INSTITUT POLYTECHNIQUE DE PARIS

NNT : 2024IPPAX021

Thèse de doctorat

Essays on Market Dynamics and Stability in Decentralized Finance

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'Ecole Polytechnique

École doctorale n°626 Dénomination (EDIPP)
Spécialité de doctorat : Sciences économiques

Thèse présentée et soutenue à Palaiseau, le 10 Juin 2024, par

MYRIAM KASSOUL

Composition du Jury :

Matthieu Bouvard Professeur, Toulouse School of Economics	Président
Hanna Halaburda Professeure Associée, New York University Stern School of Business	Rapporteur
Arash Aloosh Professeur Associé, Léonard de Vinci Business School	Examinateur
Julien Prat Professeur, CREST-Ecole Polytechnique	Directeur de thèse
Yackolley Amoussou-Genou Maître de conférence, Panthéon-Assas Université	Invité

To my wonderful parents, Mahfoud and Georgette,

To my cherished siblings, Abderrahmane and Ahlème,

And to my beloved husband, Anouar.

Aknowledgement

Embarking on a Ph.D. journey goes beyond simply gaining knowledge; it becomes a profound voyage of self-discovery intertwined with the highs and lows of academic exploration. Along this profound path, I've been fortunate to receive invaluable support from exceptional people.

First and foremost, I am sincerely grateful to my supervisor, Julien Prat, for his encouragement, patience, and the countless hours invested in refining my work. His constructive feedback has played a crucial role in shaping my development as a researcher. This dissertation stands not only as a culmination of my academic efforts but also as a testament to Julien's belief in my abilities and his commitment to fostering intellectual curiosity. Furthermore, I am immensely thankful to Julien for the opportunity he provided me to pursue my visiting period at NYU Stern school of business. Throughout my PhD journey, his guidance has been a constant driving force, always encouraging me to strive for excellence and pushing me to deliver my best.

I am also deeply grateful to the Blockchain@X center, CREST, and the ECODEC label for their financial support throughout my doctoral journey, making this thesis possible. Their investment in my academic endeavors has been invaluable, enabling me to pursue my scholarly goals and contribute to the field.

My gratitude also extends to my supportive team at Blockchain@X center. Natkamon, Yiyun, Michele, Daniel, and Yackolley, your fellowship and encouragement created a positive and motivating environment where we all cheered each other on. I extend my gratitude to Fatou and Lyza for our delightful coffee breaks and conversations. Meeting David during my Ph.D. journey has been an immense pleasure; his six months in our team exemplify the true meaning of camaraderie. Even after his departure, David has continued to be an unwavering pillar of support and friendship, transcending the bounds of a mere colleague to become a dear friend.

To Fanda, Jalel, Djamila, Ziad, Leyla, Murielle, and Teddy, I am immensely grateful for the wonderful moments we shared over lunch. Your support and encouragement have been invaluable to me, especially during those tough times. Thank you for being an integral part of my PhD experience and for the unwavering support and joy you brought into my life. You are more than friends; you are family, and I appreciate each of you for being there for me.

I also want to express my gratitude to my office mates - Guillaume, Morgane, Théo, and Elio - for the stimulating discussions, shared experiences, and moments of laughter that enriched our time together.

My heartfelt thanks go to Hanna Hallaburda for welcoming me at NYU Stern School of Business and to the entire blockchain team there for our fruitful discussions around blockchain topics. Additionally, I want to acknowledge Francesco Violante for generously taking the time to discuss and review my papers, providing valuable advice that has greatly contributed to my academic growth.

I am deeply thankful to my friends, Louisa and Clara, for their unwavering support and encouragement. Throughout this journey, our countless discussions, and our shared moments of joy have been a constant source of strength.

Last but certainly not least, I want to express my deepest gratitude to my parents, Mahfoud and Georgette, and my siblings, Abderrahmane and Ahlème, for their unwavering support since my childhood. You have consistently been a source of guidance and encouragement. Moreover, you have been exemplary role models, inspiring me to strive for excellence in all aspects of life. Your love, wisdom, and unwavering belief in me have been the cornerstone of my journey, and for that, I am eternally grateful. I also want to extend my deepest thanks to my husband, Anouar, who supported me even when we were just friends and continued to do so throughout our journey together. Through the ups and downs, you have stood by my side, believing in me and encouraging me to pursue my dreams. Your unwavering support has been a source of strength and inspiration, and I am truly grateful for your presence in my life.

Résumé

Alors que le monde de la finance continue d'évoluer rapidement, l'émergence des marchés décentralisés et de la technologie Blockchain a présenté à la fois des opportunités et des défis. Les marchés décentralisés, opérant de manière autonome et auto-régulée, introduisent un changement de paradigme qui exige une compréhension approfondie. Cette thèse se concentre sur la compréhension de ces nouveaux marchés et l'évaluation de leur stabilité, avec pour objectif plus large de contribuer à la stabilité de l'ensemble de l'écosystème financier.

Le premier chapitre, co-écrit avec Michele Fabi et Julien Prat, offre un aperçu complet sur Automated Market Makers (AMM) pour les échanges décentralisés (DEXs). L'ambition principale de la Finance Décentralisée (DeFi) est d'introduire des opérations financières sans tiers de confiance grâce à l'application de la technologie Blockchain et des contrats intelligents. Dans ce paysage émergent de la DeFi, les AMM servent de composant fondamental, employant des protocoles algorithmiques pour remplacer le besoin d'intermédiaires financiers traditionnels. Ce chapitre plonge dans les origines historiques des AMM, mettant en évidence leur évolution depuis leur création dans les marchés de prédiction jusqu'à leur utilisation moderne dans des environnements décentralisés. Notamment, nous explorons la transition des carnets d'ordres complexes vers des Constant Function Market Makers (CFMMs), qui déterminent les prix des actifs grâce à des fonctions de trading déterministes. Ce chapitre met également l'accent sur les fondements théoriques et les résultats empiriques concernant les CFMMs, éclairant leur capacité à reproduire les instruments financiers traditionnels et à optimiser le surplus des traders et des fournisseurs de liquidités. De plus, nous examinons la dynamique concurrentielle entre les plateformes d'échanges centralisées et décentralisées, en étudiant les incitations au trading et à la fourniture de liquidités. Enfin, ce chapitre examine de manière approfondie la question cruciale de la valeur maximale extractible (MEV) et son impact sur les DEXs, proposant des solutions envisageables et identifiant des voies de recherche prometteuses dans ce domaine en constante évolution.

Le deuxième chapitre explore la microstructure des DEXs, et plus particulièrement celle d'Uniswap. Lancée en 2018, Uniswap est rapidement devenue la plus grande plateforme d'échange décentralisée, permettant des échanges pair-à-pair de cryptomonnaies, de manière sécurisée et sans intermédiaire. Uniswap fonctionne sur la base de pools de liquidités automatisés et utilise des contrats intelligents pour faciliter les échanges entre les différentes cryptommonnaies. Le mécanisme de tarification, régi par la formule du produit constant, garantit que le produit des réserves reste inchangé avant et après un échange. Ce chapitre examine la capacité d'Uniswap à aligner ses prix avec ceux des marchés centralisés – consiférés comme marchés de références – échangeant les mêmes paires de crypto-actifs. En m'inspirant de la littérature sur la microstructure des marchés financiers traditionnels, je postule que les coûts d'inventaire peuvent influencer de

manière significative la dynamique et la précision des prix sur Uniswap. Pour évaluer l'impact de ces coûts sur la précision des prix d'Uniswap, je développe un modèle de microstructure puis le calibre en utilisant les données de prix de clôture par minute sur les marchés d'Uniswap et de Binance de mai 2020 à décembre 2022. L'analyse se concentre initialement sur la paire ETH-BTC, avant de s'étendre à d'autres paires. Mes résultats mettent en évidence l'impact significatif des coûts de détention d'inventaire, notamment lorsque les tailles de pool augmentent, entraînant une réactivité réduite des traders aux écarts de prix d'Uniswap par rapport aux prix des marchés de référence. Cependant, en examinant les paires stablecoin-stablecoin, le modèle révèle que les traders sont enclins à exploiter plus rapidement les opportunités d'arbitrage lorsqu'ils ne sont pas confrontés à des coûts d'inventaire. Cela conduit alors à l'alignement des prix d'Uniswap avec ceux des marchés centralisés.

Le troisième chapitre, co-écrit avec Natkamon Tovanich, Julien Prat et Simon Weidenholzer, propose une analyse complète de la contagion financière au sein de Compound V2, un protocole de prêt décentralisé sur la blockchain Ethereum. Nous proposons une méthodologie pour construire les bilans comptables des pools de liquidité de Compound et utilisons cette méthodologie pour caractériser son réseau financier complexe. Notre analyse révèle que la majorité des utilisateurs de ce réseau s'engagent principalement dans deux activités: l'emprunt de stablecoins et la participation à l'extraction de liquidité. Pour évaluer la robustesse du protocole, nous réalisons une série de tests de résistance, simulant des scénarios où le défaut d'un pool ou des chocs important du prix du Bitcoin et de l'Ether pourraient déclencher des défauts en cascade. Nos résultats suggèrent qu'une défaillance en cascade est une possibilité, mais qu'il nécessite des chocs de prix importants. De plus, les pools de stablecoins sont plus susceptibles de faire défaut, tandis que les pools de Bitcoin et d'Ether sont plus susceptibles d'initier une réaction en chaîne.

À travers ces trois chapitres, mon étude vise à approfondir notre connaissance de l'écosystème de la DeFi, en examinant sa stabilité, la microstructure du marché et les risques de contagion. En explorant ces questions de recherche cruciales, mon objectif est d'enrichir la discussion sur la DeFi, contribuant ainsi à promouvoir un écosystème financier plus résistant et pérenne.

Domaine: Economie

mots clés: Economie Digitale, Blockchain; Cryptomonnaie; Finance Décentralisée; FinTech; Microstructure de Marché, Risques Systémiques.

Summary

As the world of finance continues to evolve rapidly, the emergence of decentralized markets and Blockchain technology presents both opportunities and challenges. Decentralized markets, operating autonomously and self-regulated, introduce a paradigm shift that demands a thorough comprehension. This thesis focuses on understanding these new markets and assessing their stability, with a broader aim of contributing to the stability of the entire financial ecosystem.

In the first chapter, co-written with Michele Fabi and Julien Prat, we provide a comprehensive overview of Automated Market Makers (AMMs) for Decentralized Exchanges (DEXs). The main ambition of Decentralized Finance (DeFi) is to introduce trustless financial operations through the application of Blockchain technologies and smart contracts. Within this evolving DeFi landscape, AMMs serve as a fundamental component, employing algorithmic protocols to replace the need for traditional financial intermediaries. This chapter dives into the historical origins of AMMs, showcasing their evolution from their inception in prediction markets to their modernized use in decentralized environments. Notably, it explores the transition from complex Limit Order Books (LOBs) to the prevalent Constant Function Market Makers (CFMMs), which determine asset prices through deterministic trading functions. The literature review emphasizes the theoretical underpinnings and empirical findings regarding CFMMs, shedding light on their ability to replicate traditional financial instruments and optimize the surplus of traders and liquidity providers. Furthermore, it delves into the competitive dynamics between centralized and decentralized exchanges, addressing motives for trading and liquidity provision. Lastly, the chapter addresses the critical issue of Miner Extractible Value (MEV) and its significance for DEXs, presenting potential solutions.

The second chapter explores the microstructure of DEXs, with a particular emphasis on Uniswap. Uniswap has emerged as a prominent platform for direct peer-to-peer cryptocurrency exchanges, eliminating the need for intermediaries. Uniswap operates on the principle of CFMM, using liquidity pools rather than order books to facilitate trades. Within each pool, prices are algorithmically determined and self-regulated through a constant product formula, ensuring that the reserves' product remains consistent both prior to and following a transaction. This chapter examines Uniswap's ability to synchronize its quoted prices with those of centralized markets – considered as reference markets – trading the same asset pairs. Drawing inspiration from the market microstructure literature of traditional financial markets, I postulate that inventory holding costs can significantly influence Uniswap's price dynamics and accuracy. To assess the influence of these costs on Uniswap's price accuracy, I develop a microstructure model and calibrate it using 1-minute closing price data of Uniswap and Binance markets from May, 2020 to December, 2022. The analysis initially focuses on the ETH-BTC pair, before broadening the scope to include other pairs. My results highlight the

significant impact of inventory carrying costs, particularly as pool sizes expand, resulting in reduced trader responsiveness to deviations in Uniswap's prices from reference market prices. However, upon examining stablecoin-stablecoin pairs, the model reveals that traders are inclined to exploit arbitrage opportunities more promptly when they do not face inventory holding costs. Consequently, the price on Uniswap closely aligns with that on CEXs.

The third chapter, co-written with Natkamon Tovanich, Julien Prat and Simon Weidenholzer, provides a comprehensive analysis of financial contagion within Compound V2, a decentralized lending protocol on the Ethereum blockchain. We outline a methodology for constructing the balance sheets of Compound's liquidity pools and use this methodology to characterize its intricate financial network. Our analysis uncovers that the majority of users within this network primarily engage in two activities: borrowing stablecoins and participating in liquidity mining. To assess the robustness of the protocol, we conduct a series of stress tests, simulating scenarios where a pool's default or significant price shocks in Bitcoin and Ether could trigger cascading defaults. Our findings suggest that while a cascading failure is a possibility, it requires substantial price shocks. Notably, pools of stablecoins are more susceptible to default, while Bitcoin and Ether pools are more likely to initiate a chain reaction.

Through these three chapters, my research endeavors to advance our comprehension of the DeFi ecosystem, offering insights into its stability, market microstructure, and contagion risks. By addressing key research questions, I aim to contribute to the broader discourse on decentralized finance, ultimately fostering a more resilient and sustainable financial ecosystem.

Field: Economics

Key words: Digital Economics, Blockchain; Cryptocurrency; Decentralized Finance; FinTech; Market Microstructure, Systemic Risks.

Contents

List of Figures

List of Tables

Introduction générale

Cette thèse explore le domaine de la Finance Décentralisée (DeFi), examinant le paradigme des market makers automatisés (AMM), la microstructure du marché des échanges décentralisés (DEX) et les risques systémiques au sein des protocoles de prêt. Cette introduction générale familiarise les lecteurs avec la technologie de la blockchain et la DeFi, offrant un aperçu des différents chapitres et soulignant leur contribution collective à la compréhension et à l'avancement du domaine de la DeFi.

La crise financière mondiale de 2008 a été une rupture de confiance dans les institutions financières traditionnelles. Les institutions jugées "trop grosses pour faire faillite" se sont révélées vulnérables, pratiquant des méthodes risquées qui ont mis en péril la stabilité du système financier mondial. La crise a mis en évidence la défaillance de la surveillance centralisée et souligné le besoin de transparence, de responsabilité et de résilience dans l'infrastructure financière.

Dans ce climat de mécontentement et de désillusion vis-à-vis des systèmes centralisés, un individu, sous le nom de Satoshi Nakamoto, a publié un livre blanc intitulé "Bitcoin: Un Système de Trésorerie Électronique de Pair à Pair". Ce document a introduit un concept révolutionnaire : une monnaie numérique décentralisée, le Bitcoin, pouvant être échangée directement entre utilisateurs sans l'intervention d'une autorité centrale. Le réseau Bitcoin a été officiellement lancé le 3 janvier 2009, lorsque Nakamoto a miné le premier bloc, connu sous le nom de "bloc de genèse" ou "Bloc 0". La proposition de Nakamoto a remis en question les notions traditionnelles de monnaie et de systèmes financiers, offrant une vision d'un réseau monétaire plus inclusif et transparent. Ceci a été rendu possible grâce à la technologie de la blockchain.

Une blockchain est un registre distribué qui enregistre les transactions sur un réseau d'ordinateurs de manière sécurisée et transparente. Elle se compose d'une chaîne de blocs, où chaque bloc contient une liste de transactions. Ces blocs sont liés ensemble à l'aide de techniques de hachages cryptographiques, garantissant l'intégrité et l'immutabilité des données. Une fois qu'un bloc est ajouté à la chaîne, il devient impossible de modifier les informations qu'il contient.

Les principales caractéristiques de la blockchain incluent la décentralisation, la transparence, l'immutabilité et la sécurité. La décentralisation signifie qu'il n'y a pas de point de contrôle unique, des mécanismes de con-

sensus (comme la preuve de travail ou la preuve d'enjeu) sont utilisés pour valider et vérifier les transactions. Cela fait de la blockchain un système résistant à la manipulation et sans besoin de tiers de confiance. La transparence vient du fait que n'importe qui peut consulter les transactions enregistrées sur la blockchain. L'immutabilité garantit qu'une fois qu'une transaction est enregistrée, elle ne peut pas être modifiée ou supprimée. Et la sécurité est maintenue grâce à des techniques cryptographiques.

Alors que le Bitcoin a introduit le concept de monnaie numérique décentralisée et a popularisé la technologie de la Blockchain à l'échelle mondiale, c'est Ethereum qui a véritablement révolutionné l'espace. Proposé par Vitalik Buterin à la fin de 2013 et officiellement lancé en 2015, Ethereum a poussé le concept de la blockchain au-delà des transactions de pair à pair. Ethereum a introduit l'idée des contrats intelligents, des contrats qui s'executent automatiquement selon des termes prédéfinis par les parties et inscrits dans le code. Cela a permis aux développeurs de créer des applications décentralisées (DApps) capables d'automatiser et d'appliquer une large gamme d'accords et de fonctions, éliminant le besoin d'une autorité centrale. La blockchain d'Ethereum est devenue une plateforme programmable, offrant un écosystème plus polyvalent et étendu que la blockchain Bitcoin. Cette programmabilité a jeté les bases de la révolution de la DeFi, marquant un changement de paradigme vers un système financier ouvert, sans permission et transparent qui fonctionne sans avoir besoin d'autorités centrales.

Il existe plusieurs catégories de protocoles DeFi, chacune offrant des fonctionnalités spécifiques et des services différents:

- **Protocoles de prêt** : Les protocoles de prêt permettent aux utilisateurs de prêter leurs crypto-actifs pour percevoir des intérêts ou d'emprunter des crypto-actifs en fournissant des garanties, sans qu'il soit nécessaire de recourir à des intermédiaires.

- **Les échanges décentralisés (DEXs)** : Ces plateformes facilitent l'échange de crypto-actifs de pair à pair sans qu'une autorité centralisée ne soit nécessaire, offrant ainsi une plus grande autonomie et une meilleure sécurité aux traders.

- **Gestion d'actifs** : Les protocoles DeFi pour la gestion des actifs permettent aux utilisateurs d'investir, de diversifier et de gérer leurs portefeuilles de crypto-actifs de manière décentralisée, souvent par le biais de stratégies automatisées ou d'échanges algorithmiques.

- **Les Marchés de produits dérivés** : Ces plateformes proposent des échanges décentralisés de produits dérivés, y compris des options, des contrats à terme et d'autres instruments financiers, permettant aux utilisateurs de couvrir les risques ou de spéculer sur les mouvements de prix sans dépendre des intermédiaires financiers traditionnels.

- **Les Organisations autonomes décentralisées (DAO)** : Il s'agit d'entités régies par des contrats intelligents et fonctionnant de manière autonome sur des réseaux de blockchain, permettant une prise de décision et une allocation des ressources décentralisées au sein de l'écosystème DeFi.

- **Paiements et transferts de fonds** : Les solutions DeFi dans cette catégorie visent à fournir des services de paiement et de transfert de fonds efficaces, peu coûteux et sans frontières, en s'appuyant sur la technologie blockchain pour faciliter les transactions rapides et sécurisées à l'échelle mondiale.

La DeFi a connu une croissance exponentielle ces dernières années, devenant l'une des tendances les plus importantes de l'écosystème des cryptomonnaies et de la blockchain. A l'heure que j'écris ces lignes, la valeur totale bloquée (TVL) en DeFi s'élève à 94,276 milliards de dollars. Au sein de l'écosystème DeFi, les DEXs et les protocoles de prêt dominent, représentant respectivement $30,9\%$ et $17,1\%$ de la part de marché (voir Figure 1).

L'un des facteurs clés de la popularité de la DeFi est sa décentralisation. Contrairement aux systèmes financiers traditionnels où des intermédiaires tels que les banques, les chambres de compensation et les courtiers supervisent les transactions et la gestion des actifs, les protocoles DeFi fonctionnent sur des réseaux de blockchain, permettant des interactions de pair-à-pair sans besoin d'intermédiaires. Cela permet non seulement de réduire le risque de contrepartie, mais aussi d'améliorer la transparence et l'accessibilité, car toute personne disposant d'une connexion internet peut y participer. Cette inclusivité ouvre des opportunités financières aux individus qui sont mal desservis ou exclus du système bancaire traditionnel. De plus, la DeFi offre une plus grande transparence par rapport à la finance traditionnelle. Les transactions sur les réseaux blockchain sont immuables et vérifiables publiquement, ce qui permet aux utilisateurs d'auditer les transactions et de suivre les fonds en temps réel. Un autre aspect important du DeFi est sa composabilité. Les protocoles de DeFi sont conçus pour être interopérables, ce qui signifie que différents protocoles peuvent être combinés ou empilés pour créer de nouveaux produits et services financiers. Cette composabilité favorise l'innovation et permet aux développeurs d'itérer et d'expérimenter rapidement de nouveaux produits financiers. Les protocoles de DeFi offrent également aux détenteurs de cryptomonnaies la possibilité de maximiser leurs rendements en fournissant des liquidités à certains protocoles DeFi, souvent sous forme de prêts ou de fourniture de liquidité dans des pools de liquidités.

Bien que la DeFi offre de nombreuses opportunités, elle représente un tout nouveau paradigme financier, doté de ses propres règles et fonctionnant de manière autonome. Par conséquent, il est impératif de mener une analyse approfondie de la microstructure des marchés de ces plateformes. Il est essentiel de comprendre leur fonctionnement dans des conditions variées et d'évaluer les risques systémiques qui existent au sein de ces marchés. Sans une surveillance et une évaluation minutieuses, le potentiel attrayant de la DeFi

pourrait être obscurci par des conséquences et des vulnérabilités imprévues. Ainsi, l'objectif de cette thèse est d'approfondir la compréhension de ces marchés émergents et de contribuer à cette littérature en plein essor sur la DeFi.

Le premier chapitre, intitulé *SoK : Constant Function Market Makers*, offre une analyse exhaustive des teneurs de marché automatisés – *Automated Market Makers* (AMM) – dans le contexte des DEXs. Dans le paysage en constante évolution de la DeFi, les AMM occupent une place centrale, en exploitant des protocoles algorithmiques pour contourner les intermédiaires financiers traditionnels. Ce chapitre explore l'évolution historique des AMM, retraçant leur évolution depuis leur création sur les marchés de prédiction jusqu'à leur application actuelle dans des environnements décentralisés. Il met en lumière les travaux théoriques et empiriques concernant les Constant Function Market Makers (CFMM), mettant en évidence leur capacité à reproduire les instruments financiers traditionnels et à améliorer les avantages pour les traders et les fournisseurs de liquidité. Enfin, ce chapitre aborde la question cruciale de la valeur extractible par les mineurs (MEV) et ses implications pour les DEX, en proposant des solutions potentielles et en évoquant des directions prometteuses pour de futures explorations dans ce domaine en constante évolution.

Dans le deuxième chapitre, nous nous concentrons sur le CFMM le plus populaire dans l'écosystème DeFi : Uniswap. Uniswap est la plus grande plateforme d'échange décentralisée, facilitant les échanges directes de crypto-monnaies de pair à pair, évitant ainsi le recours à des intermédiaires. Il est composé de plusieurs pools de liquidités, chacun correspondant à une paire de crypto-actifs. Le prix auquel les deux actifs peuvent être échangés est fixé algorithmiquement selon la formule du produit constant. Selon cette formule, après chaque échange, le produit des nouvelles réserves doit demeurer égal au produit initial. La précision des prix d'Uniswap repose sur la condition de non-arbitrage. Les prix des DEX devraient toujours tendre à s'aligner avec ceux des CEX – considérés comme les prix de référence – car s'il y a une déviation, les traders peuvent réaliser un profit en achetant un actif sur le marché d'échange où le prix est bas et en le vendant sur le marché où le prix est élevé. Cependant, la mise en œuvre de ces stratégies d'arbitrage peut contraindre les traders à ajuster leurs inventaires, les éloignant ainsi de leurs positions optimales. De tels ajustements peuvent exposer les arbitrageurs à des risques et à des coûts qui peuvent impacter leur rentabilité globale. Cette situation est d'autant plus marquée lorsque la taille du pool augmente, réduisant ainsi l'impact des transactions sur les prix et entraînant une augmentation significative des coûts d'inventaire. Si l'impact des coûts des stocks sur les prix de transaction a été largement étudié dans la littérature traditionnelle sur la microstructure des marchés, à ma connaissance, cette question n'a pas encore été explorée dans le contexte de la DeFi. Ce chapitre développe un modèle de microstructure pour examiner l'impact des coûts d'inventaire sur la précision des prix au sein d'Uniswap. Le modèle est ensuite testé à l'aide de données de prix de clôture par minute sur les marchés Uniswap et Binance entre mai 2020 et décembre 2022, en se concentrant

principalement sur la paire ETH-BTC. Les résultats confirment le modèle théorique, indiquant qu'au fur et à mesure que le pool s'élargit, les traders deviennent moins réactifs aux opportunités d'arbitrage en raison des risques d'inventaire accrus. Des volumes de transactions plus importants sont nécessaires pour influer sur les prix Uniswap, ce qui nécessite des stocks d'inventaire plus importants et expose par conséquent les traders à des risques de marché plus élevés. En étendant l'analyse à d'autres paires, il devient évident que les traders opérant sur des paires de stablecoins, qui sont généralement indexées sur le dollar américain, présentent généralement peu d'exposition aux risques de marché. Dans ce cas, les traders sont enclins à exploiter plus facilement les opportunités d'arbitrage, ce qui conduit à une convergence des prix entre Uniswap et les marchés centralisés (CEX). Dans ce chapitre, je vise à offrir une compréhension plus complète de la microstructure du marché sur Uniswap et de ses conséquences potentielles pour les applications DeFi, afin d'améliorer l'efficacité et la stabilité de l'écosystème DeFi.

Enfin, le troisième chapitre de cette thèse s'intéresse aux protocoles de prêt, en particulier à Compound V2. S'il est largement admis que les systèmes financiers traditionnels sont vulnérables à la contagion par divers canaux, tels que les ruées bancaires et les cascades de défaillances, les risques de contagion potentiellement présents dans les protocoles DeFi restent relativement inexplorés. Dans ce contexte, le troisième chapitre se concentre sur l'évaluation des risques systémiques dans Compound. Le protocole de Compound gère plusieurs pools de liquidités, chacun étant dédié à des jetons spécifiques. Dans ce contexte, les fournisseurs de liquidités ajoutent des fonds aux pools, tandis que les emprunteurs retirent des fonds en déposant des crypto-actifs en garantie dans d'autres pools. Ces interactions établissent un réseau de passifs financiers interconnectant les différents pools de liquidités. Dans ce chapitre, nous décrivons une méthodologie pour construire les bilans des pools de liquidité de Compound et utilisons cette méthodologie pour caractériser son réseau financier complexe. Notre analyse révèle que les utilisateurs empruntent principalement des stablecoins et participent à des activités de minage de liquidité (liquidity mining) pour gagner des jetons de gouvernance de Compound. Dans un second temps, nous évaluons la résilience du protocole en nous inspirant de la littérature contemporaine sur la contagion financière. Nous étudions la façon dont les chocs se propagent dans le réseau financier de Compound via une série de tests de résistance. Nos résultats suggèrent que si les défaillances en cascade restent possibles, elles nécessitent des chocs de prix importants. Notamment, les pools contenant des stablecoins sont les plus susceptibles de faire défaut, tandis que ceux impliquant des Bitcoins et des Ethers sont plus susceptibles de déclencher un effet domino. Cette étude vise à combler un vide majeur dans la littérature sur les risques de contagion au sein des protocoles DeFi. Alors que les recherches existantes se concentrent sur la décentralisation et l'interconnexion des divers protocoles DeFi, notre étude se distingue en étant la première à analyser les risques de contagion et les effets de réseau au sein des protocoles de prêt. Ce faisant, nous apportons une contribution unique au champ de la recherche en

évolution constante sur la DeFi. De plus, notre approche s'aligne sur les recherches menées dans le contexte plus général de la contagion financière, en s'appuyant sur des études approfondies sur la transmission des chocs et de la détresse à travers les marchés financiers.

General Introduction

This dissertation delves into the realm of Decentralized Finance (DeFi), examining the automated market maker (AMM) paradigm, the market microstructure of decentralized exchanges (DEXs), and the systemic risks within lending protocols. This introduction familiarizes readers with Blockchain Technology and DeFi. It provides an overview of the subsequent chapters and highlights their collective contribution to understanding and advancing the field of DeFi.

The 2008 global financial crisis was a breakdown of trust in traditional financial institutions. Institutions deemed "too big to fail" were revealed to be vulnerable, engaging in risky practices that jeopardized the stability of the global financial system. The crisis exposed the failure of centralized oversight and underscored the need for transparency, accountability, and resilience in financial infrastructure.

In this climate of discontent and disillusionment with centralized systems, an individual, under the name of Satoshi Nakamoto, published a white paper titled "Bitcoin: A Peer-to-Peer Electronic Cash System". This document introduced a revolutionary concept: a decentralized digital currency, the Bitcoin, that could be exchanged directly between users without the involvement of any central authority. The Bitcoin network was officially established on January 3, 2009, when Nakamoto mined the first block, known as the "genesis block" or "Block 0". Nakamoto's proposal challenged traditional notions of currency and financial systems, offering a vision of a more inclusive and transparent monetary network. This has been made possible through the foundational technology of blockchain.

A blockchain is a distributed ledger that records transactions across a network of computers in a secure and transparent manner. It consists of a chain of blocks, where each block contains a list of transactions. These blocks are linked together using cryptographic hashes, ensuring the integrity and immutability of the data. Once a block is added to the chain, any information within it cannot be modified.

The key features of blockchain include decentralization, transparency, immutability, and security. Decentralization means there is no single point of control, consensus mechanisms (like proof-of-work or proof-of-stake) are used to validate and verify transactions. This makes blockchain a tamper-resistant and trustless system. Transparency comes from the fact that anyone can view the transactions recorded on the blockchain.

Immutability ensures that once a transaction is recorded, it cannot be altered or deleted. And security is maintained through cryptographic techniques.

While Bitcoin pioneered the concept of decentralized digital currency and introduced the world to blockchain technology, it was Ethereum that truly revolutionized the space. Proposed by Vitalik Buterin in late 2013 and officially launched in 2015, Ethereum took the concept of blockchain beyond peer-to-peer transactions. Ethereum's blockchain became a programmable platform, offering a more versatile and expansive ecosystem compared to Bitcoin blockchain. Ethereum introduced the idea of smart contracts, self-executing contracts with the terms of the agreement directly written into code. This allowed developers to create decentralized applications (DApps) that could automate and enforce a wide range of agreements and functions without the need for a central authority. This programmability laid the foundation for the DeFi revolution, marking a paradigm shift towards an open, permissionless, and transparent financial system that operates without the need for central authorities.

DeFi protocols be categorized into several distinct areas, each offering unique financial services and solutions:

- **Lending protocols**: Lending protocols enable users to lend their digital assets to earn interest or borrow assets by providing collateral, without the need for intermediaries.
- **Decentralized Exchanges (DEXs)**: These platforms facilitate the peer-to-peer exchange of digital assets without the need for a centralized authority, providing greater autonomy and security for traders.
- **Asset Management**: DeFi protocols for asset management allow users to invest, diversify, and manage their digital asset portfolios in a decentralized manner, often through automated strategies or algorithmic trading.
- **Derivatives Markets**: These platforms offer decentralized derivatives trading, including options, futures, and other financial instruments, allowing users to hedge risk or speculate on price movements without relying on traditional financial intermediaries.
- **Decentralized Autonomous Organizations (DAOs)**: These are entities governed by smart contracts and run autonomously on blockchain networks, enabling decentralized decision-making and resource allocation within the DeFi ecosystem.
- **Payments and Remittances**: DeFi solutions in this category aim to provide efficient, low-cost, and borderless payment and remittance services, leveraging blockchain technology to facilitate fast and secure transactions globally.

DeFi has experienced an explosive surge in popularity over the past few years, emerging as one of the most significant trends within the cryptocurrency and blockchain space. At the time of writing, the Total Value Locked (TVL) in DeFi stands at $94.276 billion. Within the DeFi ecosystem, DEXs and lending protocols dominate, accounting for 30.9% and 17.1% of the market share, respectively (see Figure 1).

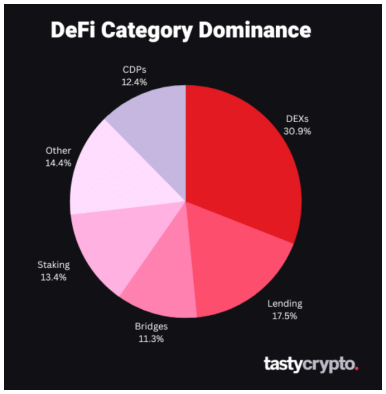

Figure 1: Total Value Locked (TVL) by category (%) in March 2024.
Source: https://www.tastycrypto.com/blog/top-defi-tokens.

