Anthony Lee Zhang's Research Statement

January 25, 2024

1 Unifying Themes

In my research, I use simple economic models to analyze unique institutional settings, and show how they can be redesigned and regulated to work better. In some of my best papers, I look at new or understudied institutions which we do not have a precise economic understanding of; for each institution, I try to create a framework which aims to be the "benchmark" lens through which the institution will in the future be viewed by academics, policymakers and practitioners.

I have worked on three main areas: cryptocurrencies and decentralized finance; finance and industrial organization; and US and Asian real estate markets.

2 Cryptocurrencies and Decentralized Finance

What are cryptocurrencies? To buy the Ethereum cryptocurrency, I can send US dollars via a bank transfer to Coinbase, a cryptocurrency exchange operating in the US, and submit a market order to buy ETH tokens for USD. I can then "withdraw" these ETH tokens, ordering Coinbase to send the tokens to a wallet I control. At this point, I own these ETH tokens in my "wallet," in the sense that I can send them to any other wallet using a "private key" password I hold.

As a US household, this functionality has very few uses. Sending ETH takes around 12 seconds and costs around $10 USD; Venmo, Zelle, and many options are essentially strictly dominating alternatives. What is interesting about cryptocurrencies is that these wallets can be operated by anyone with an internet connection: thus, any individual in any country with internet access can use a moderately costly, moderately fast form of e-money, on the same "payment rails", with no jurisdictional restrictions. Once I hold ETH in my "wallet", I can send it to anyone anywhere, in 12 seconds for $10. Cryptocurrencies thus have seen nontrivial adoption in places where traditional financial systems are inefficient or extractive; for example, if inflation is high, or property rights are insecure.

The global crypto market is nontrivially large: the total market cap of all major cryptocurrencies is slightly above $1 trillion: for reference, the market cap of all S&P 500 companies is around $42 trillion. The crypto ecosystem has developed a number of unique institutions, which play roles similar to institutions in traditional finance like market makers, banks, and exchanges, but often with peculiar differences to traditional entities. Together, these institutions are sometimes referred to as the "decentralized finance", or "defi", ecosystem.

In four papers, I show how these "defi" institutions can be understood through analogies to traditional concepts and institutions in economics and finance. My work shows how policymakers can more effectively regulate these institutions, as well as how industry participants can build improvements to these systems.

2.1 Automated Market Makers

In traditional financial markets, if you use the Robinhood app and spend USD to buy a share of Apple, you are generally buying from a market making firm, such as Citadel or Jump Trading. If you want to spend, say, the ETH token to buy some quantity of, say, the MKR token, another option is to trade with an *automated market maker* (AMM). AMMs can be thought of like programmatic robots which play the role of Citadel and Jump in blockchain financial markets: at any point in time, the Uniswap v2 AMM quotes prices for trading ETH for MKR, in a manner that is fully automated, with no human intervention or judgment. AMMs are quite successful: as of January 4th 2024, Uniswap traded $2.3 trillion USD in cumulative volume, with $3bil USD in daily volume, nontrivial relative to the $244bil traded on the NYSE on January 3rd. This raises the question: what is the core difference between AMMs and traditional market making mechanisms? Are AMMs more or less efficient than traditional market making mechanisms, and how big is the difference?

We address this question in my paper *Automated Market Making and Loss-Versus-Rebalancing* (Milionis, Moallemi, Roughgarden & Zhang, 2023). Automated market makers, incredibly, make markets without observing prices of assets like MKR on deeper centralized exchanges. As a result, in the manner of Budish, Cramton, & Shim (2015) AMMs suffer systematic *price slippage*: when big exchange prices move, AMM quotes become "stale", leaving profits to arbitrageurs who "snipe" AMMs' stale quotes against big exchange quotes. A natural definition of AMMs' losses is the volume-weighted cumulative gap between AMM trade prices and big exchange prices; we call this gap *loss versus rebalancing* (LVR). This is a neat technical concept which can be connected not only to microstructure ideas, but also to options theory, and even classical demand theory and the consumer expenditure minimization problem.

