JUST TRADING OPPORTUNISTS? HOW CRYPTO TOKEN AIRDROPS SHAPE OPEN PLATFORM CONTRIBUTIONS

This draft: April 2026

ABSTRACT

Governance token airdrops have emerged as a prominent mechanism to transition from proprietary to collective platform governance. To incentivize external contributions, platform sponsors distribute ("airdrop") governance tokens, some worth billions of dollars, across platform participants and contributors. However, such transitions raise concerns not just to reduce platform sponsor opportunism but to enable contributor opportunism instead. Hence, we study token airdrops in blockchain-based open-source platforms to examine how this form of transition to collective governance shapes platform contributions. Using a staggered difference-in-differences design that exploits airdrops as governance interventions, we analyze a panel of 615,979 quarterly contributor–platform observations that contain more than 10 million open-source platform contributions between 2015 and 2025. We find that governance token airdrops causally increase external contributions by 13 percentage points. Contrary to concerns, we find no evidence that airdrops increase opportunistic contributor behavior. We further examine how specific design features of governance token airdrops may have mitigated opportunism and discuss the dual incentive role of governance power and token value as a mitigation mechanism in platform governance transitions. Our findings contribute to property rights theory and dynamic agency theory and offer practical implications for platform strategies aimed at fostering sustained, value-enhancing external contributions.

Keywords: Governance Token Airdrops; DAOs; Open Platforms; Property Rights Theory; Dynamic Agency Theory

INTRODUCTION

With governance token "airdrops", a new form of transitioning from proprietary to collective platform governance has emerged (O'Mahony & Karp, 2022; Yue & Nagle, 2025). Through the airdrop, an open-source platform sponsor distributes governance tokens to external contributors free of charge. Thereby, sponsors reallocate parts of their decision rights to formally include contributors into platform governance (Allen, Berg, & Lane, 2023; Halaburda, Livshits, & Yaish, 2025; Makridis, Fröwis, Sridhar, & Böhme, 2023). In blockchain-based platforms, airdrops have become one of the most visible and consequential governance innovations. They often involve the transfer of substantial economic value and formal authority to external actors within the initiated transition phase (Allen, 2024; Uniswap Labs, 2020). Increasingly, platform sponsors beyond web3 and airdrops deliberately forfeit governance control as contemporary transition cases show, e.g., Meta's PyTorch (AI), Google's Kubernetes (Cloud Computing), and IBM's Eclipse (Integrated Development Environments) (CNCF, 2016; O'Mahony & Karp, 2022; Yue & Nagle, 2025).

By transitioning to collective governance, these sponsors intend to address a core strategic challenge: sustaining value-enhancing external contributions over time. In settings where value creation depends on voluntary, asset-specific investments of external contributors (e.g., open source and platforms), concentrated control by a dominant sponsor can undermine contributors' incentives by exposing them to the risk of ex-post appropriation (O'Mahony & Karp, 2022; Yue & Nagle, 2025; O'Mahony & West, 2005; Cutolo & Kenney, 2021; Rietveld, Ploog, & Nieborg, 2020; Jacobides, Cennamo, & Gawer, 2024). By reallocating decision rights to contributors, sponsors credibly commit to collective governance and acknowledge the central role of external contributors in shaping the platform's future (O'Mahony & Karp, 2022). In this sense, governance transitions like airdrops represent a deliberate governance choice rather than a purely financial incentive. Consistent with this logic, nascent research shows that reallocating control rights from dominant platform sponsors to external contributors can increase their contributions in firm-sponsored open-source platforms (O'Mahony & Karp, 2022; Yue & Nagle, 2025).

However, while transitions to collective governance can strengthen contributors' incentives to create value, it simultaneously expands their ability to pursue private benefits through governance influence. Thereby, new forms of opportunism emerge under collective control (O'Mahony & Karp, 2022). Are governance transitions just trading opportunists?

The purpose of this paper is to examine how transitions from proprietary to collective governance in firm-sponsored open-source platforms through governance token airdrops influence external contributions. Drawing on property rights theory, we theorize that governance transitions through token airdrops reflect a change in the ownership structure of residual control rights over a platform asset. The distribution of residual control rights strengthens external contributor's ex-post value capture prospects and hence increases their ex-ante contributions incentives (Grossman & Hart, 1986; Hart & Moore, 1990). Further, by drawing on dynamic agency literature, we generalize the governance transition phase as a setting of contracting under moral hazard. We theorize that external contributors will act more opportunistically once the platform is collectively governed. The new governance system grants access to influential positions which they can exploit for additional private benefits without sanctions (Holmstrom, 1979; Holmstrom & Milgrom, 1987, 1991).

We build a novel data set with a detailed coverage of platform contribution activities in web3 between 2015 and 2025. Our sample contains more than 300,000 unique external contributors and we observe more than 10 million open-source platform contributions overall. To analyze effects from transitions to collective governance on external platform contributions, we implement a staggered difference-in-differences design that exploits airdrops as governance interventions

(Callaway & Sant'Anna, 2021). In our panel of 615,979 quarterly contributor–platform observations, the share of external contributions among all platform contributions shows a causal increase by 13-percentage points following the airdrop. The effects are statistically significant and we detect no signs of violations of the parallel trends assumption (Abadie, 2005). Contrary to concerns about opportunism, we do not find evidence of increased contributor opportunism through the governance transition.

A core strength of our study is the depth of its robustness architecture. We subject our main findings to a comprehensive battery of 19 robustness and sensitivity analyses spanning four domains: design robustness of the control group definition, specification and sample robustness of modeling choices and panel composition, sensitivity to unobserved confounding, and falsification through placebo and permutation tests.

While much strategy research has focused on how platform sponsors can expand their weight over the community of contributors (Gawer & Cusumano, 2002; Iansiti & Levien, 2004; Jacobides, Knudsen, & Augier, 2006; Jacobides & Tae, 2015), we examine the implications of a decreasing influence of the platform sponsor on external contributors. In this regard, one emerging stream has focused on forced decreases in influence through antitrust measures (Thatchenkery & Katila, 2023), we contribute to an emerging stream that analyses a deliberate forfeiture of platform power by the sponsor and its implications for contributor participation strategies (Boudreau, 2010; O'Mahony & Karp, 2022; Yue & Nagle, 2025). We add insight on how external contributors respond to change in ownership structures of residual control rights over a platform. Further, we add early insights into the implications of governance tokens as a dual-incentive mechanism that inseparably bounds platform periphery value capture incentives through control with platform core value capture incentives through token value. Our findings add to the boundaries of both property rights theory as well as dynamic agency theory and offer practical implications for platform strategies aimed at fostering sustained, value-enhancing external contributions.

THEORY AND HYPOTHESES

Our main theoretical research goal is to identify causal effects of how a transition from proprietary to collective governance influences external contributions to open platforms. Governance token airdrops as governance transition events serve as our research setting. Two research streams are relevant for our research question. First, property rights theory explains why externals invest in response to different ownership structures of residual control rights over a platform. Second, research on dynamic agency issues covers problems of contracting under moral hazard that helps to explain opportunistic behavior by externals under collective governance.

Airdrops, governance tokens and DAOs

"*The introduction of UNI (ERC-20) serves this purpose, enabling shared community ownership and a vibrant, diverse, and dedicated governance system, which will actively guide the protocol towards the future.*" (Uniswap Labs, 2020)

"Airdrops" are a mechanism to distribute platform governance decision-making rights to different types of platform participants and contributors. Three features make airdrops particularly special. First, granted governance rights are tamper-proof tokenized on public blockchains. Second, tokens are tradeable and attributed a value. Lastly, tokens are distributed for free (Allen, 2024; Allen et al., 2023; Halaburda et al., 2025; Hsieh & Vergne, 2023; Makridis et al., 2023).

In the absence of regulation and industry standards, platform providers determine token recipients and allocations unilaterally. Usually, members of the platform sponsor and investors receive an allocation of platform tokens. In addition, tokens are typically distributed to demand-side users and open-source (OS) developers that contribute to the core platform (i.e., platform providers) or through complementary innovation (i.e., complementors) (Abel & Kindermann, 2023; Allen, 2024; Allen et al., 2023; Halaburda et al., 2025). Finally, a share of tokens is usually allocated to the Decentralized Autonomous Organization (DAO) that henceforth governs the platform (Abel & Kindermann, 2023; Allen et al., 2023; Eisenmann, Parker, & van Alstyne, 2009; Halaburda et al., 2025; Lommers, Makridis, & Verboven, 2023). At its core, a DAO is a neutral foundation that hosts the platform's governance bodies and is controlled by the collective of governance token holders. Because of the inherent governance power, these tokens are also called governance tokens (Beck, Müller-Bloch, & King, 2018; Hassan & De Filippi, 2021; Hsieh, Vergne, Anderson, Lakhani, & Reitzig, 2018; Makridis et al., 2023; Murray, Kuban, Josefy, & Anderson, 2021). While insiders, i.e., the platform sponsor and investors, receive governance tokens, too, significant shares are usually distributed to external participants and contributors (Abel & Kindermann, 2023; Allen et al., 2023; Halaburda et al., 2025; Lommers et al., 2023). For example, Uniswap, one of the largest transaction platforms in decentralized finance, allocated approx. 60% of all their tokens to external participants and contributors. This allocation was worth 3.4 billion dollars at the time of the airdrop, vested over five years (DeFiLlama, 2026a; Uniswap Labs, 2020).

A governance token airdrop initiates the platform governance transition. The ownership of governance tokens allows to participate in governance decision processes, including policy discussions and decision-making. DAOs usually make decisions in proposal elections, where one governance token grants one vote to the holder. The decision-making power of governance tokens has a significantly larger scope than rights of traditional equity shares. Proposal elections in DAOs can deal with decisions of high strategic and operational relevance potentially daily. For example, DAO governance token holders may be asked to cast their vote to grant or withdraw access of a particular complement to the focal platform; decide on whether the platform should charge fees, from whom fees are to be charged, how much and who should in turn receive these profits. Hence, the platform control power of governance token holders is far reaching (Beck et al., 2018; Hassan & De Filippi, 2021; Hsieh et al., 2018; Makridis et al., 2023; Murray et al., 2021).

From a platform provider perspective, airdrops mainly serve the purpose to conduct to "decentralize the community", i.e., to grant relevant contributors a voice in shaping the future platform as an additional incentive for continued contributions (Uniswap Labs, 2020). Airdrops further serve marketing and funding purposes. Due to the attributed value, airdrops typically generate high attention among users and enable the platform sponsor to access capital by selling tokens that were not distributed as part of the airdrop (Abel & Kindermann, 2023; Allen et al., 2023; Halaburda et al., 2025; Lommers et al., 2023; Makridis et al., 2023).

Airdrops differ from other token distribution mechanisms like ICOs in which tokens are not distributed based on contributions but sold to investors (Bakos & Halaburda, 2022; Chod & Lyandres, 2021; Chod, Trichakis, & Yang, 2022; Malinova & Park, 2023). Hence, airdrops are a retrospective reward for merit, i.e., previous usage of or contributions to a platform technology. Consequently, airdrops transition platform governance from a traditional closed to a collective form by including relevant external actors based on the relevance of their contributions.