One of the key factors driving the popularity of DeFi is its decentralization. Unlike traditional financial systems where intermediaries such as banks, clearinghouses, and brokers oversee transactions and asset management, DeFi protocols operate on blockchain networks, enabling peer-to-peer interactions without the need for intermediaries. This not only reduces counterparty risk but also enhances transparency and accessibility, as anyone with an internet connection can participate in DeFi activities. This inclusivity opens up financial opportunities to individuals who are underserved or excluded from the traditional banking system. Moreover, DeFi offers greater transparency compared to traditional finance. Transactions on blockchain networks are immutable and publicly verifiable, allowing users to audit transactions and track funds in real-time. Another significant aspect of DeFi is its composability. DeFi protocols are designed to be interoperable, meaning that different protocols can be combined or stacked together to create new financial products and services. This composability fosters innovation and allows developers to rapidly iterate and experiment with new financial products. DeFi protocols also offer cryptocurrency holders the opportunity to maximize their returns on their crypto assets by providing liquidity to decentralized exchanges or participating in liquidity pools.

While DeFi presents numerous opportunities, it is a completely new financial paradigm with its own set of rules and self-regulated. Therefore, it is imperative to thoroughly analyze the market microstructure of these platforms. Understanding how their properties function under various conditions and assessing the systemic risk within these markets is crucial. Without careful examination and oversight, the allure of DeFi's potential could be overshadowed by unforeseen consequences and vulnerabilities. Therefore, this thesis endeavors to attain a deeper comprehension of these emerging markets and to enrich the burgeoning

literature on DeFi.

The first chapter, entitled *SoK: Constant Function Market Makers*, provides a comprehensive examination of Automated Market Makers (AMMs) within the realm of DEXs. In the dynamic landscape of DeFi, AMMs play a pivotal role, leveraging algorithmic protocols to supplant traditional financial intermediaries. This chapter delves into the historical genesis of AMMs, tracing their evolution from their inception in prediction markets to their contemporary application in decentralized environments. The chapter underscores the theoretical foundations and empirical insights concerning CFMMs, elucidating their capacity to replicate traditional financial instruments and enhance the benefits for traders and liquidity providers. Moreover, it explores the competitive dynamics between centralized and decentralized exchanges, addressing the motivations for trading and liquidity provision. Lastly, the paper addresses the crucial issue of Miner Extractible Value (MEV) and its implications for DEXs, proposing potential solutions and outlining promising avenues for further exploration in this evolving domain.

In the second chapter, *The Market Microstructure of Uniswap*, our focus shifts towards the most popular CFMM in the DeFi ecosystem: Uniswap. Uniswap is the largest decentralized exchange, facilitating direct peer-to-peer cryptocurrency transactions, effectively bypassing the need for intermediaries. It is composed of multiple liquidity pools, each corresponding to a pair of crypto assets. The price at which the two assets can be exchanged is set algorithmically according to the constant product formula. This formula states that once an exchange is done, the product of the new reserves should remain at its pre-trade value. The accuracy of Uniswap's prices hinges on the fulfillment of the no-arbitrage condition. The prices on DEXs should always converge with those on CEXs – considered as the reference market prices – because any deviation presents an opportunity for traders to capitalize on price differentials by buying assets where prices are lower and selling them where prices are higher. However, executing such arbitrage strategies may require traders to adjust their inventory levels, potentially deviating from their desired positions. Such discrepancies can expose arbitrageurs to risks and expenses that impact their overall profitability. This effect becomes especially pronounced as the pool size increases, leading to a reduction in the price impact of trades and a significant increase in inventory costs. While the impact of inventory costs on transaction prices has been extensively explored in traditional market microstructure literature, to the best of my knowledge, there has been no prior attempt to investigate this phenomenon within the context of DeFi. This chapter develops a microstructure model to examine the influence of inventory holding costs on price accuracy within Uniswap. The model is then estimated using 1-minute closing price data from Uniswap and Binance markets spanning from May 2020 to December 2022, primarily focuses on the ETH-BTC pair. The findings support the theoretical model, indicating that as the trading pool expands, traders become less responsive to arbitrage opportunities across markets due to heightened inventory risks. Larger trades are required to impact Uniswap

prices, necessitating increased inventory holdings and consequently exposing traders to greater market risks. When broadening the scope of analysis to include other pairs, it becomes evident that traders operating in markets with stablecoin pairs, which are pegged to the US dollar, typically encounter minimal exposure to market risks. In such cases, traders are inclined to exploit arbitrage opportunities more readily, leading to a convergence of prices between Uniswap and centralized exchanges (CEXs). In this chapter, I aim to offer a more comprehensive understanding of the market microstructure on Uniswap, and its potential consequences for DeFi applications, ultimately enhancing the efficiency and stability of the DeFi ecosystem.

Finally, the third chapter, entitled *Contagion in Decentralized Lending Protocols: A Case Study of Compound*, directs its attention towards lending protocols, evaluating systemic risks within Compound V2. While it is widely acknowledged that traditional financial systems are vulnerable to contagion through various channels,such as bank runs and default cascades, the contagion risks potentially present in DeFi protocols remain relatively unexplored. Compound protocol manages multiple liquidity pools, each dedicated to specific tokens. Within this framework, lenders contribute liquidity to pools, while borrowers withdraw liquidity by pledging collateral in other pools. These interactions establish a network of financial liabilities interconnecting the various liquidity pools. In this chapter, we outline a methodology for constructing the balance sheets of Compound's liquidity pools and use this methodology to characterize its intricate financial network. Our analysis reveals that users predominantly engage in borrowing stablecoins and participating in liquidity mining activities to earn Compound's governance token. Subsequently, we evaluate the protocol's resilience by drawing inspiration from contemporary literature on financial contagion We explore how shocks propagate through Compound's financial network via a series of stress tests. Our findings suggest that while cascading failure remains a possibility, it necessitates substantial price shocks. Notably, pools containing stablecoins are most susceptible to default, while those involving Bitcoins and Ethers are more likely to trigger a domino effect. This research addresses a critical gap in the literature regarding the potential for contagion risks within DeFi protocols. While existing studies explore the decentralization and interconnectedness of various DeFi protocols, our paper pioneers the investigation of contagion risks and network effects within lending protocols, adding a unique dimension to the evolving field of DeFi research. Additionally, our approach aligns with research in the broader context of financial contagion, drawing on extensive studies on the transmission of shocks and distress across financial markets.

Chapter 1

SoK: Constant Function Market Makers[1]

Co-written with Michele Fabi and Julien Prat.

Abstract

We provide an overview of the academic literature on Automated Market Makers for Decentralized Exchanges. Our review puts an emphasis on contributions from researchers in economics and finance. We cover papers that study the optimal design of Automated Market Makers. Then we discuss models that leverage the insights from the literature on two-sided markets to characterize the equilibrium size of liquidity pools and the incentives of liquidity providers. Finally, we review recent research on the interactions between Miner Extractible Value and Decentralized Exchanges.

Keywords: *constant function market makers, automated market makers, decentralized exchanges, miner extractable value.*

JEL Classification: *G14, D82*

[1]Forthcoming: Chapter 10 in "A Companion to Decentralized Finance, Digital Assets, and Blockchain Technologies" edited by Edward Elgar Publishing Ltd.

1.1 Introduction

The advent of blockchain technologies has given rise to a fast-growing stream of financial innovations gathered under the term Decentralized Finance (DeFi). They rely on blockchains and smart contracts to provide decentralized implementations of financial services. The main promise of DeFi is to replace human interventions, and all the moral hazard that they entail, with algorithmic protocols. At its most ambitious, DeFi attempts to do away with the trust in financial intermediaries upon which the legacy system is built.

One the most important building blocks of DeFi are Automated Market Makers (AMMs), algorithms that pool resources to run decentralized exchanges. Although the first AMMs were invented in the 80's, long before the popularization of blockchains, they are currently receiving much attention as their modern upgrades are particularly suited to running decentralized exchanges. Due to their growing adoption, AMMs are also becoming the focus of renewed academic scrutiny.

This chapter provides an overview of the emerging literature on AMMs. It looks at the history of automated market-making to provide some perspective, and then reviews recent innovations. AMMs were first introduced in prediction markets to substitute human intermediaries with algorithmic rules. Their aim was to elicit information about the beliefs of participants, rewarding them for making correct predictions. We explain how researchers were able to derive clear connections between the AMM prices, the cost function of the market maker, and risk-neutral security pricing.

AMMs for prediction markets were theoretically motivated, their practical implementation being grounded in explicit principles. By contrast, AMMs for DeFi were introduced by practitioners who adopted heuristic solutions to solve the challenges of decentralized market-making. The new type of AMMs that have nowadays almost completely superseded limit order books (LOBs) in decentralized environments are Constant Function Market Makers (CFMMs). Instead of relying on complex matching mechanisms, CFMMs determine asset prices using a deterministic trading function of their inventory. Liquidity providers adjust the level of this trading function by supplying or withdrawing liquidity. Traders swap their assets with the CFMM in a way that keeps the trading function constant.

We present the burgeoning literature that studies the theoretical properties of CFMMs as well as several empirical papers testing its predictions. Our review illustrates that, in spite of their simplistic approach, CFMMs possess several desirable features. First, they satisfy the oracle property as their equilibrium in the absence of arbitrage opportunities tracks market prices. Also, under a set of reasonable design restrictions, CFMMs can be fine-tuned so as to replicate the payoffs that liquidity providers would obtain from investing in traditional financial instruments. The joint surplus of traders and liquidity providers can be optimized by adjusting the curvature of the trading function: Curved functions are more suited to pools of volatile

assets, whereas flatter functions are more suited to stable, mean-reverting assets such as stablecoins. We explain how these findings follow from standard microeconomic principles since the problem of arbitrageurs is isomorphic to the derivation of the compensated demand function in consumer theory (Fabi and Prat, 2022).

After having discussed the optimal design of CFMMs, we survey the literature that studies the competition between centralized and decentralized exchanges. Decentralized protocols have grown exponentially over the last years, becoming serious competitors to centralized exchanges that operate mostly through limit order books. A recent line of research investigates how traders and liquidity providers allocate their assets across these two types of exchanges. It builds on the market microstructure literature and identifies several motives and counter-motives for liquidity provision and trade. Trade is motivated either by exogenous motives, resulting in noise trading, or by arbitrage. Liquidity provision is motivated by returns from fees and capital gain on reserves, but discouraged by adverse selection costs, intrinsic to all types of AMMs. Adverse selection arises because providers lock their funds in the AMM, which prevent them from swiftly reacting to price shocks. Adverse selection costs depend on the relative amount of informed to noise trading. In Centralized Exchanges (CEXes), adverse selection is measured by the signal-to-noise ratio. In Decentralized Exchanges (DEXs), adverse selection is measured by the impermanent loss, a function of the curvature of the trading function. The literature highlights that competition among trading platforms leads to a stable equilibrium distribution of liquidity and explains empirical distribution of liquidity in popular DEXs such as Uniswap.

The last branch of literature addresses issues related to privacy and, more specifically, to Miner Extractible Value (MEV), a fundamental challenge for all trades processed on a public blockchain. MEV is the profit that malicious bots can make by defrauding users using insider information. Doing so on a blockchain is particularly easy since transactions are stored in a public memory pool. We present the literature documenting the significance of MEV for popular DEXs and the solutions that have been proposed to alleviate this problem. The last section concludes by discussing promising directions for further research.

1.2 AMMs For Prediction Markets

AMMs predate the advent of decentralized finance. They were first introduced to streamline prediction markets, i.e. securities markets whose primary objective is to aggregate information about probabilities of future events, such as election results or weather forecasts.[2] Algorithmic solutions are deemed promising because they have the potential to alleviate the thin market problem that often beleaguers prediction markets.

[2]Prediction markets are also called information markets. Successful examples include Iowa electronic markets and the Hollywood stock exchange.

The traditional way to foster market liquidity is to rely on the intermediation of professional market makers who act as counterparties to buy and sell orders. AMMs raise the possibility of replacing human market makers with automated algorithms. In order to provide a workable solution, the algorithm has to solve the two following challenges. First, information must be elicited. Traders who interact with the AMM should find it optimal to reveal their beliefs about the probability distribution of future events. Second, the algorithm has to aggregate opinions by producing an estimate which summarizes the information revealed by all the previous trades.

1.2.1 Information elicitation

The information elicitation problem was analyzed long before the creation of prediction markets. Consider an agent who entertains the belief p about the distribution of the random variable X. The elicitation mechanism asks the agent to report a distribution r based on the knowledge that she will receive the reward $c_i = s(r, X_i)$ if X_i turns out to be the true state of X. The design problem consists in finding a scoring rule $s(r, X_i)$ such that

$$p = arg \max_r \mathbb{E}_p[s(r, X_i)], \tag{1.1}$$

where E_p is the expectation operator associated to the agent's belief p. A rule $s(r, X_i)$ that satisfies eq. 1.1 is called a proper scoring rule since the best response of the agent is to truthfully report her belief p.

A naive approach would consist in replacing the score with the reported probability, so that $s(r, X_i) = r_i$. However, this rule does not induce truthful reporting. It is not a proper scoring rule because its linearity induces the forecaster to place all probability weights on the events that she deems the most likely to occur.

Truthful reporting therefore excludes linear rules. But does it identify a unique solution? Unfortunately no. The quadratic scoring rule proposed by Brier (1950) was quickly followed by Good (1952)'s introduction of the logarithmic scoring rule. These two seminal contributions spurred a research agenda whose objective was to provide additional criteria for the selection of scoring rules. For instance, Winkler (1969) showed that Good's logarithmic rule is the only rewarding scheme which allows the designer to use standard likelihood methods for the evaluation of the forecaster's performance. As we will show below, the need to come up with additional criteria to narrow down the design space remains one of the main motivations for ongoing research on AMMs.

1.2.2 Information aggregation

Since proper scoring rules elicit beliefs, why shouldn't we harvest the wisdom of crowds by combining beliefs elicited from multiple forecasters? The answer from the literature on opinions pooling is full of caveats

(see, for instance, Genest and Zidek, 1986), showing that more often than not, the pooled distribution is indistinguishable from one of the individual distributions.

The difficulty to pool opinions explains why it remains common for forecasts to rely on the advice of a single expert. Prediction markets, on the other hand, do not shy away from the information aggregation challenge. Instead of trying to compute one summary statistics out of disparate beliefs, they leverage the economic incentives of forecasters. Betting markets indicate that individuals are willing to place wagers on future events and that they use each other's bets to update their priors, resulting in reliable probability estimates (Hausch, 1994).

A market structure similar to that of betting markets can be algorithmically emulated. As originally proposed by Hanson (2003, 2007), the first step consists in using scoring rules sequentially so as to construct a market scoring rule. The state of the market scoring rule is equal to the report r of the last person that has interacted with the rule. The state is public information and anyone can update it by submitting a new report. Once the actual value X_i is revealed, the payout Π to the agent that made the n-th update is equal to $\Pi(r^n, X_i|r^{n-1}) = s(r^n, X_i) - s(r^{n-1}, X_i)$. Given that agents cannot modify the previous report r^{n-1}, maximizing one's expected payoff amounts to maximizing the expected value of $s(r^n, X_i)$, which is equivalent to solving eq. 1.1. Hence, when the scoring rule s is proper, agents whose beliefs differ from the last report are incentivized to update it by submitting their own truthful report.

Although market scoring rules achieve our goal of sequentially eliciting truthful forecasts, their implementation is not very user-friendly. Reports are not standard bets as agents have to submit probability distributions.[3] A more natural implementation would allow traders to buy and sell securities whose payoffs are attached to the realization of specific outcomes. This is where the introduction of cost functions becomes useful. More specifically:

1. Let N denote the number of mutually exclusive and exhaustive outcomes of the random variable X to be predicted. The AMM offers to trade N distinct securities, each paying \$1 per share if its corresponding outcome is realized.

2. The AMM cost function $C(q)$ returns the total amount of money spent by traders as a function of the vector q that keeps track of the overall number of shares held for each of the N securities. The AMM quotes a marginal price $p_i(q)$ for security i that is equal to $\partial C(q)/\partial q_i$.

3. Anyone can at any time buy or sell shares. The AMM charges $C(q') - C(q)$ for a trade that changes the number of outstanding shares from q to q'.

[3] Another implementation challenge arises from the fact that payoffs are negative in some states of the world. The market scoring rule therefore has to be complemented with an escrow mechanism requesting from each forecaster that they deposit an amount equal to their losses in the worst-case scenario, i.e. $min_{X_i}\Pi(r^n, X_i|r^{n-1})$.

4. Once the outcome is realized, the traders that bought the winning security receive \$1 per share from the AMM, while those that did not receive nothing.

The connection between cost functions and market scoring rules is not obvious. Yet, Chen and Pennock (2012) shows that the cost function can always be chosen so as to replicate the incentives of participants interacting with a market scoring rule. The equivalence holds under the additional constraint that the vector of marginal prices belongs to its probability simplex, i.e. $\sum_{i=1}^{N} p_i(q) = 1$ for all $q \in R_+^N$. This restriction ensures that the price of a security indicates the probability of its predicted outcome.

The most popular cost function is the one associated to the logarithmic market scoring rule $s(r, X_i) = b \cdot log(r_i)$. It reads $C(q) = b \cdot log(\sum_{i=1}^{N} \exp(\frac{q_i}{b}))$. The coefficient $b > 0$ is called the liquidity parameter because it controls the elasticity of the price response of the AMM: The larger b is, the less reactive prices are, and so, the more liquid the AMM is. Then, given that the owner of the AMM chooses the value of b, shouldn't she set it as high as possible? The reason why one should refrain from doing so is that the liquidity parameter also controls the amount of money that the AMM can lose. A fundamental property of market scoring rules is that, although running an AMM might be costly, the losses can usually be bounded ex-ante. For the logarithmic market scoring rule, the worst-case loss, $b \cdot log(N)$, is linearly increasing in b.[4] Hence the owner of the AMM faces a fundamental trade-off between risk exposure and market liquidity.

Another result in Chen and Pennock (2012) that yields intriguing connections between prediction markets and DeFi is their construction of utility-based market makers. Consider a market maker with utility function U for money and subjective prior about the distribution of the random variable X. A constant utility cost function charges users the exact amount that keeps the market maker indifferent between accepting and declining the trade. Accordingly, "the market maker starts the market with some initial expected utility and then keeps this expected utility level during the whole process of trading" (Chen and Pennock, 2012, p.51). It turns out that a broad class of market scoring rules[5] satisfies this constant utility requirement when the market maker displays hyperbolic absolute risk aversion. This equivalence result provides a microfoundation for the parametric specification of the AMM, making it possible to directly import the market maker's priors into prices. We will see below that obtaining similar results for CFMMs is one of the main open questions of DeFi.

[4] Note that the worst-case loss is independent of the number of forecasts because the AMM only has to pay the last trader. This result immediately follows from the definition of the AMM's overall payout: $\sum_{n=1}^{T} s(r^n, X_i) - s(r^{n-1}, X_i) = s(r^T, X_i) - s(r^0, X_i)$, where T is the index of the last trader.

[5] More precisely, weighted pseudospherical scoring rules, as defined by Jose et al. (2008), which include logarithmic scoring rules as one of their limit cases.

1.3 CFMMs and Decentralized Finance

The previous section illustrates how research on AMMs for prediction markets progressed in a systematic way. Starting from a well-defined problem, researchers used scoring rules as building blocks to clarify and streamline the design of AMMs. The emergence of AMMs in decentralized finance followed a completely different process. To fully grasp the motivation behind the introduction of CFMMs, it is instructive to go back to their genesis in a proposal first floated by Buterin (2016). The rationale for introducing CFMMs was not to tackle a clear conceptual problem but rather to expediently solve some practical challenges posed by market-marking on blockchains that limit-order books were not able to address.

1.3.1 Motivation for CFMMs

Blockchains have many promises but, as of today, most applications are centered on DeFi. Their fundamental benefit is to algorithmically handle the settlement and enforcement of contracts, thereby avoiding the risks and costs associated to financial intermediation. Decentralized exchanges such as Uniswap, Sushiswap and Curve constitute the core of DeFi. They leverage blockchain infrastructures to allow their users to swap digital assets "atomically" (i.e., without counterparty risk).

Traditional AMMs and LOBs are not suited to these limitations because they require many iterations to perform simple asset transfers (Angeris et al., 2022).[6] The inefficient usage of resources entailed by traditional market structures is not only costly to the traders directly involved in a transaction, but also generates a negative externality on all blockchain users.

CFMMs use a simple, minimalist approach to improve resource efficiency: They only need to keep track of a *liquidity pool* and its reserve vector $R \in R_+^N$ for the N tradable assets. The fact that CFMMs' internal states are small relative to LOBs makes them extremely light, as can be seen from their publicly accessible codes.[7] We now explain how the smart contracts of CFMMs price swaps using a deterministic function of the reserves.

1.3.2 Functioning of CFMMs

The market actors that interact with a CFMM classify broadly according to two roles. On the one hand, liquidity providers (LPs) determine the size of the liquidity pool. They can inject additional reserves but do not affect the relative share of the assets listed in the pool. In exchange, they obtain shares of the liquidity

[6] A CFMM takes constant memory to operate and constant time to process each trade. On the contrary, the memory and time taken by a LOB to process a set of trades scale linearly and quadratically in the number of trades to be processed.

[7] The codes for Uniswap's smart contract can be found at https://github.com/Uniswap.

pool in the form of LP tokens. These LP tokens, besides giving liquidity providers the right to withdraw the pool's reserves in proportion to their shares, also allow them to collect trading fees.

On the other side of the market, traders or users interact with the CFMM by supplying some assets and withdrawing others. The amount of output $O \in R_+^N$ a trader is entitled to receive is determined by his input, $I \in R_+^N$, and the CFMM's reserves R so as to keep the trading function constant at its pre-trade value. That is,

$$f(R - O + I) = f(R).$$

The trading function f is usually positive, smooth, strictly increasing, and strictly quasi-concave. These are very convenient properties as we will see in the next paragraphs.

The level curve of f attained at current reserves is often called bonding curve. In relation to the trading function, the division of interactions between the CFMM and market actors can be described as follows: Liquidity providers shift the CFMM to a different bonding curve, while traders move reserves along a fixed bonding curve (see Figure 1.1).

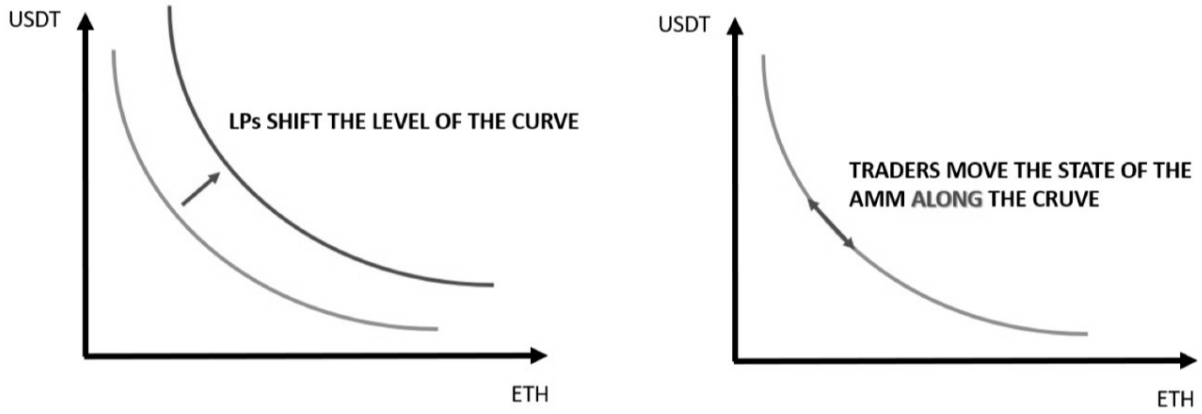

Figure 1.1: Bonding curve movements.

In practice, some CFMMs implement additional operations besides liquidity provision and swap as described above, making the classification of market actors blurrier.[8] It is nevertheless useful to keep in mind the provider-trader dichotomy as a benchmark.

CFMMs usually charge traders with a proportional fee, $1 - \gamma \in [0, 1)$, on their trades. In this case, the output O the traders receive for an input of I assets will be determined by counting only γI of the input in the trading-function equality, so that $f(R - O + \gamma I) = f(R)$. The other $(1 - \gamma)I$ will directly feed the

[8]For example, some CFMMs (e.g., Bancor) implement a "unilateral liquidity provision" functionality for supplying reserves only of part of the supported assets. We can treat unilateral liquidity according to our binary classification as a composed operation, made of balanced liquidity provision and subsequent swap. Also, as explained below, a swap that charges transaction fees can be thought of as a composed operation.

liquidity pool's reserves, which after the swap will amount to $R' = R - O + I > R - O + \gamma I$ for $\gamma > 0$. Notice that for f strictly increasing, $f(R') > f(R - O + \gamma I) = f(R)$. Transaction fees therefore shift the bonding curve upwards owing to the increased liquidity in the pool.

CFMM	Launch	Trading function	#pools	TVL	Blockchain(s)
Curve	Jan. 2020	Stableswap[i]	257	\$5.88b	Ethereum, Polygon, Arbitrum, Optimism,...
Uniswap	Nov. 2018	CPMM[ii]	1239	\$4.3b	Ethereum, Polygon, Arbitrum, Optimism, Celo
Pancakeswap	Sept. 2020	CPMM	108	\$2.92b	BSC
Balancer	March. 2020	G3MM[iii]	186	\$1.83b	Ethereum, Polygon, Arbitrum
Sunswap	Jan. 2022	CPMM and Stableswap (for stablecoin pools)	237	\$962.59m	TRON
Sushiswap	Sept. 2020	CPMM	688	\$831.22m	Ethereum, Arbitrum, Polygon, Harmony, Gnosis,...
VVS Finance	Nov. 2021	CPMM	18	\$451.47m	Cronos
DeFichain DEX	May 2020	CPMM	57	\$314.3m	DefiChain
BiSwap	May 2021	CPMM	87	\$273.11m	BSC
Osmosis	June 2021	Customizable[iv]	261	\$220.68m	Osmosis

[i] Stableswap:$f(R) = \alpha \sum_{i=1}^{N} R_i - \beta(\Pi_{i=1}^{N} R_i)^{-1}$, for $\alpha \leq 0, \beta \leq 0$.
[ii] CPMM:$f(R) = R_1 \cdot R_2$.
[iii] G3MM:$f(R) = \Pi_{i=1}^{N} R_i^{w_i}$, for $\sum_{i=1}^{N} w_i = 1$.
[iv] Customizable: Can be customized in each liquidity pool.

Table 1.1: Top-10 CFMMs by Total Value Locked (TVL).

Example: Uniswap. The most popular type of CFMM is the Constant Product Market Maker (CPMM) introduced by Uniswap, a DEX organized as a collection of liquidity pools for asset pairs. CPMMs use the trading function $f(R) = R_1 \cdot R_2$. The quantity I_1 of asset 1 required by the CFMM to output O_2 units of asset 2 is therefore given by

$$I_1 = O_2 \frac{R_1}{R_2 - O_2}. \tag{1.2}$$

The CFMM determines the terms of swap among assets though the quantity function $q : \mathbb{R} \times \mathbb{R}_+^N \to \mathbb{R}$ (Angeris and Chitra, 2022). $q(\Delta_2; R)$ is the purchase cost of Δ_2 units of asset 2; i.e., in eq. (1.2), $I_1 = q(O_2; R)$.[9] Conversely, $q(-\Delta_2; R)$ is the sale revenue for the converse operation.[10]

Notice that in eq. 1.2, $I_1 = q(O_2; R)$ is increasing in R_1 and O_2, but decreasing in R_2. This is on purpose since the CFMM has to incentivize users to keep the pool balanced, making it expensive to drain reserves and cheap to refurnish them when they are scarce. Moreover, the quantity function is convex in O_2 and

[9] In the two-asset case, some authors define q using asset 1 rather than asset 2 as the output asset.
[10] The quantity function for a swap among generic assets ij has additional inDEXs $q_i j(\Delta_j)$.

makes total reserve depletion impossible as

$$\frac{d}{dO_2}q(0_2;R) > 0, \quad \frac{d^2}{dO_2^2}q(0_2;R) > 0, \tag{1.3}$$

$$\text{and} \lim_{O_2 \to R_2} q(0_2;R) = \infty$$

The properties in eq. 1.3 can be summarized as *convex pricing*. The marginal price, or simply price dI_1/dO_2 (of asset 2) is given by:

$$p(R) \equiv p = \lim_{O_2 \to 0} \frac{q(O_2;R)}{O_2} = \frac{R_1}{R_2}.$$

More generally, the vector of prices quoted by a CFMM is proportional to the gradient of the trading function, i.e., $p = \lambda \nabla f(R)$, where $\lambda > 0$ is a scaling constant. This means that marginal prices are proportional to the slope of the supporting hyperplane of the trading set $\{(I,O) : f(R+I-O) \leq f(R)\}$; i.e. the set of trades that keep the market maker weakly above its bonding curve (see Figure 1.2 that shows the trading set in the shaded region and its supporting hyperplane at given reserves). Taking asset k as numeraire, so that $p_k = 1$, and setting $\lambda = 1/\nabla_k f(R)$, we get $p_i = \nabla_i f(R)/\nabla_k f(R)$ for all $i = 1, \dots N$ (Angeris et al., 2020; Fabi and Prat, 2022).

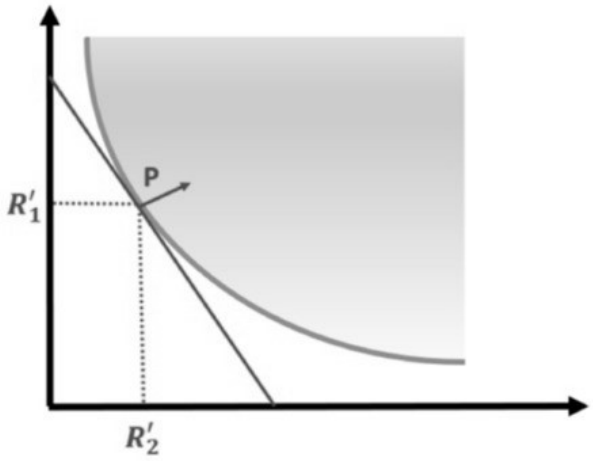

Figure 1.2: CFMM pricing.

1.3.3 Properties of CFMMs

The organizing principles of CFMMs are still eluding researchers. This situation is largely explained by the fact that CFMM were not introduced to solve a clear problem but instead to expediently address some practical challenges. Despite research being still in its infancy, some promising contributions have been made over the last few years.

The pioneers of research on CFMMs are Angeris et al. (2020, 2021a, 2021b, 2021c, 2021d, 2021e, 2022). They point out that much can be gathered about the behavior of CFMMs by studying the shape of their trading sets.

Price Oracle

The first topic these authors analyze is the decentralized oracle problem. A CFMM solves the decentralized oracle problem if its price closely tracks the fundamental price, usually assimilated to the price on a centralized exchange such as Binance. Decentralized oracles are of great value for the ecosystems hosted by a blockchain. Thanks to them, other DeFi applications can read prices directly on-chain, thereby avoiding a potential single-point-of-failure originating from the interaction with external applications. This is why major DEXs implement a price oracle API that retrieves their historical price averages.[11]

Angeris and Chitra (2020) show that if the trading set is convex, i.e. if the trading function is strictly quasi-concave, the CFMM does have a unique no-arbitrage equilibrium in which the CFMM tracks the reference price.

To illustrate how the no-arbitrage condition ensures that the CEX and the DEX quote the same price, consider Uniswap's CFMM with reserves R_1,R_2. Suppose that asset 1 is the numeraire and the CEX quotes a market price of p for asset 2. Suppose also that both DEX and CEX charge no fees. Now, whenever $p \neq R_1/R_2$, an arbitrageur can profit from trading across the DEX and the CEX. For $p > R_1/R_2$, asset 2 is underpriced in the DEX. The arbitrageur can then swap asset 1 for asset 2 on the DEX and sell it to the CEX, cashing in the price difference. If $p < R_1/R_2$, then asset 1 is underpriced in the DEX, leading to a symmetric arbitrage opportunity. A strictly convex trading set guarantees that there exists a unique post-arbitrage reserve level with the implied price matching that of the CEX. Note that the oracle property is robust to the addition of transaction fees γ in the DEX. More precisely, fees cause the DEX price p^{DEX} to track the CEX price p^{CEX} within tight bounds, $p^{DEX} \in [\gamma p^{CEX}, \gamma^{-1} p^{CEX}]$.