The paper, while still unpublished, has had a fairly large impact on industry discourse and direction. We believe loss versus rebalancing (LVR), a term we coined in this paper, is now the main term and measure used in industry to measure AMM losses. Dan Robinson, general partner and head of research at Paradigm, a leading crypto VC firm, has stated: "LVR is obviously the correct metric for measuring "impermanent loss" as a running cost." Many product teams are trying to build protocols to "fix" LVR,[1] and many industry participants have written about LVR.[2]

2.2 Stablecoins

Stablecoins are crypto tokens designed to be worth a dollar each. Stablecoins are useful because they allow users to use crypto "payment rails", while maintaining the price stability of US dollars. For example, if an individual in a developing country would like to hold USD to avoid capital controls or inflation, stablecoins are in some ways more convenient than physical cash, and have essentially no jurisdictional restrictions, being accessible by anyone with internet access. The total outstanding market cap of USD stablecoins is around $130bil USD as of January 4th, 2024. Stablecoins are created by issuers through a process very analogous to how banks create deposits. I send a stablecoin issuer $1 USD via wire transfer; the issuer sends me a "crypto token" which is essentially a tokenized deposit; and the issuer buys USD assets with positive yield to back the tokens. The issuer profits from this, since the assets pay interest, which is currently not passed on to stablecoin holders.

In *Stablecoin Runs and the Centralization of Arbitrage* (Ma, Zeng & Zhang, 2023), we analyze an interesting feature of stablecoin design: like ETFs, stablecoins generally cannot be created or redeemed in primary markets at will. Instead, if I want to buy or sell stablecoins, I trade on *secondary* markets from other holders. A special set of arbitrageurs, like "authorized participants" in ETF markets, have the ability to create and redeem stablecoins for cash. Arbitrage is imperfectly efficient: there are very few arbitrageurs for many coins, so prices can deviate significantly from $1 before arbs trade to revert these price fluctuations. Why would issuers do this? We argue that issuers *purposefully* limit arbitrage efficiency to limit *run risk*. Inefficient arbitrage implies that, when holders panic and try to sell, *prices* move significantly below $1, but this translates into little redemption *quantities*, so the issuer doesn't have to sell as many assets, making runs less likely. This is somewhat analogous to traditional banks suspending withdrawals, or imposing redemption fees, as a mechanism to combat runs.

Our results help show how stablecoins compare to well-understood traditional financial intermediaries like banks and ETFs. Our results also inform ongoing stablecoin policy debates: whether stablecoins should be considered securities, whether they should be able to pay dividends, and what metrics regulators should track to monitor stablecoin run risks.

[1] Some examples are CoW Protocol, Aori, Cata Labs, ALGEBRA protocol, Poolshark, and VoLVER.
[2] See, for example, Delphi Digital, Blockworks, Titania, and Frontier Research.

2.3 Cryptocurrency Exchanges

Individuals buy cryptocurrencies with fiat using crypto exchanges, such as Coinbase and Binance. There are a surprisingly large number of such exchanges: why are there so many, and how do they compete with each other? In *Competition in the Cryptocurrency Exchange Market* (Hu and Zhang, 2023), we argue theoretically and empirically that small crypto exchanges behave basically like "brokers" and large exchanges like "inter-broker clearinghouses". A small exchange in, say, Australia, figures out how to interface with Australia's banking and payments systems and financial regulators, thus allowing Australian customers to legally deposit Australian currency and purchase crypto; a larger exchange like Binance may be unable or unwilling to work with local payments rails and regulators. The Australian exchange then links to large global exchanges, like Binance or Coinbase, through arbitrage, so Australian customers – net of trading fees – trade in the global crypto marketplace.