Owner opportunism and external investments in open-source platforms

Airdrops are a phenomenon recently observed to "decentralize communities" of open-source platforms to incentivize external contributions. OS platforms combine platform business models with a firm-sponsored OS development model. Platform owners forfeit parts of their proprietary IP to allow external actors to access their platform technologies (e.g., common components, tools, and interfaces) and to contribute code-based extensions or complementary innovations. By opening access, platform owners aim to increase the platform value which they capture later based on the control rights they usually retain (Boudreau, 2010; Eisenmann et al., 2009; O'Mahony, 2003;

O'Mahony & West, 2005). Opening access can also incentivize external contributors, since the platform offers a non-proprietary channel to monetize their developments or services (Dahlander & Magnusson, 2005; Von Krogh, Haefliger, Spaeth, & Wallin, 2012).

Extant research shows, enabling externals to co-develop the platform enhances scale and innovativeness (Cennamo, 2018; Parker & Van Alstyne, 2005; West, 2003). Opening platform access broadens adoption, stimulates diverse experimentation and creativity of the crowd, reduces innovation costs for all participants, and improves the relevance of offerings by drawing on the localized knowledge of external contributors (Baldwin & Von Hippel, 2011; Franke & Hippel, 2003; Hertel, Niedner, & Herrmann, 2003; Lakhani & Von Hippel, 2003; Lerner & Tirole, 2005; West, 2003).

However, research from both streams on platforms and firm-sponsored OS projects show that, despite opened access, control being concentrated in the firm sponsor (i.e., platform owner or platform sponsor) deters external contributions. In classic platforms, owners usually hold full IP and governance rights. This grants platform owners "bouncer's rights" (Bresnahan & Greenstein, 2014), which have led to multiple forms of opportunistic behavior by the platform owner vis-à-vis external contributors once they have contributed to value creation, e.g., bundling, profit squeezing, selective promotion, blocking, and vertical integration (Boudreau, 2010; Cutolo & Kenney, 2021; Rietveld et al., 2020; Thatchenkery & Katila, 2023). This ex-post opportunism by the owner reduces externals' value capture prospects ex ante and correspondingly their contributions (Cutolo & Kenney, 2021; Jacobides et al., 2024; O'Mahony & Karp, 2022; Rietveld et al., 2020). Firm-sponsored OS projects are founded as a profit-maximizing firm releases a technology that was created under proprietary conditions (Dahlander & Magnusson, 2005; O'Mahony & West, 2005). Though, multiple studies report about conflicts between the sponsoring firm and the community of external contributors about ownership and control. Where sponsoring firms, despite forfeited IP, fully control or dominate governance an exit of contributors is promoted (e.g., retain proprietary IP for innovations built on top of the public good or unilaterally set the strategic direction of the project and distribute value) (Dahlander & Magnusson, 2005; O'Mahony & West, 2005; Shah, 2006; Von Krogh et al., 2012). Extant research provides evidence for such conflicts even for collective forms of OS project governance. The researchers show, only informally implemented collective governance promotes opaque power balances, benevolent dictators & "in the moment"-authority, which eventually disincentivize external contributions (O'Mahony, 2007; O'Mahony & Ferraro, 2007).

Given formalized governance decision-making processes and rights, collective forms of OS platform governance may, however, reduce externals' perceived threat of platform owner opportunism. Projects with collective governance forms are more likely to attract contributions than those with a dominating firm (Dahlander & Magnusson, 2005; Franck & Jungwirth, 2003; Markus, 2007; O'Mahony, 2007; O'Mahony & West, 2005; Shah, 2006; Von Krogh et al., 2012). Open and transitions to open platform governance structures have a positive influence on the contribution levels of external participants (Chen, Richter, & Patel, 2021; O'Mahony & Karp, 2022; Yue & Nagle, 2025).

Property Rights Theory (PRT) is particularly suitable to explain ex-post platform sponsor opportunism and related ex-ante underinvestment by external contributors in the context of OS platforms. PRT states that the allocation of ownership and residual control rights determines investment incentives and shapes socially efficient collaboration among actors that are interdependent. A core assumption is a setting in which contracts are inevitably incomplete. That is, it is impossible or too costly to specify all future contingencies. This gives rise for *residual* control rights, which are ex-ante granted rights to make decisions in the future about non-contracted matters. These can include among others the usage of shared assets and the distribution of surplus.

In addition to incomplete contracts, contributions are asset-specific, irreversible investments whose value depends on continued access to shared assets and ex-post value capture – both depending on residual control rights. PRT predicts that an efficient outcome (i.e., optimal contribution level) requires an ownership structure that aligns the decision right distribution with actors' relative importance for joint output (Grossman & Hart, 1986; Hart & Moore, 1990).

Under such conditions, the distribution of residual control rights defines each actor's bargaining position over future value capture. Ownership of scarce and valuable assets strengthens this position, whereas lacking control exposes actors to surplus diversion or hold-ups, such as blocked access or discontinued interoperability. Thus, ownership influences contributors' expected return on investment ex post, which in turn determines ex-ante productive effort (Grossman & Hart, 1986; Hart & Moore, 1990).

The central problem PRT addresses is opportunistic behavior that arises when control rights are distributed disproportionately to actors' relative importance for joint production. A controlling actor with residual rights may exploit this discretion by diverting surplus ex post, by holding up other parties through discontinuing access to the asset or through withholding interoperability support. Opportunistic actions discourage specific investments because contributors anticipate the possibility of future ex-post reallocation. Such actions have a stronger effect, if dependent contributors lack sufficient outside options (Grossman & Hart, 1986; Hart & Moore, 1990).

In OS platforms, governance is the function to continuously decide on ex-ante non-contractible matters, e.g., adjusting the strategic direction and setting or adapting rules to distribute value. Hence, residual control rights reflect decision-making rights for governance processes in OS platforms. Typical governance processes in OS platforms are fully controlled or highly dominated by the platform sponsor (i.e., proprietary) who holds all or most governance decision-making rights (O'Mahony & West, 2005; West, 2003). Airdrops transition OS platform governance from a proprietary to a collective mode. Specifically, not just an allegedly neutral, third-party foundation is empowered (Yue & Nagle, 2025), but actual contributors. That is, a transition reallocates formal and tamper-proof ownership of residual control rights over a platform from a dominating platform sponsor to important contributors to the platform. This increases externals' bargaining position for future value capture and hence reduces their threat of ex-post platform sponsor opportunism. Increased ex-post value capture prospects lead to higher ex-ante efforts for asset-specific and irreversible platform investments by externals, i.e., contributions to specific, open-sourced code. In line with PRT, we therefore predict.

Hypothesis 1: Transitions from proprietary to collective governance in firm-sponsored OS platforms through governance token airdrops relate positively to platform contributions from external participants.

Airdrops, dynamic agency and contributor opportunism

According to PRT, changes in the ownership structure of residual control rights do not only incentivize investments but increase welfare subsequently (Hart & Moore, 1990). Interestingly, O'Mahony & Karp (2022) discover that governance transitions – reallocating residual control rights to externals to reduce platform owner opportunism threats – can introduce opportunistic behavior by these externals *through* contributions. The researchers find that some externals, besides others that act cooperatively, deliberately engage in distinct "contribution" strategies that increased their competitive advantage without improving the platform itself. This is outside of PRT's explanatory scope. PRT predicts opportunistic behavior and underinvestment, when control rights are disproportionately distributed compared to the relative importance of contributors (Grossman & Hart, 1986; Hart & Moore, 1990). However, the investigated governance transition introduces underinvestment in quality precisely because of this transition to collective governance. The researchers show that transitions enable external "contributors" to take over influential positions which they can exploit to either slow down or omit feature implementations, contribute inferior code, or even redirect the platform entirely to meet their own interests (O'Mahony & Karp, 2022). Therefore, transitions to collective platform governance can introduce new issues of moral hazard that go beyond or even contradict predictions of PRT: while changes in control structures may incentivize more investments, these investments may not always be complementary to the platform, i.e., increase overall welfare.

This discovery also differs from other scholarly findings where contributors harm a platform. For example, platform research already provides evidence that too high quantities of contributions can crowd out externals (Boudreau, 2012; Boudreau & Jeppesen, 2015; Li, Shen, & Bart, 2018) or dilute the overall value proposition where free riders underinvest and join the platform with inferior offerings (Cennamo, 2018; Cennamo & Santaló, 2019; Geva, Barzilay, & Oestreicher-Singer, 2019; Wang, Li, & Singh, 2018). While these contributions have negative implications for the focal platform or other participants, corresponding contributors do not intend to increase their own competitive differentiation by harming the platform. The governance-based opportunistic behavior that O'Mahony & Karp (2022) discover, in turn, deliberately limits the platform's competitive advantage for the private benefit of the focal "contributor". This private benefit of residual control, not considered in PRT, cannot be internalized by other participants and is only enabled by collective governance.

Beyond this study, our understanding of how transitions to collective platform governance may promote opportunistic behavior by external participants and contributors is very limited. Some research focusing on airdrops show instances of opportunistic behavior by externals. Allen et al. (2023) and Halaburda et al. (2025) provide empirical evidence for a phenomenon called "airdrop farming". In anticipation of the airdrop, fake users engage with a platform to qualify for a token reward. If successful, farmers monetize the reward on token markets once airdropped ("token dumping") and stop using the platform largely. Interestingly, Halaburda et al. (2025) find that, depending on the airdrop allocation design, airdrop farming can still be mutually beneficial.

We draw on contracting under moral hazard literature to attempt an explanation for these forms of opportunism. The contracting under moral hazard literature examines how incentives shape agents' effort when actions are valuable but imperfectly verifiable and uncontactable (Georgiadis, 2022). Agency theory predicts that moral hazard in the form of sequencing choices (strategic delay of upgrades), hidden effort margins (omission of welfare-critical improvements or full discontinuation of participation), and privately rewarded but principal-misaligned actions (inferior or misdirected code contributions) arises when performance cannot be credibly conditioned on quality or continuation (Holmstrom, 1979; Holmstrom & Milgrom, 1987, 1991). In the context of collective governance transitions, governance rights are granted retrospectively and become irreversible, so contributors can capture immediate private benefits while downstream quality and timing externalities of their actions are realized later by the collective and cannot be sanctioned ex post (Allen, 2024; Halaburda et al., 2025; Makridis et al., 2023). This prevails in dynamic settings without "commitment", e.g. autonomous contributions to open-source projects (O'Mahony & Karp, 2022; Yue & Nagle, 2025). In other words, the retrospective allocation of governance power, and externalities among interdependent contributions allow single actors to realize private benefits. These private benefits can be realized through short-term exploits (e.g., airdrop farming and token dumping) or long-term actions to redirect the platform (e.g., intertemporal manipulation including delaying or omitting value-enhancing changes). In line with the dynamic agency theory, we therefore predict:

Hypothesis 2a: Transitions from proprietary to collective governance in firm-sponsored OS platforms through governance token airdrops are related with increased opportunistic behavior in form of airdrop farming by external contributors.

Hypothesis 2b: Transitions from proprietary to collective governance in firm-sponsored OS platforms through governance token airdrops are related with increased opportunistic behavior in form of platform redirection by external contributors.