Angeris et al. (2021e) simulate the price-oracle property for Uniswap using agent-based techniques. Other authors provide empirical evidence with real-world data. Lo and Medda (2021), use an ARDL-ECM (Autoregressive Dynamic Lag Error Correction Model) to estimate the cointegrating relationship between Uniswap's ETH-USDT exchange rate and the exchange rates for the same pair offered by several centralized exchanges. Lehar and Parlour (2021) find that price differences between Binance and Uniswap are usually below 1% when liquidity is abundant.

The empirical literature also investigates the discovery of fundamental prices directly on DEXs. A market discovers the fundamental price of an asset when its quoted price is efficient (i.e., it reflects the value of the

[11] For Uniswap the price oracle API can be retrieved at: https://docs.uniswap.org/protocol/concepts/V3-overview/oracle.

asset). When a DEX quotes fundamental prices, not only those prices and the ones posted by coexisting CEXes are close under no arbitrage, but the DEX itself becomes the reference market. In other words, the CEXes adjust their prices towards the ones quoted by the DEX, not the other way around. Han et al. (2021) argue that Binance traders do indeed consider Uniswap prices as efficient. Specifically, they show that when liquidity provision on Uniswap increases, Binance traders tend to gravitate more towards Uniswap's price, whereas Uniswap traders become less sensitive to Binance's price.

Fabi and Prat (2022) highlight a striking equivalence among the CFMM's no-arbitrage equilibrium and the expenditure minimization problem (EMP) in microeconomic theory. For a bonding curve set at $f(R) = K$, they show that equilibrium reserves equal the Hicksian demand $h(p; K)$ of a consumer given indirect utility K and commodity prices p. The LP value function (i.e., the value of the pool) equals the expenditure function $V(p; K)$ under the same interpretation as before. From now on, we omit to state the level of the trading function (i.e., the indirect utility in the EMP) when unnecessary, and just write $h(p)$ and $V(p)$.

CFMMs equivalence

Angeris and Chitra (2020) study the classification of CFMMs. They find that two CFMMs are equivalent if their trading functions generate the same trading sets. This is intuitive since trading functions and their bonding curves are designed to represent trading sets, in the same way that utility functions represent preferences or transformation functions represent production processes in microeconomic theory (Mas-Colell et al., 1995, chapters 3 and 5). More formally, two CFMMs with trading functions f and g are equivalent if $f = v \circ g$ where v is a monotonically increasing function.

Accordingly, under some specific parametrization, Uniswap and its multi-asset generalization Balancer are actually indistinguishable. Balancer uses the trading function

$$f(R) = \Pi_{i=1}^{N} R_i^{w_i} \tag{1.4}$$

with weights $w_i \in (0, 1) : \sum w_i = 1$. For $N = 2$ and $w_1 = w_2 = \frac{1}{2}$, the Uniswap trading function g(R) can be written as

$$g(R) = [R_1^{\frac{1}{2}} R_2^{\frac{1}{2}}]^2 = [f(R)]^2,$$

thus proving the equivalence with Balancer for this parametrization, while, for $w_i \neq \frac{1}{2}$, the equivalence breaks down.

Optimal curvature

As previously stated, CFMM pricing should discourage traders to create excessive movements in the pools' reserves. This is especially important to limit the cost of adverse selection for liquidity provision, originating whenever an arbitrageur takes advantage of a discrepancy between the CFMM's prices and the prices quoted in the reference market. In DeFi jargon, the cost of adverse selection is called Impermanent Loss (IL).

A way to mitigate excessive arbitrage, and thus impermanent losses, is by having trading functions that induce convex pricing, as in the case of Uniswap (eq. 1.3). This requires positive slippage, defined as the derivative of the price with respect to output; i.e. for a two-assets pool,

$$\frac{d}{dO_2} p(R_1 + q(O_2; R), R_2 - O_2) > 0.$$

Angeris et al. (2022) and Fabi and Prat (2022) show that slippage is positive if the trading sets are convex, a property that is satisfied if and only if the trading function is strictly quasi-concave or if it exhibits positive Gaussian curvature.

Some curvature in the trading function is thus desirable to create slippage and limit the harms of arbitrage, but excessive curvature can discourage traders. What is then the optimal curvature of a trading function? The answer that Angeris et al. (2022) provide is that low-curvature CFMMs are more suited to low-volatility, mean-reverting assets, such as stablecoins, whereas high-curvature CFMMs are most suited to volatile assets. Trading of the first asset type is not information-driven, hence reserves are likely to move according to small random shocks that pose little risk of depletion. In this context, it is natural to let swaps take place at low slippage. On the other hand, trading of volatile assets is more susceptible to arbitrage, which can cause large reserves movements and harm liquidity providers. In this case, setting a higher slippage via curvature shields LPs. Optimal CFMM curvature resembles classical results in the order-book literature (Glosten and Milgrom (1985)). In that context, to reduce adverse selection, market makers eliminate excessive orders. Analogously, in CFMMs, the smart contract discourages excessive trading by implementing a steeply curved trading function.

Given that the slippage of any convex price function can be bounded from below using a linear function of the traded tokens, Angeris et al. (2022) uses this lower bound to compute the maximum trade size that makes liquidity provision profitable.

The main takeaway from analyzing CFMM curvature and LP returns is that an optimal CFMM should adjust curvature dynamically based on the volatility of their traded assets (Krishnamachari et al., 2022). To the best of our knowledge, Krishnamachari et al.'s design has not yet been implemented. However,

alternative dynamic curves are already used by some recent AMMs.[12]

Replication

The other fundamental aspect covered in the series of papers by Angeris and coauthors is payoff replication via CFMMs (Angeris et al., 2021b). The problem of payoff replication deals with the inversion of the liquidity providers' value function. We have already explained how to compute LPs' value under no-arbitrage for a given trading function (see section 1.3.2). Now we address the converse question: Starting from a desired payoff profile for liquidity providers, can we find a trading function which replicates that specific payoff? The authors manage to do so for a comprehensive class of LP payoff functions. This problem has the same flavor as the one addressed by Chen and Pennock (2012) for prediction markets since they showed how to derive the AMM associated to a particular cost function.

To solve the replication challenge, Angeris et al. (2021b) exploit duality theory. They prove that the dual problem associated to the maximization of arbitrageurs' gains corresponds to the minimization of LPs' value. In analogy with demand theory, the process of finding the trading set that achieves a given LP payoff is analogous to recovering consumer preferences from a given expenditure function (Fabi and Prat, 2022).

Angeris et al. (2021b) establish a bijection between the family of LP value functions that are concave, positive, increasing, homogeneous of degree one in the market price and CFMM with convex trading sets. Each property of the class of payoff functions that can be replicated has economic interpretations. Concavity of the value function is analogous to requiring that LPs face impermanent losses; positivity implies that the LP cannot run a deficit; increasing functions ensure that the value of the portfolio grows as the assets appreciate. Homogeneity of degree 1 is simply a scaling property: If all prices are scaled proportionally, so does the value of the LPs' portfolio.

Privacy

So far, we have seen that convexity of the trading set is a desirable property. However, it also has the drawback of causing a loss of privacy for users, as explained by Angeris et al. (2021d). Since the history of pool reserves is public, the quantity function derived from a convex trading set can be inverted so as to recover the exact sequence of trades that led a liquidity pool to its current state. Transactions batching is not a viable solution as it considerably worsens users' terms of trade owing to the price impact. Randomizing also worsen the price impact because of convex pricing. Moreover, the Central Limit Theorem implies that one can efficiently extract the random noise component from prices.

[12] For example, Dodo uses a dynamic bonding curve (https://docs.dodoex.io/english/).

1.4 Decentralized Platform Economics

We have covered the main properties and optimal design of CFMMs. We now turn our attention to the market structure that they generate. The economics literature on platforms uses the term two-sided market to refer to markets that have two distinct sides exerting direct positive externalities on each other (Rochet and Tirole, 2003). This is exactly the case for DEXs that operate through a CFMM. As mentioned in the previous sections, the two market sides are liquidity providers and traders. Liquidity providers benefit from higher trade volumes by earning more fees. Traders benefit from higher liquidity by getting lower slippage from the CFMM.

A small but fast-growing literature studies the platform economics of decentralized exchanges. The earliest contributions on this topic are by Aoyagi and Ito (2021), Capponi and Jia (2021), and Lehar and Parlour (2021). These papers propose variations of the same benchmark model to study the rents of traders and liquidity providers who participate to a constant-product DEX or a CEX, together with the resulting allocation of liquidity among the competing exchanges.

1.4.1 Comparison of CFMM and LOB exchanges

Although this strand of literature outlines many similarities with classical market microstructure models, it also highlights important differences. For example, in LOBs, LPs choose both the quantity they want to supply as well as the minimum sale price of their assets. By contrast, when interacting with a CFMM, LPs choose only their supply while the price of the liquidity is given by a function of the CFMM's reserves. Once liquidity is supplied, it cannot be conditioned on the price. Due to this feature, the rents from liquidity provision in CFMMs are mutualized, whereas in LOBs liquidity provision is competitive. This feature creates an important difference in the composition of LPs' rents for the two market mechanisms. In LOBs, if LPs compete on prices, their revenues come mainly from setting price slippage strategically. On the other hand, in CFMMs, slippage revenues are captured by arbitrageurs, so LPs revenues come from transaction fees (Lehar and Parlour, 2021).

Interestingly, Uniswap's latest upgrade (Uniswap V3) closes part of the gap between CFMMs and LOBs by introducing concentrated liquidity. This functionality allows LPs to supply reserves only within specific price ranges of the traded pair. The advantage of this additional feature for a liquidity provider is that, when the exchange rate matches the range at which she supplies liquidity, she can achieve the same fee income as in a standard constant-product pool supplying lower liquidity. The drawback of liquidity concentration is that exchange rates that fall outside of the specified price range can cause substantial losses for LPs. Liao and Robinson (2022) show that, despite its recent launch in 2021, Uniswap V3 was able to collect more

liquidity than some of the most popular CEXes (Binance, Coinbase, Kraken, Gemini).

1.4.2 Incentives for liquidity providers

The economic literature on two-sided markets underlines that platforms must attract the market side whose demand is more elastic to rents (Rochet and Tirole, 2003). In the case of DEXs, the elastic side of the market are liquidity providers. In line with theory, existing DEXs strive to attract liquidity providers, competing against each other and against CEXes.

Yield Farming

The process through which LPs harvest the rewards from competing DEXs is known as yield farming. Its revenues are mostly made of two components. First, LP tokens, which allow LPs to accrue transaction fees and claim a fraction of the DEX that they furnish (see section 1.3.2). Second, governance tokens, that carry voting rights on DEXs' strategic decisions (for example, the level of trading fees associated to a given liquidity pool).

The open, unregulated nature of DeFi makes liquidity competition particularly harsh for DEXs. Such competition gives rise to frequent mushrooming of yield farming opportunities. Their landscape is so vast and transient that LPs rely on aggregators, such as Yearn or 1inch, to optimize their liquidity positions across multiple DEXs. To the best of our knowledge, yield farming has received limited academic attention. We therefore only present two major cases of yield farming opportunities that could inspire future research.

Example: Vampire attacks

Some yield farming opportunities originate from DEXs' attempts to enter a market. Network effects should make it difficult for new DEXs to attract liquidity towards them rather than towards well-established protocols. Yet, the open nature of smart contracts combined with traders seeking yield-farming opportunities open the way for newcomers to use cloning strategies. The most common of them is called vampire attack. It entails a DEX siphoning liquidity from another protocol by offering its liquidity providers attractive deals.

The first example is the SushiSwap attack, which allowed Sushiswap (a Uniswap fork) to appropriate over $\$1B$ of liquidity from Uniswap in less than a week. Sushiswap offered SUSHI tokens to Uniswap LPs as extra rewards in exchange for migrating their Uniswap LP tokens on Sushiswap through the Masterchef smart contract.[13] In response, Uniswap launched in September 2020 its governance token, the UNI token, along with an airdrop of 400 UNI to users who had previously used Uniswap V1 or V2.

[13]The MasterChef contract can be found here: https://docs.sushi.com/docs/Developers/Sushiswap/MasterChef.

Vampire attacks have been common in DeFi space and, unfortunately for the victim protocol, give to LPs attractive profit opportunities, at time combined with the entertainment of gamification.[14]

Example: Curve wars.

Other yield farming opportunities originate from DEXs' efforts to establish a dominant market position, as in the "Curve wars" that started around May 2020 and which is still ongoing. The Curve wars are competitions among DeFi protocols for the control of Curve's liquidity. Curve being the biggest DEX by total value locked ($5.83b at the time of writing), DeFi protocols have a strong interest in having their tokens traded on Curve at competitive rates. This enables them to attract yield farmers and thus to establish prominence over their competitors.

A protocol can achieve such goal by holding a big share of Curve's liquidity tokens, the CRV. These can be locked into Curve and converted into Curve's svCRV governance tokens. Besides providing passive yields, svCRV allow their owners self-vote for pool gauging; that is, directing new emissions of CRV tokens to themselves. Generous CRV emissions attract yield farmers and explain the importance of controlling the outstanding token basis of CRVs.

In a similar fashion to Vampire attacks, Convex Finance (Convex for brevity) managed to control the majority of CRV tokens through a cloning mechanism. It devised a protocol that allowed CRV owners to sell their voting rights on Curve for more liquid assets with boosted financial yield, the CvxCRV tokens minted by Convex. Then Convex allowed its LPs to vote on Curve's governance by locking its native tokens, the CVX. In this way, Convex effectively took control over Curve.

Despite Convex being in control of most CVXs in circulation, Curve's governance is to some extent still contestable. Vote buying (or "bribing") protocols such as Bribe.cvx or Votium allow owners of CRV or CVX tokens to delegate their votes in exchange for side payments. Some protocols such as Llama Airforce Union even optimize vote buying for CVX holders.

Adverse selection costs

The downside of liquidity provision in DEXs are adverse selection costs. These arise whenever reference prices change because LPs have their funds locked in liquidity pools, and so cannot react before arbitrageurs take advantage of the price discrepancy.

The usual metric of adverse selection in CFMMs is the impermanent loss (IL) as already mentioned in Section 1.3.3. Fabi and Prat (2022) use a static model to compute the IL as the arbitrage profit from a price

[14] Enzo finance recently carried out a vampire attack against six Ethereum exchanges masked as a videogame (https://decrypt.co/87778/enso-finance-vampire-attack-six-ethereum-defi-competitors).

shock. They show that the impermanent loss caused by a price move from p^0 to p^1 is given by:

$$IL(p^0, p^1) = p^1 \cdot [h(p^0) - h(p^1)] \tag{1.5}$$

$$= p^1 \cdot \nabla_p V(p^0) - V(p^1) \geq 0. \tag{1.6}$$

The second equality in eq.(1.6) follows from $h(p) = \nabla_p V(p)$ (Shephard's lemma), and the final inequality from the concavity of V in p. Fabi and Prat (2022) demonstrate through eq. (6) that IL is convex in p^1 with global minimum $IL = 0$ at $p^0 = p^1$.

Milionis et al. (2022) develop a continuous-time model of liquidity provision and propose a new measure of adverse selection costs called *loss-versus-rebalancing* (LVR), which is simply the sum of IL in eq. (1.5) over all price changes in a given time period. They call their measure loss-versus-rebalancing rather than impermanent loss to underline that they do not keep reserves fixed at their initial value throughout the history of arbitrages. In a static model, LVR and IL coincide. However, in a dynamic model, IL do not take into account the impact of reserves rebalancing after each arbitrage. Formally, assuming that prices evolve according to a sequence $p^0, p^1, \ldots, p^{n+1}$ and that arbitrageurs trade with the CFMM in periods 1 to n,

$$\begin{aligned} LVR &= \sum_{k=1}^{n} IL(p^{k-1}, p^k) \\ &= \sum_{k=0}^{n} h(p^k)[p^k - p^{k-1}] - [V(p^{n+1}) - V(p^0)]. \end{aligned} \tag{1.7}$$

Milionis et al. (2022) solve the continuous-time limit of eq. (1.7) assuming a geometric Brownian motion for p. For a timespan of length T, they demonstrate that

$$LVR \longrightarrow \int_0^T l(p_t)dt,$$

and they solve in closed-form the instantaneous LVR rate $l(p_t)$, expressing it as a function of price volatility and curvature of the LP expenditure function. For a single risky asset with price volatility σ, the instantaneous LVR reads

$$l(p) = -\frac{\sigma^2 p^2}{2} \frac{d^2}{dp^2} V(p). \tag{1.8}$$

Milionis et al. (2022) estimate LVR applying eq. (1.8) to Uniswap and Balancer (see eq. 1.4). They find that daily TVL is about 3bp and that LPs' annual yield is about 8%.

1.4.3 Incentive for traders

The recent models of competing exchanges distinguish traders as either noise or informed traders (i.e., arbitrageurs). Noise traders have exogenous reasons to trade; they either swap-in or -out a fixed quantity of assets with equal probability. Informed traders trade with the CFMM to profit from price discrepancies among exchanges.

The interactions between liquidity providers and traders determine the inventory of liquidity pools as well as the distribution of liquidity among competing DEXs and CEXes. The literature keeps CEX and DEX platforms passive but endogenizes the allocation of trades and liquidity across such platforms. It shows that the impact of trader participation on the liquidity of DEXs is in line with established results in the market microstructure literature (Kyle, 1985; Glosten and Milgrom, 1985; Easley and O'hara, 1987). Informed trading discourages liquidity provision as it causes impermanent losses. Noise trading attracts LPs as it allows them to profit from fees while avoiding large movements in reserves. In the case of DEXs, informed trading also causes temporary capital gains thanks to convex CFMM pricing: If noise trades net-out in expectation, then expected pool reserves increases due to the (strict) convexity of the bonding curve. Formally, for a distribution over trades Δ with $\mathbf{E}[\Delta] = 0$, Jensen's inequality implies that

$$\mathbf{E}[q(\Delta; R)] \geq q(\mathbf{E}[\Delta]; R) = 0. \tag{1.9}$$

The intuition behind eq. (1.9) is that for convex bonding curves, buying assets from a CFMM is more expensive than selling them.

The problem with such capital gains is that they are completely cancelled out after a single arbitrage. Hence the only lasting factors for liquidity providers originate from transaction fees and asset volatility. In models such as Capponi and Jia (2021) and Lehar and Parlour (2021) where traders' platform choice is exogenous, both fees and volatility have very clear *direct* effects on equilibrium liquidity: Positive for fees, negative for volatility. However, other contributions point out that equilibrium effects that also take into account traders' reactions can be subtler.

Saleh et al. (2022) argue that only moderate fee rates stimulate liquidity provision. Excessive rates instead fire back on LPs due to the disincentive that they entail for traders. Surprisingly enough, Saleh et al. (2022) show that charging low trading fees rather than none can attract traders. Since collected fees are poured back in the liquidity pool, they reduce slippage, in turn making the CFMM more attractive to traders if the lower slippage compensates the higher fees.

Aoyagi and Ito (2021) model a CEX-DEX duopoly. They highlight that higher DEX liquidity generates positive spillovers on the CEX. In their model, an exogenous increase in DEX liquidity causes both types

of traders to migrate from the CEX to the DEX, but not to the same extent. Informed traders are more likely to migrate than noise traders due to their higher sensitivity to slippage. The signal-to-noise ratio of the CEX therefore decreases, leading to lower bid-ask spreads in its LOB.

The theoretical results discussed above are supported by empirical evidence. Lehar and Parlour (2021) conduct a regression analysis on Uniswap daily data for 1,376 pools. They show that, as predicted by their model, liquidity is negatively affected by volatility and positively affected by noise trading. Lehar and Parlour (2021) argue that the stability of reserves is guaranteed by Ethereum's gas fees, which prevent LPs from draining pools in response to extreme market events (contrarily to what is usually observed in LOBs). They support their stability hypothesis by observing that the 41% ETH price decline which occurred in May 19 of 2021 only caused a withdrawal of 17% of Uniswap's liquidity.

Capponi and Jia (2021) analyze a dataset ranging from April to December 2021, collected from major CEXes and 12 liquidity pools, half from Uniswap, half from Sushiswap. They also find empirical support for a negative relationship between volatility and liquidity, and a positive relationship between liquidity and noise trading. They test and confirm that gas fees and asset volatility are positively correlated.

Finally, Heimbach et al. (2021) also analyze Uniswap data from its deployment in May 2020 to January 2021 in order to describe liquidity distribution across pools. They find that pool participation follows a power law: Most LPs participate to a single pool, and 96.5% of them participate to at most 5 pools at once. They argue that LPs' cautiousness is the most plausible explanation for liquidity concentration. LPs tend to participate only to pools whose risks they can assess with confidence. Such an evaluation is easier for the few popular pools to which most LPs participate, rather than for the many small pools in the tail of the liquidity distribution.

1.5 Miner Extractible Value

We have seen that CFMM is a simple and versatile technology. Nonetheless, its minimalistic design is not without drawbacks. The most notorious one is Miner Extractible Value (MEV).

Technically, MEV are profits that blockchain miners or speculators can earn through the reordering of transactions within blocks. Transaction reordering is harmless to blockchains featuring simple transactions such as Bitcoin, but not for CFMMs and DeFi applications. For instance, when two buy orders are sent to a CFMM, one after another, the pricing formula results in a better deal for the first order.[15]

The name MEV comes from the fact that miners are in control of transaction ordering within blocks and can freely select it according to strategic or even malicious motives. Most of the times, however, aggressive

[15] Exposure to MEV is another reason why LOBs are unsuited to blockchains

block reordering is not performed by miners but rather arbitrage bots that continuously scan the memory pool (mempool for brevity) to spot profit opportunities, and engage in gas auctions to claim them (see for example Figure 1.3). This behavior indirectly benefits miners who ultimately get the rents from bots' auctions (Capponi and Jia, 2021). MEV priority auctions are problematic for the whole ecosystem as all participants bear the congestion externality they generate. Ethereum's mempool is teaming with bots that profit from MEV opportunities at the expenses of users. This is why practitioners refer to it as a "dark forest", a term originally coined in the sci-fi novel Remembrance of the Earth's Past (Liu, 2008) and popularized in the context of blockchains by a blog article (Robinson and Konstantopoulos, 2020).[16]

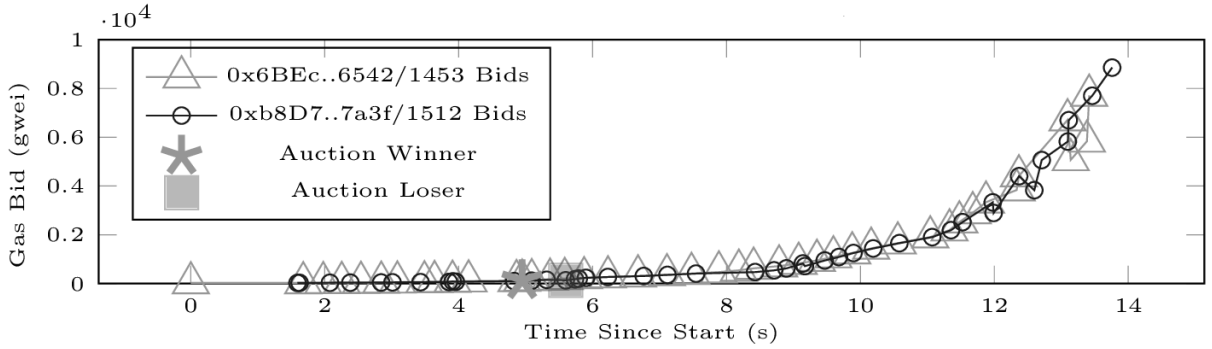

Seconds Elapsed	Quantity @ Price Bid	Ethereum Transaction Origin (Public Key Hash)	Nonce	Transaction Hash
0.000	192085 @ 25.10	0x6BEcAb24Ed88Ec13D0A18f20e7dC5E4d5b146542	1453	0xd32653ca9694a6d1299335f3c04f74cc159bee48c1d32d3a421db08c638ffc78
1.593	231520 @ 25.00	0xb8D76f4BC2518F8eb508bf0Ccca76f8F9DD57a3f	1512	0xb901e6dc2c229fd9105448fcc23eaebdedb476c21b6c6e7ddf8d2df4e838d2c7
1.624	231520 @ 28.75	0xb8D76f4BC2518F8eb508bf0Ccca76f8F9DD57a3f	1512	0x9f592504eb71a7452b7a395a7f5ecd34eaa5d090da1162e74221562af54c8f67
1.679	227534 @ 28.81	0x6BEcAb24Ed88Ec13D0A18f20e7dC5E4d5b146542	1453	0x83e2a6774654a9540c3fad8837afcc88b4c932ab2374819254f887305c3a4b22
...	...	...	...	...
4.949	227534 @ 134.02	0x6BEcAb24Ed88Ec13D0A18f20e7dC5E4d5b146542	1453	0xc889bd13594f75e4dd824f04f0c2ad03896cb7ec6518df02455e9560367bb9c4
5.599	231520 @ 133.76	0xb8D76f4BC2518F8eb508bf0Ccca76f8F9DD57a3f	1512	0xaa86d782328c0c9c422e3f2a3170ff41ae21a27ad395c48db76b0080898f85db
...	...	...	...	...
13.383	227534 @ 5834.77	0x6BEcAb24Ed88Ec13D0A18f20e7dC5E4d5b146542	1453	0xb0dc97140394c5f65332ebc459d5e66f89099dbb4d335c866b32280270102858
13.416	227534 @ 7716.48	0x6BEcAb24Ed88Ec13D0A18f20e7dC5E4d5b146542	1453	0x1825be6951577e72a1dafc8de564ce1ccfe5d284173e11e77b2e7f6b1b44571c
13.462	231520 @ 7701.08	0xb8D76f4BC2518F8eb508bf0Ccca76f8F9DD57a3f	1512	0xa9823358c99149f0e6343c604c35988468d01d02868437d8251b3cee282dc92b
m13.759	231520 @ 8856.24	0xb8D76f4BC2518F8eb508bf0Ccca76f8F9DD57a3f	1512	0x366c30a534b5f3d8a6d251f97d401997624d1fe8d3af07ede4d19105dc970942

Figure 1.3: Example of Price Gas Auction (PGA) that was observed over the Ethereum peer-to-peer network. The top graph shows the gas bids of two observed bots over time, while the bottom table details the bots' initial and final bids and the two mined bids (enclosed within continuation dots). Source: Daian et al. (2019).

1.5.1 Front-running

Daian et al. (2019) is the first academic contribution that documents the significance of MEV. It studies its consequences and simulates bots' priority gas auctions. Analyzing bot activity on Ethereum, Daian et al. (2019) find that miners often earn more from MEV than from block rewards. They further show that, to fully exploit MEV rents, miners should rewind and fork the blockchain, committing a so-called time-bandit

[16] "Dark Forest" is the name of a fictional hypothesis on the state of the universe on which the novel fantasizes. According to the dark forest hypothesis, the universe is an adversarial environment populated by advanced civilizations. They hide from each other as an encounter would trigger devastating wars. By analogy, Ethereum's mempool is said to be a dark forest because it is populated by arbitrage bots that engage in costly bid auctions if they detect each other.

attack. There is no account of a time-bandit attack having happened on a popular blockchain, but its profitability represents a concrete consensus layer threat. Ironically enough, the evidence collected in Daian et al. (2019) was partly a side-effect of the same authors publishing a preliminary report on arbitrages in DEXs (Bentov et al., 2017). Its publication attracted the attention of developers who designed bots that could exploit the arbitrage opportunities identified in the report.

Qin et al. (2021) quantify the MEV on Ethereum from December 2018 to August 2021. They find it to be over four times the value of block rewards and transaction fees combined. The highest MEV opportunity was worth 4.1M USD; more than 600 times the sum of block reward and transaction fees for an average block. They argue that this MEV opportunity was sufficient to incentivize a miner with only 0.1% of the network mining capacity to commit a time-bandit attack.

1.5.2 Sandwich attacks

The most common form of MEV on DEXs are sandwich attacks. A sandwich attack can be performed on both buy and sell orders. It involves the combined front-running and back-running of a target transaction. When a bot (or miner) is notified of some pending buy order on a CFMM, it can forge a buy and a sell order of equal amounts and squeeze the user's transaction within these two. Doing so, the attacker earns a profit as convex bonding curves make it profitable for the bot to buy assets and sell them back when their reserves are higher. These sorts of manipulations take place also in traditional finance, especially through algorithmic trading (Michel, 2020). However, CFMMs makes the attackers' life easier as all pending transactions are public. The seriousness of MEV is also highlighted by Park (2021), who compares constant-product pricing

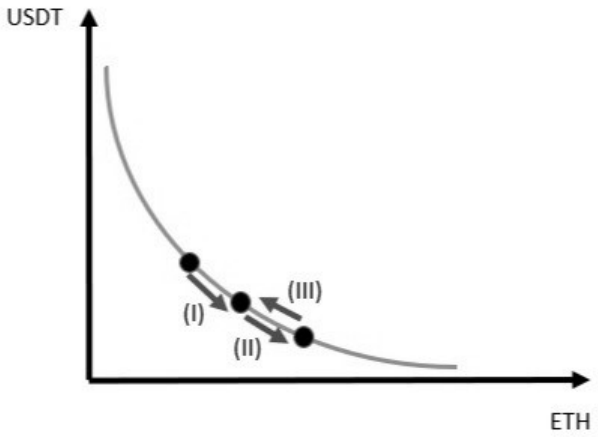

Figure 1.4: Sandwich attack.

with traditional LOBs. Although he finds that CFMMs satisfy the desirable feature of not inducing users

to break their orders into smaller ones nor split them across multiple DEXs, Park also proves that CFMMs with constant-product pricing are always vulnerable to sandwich attacks.

To see why, consider the CFMM moving as in Figure 1.4 suppose that the initial reserves of the two assets are R. The attacker observes the mempool and notices that a trader wants to buy O_2 units of asset 2. The attacker first front-runs the trader with the same operation and then back-runs the trader by inverting it. To front-run the trader, the attacker pays $I_1 = q(O_2; R)$ of asset 1, moving reserves to $R^I = R + \Delta$, with $\Delta = (q(O_2; R), -O_2)$ (movement I). Then, the original trade is executed, but now, the trader pays $q(O_2; R + \Delta)$ instead of $q(O_2; R)$ to receive output O_2, moving the CFMM reserves to $R^{II} = R + (q(O_2; R) + q(O_2; R + \Delta), -2O_2)$ (movement II). Finally, by back-running the trade, the attacker returns the reserves to R^I, thereby paying back O_2 of asset 2 and receiving $q(O_2; R+\Delta)$ of asset 1 (movement III). At the end of the sandwich attack, the attacker has 0 units of asset 2 but ends up with a positive balance

$$q(O_2; R + \Delta) - q(O_2; R) > 0$$

of asset 1, which corresponds to the extra input price the user pays owing to the sandwich attack. Given that the inequality follows from the convexity of the quantity function, sandwich attacks are always profitable for any CFMM with convex trading sets (Fabi and Prat, 2022).

The practical consequence of sandwich attacks for users is that they cause involuntary price slippage. In other words, the price at which an order is meant to be executed might not correspond to the actual execution price. To defend against sandwich attacks, DeFi platforms allow users to set a maximum slippage tolerance. Tuning the tolerance threshold is not trivial as it makes users vulnerable to arbitrage if set too low, but stops most of their trades if set too high. Heimbach and Wattenhofer (2022) compute the optimal slippage tolerance as a function of the size of the trade and the characteristic of the pool. Their simulations indicate that customized tolerance rules outperform the constant slippage tolerance implemented by popular DEXs.

1.5.3 Solutions to MEV

Flashbots is the leading proposal to alleviate MEV. Its key idea is to remove block-building capabilities from miners and to transfer them to dedicated agents. In Flashbots, block builders are called searchers. They scan the mempool in search for transactions carrying profitable MEV opportunities, bundle them together, and auction these bundles to miners through a private relay channel.[17] This design reduces the dangers of MEV in two important ways. First, as miners cannot decide anymore the internal block ordering, they have

[17]Private relay channels serve the same function as dark pools in the context of centralized LOBs.

no incentive to commit a time-bandit attack. Second, the private communication channel between builders and miners eliminates the congestion costs generated by priority auctions.

The separation between builders and miners is part of the design proposal for the new Ethereum V2 running on proof-of-stake. It entails important consensus layer modifications which Ethereum's developers have not yet finalized. However, they are already proposing a temporary solution through the Builder API. Ethereum's consensus-layer clients can install it to request data directly from the network of block builders, which operates through execution layer clients.

1.6 Conclusion

We have presented the state of research on AMMs in Decentralized Finance. Born as tools to facilitate information elicitation in prediction markets, AMMs have morphed into the CFMMs now used in decentralized exchanges. These modern variants are suited to blockchains thanks to their lightweight design. We have explained how decentralized exchanges operated by AMMs give rise to two-sided markets whose success crucially depends on their ability to attract liquidity. Finally, we have discussed the problems raised by miner extractible value, highlighting possible solutions and open challenges.