Liquidity thus concentrates on deep large global exchanges, and the role of small exchanges is basically to serve as "costly windows", which specialize in capturing customers and dealing with local regulations and rails, but whose liquidity basically derives from the large global exchanges. Our results imply that Binance and Coinbase's listing decisions "lead" the rest of the world: since small exchanges rely on Binance/Coinbase for a nontrivial share of their liquidity, they will tend to follow Binance/Coinbase's token listing decisions. These private, profit-motivated entities thus play a role in the global crypto marketplace which is somewhat analogous to the role the SEC plays in choosing which securities can be traded in the US, a situation regulators may want to monitor more closely.

2.4 Non-Fungible Tokens

"Non-fungible tokens" are crypto tokens that are designed to represent digital images: for example, the "Bored Ape Yacht Club" is a collection of 10,000 tokens, each representing a slightly different cartoon ape picture. The cheapest such token, as of writing, is listed for sale over $60,000 USD. These tokens convey no cash flow rights and have essentially no use. Essentially the only benefit one gets from purchasing a $60,000 cartoon ape is that one can say one spent $60,000 on a cartoon ape, and prove it on the blockchain. Yet, for most of 2021, daily NFT trading volume was in the hundreds of millions of USD (though volume is now substantially lower).

Why?

In *Digital Veblen Goods* (Oh, Rosen & Zhang, 2023), we propose that NFTs are basically *social goods*. Like a luxury watch or handbag, an overpriced cartoon ape picture is valuable to me partly because many others demand the cartoon ape. A classic paper by Becker (1991) notes that, in the presence of social influences on demand, a number of phenomena that would be puzzling in traditional private-valued models are possible. Multiple equilibria are possible; goods with the same fundamentals may be "in" or "out", since consumers value a thing if they think other consumers value it: demand begets demand. Sellers respond by underpricing assets: Rolexes and other such luxury goods are often sold in primary markets for a price well below the secondary market clearing price. This pricing strategy could never be optimal for a traditional monopolist seller; however, Rolex demand is fragile, since the Rolex seller realizes that a small change in price can potentially cause demand to collapse from "in" to "out", hence purposefully induces undersupply and rationing in primary markets. The Rolex seller's pricing strategy is exploitable by scalpers, who buy on primary and flip to secondary for a profit.

We find support for all three empirical predictions in the NFT market: demand is very bimodal, with NFT collections either "in" or "out", but little in between; sellers underprice in primary markets; and scalpers exploit sellers' primary market underpricing. From the perspective of crypto research, our paper attempts to answer the "why" question of fundamentally how to think about the NFT market's existence. From the perspective of economics research more broadly, our study uses the NFT market as a relatively clean and data-rich setting to analyze and test the Becker (1991) theory of social goods pricing; our findings are potentially relevant for markets for luxury fashion products, restaurants, and nightclubs, and other settings with social effects on demand.

2.5 Other Work

Besides my academic research, I have developed an MBA/law student class, *Blockchain, Cryptocurrencies, and Web3*, co-taught with Anup Malani; the class is mostly new content and I believe it is currently one of world's best MBA-level treatments of web3 and decentralized finance content. I have also written a number of nontechnical and semi-technical blog posts about related topics, some of which I use as teaching material, and some of which has had some publicity.

3 Industrial Organization of Financial Markets

I have a number of papers on the industrial organization of financial markets. These papers analyze imperfect competition in the context of various financial market settings: derivative markets, where derivative contract holders may exercise market power to manipulate payoffs; credit markets where lenders invest in costly technologies to screen agents; and interbank repo markets, where large dealers' market power over their customers limits the efficiency of monetary policy passthrough.

3.1 Derivative Market Manipulation

In my job market paper, *Competition and Manipulation in Derivative Contract Markets* (Zhang, 2022), published in the *Journal of Financial Economics*, I analyze the issue of manipulation in derivative contract markets.