METHODS

Empirical setting: web3 and governance token airdrops

We test our hypotheses in the emerging web3 industry between January 2015 and June 2025. 2015 was the year when the Ethereum blockchain went live, the first and until now dominant smart-contract platform (DeFiLlama, 2026b). Since then, web3 has experienced enormous growth. The total industry market capitalization has risen from a few billions dollars to approx. three trillion (per 9th of January 2026; CoinMarketCap, 2026). The invention of smart contracts enabled productive application functionalities on top of general-purpose blockchain platforms ("Layer 1s", or "L1s") like Bitcoin, Ethereum, or Solana which initiated broad value creation in the industry. Consequently, many digital platform start-ups have emerged that utilize L1s as their infrastructure layer to execute and settle transactions. Here, we distinguish between platforms on the second (software) and third (applications) layer. On the software layer, platforms extend (e.g., "L2s"), coordinate (e.g., on-chain governance) or modularize (e.g., data availability) foundational infrastructure. On the application layer, user-facing platforms have emerged to serve known or new domains of user needs. For example, the decentralized finance (DeFi) sector has emerged with applications like digital payment solutions through stablecoins as well as decentralized physical infrastructure networks (DePIN) with offerings like decentralized telecommunications services.

Across the industry, all layers are interwoven. Uniswap, one of the largest exchange platforms in web3, uses data provider and execution platforms on downstream layers to offer up-to-date exchange rates and settlement functions. Upstream, Uniswap is integrated in even higher-level application platforms like wallets, aggregators or games that leverage the exchange capabilities bundled in Uniswap's platform. Our data covers a lot of the depth and breadth of the web3 industry. The platforms in our sample represent approx. 47% of the total value held in the industry's smart contracts, and operate in 18 sub-sectors across all three industry layers (Token Terminal, 2026; DeFiLlama, 2026b).

Another key characteristic of the web3 industry is that most platforms are built open source usually via GitHub (van Pelt, Jansen, Baars, & Overbeek, 2021). GitHub is the most popular open-source collaboration platform with more than 180 million active developers and more than 630 million open-source projects (GitHub, 2025). The platform enables collaboration on software development in distinct projects, i.e., repositories. On GitHub, every organization, like web3 platforms, can open an account to jointly develop source code for their core platform and extensions. To enable external contributions, organizations publish code in their repositories. Subsequently, the code is "open sourced" under the defined OS license terms. Each organization can grant contributors different user roles with distinct rights. Usually, internal contributors receive member roles on GitHub that typically include direct writing rights, i.e., to contribute commits directly to the repositories. New users without a role or users with a contributor role are "externals" and do not have direct writing rights. Hence, externals contribute by opening a "pull request". Through pull requests, contributors can propose single or multiple code changes, additions or entirely new elements that will be reviewed by internal developers with direct writing rights. In reviews, external contributions may be accepted or rejected. When accepted, all commits of a pull request are individually merged into the official code repositories. If a pull request is rejected, no commits are stored. After a review, pull requests get closed. Hence, pull request are a technique not only to enable external contributions but implement code control, particularly to ensure quality standards. Each contribution, either through direct commits or a pull request commit, has a unique ID and metadata information. For example, user roles are assigned to each direct or pull request commit, which allows to track the role of each contributor on a commit level, including role changes. The same applies for the review status of pull requests. This enables us to analyze which internal or external developer made which type of contribution at what point in time (Chen et al., 2021; Petryk, Qiu, & Pathak, 2023).

Considering the features and dynamics of the industry, external contributions to develop focal platforms mattered among others to keep track with the industry`s innovativeness, growth trajectory and to ensure proper interoperability with upstream and downstream complements. In this regard, governance token airdrops are a rich setting to examine our research question on how governance transitions affect external platform contributions. The defining feature of airdrops as a token distribution mechanism is that allocations are received for free, unlike in other forms (e.g., private token sales, ICOs) where investors buy a token allocation as an endogenous choice (Allen, 2024; Allen et al., 2023; Chod & Lyandres, 2021; Halaburda et al., 2025). Instead, it is the endogenous choice of the platform sponsor to airdrop at all, the timing as well as recipients and their allocations. Hence, airdrops represent an exogenous governance interventions for external platform contributors. Further, the treatment timing is individual per platform, which is why we apply a staggered difference-in-differences design (Callaway & Sant'Anna, 2021).

Especially in the last two years, blockchain platforms have announced their airdrops in advance which has led to airdrop farming by users (a desired result for some; Halaburda et al., 2025). However, there is no empirical evidence for systematic and large-scale airdrop farming by external developers yet. Though, there is anecdotal evidence for individual farming attempts through

GitHub. In these few cases, potential farmers were easy to identify since they provided one-time, low value “contributions”, like highlighting typos in project descriptions. Such farming attempts are easy to filter out for platforms when they make token allocation choices (Binance, 2024). Consequently, token rewards for contributions likely have a significant higher bar in terms of effort and skill set than for platform usage allocations. However, an anticipation effect cannot be ruled out entirely ex ante and will be tested empirically.

Research design and data sources

The distribution of governance tokens through an airdrop serves as the treatment in our study. We collect all governance token airdrops that were conducted by platforms in the Ethereum or Solana ecosystem, the two largest smart-contract blockchain ecosystems. We identify airdrops through the corresponding database on DeFiLlama, despite its name, a data provider covering web3 platforms across all sub sectors. DeFiLlama does not only provide information about the airdrop date, but also granular data about the single recipient categories as well as their individual allocations and unlock schedules. We verified the main airdrop parameters and that tokens grant governance rights through primary sources, mostly information provided in the platform’s “tokenomics”. Through this data, we further validated the relevance of the treatment in our difference-in-differences estimation. In our setting, a relevant governance intervention should de facto result in a collective governance system that enables governance token holders to make policy proposals and vote on the same following the airdrop. Most of the platforms use major governance decision-making platforms like Discourse (proposal discussions), Snapshot and Tally (policy elections) for this purpose.

Through these sources, we identify 183 platform that conducted an airdrop to distribute tokens free of charge. We excluded the decentralized browser project “Brave” since we could not collect profound evidence that the platform's token (Basic Attention Token) grants governance decision-making rights. Further, the vesting details for "Synthetix", a derivatives exchange, do not provide information that tokens were also distributed to external contributors, which led to the exclusion of this platform. To ensure comparability, we only included platforms that apply an open-source development model considering the special interdependencies and dynamics in such open forms of collaboration (Baldwin & Von Hippel, 2011; Yue & Nagle, 2025), which we identify through GitHub. The open-source filter reduces our sample to 151 platforms.

The unit of treatment in our study is the unique pair of an external contributor and the platform the unique committer is contributing to. Consequently, contributors inherit the treatment date from the platform they are contributing to. In our main analysis, we apply a binary treatment, for which the day of the initial distribution of tokens to the recipients is defining the treatment date for the unique contributor-platform pair. With that date, platform governance is transferred to the DAO through which governance token holders can participate in governance decision-making (Allen et al., 2023; Beck et al., 2018; Hsieh et al., 2018). Hence, the initial token distribution is the pivotal moment initiating the governance transition. Since tokens usually unlock over a vesting schedule, i.e., only become transferable over time, the day of the last token unlock represents the end of the governance transition. In our secondary analysis with a continuous design, we define the single token unlock dates over the vesting schedule as the continuous treatment dates.

The distribution of airdrops exhibits substantial temporal heterogeneity. Too small cohorts and too few periods can cause instable estimates. For this purpose, we implement several methodological filters to ensure valid causal inference. First, we implement cohort viability filters requiring each treatment cohort to contain at least three contributor-platform pairs and at least two pre-treatment periods permit parallel trends assessment. Organizations in cohorts failing either threshold are excluded from estimation. Second, we require each time period to contain at least five control units to ensure valid counterfactual construction, trimming periods with insufficient control support. Following estimation, we trim event-study coefficients to a symmetric ±20 quarter window around treatment to focus on estimable event-times, though this affects only reported coefficients rather than the estimation sample itself. These filters yield a final sample of 615,979 contributor-platform-quarter observations spanning 346,487 unique contributor-organization pairs across 127 treated contributor-platform pairs and 42 quarters. The airdrop activity in our final estimation sample accelerated dramatically during "DeFi Summer" (2020Q3; 6 airdrops), and peaked in 2021-2022 (54 airdrops) and 2023-2024 (48 airdrops). The two most active quarters in our sample are 2022Q2 and 2024Q4, with 12 airdrops each.

For our counterfactual, we aimed to collect all platforms in the Ethereum and Solana ecosystems without an active governance token and, hence, without a decentralized governance system in place and an airdrop conducted. Again, these platforms needed to apply an open-source development model. We identify potential platform candidates through the potential platform airdrop database on DeFiLlama. We collect open-source data for these platforms also through GitHub. Through this approach, we identify a control group sample of 51 platforms.

For our main analyses, we apply a "never treated" definition for control units. Hence, in these analyses, we only consider contributor-platform pairs for which the platform never conducted an airdrop within our observation period. For robustness purposes and considering the heterogeneity in the airdrop activity within our treated sample, we also ran alternative estimations for which we apply a "not-yet treated" definition. Here, platform-contributor pairs were assigned to the control group until the last pre-treatment period. Never-treated as well as not-yet-treated analyses show consistent results.

Measures

Dependent variable (H1) – external contribution share. We measure external contributions to firm-sponsored open-source platforms as the quarterly share of pull-request commits by authors without a member role (i.e., external contributors) for the respective GitHub organization out of all commits received by that organization. The numerator captures all commits submitted via pull request by external contributors across all repositories of the organization in a given quarter (Chen et al., 2021; Petryk et al., 2023). Authors are identified through their user ID rather than their username, as user IDs are stable over time. The denominator comprises all commits received by the organization in the same quarter: direct commits and pull-request commits by internal contributors (i.e., organization members) and external contributors. The resulting ratio ranges from zero to one. It captures the composition of contribution activity, specifically, what share of all incoming commits originated from outside the organization.

Dependent variable (H2a) – airdrop farming. We measure contributor opportunism in the form of airdrop farming as the churn ratio of external contributors for a given platform. Analogously to user airdrop farming, we would expect external developers that contribute solely in anticipation of a token reward to discontinue their participation via GitHub (Allen et al., 2023; Halaburda et al., 2025). Hence, airdrop farming by external developers is measured as the share of external contributor-platform pairs active in quarter t that do not contribute in quarter t+1, computed by comparing the set of external contributors across consecutive quarters and expressed as a percentage (range: 0-100). This organization-level metric captures "hit-and-run" behavior or low contributor retention, where higher values indicate that a larger proportion of external contributors disappear after their initial contribution period.

Dependent variable (H2b) – platform redirection. We measure contributor opportunism in the form of platform redirection by observing signals for reduced quality of external contributions. In line with evidence provided by O'Mahony & Karp (2022), contributors may exploit influential positions in a collective governance mode to omit, or delay certain feature request. We expect that such a behavior would lead to lower acceptance rates of external pull requests or longer review periods. Therefore, we measure implications from platform redirection activities through a backlog rate. It measures the share of external pull requests created in quarter t that were either never resolved (merged or closed without merge) or resolved slower than the platform-quarter's own median time to merge pull requests, using a nested design where the threshold is dynamically set to each platform-quarter's median time-to-merge (range: 0-1 fraction). This load-adjusted metric (i.e., adjusted for variation in external pull request activity) captures capacity strain or processing inefficiency, with higher values indicating that a larger fraction of arriving external PRs experience delayed resolution or become stuck in the backlog.

Independent variable. Since we implement a difference-in-differences analysis, our main explanatory variable is an interaction between the treated group and post-treatment periods. A contributor-organization pair was considered as treated when the contributions were made to a platform that conducted an airdrop within our observation period. The first post-treatment period is defined as the quarter in which the initial token distribution event occurred for a given platform.

