk Be Diversified? A Cautionary Tale from the Housing Boom and Bust.” *Review of Financial Studies* 38. (3), 913–936.

Cvijanovic, D. and C. Spaenjers (2020). “‘We’ll Always Have Paris’: Out-of-Country Buyers in the Housing Market.” *Management Science* forthcoming.

Cypher, M. and J. A. Hansz (2003). “Does Assessed Value Influence Market Value Judgments?” *Journal of Property Research* 20. (4), 305–318.

DeFusco, A., C. Nathanson, and E. Zwick (2020). *Speculative Dynamics of Prices and Volume*. Working Paper 23449. National Bureau of Economic Research.

Diamond, R. and T. McQuade (2019). “Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development.” *Journal of Political Economy* 127. (3), 1063–1117.

Du, Z. and L. Zhang (2015). “Home-Purchase Restriction, Property Tax and Housing Price in China: A Counterfactual Analysis.” *Journal of Econometrics* 188. (2), 558–568.

Elinder, M. and L. Persson (2017). “House Price Responses to a National Property Tax Reform.” *Journal of Economic Behavior & Organization* 144, 18–39.

Fischer, M., R. Füss, and S. Stehle (2021). “Local House Price Comovements.” *Real Estate Economics* 49. (S1), 169–198.

Fu, Y. and W. Qian (2014). “Speculators and Price Overreaction in the Housing Market.” *Real Estate Economics* 42. (4), 977–1007.

Fu, Y., W. Qian, and B. Yeung (2016). “Speculative Investors and Transactions Tax: Evidence from the Housing Market.” *Management Science* 62. (11), 3254–3270.

Fuster, A. and J. Vickery (2015). “Securitization and the Fixed-Rate Mortgage.” *Review of Financial Studies* 28. (1), 176–211.

Gao, Z., M. Sockin, and W. Xiong (2020). “Economic Consequences of Housing Speculation.” *Review of Fiancial Studies* 33. (11), 5248–5287.

Genesove, D. and C. Mayer (2001). “Loss Aversion and Seller Behavior: Evidence From the Housing Market.” *Quarterly Journal of Economics* 116. (4), 1233–1260.

Gerardi, K., E. Rosenblatt, P. S. Willen, and V. Yao (2015). “Foreclosure Externalities: New Evidence.” *Journal of Urban Economics* 87, 42–56.

Gibbs, C. G. and M. Kulish (2017). “Disinflations in a Model of Imperfectly Anchored Expectations.” *European Economic Review* 100, 157–174.

Goolsby, W. (1997). “Assessment Error in the Valuation of Owner-Occupied Housing.” *Journal of Real Estate Research* 13. (1), 33–45.

Guerrieri, V., D. Hartley, and E. Hurst (2013). “Endogenous Gentrification and Housing Price Dynamics.” *Journal of Public Economics* 100, 45–60.

Gupta, A. (2019). “Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults.” *Journal of Finance* 74. (5), 2249–2301.

Guren, A. M. and T. J. McQuade (2020). “How Do Foreclosures Exacerbate Housing Downturns?” *Review of Economic Studies* 87. (3), 1331–1364.

Han, L. (2010). "The Effects of Price Risk on Housing Demand: Empirical Evidence from US Markets." *Review of Financial Studies* 23. (11), 3889–3928.

— (2013). "Understanding the Puzzling Risk-Return Relationship for Housing." *Review of Financial Studies* 26. (4), 877–928.

Harding, J. P., E. Rosenblatt, and V. W. Yao (2009). "The Contagion Effect of Foreclosed Properties." *Journal of Urban Economics* 66. (3), 164–178.

Hilber, C. A. (2017). "The Economic Implications of House Price Capitalization: A Synthesis." *Real Estate Economics* 45. (2), 301–339.

Hodge, T. R., D. P. McMillen, G. Sands, and M. Skidmore (2017). "Assessment Inequity in a Declining Housing Market: The Case of Detroit." *Real Estate Economics* 45. (2), 237–258.

Hoechle, D., M. Schmid, and H. Zimmermann (2020). "Does Unobservable Heterogeneity Matter for Portfolio-Based Asset Pricing Tests?" Available at SSRN: https://ssrn.com/abstract=3569485.

Jensen, M. C., F. Black, and M. S. Scholes (1972). "The Capital Asset Pricing Model: Some Empirical Tests." In: *Studies in the Theory of Capital Markets*. Ed. by Jensen, M. C. Praeger Publishers Inc.

Jones, P. (2020). "Loss Aversion and Property Tax Avoidance." Available at SSRN: https://ssrn.com/abstract=3511751.

Kallberg, J. G., C. H. Liu, and P. Pasquariello (2014). "On the Price Comovement of US Residential Real Estate Markets." *Real Estate Economics* 42. (1), 71–108.

Lambie-Hanson, L., W. Li, and M. Slonkosky (2020). *Leaving Households Behind: Institutional Investors and the U.S. Housing Recovery*. Working Paper 19-1. Federal Reserve Bank of Philadelphia.