Suppose a market participant wants to buy or sell exposure to some risk factor, such as oil prices, interest rates, or wheat prices; the simplest way to do this is often to use a derivative contract. A long position in the ICE Houston Ship Channel gas basis futures contract, at settlement, pays its holder some multiple of a gas price index, which is calculated based on prices of physical gas traded in Houston. If the gas price index accurately reflects real gas prices, then the futures contract allows its holder to hedge gas price risk.

These financial markets are naturally "derivative" of the physical goods markets they are based off, in the sense that they entitle holders to payoffs based on prices from the underlying spot markets. However, interestingly, derivative markets are often much larger than their underlying spot markets: for example, the Platts *Inside FERC* Houston Ship Channel benchmark is based on around 1.4 million MMBtus of natural gas trades per week; the ICE HSC basis future, based on this benchmark, has open interest of over 75 million MMBtus. The fact that derivative markets are so much larger than their underlying spot markets leads to the possibility of an interesting trading strategy: if a trader is long the Houston Ship Channel futures contract, she can buy large amounts of physical gas at Houston to raise the contract settlement price. If the trader has a large enough futures position, her increased futures profits may outweigh any losses in spot markets.

Regulators generally consider trading in this manner to constitute illegal manipulation, and have imposed billions of dollars of fines on market participants for manipulation in the past two decades alone. But manipulation is both vaguely defined legally, and purely understood from the perspective of economic theory. It is not clear how manipulation influences trader welfare, and whether it is in fact a market failure which can be resolved through regulation. It is also not clear what makes contract markets vulnerable to manipulation, or how to empirically measure contract market manipulation risk.

In this paper, I attempt to model manipulation and show how to measure it empirically. In the model, agents buy futures contracts to hedge risk, but cannot commit not to "manipulate" in spot markets, trying to move spot prices to influence their futures payoffs. Manipulation is a market failure in the strong sense that it can be *Pareto disimproving*: all agents may prefer if agents could commit not to manipulate, since agents' manipulation ends up creating nonfundamental risk for all agents. A regulator could in some cases create a Pareto improvement by imposing taxes on the size of agents' contract positions, a policy which approximates "position limits" imposed in practice. Moreover, I show that manipulation is more likely if liquidity mismatch is more severe, and when derivative markets are more concentrated, showing regulators how to detect markets that may be more vulnerable to manipulation.

The paper has had some impact on policy discussion: it has been covered in Regulatory Compliance Watch, and was cited in a Petition to the CFTC to ban water index futures.

3.2 Data in Consumer Credit Markets

Lenders across many consumer credit markets are increasingly purchasing data to improve lending decisions: how does this trend affect equilibrium market structure, lending decisions, and welfare of different market participants? In *Data and Welfare in Credit Markets* (Jansen, Nagel, Yannelis & Zhang, 2023), we analyze how data-based pricing affects social welfare of different market participants. The key insight is that data can be thought of as a form of *third-degree price discrimination*. With data, lenders can set different prices for two subgroups of consumers which were previously indistinguishable. However, a difference between this setting and classic results about third-degree price discrimination is that, in credit markets, data is primarily informative about consumers' *costs*, rather than consumer WTP. We show that, in contrast to classic results, more data is unambiguously good for social welfare in the setting we study. However, using a simple supply-demand argument, we show that the transfers induced by data availability between groups tend to be large relative to the social welfare gains – essentially, we are comparing consumer surplus rectangles to deadweight loss triangles. Thus, data removal is never a first-best way to transfer welfare between consumer groups, but in settings where data does not induce very large price changes, the welfare cost of data removal will tend to be small relative to the size of the transfers induced. Our methodology gives a way to quantify just how bad data removal is as a transfer tool: how many cents of social welfare are burned, per dollar transferred between groups.