Additional analyses – Airdrop design variables. We create two variables to capture airdrop design choices for additional analyses (see Discussion), i.e., 80% vesting duration and unlock share at treatment. 80% vesting duration is defined as the minimum number of quarters required for 80% of airdrop category tokens to unlock, only considering allocations not assigned to insiders (i.e., platform sponsor and investors), according to the platform's token allocation details. This metric captures the time-to-substantial-unlock threshold for external token distributions, considering that tokens unlocks usually reflect an S-shaped distribution curve (Lommers et al., 2023). Unlock share at treatment measures the percentage of airdrop tokens unlocked in the first quarter post-treatment versus the cumulative token unlocks across the full vesting schedule. This design variable quantifies initial unlock intensity for external allocations.

ANALYSIS AND RESULTS

Estimation strategy

Difference-in-differences. We employ a staggered difference-in-differences research design using the Callaway and Sant'Anna (2021) group-time average treatment effect on the treated (ATT) estimator, which addresses selection bias by controlling for time-invariant unobserved heterogeneity across individual-organization pairs while allowing treatment timing to vary across organizations. This approach corrects for potential bias through selection when treated and control groups differ in outcome levels but unobserved differences between groups remain constant over time, consistent with the parallel trends assumption (Abadie, 2005). The estimator employs panel differencing to remove both time-invariant individual-organization fixed effects and common time shocks through the double-differencing structure inherent to the method. By comparing within-unit changes in outcomes for treated contributor-platform pairs against within-unit changes for control pairs, the method effectively nets out all time-constant unobserved heterogeneity at the contributor-platform level. The estimator aggregates cohort-specific treatment effects across contributor-platform pairs into dynamic event-study parameters that capture how treatment effects evolve relative to the intervention timing, avoiding the well-documented bias from treatment effect heterogeneity under staggered adoption that afflicts two-way fixed effects estimators (Goodman-Bacon, 2021).

While governance tokens are typically unlocked gradually and with heterogeneous intensities over time, we define treatment as a binary first-exposure event because the initial token unlock (i.e., transition start) constitutes the discrete institutional change that distributes governance rights and alters the platform's governance design. Variation in subsequent unlock schedules and token intensities is therefore analyzed in robustness and extended specifications, rather than folded into the baseline treatment definition. Treatment is assigned at the organization level, but effects are measured at the individual-organization level to isolate how the same contributor responds differently across treated versus control organizational contexts. Since undoing this governance transition is, in practice, highly complicated if not impossible (and has not been seen yet), we consider the treatment as absorbing (Baker, Callaway, Cunningham, Goodman-Bacon, & Sant'Anna, 2025).

Standard errors are clustered at the organization level to account for within-organization correlation in outcomes across individual contributors and over time, recognizing that treatment assignment occurs at the organizational level.

Formally, for contributor–platform pair i in platform j, observed in quarter $\tau \in \{g - 1, t\}$, we estimate for each treatment cohort g and evaluation period t the group–time average treatment effect on the treated as:

$$ATT(g, t) \equiv \mathbb{E}[DV_{ij,t} - DV_{ij,g-1} \mid G_{ij} = g] - \mathbb{E}[DV_{ij,t} - DV_{ij,g-1} \mid C_{ij} = 1].$$

Here, $DV_{ij,\tau}$ denotes the outcome of contributor–platform pair i in platform j observed in quarter τ. The variable G_{ij} denotes the treatment cohort of the contributor–platform pair, defined as the quarter in which the platform conducts its initial governance-token airdrop. Contributor–platform pairs with $G_{ij} = g$ are first treated in cohort g, while $C_{ij} = 1$ indicates never-treated contributor–platform pairs, i.e., pairs associated with platforms that do not conduct a governance-token airdrop during the sample period. The first term captures the average within-pair change in the outcome for contributor–platform pairs first treated in cohort g, measured between the pre-treatment reference quarter $g - 1$ and quarter t. The second term captures the corresponding outcome change for never-treated contributor–platform pairs over the same calendar periods. Under an unconditional parallel trends assumption, the outcome evolution of never-treated contributor–platform pairs provides a valid counterfactual for the treated cohort, such that the difference between these two terms identifies the group–time average treatment effect on the treated, $ATT(g, t)$, in the Callaway and Sant'Anna (2021) framework. This procedure is repeated for all relevant combinations of cohorts g and post-treatment periods t, yielding a set of cohort- and time-specific treatment effects that allow treatment effects to vary flexibly across adoption timing and exposure duration, while avoiding the implicit weighting and bias issues associated with two-way fixed effects estimators in staggered adoption settings.

Parallel trends assumption. To examine whether we can detect any signs of violations of the parallel trends assumption, we investigate the sample raw data, in particular regarding the outcome dynamics between treated and control groups (Callaway, Goodman-Bacon, & Sant'Anna, 2024; Cunningham, 2021). Figures B1-B3 show the trends for our three dependent variables. We plotted quarterly means of treated versus control contributor-platform pairs. B1 provides supportive evidence of parallel trends. B2 and B3 indicate potentially different levels and trends between treated and control groups, requiring further investigation.

Figures B1-B3: Visualization of trends (quarterly means) in the external contributions share (B1), developer airdrop farming (external contributor churn; B2) and platform redirection activities (B3) with median treatment timing and +/- 1 SD deviation in treatment timing

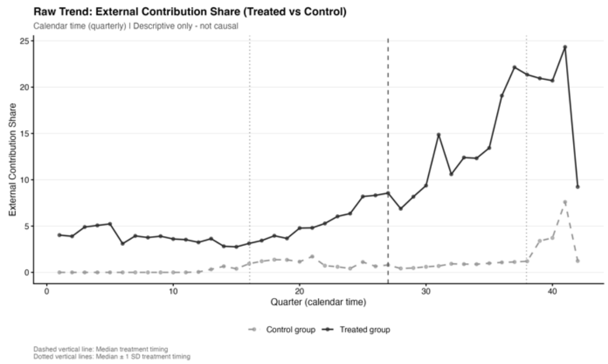

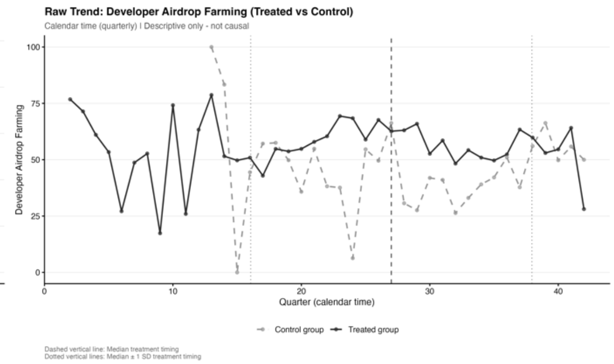

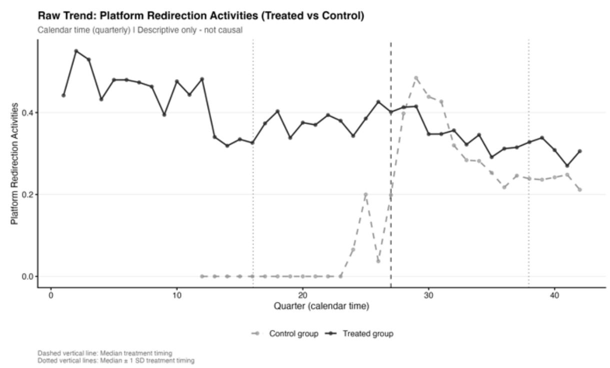

Subsequently, we conducted an event study to examine whether the assumption of comparable pre- and post-treatment dynamics between the treated group versus the control group holds. Here, we include placebos in a regression for anticipatory effects and post-treatment effects, i.e., leads and lags, respectively (Cunningham, 2021). We first generate visual event-study plots displaying point estimates and 95% confidence intervals for all pre-treatment and post-treatment periods to enable visual inspection of trend patterns (Figures B4-B6). These plots provide an intuitive assessment of whether treatment and control groups followed parallel trajectories before the individual airdrops and allow identification of anticipation effects or delayed treatment responses. The plots confirm flat pre-treatment coefficients clustered around zero with overlapping confidence intervals, visually corroborating the statistical tests of parallel pre-treatment trends for two dependent variables, i.e., external contribution share and airdrop farming. Pre-treatment coefficients for the platform redirection activities are more volatile. We achieve further support for the assumption of parallel tests through applying a Wald test to evaluate whether the lead coefficients jointly differ from zero before the airdrop. For external contribution share and airdrop farming variables, the test yields p = 0.587 and p = 0.999 respectively, indicating no statistically significant pre-treatment differences between groups and supporting the parallel trends assumption. However, confirming visual indications, pre-trend diagnostics reveal a violation of the parallel trends assumption for the estimation of effects on platform redirection activities (joint Wald p < 0.000, 8/19 significant leads), further discussed below.

Figures B4-B6: Event study plots: Estimates of airdrop's effects on the external contributions share (B4), developer airdrop farming (external contributor churn; B5) and platform redirection activities (B6) using leads and lags (95% CI)

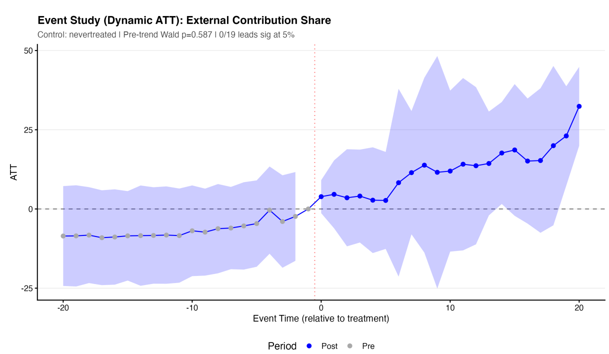

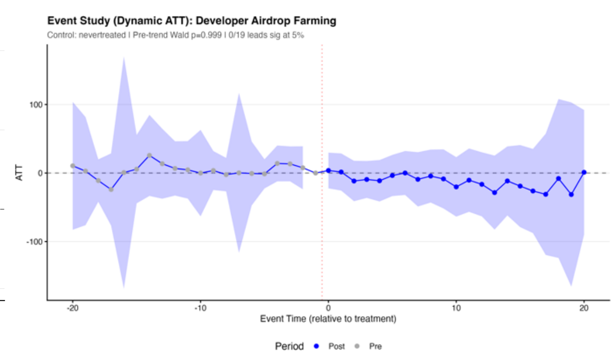

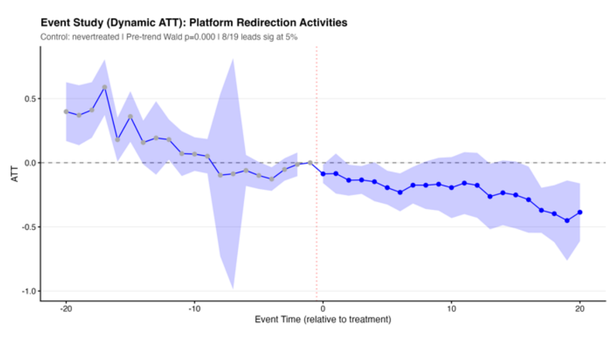

We implement different measures to ensure and verify post-treatment stability in differences between control and treatment groups (Abadie, 2005; Cunningham, 2021). Unlike traditional firm settings with stable firm populations (Thatchenkery & Katila, 2023), open-source contributions are autonomous and often task-specific, creating inherent volatility in contributor appearance and the number of contributor-platform pairs that we can observe per period (Petryk et al., 2023; Yue & Nagle, 2025). Further, the entire web3 industry showed significant growth, especially from 2020 onwards. This is also reflected in the extent of industry contributors. To address this volatility while maintaining valid counterfactual comparisons, we implement minimum control support thresholds that ensure viable comparison groups exist throughout the analysis window despite the dynamic entry and exit of contributors. Therefore, we require at least 100 contributor-platform pairs per observation period for viable comparison groups. Overall, we observe a median of 1,939 control units per period; with early periods ($t \leq 20$) showing a median of 418 units, while later periods reflect Web3's emergence as a growing industry with a median of 8,764 units.