Although a rigorous theory of decentralized market making does not exist yet, much progress has been accomplished over the last couple of years. A vocabulary has been elaborated to rigorously describe the design space. Equilibrium models have been built to study existing applications and to provide suggestions for their improvement. With these foundations laid down, the stage is set for a new wave of research on the general optimality principles of CFMMs. In this enterprise, standard microeconomic theory is likely to prove particularly useful since it accurately describes the inner workings of CFMMs (Fabi and Prat, 2022). Besides providing theoretical insights on market design, we also expect future contributions to foster our understanding of the interaction between different types of exchanges. The cut-throat competition to attract liquidity is one of the most striking feature of DeFi. Its study will require more advanced models and better empirical evidence on equilibrium liquidity provision.

Building a solid theory of CFMMs is a challenging endeavor as the technology evolves at an ever-increasing pace. Keeping up with this continuous stream of innovations may seem unfeasible, and yet, recent contributions have considerably closed the gap separating theory from practice, making it realistic to envision a new stage of DeFi where protocol design would draw on scientific insights rather than on serendipity and educated guesses.

Chapter 2

The market microstructure of Uniswap.

Abstract

Uniswap is the largest decentralized exchange operating on the Ethereum Blockchain. It allows users to exchange one cryptoasset for another through liquidity pools whose exchange rate is set algorithmically. Uniswap has gained popularity as an oracle, providing realtime price data to diverse Decentralized Finance (DeFi) protocols. To maintain its reputation as a reliable oracle, it is crucial that Uniswap's prices closely track reference market prices. This accuracy in pricing crucially relies on arbitrage activity. However, inventory holding costs can impact the ability of traders to engage in arbitrage. This paper is the first attempt to develop a microstructure model to analyze the impact of these costs on the price accuracy of Uniswap. The model is estimated using price data from Uniswap and Binance. I find that traders are less likely to close the arbitrage opportunities as the size of the pool increases, a finding that is in line with the impact of inventory holding costs. Conversely, when traders perceive the trading pair to be stable and not subjected to inventory holding risks, they are more inclined to take advantage of arbitrage opportunities.

2.1 Introduction

Uniswap is a decentralized exchange (DEX) which allows for direct peer-to-peer cryptocurrency transactions without the need for an intermediary. Uniswap is an automated protocol composed of multiple liquidity pools, with each pool corresponding to a pair of crypto assets. The price at which the two assets can be exchanged is set algorithmically according to the constant product formula. This formula states that once an exchange is done, the product of the new reserves should remain at its pre-trade value.

Launched in 2018, Uniswap has exceeded $\$1,5$T in total trading volume in 2023 against $1 billion in 2020 and $100 million in 2019. This remarkable growth has positioned Uniswap as an effective alternative to the Centralized Exchanges (CEXs) that are based on limit order books.

The accuracy of Uniswap's prices relies on the fulfillment of the no-arbitrage condition. While DEXs and CEXs operate differently, prices on DEXs should always converge to those on CEXs – considered as the reference markets – as traders can make a profit by buying an asset in the exchange market where the price is low and selling it on the other one. However, inventory holding costs can impact the ability of traders to engage in arbitrage if they entail substantial deviations from their desired inventory level. Deviating from the desired inventory can expose arbitrageurs to risks and costs that impact their overall profitability. This effect becomes particularly pronounced as the pool size increases, resulting in a decrease in the price impact of trades and a substantial rise in inventory costs.

This paper develops a microstructure model to analyze the impact of these costs on the price accuracy of Uniswap. The model is estimated using 1-minute closing price data of Uniswap and Binance markets from May, 2020 to December, 2022, mainly focusing on the ETH-BTC pair. The empirical results provide support for the theoretical model and suggest that as the pool increases, traders tend to react less to cross market arbitrage opportunities as they face inventory risks. During the early stages of the pool, traders managed to close almost the entire price discrepancy between Uniswap and Binance. However, their willingness to close the gap gradually reduced, with only around 20% in August 2020, followed by a further decline to 10% in October 2020. This pattern coincides with the significant increase in liquidity provision within the ETH-BTC pool, resulting in a reduced price impact in the Uniswap market. As the pool size expands, trades need to be larger in order to influence the Uniswap price and bridge the gap. This requires traders to hold more inventory, thereby increasing their exposure to market risks. Consequently, traders become less responsive to price differences between the two markets, despite the potential profit opportunities they offer. Starting from early 2022, traders once again began closing the gap. This coincided with a decrease in the price impact as the pool size decreased. Nonetheless, during this period, traders were able to close on average 30% of the gap, possibly due to the illiquidity and limited activity in the pool following the migration to

Uniswap V3. When extending the analysis to other pairs, it is observed that when traders are not exposed to market risks, as is often the case with stablecoin pairs due to their peg to the US dollar, they do not face inventory holding risks. In such cases, traders tend to seize arbitrage opportunities more readily. As a result, the price on Uniswap tends to align closely with that on CEXs.

This paper aims to contribute to the growing literature on DEXs. Existing studies have explored price discovery, liquidity provision, the pricing curve, and the dynamics of trader behavior on DEXs. They provide valuable insights into the functioning and challenges of DEXs in the blockchain ecosystem. Pioneering research led by Angeris et al. (2020, 2021a, 2021b, 2021c, 2021d, 2021e, 2022) has played a pivotal role in advancing our comprehension of Constant Function Market Makers (CFMMs). They underscore the critical importance of comprehending CFMM by meticulously examining the structure of their trading sets and exploring the non-arbitrage assumption, particularly in the context of Uniswap, using agent-based simulations.

Additionally, empirical studies provide invaluable insights into the operational dynamics of DEXs. For instance, Lo and Medda (2021) employ an Autoregressive Dynamic Lag Error Correction Model (ARDL-ECM) to examine the cointegrating relationship between Uniswap's ETH-USDT exchange rate and that of centralized exchanges. Lehar and Parlour (2021) find that price differences between Binance and Uniswap are typically below 1 percent in the presence of abundant liquidity. The literature also delves into the concept of DEXs discovering fundamental asset prices. Han et al. (2022) find that the contribution of Uniswap in determining the fundamental value of cryptocurrencies increases as its liquidity provision increases.

Aoyagi (2020), Capponi and Jia (2021), and Lehar and Parlour (2021) contribute to the platform economics of DEXs. They analyze the rent of traders and liquidity providers participating in either DEXs or CEXs and dissect the dynamics governing the distribution of liquidity among these competing marketplaces. Heimbach et al. (2021) analyze Uniswap data and find that liquidity provision follows a power law, with most liquidity providers participating in a few pools. Saleh et al. (2022) suggest that moderate fee rates stimulate liquidity provision, while excessive rates can deter traders. According to Krishnamachari et al. (2021), the optimal CFMMs should dynamically adjust curvature based on asset volatility.

Finally, Miner Extractable Value (MEV) has also become a critical topic in the DEX ecosystem. Daian et al. (2019) document the significance of MEV and the potential for time-bandit attacks by miners. Qin et al. (2022) quantifiy MEV's value, revealing its substantial impact on miners' earnings.

Fabi et al. (2023) offer an extensive literature review on Automated Market Makers (AMMs), consolidating insights and knowledge within this critical component of the DEX landscape. However, to the best of my knowledge, there has been no prior research investigating the impact of inventory carrying costs on the price efficiency of DEXs. The market microstructure literature has extensively documented the importance and

the impact of inventory holding costs on price movements in traditional financial markets. Smidt (1971), Garman (1976), Amihud and Mendelson (1980), Stoll (1978), Ho and Stoll (1981, 1983), among many others, have shown that market makers are not merely passive liquidity providers, as they actively seek to maintain a desired inventory position. Should they deviate from this position, they will adjust the bid-ask spread in order to return to their desired inventory level.

Therefore, it seems natural to conjecture that inventory holding costs impact Uniswap's price dynamics and its price accuracy. By incorporating these costs, one can more accurately characterize the price dynamics on Uniswap V2 and gain further insight into its market microstructure. Specifically, Uniswap's reputation as a reliable oracle, providing real-time price data to diverse DeFi applications, hinges on maintaining a close alignment of its prices with reference values. Consequently, developers can use these insights to fine-tune protocol parameters, ultimately leading to improve efficiency and stability within DeFi. Additionally, comprehending these dynamics is crucial for traders and liquidity providers, enabling them to make more informed decisions and effectively manage their risk exposure.

While the model developed in this paper draws on the traditional market microstructure literature, it differs from it notably. The traditional inventory models rely on market makers' role as price setters, while in the case of Uniswap, liquidity providers have a passive role. This is due to the protocol's design where liquidity providers do not adjust the price. They will only affect the price impact of trades, which decreases as the liquidity provision of that pool increases. The exchange rate changes only when swapping one asset for another. Consequently, the inventory risk in this model is carried by traders rather than liquidity providers as in traditional models. Therefore, this paper focuses on the trader's decision problem to characterize price dynamics on Uniswap, with the execution price reflecting the liquidity provision of the pool.

The paper is organized as follows. Section 2.2 describes the Uniswap V2 protocol. Section 2.3 presents the microstructure model of price dynamics on Uniswap with inventory. Section 2.4 describes the data and Section 2.5 presents estimation results. Section 2.6 concludes. Additional insights into other trading pairs are presented in Appendix 2.7.

2.2 Description of Uniswap

This section provides an overview of Uniswap's functioning and the market participants it interacts with. Uniswap is a decentralized cryptocurrency exchange protocol that enables peer-to-peer transactions without the need for a centralized intermediary. It operates using a Constant Product Market Maker (CPMM) model, where liquidity providers and traders interact through liquidity pools. These liquidity pools consist of pairs of cryptocurrencies, such as *USDC-ETH*, represented by the reserve vector $R \in \mathbb{R}^2_+$, where R_i denotes the

quantity of the $i-th$ asset held in the pool. Furthermore, pools are governed by a continuous trading function, also known as the invariant, denoted by $f : \mathbb{R}_+^2 \to \mathbb{R}_+$. This function maps the reserve vector R of the two assets held in the liquidity pool to a real number K, and is defined as

$$f(R) = R_1 R_2.$$

on Uniswap's CPMM model, the price of a token 1 relative to a token 2 is determined by the current ratio of the reserves of the two tokens in the liquidity pool, i.e. $p_{1,2}(R) = R_2/R_1$.

Liquidity Providers. Liquidity providers play a crucial role in determining the size of the liquidity pool and the trading function value by depositing a proportionate amount of each token in a trading pair into the liquidity pool based on the current exchange rate. In return, they earn a share of the transaction fees generated from trades that occur within the pool. Liquidity providers have the flexibility to withdraw their liquidity at any time. More formally, let $I \in \mathbb{R}_+^2$ be the vector of reserves that an agent inputs into the liquidity pool and $O \in \mathbb{R}_+^2$ the reserves that the protocol outputs to the agent. When the liquidity provider adds liquidity to the pool, they do so in proportion to the existing reserves, resulting in $I = R(\alpha - 1) > 0$ and $O = 0$, where R is the current reserve and $\alpha \geq 1$ is a scalar that determines the proportion. Consequently, the updated reserves are given by $R' = R + I = \alpha R > R$. Conversely, when the liquidity provider removes liquidity from the pool, they do so in proportion to the existing reserves, resulting in $I = 0$ and $O = R(1 - \alpha) > 0$, where $\alpha \in [0, 1]$. In this case, the updated reserves are given by $R' = R - O = \alpha R < R$. This change in reserves updates the trading function from $f(R)$ to $f(R')$, without affecting the exchange rate.

Traders. Traders, on the other hand, exchange one cryptocurrency for another against the liquidity pool. The output amount O that a trader receives when trading is calculated based on their input amount I and the reserve R such that the trading function remains *constant* at its pre-trade value, expressed as $f(R - O + I) = f(R)$. Therefore, the quantity O_1 of asset 1 a trader will receive when putting I_2 of asset 2 to the pool is given by

$$O_1 = I_2 \frac{R_1}{(R_2 + I_2)}$$

and the marginal execution price of asset 1 paid by the trader is then

$$p_{1,2}(\Delta, R) = \frac{I_2}{O_1}$$

where $\Delta \equiv I - O \in \mathbb{R}^2_+$ denotes the net output vector. Finally, the trade results in a new spot price of

$$p_{1,2}(R') = \frac{R_2 + I_2}{R_1 - O_1}.$$

The level curve of the constant product function, which represents all possible combinations of the two assets that would result in the same value of f, is commonly referred to as the bonding curve. Graphically, the interaction between users and the protocol can be summarized as traders moving reserves along the bonding curve (keeping the K constant) while liquidity providers shift the constant product function to a different bonding curve (keeping the ratio of the two reserves constant) as depicted in Figure 2.1.

Transaction fees. In practice, Uniswap charges traders a transaction fee, $1-\gamma \in [0, 1]$, that is proportional to their trades. In this case, the quantity O_1 of asset 1 a trader will receive when putting I_2 of asset 2 to the pool is determined by accounting for γI of the input asset in the trading function, meaning the trade must satisfy the following condition:

$$f(R - O + \gamma I) = f(R).$$

Therefore, the quantity O_1 of asset 1 a trader will receive when putting I_2 of asset 2 to the pool is given by

$$O_1 = \gamma I_2 \frac{R_1}{(R_2 + \gamma I_2)}.$$

As a result of the fees, the constant product is no longer fixed, but rather increases after each trade, shifting the bonding curve upwards. This is because each trade incurs a fee of $(1 - \gamma)I$, which is added to the reserves of the liquidity pool, resulting in a new reserve vector $R' = R - O + I > R - O + \gamma I$. The trading function f is strictly increasing as $\frac{\partial f(R)}{\partial R_1} = R_2 > 0$. Therefore, it follows that

$$f(R') > f(R - O + \gamma I) = f(R).$$

Let's consider an example of trading in the *DAI-ETH* pool. Assume there are 100,000 *DAI* and 100 *ETH* in the pool, and a trader wants to sell 10 *ETH* for *DAI*. This can be represented as $R = (100,000; 100)$, with the trading function value equal to $f(R) = 100,000 \times 100 = 10,000,000$.

When the trader sells $I_{ETH} = 10$ *ETH*, she will receive $O_{DAI} = \gamma 10 \frac{100,000}{100 + \gamma 100} = 9,066.11$ *DAI*, paying 0.0011 *ETH* per *DAI*[1]. The new reserves become $R' = (90,933.89; 110)$, and the new trading function value is $f(R') = 90,933.89 \times 110 = 10,002,727.9$.

[1] $p^u_{trans,t+1} = \frac{10}{9,066.11} = 0.0011$

The no-arbitrage equilibrium without market frictions. In the absence of market frictions, whenever the prices of assets on Uniswap differ from their reference market prices, traders can exploit the price differences by buying or selling assets on Uniswap. This trading activity triggers a price adjustment mechanism that drives the prices of assets on Uniswap towards their reference market values.

The no-arbitrage equilibrium in case of transaction fees. In the presence of transaction fees, the process of price convergence towards the no-arbitrage equilibrium on Uniswap is affected in the following ways: (i) Traders do not close the gap between the Uniswap price and the reference market price entirely. (ii) The degree of gap reduction is non-linear in the initial price difference, and the gap reduction is greater when the initial gap is higher. (iii) The higher the transaction fees, the less likely traders are to exploit arbitrage opportunities and reduce the price gap.

Reference Markets. CEXs prices are commonly considered as reference prices, primarily due to their substantial trading volumes and liquidity. These attributes contribute to price stability and reduce vulnerability to abrupt fluctuations and potential price manipulation, establishing CEX prices as excellent reference points for asset pricing.

2.3 Microstructure model of price dynamics on Uniswap

This section presents a theoretical model of price dynamics on Uniswap, assuming no transaction fees for simplicity, as they are negligible (0.03%) and do not significantly impact the results. I consider a multi-period framework where individuals trade an asset X for an asset Y on Uniswap at times t = 1, 2, ..., T. For simplicity, token X is assumed to be a risky security and security Y is the numeraire asset. The trading

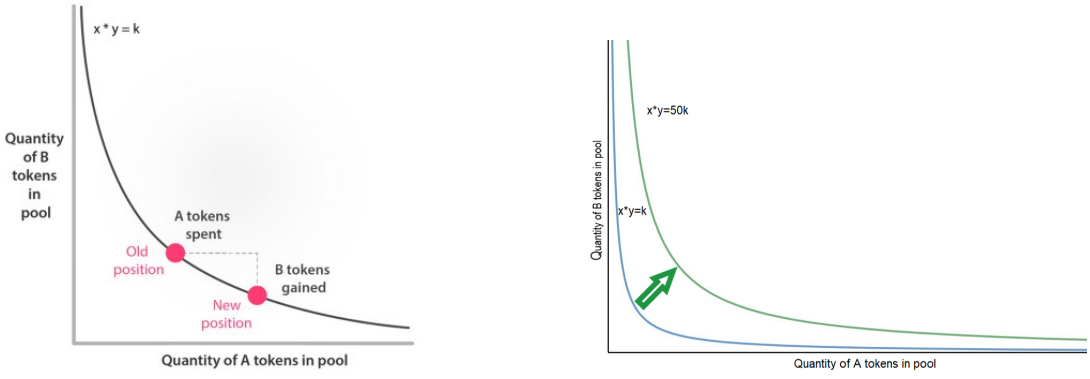

Figure 2.1: Interaction between the Uniswap V2's Constant Product Function and the different market actors.

pair $X-Y$'s underlying value, V_t, is not public and is viewed by traders as a random value, denoted by $\widetilde{V}_t$.

Specifically, it is assumed that only liquidity providers and arbitrageurs operate on Uniswap, and prior to any trade, they observe the price of the trading pair on both Uniswap and the CEX, denoted as P_t and P_t^C, respectively. The trader's information set about the fundamental value just before trading is then given by $\Omega_t = (P_t^C, P_t)$. Since the CEX price is the reference price, we have $E[\widetilde{V}_t|\Omega_t] = P_t^C$.[2]

Once the trader has observed the Uniswap and CEX prices, she submits her trade on Uniswap to maximize her utility. The utility function is expressed as a linear combination of the mean and variance of her wealth, subject to an inventory holding cost:

$$u(t, W_{i,t}) = E\left[W_{i,t}\right] - \frac{\rho}{2} Var\left[W_{i,t}\right] - \phi(X_{i,t} + Q_t - \overline{X})^2, \tag{2.1}$$

where ρ is the coefficient of absolute risk aversion of the trader, ϕ is the inventory holding cost, $X_{i,t}$ is the trader's asset inventory before the trade and $\overline{X}$ is her desired inventory level. The trader's wealth W is given by

$$W_{i,t} = (X_{i,t} + Q_t)\widetilde{V}_t - Q_t P_E(Q_t) + C_{i,t},$$

where $C_{i,t}$ denotes the cash holding (in ETH), Q_t the size of the trade and $P_E\left(Q_t\right)$ the execution price of the trade.[3]

The utility function (2.1) captures the trade-off between maximizing profits and managing inventory levels. By considering both the mean and variance of wealth, the trader takes into account not only the expected outcome of the trade but also the level of risk involved. Furthermore, the inventory holding cost in the utility function reflects the fact that traders may be reluctant to hold inventory that deviates from their desired level, with the parameter ϕ capturing the trader's sensitivity to such deviations.

Both arbitrageurs and liquidity traders aim to buy and sell assets quickly and efficiently without holding onto inventory for extended periods. Indeed, arbitrageurs seek to profit from price discrepancies in different markets for the same asset. Their strategy involves buying the undervalued asset in one market and selling it in the overvalued market, with the goal of profiting from the price difference. Since arbitrage trading is typically done on a short-term basis and the arbitrageur does not intend to hold onto the asset, they aim for zero inventory levels. This allows them to avoid taking on any additional market risk, as well as to minimize the impact of any sudden price changes in the asset. Pure liquidity traders, on the other hand, are typically not interested in holding any inventory of the assets they trade for an extended period, as their goal is simply to facilitate their own liquidity needs. Therefore, it is likely that they will both employ a zero inventory level

[2] I will show in Section 2.4.2 that this is a reasonnable assumption.

[3] Q_t is positive if he decides to buy the asset X in exchange for the asset Y and negative if he decides to sell it.

strategy (i.e. $\overline{X} = 0$).

The present model incorporates inventory holding costs in contrast to the model without market frictions presented in Angeris et al. (2021e), which considers only available arbitrage opportunities to determine the trader's utility and consequently the exchange quantity. Indeed, Angeris et al. (2021e) argue that traders can make a profit by simply buying the asset in the market where the price is low and selling it on the market where the price is high, thereby driving the Uniswap's price back to the reference market price. The present model requires traders to balance arbitrage profits with the additional costs of carrying inventory to optimize their utility.

Moreover, the current model is based on inventory models in the market microstructure literature, but it differs in that it focuses on the inventory problem of liquidity traders rather than liquidity providers to model the price dynamics on Uniswap. This is because, in contrast to traditional inventory models, liquidity providers on Uniswap are passive and do not influence the exchange rate. Instead, they only affect the price impact by adjusting the quantity supplied for the two assets in the pool, which is reflected in the execution price in my model.

Using (2.1), the trader chooses the quantity that maximizes her utility, as follows:

$$\max_{Q_t} \left\{ (X_{i,t} + Q_t)P_t^C - Q_t P_E(Q_t) + C_{i,t} - \frac{\rho\sigma_V^2}{2}(X_{i,t} + Q_t)^2 - \phi(X_{i,t} + Q_t - \overline{X})^2 \right\}. \tag{2.2}$$

Solving problem (2.2) combined with the Uniswap's properties, the FOC can be written as:

$$P_t \left(\frac{1}{1 - \frac{Q_t^*}{R_t^X}} \right)^2 = P_t^C - \rho\sigma_V^2(X_{i,t} + Q_t^*) - 2\phi(X_{i,t} + Q_t^* - \overline{X}) \tag{2.3}$$

where R_t^X is the reserve size of asset X on Uniswap's pool X-Y.

Equation (2.3) is non-linear. To simplify it, I take a Taylor expansion around zero as $Q_t^*/R_t^X \approx 0$ in practice:[4]

$$\left(\frac{1}{1 - \frac{Q_t^*}{R_t^X}} \right)^2 \approx 1 + 2\frac{Q_t^*}{R_t^X} \quad \text{when } \frac{Q_t^*}{R_t^X} \to 0.$$

Explicit solution. Using the above linear approximation, the FOC (2.3) can be rewritten as

$$P_t \left(1 + 2\frac{Q_t^*}{R_t^X} \right) = P_t^C - \rho\sigma_V^2(X_{i,t} + Q_t^*) - 2\phi(X_{i,t} + Q_t^* - \overline{X}),$$

[4]See appendix 2.7 for an empirical justification of this assumption.

hence,

$$Q_t^*\left(P_t^C, P_t, X_{i,t}, \overline{X}, R_t^X\right) = \frac{P_t^C - P_t - \rho\sigma_V^2 X_{i,t} - 2\phi(X_{i,t} - \overline{X})}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}.$$

Putting this expression back into the FOC, I finally get the following characterization of price dynamics on Uniswap:

$$\Delta P_{t+1} \quad = \quad \beta_{0,t}\overline{X} + \beta_{1,t}\left(P_t^C - P_t\right) + \beta_{2,t}(X_{i,t} - \overline{X}) \tag{2.4}$$

where $\beta_{1,t} \equiv 1 - \left(\frac{\rho\sigma_V^2 + 2\phi}{\rho\sigma_V^2 + 2\frac{P_t}{R_{X,t}} + 2\phi}\right)$, $\beta_{0,t} \equiv -\rho\sigma_V^2\beta_{1,t}$, and $\beta_{2,t} \equiv -(\rho\sigma_V^2 + 2\phi)\beta_{1,t}$.

In equation (2.4), $\beta_{1,t}$ captures the arbitrage opportunity and is positive since traders will close the price gap between the centralized and decentralized markets. Everything else being equal, when the asset's price on the CEX is higher than on Uniswap, the trader will buy on Uniswap, increasing the Uniswap price ($\Delta P_{t+1} > 0$). Conversely, when it is lower, the trader will sell on Uniswap, decreasing the Uniswap price ($\Delta P_{t+1} < 0$).

However, $\beta_{1,t}$ decreases as the inventory carrying cost ϕ increases, indicating that traders are less likely to reduce the price gap between the CEX and Uniswap. This suggests that arbitrageurs may be reluctant to take advantage of arbitrage opportunities that require them to deviate significantly from their desired inventory level. Furthermore, the less risk averse the trader (low ρ) and/or the less volatile the asset (low σ_V^2), the more the trader will close the gap between the CEX and Uniswap.

Additionally, as the liquidity provision in the pool increases, the trade's impact on the Uniswap price decreases. Therefore, the trader needs to sell or buy a larger amount of the asset to bring the Uniswap price back to the reference market price. The inventory carrying costs cause the trader to marginally close the gap, even though there is an arbitrage opportunity. Indeed, as R_t^X goes to ∞, $\beta_{1,t}$ converges to 0. It is worth noting that in a frictionless market, traders should consistently close the price gap entirely to exploit arbitrage opportunities. Consequently, it would be expected that the parameter $\beta_{1,t}$ equals 1. However, in the present model, they do not achieve full closure, as indicated by $\beta_1 < 1$.

The coefficient $\beta_{2,t}$ represents the inventory effect and is therefore negative. Intuitively, if the trader's inventory before trading is greater than her desired level, she will sell some of asset X to adjust her inventory towards its target level, causing the Uniswap price to decrease, i.e., $\Delta P_{t+1} < 0$. Conversely, if her inventory is lower than the desired level, she will buy some of asset X, increasing the Uniswap price, i.e., $\Delta P_{t+1} > 0$. Similar to traditional markets, if traders on Uniswap deviate from their desired inventory level, they adjust the price but they will do so through the exchanged quantity. The magnitude of this inventory effect is

amplified with higher inventory costs (represented by a larger ϕ), as well as with greater risk aversion on the trader's part and higher volatility of the asset.

2.4 Empirical Evidence

2.4.1 The Data

Two datasets are used in this paper: one-minute candlestick data from Binance and transaction data from Uniswap. The data have been collected for multiple trading pairs (*ETH-BTC, AAVE-ETH, USDC-USDT, USDT-DAI, USDT-ETH*). The primary emphasis within the paper revolves around one specific trading pair, namely, the wrapped Ether (*ETH*) - wrapped Bitcoin (*BTC*) pair. The findings for other pairs are presented in the appendix 2.7, serving as a means to validate and corroborate the results observed for this particular pair.

Uniswap Data

The Uniswap data were extracted from The Graph, a decentralized API that inDEXs and queries Ethereum data. The dataset contains all transactions which occurred within the *ETH-BTC* pool between May 2020 and December 2022.

There are three types of transactions on Uniswap:

- Mint event, referring to a liquidity provision by a liquidity provider within a pool.
- Burn event, referring to a liquidity withdrawal by a liquidity provider within a pool.
- Swap event, referring to an exchange of an amount of one token for another by a trader within a pool.

When a user submits a transaction to Uniswap, it is first broadcast to the Ethereum network and included in a pool of pending transactions, commonly known as the mempool. Miners on the Ethereum network then pick up transactions from the pool and include them in the next block they mine. The block interval on the Ethereum network is currently around 13-15 seconds, meaning that new blocks are added to the blockchain roughly every 13-15 seconds. The speed at which a Uniswap transaction is processed and confirmed depends on several factors, including the gas price paid for the transaction and the current level of network congestion. During periods of high network activity, transaction processing times may be longer.

For each transaction type within a specific trading pair, the data collected include the issuer, the amount exchanged/minted/burnt, the timestamp, the block number, and the log-index, which represents the order

of the transaction within the block where it was validated.[5] Using this data, the reserves of the two tokens in the *ETH-BTC* pool were reconstructed and updated after every event, inferring the adjustment of the Uniswap price following a swap.

Binance Data

For the centralized market, Binance was selected as it is the most popular CEX. The 1-minute price data for this market, including the open price, closing price, lowest price, highest price, and total traded volume, were obtained from Kaggle.

The data from Uniswap and Binance were then matched by using the minute closing price for both markets. In the case of Uniswap, the minute closing price corresponds to the price that followed the last transaction of each minute.

2.4.2 Descriptive Statistics

Trading activity within Uniswap

From August 2020 to September 2021, there was a notable increase in trading activity on Uniswap V2, as demonstrated by Figure 2.2a and Figure 2.2b. Both transaction volume and value (measured in USD) experienced significant growth during this period. Additionally, the intensity of trading, as indicated by the decrease in intertrade duration, also showed a continuous rise until September 2021. Prior to September 2020, the average intertrade duration ranged from 10 to 45 minutes. It subsequently dropped to less than 1 minute and remained at that level until September 2021 (see Figure 2.6 in Appendix 2.7). The primary driver behind this surge in trading activity can be attributed to the UNI token airdrop that occurred in September 2020. Uniswap launched its governance tokens, the UNI tokens, and distributed them via airdrop to users who had interacted with the protocol before September 1st, 2020. The airdrop started on September 18th, and as users claimed their UNI tokens, many chose to trade them on Uniswap. The airdrop was a pivotal event for Uniswap, as it motivated users to engage with the platform and contributed to its growth and adoption. However, starting from September 2021, there has been a decline in trading activity and pool reserves on Uniswap V2, with levels returning to those seen before August 2020. Despite this decline, the volume exchanged remained ten times higher than before August 2020. The decrease in trading activity and pool reserves post-September 2021 can be attributed to the launch of Uniswap V3 in May 2021. Uniswap V3 introduced concentrated liquidity pools, offering potential advantages over the V2 model.[6] Consequently,

[5] As some transactions happen simultaneously, relying solely on the timestamp is inadequate to sequence them. To arrange transactions that occur concurrently, one needs to use both the block number and the log-index.

[6] Concentrated liquidity enables LPs to offer reserves exclusively within defined price ranges of the trading pair. The benefit for an LP using this option is that, when the exchange rate aligns with the range in which they provide liquidity, they can generate the same fee income as in a standard constant-product pool with lower liquidity supply. For further details, refer to

some liquidity providers may have chosen to migrate their funds from V2 to V3 pools. This migration resulted in reduced activity on Uniswap V2, contributing to the observed decline in pool reserves.

Figure 2.7 in Appendix 2.7 shows a candlestick chart of 1-minute Uniswap price for the ETH-BTC pair where green spikes indicate a higher closing price than the opening price, and red spikes indicate the opposite. The size of the spikes reflects the magnitude of the price movement within that minute. Prior to September 2020, the chart displays large spikes with alternating red and green spikes, indicating a market with high volatility and substantial price movements. However, after September 2020, there was a shift towards smaller spikes, indicating a more stable market. This shift can be attributed to the increase in pool size, resulting in a decrease in price impact and hence smaller price movements (Figure 2.8a in Appendix 2.7). Notably, there was a substantial increase in liquidity provision in mid-September 2020, with the liquidity tripling from its previous levels. It is interesting to note that despite the decrease in the size of the spikes, the alternating pattern of red and green spikes persisted, indicating continued market activity and opportunities for traders. After July 2022, there is a reemergence of punctual high volatility and substantial price movements. This occurrence can be attributed to the fact that the pool was smaller and less popular during this period.

In summary, the candlestick chart in Figure 2.7 demonstrates the transition from a highly volatile market with significant price movements to a more stable market environment with smaller spike sizes, primarily driven by an increase in pool size. However, after July 2022, there is a resurgence of volatility, possibly due to a smaller and less popular pool as users migrated to Uniswap V3.

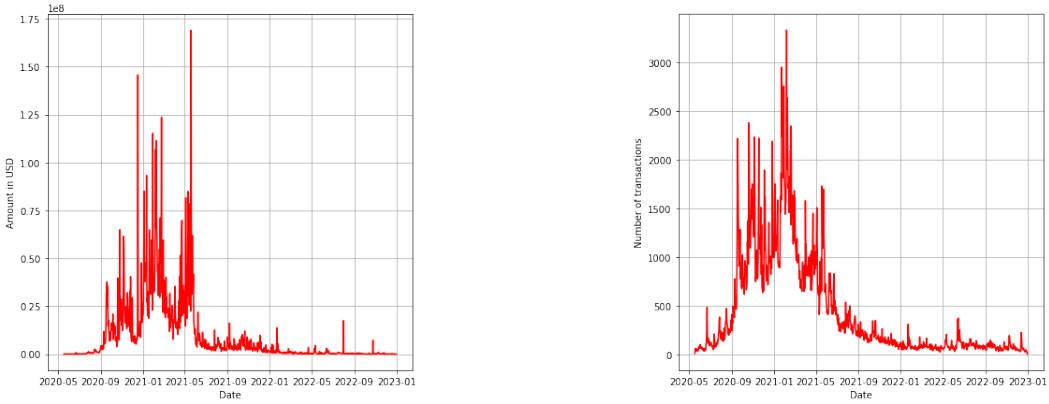

(a) Evolution of volume exchanged in USD

(b) Evolution of the daily number of transactions

Figure 2.2: Trading activity on Uniswap

Uniswap V3 write paper at https://uniswap.org/whitepaper-v3.pdf.