In *Competition and Selection in Credit Markets* (Yannelis & Zhang, 2023), published in the *Journal of Financial Economics* in 2023, I analyze how data affects competition between lenders. We posit that data technologies are basically fixed-cost: it is costly to develop a model, but the cost of applying a developed model to a new customer is fairly low. This leads to a counterintuitive prediction about how competition influences market outcomes: when markets are more competitive, lenders capture smaller market shares, meaning that their fixed data costs are amortized over a smaller customer base, so they have lower incentives to invest in data acquisition. We find empirical support for the model's predictions.

In ongoing work, we are looking for ways to develop richer and more detailed models of the welfare effects of increased data use in credit markets. We hope that our findings in this stream of work can be useful for policymakers deciding on, for example, whether to enact data use restriction policies, such as the GDPR, FCRA, and CCPA, and how to think about competition policy in credit markets where lenders are making intensive investments in data use.

3.3 Market Power and Competition Policy in Financial Markets

In *Monetary Policy Transmission in Segmented Markets* (Eisenschmidt, Ma & Zhang, 2024), published in the *Journal of Financial Economics* in 2024, we analyze how dealer market power inhibits monetary policy transmission in European repo markets. Interbank markets are an important first stage in the transmission of monetary policy: in Europe, the ECB sets the deposit facility rate, which influences the rate at which large dealer banks trade "repo" with each other, which then propagates to the short-term rates facing dealer banks' over-the-counter (OTC) customers: other banks, insurance and pension funds, and hedge funds, among other entities. Ideally, when the ECB raises rates, dealer banks should raise the rates they offer to their OTC customers, allowing ECB policy shifts to "pass through" efficiently to these institutions. However, the European interbank market is surprisingly concentrated: the majority of OTC customers interact with only one or two dealer banks. In classic theory, firms with market power may not pass through cost shocks one-to-one to their customers: analogously, we find that dealers do not pass through repo rate shocks one-to-one to their OTC customers, and pass-through is worse for firms which are more poorly connected to dealers. We estimate a structural model to show how dealer market power influences monetary policy pass-through to repo customers. Our model implies that imperfect competition among repo dealers meaningfully inhibits monetary policy transmission in the Euro area. Our results also show how different policy interventions, such as giving customers direct access to inter-dealer markets or an ECB repo facility, could improve monetary policy pass-through. The paper was referenced in a speech by Isabel Schnabel, member of the Executive Board of the ECB, at the ECB conference on money markets, and also received coverage from Central Banking magazine.

4 Real Estate

I have a number of papers on real estate markets. In two papers, I analyze large policy-relevant and understudied issues in Asian real estate markets. In two other papers, I analyze liquidity measurement and the consequences of illiquidity in US real estate markets.

4.1 The Chonsei System

In *The Credit Channel of Monetary Policy Transmission: Evidence From the Chonsei System* (Jing, Park & Zhang, 2022), we analyze a unique housing finance system in Korean housing markets. In a Chonsei arrangement, a tenant makes a very large "security deposit" – often between 40% to 70% of the price of the property – to their landlord, and in return pay no rent to landlords. The tenant is basically making a zero-interest loan to her landlord; the tenant thus trades off the rent savings with the foregone interest on the Chonsei deposit. The landlord is essentially taking a mortgage loan from her tenant: unlike security deposits in the US, the Chonsei funds are unconstrained, so landlords can use them for essentially any purpose. As unusual as this arrangement seems, Chonsei arrangements are extremely popular: for much of the past 20 years, more tenants have used Chonsei agreements than standard flow-payment rental agreements in Korea. These loans are also short-term – mostly two years – and landlords tend to "roll" the loans, returning money borrowed from one tenant by borrowing from the next.