Further, we examine pre-treatment period coefficients to detect anticipatory behavioral responses that would violate the stable groups assumption by indicating firms altered behavior before formal treatment implementation (Wing, Simon, & Bello-Gomez, 2018). If anticipation effects were present, we would observe statistically significant negative or positive coefficients in periods immediately preceding treatment, contaminating the clean identification of post-treatment effects. The systematic absence of significant lead coefficients (zero significant leads across 19 pre-treatment periods for external pull requests, with p-values exceeding 0.60 for periods $t=-2$ through $t=-4$; $t=-1$ is the last pre-treatment period, hence, serving as the universal base period) confirms no anticipatory responses, supporting the stable groups assumption. Also, none of the 19 pre-treatment periods show significant coefficients for airdrop farming. We detect eight significant pre-treatment periods for platform redirection activities, however, in the first observation periods, not close to treatment.

Estimation results

Descriptive statistics are reported in Table A1. External contributors are responsible for on average 8.94% of total contributions to a given organization made through pull requests per quarter during the study period (2015–2025), though the standard deviation is substantial (27.03). External contributor churn to measure airdrop farming averaged 54.70% per quarter, indicating that more than half of external contributors in a given quarter do not return in the subsequent period. Regarding platform redirection activities, organizations exhibited an average backlog rate of 0.32 (32%), meaning approximately one-third of external pull requests were resolved slower than the organization's median merge time or remained unresolved. Airdrop design variables show that the average organization required 12.49 quarters (approximately 3.1 years) for 80% of airdrop tokens to unlock, with a mean upfront unlock share of 0.22 (22%) in the first post-treatment quarter.

Correlations among variables are predominantly low, with the notable exception of a moderate negative correlation between vesting duration and upfront unlock share ($r = -0.33$, $p < .05$), intuitively indicating that organizations with longer vesting schedules tend to distribute smaller initial token allocations. All other pairwise correlations fall below 0.17 in absolute magnitude, suggesting minimal multicollinearity concerns. The correlation structure is consistent with theoretical expectations: design variables (vesting duration and upfront share) exhibit interdependence reflecting strategic token allocation choices, while outcome variables (external contribution share, airdrop farming, platform redirection) show weak correlations

Table A1: Descriptive statistics and correlations

	Variable	Mean	SD	1	2	3	4	5
1	External Contributions (Share)	8.94	27.03		0.15	0.16	0.09	0.01

	Variable	Mean	SD	1	2	3	4	5
2	Airdrop farming	54.70	27.10			0.08	0.06	-0.07
3	Platform redirection	0.32	0.21				0.09	0.07
4	80% Vesting Duration	12.49	10.82					-0.33
5	Unlock Share at Treatment	0.22	0.27					

Note: 615,979 contributor-platform-quarters (External Contributions measured at IxO-quarter level). Correlations calculated at org-quarter level (N = 1,746). Correlations above .10 are significant at $p < .05$.

Hypothesis H1 predicted that externals would increase their contributions to platforms that conducted an airdrop more than to platforms that never conducted an airdrop. The difference-in-differences estimator for the average post-treatment effect is positive and statistically significant for the quarterly external pull request share ($\beta = 12.95$ percentage points, $p = .000$). On average across post-treatment periods, the share of external contributions among all contributions increases by approximately 13-percentage points, supporting H1.[1]

In line with dynamic agency theory, we made two additional predictions about contributor opportunism following the treatment. In hypothesis H2a, we predicted that governance token airdrops would lead to higher external contributor churn, indicating airdrop farming, than for platforms that never conduct an airdrop. The difference-in-differences estimator for the average post-treatment effect is negative and statistically significant at 10%-level ($\beta = -12.91$, $p = .068$). On average across post-treatment periods, the share of external contributors who disappeared in the following quarter decreased by 13 percentage points for treated organizations relative to controls. This rather suggests improved external contributor retention post-unlock instead of contributor opportunism through airdrop farming. Further, pseudo leads and lags in the event study

[1] Results confirmed for the simple count measure, showing a statistically significant increase by 4.90 external contributions per quarter relative to controls ($\beta = 4.90$, $p = .002$).

estimation do not show any significant coefficients when focusing on four periods pre and post treatment (i.e., +/- one year). These leads and lags close to treatment would have been a more granular sign for anticipation effects and sharp exits after treatment. Hence, we must reject H2a.

Hypothesis H2b predicted that the backlog rate of external pull requests would increase in platforms that conducted an airdrop relative to never-airdropping platform organizations, reflecting signs of contributor opportunism through platform redirection activities. The difference-in-differences estimator for the average post-treatment effect is negative and highly statistically significant ($\beta = -0.231$, $p < .000$). However, pre-trend diagnostics reveal a violation of the parallel trends assumption (joint Wald $p < .000$, 8/19 significant leads). The magnitude assessment ratio (0.76) suggests post-treatment effects are larger than pre-treatment deviations. Further, the six statistically significant pre-treatment coefficients are the earliest observation periods with lower observation coverage. Yet, this result cannot be interpreted causally. Nevertheless, since the results rather suggest improved than degraded external pull request processing, we do not find evidence in support of H2b and, hence, reject it. Estimation results for H1, H2a, and H2b are jointly reported in Table A2.

Table A2: Difference-in-Differences Estimates (Never-Treated Control Group)

DV focus	**External Contributions (H1)**	**Airdrop Farming (H2a)**	**Platform Redirection (H2b)**
Airdrop	12.950****	-12.906*	-0.231****
	{0.000}	{0.068}	{0.000}
	(2.466)	(7.061)	(0.024)
N (observations)	615,979	344,390	377,499
Joint Wald test (p-value)	0.587	0.999	0.000****

Note: This table presents average post-treatment ATT estimates from staggered difference-in-differences analyses using the Callaway & Sant'Anna (2021) estimator. The estimation uses never-treated units as controls. p-values in curly brackets, standard errors in parantheses. Significance levels: **** p<0.001, *** p<0.01, ** p<0.05, * p<0.10. Joint Wald tests assess the null hypothesis of no pre-treatment differential trends. N deviates due to cohort viability criteria.

Robustness and Sensitivity Analyses

We subject our main findings to a comprehensive battery of robustness and sensitivity analyses. Table A4 summarizes 19 checks addressing baseline estimation, functional form, sample restrictions, specification sensitivity, omitted variable bias, and falsification. Below, we describe each category and highlight key results.

Design robustness: alternative control group definitions. We implement two control group definitions to assess sensitivity to the composition of the comparison group and the strength of the required parallel trends assumption, following standard practice in staggered difference-in-differences designs (Callaway & Sant'Anna, 2021). The never-treated control group consists exclusively of individual-organization pairs in organizations that never receive treatment throughout the observation period, providing a fixed counterfactual across all time periods but requiring that these permanently untreated units follow parallel trends with each treated cohort; i.e., an assumption that may be strained if never-treated units differ systematically from eventually-treated ones. The not-yet-treated control group includes both never-treated units and units that have not yet been treated at the relevant time period, expanding the comparison set with units that are arguably more comparable to treated units (since they too eventually adopt treatment), but imposing a more restrictive parallel trends assumption: that later-treated units would have evolved in parallel with earlier-treated units during their respective pre-treatment windows. A systematic comparison across both specifications assesses whether our findings depend on the choice of counterfactual population and, by extension, on the specific form of the parallel trends assumption invoked.

Post-treatment effect estimates are mostly consistent across control group definitions for all three dependent variables. For external contributions, effects and parallel trends diagnostics are substantively identical across both specifications. For airdrop farming, the never-treated specification yields larger estimated reductions. However, pre-treatment diagnostics show violations under not-yet-treated controls. The coefficient estimates for platform redirection activities are smaller, yet highly significant. The pre-trend violations are also prevalent under not-yet-treated control specifications, though. Table A3 reports estimation results for H1, H2a, and H2b based on alternative control group definitions.

Table A3: Difference-in-Differences Estimates (Not-Yet-Treated Control Group)

DV focus	External Contributions (H1)	Airdrop Farming (H2a)	Platform Redirection (H2b)
Airdrop	11.100****	-7.916*	-0.111***
	{0.000}	{0.263}	{0.000}
	(2.486)	(7.072)	(0.022)
N (observations)	615,979	344,390	377,499
Joint Wald test (p-value)	1	0.000****	0.000****

Note: This table presents average post-treatment ATT estimates from staggered difference-in-differences analyses using the Callaway & Sant'Anna (2021) estimator. The estimation uses not-yet-treated units as controls. p-values in curly brackets, standard errors in parantheses. Significance levels: **** p<0.001, *** p<0.01, ** p<0.05, * p<0.10. Joint Wald tests assess the null hypothesis of no pre-treatment differential trends. N deviates due to cohort viability criteria.

Specification and sample robustness I: functional form. Open-source contribution data is heavily right-skewed with a substantial mass at zero on a quarterly level, i.e., most developers only contribute once. To ensure our results are not artifacts of the level specification, we re-estimate all models using a log(1+y) transformation. The log specification yields qualitatively identical results for external contribution share ($\beta = 0.72$, $p < .001$) and platform redirection ($\beta = -0.18$, $p < .001$), confirming that effect directions and approximate magnitudes are robust to functional form (Baker et al., 2025; Roth & Sant'Anna, 2023). Airdrop farming remains statistically insignificant ($\beta = -0.23$, $p = .588$), consistent with the baseline null.

We further decompose effects into extensive and intensive margins (Loh & Kretschmer, 2023). The extensive margin, i.e., the probability of any contribution, shows a significant increase for external contribution share ($\beta = 0.17$, $p < .001$), suggesting that airdrops attract new contributors rather than merely increasing volume from incumbents. The intensive margin, conditional on positive participation, produces insignificant estimates across all outcomes, indicating that the aggregate treatment effects are primarily driven by changes in participation rates rather than individual-level quantity adjustments.

Specification and sample robustness II: sample restrictions. Our panel includes contributor-platform pairs with heterogeneous observation frequencies (mean 1.78 quarters per pair), a typical contribution pattern in open-source software (Hoffmann, Nagle, & Zhou, 2024; Loh & Kretschmer, 2023; Nagle, 2018; Nagle, Dana, Hoffman, Randazzo, & Zhou, 2022). Hence, we test whether our results are driven by one-time observations by restricting to contributors with at least two panel observations. Estimates remain stable: external contribution share ($\beta = 12.16$, $p < .001$) and platform redirection ($\beta = -0.23$, $p < .001$) retain significance and direction, while airdrop farming remains null ($\beta = -11.07$, $p = .141$).