Uniswap Vs. Binance

Figure 2.3 illustrates an interesting pattern prior to January 2022, where price divergences between Uniswap and Binance consistently result in the Uniswap price reverting back to the Binance price. This observation suggests that the Binance price is leading the Uniswap price, indicating that traders perceive the centralized market as the reference market. However, a noteworthy change occurs starting from January 2022. The Uniswap and Binance prices exhibit the same trend but differ in level. This persistence in price discrepancy can be attributed to market inefficiencies, which may stem from various factors including transaction costs and liquidity constraints. These factors hinder traders from efficiently arbitraging price discrepancies between Uniswap and Binance, preventing price convergence and leading to persistent inefficiencies between the two markets. The results of a Granger Causality test presented in Table 2.2 in the Appendix 2.7 provide support

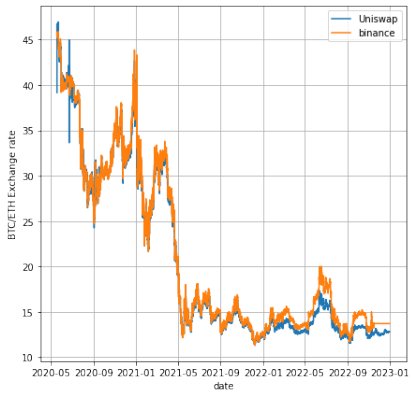

Figure 2.3: Evolution of BTC/ETH exchange rate

for the hypothesis that the Binance price leads the Uniswap price. The test indicates that the Binance price Granger causes the Uniswap price at a significance level of 5%, but the reverse is not true. These findings corroborate the assumption that traders view the centralized market as the reference market and only consider the price on this market as informative about the fundamental value $\widetilde{V}_t$.

The frictionless arbitrage model

Following the approach outlined in Angeris et al. (2019), I fit an arbitrage model without market frictions, where traders aim to maximize their profits by exploiting price discrepancies between Uniswap and the reference market.[7] Subsequently, I compared the model's predictions to the observed prices over time (see Figure 2.8c in Appendix 2.7). It is evident that the frictionless arbitrage model does not closely align with the data. While it does capture the overall trading pattern, it fails to fully explain the price dynamics on Uniswap.

[7]The decision problem faced by the trader is detailed in 2.7.

This suggests that arbitrage opportunities alone are insufficient to account for the price dynamics observed on Uniswap. One potential explanation for this discrepancy is the neglect of market frictions, such as inventory carrying costs. I show in the next section that these frictions play a significant role in shaping the price dynamics on Uniswap and need to be considered for a more accurate analysis.

2.5 Estimation

2.5.1 Reduced form estimation

The microstructure model outlined in Section 2.3 characterizes the price dynamics on Uniswap, which takes into account arbitrage opportunities and the carrying cost associated with the trader's inventory. Unfortunately, we can neither directly observe the trader's inventory $X_{i,t}$ nor her desired level $\overline{X}$ in the available data. Furthermore, deviations from the desired level, denoted as $(X_{i,t} - \overline{X})$, are likely to be correlated with the price difference between Uniswap and the reference markets, as β_2 is a function of β_1. To adress this issue and estimate β_1 and β_2 without introducing any bias, I employ a first-difference approach.

First, I assume that the parameters $\beta_{0,t}$, $\beta_{1,t}$ and $\beta_{2,t}$ are constant over time. This is not a restrictive assumption because the only time-varying component of the parameters is the price-reserve ratio $\frac{P_t}{R_t}$, which, as shown in Figure 2.8b, tends to stay relatively constant over extended periods and hovers near zero. As a result, parameters can be expressed as $\beta_{0,t} = \beta_0$, $\beta_{1,t} = \beta_1$, and $\beta_{2,t} = \beta_2$.

Equation (2.4) can then be rewritten as:

$$\Delta P_t = \beta_0 \overline{X} + \beta_1 \left(P_t^C - P_t\right) + \beta_2 (X_{i,t} - \overline{X}) + \varepsilon_t.$$

wher ε_t is distributed normal with a mean vector of zero and covariance matrix Ω.

Taking the first-difference, I obtain:

$$\Delta P_{t+1} - \Delta P_t = \beta_1 [\Delta P_t^C - \Delta P_t] + \beta_2 [X_{i,t} - X_{i,t-1}] + \Delta \varepsilon_t. \tag{2.5}$$

Finally, $[X_{i,t} - X_{i,t-1}]$ is the quantity $Q_{i,t}$ exchanged on Uniswap at period t. As I consider 1-minute data, I assume that the quantity exchanged between two consecutive periods is the total exchanged quantity over that minute. Therefore, equation (2.5) is rewritten as:

$$\Delta P_{t+1} - \Delta P_t = \beta_1 [\Delta P_t^C - \Delta P_t] + \beta_2 Q_t + \Delta \varepsilon_t. \tag{2.6}$$

To assess the presence of autocorrelation in the error terms, a Ljung-Box Test was conducted (see Table 2.2). The null hypothesis that the errors are not correlated is rejected at the 5% level, indicating autocorrelation. To address this issue, the model is estimated using OLS with Newey-West robust estimation of the variance-covariance matrix. Rolling window estimation on daily basis of the reduced form model (2.6) was performed to observe the potential changes in the arbitrage opportunity effect over time. The results are presented in Figure 2.4 and, as explained below, strongly support the theoretical model developed in Section 2.3.

The arbitrage opportunity effect

The coefficient associated with the price difference between Binance and Uniswap is statistically lower than 1, rejecting the model without inventory costs of Angeris et al. (2019) in favor of model (2.4). Additionally, the effect of the arbitrage opportunity is decreasing over time and gradually approaching zero.

In the initial stages of the pool, traders effectively closed almost the entire price gap between Binance and Uniswap. However, over time, their ability to close the gap gradually declined. By August 2020, they were only able to close about 20% of the price gap, which further declined to 10% by October 2020. This pattern corresponds with the substantial increase in liquidity provided within the ETH-BTC pool, which, in turn, led to a decrease in the price impact within the Uniswap market. These findings are in line with the conclusions drawn from the inventory cost model discussed in Section 2.3. This model suggests that as the size of the pool grows, trades must be more substantial to influence the Uniswap price and narrow the gap. Consequently, traders are required to hold larger inventories, which increases their exposure to risks. Consequently, traders become less inclined to respond to price disparities between the two markets, even when profit opportunities arise.

At the beginning of 2022, traders once again made efforts to reduce the price gap. This resurgence coincided with an increase in the price impact as the pool size decreased. However, during this period, they were only able to, on average, close approximately 30% of the gap. This outcome is likely a consequence of reduced trading activity on Uniswap V2 due to the transition to Uniswap V3.

The inventory effect

The daily estimation of β_2 is displayed in Figure 2.4b. Initially, when Uniswap V2 was launched, a trader with an inventory surplus of one unit would engage in a sale, resulting in a 0.002-point drop in the exchange rate. Nevertheless, as time passed, this impact waned, becoming negligible. Even if the trader were to conduct trades to bring her inventory in line with her desired level, the influence on the price would remain minimal due to the large size of the pool, resulting in a negligible price impact from her trades.

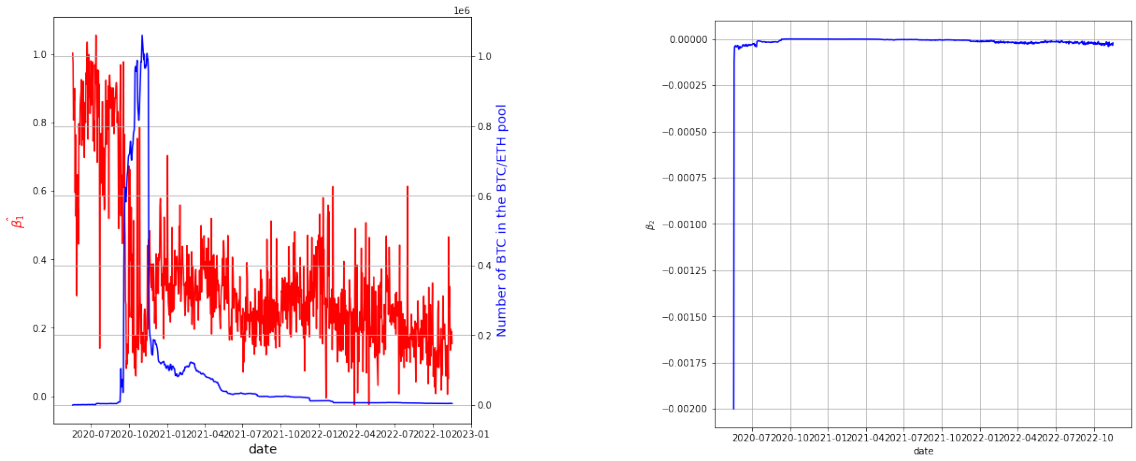

(a) Daily estimation of β_1

(b) Daily estimation of β_2

Figure 2.4: Daily estimation of parameters β_1 and β_2.

2.5.2 Measuring the implicit inventory cost on the no-arbitrage assumption

In section 2.5, my primary focus has been on estimating the arbitrage effect, a crucial element that ensures the Uniswap price closely tracks the reference price. The results of this estimation have revealed an important trend: the estimated arbitrage effect is consistently below one and displays a decreasing pattern over time, suggesting that traders are not effectively closing the entire price discrepancies between Binance and Uniswap.

This pattern coincides with the growth of the liquidity pool on Uniswap. Therefore, I have postulated that this phenomenon can be attributed to the presence of inventory holding costs. As the Uniswap liquidity pool grows in size, larger trades are necessary to exert a meaningful influence on the Uniswap price and, consequently, to narrow the price gap. Accommodating these larger trades necessitates traders to maintain more substantial inventory positions. This heightened exposure to inventory holding risk seems to make traders less responsive to price disparities between the two markets, even when such disparities present profitable opportunities. However, it is important to note that this assertion is based on a conjecture stemming from the reduced form estimation of β_1 and the simultaneous evolution of the pool size.

In this section, my focal point is to dissect each element within β_1 and isolate the influence of inventory costs ϕ, aiming to validate my initial hypothesis. I will delve into whether these inventory costs indeed serve as a deterrent for traders, hindering their active participation in arbitrage opportunities.

Estimating the volatility of Binance Price

I estimate the Binance Price volatility, σ_v, using the local average version of the Optimal candlesticK (OK) estimator proposed by Jia Li and Zhang (2021). Their method takes into account all information in a candlestick, including the range between high and low prices in addition to the return calculated by the

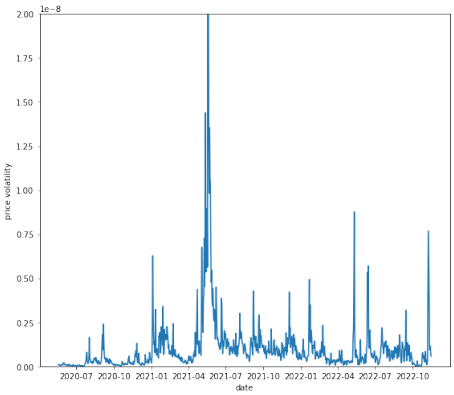

Figure 2.5: OK estimator of the Binance Price Volatility

closing and opening prices, resulting in a more accurate estimator than those that rely solely on returns. The single-candlestick estimator is computed using the following formula:

$$\frac{0.811 \times (\text{High} - \text{Low}) - 0.369 \times |\text{Close} - \text{Open}|}{\text{Duration of the trading session}^{0.5}}.$$

Additionally, when higher frequency candlesticks are available over fixed time interval $[0, T]$, they suggest improving this estimator with a local average method over a time window k. The local average estimator is expressed as:

$$\widehat{\sigma}_t(k) = \frac{1}{k} \sum_{t=1}^{k} \frac{0.811 \times w_{i+t} - 0.369 \times |r_{i+t}|}{\Delta_n^{0.5}}, \text{ for t} \in [(i-1)k\Delta_n, ik\Delta_n], i \in \left[1, \left\lfloor \frac{T}{k} \right\rfloor\right]$$

where Δ_n is the duration of the trading session, $w_{i+j} = \text{High price}_{i+t} - \text{Low price}_{i+t}$ and $r_{i+t} = \text{Closing price}_{i+t} - \text{Open price}_{i+t}$.

Given that I have 1-minute candlesticks for Binance price data, I apply this estimator with a daily estimation window by setting $\Delta_n = 1$ minute and $k =$ daily number of observations. Figure 2.5 depicts the volatility estimation of the *ETH-BTC* trading pair on Binance over time. The plot reveals that the exchange rate of the *ETH-BTC* pair is subject to many fluctuations, but with a relatively low volatility level. The magnitude of volatility is around 1e-8, suggesting that the changes in price for the *ETH-BTC* pair on Binance are relatively small[8]. Furthermore, comparison with the candlestick chart of Uniswap price illustrated in Figure 2.7 clearly indicates that the price of the *ETH-BTC* pair is significantly more stable on Binance than on Uniswap.

[8]Please refer to Figure 2.3 for the *ETH-BTC* price's range.

Deriving an estimate for the inventory costs

By using the expression for $\beta_{1,t}$ and employing a commonly used risk-aversion parameter $\rho = 4$ as detailed in the literature in the context of mean-variance utility (Ang 2014), it is possible to obtain estimates for the inventory costs ϕ. This estimation is accomplished through ordinary least squares (OLS) regression on the following equation:

$$y_t = 2\phi + \rho\sigma^2_{V_t} \tag{2.7}$$

where $y_t = 2\frac{P_t}{R_{X,t}} \times (1 - \beta_{1,t})/\beta_{1,t}$.

The estimate of ϕ is presented in table 2.1. It is statistically significant and is positive. Therefore, inventory costs prevent traders from exploiting arbitrage opportunities and, in turn, from narrowing the gap between Uniswap and the reference market.

Variable	Estimate	Std. Error
ϕ	0.051587	0.006658 ***
$n = 913$		

Table 2.1: Estimation Results for equation (2.7)

The analysis and model estimation for additional top pairs during the same period can be found in Appendix 2.7. The findings related to these pairs corroborate and complement the conclusions derived from the analysis of the ETH-BTC pair. In particular, when traders are not exposed to market risks, as is the case with stablecoin pairs due to their peg to the US dollar, they do not face inventory holding risk. Consequently, they are more inclined to exploit arbitrage opportunities.

2.6 Conclusion

Uniswap is one of the most widely used decentralized application on Ethereum with more than $1 trillion in total value locked. Yet, its market microstructure and its trading process are not fully understood as Uniswap is an AMM and thereby works differently to traditional exchange markets. Therefore, it is crucial to investigate its market microstructure and its trading process, taking into account market frictions.

Specifically, the price efficiency of Uniswap relies on the arbitrage assumption which implies that any price differences between Uniswap and other exchanges are eliminated through arbitrage activities, resulting in the Uniswap price reflecting the true market value of an asset. However, traders may face challenges when trying to profit from cross-market arbitrage opportunities due to market frictions such as inventory holding costs. In this context, this paper attempts to complement the existing literature on DEXs and gives further insight about DEXs by developing and testing a microstructure model of Uniswap's price dynamics.

My microstructure model not only considers arbitrage opportunities, but also takes into account inventory holding costs, which have been found to be a significant factor affecting traders' decision-making processes. As the inventory costs increases, traders become less willing to take advantage of arbitrage opportunities and are less likely to close the price gap between Uniswap and reference markets. Additionally, our empirical findings provide support for the theoretical model and suggest that as the pool increases, resulting in a decrease in price impact, traders tend to mitigate the price discrepancy between Uniswap and centralized exchanges less, as they face greater exposure to inventory risks. Conversely, when traders operate in pools with minimal market risks, such as stablecoin pairs, they do not encounter inventory holding risks. Consequently, they are more inclined to exploit arbitrage opportunities.

In future work, I aim to incorporate other market frictions into the model. On top of trader's inventory, private information and processing costs are also likely to affect the trading process and thereby, the no-arbitrage assumption on Uniswap. Indeed, when trading on Uniswap, investors may have private information regarding the fundamental value of the asset or observe a signal of it through trades of others. They may then update their belief and trade in consequence. Furthermore, in addition to the 0.3% transaction fees, investors also need to pay mining fees to have their transaction validated and added to the Ethereum Blockchain. These fees are fixed and do not depend on the amount traded. As a result, traders may decide not to take advantage of an arbitrage opportunity if the transaction costs exceed the potential gains, leading to inaction.

Moreover, it is crucial to consider the impact of Miner Extractable Value (MEV) as another significant factor. Traders on Uniswap may adopt a strategy of splitting their trades into smaller segments as a precaution against front-running. Consequently, they may not fully close the price gap in a single transaction.

In summary, my research agenda encompasses the incorporation of market frictions beyond trader inventory, including private information, processing costs, and MEV to gain a more comprehensive understanding of how they collectively shape trading behavior and influence the completeness of arbitrage opportunities within the Uniswap ecosystem.

2.7 Appendix

The no-arbitrage equilibrium

In the context of the optimal arbitrage problem, there are two cryptocurrencies, $X$and Y, which can be exchanged within either a reference market or the Uniswap platform. Within this problem, traders endeavor to optimize their profit by exploiting the difference between the reference market price, denoted as P_t^C, and

the Uniswap price, denoted as P_t. The optimal arbitrage problem can be formulated as follows:

Hence,

$$\begin{aligned} &\max_{\Delta X_t \in \mathbb{R}^+} g(\Delta X_t) \\ &\text{s.t } \Delta Y_t \geq 0 \\ \Delta Y_t \leq R_t^Y \quad &\text{if he sells } \Delta X_t \\ \Delta X_t \leq R_t^X \quad &\text{if he buys } \Delta X_t \end{aligned} \tag{2.8}$$

where

$$g(\Delta X_t) = \begin{cases} g_1(\Delta X_t) = \left[P_E(\Delta X_t) - P_t^C\right] \Delta X_t & \text{if he sells } \Delta X_t \\ \\ g_2(\Delta X_t) = \left[P_t^C - P_E(\Delta X_t)\right] \Delta X_t & \text{if he buys } \Delta X_t \end{cases}$$

The two last constraints require that the quantity bought should be lower than the reserve size.

Since $P_E = \frac{\Delta Y_t}{\Delta X_t}$ and the quantity exchanged on Uniswap is:

$$\Delta Y_t = \begin{cases} \frac{\gamma \Delta X_t R_t^Y}{R_t^X + \gamma \Delta X_t} & \text{if he sells } \Delta X_t \\ \\ \frac{\Delta X_t R_t^Y}{(R_t^X - \Delta X_t)\gamma} & \text{if he buys } \Delta X_t \end{cases}$$

$g(\Delta X_t)$ can be rewritten as follows :

$$g(\Delta X_t) = \begin{cases} g_1(\Delta X_t) = \frac{\gamma \Delta X_t R_t^Y}{R_t^X + \gamma \Delta X_t} - P_t^C \Delta X_t & \text{if he sells } \Delta X_t \\ \\ g_2(\Delta X_t) = P_t^C \Delta X_t - \frac{\Delta X_t R_t^Y}{(R_t^X - \Delta X_t)\gamma} & \text{if he buys } \Delta X_t \end{cases}$$

Solutions to problem (1) are:

$$\Delta X_t^* = \begin{cases} \Delta X_{t,sell}^* = \frac{-R_t^X P_t^C + \sqrt{\gamma P_t^C R_t^Y R_t^X}}{\gamma P_t^C} & \text{if } P_t^C < \gamma P_t \\ \\ 0 & \text{if } \gamma P_t \leq P_t^C \leq \frac{P_t}{\gamma} \\ \\ \Delta X_{t,buy}^* = \frac{\gamma P_t^C R_t^X - \sqrt{\gamma P_t^C R_t^Y R_t^X}}{\gamma P_t^C} & \text{if } P_t^C > \frac{P_t}{\gamma} \end{cases}$$

PROPOSITION 1. When there is no fees and the prices are misaligned, the trader will always take advantage of pure arbitrage opportunities. However, in presence of fees, even when the prices are misaligned, the trader will trade only if the difference is high enough to recover the fees, otherwise, they will not respond

to the price divergence between Uniswap and the centralized exchange, giving rise to inaction.

Using the constant product formula, the new price after the trade is:

$$P_{t+1}^{*} = \begin{cases} \frac{\gamma R_t^Y P_t^C}{(\gamma-1)\sqrt{\gamma P_t^C R_t^Y R_t^X}+\gamma R_t^Y} & \text{if } P_t^C < \gamma P_t \\ P_t & \text{if } \gamma P_t \leq P_t^C \leq \frac{P_t}{\gamma} \\ P_t^C - \frac{(1-\gamma)\sqrt{\gamma P_t^C R_t^Y R_t^X}}{\gamma R_t^X} & \text{if } P_t^C > \frac{P_t}{\gamma} \end{cases}$$

Finally, the ratio between the Uniswap price and the fundamental value (or centralized price) is :

$$\frac{P_t^C}{P_{t+1}^{*}} = \begin{cases} 1-(1-\gamma)\sqrt{\frac{P_t^C}{\gamma P_t}} & \text{if } \frac{P_t^C}{P_t} < \gamma \\ \frac{P_t^C}{P_t} & \text{if } \gamma \leq \frac{P_t^C}{P_t} \leq \frac{1}{\gamma} \\ \left(1-(1-\gamma)\sqrt{\frac{P_t}{\gamma P_t^C}}\right)^{-1} & \text{if } \frac{P_t^C}{P_t} > \frac{1}{\gamma} \end{cases}$$

PROPOSITION 2. (i) When there is no fees, the trader closes the gap between the two markets, no matter the initial price difference. (ii) When there are fees, the trader does not close the gap entirely. Furthermore, the higher the gap (in one direction or another) before the trade, the greater the gap reduction by the trader. (iii) The gap reduction is not linear in the initial price difference. (iv) The higher the fees, the less the trader will exploit arbitrage opportunities and reduce the gap.

Deriving Uniswap's price dynamics

The trader chooses the quantity that maximizes their utility subject to inventory holding costs, as follows:

$$\max_{Q_t} \left\{ (X_{i,t}+Q_t)P_t^C - P_E(Q_t)Q_t + C_{i,t} - \frac{\rho\sigma_V^2}{2}(X_{i,t}+Q_t)^2 - \phi \times (X_{i,t}+Q_t-\overline{X})^2 \right\}. \tag{2.9}$$

where $P_t^C = E\left[\widetilde{V}_t\right]$.

The FOC gives:

$$P'_E(Q_t)Q_t + P_E(Q_t) = P_t^C - \rho\sigma_V^2(X_{i,t} + Q_t) - 2\phi(X_{i,t} + Q_t - \overline{X}).$$

Note that equality [9]

$$P'_E(Q_t)Q_t + P_E(Q_t) = P_{t+1}, \tag{2.10}$$

where P_{t+1} is the new Uniswap price after that trade. Furthermore,

$$P_{t+1} = P_t \left(\frac{1}{1 - \frac{Q_t}{R_t^X}} \right)^2. \tag{2.11}$$

where R_t^X is the liquidity size of asset X at time t. Combining (2.10) and (2.11), the FOC can be rewritten as:

$$P_t \left(\frac{1}{1 - \frac{Q_t}{R_t^X}} \right)^2 = P_t^C - \rho\sigma_V^2(X_{i,t} + Q_t^*) - 2\phi(X_{i,t} + Q_t^* - \overline{X}) \tag{2.12}$$

The expression (2.11) is highly non-linear. To simplify it, I take a Taylor expansion around zero as $Q_t/R_t^X \approx 0$ in practice:

$$\left(\frac{1}{1 - \frac{Q_t}{R_t^X}} \right)^2 \approx 1 + 2\frac{Q_t}{R_t^X} \quad \text{when } \frac{Q_t}{R_t^X} \to 0.$$

Explicit solution. Let us now solve for Q_t explicitly. The FOC (2.12) can be rewritten as

$$P_t \left(1 + 2\frac{Q_t^*}{R_t^X} \right) = P_t^C - \rho\sigma_V^2(X_{i,t} + Q_t^*) - 2\phi(X_{i,t} + Q_t^* - \overline{X}),$$

or

$$Q_t^* \left(P_t^C, P_t, X_{i,t}, \overline{X}, R_t^X \right) = \frac{P_t^C - P_t - \rho\sigma_V^2 X_{i,t} - 2\phi(X_{i,t} - \overline{X})}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}.$$

[9] (2.10) & equality (2.11) come from the constant function properties.

Putting it back in the FOC, we finally get:

$$\begin{aligned}
P_t^C - P_{t+1} &= (\rho\sigma_V^2 + 2\phi)X_{i,t} - 2\phi\overline{X} + (\rho\sigma_V^2 + 2\phi)\frac{P_t^C - P_t - \rho\sigma_V^2 X_{i,t} - 2\phi(X_{i,t} - \overline{X})}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi} \\
&= \left(\rho\sigma_V^2 + 2\phi - \frac{(\rho\sigma_V^2 + 2\phi)^2}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}\right) X_{i,t} + 2\phi\left(\frac{\rho\sigma_V^2 + 2\phi}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi} - 1\right)\overline{X} + \left(\rho\sigma_V^2 + 2\phi\right)\frac{P_t^C - P_t}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi} \\
&= \frac{2\frac{P_t}{R_t^X}(\rho\sigma_V^2 + 2\phi)}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi} X_{i,t} - \frac{4\phi\frac{P_t}{R_t^X}}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}\overline{X} + \left(\rho\sigma_V^2 + 2\phi\right)\frac{P_t^C - P_t}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}.
\end{aligned}$$

Rearranging terms and subtracting both side by P_t, we then get:

$$\begin{aligned}
P_{t+1} - P_t &= \left(1 - \frac{\rho\sigma_V^2 + 2\phi}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}\right)(P_t^C - P_t) - \frac{2\frac{P_t}{R_t^X}(\rho\sigma_V^2 + 2\phi)}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi} X_{i,t} + \frac{4\phi\frac{P_t}{R_t^X}}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}\overline{X} \\
&= \frac{2\frac{P_t}{R_t^X}}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}(P_t^C - P_t) - \frac{2\frac{P_t}{R_t^X}(\rho\sigma_V^2 + 2\phi)}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}(X_{i,t} - \overline{X}) + \frac{4\phi\frac{P_t}{R_t^X} - 2\frac{P_t}{R_t^X}(\rho\sigma_V^2 + 2\phi)}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}\overline{X} \\
&= \frac{2\frac{P_t}{R_t^X}}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}(P_t^C - P_t) - \frac{2\frac{P_t}{R_t^X}(\rho\sigma_V^2 + 2\phi)}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}(X_{i,t} - \overline{X}) - \frac{2\frac{P_t}{R_t^X}\rho\sigma_V^2}{\rho\sigma_V^2 - 2\frac{P_t}{R_t^X} + 2\phi}\overline{X}.
\end{aligned}$$

We finally get the following characterization of price dynamics on Uniswap:

$$P_{t+1} - P_t = \beta_{1,t}(P_t^C - P_t) - (\rho\sigma_V^2 + 2\phi)\beta_{1,t}(X_{i,t} - \overline{X}) - \rho\sigma_V^2\beta_{1,t}\overline{X}$$

where $\beta_{1,t} = \dfrac{2\frac{P_t}{R_t^X}}{\rho\sigma_V^2 + 2\frac{P_t}{R_t^X} + 2\phi}$.

Additional graphs and statistical tests for the ETH-BTC pool

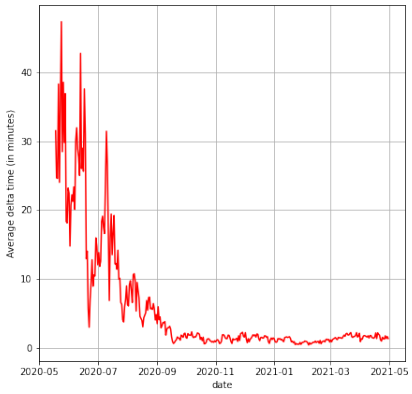

Figure 2.6: Daily average delta time (in minutes) between two transactions on Uniswap

Note: Trading intensity is measured as the time length between two consecutive trades.

Figure 2.7: Candlestick Chart of the BTC-ETH minute exchange rate on Uniswap

Test	Nb of lags	Null hypothesis	Stat	p-value
Ljung-Box test		Residuals are independently distributed	18588.225853	0.0
Granger causality test	3	Binance price does not Granger causes Binance price	624.0059	0.0000
Granger causality test	3	Uniswap price does not Granger causes Binance price	1.3501	0.2561

Table 2.2: Ljung-Box test and Granger causality tests

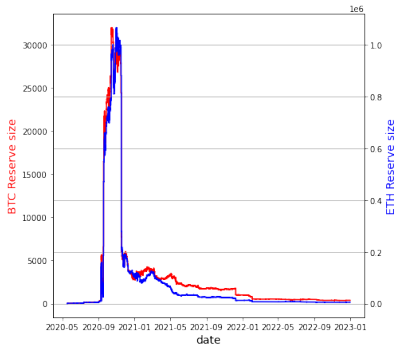

(a) Evolution of ETH and BTC Reserves.

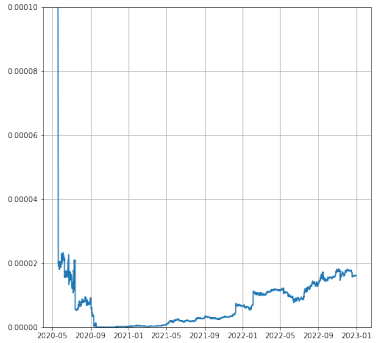

(b) Evolution of the price-reserve ratio $\frac{P_t}{R_t}$.

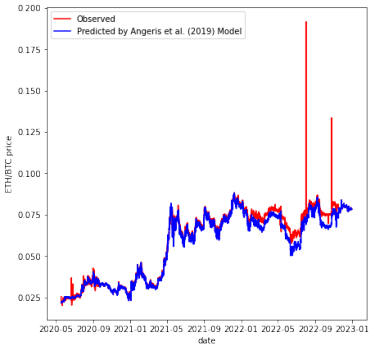

(c) Observed Vs Predicted ETH-BTC price on Uniswap

Figure 2.8: Evolution of the ETH-BTC pool over time.

Other crypto pairs on Uniswap

In this appendix, I broaden the scope of my analysis by examining additional cryptocurrency pairs. The goal is to assess the trading activity within various pools, compare their exchange rates to those on Binance, and apply the model (2.6) to these pairs. This extension aims to determine if the insights derived from the ETH-BTC pair can be extrapolated to other pairs or if the inventory holding cost varies depending on the characteristics of the pair. To achieve this, I gather data for the top pairs during the covered period, specifically: AAVE-ETHER, USDC-USDT, USDT-DAI, and USDT-ETHER. It is worth noting that USDC, USDT, and DAI are U.S. stablecoins, cryptocurrencies designed to maintain their value as closely as possible to the U.S. dollar. The data collection process for those pairs follows the same methodology detailed in section 2.4.1 for the ETH-BTC pair.

Trading activity

Similarly to the ETH-BTC pair, the various cryptocurrency pairs experienced a surge in trading activity beginning in November 2020, followed by a decline starting in July 2021, as depicted in Figure 2.9. The initial increase can be attributed to the Uniswap airdrop, whereas the subsequent decline can be traced back to users migrating towards Uniswap V3. A corresponding shift in intertrade duration for these pairs is illustrated in Figure 2.10, showing a consistent decrease until July 2021, reflecting an increasing trading intensity. However, this trend reversed with the transition to Uniswap V3. The AAVE-ETHER pair stands as an exception, with a declining trading intensity even during the peak of Uniswap V2, before becoming inactive with the launch of Uniswap V3. It is also noteworthy that the trading intensity for the ETH-BTC trading pair remained higher, even in the early stages of Uniswap, compared to these other pairs. The average intertrade duration never exceeded 40 minutes for the ETH-BTC pair, while for the other pairs,

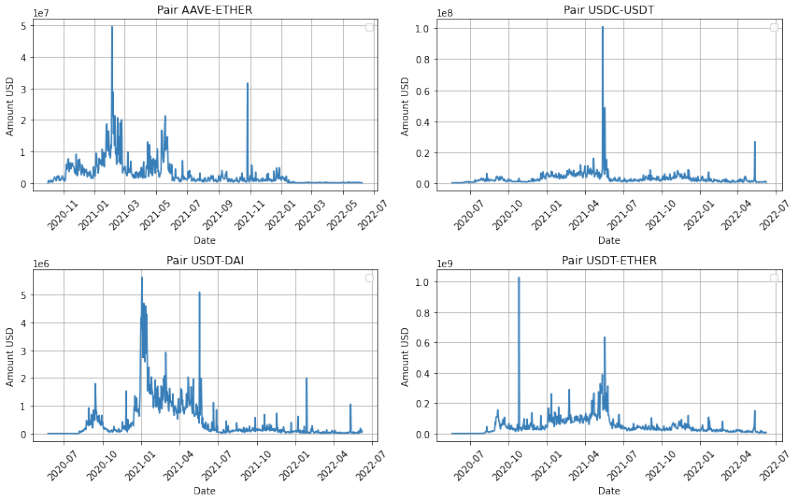

Figure 2.9: Daily Trading Volume (USD).