What determines the equilibrium size of a Chonsei loan? The intuition is simple: since the tenant chooses between foregone interest on the loan and paying rent, the loan size must be such that the foregone interest is roughly equal to rent. But this implies that interest rates dramatically affect loan size. If rent is $10k USD a year and renter-facing interest rates are 5%, Chonsei loan size is $200k; if interest rates decrease to 2%, Chonsei loan size skyrockets to $500k. The landlord is quite happy if rates decrease: she owes a leaving tenant $200k, but this is more than paid off by the $500k she gets from her next tenant. But when rates increase, landlords face the problem of returning the $500k loan to the leaving tenant, when the market loan size has decreased to $200k. If the landlord defaults, the tenant keeps the house, but this has potentially also declined in value due to rate increases. Monetary policy thus has perverse effects on loan sizes, and thus house prices, in the Chonsei system. While this effect is ex-post fairly intuitive, to our knowledge we are the first to analyze monetary policy passthrough in the Chonsei system.

The Chonsei system is currently a pressing policy issue in Korea: the Bank of Korea raised rates substantially after COVID, leading to a large wave of Chonsei defaults. In the paper, we construct an equilibrium model of Chonsei credit determination and its effect on house prices. Our paper gives concrete policy recommendations, showing how the Chonsei system could be redesigned to improve financial stability. Our work has been covered in the Chosun Ilbo, one of the largest newspapers in Korea.

4.2 Chinese Local Governments' Land Sale Decisions

In *Zoning for Profits: How Public Finance Shapes Land Supply in China* (He, Nelson, Su, Zhang & Zhang, 2023), we analyze Chinese local governments' land sale decisions. A large part of Chinese local government revenues come from selling land development rights to private parties. Land can be zoned for "residential" use (condos), or "industrial" use (factories). Puzzlingly, the "law of one price" is substantially violated: matched residential land parcels sell for over ten times higher prices than industrial parcels! However, industrial land sales to firms eventually increase the output taxes paid to local governments, whereas residential land sales generate no direct tax revenues, since there are currently no property taxes in China. We thus argue that local governments' decision to sell residential vs industrial land is an intertemporal choice problem: residential pays more upfront, industrial pays more in the future. This is interesting because it implies that Chinese local governments' finances are intertwined with land sale decisions: a financially stressed government, being impatient, may sell large amounts of residential land even when there is little market demand for this land. In follow-on work in progress, I am working on a project to understand the role of land development companies in Chinese land markets, and how they may have contributed to the runup in Chinese real estate prices.

4.3 Liquidity in US Residential Real Estate Markets

I have two papers on the topic of liquidity in US housing markets. Houses are a large component of many households' wealth. Housing markets are also very illiquid: houses take time to sell and individual houses' sale prices are very hard to predict. Housing market illiquidity has important aggregate and distributional consequences for households. In *Collateral Value Uncertainty and Mortgage Credit Provision* (Jiang & Zhang, 2023), which is R&R at the *Journal of Financial Economics*, we show that house price uncertainty influences mortgage credit provision. Suppose a bank lends against a house, and the buyer defaults, so the bank has to sell the house in foreclosure. If the house sells for more than the outstanding debt, the bank does not keep the excess; if it sells for less, the bank is often on the hook for the losses. Banks thus should dislike lending against houses that are harder to price. We find empirical evidence supporting this: houses which are less standardized, and thus harder to price, tend to have mortgages with higher interest rates, lower LTVs, and higher rates of mortgage rejections. Interestingly, poorer people also tend to live in houses which are harder to price; thus, the dispersion channel implies mortgage credit is most difficult to provide to households who need credit most to purchase houses.

In *Liquidity in Residential Real Estate Markets* (Kotova, Jiang & Zhang, 2024), which is R&R at the *Review of Financial Studies* , we analyze how housing market liquidity should be measured. We show that two commonly used liquidity measures, time-on-market (TOM) and price dispersion (PD), can be thought of as equilibrium outcomes from a "supply and demand system" for liquidity. An aggregate mass of sellers on the "demand side" attempts to purchase liquidity from an aggregate mass of buyers on the "supply side". The key tradeoff sellers face is that selling quickly is costly: if sellers choose to sell faster, prices decrease and price dispersion increases. When liquidity *supply* increases – that is, the mass of available buyers increases – the entire menu of available TOM and PD options available to sellers improves. Hence, liquidity supply increases will cause both price dispersion and time-on-market to decrease, so housing markets look more liquid along both measures. However, if liquidity *demand* increases – if sellers become in aggregate more impatient – sellers will shift along the menu of TOM-PD options, selling more quickly, but incurring lower and less stable prices as a result. In equilibrium, TOM will decrease, but prices may increase. This implies that looking at TOM or PD alone gives a misleading measure of housing market liquidity: a market which has low TOM because liquidity demand is high may in fact have lowered prices and higher price dispersion, since TOM is low because sellers are demanding a larger fraction of available liquidity, rather than because the supply of liquidity is greater. We find empirical support for our hypotheses.