Several treatment cohorts contain few organizations, raising concerns about idiosyncratic variance. Dropping singleton cohorts ($\beta = 11.27$, $p < .001$) and cohorts with fewer than three organizations ($\beta = 9.81$, $p = .003$) attenuates precision for external contribution share but preserves significance for platform redirection ($p < .001$ for both). Restricting to single-organization contributors, i.e., who participate in exactly one platform, addresses stable unit-treatment value assumption (SUTVA) concerns arising from multi-platform spillovers (Rubin, 1980). With 84% of contributor-platform pairs involving multi-org contributors, this restriction substantially reduces the sample. Platform redirection retains its direction and significance ($\beta = -0.18$, $p = .006$), while external contribution share remains directionally consistent ($\beta = 13.05$, $p = .071$) with wider confidence intervals.

Specification and sample robustness III: coarsened exact matching. To address potential covariate imbalance between treated and control organizations, we apply Coarsened Exact Matching (CEM; Iacus, King, & Porro, 2011, 2012) on baseline platform characteristics, i.e., platform age, sector, and OS license regime. CEM-matched estimates corroborate the main findings under parsimonious matching: external contribution share shows a positive effect ($\beta = 13.05$ $p < .001$) and platform redirection a negative effect ($\beta = -0.23$, $p < .001$) when matching based on platform age, with similar magnitudes when matching based on platform sector. Point estimates remain directionally consistent across all five specifications, though precision decreases as matching constraints tighten and matched sample sizes shrink; i.e., from 154 organizations under age-only to 102 under the full three-covariate design. The most restrictive CEM specification yields directionally consistent but statistically insignificant estimates ($\beta = 11.72$, $p = .151$ for contribution share; $\beta = 0.05$, $p = .103$ for platform redirection), reflecting the efficiency cost of exact matching on multiple covariates simultaneously rather than a substantive change in the treatment effect.

Specification and sample robustness IV: specification sensitivity. We test sensitivity to anticipatory behavior by allowing one-quarter anticipation in the Callaway and Sant'Anna (2021) estimator. If contributors adjust behavior before the formal airdrop in response to announcements or speculation, failing to account for anticipation biases group-time ATTs. Results with anticipation

= 1 are virtually unchanged: external contribution share ($\beta = 14.18$, $p < .001$) and the point estimates remain within the baseline confidence interval, suggesting minimal pre-treatment contamination.

Event-window length determines which post-treatment periods contribute to the dynamic ATT and introduces compositional imbalance as longer windows include fewer late-treated cohorts. We estimate dynamic ATTs across four windows: ±8, ±12, ±16, and ±20 quarters. External contribution share grows from 6.41 ($p = .060$, ±8Q) to 12.95 ($p < .001$, ±20Q), with the effect strengthening as the window lengthens. Platform redirection is consistently significant across all windows ($p < .001$ for all), ranging from −0.16 (±8Q) to −0.23 (±20Q). This pattern is consistent with treatment effects that accumulate over time rather than artifacts of window selection.

Sensitivity to unobserved confounding. Unobserved time-varying confounders, i.e., changes in project quality, team composition, or ecosystem-level shocks, may correlate with both airdrop timing and contribution outcomes. We assess this threat using (Oster, 2019) δ^* bounds, computed from a pooled two-way fixed effects approximation of our main specification. The bounds quantify how much stronger selection on unobservables would need to be, relative to selection on observables, to fully explain the estimated treatment effect.

For our primary outcomes, δ^* values are: external contribution share ($\delta^* = -18.88$), platform redirection ($\delta^* = -267.78$), and airdrop farming ($\delta^* = -0.54$). The conventional threshold for robustness is $|\delta^*| > 1$, meaning unobservables would need to be more influential than observables to explain away the effect (Oster, 2019). For external contribution share, unobservables would need to be nearly 19 times more influential than our baseline controls (platform age, sector, OS license regime) to nullify the treatment effect. For platform redirection, the threshold exceeds 267 times. Airdrop farming, our null result, shows $|\delta^*| < 1$, consistent with its lack of statistical significance in the baseline specification. These bounds indicate that our significant results are highly unlikely to be driven by omitted variable bias.

Falsification and randomization inference I: permutation tests and Fisher randomization.

Model-based inference in difference-in-differences designs may understate uncertainty when the number of treated cohorts is small (Bertrand, Duflo, & Mullainathan, 2004; MacKinnon & Webb, 2020; Young, 2019), and p-values may be sensitive to the aggregation method and systematically too small. We therefore verify our findings using randomization inference, which derives p-values from the permutation distribution rather than from analytical standard error formulae.

We randomly reassign the full treatment regime, i.e., which organizations received airdrops and when, including never-treated status, across organizations 500 times, preserving the marginal distribution of cohort sizes. For each reassignment, we re-estimate the Callaway and Sant'Anna (2021) model, generating a distribution of estimates under the null hypothesis of zero treatment effect. If our observed effect is large relative to this distribution, it is unlikely to arise by chance alone. We implement two complementary variants. The Fisher randomization test compares a studentized test statistic (t = ATT/SE) against its permutation distribution, providing exact, distribution-free p-values (Athey & Imbens, 2022; Fisher, 1936; MacKinnon & Webb, 2020). The permutation test compares the raw ATT magnitude against the distribution of permuted ATTs, providing a fully standard-error-free benchmark (Bertrand et al., 2004). The key distinction between these tests and our baseline concerns the identifying assumption. Our baseline requires only that untreated potential outcomes follow parallel trends across treatment cohorts. Randomization inference requires the stronger condition that treatment assignment is as good as random (Athey & Imbens, 2022). When results remain significant under this more demanding condition, they provide evidence beyond what parallel trends alone can deliver.

Both tests confirm our main findings. The external contribution share is significant under Fisher randomization ($p = .002$) and the permutation test ($p = .004$). Platform redirection is significant under both approaches (Fisher $p = .032$; permutation $p = .002$). The airdrop farming proxy remains nonsignificant (Fisher $p = .553$; permutation $p = .118$), consistent with the baseline null. These results are particularly informative for platform redirection, where pre-trend tests detect parallel trends violations. While the randomization tests do not correct for potential bias in the point estimate arising from differential pre-trends, they establish that the observed effect is too large to be attributed to random variation in treatment assignment alone – providing evidence that the platform redirection effect reflects a genuine treatment response rather than an artifact of model-based inference. The convergence of both tests across all outcome variables, with identical significance conclusions, indicates that our baseline estimates are not artifacts of distributional assumptions, the choice of aggregation method, or finite-sample bias in the standard error computation. The external contribution share and platform redirection effects survive escalation from the weaker parallel trends assumption to the stronger condition of random treatment assignment.

Falsification and randomization inference II: placebo tests. If our estimates reflect pre-existing trends rather than treatment effects, shifting treatment dates into the pre-period should produce non-null results (Wing et al., 2018). We construct a shifted-treatment placebo by moving each organization's treatment date four quarters earlier. The dynamic ATT under the shifted placebo is null for external contribution share ($p = .383$), and airdrop farming ($p = .692$), confirming that the pre-treatment period does not exhibit dynamics that could masquerade as treatment effects for these outcomes. For platform redirection, however, the shifted placebo yields a significant positive coefficient ($p = .001$), consistent with the pre-trend violations documented in the parallel trends diagnostics. This reinforces our earlier caveat that causal claims for platform redirection should be interpreted with care, even though randomization inference results confirm that the general effect is unlikely to arise by chance.

A distinct concern in the cryptocurrency context is that major market events, e.g., DeFi Summer 2020 and the 2022 crypto crash, coincide with several airdrop dates, potentially confounding treatment effects with market-wide dynamics. We implement a concurrent-shock placebo by assigning fake treatment dates corresponding to these market events to never-treated organizations. The pooled dynamic ATT across both shock events is null for external contribution share ($p = .880$), airdrop farming ($p = .753$), and the supporting count measure ($p = .985$). For platform redirection, the concurrent-shock placebo is marginally significant ($p = .094$), again consistent with the known pre-trend fragility for this outcome. Taken together, the placebo tests confirm that our results for external contribution share and airdrop farming are not driven by concurrent crypto market dynamics, while the platform redirection estimates warrant the interpretive caution noted above.

Table A4 consolidates all 19 robustness checks. The overall pattern is clear: the positive effect of token airdrops on external contribution share (H1) is robust across functional forms, matching specifications, and specification variants, and survives demanding tests of omitted variable bias ($\delta^* = -18.88$) and randomization inference (Fisher $p = .002$). The platform redirection effect (H2b) maintains statistical significance across virtually all specifications, windows, and randomization inference (Fisher $p = .032$), though placebo tests for this outcome reflect the pre-trend fragility documented in our parallel trends diagnostics. The null finding for airdrop farming (H2a) is also robust: no specification or sample restriction produces a significant churn effect, and the Oster bound confirms that the null is not merely a consequence of low power ($\delta^* = -0.54$). The falsification tests, i.e., shifted-treatment placebos and concurrent-shock placebos, produce null effects for external contribution share and airdrop farming, strengthening the causal interpretation for these outcomes.

Table A4: Summary of robustness checks

Concern	Test
Design robustness	
Results sensitive to control group definition	Never-treated vs. not-yet-treated control groups
Specification and sample robustness	
Results are artifacts of level specification	Re-estimate with log transformation
Effects driven by participation vs. quantity changes	Extensive margin $P(y > 0)$ and intensive margin $(y \mid y > 0)$ decomposition
Results driven by one-time observations	Restrict to contributors with ≥ 2 panel observations
Idiosyncratic variance from small cohorts	Drop singleton cohorts; drop cohorts with < 3 organizations
SUTVA violation from multi-platform contributors	Restrict to single-organization contributors
Covariate imbalance between treated and control	CEM on baseline platform characteristics (age, sector, portfolio); 5 specifications
Pre-treatment contamination from anticipation	Allow 1-quarter anticipation in CS estimator
Results sensitive to event-window length	Dynamic ATTs across ±8, ±12, ±16, ±20 quarter windows
Sensitivity to unobserved confounding	
Omitted variable bias from unobserved confounders	Oster (2019) δ^* bounds; threshold $\lvert\delta^*\rvert > 1$
Falsification and randomization inference	
Model-based standard errors may understate uncertainty with sparse cohorts	Permutation test (500 reshuffles) and Fisher randomization of studentized t-statistic
Effects reflect pre-existing trends	Shifted-treatment placebo (−4 quarters)
Effects confounded by concurrent market events	Concurrent-shock placebo (DeFi Summer 2020 & 2022 Crash)

DISCUSSION AND CONCLUSION

This study aimed to examine how platform contributions by autonomous externals change when a focal platform initiates a transition from proprietary to collective governance. Drawing on property rights theory and evidence from platform governance literature, we introduced a seeming tension: platform governance transitions may incentivize external contributions as they reduce externals' threat of ex-post platform owner opportunism. However, it can also harm the platform when external contributors exploit their new influence through collective governance for private benefits. We investigated governance token airdrops as a particular mechanism to transition from proprietary to collective governance. Using a novel data set on web3's blockchain-based platforms between 2015 and 2025, we show how reallocating platform control rights from platform sponsors to external participants and contributors resulted in increased external contributions. However, our analyses did not provide evidence for opportunistic behavior by externals enabled through collective platform governance. We discuss this contradiction with predictions from dynamic agency theory against the unique features of governance tokens and airdrops.