Landier, A., D. Sraer, and D. Thesmar (2017). "Banking Integration and House Price Co-Movement." *Journal of Financial Economics* 125. (1), 1–25.

Landvoigt, T. (2017). "Housing Demand During the Boom: The Role of Expectations and Credit Constraints." *Review of Financial Studies* 30. (6), 1865–1902.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). "The Housing Market(s) of San Diego." *American Economic Review* 105. (4), 1371–1407.

Levy, D., Z. Dong, and J. Young (2016). "Unintended Consequences: The Use of Property Tax Valuations as Guide Prices in Wellington, New Zealand." *Housing Studies* 31. (5), 578–597.

Livy, M. R. (2018). "Intra-School District Capitalization of Property Tax Rates." *Journal of Housing Economics* 41, 227–236.

Loutskina, E. and P. E. Strahan (2015). “Financial Integration, Housing, and Economic Volatility.” *Journal of Financial Economics* 115. (1), 25–41.

McMillen, D. P. (2013). “The Effect of Appeals on Assessment Ratio Distributions: Some Nonparametric Approaches.” *Real Estate Economics* 41. (1), 165–191.

Mian, A. and A. Sufi (2019). *Credit Supply and Housing Speculation.* Working Paper 24823. National Bureau of Economic Research.

Mills, J., R. Molloy, and R. Zarutskie (2019). “Large-Scale Buy-to-Rent Investors in the Single-Family Housing Market: The Emergence of a New Asset Class.” *Real Estate Economics* 47. (2), 399–430.

Murfin, J. and R. Pratt (2019). “Comparables Pricing.” *Review of Financial Studies* 32. (2), 688–737.

Nathanson, C. G. and E. Zwick (2018). “Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market.” *Journal of Finance* 73. (6), 2587–2633.

Ngai, L. R. and S. Tenreyro (2014). “Hot and Cold Seasons in the Housing Market.” *American Economic Review* 104. (12), 3991–4026.

Northcraft, G. and M. Neale (1987). “Experts, Amateurs, and Real Estate: An Anchoring-and-Adjustment Perspective on Property Pricing Decisions.” *Organizational Behavior and Human Decision Processes* 39. (1), 84–97.

Palmon, O. and B. A. Smith (1998). “New Evidence on Property Tax Capitalization.” *Journal of Political Economy* 106. (5), 1099–1111.

Peng, L. (2016). “The Risk and Return of Commercial Real Estate: A Property Level Analysis.” *Real Estate Economics* 44. (3), 555–583.

Peng, L. and T. G. Thibodeau (2017). “Idiosyncratic Risk of House Prices: Evidence from 26 Million Home Sales.” *Real Estate Economics* 45. (2), 340–375.

Peng, L. and L. Zhang (2019). “House Prices and Systematic Risk: Evidence from Microdata.” *Real Estate Economics* forthcoming.

Rossi-Hansberg, E., P.-D. Sarte, and R. Owens III (2010). “Housing Externalities.” *Journal of Political Economy* 118. (3), 485–535.

Saiz, A. (2010). “The Geographic Determinants of Housing Supply.” *Quarterly Journal of Economics* 125. (3), 1253–1296.

Sirmans, S., D. Gatzlaff, and D. Macpherson (2008). “Horizontal and Vertical Inequity in Real Property Taxation.” *Journal of Real Estate Literature* 16. (2), 167–180.

Szumilo, N. (2020). “Prices of Peers: Identifying Endogenous Price Effects in the Housing Market.” *Economic Journal* forthcoming.

Tu, Y., S.-M. Yu, and H. Sun (2004). “Transaction-Based Office Price Indexes: A Spatiotemporal Modeling Approach.” *Real Estate Economics* 32. (2), 297–328.

Tversky, A. and D. Kahneman (1974). “Judgment Under Uncertainty: Heuristics and Biases.” *Science* 185. (4157), 1124–1131.

Wassmer, R. W. (1993). “Property Taxation, Property Base, and Property Value: An Empirical Test of the “New View”.” *National Tax Journal* 46. (2), 135–159.

Eigenabgrenzung

Das erste Kapitel, „*Local House Price Comovements*“, ist in Zusammenarbeit mit Marcel Fischer, Professor für Finanzwirtschaft an der Universität Konstanz, und Roland Füss, Professor für Real Estate Finance an der Universität St. Gallen, entstanden. Mein individueller Anteil an diesem Projekt beträgt 60%.

Das zweite Kapitel, „*How Do Assessed Values Affect Transaction Prices of Homes?*“, habe ich in Eigenarbeit erstellt.

Das dritte Kapitel, „*Investors in the Housing Market*“, ist in Kooperation mit Marcel Fischer und Roland Füss, sowie Daniel Ruf, Juniorprofessor für Real Estate Finance an der Goethe Universität Frankfurt, entstanden. Mein individueller Anteil an diesem Projekt beträgt 50%.