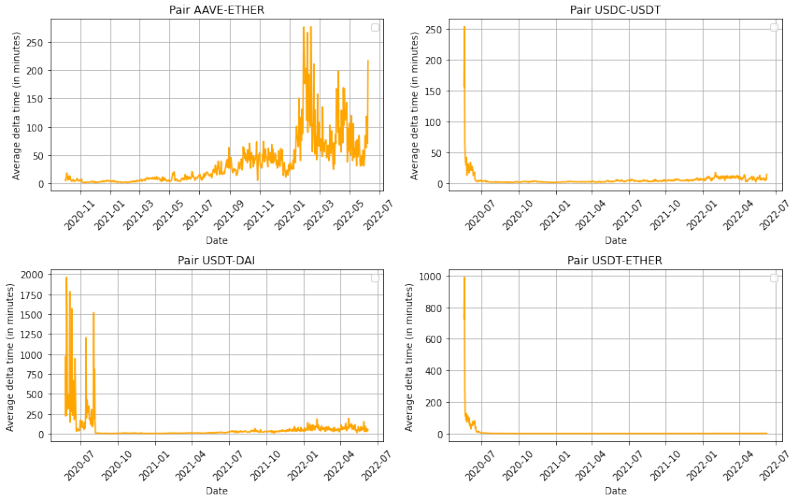

Figure 2.10: Daily inter-trade duration in minutes.

it extended to hours and even days. Shifting our focus to Figure 2.11, it offers insight into the exchange rates of these different crypto pairs, comparing their evolution on both Uniswap and Binance. For pairs like USDC-USDT and USDT-DAI, where stablecoins pegged to the US dollar are involved, the exchange rates should ideally remain at 1. Binance demonstrates remarkable stability in these rates, consistently hovering around 1. In contrast, Uniswap shows some fluctuations around this value. When we focus on the USDT-ETHER pair, we observe a close alignment in exchange rates between Uniswap and Binance. However, delving into the AAVE-ETHER pair, a substantial disparity becomes apparent. Uniswap's exchange rate for the AAVE-ETHER pair deviates significantly from the corresponding rate on Binance. To put this into context, the magnitude of this difference is quite substantial, with Uniswap's exchange rate registering at approximately 1e-19, while Binance's rate remains at 0.1. This marked contrast raises questions about the

potential influence of Uniswap's decentralized nature on the behavior of exchange rates. In summary, trading activity in these pools follows the same pattern as observed in the ETH-BTC pool, with the exception of the AAVE-ETHER pair.

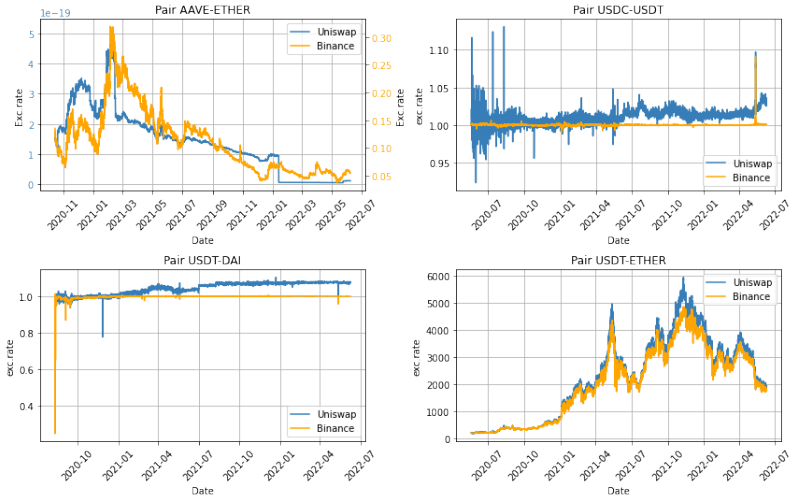

Figure 2.11: Exchange Rate on Uniswap over time.

Model estimation

I estimate equation (2.6) for these trading pairs. The results are presented in Figure 2.12.

Arbitrage opportunity effect

As one might expect, in the case of the AAVE-ETHER trading pair, traders are reluctant to narrow the price disparity between Uniswap and Binance. This is due to the unreliability of the Uniswap price, as evidenced in Figure 2.11, where the Uniswap price significantly diverges from the Binance price. This divergence is exacerbated by a decline in trading activity within this pool, as depicted in Figure 2.9. Similar to the ETH-BTC pair, for the USDT-ETHER pair, traders initially fully close the price gap, but later opt to partially close it, typically narrowing it by around 50%. When it comes to US-stablecoin pairs, traders generally aim to completely close the price gap between Uniswap and Binance. This outcome is not suprising, given that stablecoins are expected to maintain their stability and not exhibit significant volatility. Traders might be more willing to hold inventory, confident that the prices should remain stable within these pools. However, in some instances, they tend to overreact to the price gap and end up narrowing it by even more, up to 120%.

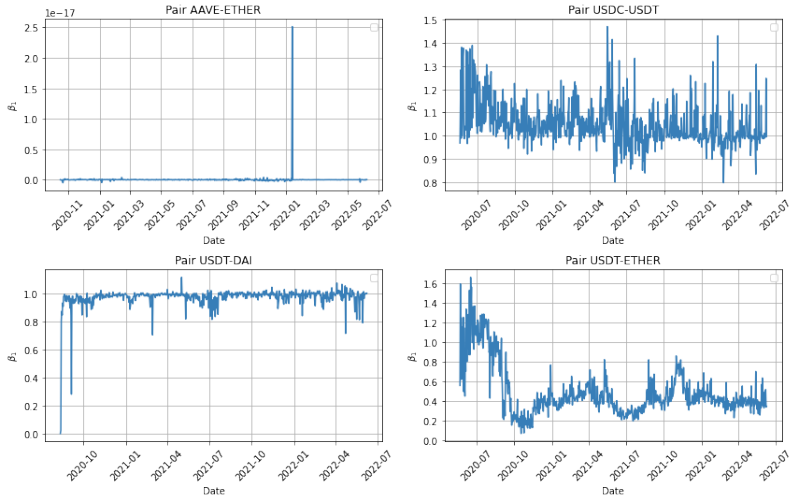

Figure 2.12: Daily estimation of β_1 using OLS with Newey-West robust estimation of the variance-covariance matrix.

Binance price volatility

I employ the OK estimator to compute the Binance Price volatility, represented as σ_v, for the designated pairs, following the same methodology outlined in section 2.5.2 for the BTC-ETHER trading pair. The results are presented in Figure 2.13. It is evident from the Figure that, with the exception of the USDT-ETHER pair, the prices on Binance are stable. This remarkable stability suggests that US-stablecoin pairs are effectively pegged to the US Dollar on Binance, highlighting the reliability of Binance as a pricing reference.

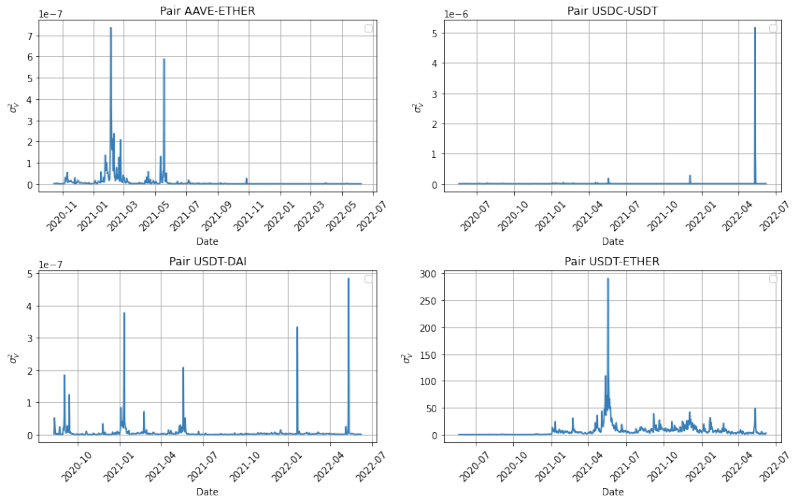

Figure 2.13: OK estimator of Binance price volatility.

Estimating the inventory costs

The inventory cost estimates for each pair are provided in Table 2.3. The effect is statistically significant at the 5% level for all pairs under consideration, except for the DAI-USDT trading pair. Furthermore, for the stable-coin pairs, the effect is nearly approaching zero, consistent with the estimator I found for β_1 (see Figure 2.12). Traders effectively eliminate the price gap between Uniswap and Binance because they do not encounter inventory costs when dealing with these pairs. As previously hypothesized, this phenomenon can be attributed to the stability of these pairs. Stablecoin pairs are expected to exhibit low volatility, enabling traders to take advantage of arbitrage opportunities without being exposed to market risks associated with holding inventory.

In summary, this section complements the findings for the ETH-BTC pair, demonstrating that when the market risk is low, traders do not encounter inventory risks. Consequently, they seize arbitrage opportunities and effectively narrow the price gap between Uniswap and Binance.

pair	Estimate	Std. Error
AAVE-ETHER	-1.982e-08	3.884e-09 ***
USDC-USDT	-5.692e-08	1.741e-08 **
ETH-USDT	-14.198	1.105 ***
DAI-USDT	1.662e-06	1.474e-06

Table 2.3: Estimation of the inventory holding cost ϕ.

Chapter 3

Contagion in Decentralized Lending Protocols: A Case Study of Compound.[1]

Co-written with Natkamon Towanich, Julien Prat, and Simon Weidenholzer.

Abstract

We study financial contagion in Compound V2, a decentralized lending protocol deployed on the Ethereum blockchain. We explain how to construct the balance sheets of Compound's liquidity pools and use our methodology to characterize the financial network. Our analysis reveals that most users either borrow stablecoins or engage in liquidity mining. We then study the resilience of Compound v2 through a series of stress tests, identifying the pools that are most likely to set off a cascade of defaults.

[1]DeFi '23: Proceedings of the 2023 Workshop on Decentralized Finance and Security, November 30, 2023, Copenhagen, Denmark. https://doi.org/10.1145/3605768.3623544.
This project has been funded under the Blockchain@X Research Center. Weidenholzer acknowledges support from the Economic and Social Research Council [grant number ES/T015357/1]. We thank Riho Marten Pallum for helping us develop the code to extract Compound V2 data from The Graph.

3.1 Introduction

Smart contracts have enabled the rise of Decentralized Finance (DeFi) protocols that offer financial services without relying on an intermediary such as a bank or brokerage house. While it is widely acknowledged that traditional financial systems are vulnerable to contagion through various channels, including bank runs Diamond and Dybvig (1983) and default cascades Eisenberg and Noe (2001), little is known about the contagion risks potentially present in DeFi protocols.

To study this question, we focus on Compound V2 but note that alternative protocols, such as AAVE and MakerDAO, share a comparable architecture, potentially exposing them to similar forces. Compound is a decentralized lending protocol built on the Ethereum blockchain (see Leshner and Hayes, 2019). The protocol manages multiple liquidity pools, each dedicated to a specific token. Lenders can add liquidity to any pool, while borrowers can withdraw liquidity by providing collateral in the form of deposits in other pools. These operations connect the various liquidity pools through a network of financial liabilities.

Our first contribution lies in proposing a methodology for the description of Compound's financial network. We do so by characterizing the balance sheet of its liquidity pools and identifying how they are connected by the borrowing and collateral obligations of users. Leveraging the public availability of Ethereum's transaction history, we reconstruct the balance sheet of each pool at any given point in time and without measurement errors. The resulting financial network sheds light on the key functionalities of Compound. Specifically, it indicates that users predominantly utilize Compound for two types of financial operations: borrowing stablecoins and participating in liquidity mining of Compound's governance token.

Then, we assess the resilience of the protocol. Inspired by the recent literature on financial contagion (e.g., Elliott et al., 2014; Jackson and Pernoud, 2021), we investigate how shocks propagate through the financial network. Our first set of stress tests simulates the aftermath of a pool's default, identifying the pools that pose the highest level of systemic risk. In a second set of simulations, we characterize which liquidity pools default in response to a drop in the price of Bitcoin and Ether. We find that cascading failure is a distinct possibility, albeit requiring fairly sizeable price shocks. The pools of stablecoins are the most likely to default, whereas the pools of Bitcoins and Ethers are the most likely to set off a domino effect.

Related Literature. A growing body of research investigates DeFi protocols in order to assess their robustness and vulnerabilities. Formal analyses of lending pools can be found in Bartoletti et al. (2021, 2022) and Gudgeon et al. (2020a). Other studies simulate crash scenarios to explore how lending protocols respond to market price fluctuations (Kao et al., 2020; Qin et al., 2021). Additionally, researchers have analyzed the resilience of lending protocols to significant market events, such as the Ethereum Merge (Heimbach et al., 2023) and governance attacks (Gudgeon et al., 2020).

The examination of participants' behavior highlights severe liquidation risks due to their leverage (Heimbach and Huang, 2023) and risk appetite (Perez et al., 2021). Darlin et al. (2022) and Saengchote (2023) show that using debt-financed collateral fosters interconnectivity, whereas Chiu et al. (2022) explains why rigid haircut rules are likely to cause price-liquidity feedback loops. Empirical studies of liquidations reveal vulnerabilities leading to fire sales (Qin et al., 2021) and liquidation spirals (Warmuz et al., 2023) that could potentially endanger the stability of the DeFi ecosystem (Lehar and Parlour, 2022).

Instances of illiquidity have been documented, especially in newly established platforms (Gudgeon et al., 2020; Hafner et al., 2023, Sun et al., 2023). Our paper investigates whether liquidity pools are likely to become illiquid and explores how such an occurrence might propagate across Compound's network. To model these scenarios, we draw upon the extensive research that studies the mechanisms governing the transmission of shocks and distress across financial markets (Acemoglu et al., 2015; Elliott et al., 2014; Gai et al., 2011; Gai and Kapadia, 2010; Haldane and May, 2011). Given the extensive scope of the literature on financial contagion, we direct interested readers to two comprehensive surveys (Glasserman and Young, 2016; Jackson and Pernoud, 2021). Our approach, like these studies, uses the balance sheets of financial institutions to capture their connections and derive the corresponding network structure. In this context, Ao et al. (2023) and Saengchote and Castro-Iragorri (2023) are closely related to our research, as they use network analysis to assess the decentralization of DeFi and the interconnectedness of various protocols. By contrast, our paper pioneers the investigation of contagion risks and network effects *within* lending protocols.

Structure of paper. The design of Compound's protocol is presented in Section 3.2. Section 3.3 describes the structure of Compound's financial network. Section 3.4 explains how we use data about users' accounts to build the balance sheets of the lending pools. Section 3.5 analyzes the resilience of the protocol to cascades of failures. Section 3.6 concludes while the Appendices contain additional material about our algorithms and data.

3.2 Description of Compound

We start our analysis by describing how users interact with Compound since their actions determine the structure of the financial network.

3.2.1 Lending

When a user adds liquidity to a pool by depositing tokens, she receives an equivalent amount of cTokens in return. Essentially, cTokens are tokenized proofs of the deposit that can be redeemed at any time. Users are incentivized to provide liquidity because cTokens are deflationary and tend to increase in value relative

to the underlying asset, offering a rate of return on deposits.

For ease of notation, we will omit time inDEXs for all variables. We use $d_i^u \geq 0$ to denote the amount of tokens held by user u as deposit in pool i and gather all the deposits in the matrix d.[2] The row vector $d^u = (d_1^u, d_2^u, \ldots, d_k^u)$, where k is the number of pools managed by Compound, represents user u's deposits across the different pools. The column vector $d_i = (d_i^1, d_i^2, \ldots, d_i^n)$, where n is the number of active users, lists all the deposits in pool i. Consequently, the sum of deposits in pool i can be calculated as $\bar{d}_i = \sum_u d_i^u$.

To compare deposit values across pools, we need to convert them into a common unit of account. Let p_i represent the price of token i in US Dollars.[3] By stacking all prices into the vector $p = (p_1, \ldots, p_k)$, we can express the value of user u's deposits as $v(d^u, p) = d^u p^T = \sum_{i=1}^k d_i^u p_i$. Additionally, the vector of deposit values *across all* users is given by $v(d, p) = (v(d^1, p), \ldots, v(d^n, p)) = dp^T$.

3.2.2 Borrowing

Users have two ways of withdrawing tokens from liquidity pools. As mentioned before, they can redeem their cTokens, which are then burned by the protocol, effectively releasing their deposited tokens. Alternatively, they can borrow tokens by using a portion of their deposits as collateral. When borrows are backed by cTokens from different pools, the loans create a web of liabilities interconnecting the various liquidity pools.

Using a notation similar to that used for deposits, we denote the amount of asset i borrowed by user u as b_i^u and collect these borrow amounts in the matrix b. Before borrowing an asset, users must select which cTokens they wish to use as collateral from the various tokens they have supplied. When a user enters a market, all cTokens they hold in that specific asset class are considered collateral. Let e be the matrix of dimension $n \times k$, where the element e_i^u represents the "enterMarket" option chosen by user u for token i. It takes a value of 1 if the user intends to use this asset class as collateral and 0 otherwise. Consequently, the collateral matrix c is given by $c = e \odot d$ where $\odot$ denotes the Hadamard product.

Each cToken has its own collateral factor, indicating the proportion of the underlying asset value that can be borrowed. These collateral factors are determined and set by the governance of the protocol. In general, tokens with a small market capitalization tend to have a low collateral factor. Formally, let $\kappa = (\kappa_1, \ldots, \kappa_k)$ where $\kappa_i \in [0, 1)$ for all i, be a vector representing the collateral factors associated with each pool.[4] The maximal collateral value available to user u for borrowing is given by $v(c^u, \kappa p) = c^u(\kappa \odot p)^T = \sum_{i=1}^k c_i^u \kappa_i p_i$. This value represents the borrowing capacity of user u since she faces the credit constraint $v(b^u, p) \leq v(c^u, \kappa p)$.

[2] For simplicity, we present deposits in terms of the underlying token. In practice, when you deposit funds into Compound V2, the protocol internally converts the underlying tokens into an equivalent amount of cTokens based on the current exchange rate.

[3] Compound primarily relies on Chainlink's *Open Price Feed* as its price oracle. The protocol preforms sanity checks by comparing Chainlink's price feeds to the prices quoted by Uniswap V2.

[4] See: https://docs.compound.finance/v2/comptroller/#collateral-factor

The financial health of each user is quantified by the ratio $h^u(p) = v(c^u, \kappa p)/v(b^u, p)$. Whenever a user's health ratio falls below 1, which can occur due to various reasons, such as an increase in the price of the borrowed tokens, her collateral assets become eligible for liquidation.

3.2.3 Borrowing and lending rates

Borrowers pay interest on their borrowed tokens, while lenders receive interest for providing their assets. The interest rates for borrowing and lending are determined algorithmically based on the utilization rates of each pool, i.e., the proportion of the pool's tokens that have been borrowed (Leshner and Hayes, 2019). The protocol maintains a positive interest rate difference between borrowing and lending to reward stakeholders and create liquidity reserves. Furthermore, Compound introduced a liquidity mining program on June 16, 2020, incentivizing user engagement through the distribution of its governance token (COMP). The details of this distribution mechanism are subject to governance control and may vary over time.[5]

3.3 Compound's financial network

The liquidity pools within Compound's financial network are interconnected through loans, as users lock tokens in one pool to borrow from another. This mechanism is similar to repurchase agreements (repos) between banks. Liquidity pools replace financial institutions, with each loan representing a claim from the pool of the borrowed asset towards the pool(s) providing collateral.

Compound's financial network comprises a set $\mathcal{K} = \{1, ..., k\}$ of pools or nodes. The balance sheets of the pools are constructed as follows:

- On the liabilities side, we allocate the certificates of deposit (cTokens) across three categories. Firstly, we collect the market value of all deposits (D_i) that have not entered any market and are, therefore, external liabilities. Secondly, the interpool liabilities of pool i are the sum of all liabilities towards other pools ($\sum_{j=1}^{K} L_{ij}$). Specifically, when a user decides to lock their cTokens as collateral, an amount corresponding to the debt becomes a liability of the pool towards the pool from which the tokens have been borrowed. Thirdly, the remaining portion of cTokens that have entered the market but are in excess of the value of the debt is allocated to the pool's buffer (B_i). We separate these cTokens from deposits because they are not external liabilities. Instead, the buffer can be mobilized to secure the debt of users with a deteriorating health ratio.

[5] See: https://compound.finance/governance/comp

- On the assets side, we aggregate the market value of all the available tokens in each pool (T_i),[6] including the reserves set aside by the protocol (R_i) and the value of the interpool assets ($\sum_{j=1}^{K} L_{ji}$), i.e., the value of all tokens from other pools locked as collateral for loans originating from pool i.

We use the following procedure to identify undercollateralized claims. First, we define $\alpha^u = v(b^u, p)/v(c^u, p)$ as user u's borrowing-to-collateral ratio.[7] Assuming proportional allocation of collateral, if the user has non-negative net worth ($\alpha^u \leq 1$), then the effective liability of pool i towards pool j corresponds in value to the nominal debt, implying $l_{ij}^u = \alpha^u \beta_j^u c_i$ where $\beta_j^u \equiv p_j b_j^u / v(b^u, p)$ is the share of u's borrowing in pool j. However, if the user has a negative net worth ($\alpha^u > 1$), the pool can only recover $l_{ij}^u = \beta_j^u c_i^u$ units of asset i before depleting all the funds set aside in u's buffer. Hence, the liabilities and buffer of user u read

$$l_{ij}^u = \min\{1, \alpha^u\} \beta_j^u c_i^u \text{ and } B_{ij}^u = (1 - \min\{1, \alpha^u\}) \beta_j^u c_i^u. \tag{3.1}$$

The matrices $\bar{L}$ and L summarize the nominal and effective interpool liabilities. Nominal values represent the *promised payments* associated with each claim, while effective values take into account users' solvency by adjusting for all undercollateralized claims. Therefore, $\bar{L}_{ij} \in \mathbb{R}^+$ ($L_{ij} \in \mathbb{R}^+$) represents the nominal (effective) liabilities of pool i towards pool j. These values can be calculated by aggregating the nominal and effective liabilities of all users as follows: $\bar{L}_{ij} = \sum_{u=1}^{n} \alpha^u \beta_j^u c_i^u p_i$ and $L_{ij} = \sum_{u=1}^{n} l_{i,j}^u p_i$. Finally, the buffer of each pool in $\mathcal{K}$ is given by $B_i = \sum_j \sum_u B_{ij}^u p_i$.

With this at hand, we can express the net worth of pool i as

$$V_i = \sum_j L_{ji} + R_i + T_i - \sum_j L_{ij} - B_i - D_i. \tag{3.2}$$

Figure 3.1 contains a schematic financial network with only two liquidity pools. The arrows connecting the pools represent the direction in which payments flow.

3.4 Data and Descriptive Statistics

3.4.1 Data

The public availability of Ethereum's transaction history allows us to reconstruct the pools' balance sheets at any desired time point. By collecting the daily snapshots of users' positions from Compound V2's subgraph,[8]

[6] More precisely, we calculate the total market value by multiplying the number of tokens T_i available in pool i by their market price p_i.

[7] We define α as the borrowing-to-collateral ratio, rather than the collateral-to-borrowing ratio, in order to prevent infinite values for pure lenders.

[8] https://thegraph.com/hosted-service/subgraph/graphprotocol/compound-v2

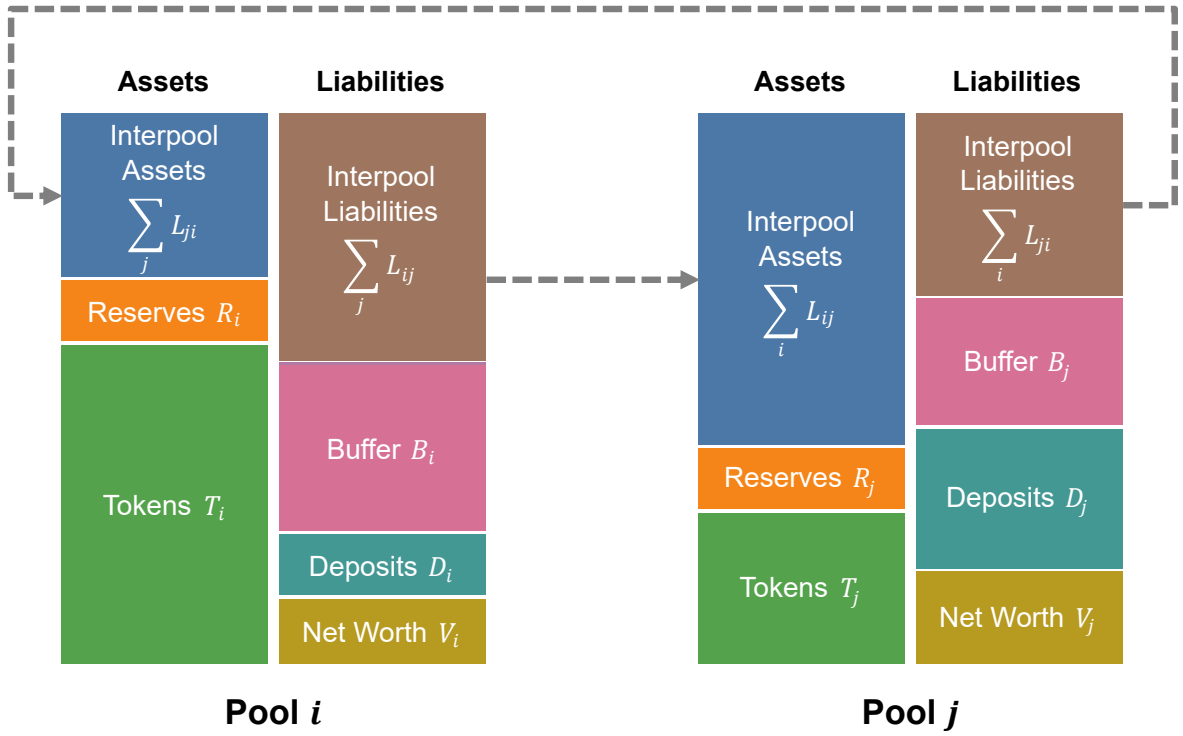

Figure 3.1: Pools' balance sheets and interpool linkages.

we obtain a comprehensive list of all users, denoted as $\mathcal{U}$, along with their respective lending (d_i^u) and borrowing (b_i^u) balances for each asset, including accrued interests from both lending and borrowing. In order to convert all balances to US Dollars, we rely on the market prices (p) of the tokens provided by Compound's oracle. This results in a sizable dataset containing the daily positions of 422,459 users across 19 pools, spanning from January 1, 2020, to June 30, 2023. Using this information, we construct the liability matrices and balance sheets for each daily snapshot, following the procedure outlined in section 3.3.

3.4.2 Financial Network

Our procedure generates daily snapshots of the balance sheet of each liquidity pool. Figure 3.2 presents a cross-section of the top 10 pools on September 7, 2021.[9] Two observations stand out. Firstly, the majority of all deposits are concentrated within five main pools, which can be categorized into two groups: (i) stablecoin-pools (cUSDC, cUSDT, cDAI), and (ii) crypto-pools (cETH, cWBTC2).[10] Secondly, there is a

[9] We selected this snapshot because it is the day on which Compound reached its highest Total Value Locked (TVL). We show below that the network structure exhibits persistent features, allowing us to extrapolate general insights from this specific day.

[10] WBTCs are wrapped bitcoins, i.e., ERC-20 tokens on the Ethereum blockchain pegged to Bitcoin. The cWBTC2 pool replaced cWBTC in April 2021, following an upgrade of the cToken contract implementation (See: https://compound.finance/governance/proposals/41).

consistent difference in balance sheet composition between the crypto-pools and stablecoin-pools. Specifically, stablecoin-pools hold significant interpool-assets, whereas crypto-pools have minimal interpool-assets. On the liabilities side, most cTokens in crypto-pools are used as collateral, while a substantial portion is held as deposits in stablecoin-pools. These observations suggest that users deposit cryptoassets in Compound primarily to borrow stablecoins.

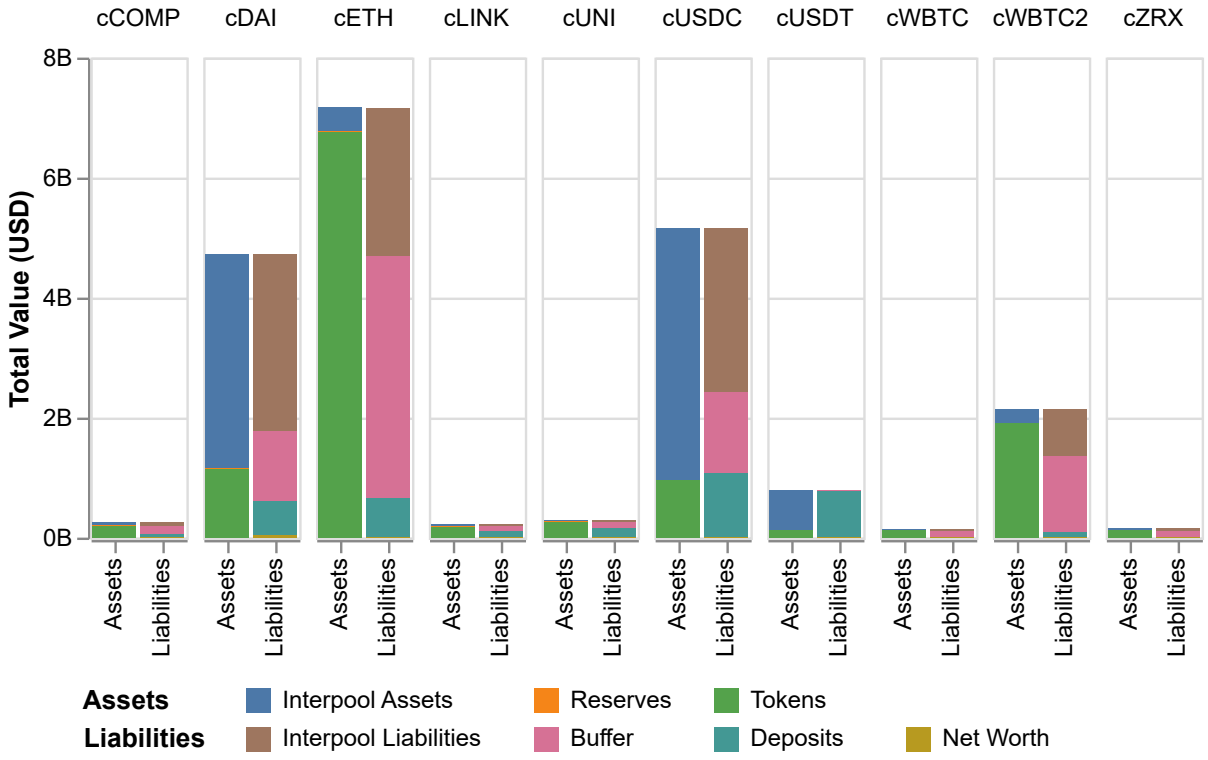

Figure 3.2: Balance sheets of the top 10 pools on Sept. 7, 2021.

This intuition is confirmed by Figure 3.3, which presents the financial network on our reference day (September 7, 2021). The size of the circles in the graph is proportional to the values in US Dollars of the pools' interpool assets. The arrows connecting the pools indicate the direction of payment flows, while the size of the arrows is proportional to the value of the claims. Upon analyzing Figure 3.3, it becomes evident that the majority of interpool links originate from the crypto-pools and connect to the stablecoin-pools. This observation supports the notion that most borrowers utilize Compound as a protocol to engage in repurchase agreements (repos) to trade their cryptoassets for stablecoins.