5 Earlier Work

I have a number of earlier papers, broadly in the area of industrial organization.

5.1 Concentration in Product Markets

In *Concentration in Product Markets* (Benkard, Yurukoglu & Zhang, 2023), we analyze trends in concentration at the level of *product markets*. Industry concentration measures are a key input used in antitrust enforcement and a barometer that many economists employ for assessing the level of competition in a market. A key challenge in measuring concentration is market definition. Due to data availability, the vast majority of the prior literature has relied on *production-based* market definitions from the Census. We instead construct concentration measures that reflect *product markets* – markets for products which are substitutes to consumers – using a market research survey covering a broad set of consumer brands. Contrary to evidence from prior papers, we find concentration levels have been declining over the past 25 years, providing a counterpoint to the narrative that there has been a broad-based rise in concentration and thus market power across industries. This paper has been covered in The Times and the NBER Digest.

5.2 Bargaining in Used Car Markets

In *Quantifying Bargaining Power Under Incomplete Information: A Supply-Side Analysis of the Used-Car Industry* (Larsen & Zhang, 2021), which is R&R at the *Review of Economic Studies*, we develop a generalization of Nash bargaining weights to settings with two-sided asymmetric information; our weights can capture variation in both the division of total surplus, and the efficiency of trade. We also develop a identification and estimation approaches for these weights, using them to evaluate the efficiency of trade in supply-side used car markets. In *Intermediaries in Bargaining: Evidence from Business-to-Business Used-Car Inventory Negotiations* (Larsen, Lu & Zhang, 2021), which is R&R at the *Journal of Political Economy*, we show that mediators in used-car markets have a surprising amount of influence on trade outcomes. This is somewhat surprising, given the mediators play a pure communicative role, and do not have the ability to credibly vary the "sharp" incentives faced by buyers and sellers. We structurally estimate the extent to which mediators influence aggregate gains from trade.

5.3 Depreciating Licenses

In *Depreciating Licenses* (Weyl & Zhang, 2022), published in the *AEJ: Economic Policy* in 2022, I propose a novel "depreciating property rights" design for radio spectrum and other such resources. Optimal resource utilization has two important dimensions. Efficient use often requires users to make costly investments: to maximize the social value of a fishery, fishing firms must invest to prevent fishery pollution and preserve fish stocks. Efficient use also requires that the resource is allocated to the highest-value user: some fishing firms may have lower operating costs than others, and efficient use requires that fishing licenses can be quickly reallocated to low-cost firms when they enter the market. Resource designers face a tradeoff: no existing resource license design achieves full efficiency in both allocation and investment. Perpetual, long-term licenses encourage innovation, but distort reallocation, since low-value license holders have incentives to hold out for high values in secondary markets. Short-term licenses improve reallocation, but lower investment incentives.

We propose a novel license design, which we call a *depreciating license*. A depreciating license lasts forever, but decays over time: each period, a fraction of the license reverts to the resource administrator, and the license holder must repurchase this share to keep using the asset. This license is simple to use in practice – it can be implemented simply by running periodic auctions with slightly modified payment rules – and we show that the depreciating license trades off allocative efficiency and investment incentives more effectively than existing license designs.