In dynamic settings without commitment, agents' actions are valuable yet non-verifiable, and principals cannot credibly condition future rewards or authority on realized performance through enforceable contracts. Once rewards or control rights are granted, they are typically irreversible, and ex-post discipline is infeasible (Fong & Li, 2017; Levin, 2002, 2003; MacLeod & Malcomson, 2023, 1989; Rayo, 2007). This gives rise to intertemporal moral hazard: agents may undertake actions that generate immediate private benefits while shifting quality or timing costs into the future, anticipating that such distortions cannot be sanctioned once authority or rewards have been allocated (Georgiadis, 2022).

In response to this problem, the contracting-under-moral-hazard literature emphasizes backloading and continuation-value mechanisms as central incentive instruments in dynamic environments, even in the absence of formal commitment (Georgiadis, 2022). Rather than conditioning whether rewards are paid, optimal designs condition when and how fully payoffs can be realized, so that current behavior affects the agent's future payoff from the relationship (Georgiadis, 2022). By delaying monetization or effective payoff, such mechanisms make short-term opportunism less attractive when it undermines future value.

Applied to uncontracted contributions to firm-sponsored open-source platforms, governance token airdrops may contain features of such optimal contracts according to contracting under moral hazard literature and hence prevent opportunism. What governance transitions through airdrops distinguishes from other cases examined by strategy scholars (O'Mahony & Karp, 2022; Yue & Nagle, 2025) is the inseparable dual-incentive role of governance tokens: they grant both governance power and monetary incentives. For contributors, the governance power granted through the token increases value capture prospects of their investments usually monetized through developments or services being add-ons to the platform. In addition, since the token reflects a portion of the platform value, the airdrop introduces another value capture stream, i.e., through the platform itself. Therefore, the governance token airdrop transforms the initial principal-agent setting into a principal-principal scenario. Governance tokens need to be hold (and often staked) to utilize the granted governance rights. And because token prices aggregate expectations about platform value (Makridis et al., 2023), contributors internalize at least temporarily the consequences of their actions for future wealth: harmful behavior by contributors would, hence, lead to a token price decrease that negatively compensates for potential private benefits. Consequently, incentives among contributors holding governance tokens are aligned through the market value of governance tokens which eventually prevents opportunistic behavior by contributors. In other words, dynamic agency theory predicts opportunism through granted power.

In turn, the value alignment with the core platform might have prevented opportunism in line with the literature on optimal contracting under moral hazard.

In particular, token vestings operationalize the logic of optimal contracts in dynamic settings without commitment through delayed monetization rather than discretionary performance pay. Governance tokens are allocated retrospectively and ownership itself cannot be revoked (Allen et al., 2023), but vesting constrains their tradability, delaying contributors' ability to monetize their stake (Abel & Kindermann, 2023). During the lock-up period, contributors are therefore exposed to the market valuation of the jointly produced platform, as reflected in the token price. Actions that depress platform value reduce expected payoffs at unlock, while value-adding behavior supports it.

Vesting thus introduces incentive backloading in the form of delayed realization of value, while the strength of this intertemporal discipline declines mechanically as vesting progresses and liquidity becomes available. Early in the vesting period, when a larger share of the stake is locked, contributors have stronger incentives to avoid short-term, platform-harming actions whose costs materialize later; as unlocks accumulate, the remaining continuation value shrinks and intertemporal moral hazard becomes more salient again (Georgiadis, 2022). The higher the share of upfront allocations, the earlier decrease incentives to avoid short-term exploits.

In line with these theoretical arguments, we conducted post-regression tests to explain heterogeneity in ATT estimates. The estimates show significant negative correlations between higher upfront unlock shares on external contribution shares ($\beta = -3.61$, $p = .045$). Further, we find a significant negative correlation ($\beta = -0.45$, $p = .0012$) of longer vesting length to the 80% vesting threshold on developer airdrop farming. Though, we cannot detect economically meaningful effects of airdrop design factors on platform redirection activities.

Figures B7-B8: Visualization of regressions on post-treatment outcome means (on the organization level; N=127) showing correlations between upfront unlock shares and external contribution shares (B7) and between the duration until the 80%-vesting threshold and developer airdrop farming

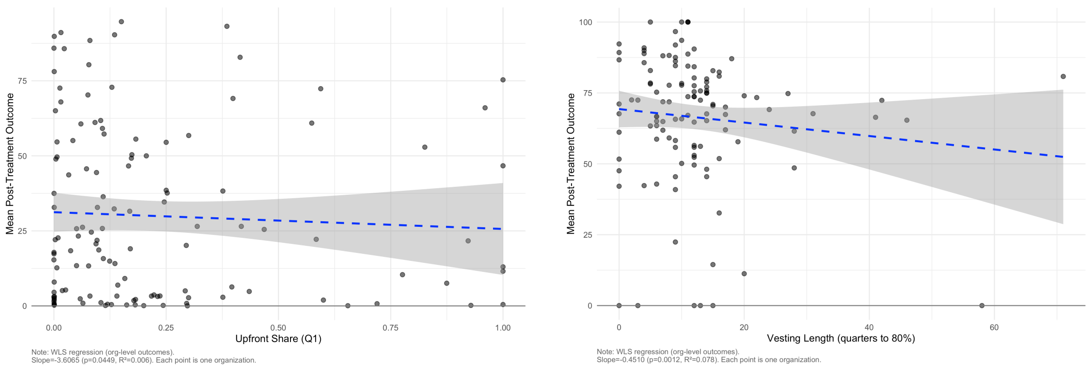

Overall, these post-regression estimates indicate the certain strategic choices of how to design the transition from proprietary to collective governance can have an effect on external contribution behavior and might have prevented opportunism once the corresponding platforms were collectively governed. Consistent with dynamic agency theory, governance token vesting therefore may mitigate but do not eliminate intertemporal distortions in a distributed, collective production setting with non-separable output (Holmström, 1982) and multitask moral hazard (Holmström & Milgrom, 1991), partially aligning contributor incentives with platform welfare while leaving scope for opportunistic behavior to re-emerge as vesting constraints relax.

Besides these dynamic aspects, also static factors may influence how control allocations incentivize external contributions. For example, the design of the created governance system may vary. Single DAOs experiment with sub-DAO structures, where topic-specific governance tokens instead of global governance tokens restrict control rights to sub-elements of the platform. According to dynamic agency theory, this task-specific authority allocations may prevent moral hazard in multi-tasking settings (Georgiadis, 2022). Further, as a specific form of governance right vesting, some DAOs require token holders to "stake" their governance tokens (i.e., fixed-time, temporary deposit into a trust account) in order to utilize the governance rights embedded in the tokens. During the staking period, token holders cannot trade the token. Hence, governance staking may create additional incentives alignment between contributors and the platform and offer a valuable opportunity for further research.

REFERENCES

Abadie, A. 2005. Semiparametric Difference-in-Differences Estimators. ***The Review of Economic Studies***, 72(1): 1–19.

Abel, D., & Kindermann, B. 2023. Signals for Governance Token Performance: Role of DAOs' Governance Token Design and Distribution. ***Academy of Management Proceedings***, 2023(1): 14886.

Allen, D. W. E. 2024. Crypto airdrops: An evolutionary approach. ***Journal of Evolutionary Economics***, 34(4–5): 849–872.

Allen, D. W. E., Berg, C., & Lane, A. M. 2023. Why airdrop cryptocurrency tokens? ***Journal of Business Research***, 163: 113945.

Athey, S., & Imbens, G. W. 2022. Design-based analysis in Difference-In-Differences settings with staggered adoption. ***Journal of Econometrics***, 226(1): 62–79.

Baker, A., Callaway, B., Cunningham, S., Goodman-Bacon, A., & Sant'Anna, P. H. C. 2025. Difference-in-Differences Designs: A Practitioner's Guide. ***Journal of Economic Literature***.

Bakos, Y., & Halaburda, H. 2022. Overcoming the Coordination Problem in New Marketplaces via Cryptographic Tokens. ***Information Systems Research***, 33(4): 1368–1385.

Baldwin, C., & Von Hippel, E. 2011. Modeling a Paradigm Shift: From Producer Innovation to User and Open Collaborative Innovation. ***Organization Science***, 22(6): 1399–1417.

Beck, R., Müller-Bloch, C., & King, J. L. 2018. Governance in the Blockchain Economy: A Framework and Research Agenda. ***Journal of the Association for Information Systems***, 19(10): 1020–1034.

Bertrand, M., Duflo, E., & Mullainathan, S. 2004. How Much Should We Trust Differences-in-Differences Estimates? ***The Quarterly Journal of Economics***, 119(1): 249–275.

Binance. 2024, April 20. Airdrop farmers spam GitHub after Starknet and Celestia distributions. ***Binance Square***. https://www.binance.com/en/square/post/4352734181730.

Boudreau, K. 2010. Open Platform Strategies and Innovation: Granting Access vs. Devolving Control. ***Management Science***, 56(10): 1849–1872.

Boudreau, K. J. 2012. Let a Thousand Flowers Bloom? An Early Look at Large Numbers of Software App Developers and Patterns of Innovation. ***Organization Science***, 23(5): 1409–1427.

Boudreau, K. J., & Jeppesen, L. B. 2015. Unpaid crowd complementors: The platform network effect mirage. ***Strategic Management Journal***, 36(12): 1761–1777.
Bresnahan, T., & Greenstein, S. 2014. Mobile Computing: The Next Platform Rivalry. ***American Economic Review***, 104(5): 475–480.
Callaway, B., Goodman-Bacon, A., & Sant'Anna, P. H. C. 2024. Difference-in-differences with a Continuous Treatment. ***Working Paper***.
Callaway, B., & Sant'Anna, P. H. C. 2021. Difference-in-Differences with multiple time periods. ***Journal of Econometrics***, 225(2): 200–230.
Cennamo, C. 2018. Building the Value of Next-Generation Platforms: The Paradox of Diminishing Returns. ***Journal of Management***, 44(8): 3038–3069.
Cennamo, C., & Santaló, J. 2019. Generativity Tension and Value Creation in Platform Ecosystems. ***Organization Science***, 30(3): 617–641.
Chen, Y., Richter, J. I., & Patel, P. C. 2021. Decentralized Governance of Digital Platforms. ***Journal of Management***, 47(5): 1305–1337.
Chod, J., & Lyandres, E. 2021. A Theory of ICOs: Diversification, Agency, and Information Asymmetry. ***Management Science***, 67(10): 5969–5989.
Chod, J., Trichakis, N., & Yang, S. A. 2022. Platform Tokenization: Financing, Governance, and Moral Hazard. ***Management Science***, 68(9): 6411–6433.
CNCF. 2016, March 10. Cloud Native Computing Foundation accepts Kubernetes as first hosted project; Technical Oversight Committee elected. ***CNCF***. https://www.cncf.io/announcements/2016/03/10/cloud-native-computing-foundation-accepts-kubernetes-as-first-hosted-project-technical-oversight-committee-elected/.
CoinMarketCap. 2026, January 9. Crypto Market Overview. ***CoinMarketCap***. https://coinmarketcap.com/charts/.
Cunningham, S. 2021. ***Causal Inference: The Mixtape***. Yale University Press.
Cutolo, D., & Kenney, M. 2021. Platform-Dependent Entrepreneurs: Power Asymmetries, Risks, and Strategies in the Platform Economy | Academy of Management Perspectives. ***Academy of Mangement Perspectives***, 35(4): 584–605.
Dahlander, L., & Magnusson, M. G. 2005. Relationships between open source software companies and communities: Observations from Nordic firms. ***Research Policy***, 34(4): 481–493.
DeFiLlama. 2026a. Uniswap UNI Token Unlocks & Vesting Schedules. ***DefiLlama***. https://defillama.com/unlocks/uniswap.
DeFiLlama. 2026b, January 9. All Chains DeFi TVL. ***DefiLlama***. https://defillama.com/chains.
Eisenmann, T. R., Parker, G., & van Alstyne, M. 2009. Opening Platforms: How, When and Why? ***Platforms, Markets and Innovation***.
Fisher, R. A. 1936. Design of Experiments. ***British Medical Journal***, 1(3923): 554.
Fong, Y., & Li, J. 2017. Relational contracts, limited liability, and employment dynamics. ***Journal of Economic Theory***, 169: 270–293.
Franck, E., & Jungwirth, C. 2003. Reconciling Rent-Seekers and Donators – The Governance Structure of Open Source. ***Journal of Management and Governance***, 7: 401–421.
Franke, N., & Hippel, E. V. 2003. Satisfying heterogeneous user needs via innovation toolkits: The case of Apache security software. ***Research Policy***, 32(7): 1199–1215.
Gawer, A., & Cusumano, M. A. 2002. ***Platform Leadership: How Intel, Microsoft, and Cisco Drive Industry Innovation***, vol. 5. Boston: Harvard Business School Press.
Georgiadis, G. 2022. Contracting with Moral Hazard: A Review of Theory & Empirics. ***Available at SSRN***, 4196247.