Figure 3.3 further illustrates an intriguing pattern: users often engage in borrowing from the same pool they have previously lent to. This self-borrowing behavior is captured by the color of the nodes, with darker shades indicating a higher proportion of cTokens being used for self-borrowing. Such a strategy can be financially advantageous, despite the gap separating the borrowing from the lending interest rates, because

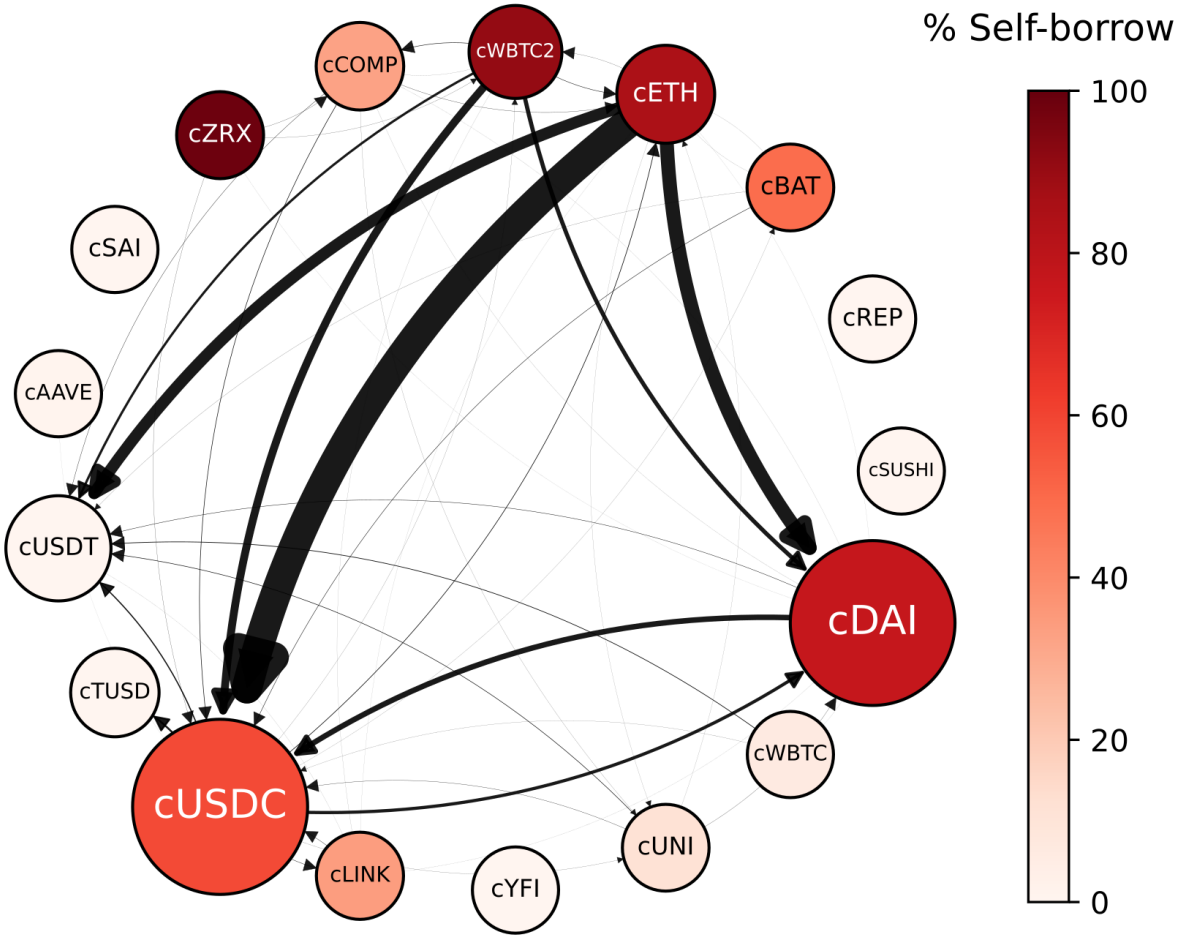

Figure 3.3: Interpool liability network on Sept. 7, 2021.
The percentage of self-borrow represents the proportion of interpool assets that consist of tokens from the same pool $\left(L_{ii}/\sum_{j=1}^{K} L_{ji}\right)$.

Compound encourages liquidity provision by distributing its governance token (COMP). Additionally, the benefits of liquidity mining motivate certain users to utilize one stablecoin to borrow another, with cUSDC and cDAI being particularly notable in this regard.

3.4.3 Centrality

The node-link diagram presented in Figure 3.3 may not readily apply to generalization across multiple snapshots. To achieve this, we build a scalar measure for each pool that effectively encapsulates their centrality within the financial network. Centrality measures can be computed from the liabilities matrix

$\bar{L}$.[11] Specifically, there are two measures of eigenvector centrality of node j which can be obtained in the following way (Glasserman and Young, 2016):

$$\lambda \nu_j^L = \sum_{i=1}^{k} \nu_j^L \bar{L}_{ij} \text{ and } \lambda \nu_j^R = \sum_{i=1}^{k} \bar{L}_{ij} \nu_j^R, \tag{3.3}$$

where λ represents the dominant eigenvalue of $\bar{L}$. The left eigenvector ν^L measures *funding centrality*, attributing more centrality to nodes that hold claims on nodes with higher centrality. Conversely, the right eigenvector ν^R measures *borrowing centrality*, assigning more centrality to nodes that hold obligations towards nodes with greater centrality.

Figure 3.4 reports the evolution of the two centrality measures for the dominant pools over time. Although we observe some variations in the relative importance of each pool, an underlying pattern consistent with the earlier snapshot becomes apparent: Borrowing centrality attributes most of the weights to crypto-pools (ETH and either WBTC1 or WBTC2) while funding centrality is concentrated within the stablecoin-pools (DAI, USDC, and USDT).

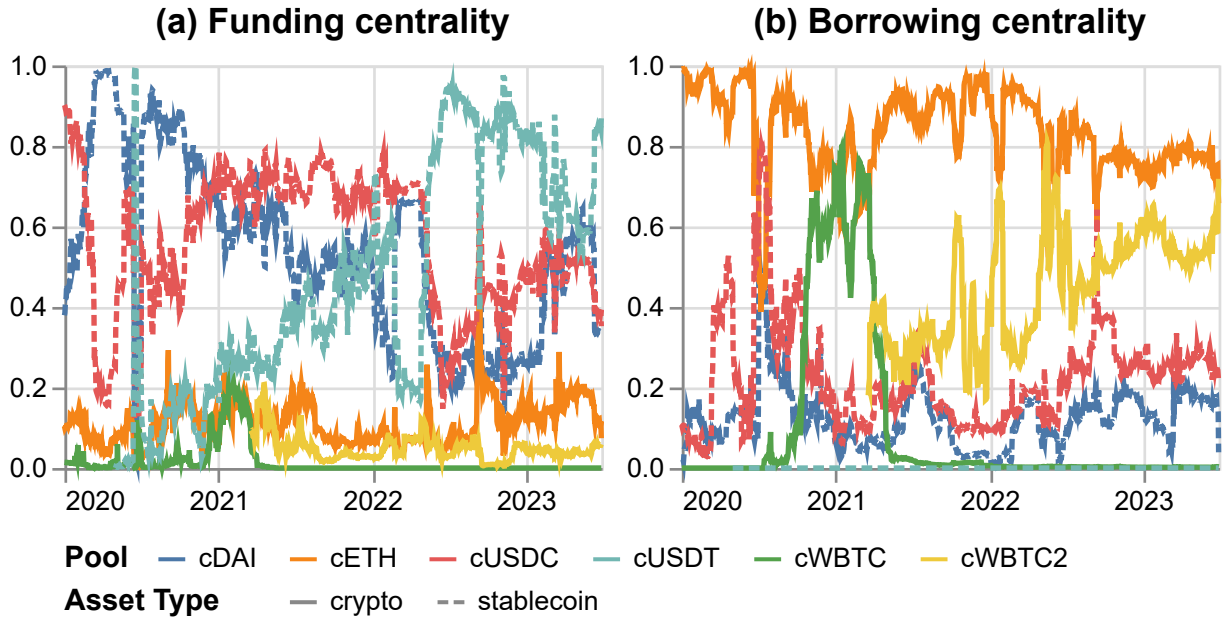

Figure 3.4: Centrality of the six main pools over time.

We identified only two exceptions to this rule, occurring in July 2020 and September 2022, where stablecoins gained a higher borrowing centrality than cryptoassets. In June 2020, following the launch of the liquidity mining program, cDAI and cUSDC experienced increased usage to borrow other stablecoins. Most

[11] Since we focus on interpool links, we exclude self-borrowing by setting the diagonal entries of the liability matrix to zero, i.e., $\bar{L}_{ii} = 0$ for all $i \in \mathcal{K}$

of these loans were repaid within a span of one to two weeks. In September 2022, a significant ETH price drop triggered a flight to stable assets, temporarily making cUSDC the main collateral pool. Towards the end of September, as ETH stabilized, cUSDC's borrowing centrality returned to its usual level.

To summarize, a descriptive analysis of Compound's network reveals that it serves two primary purposes. Firstly, it provides a decentralized protocol enabling users to issue repurchase agreements of cryptoassets against stablecoins. Consequently, the main risk to the stability of the protocol is a decline in the value of cryptoassets, particularly ETH and BTC, as it would undermine the value of the collateral supporting the majority of loans. Secondly, Compound allows its investors to engage in liquidity mining, either by borrowing from the same pool they have lent to or by utilizing one stablecoin to borrow another. The latter strategy carries little risk, except in cases where the stablecoin used as collateral undergoes a depegging episode. We now investigate whether these two sources of risks are likely to propagate across the network.

3.5 Contagion

Financial networks are prone to contagion episodes, wherein the default of a financial intermediary triggers a cascade of failures (Glasserman and Young, 2016. Jackson and Pernoud, 2021; Upper, 2011). We focus on the following contagion mechanism which proceeds in two steps. Initially, a wave of liquidations is set off whenever a pool defaults on its obligations by suspending the convertibility of its cTokens. If the shock is large enough, the liquidation process fails to restore the value of all interpool assets, thereby burdening the balance sheets of connected pools with bad loans. This may lead to other pools becoming insolvent, further amplifying the initial shock and causing a domino effect. We now describe how one can model this contagion process.

3.5.1 Liquidations

Compound's liquidation process safeguards lenders by maintaining an adequate level of collateralization. When a borrower's health ratio falls below 1, the liquidation of her position is automatically triggered. Liquidators repay a portion of the debt, known as the "close factor" (denoted by γ), and receive collateral at the current price plus a liquidation discount factor $\lambda \in (0, 1)$. Multiple rounds of liquidation may occur until the borrower's health ratio is restored above 1.

Let's consider a scenario where user u becomes eligible for liquidation after a change in price from p to p'. For tractability, we assume that liquidators follow a proportional rule, seizing all borrowed assets based on their share of the user's total debt.[12] Under this assumption, the user's borrowing balance is reduced

[12]In practice, liquidators have the flexibility to choose the token for repayment, the amount to repay within the factor limit,

by $\gamma v(b^u, p')$ and the liquidator acquires collateral worth $(1+\lambda)\gamma v(b^u, p')$ in return. As a result, all asset holdings of user u after liquidation are diminished by $\psi^u(p')c_j^u$, where $\psi^u(p') \equiv (1+\lambda)\gamma v(b^u, p')/v(c^u, p')$. The liquidation process is explained in more detail in section 3.7 where we describe the algorithm used for its simulation.

The health ratio of user u after t rounds of liquidation, which we denote by $h_t^u(p')$, obeys the following law-of-motion

$$h_{t+1}^u(p') = \frac{1-\psi^u(p')}{1-\gamma} h_t^u(p'). \tag{3.4}$$

This expression yields an intuitive threshold condition:

1. $v(c^u, p') > (1+\lambda)v(b^u, p')$: In this case, the collateral value exceeds the borrowed value multiplied by one plus the liquidation discount. Here, liquidation improves the health factor as the reduction in borrowing exceeds the reduction in collateral ($\gamma > \psi^u$). After potentially multiple rounds of liquidation, the health of the account will eventually be restored ($h^u(p') > 1$).

2. $v(c^u, p') < (1+\lambda)v(b^u, p')$: Here, the collateral value is lower than the borrowed value multiplied by one plus the liquidation discount. Then liquidation *worsens* the health factor, triggering multiple rounds of liquidation until all of user u's collateral is liquidated Warmuz et al. (2023). Consequently, there will be $\left(1 - \frac{v(c^u,p')}{v(b^u,p')(1+\lambda)}\right) b^u$ bad loans left in the various pools where u chose to invest.

3.5.2 Cascades

The liquidation process has the following impact on the pools' balance sheets. On the liabilities side, collateral assets are transferred from borrowers to liquidators, effectively converting the interpool liabilities of the pools where the borrowers held collateral into deposits. On the assets side, all pools in which users borrow are replenished in proportion to the amounts borrowed, thereby converting the interpool assets into tokens. However, when the shock is so severe that condition 2 mentioned above holds, a share of the interpool assets is not fully repaid, which induces a fall in the net worth of the lending pool.

The shortfall may result in a negative net worth, indicating that the pool owes more than it owns. We follow the literature on financial contagion in assuming that this signal triggers a run wherein depositors try to withdraw all the available funds Diamond and Dybvig (1983). If borrowers can mobilize external cash to repay their loans, they are able to withdraw all the liabilities of the pool. But, since the pool has a negative net worth, it cannot honor its commitment and must suspend the convertibility of its cTokens.[13]

and the collaterals to seize. We impose proportional liquidations because it enables us to derive analytical results. We assess the impact of this assumption in Section 9.

[13]The suspension of convertibility is explicitly handled on lines 518-521 of the smart contract CToken.sol for pool management.

To avoid overestimating the likelihood of cascades, we consider a more conservative scenario where users do not have access to external funds.[14] In this scenario, users can only redeem a portion of the pool's liabilities, specifically the cTokens held as deposit and excess buffer.[15] Then, the run on cTokens leads to default solely when the combined value of the deposits and excess buffers redeemed by users surpasses the value of the tokens and reserves held by the pool.

We emphasize that our simulations should be seen as offering a conservative estimate of the impact of liquidations. This is because we do not account for the potential further price declines of the liquidated asset. Empirical studies (Chiu et al., 2022; Yaish et al., 2023, Lehar and Parlour, 2022) indicate that this is an optimistic scenario, as markets do not always have sufficient liquidity to absorb selling pressure effectively.

3.5.3 Simulations

We combine the mechanisms discussed in the preceding subsections to design an algorithm, outlined in Section 3.7, that simulates the spread of a pool default throughout the entire network. Running this algorithm for each pool identifies those that pose the highest level of systemic risk. We will defer the examination of the factors that might initiate a particular pool's default until the conclusion of this subsection.

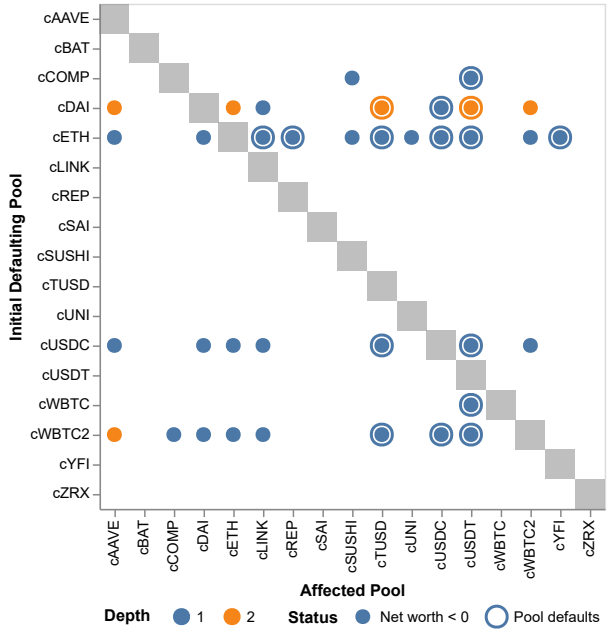

Figure 3.5: Default cascades on Sept. 7, 2021.
Initial defaulting pool indicated by the row, affected pools indicated by the columns. The pools with a negative net worth are encoded as a dot, while the defaulting pools are encoded as an outer circle. The color indicates the round in which a pool is affected.

[14] This assumption also has the advantage of being consistent with the premise that users do not prevent liquidations by recapitalizing their undercollateralized positions.

[15] The excess buffer refers to the cTokens that have entered the market but can be redeemed without triggering additional liquidations. In other words, the excess buffer equals the buffer minus the haircut associated with the collateral factor of the loans.

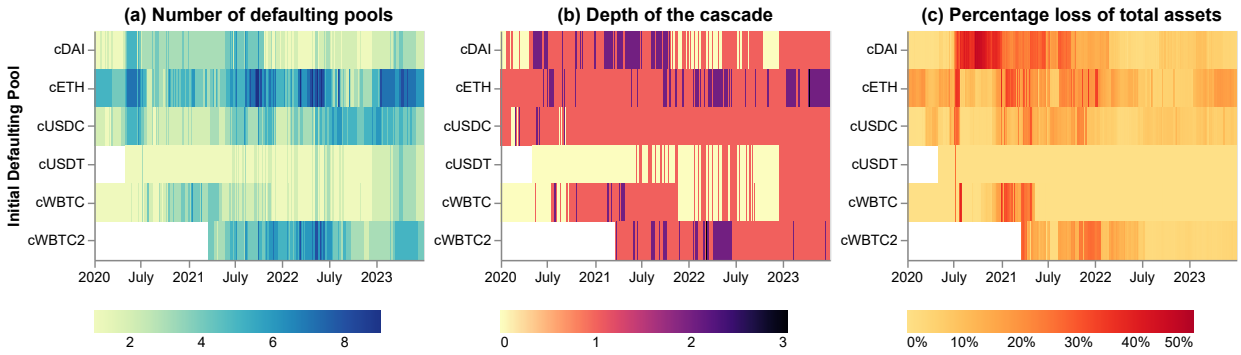

Figure 3.6: Daily snapshots of default cascades for the top six pools.

Focusing on our reference day (September 7, 2022), we present the predictions of the contagion algorithm in Figure 3.5. Each row tracks the contagion resulting from the default of a specific pool. As expected, contagion risks are concentrated within the main pools, except for cUSDT because its collateral factor is set to zero. Among these pools, the crypto-pools (cETH and cWBTC2) are the primary sources of contagion because they account for the bulk of collateral assets. Meanwhile, the stablecoin-pools exhibit a higher likelihood of default due to their elevated utilization rates.

Additionally, the stablecoin-pools also pose systemic risks to other stablecoin-pools, as users partake in liquidity mining by utilizing one stablecoin to borrow another. These strategies are generally considered riskless, unless one stablecoin depegs and, as captured by our simulations, triggers the default of its borrowing pools.

We repeated these simulations for each day in our sample and compiled the results in Figure 3.6. The left-hand panel summarizes the number of pools affected by the cascades. Consistent with the snapshot reported in Figure 3.5, cETH and cWBTC/cWBTC2 usually trigger the largest number of defaults. However, by the end of 2022, cUSDC had emerged as a significant threat because users responded to the significant decline in cryptocurrency prices by diversifying their sources of collateral.[16] The middle panel illustrates the depth of the cascades, revealing that most of them unfolded within one or two rounds before reaching a state where no further defaults occurred. By the third round, all cascades had been resolved. The right-hand panel depicts the percentage of total asset loss caused by the cascades. The chart highlights episodes of significant losses, mostly involving the two crypto-pools (cETH and cWBTC). It also indicates significant losses resulting from a default of the cDAI pool towards the end of 2020, when self-borrowing strategies were particularly widespread.[17] It is worth noting that the protocol's robustness improved over time, as indicated by the diminishing losses in recent years. This improvement is explained by the accumulation of reserves,

[16] See Section 9 for further details on the composition of loans over time.
[17] See Figure 3.10 in Section 9 for supporting evidence.

resulting in a higher net worth for the pools.

The aforementioned simulations follow the cascades that result from the default of a specific pool without providing an explanation for why the default occurred initially. In practice, the initial default can be triggered by a decline in the market price of the pool's collateral. To evaluate the likelihood of this scenario, we subjected the prices of the main cryptoassets (ETH and BTC) to negative shocks. We use δ to denote the magnitude of the price shock, so that $p'_{\{ETH,BTC\}} = (1-\delta)p_{\{ETH,BTC\}}$.

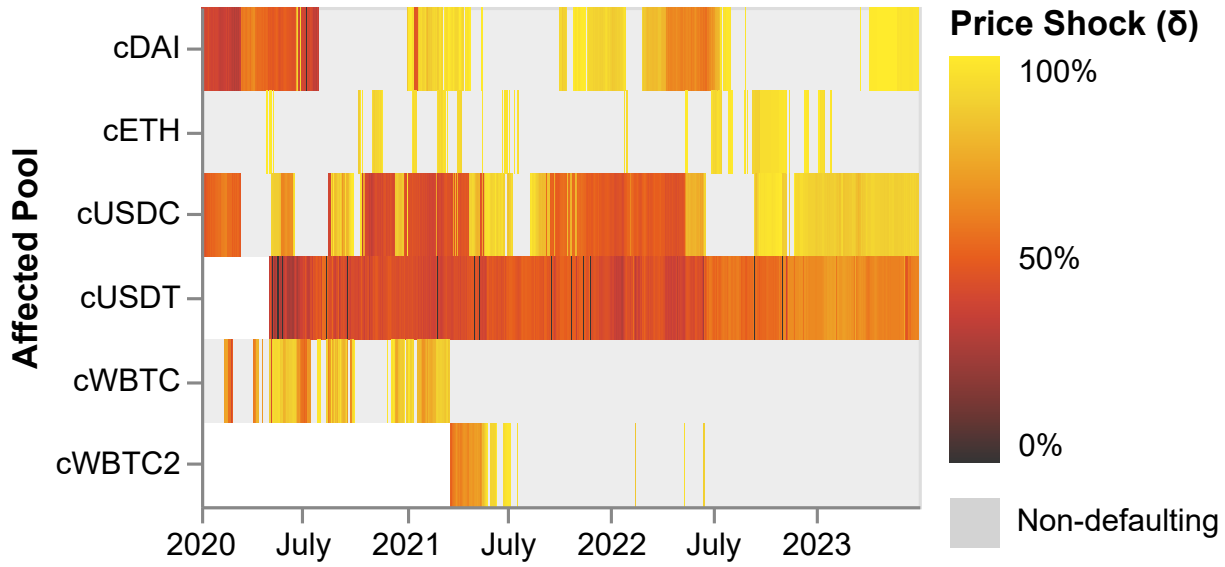

Figure 3.7: Minimal price shocks triggering default.

The outcomes of these experiments are depicted in Figure 3.7. It shows that default cascades are more likely to originate from stablecoin pools. This fragility can be primarily attributed to two factors: high utilization rates and the reliance on cryptoassets as collateral for the majority of stablecoin loans. Additionally, we observe that cUSDT is the most prone to default, which rationalizes the protocol's decision of curtailing the contagion that may originate from cUSDT by setting its collateral factor to zero. Overall, we find that to endanger the protocol, the shocks have to be fairly consequential, involving a decline of 50% or more in market prices.

3.6 Conclusion

Smart contracts reduce counterparty risks, but, as illustrated by the Terra debacle,[18] they do not eliminate contagion risks resulting from flaws in the economic design of their protocol. Fortunately, the transparency of blockchains offers an opportunity to develop sophisticated supervisory tools. The availability of exhaustive

[18] See https://www.coindesk.com/learn/the-fall-of-terra-a-timeline-of-the-meteoric-rise-and-crash-of-ust-and-luna/

and exact data makes it possible to monitor the sources of systemic risks in real time with a precision that far exceeds what is possible in the traditional financial sector. Our paper, by applying classical methods for the analysis of financial networks, provides an example of such an endeavor. It characterizes how contagion might spread through Compound's network, identifying the pools that are more likely to set off or propagate a domino effect. Moving forward, we aim to further explore this research avenue by delving deeper into the trove of data accumulated during this study. In particular, we intend to study the bipartite structure of the financial network and interact it with the empirical behavior of users in order to identify those that are more likely to endanger the stability of the protocol.

3.7 Appendix

Description of Algorithms

This appendix outlines the logic of the algorithms used to simulate default cascades. We first present in algorithm 1 how liquidations are simulated. All borrowers with an health factor h^u below one are liquidated up to the maximum amount determined by the close factor (γ). The liquidator repays the borrowed assets and seizes the collateral assets proportionally.[19] Under proportional liquidations, the liquidator acquires collateral worth $(1+\lambda)\zeta_j^u(p)\gamma v(b^u, p')$ for each asset j used as collateral by user u. Here, $\zeta_j^u(p) \equiv p_j c_j^u / v(c^u, p)$ represents the value of assets j relative to the user's total collateral value.

For each liquidation round, the algorithm updates the new borrow (b^u) and collateral (c^u) positions of the borrower.[20] Moreover, the deposits of the collateral pools (D_j) and tokens of the borrowing pools (D_i) are also increased as a result of liquidation. The algorithm liquidates the borrower's position until $h^u > 1$. We set the gas costs of liquidations equal to the median transaction fee on that snapshot date (i.e., $median(gasFees) \times gasUsed$, where $gasUsed = 500,000$).

Algorithm 1: Simulation of liquidation process

for $u \in \mathcal{U}$ **do**
  while $h^u < 1$ **do**
    `// 1. Calculate % borrow and collateral.`
    $\beta^u \leftarrow pb^u / v(b^u, p)$;
    $\zeta^u \leftarrow pc^u / v(c^u, p)$;
    `// 2. Calculate repaid and seized amounts.`
    $repay^u \leftarrow min\{v(b^u, p) \cdot \gamma, v(c^u, p)/(1+\lambda)\}$;
    $seize^u \leftarrow (1+\lambda) * repay^u$;
    `// 3. Stop liquidating if profits` $\leq txFees$.
    if $seize^u - repay^u \leq txFees$ **then**
      break;
    end
    `// 4. Repay the borrowed assets.`
    $pb^u \leftarrow pb^u - repay^u * \beta^u$;
    $T \leftarrow T + repay^u * \beta^u$;
    `// 5. Seize the collateral assets.`
    $pc^u \leftarrow pc^u - seize^u * \zeta^u$;
    $D \leftarrow D + seize^u * \zeta^u$;
    `// 6. Update the health of user` u
    $h^u \leftarrow v(c^u, \kappa p)/v(b^u, p)$;
  end
end

Algorithm 2 explains how we simulate the propagation of a pool's default across the network. Since the cTokens of the defaulting pool are not anymore liquid, we set their collateral factor κ_i to zero. This captures

[19] We assess the impact of this assumption in section 9.
[20] The updating process for balances is detailed in Section 9.

the fact that the cTokens cannot be used to repay the loans. The algorithm calls two procedures. The first one, called *simulate_liquidation*, follows the logic outlined in algorithm 1. It returns the updated collateral (c) and borrowing (b) matrices defined in section 3.2, along with the deposits (D) and token holdings (T) of each pool. The second function, called *update_balance_sheet*, follows the methodology presented in section 3.3 to compute the interpool liabilities (L), reserves (R), buffer (B), excess buffer (X), and net worth (V) of each pool. As explained in footnote 15, the excess buffer collects all the cTokens in the buffer that can be redeemed without triggering additional liquidations. The simulation iterates until the cascade reaches a state where no further pools default.

Algorithm 2: Simulation of cascading default of the pools

Input: *init_pool*: initial defaulting pool
Output: D_r: list of defaulting pools at depth r.

$D_0 \leftarrow \{init\_pool\}$; `// list of defaulting pools`
r $\leftarrow 0$; `// depth of liquidation cascade`
`// Continue liquidation if any new pool defaulted in the previous round`
while $r > 0$ *and* $|D_r \cap D_{r-1}| > 0$ **do**
  $\kappa_{\{d \in D_r\}} \leftarrow 0$; `// set` $\kappa = 0$ `for defaulting pools`
  r $\leftarrow$ r + 1 ; `// increment the cascade depth`
  `// Liquidate borrowers and update the balance sheet`
  (c, b, D, T) $\leftarrow$ simulate_liquidation(c, b, D, T, p, κ);
  (L, R, B, X, V) $\leftarrow$ update_balance_sheet(c, b);
  `// Add default pools after liquidation`
  $D_r \leftarrow D_{r-1} \cup \{k \in \mathcal{K} | V_k < 0 \text{ and } D_k + X_k > T_k + R_k\}$;
end

Updating balance sheets following Liquidations

After liquidation, Compound adjusts the balances for both the borrower and the liquidator, thus impacting the balance sheets of the pools associated with the liquidated asset i and seized collateral j. However, the effects of the process extend beyond those specific pools. The pools associated with assets that borrower i has used as collateral are also impacted.

In this Section, we provide a detailed explanation of the methodology employed to update balance sheets following each liquidation event.

Updating the liability side of collateral pools

The seized collateral assets are transferred from the borrower's wallet to the liquidator's wallet. This process effectively converts the interpool liabilities of the pools, where the borrowers had placed their collateral, into

deposits. Specifically, we have

$$\begin{aligned} D'_j &= D_j \frac{p'_j}{p_j} + \Delta c^u_j p'_j, \quad \forall j \in \mathcal{K} \\ &= D_j \frac{p'_j}{p_j} + (1+\lambda)\zeta^u_j(p')\gamma v(b^u, p') \end{aligned}$$

Prior to the liquidation process, the interpool liabilities for pool j on the behalf of user u were $\sum_i l^u_{ji} = min\{1, \alpha^u(p)\}c^u_j p_j$. After the liquidation, the pool j's new interpool liabilities is:

$$min\{1, \alpha^u(p')\}p'_j(c^u_j - \Delta c^u_j), \quad \forall j \in \mathcal{K})$$

Therefore, it has decreased by:

$$min\{1, \alpha^u(p')\}\Delta c^u_j p'_j + \left[min\{1, \alpha^u(p)\}p_j - min\{1, \alpha^u(p')\}p'_j\right] c^u_j.$$

Furthermore, the new pool j's buffer fund is:

$$min\{1, \alpha^u(p')\}p'_j(c^u_j - \Delta c^u_j)$$

Hence, after the liquidation, the pool j's buffer fund has decreased by:

$$\left[1 - min\{\alpha^u(p')\}\right]\Delta c^u_j p'_j - \left[min\{1, \alpha^u(p)\}p_j - min\{1, \alpha^u(p')\}p'_j\right] c^u_j.$$

Updating the asset side of borrowing pools

On the other hand, all pools in which user u has borrowed are replenished in proportion to the amount borrowed, leading to the conversion of interpool assets into tokens. More specifically,

$$\begin{aligned} T'_i &= T_i \frac{p'_i}{p_i} + \Delta b^u_i p'_i, \quad \forall i \in \mathcal{K} \\ &= T_i \frac{p'_i}{p_i} + \beta^u_i(p')\gamma v(b^u, p') \end{aligned}$$

Moreover, the new interpool assets for pool i ($i \in \mathcal{K}$) on the behalf of user u is given by:

$$\begin{aligned}\sum_j L_{ji} p'_j &= min\{1, \alpha^u(p')\}\beta_i^{u'} \left[\sum_j (c_j - \Delta c_j^u) p'_j\right] \\ &= min\{1, \alpha^u(p')\}\beta_i^{u'} \left[\sum_j c_j p'_j - (1+\lambda)\gamma v(b^u, p')\right]\end{aligned}$$

Therefore, the change in interpool assets of pool i ($i \in \mathcal{K}$) after the liquidation of account u is then given by:

$$\begin{aligned}\sum_j c_j \left[min\{1, \alpha^u(p')\}\beta_i^{u'} p'_j - min\{1, \alpha^u(p)\}\beta_i^u p_j\right] \\ - \sum_j min\{1, \alpha^u(p')\}\beta_i^{u'}(1+\lambda)\gamma v(b^u, p')\end{aligned}$$

If the user u's new health ratio is above 1, the liquidation process stops there. However, for some unhealthy accounts, multiple rounds of liquidation may be necessary to raise the health ratio above 1. Moreover, in cases where the liquidation fails to improve the user's health ratio and instead results in a borrowing balance larger than the collateral balance ($\alpha^u > 1$), Compound will only be able to recover $(1-\lambda)v(c', p')$ (and not $v(c', p')$) of user u's debt since the protocol allocates a portion λ of user u's collateral to the liquidator as a reward. The remaining debt $v(b, p') - (1-\lambda)v(c', p')$ will be left unpaid.

The liquidation process affects both the liability side of each of user u's collateral assets and the asset side of each of her borrowed assets.

Before the liquidation but after the price change, the budget share of u's borrowing in asset i ($i \in \mathcal{K}$) is $\beta_i^u(p') = \frac{p'_i b_i^u}{v(b^u, p')}$ and her budget share of u's collateral in asset i is $\zeta_i^u(p') = \frac{p'_i c_i^u}{v(c^u, p')}$.

The borrower u's new borrowing and deposit balances after the liquidation are:

$$b_i^{u'} = b_i^u - \Delta b_i^u, \forall i \in \mathcal{K}$$

$$d_j^{u'} = d_i^u - \Delta c_j^u, \forall j \in \mathcal{K}$$

where $\Delta b_i^u = \frac{\beta_i^u(p')\gamma v(b^u, p')}{p'_i}$ and $\Delta c_j^u = (1+\lambda)\frac{\zeta_j^u(p')\gamma v(b^u, p')}{p'_j}$.

The liquidator l's new deposit balance is:

$$d_j^{l'} = d_j^l + \Delta c_j^u, \forall j \in \mathcal{K}$$

As a consequence, the borrower u has a new debt-to-collateral ratio

$$\alpha^u(p') = \frac{v(b', p')}{v(c', p')} = \frac{\sum_i (b_i^u - \Delta b_i^u) p_i'}{\sum_j (c_j^u - \Delta c_j^u) p_j'},$$

Robustness Check on Simulations of liquidations

In this appendix, we explore an alternative scenario where the liquidator's approach to asset seizure differs from the one used in the main analysis. Instead of seizing assets in proportion to their share of the user's total debt, we employ a *sequential rule* in which the liquidator prioritizes the assets with the highest share of the user's total debt.

To assess the impact of this alternative liquidation strategy, we rely on the Jaccard similarity index to compare the lists of defaulting pools between our benchmark simulations and the simulations based on sequential liquidations. The defaulting pools predicted by the two simulations agree (i.e., Jaccard index = 1) 96.72% of the time. Figure 3.8 illustrates the evolution over time of the Jaccard index for the top six pools. It indicates that, in general, the outcomes of the two simulations are either identical or very similar.