We discuss this design in a policy piece, *Redesigning Spectrum Licenses* (Milgrom, Weyl & Zhang, 2017), published in *Regulation* in 2017. This paper has seen some interest from policymakers. OFCOM, the UK communications regulator, has discussed this proposal in a one-day conference, and also in a policy whitepaper.

References

BENKARD, C LANIER, ALI YURUKOGLU, AND ANTHONY LEE ZHANG (2023). "Concentration in product markets." National Bureau of Economic Research Working Paper.

EISENSCHMIDT, JENS, YIMING MA, AND ANTHONY LEE ZHANG (2024). "Monetary policy transmission in segmented markets." *Journal of Financial Economics*, 151: 103738.

HE, ZHIGUO, SCOTT NELSON, YANG SU, ANTHONY LEE ZHANG, AND FUDONG ZHANG (2023). "Zoning for Profits: How Public Finance Shapes Land Supply in China." National Bureau of Economic Research Working Paper.

HU, JUNYI AND ANTHONY LEE ZHANG (2023). "Competition in the Cryptocurrency Exchange Market." Available at SSRN 4472475.

KOTOVA, NADIA, ERICA XUEWEI JIANG, AND ANTHONY LEE ZHANG (2024). "Liquidity in residential real estate markets." R&R, *Review of Financial Studies*.

JANSEN, MARK, FABIAN NAGEL, CONSTANTINE YANNELIS, AND ANTHONY LEE ZHANG (2023). "Data and welfare in credit markets." R&R, *Journal of Financial Economics*.

JIANG, ERICA XUEWEI AND ANTHONY LEE ZHANG (2023). "Collateral value uncertainty and mortgage credit provision." R&R, *Journal of Financial Economics*.

JING, BAIYUN, SEONGJIN PARK, AND ANTHONY LEE ZHANG (2022). "The Credit Channel of Monetary Policy Transmission: Evidence From the Chonsei System." Available at SSRN 4019774.

LARSEN, BRADLEY AND ANTHONY LEE ZHANG (2021). "Quantifying bargaining power under incomplete information: A supply-side analysis of the used-car industry." Available at SSRN 3990290.

LARSEN, BRADLEY, CAROL HENGHENG LU, AND ANTHONY LEE ZHANG (2021). "Intermediaries in bargaining: Evidence from business-to-business used-car inventory negotiations." National Bureau of Economic Research.

MA, YIMING, YAO ZENG, AND ANTHONY LEE ZHANG (2023). "Stablecoin Runs and the Centralization of Arbitrage." Available at SSRN 4398546.

MILGROM, PAUL, E. GLEN WEYL, AND ANTHONY LEE ZHANG (2017). "Redesigning spectrum licenses." *Regulation*, 40: 22.

MILIONIS, JASON, CIAMAC C MOALLEMI, TIM ROUGHGARDEN, AND ANTHONY LEE ZHANG (2023). "Automated market making and loss-versus-rebalancing." arXiv preprint arXiv:2208.06046.

OH, SEBEOM, SAMUEL ROSEN, AND ANTHONY LEE ZHANG (2022). "Digital Veblen goods." Available at SSRN 4042901.

WEYL, E. GLEN AND ANTHONY LEE ZHANG (2022). "Depreciating licenses." *American Economic Journal: Economic Policy*, 14(3): 422–448.

Yannelis, Constantine and Anthony Lee Zhang (2023). "Competition and selection in credit markets." *Journal of Financial Economics*, 150(2): 103710.

Zhang, Anthony Lee (2022). "Competition and Manipulation in Derivative Contract Markets." *Journal of Financial Economics*, 144: 396-413.

Becker, Gary S. (1991). "A note on restaurant pricing and other examples of social influences on price." *Journal of Political Economy*, 99(5): 1109–1116.

Budish, Eric, Peter Cramton, and John Shim (2015). "The high-frequency trading arms race: Frequent batch auctions as a market design response." *The Quarterly Journal of Economics*, 130(4): 1547–1621.