Geva, H., Barzilay, O., & Oestreicher-Singer, G. 2019. A Potato Salad with a Lemon Twist: Using A Supply-Side Shock to Study the Impact of Opportunistic Behavior on Crowdfunding Platforms. ***MIS Quarterly***, 43(4): 1227-A9.

GitHub. 2025, October 28. Octoverse: A new developer joins GitHub every second as AI leads TypeScript to #1. ***The GitHub Blog***. https://github.blog/news-insights/octoverse/octoverse-a-new-developer-joins-github-every-second-as-ai-leads-typescript-to-1/.

Goodman-Bacon, A. 2021. Difference-in-differences with variation in treatment timing. ***Journal of Econometrics***, 225(2): 254–277.

Grossman, S. J., & Hart, O. D. 1986. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. ***Journal of Political Economy***, 94(4): 691–719.

Halaburda, H., Livshits, B., & Yaish, A. 2025. Platform Building With Fake Consumers: On Double Dippers and Airdrop Farmers. ***NYU Stern School of Business Research Paper Forthcoming***.

Hart, O., & Moore, J. 1990. Property Rights and the Nature of the Firm. ***Journal of Political Economy***, 98(6): 1119–1158.

Hassan, S., & De Filippi, P. 2021. Decentralized Autonomous Organization. ***Internet Policy Review***, 10(2). https://doi.org/10.14763/2021.2.1556.

Hertel, G., Niedner, S., & Herrmann, S. 2003. Motivation of software developers in Open Source projects: An Internet-based survey of contributors to the Linux kernel. ***Research Policy***, 32(7): 1159–1177.

Hoffmann, M., Nagle, F., & Zhou, Y. 2024. The Value of Open Source Software. ***SSRN Electronic Journal***. https://doi.org/10.2139/ssrn.4693148.

Holmstrom, B. 1979. Moral Hazard and Observability. ***The Bell Journal of Economics***, 10(1): 74.

Holmstrom, B., & Milgrom, P. 1987. Aggregation and Linearity in the Provision of Intertemporal Incentives. ***Econometrica***, 55(2): 303.

Holmstrom, B., & Milgrom, P. 1991. Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. ***The Journal of Law, Economics, and Organization***, 7: 24–52.

Hsieh, Y.-Y., & Vergne, J.-P. 2023. The future of the web? The coordination and early-stage growth of decentralized platforms. ***Strategic Management Journal***, 44(3): 829–857.

Hsieh, Y.-Y., Vergne, J.-P., Anderson, P., Lakhani, K., & Reitzig, M. 2018. Bitcoin and the rise of decentralized autonomous organizations. ***Journal of Organization Design***, 7(1): 14.

Iacus, S. M., King, G., & Porro, G. 2011. Multivariate Matching Methods That Are Monotonic Imbalance Bounding. ***Journal of the American Statistical Association***, 106(493): 345–361.

Iacus, S. M., King, G., & Porro, G. 2012. Causal Inference without Balance Checking: Coarsened Exact Matching. ***Political Analysis***, 20(1): 1–24.

Iansiti, M., & Levien, R. 2004. Strategy as ecology. ***Harvard Business Review***, 82(3): 68–81.

Jacobides, M. G., Cennamo, C., & Gawer, A. 2024. Externalities and complementarities in platforms and ecosystems: From structural solutions to endogenous failures. ***Research Policy***, 53(1): 104906.

Jacobides, M. G., Knudsen, T., & Augier, M. 2006. Benefiting from innovation: Value creation, value appropriation and the role of industry architectures. ***Research Policy***, 35(8): 1200–1221.

Jacobides, M. G., & Tae, C. J. 2015. Kingpins, Bottlenecks, and Value Dynamics Along a Sector. ***Organization Science***, 26(3): 889–907.

Lakhani, K. R., & Von Hippel, E. 2003. How open source software works: "Free" user-to-user assistance. ***Research Policy***, 32: 923–943.

Lerner, J., & Tirole, J. 2005. The scope of open source licensing. ***Journal of Law, Economics and Organization***, 21(1): 20–56.

Levin, J. 2002. Multilateral Contracting and the Employment Relationship. ***The Quarterly Journal of Economics***, 117(3): 1075–1103.

Levin, J. 2003. Relational Incentive Contracts. ***The American Economic Review***, 93(3): 835–857.

Li, H., Shen, Q., & Bart, Y. 2018. Local Market Characteristics and Online-to-Offline Commerce: An Empirical Analysis of Groupon. ***Management Science***, 64(4): 1860–1878.
Loh, J., & Kretschmer, T. 2023. Online communities on competing platforms: Evidence from game wikis. ***Strategic Management Journal***, 44(2): 441–476.
Lommers, K., Makridis, C. A., & Verboven, L. 2023. Designing Airdrops. ***Available at SSRN 4427295***.
MacKinnon, J. G., & Webb, M. D. 2020. Randomization inference for difference-in-differences with few treated clusters. ***Journal of Econometrics***, 218(2): 435–450.
MacLeod, W. B., & Malcomson, J. 2023. ***Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment: Thirty Years On***.
MacLeod, W. B., & Malcomson, J. M. 1989. Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment. ***Econometrica***, 57(2): 447.
Makridis, C. A., Fröwis, M., Sridhar, K., & Böhme, R. 2023. The rise of decentralized cryptocurrency exchanges: Evaluating the role of airdrops and governance tokens. ***Journal of Corporate Finance***, 79: 102358.
Malinova, K., & Park, A. 2023. Tokenomics: When Tokens Beat Equity. ***Management Science***, 69(11): 6568–6583.
Markus, M. L. 2007. The governance of free/open source software projects: Monolithic, multidimensional, or configurational? ***Journal of Management & Governance***, 11(2): 151–163.
Murray, A., Kuban, S., Josefy, M., & Anderson, J. 2021. Contracting in the Smart Era: The Implications of Blockchain and Decentralized Autonomous Organizations for Contracting and Corporate Governance. ***Academy of Management Perspectives***, 35(4): 622–641.
Nagle, F. 2018. Learning by contributing. ***Organizatino Science***, 29(4): 569–587.
Nagle, F., Dana, J., Hoffman, J., Randazzo, S., & Zhou, Y. 2022. ***Census II of Free and Open Source Software—Application Libraries***. The Linux Foundation. https://doi.org/10.70828/KHEH5209.
O'Mahony, S. 2003. Guarding the commons: How community managed software projects protect their work. ***Research Policy***, 32(7): 1179–1198.
O'Mahony, S. 2007. The governance of open source initiatives: What does it mean to be community managed? ***Journal of Management & Governance***, 11(2): 139–150.
O'Mahony, S., & Ferraro, F. 2007. The Emergence of Governance in an Open Source Community. ***Academy of Management Journal***, 50(5): 1079–1106.
O'Mahony, S., & Karp, R. 2022. From proprietary to collective governance: How do platform participation strategies evolve? ***Strategic Management Journal***, 43(3): 530–562.
O'Mahony, S., & West, J. 2005. ***What makes a project open source? Migrating from organic to synthetic communities***.
Oster, E. 2019. Unobservable Selection and Coefficient Stability: Theory and Evidence. ***Journal of Business & Economic Statistics***, 37(2): 187–204.
Parker, G. G., & Van Alstyne, M. W. 2005. Two-Sided Network Effects: A Theory of Information Product Design. ***Management Science***, 51(10): 1494–1504.
Petryk, M., Qiu, L., & Pathak, P. 2023. The Impact of Open-Source Community on Cryptocurrency Market Price: An Empirical Investigation. ***Journal of Management Information Systems***, 40(4): 1237–1270.
Rayo, L. 2007. ***Relational Incentives and Moral Hazard in Teams***.
Rietveld, J., Ploog, J. N., & Nieborg, D. B. 2020. Coevolution of Platform Dominance and Governance Strategies: Effects on Complementor Performance Outcomes | Academy of Management Discoveries. ***Academy of Management Discoveries***, 6(3): 488–513.
Roth, J., & Sant'Anna, P. H. C. 2023. When Is Parallel Trends Sensitive to Functional Form? ***Econometrica***, 91(2): 737–747.

Rubin, D. B. 1980. Randomization Analysis of Experimental Data: The Fisher Randomization Test Comment. ***Journal of the American Statistical Association***, 75(371): 591–593.
Shah, S. K. 2006. Motivation, Governance, and the Viability of Hybrid Forms in Open Source Software Development. ***Management Science***, 52(7): 1000–1014.
Thatchenkery, S., & Katila, R. 2023. Innovation and profitability following antitrust intervention against a dominant platform: The wild, wild west? ***Strategic Management Journal***, 44(4): 943–976.
Token Terminal. 2026, January 9. ***Total value locked***. https://tokenterminal.com/explorer/metrics/tvl.
Uniswap Labs. 2020. Introducing UNI. ***Uniswap Labs***. https://blog.uniswap.org/uni.
van Pelt, R. V., Jansen, S., Baars, D., & Overbeek, S. 2021. Defining Blockchain Governance: A Framework for Analysis and Comparison. ***Information Systems Management***, 38(1): 21–41.
Von Krogh, Haefliger, Spaeth, & Wallin. 2012. Carrots and Rainbows: Motivation and Social Practice in Open Source Software Development. ***MIS Quarterly***, 36(2): 649.
Wang, Q., Li, B., & Singh, P. V. 2018. Copycats vs. Original Mobile Apps: A Machine Learning Copycat-Detection Method and Empirical Analysis. ***Information Systems Research***, 29(2): 273–291.
West, J. 2003. How open is open enough? ***Research Policy***, 32(7): 1259–1285.
Wing, C., Simon, K., & Bello-Gomez, R. A. 2018. Designing Difference in Difference Studies: Best Practices for Public Health Policy Research. ***Annual Review of Public Health***, 39(1): 453–469.
Young, A. 2019. Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results*. ***The Quarterly Journal of Economics***, 134(2): 557–598.
Yue, D., & Nagle, F. 2025. Igniting Innovation: Evidence from PyTorch on Technology Control in Open Collaboration. ***Working Paper***.