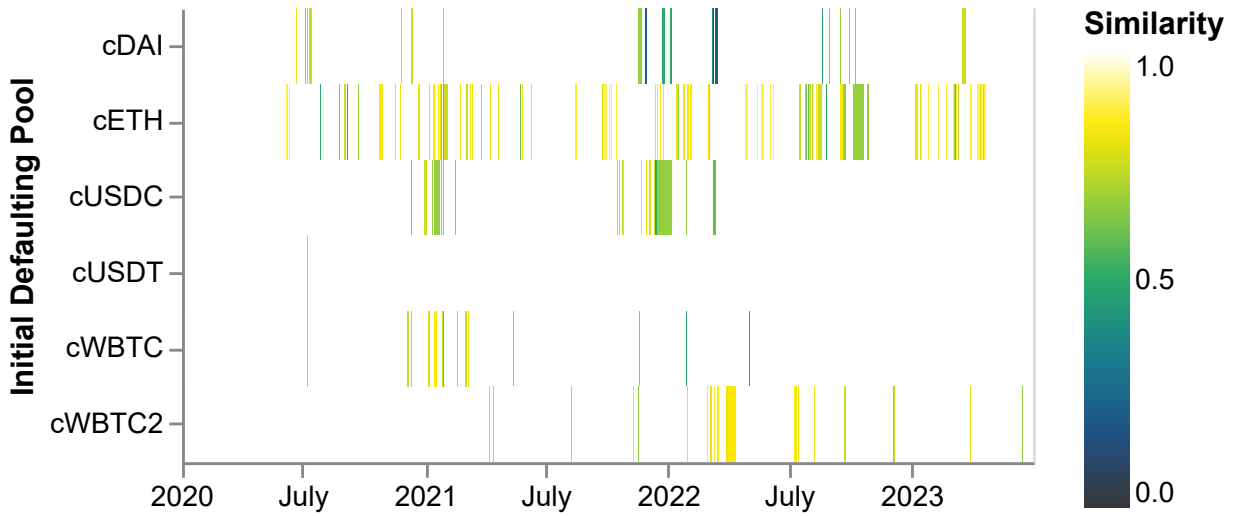

Figure 3.8: Jaccard similarity index of defaulting pools between benchmark and sequential simulations of liquidations.

Despite the stability of our results, we do observe a few short-lived periods of divergence. Particularly, on March 12th and 13th, 2022, the Jaccard index for the cDAI pool suddenly dropped to a very low value of 0.1429. To better understand this divergence, we report in Figure 3.9 a comparison of the default cascades predicted by the two simulation methods on March 13, 2022. It shows that a default of the cDAI pool did not trigger a domino effect in the benchmark simulation. However, in the sequential simulation, it triggered a default of the cUSDC pool, which then propagated to other pools heavily reliant on cUSDC, such as cLINK,

cTUSD, cUSDT, and cYFI.

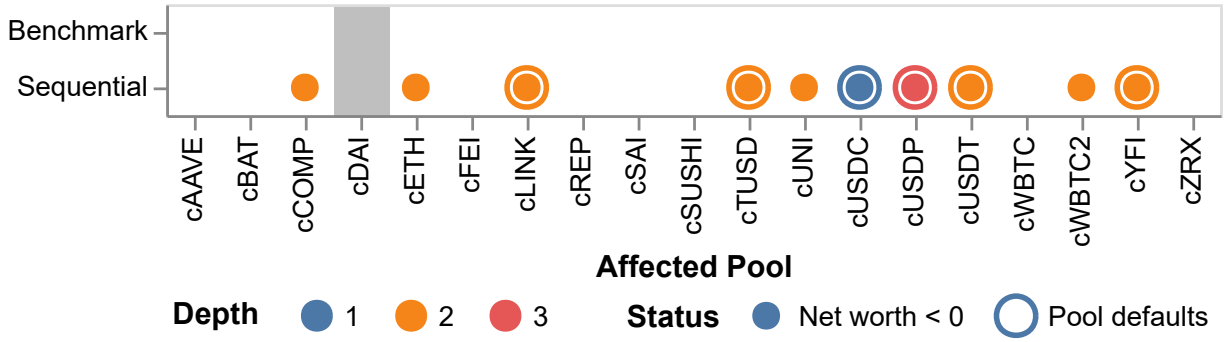

Figure 3.9: Comparison of cascades triggered by a default of the cDAI pool on March 13, 2022.

This difference in outcomes can be attributed to the fact that liquidators now prioritize USDCs. This concentration of liquidations on a specific asset can exacerbate the fragility of the network, particularly when the utilization rate and borrowing centrality of this asset are high. In summary, our findings are usually robust across the two specifications of the liquidation process. However, there are specific periods where the selection rule followed by the liquidators has an impact on the risk of contagion within the network. An interesting avenue for future research would therefore consist in studying the behavior of liquidators and devising a liquidation algorithm that reflects their actions.

Composition of Loans

This appendix provides additional information regarding the composition of loans. We categorize them into classes based on the nature of the asset pairs, distinguishing between cryptoassets and stablecoins. Additionally, we differentiate between self-borrowing and interpool loans. Figure 3.10 confirms that self-borrowing is primarily motivated by liquidity mining since pairs of self-borrowed stablecoins gained prominence shortly after the introduction of Compound's liquidity mining program on June 16, 2020.

Figure 3.10 further corroborates the findings presented in subsection 3.4.3, according to which a majority of interpool loans use cryptoassets as collateral for borrowing stablecoins. During the bear market from late 2021 to mid-2022, the weight of cryptoasset–stablecoin pairs significantly increased, which aligns with our simulation results in Figure 3.6 where we observe a heightened contagion risk emanating from cETH and cWBTC over this period.

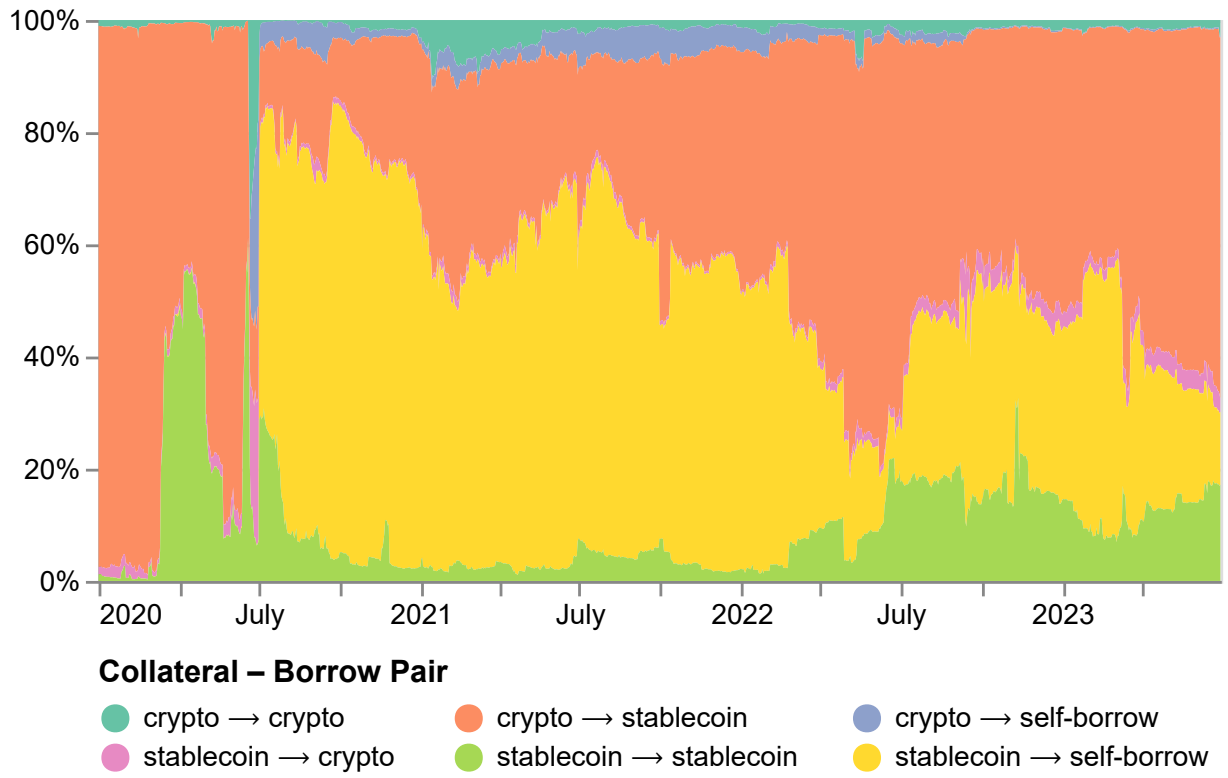

Figure 3.10: Distribution of nominal liabilities over time.

Subsequently, after the bear market phase, we notice a rise in the proportion of loans associated with stablecoin-stablecoin pairs, as well as an increase in self-borrowing of stablecoins. This trend amplified the borrowing centrality and, consequently, the risk of contagion associated with the cUSDC pool, as indicated in Figure 3.6.

Conclusion

This thesis has delved into the intricate ecosystem of DeFi, particularly focusing on the market microstructure of DEXs and the systemic risks inherent in lending protocols. Through a comprehensive analysis of existing literature, theoretical frameworks and empirical data, several critical insights have emerged.

Chapter 1 offered a systematic overview of CFMMs, emphasizing their role as fundamental components of decentralized exchanges. By tracing the evolution of AMMs from their theoretical origins to their practical implementations in DeFi, we highlighted the theoretical underpinnings and empirical observations surrounding CFMMs. Through rigorous examination, we uncovered their desirable properties, such as the oracle property and their ability to replicate payoffs from traditional financial instruments. Additionally, we explored the competition between centralized and decentralized exchanges, shedding light on the motivations behind liquidity provision and trade, ultimately contributing to a more comprehensive understanding of market microstructure dynamics.

In Chapter 2, we investigate the impact of inventory holding costs on the accuracy of Uniswap's pricing mechanism. By developing a microstructure model and leveraging price data from Uniswap and Binance, the study reveals how inventory holding costs impact traders' decision-making processes and subsequently, the price accuracy of Uniswap. This chapter underscores the importance of understanding market microstructure in decentralized exchanges and the role of inventory holding costs in shaping price dynamics. By incorporating additional market frictions and exploring the impact of factors such as private information and MEV, future studies can enhance our understanding of decentralized finance and contribute to the development of more robust and efficient decentralized exchange protocols.

In Chapter 3, we turned our attention to financial contagion within Compound V2, a decentralized lending protocol on the Ethereum blockchain. Leveraging network analysis and stress testing methodologies, we unveil the interconnectedness of Compound's liquidity pools and evaluate their susceptibility to contagion risks. By simulating various stress scenarios, we identify pivotal pools and potential triggers for cascading defaults, offering insights into the robustness and resilience of decentralized lending protocols. Our findings underscore the importance of transparency and real-time monitoring in mitigating systemic risks within the

DeFi ecosystem.

In conclusion, this thesis has provided a comprehensive examination of DeFi, spanning from the theoretical foundations of AMMs to the empirical analysis of market microstructure and systemic risk within lending protocols. By elucidating the intricacies of CFMMs, exploring the dynamics of price accuracy in Uniswap, and uncovering the potential for financial contagion within Compound V2, we have advanced our understanding of the decentralized financial landscape. Moving forward, addressing the challenges identified in this thesis will be essential for fostering the resilience and sustainability of DeFi ecosystems, paving the way for further innovation and growth in DeFi.

Bibliography

Acemoglu, Daron et al. (2015). "Systemic risk and stability in financial networks". In: *American Economic Review* 105.2, pp. 564–608. DOI: 10.1257/aer.20130456.

Adams, H. et al. (2020). "Uniswap v2 Core." In.

Adhavan, A. and Smidt S. (1991). "A Bayesian model of intraday specialist pricing". In: 30.1, pp. 99–134.

— (1993). "An Analysis of Changes in Specialist Inventories and Quotations". In: 48.5, 1595–1628 (34 pages).

Amihud, Y. and H. Mendelson (1980). "Dealership market: Market-making with inventory." In: 8.1, pp. 31–53.

Ang, Andrew (2014). *Asset management: A systematic approach to factor investing*. Oxford University Press.

Angeris, Guillermo and Tarun Chitra (2020). "Improved price oracles: Constant function market makers". In: *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pp. 80–91.

Angeris, Guillermo et al. (2022a). "Constant function market makers: Multi-asset trades via convex optimization". In: *Handbook on Blockchain*. Springer, pp. 415–444.

Angeris, Guillermo et al. (2022b). "When does the tail wag the dog? Curvature and market making". In.

Angeris, Guillermo et al. (2021e). "An analysis of Uniswap markets". In.

Angeris, Guillermo et al. (2021d). "Differential privacy in constant function market makers". In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 149–178.

Angeris, Guillermo et al. (2021b). "Replicating monotonic payoffs without oracles". In: *arXiv preprint arXiv:2111.13740*.

Ao, Ziqiao et al. (2023). *Is decentralized finance actually decentralized? A social network analysis of the Aave protocol on the Ethereum blockchain*. arXiv: 2206.08401 [econ.GN].

Aoyagi, J. (2020). "Liquidity Provision by Automated market makers." In.

Aoyagi, Jun and Yuki Ito (2021). "Coexisting exchange platforms: Limit order books and automated market makers". In.

Aymanns, Christoph et al. (2018). "Models of financial stability and their application in stress tests". In: *Handbook of Computational Economics*. Vol. 4. Elsevier, pp. 329–391. DOI: 10.1016/bs.hescom.2018.04.001.

Bardoscia, Marco et al. (2015). "DebtRank: A microscopic foundation for shock propagation". In: *PloS one* 10.6, e0130406. DOI: 10.1371/journal.pone.0130406.

Bartoletti, Massimo et al. (2021). "SoK: Lending Pools in Decentralized Finance". In: *Financial Cryptography and Data Security. FC 2021 International Workshops*. Springer, pp. 553–578. ISBN: 978-3-662-63958-0. DOI: 10.1007/978-3-662-63958-0_40.

Bartoletti, Massimo et al. (2022). "Formal analysis of lending pools in decentralized finance". In: *International Symposium on Leveraging Applications of Formal Methods*. Springer, pp. 335–355. DOI: 10.1007/978-3-031-19759-8_21.

Battiston, Stefano et al. (2012). "Debtrank: Too central to fail? Financial networks, the fed and systemic risk". In: *Scientific Reports* 2.1, p. 541. DOI: 10.1038/srep00541.

Bentov, Iddo et al. (2017). *The cost of decentralization in 0x and EtherDelta*.

Brier, Glenn W (1950). "Verification of forecasts expressed in terms of probability". In: *Monthly weather review* 78.1, pp. 1–3.

Buterin, Vitalik (2016). "Let's run on-chain decentralized exchanges the way we run prediction markets". In: *URL: https://www. reddit. com/r/ethereum/comments/55m04x/lets_run_onchain_ decentralized_exchanges_the_way*.

Capponi, Agostino and Ruizhe Jia (2021). "The adoption of blockchain-based decentralized exchanges". In: *arXiv preprint arXiv:2103.08842*.

Castro-Iragorri, Carlos et al. (2021). *Financial intermediation and risk in decentralized lending protocols*. arXiv: 2107.14678 [q-fin.GN].

Chen, Yiling and David M Pennock (2012). "A utility framework for bounded-loss market makers". In: *arXiv preprint arXiv:1206.5252*.

Chitra, Tarun et al. (2021c). "How Liveness Separates CFMMs and Order Books". In.

Chiu, Jonathan et al. (Nov. 2022). "On the Fragility of DeFi Lending". In: *Available at SSRN 4328481*. URL: https://ssrn.com/abstract=4328481.

Cifuentes, Rodrigo et al. (2005). "Liquidity risk and contagion". In: *Journal of the European Economic Association* 3.2-3, pp. 556–566. DOI: 10.1162/jeea.2005.3.2-3.556.

Daian, Philip et al. (2019). "Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges". In: *arXiv preprint arXiv:1904.05234*.

Darlin, Michael et al. (2022). "Debt-financed collateral and stability risks in the DeFi ecosystem". In: *2022 4th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS)*. IEEE, pp. 5–12. DOI: 10.1109/BRAINS55737.2022.9909090.

DeFi Pulse (n.d.). URL: https://www.defipulse.com/.

Diamond, Douglas W and Philip H Dybvig (1983). "Bank runs, deposit insurance, and liquidity". In: *Journal of Political Economy* 91.3, pp. 401–419. DOI: 10.1086/261155.

Easley, D. et al. (1998). "Option Volume and Stock Prices: Evidence on Where informed traders Trade." In: Vol. 53.2, 431–465 (35 pages).

Easley, David and Maureen O'hara (1987). "Price, trade size, and information in securities markets". In: *Journal of Financial economics* 19.1, pp. 69–90.

Eisenberg, Larry and Thomas H Noe (2001). "Systemic risk in financial systems". In: *Management Science* 47.2, pp. 236–249. DOI: 10.1287/mnsc.47.2.236.9835.

Elliott, Matthew et al. (2014). "Financial networks and contagion". In: *American Economic Review* 104.10, pp. 3115–3153. DOI: 10.1257/aer.104.10.3115.

Elsinger, Helmut et al. (2006). "Risk assessment for banking systems". In: *Management science* 52.9, pp. 1301–1314. DOI: 10.1287/mnsc.1060.0531.

Fabi, M. and J. Prat (2022). "The Economics of Constant Function Market Makers". In: DOI: https://blockchain-polytechnique.com/the-economics-of-constant-function-market-makers/.

Fabi, Michele et al. (2023). *A Companion to Decentralized Finance, Digital Assets, and Blockchain Technologies*. Eward Elgar Publishing Ltd. Chap. SoK: Constant Function Market Makers.

Gai, Prasanna and Sujit Kapadia (2010). "Contagion in financial networks". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 466.2120, pp. 2401–2423. DOI: 10.1098/rspa.2009.0410.

Gai, Prasanna et al. (2011). "Complexity, concentration and contagion". In: *Journal of Monetary Economics* 58.5, pp. 453–470. DOI: 10.1016/j.jmoneco.2011.05.005.

Garman, M. B. (1976). "Market microstructure". In: 3.3, pp. 257–275.

Genest, Christian and James V Zidek (1986). "Combining probability distributions: A critique and an annotated bibliography". In: *Statistical Science*, pp. 114–135.

Glasserman, Paul and H Peyton Young (2016). "Contagion in financial networks". In: *Journal of Economic Literature* 54.3, pp. 779–831. DOI: 10.1257/jel.20151228.

Glosten, Lawrence R and Paul R Milgrom (1985a). "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders". In: *Journal of financial economics* 14.1, pp. 71–100.

Glosten, Lawrence R. and Paul R. Milgrom (1985b). "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders". In: Vol. 14, issue 1 (1), pp. 71–100.

Gonzalo, J. and C. Granger (1995). "Estimation of common long-memory components in cointegrated systems". In: 13, pp. 27–35.

Good, Irving John (1952). "Rational decisions". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 14.1, pp. 107–114.

Gudgeon, Lewis et al. (2020a). "DeFi protocols for loanable funds: Interest rates, liquidity and market efficiency". In: *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pp. 92–112. DOI: 10.1145/3419614.3423254.

Gudgeon, Lewis et al. (2020b). "The decentralized financial crisis". In: *2020 Crypto Valley Conference on Blockchain Technology (CVCBT)*. IEEE, pp. 1–15. DOI: 10.1109/CVCBT50464.2020.00005.

Hafner, Matthias et al. (2023). "DeFi Lending Platform Liquidity Risk: The Example of Folks Finance". In: *The Journal of The British Blockchain Association*. DOI: 10.31585/jbba-6-1-(5)2023.

Haldane, Andrew G and Robert M May (2011). "Systemic risk in banking ecosystems". In: *Nature* 469.7330, pp. 351–355. DOI: 10.1038/nature09659.

Han, J. et al. (2021). "Trust in DeFi: an empirical study of the decentralized exchange". In.

Han, Jianlei et al. (2022). "Trust in defi: an empirical study of the decentralized exchange". In: *Available at SSRN 3896461*.

Hanson, Robin (2003). "Combinatorial information market design". In: *Information Systems Frontiers* 5, pp. 107–119.

— (2007). "Logarithmic markets coring rules for modular combinatorial information aggregation". In: *The Journal of Prediction Markets* 1.1, pp. 3–15.

Hausch Lo, Ziemba (1994). "Efficiency of Racetrack Betting Markets". In.

Heimbach, Lioba and Wenqian Huang (May 2023). "DeFi Leverage". In: *Available at SSRN 4459384*. URL: https://ssrn.com/abstract=4459384.

Heimbach, Lioba and Roger Wattenhofer (2022). "Eliminating sandwich attacks with the help of game theory". In: *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, pp. 153–167.

Heimbach, Lioba et al. (2021). "Behavior of liquidity providers in decentralized exchanges". In: *arXiv preprint arXiv:2105.13822*.

Heimbach, Lioba et al. (2023). *DeFi Lending During The Merge*. arXiv: 2303.08748 [q-fin.GN].

Hermans, L. et al. (2022). "Decrypting financial stability risks in crypto-asset markets." In.

Ho, T. and H. Stoll (1981). "Optimal dealer pricing under transactions and return uncertainty". In: *Journal of Financial economics* 9.1, pp. 47–73.

— (1983). "The dynamics of dealer markets under competition". In: *The Journal of finance* 38.4, pp. 1053–1074.

Jackson, Matthew O and Agathe Pernoud (2021). "Systemic risk in financial networks: A survey". In: *Annual Review of Economics* 13, pp. 171–202. DOI: 10.1146/annurev-economics-083120-111540.

— (Mar. 2019). "Optimal Regulation and Investment Incentives in Financial Networks". In: *Available at SSRN 3311839*. URL: https://ssrn.com/abstract=3311839.

Jose, Victor Richmond R et al. (2008). "Scoring rules, generalized entropy, and utility maximization". In: *Operations research* 56.5, pp. 1146–1157.

Kao, Hsien-Tang et al. (2020). "An analysis of the market risk to participants in the compound protocol". In: *Third International Symposium on Foundations and Applications of Blockchains*. URL: https://scfab.github.io/2020/FAB2020_p5.pdf.

Krishnamachari, Bhaskar et al. (2021). "Dynamic curves for decentralized autonomous cryptocurrency exchanges". In: *arXiv preprint arXiv:2101.02778*.

Kyle, Albert S (1985). "Continuous auctions and insider trading". In: *Econometrica: Journal of the Econometric Society*, pp. 1315–1335.

Lehar, Alfred and Christine A Parlour (2021). "Decentralized exchanges". In: *Available at SSRN* 3905316.

— (Dec. 2022). *Systemic fragility in decentralized markets*. BIS Working Papers 1062. Bank for International Settlements. URL: https://www.bis.org/publ/work1062.pdf.

Leshner, Robert and Geoffrey Hayes (2019). *Compound: The money market protocol*. Tech. rep. URL: https://compound.finance/documents/Compound.Whitepaper.pdf.

Lewis, Michael (2014). *Flash boys: a Wall Street revolt*. WW Norton & Company.

Li, J. et al. (2022). "Reading the Candlesticks: An OK Estimator for Volatility". In: pp. 1–45.

Liao, Gordon and Dan Robinson (2022). *The dominance of uniswap v3 liquidity*.

Liu, C (2008). *The Dark Forest*. Tor Books.

Lo, Yuen C and Francesca Medda (2021). "Uniswap and the Emergence of the Decentralized Exchange". In: *Journal of financial market infrastructures* 10.2, pp. 1–25.

Mas-Colell, Andreu et al. (1995). *Microeconomic theory*. Vol. 1. Oxford university press New York.

Michel, C. (2022). *Pricing LP Tokens—Warp Finance Hack*. Accessed: 2022-07-30. URL: https://cmichel.io/pricing-lp-tokens.

Milionis, Jason et al. (2022). "Automated market making and loss-versus-rebalancing". In: *arXiv preprint arXiv:2208.06046*.

Nakamoto, Satoshi and A Bitcoin (2008). "A peer-to-peer electronic cash system". In: *Bitcoin.–URL: https://bitcoin. org/bitcoin. pdf* 4.2, p. 15.

Park, Andreas (2021). "The conceptual flaws of constant product automated market making". In: *Available at SSRN* 3805750.

Perez, D. et al. (2021a). "Liquidations: DeFi on a Knife-edge." In: vol. 457-476. Springer. International Conference on Financial Cryptography and Data Security.

Perez, Daniel et al. (2021b). "Liquidations: DeFi on a Knife-Edge". In: *Financial Cryptography and Data Security*. Ed. by Nikita Borisov and Claudia Diaz. Springer, pp. 457–476. ISBN: 978-3-662-64331-0. DOI: 10.1007/978-3-662-64331-0_24.

Qin, Kaihua et al. (2021). "An empirical study of DeFi liquidations: Incentives, risks, and instabilities". In: *Proceedings of the 21st ACM Internet Measurement Conference*, pp. 336–350. DOI: 10.1145/3487552.3487811.

Qin, Kaihua et al. (2022). "Quantifying blockchain extractable value: How dark is the forest?" In: *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, pp. 198–214.

Queries - Uniswap Docs. (N.d.). URL: https://docs.uniswap.org/protocol/V2/reference/API/queries..

Robinson, Dan and Georgios Konstantopoulos (2020). "Ethereum is a dark forest". In: *Medium, Aug.*

Rochet, Jean-Charles and Jean Tirole (2003). "Platform competition in two-sided markets". In: *Journal of the european economic association* 1.4, pp. 990–1029.

Saengchote, Kanis (2023). "Decentralized lending and its users: Insights from Compound". In: *Journal of International Financial Markets, Institutions and Money* 87, p. 101807. DOI: 10.1016/j.intfin.2023.101807.

Saengchote, Kanis and Carlos Castro-Iragorri (June 2023). "Network Topology in Decentralized Finance". In: *Available at SSRN 4469783*. URL: https://ssrn.com/abstract=4469783.

Saleh, Fahad et al. (2022). "The need for fees at a dex: How increases in fees can increase dex trading volume". In: *Available at SSRN*.

Smidt, S. (1971). "Which road to an efficient stock market: free competition or regulated monopoly?" In: 27.5, pp. 18–20.

Stoll, H. (1978). "The supply of dealer services in securities markets". In: *The Journal of Finance* 33.4, pp. 1133–1151.

Sun, Xiaotong et al. (2023). *Liquidity Risks in Lending Protocols: Evidence from Aave Protocol*. arXiv: 2206.11973 [q-fin.RM].

"The Graph" (n.d.). In: (). URL: https://thegraph.com/en/.

Upper, Christian (2011). "Simulation methods to assess the danger of contagion in interbank markets". In: *Journal of Financial Stability* 7.3, pp. 111–125. DOI: 10.1016/j.jfs.2010.12.001.

Warmuz, Jakub et al. (2023). *Toxic Liquidation Spirals*. arXiv: 2212.07306 [econ.GN].

Winkler, Robert L (1969). "Scoring rules and the evaluation of probability assessors". In: *Journal of the American Statistical Association* 64.327, pp. 1073–1078.

Yaish, Aviv et al. (2023). *Suboptimality in DeFi*. Cryptology ePrint Archive, Paper 2023/892. URL: https://eprint.iacr.org/2023/892.

Yuen, L. and M. Francesca (2020). "Uniswap and the Emergence of the Decentralized Exchange." In.

Zhou, L. et al. (2021). "An Empirical Study of DeFi Liquidations: Incentives, Risks, and Instabilities". In: vol. ACM, New York, NY, USA, 15 pages. ACM Internet Measurement Conference (IMC '21), November 2–4, 2021, Virtual Event, USA.

ECOLE DOCTORALE

Titre : Essais sur les dynamiques de marché et la stabilité en Finance Décentralisée

Mots clés : Blockchain, Finance Décentralisée, Cryptomonnaie, Microstructure de Marché, Risques Systémiques, Économie Digitale

Résumé : Alors que le monde de la finance continue d'évoluer rapidement, l'émergence des marchés décentralisés et de la technologie Blockchain présente à la fois des opportunités et des défis. Opérant de manière autonome et autorégulées, les marchés décentralisés introduisent un changement de paradigme qui exige une compréhension approfondie. Cette thèse se concentre sur la compréhension de ces nouveaux marchés et sur l'évaluation de leur stabilité, avec pour objectif plus large de contribuer à la stabilité de l'ensemble de l'écosystème financier. Le premier chapitre propose une analyse approfondie des "Automated Market Makers" (AMM) pour les échanges décentralisés (DEXs). Il retrace l'évolution historique des AMM depuis leurs origines dans les marchés de prédiction jusqu'à leur application actuelle dans la finance décentralisée (DeFi). Le chapitre met en lumière le passage de cahiers d'ordres (Limit Order Books) complexes aux "Constant Function Market Makers" (CFMM), qui utilisent des fonctions de trading déterministes pour fixer le prix des actifs. Ce chapitre examine les recherches théoriques et empiriques sur les CFMM, illustrant leur capacité à reproduire les instruments financiers traditionnels et à optimiser les surplus des traders et des fournisseurs de liquidité. De plus, ce chapitre explore la dynamique concurrentielle entre les échanges centralisés et décentralisés et aborde la question critique de la valeur maximale extractible (MEV), proposant des solutions potentielles. Le deuxième chapitre se concentre sur la microstructure des DEXs, avec un accent particulier sur Uniswap. Uniswap facilite les échanges de cryptomonnaies pair-à-pair en utilisant un CFMM et des pools de liquidité plutôt que des livres d'ordres. Ce chapitre examine comment les coûts d'inventaire supportés par les traders affectent la capacité d'Uniswap à synchroniser ses prix avec ceux des marchés centralisés, considérés comme marchés de référence. Un modèle de microstructure est développé et calibré à l'aide des données de prix de clôture par minute des marchés Uniswap et Binance, de mai 2020 à décembre 2022. L'analyse révèle que les coûts d'inventaire réduisent la réactivité des traders aux écarts de prix. Le troisième chapitre présente un examen détaillé de la contagion financière au sein de Compound V2, un protocole de prêt décentralisé. En construisant les bilans des pools de liquidité de Compound, le chapitre caractérise son réseau financier et identifie les activités prédominantes des utilisateurs, comme l'emprunt de stablecoins et la participation à l'extraction de liquidité. Des tests de stress simulent des scénarios de défauts de pool et de chocs de prix significatifs sur Bitcoin et Ether, révélant que bien que des défaillances en cascade soient possibles, elles nécessitent des chocs de prix importants. Les pools de stablecoins sont plus susceptibles de faire défaut, tandis que les pools de Bitcoin et Ether sont plus susceptibles de déclencher des réactions en chaîne. Dans l'ensemble, cette thèse vise à améliorer la compréhension de l'écosystème DeFi en explorant sa microstructure de marché et les risques de contagion. Les conclusions tirées de cette recherche contribuent au débat sur la finance décentralisée, promouvant un écosystème financier plus résilient et durable.

Title : Essays on Market Dynamics and Stability in Decentralized Finance

Keywords : Blockchain, Decentralized Finance, Cryptocurrency, Market Microstructure, Systemic Risks, Digital Economy

Abstract : As the world of finance continues to evolve rapidly, the emergence of decentralized markets and Blockchain technology presents both opportunities and challenges. Operating autonomously and self-regulated, decentralized markets introduce a paradigm shift that demands a thorough comprehension. This thesis focuses on understanding these new markets and assessing their stability, with a broader aim of contributing to the stability of the entire financial ecosystem. The first chapter provides an in-depth analysis of Automated Market Makers (AMMs) for Decentralized Exchanges (DEXs). It traces the historical evolution of AMMs from their origins in prediction markets to their current application in Decentralized Finance (DeFi). The chapter highlights the shift from complex Limit Order Books (LOBs) to Constant Function Market Makers (CFMMs), which use deterministic trading functions to set asset prices. It reviews theoretical and empirical studies on CFMMs, demonstrating their ability to replicate traditional financial instruments and optimize trader and liquidity provider surpluses. Additionally, this chapter explores the competitive dynamics between centralized and decentralized exchanges and addresses the critical issue of Miner Extractable Value (MEV), proposing potential solutions. The second chapter focuses on the microstructure of DEXs, with a particular emphasis on Uniswap. Uniswap facilitates peer-to-peer cryptocurrency exchanges using a CFMM and liquidity pools instead of order books. This chapter examines how inventory holding costs incurred by traders affect Uniswap's ability to synchronize its quoted prices with those of centralized markets, considered as reference markets. A microstructure model is developed and calibrated using 1-minute closing price data from Uniswap and Binance markets from May 2020 to December 2022. The analysis reveals that inventory holding costs reduce trader responsiveness to price deviations. The third chapter presents a detailed examination of financial contagion within Compound V2, a decentralized lending protocol. By constructing the balance sheets of Compound's liquidity pools, the chapter characterizes its financial network and identifies predominant user activities, such as borrowing stablecoins and liquidity mining. Stress tests simulate scenarios of pool defaults and significant price shocks in Bitcoin and Ether, revealing that while cascading failures are possible, they require substantial price shocks. Stablecoin pools are more susceptible to default, whereas Bitcoin and Ether pools are more likely to initiate chain reactions. Overall, this thesis aims to enhance the understanding of the DeFi ecosystem by exploring its market microstructure, and contagion risks. The insights gained from this research contribute to the broader discourse on decentralized finance, promoting a more resilient and sustainable financial ecosystem.

Institut Polytechnique de Paris
91120 Palaiseau, France